

ABSTRACT

Title of dissertation: **VARIATIONAL ALGORITHMS AND
RESOURCES FOR NEAR-TERM QUANTUM
SIMULATION**

**Troy Jacob Sewell
Doctor of Philosophy, 2023**

Dissertation directed by: **Professor Stephen P. Jordan
Department of Physics**

The difficulty of efficiently simulating quantum many-body systems was one of the first motivations for developing quantum computers and may also be one of the first applications to find practical computational advantage on real quantum hardware. With the relatively recent advent of publicly available quantum technologies, we have now entered the era of noisy intermediate-scale quantum (NISQ) computing. The capabilities of these technologies are evolving rapidly, and with them the computational affordances to which we have access. The time is now ripe to test the capabilities of existing quantum hardware and leverage them to the best of our ability toward achieving a practical quantum computational advantage.

In this dissertation, we address these aims by benchmarking a class of variational multi-scale quantum circuits for state preparation which are locally robust against noise. We demonstrate the advantages of these multi-scale circuits compared to a purely local circuit ansatz using the critical transverse-field Ising model to optimize circuit parameters, numerically test the noise resilience of observables and customized error mitigation techniques using a local gate noise model, and demonstrate the robustness of local subregion

preparation on an existing ion-trap quantum computer. We then show how the ground state optimized circuit can be simply extended to an ansatz for thermal state preparation using the separation of energy scales afforded by the multi-scale circuit structure.

Additionally, we evaluate the quantum resources needed for some quantum simulation tasks. We estimate the gate complexity of the site-by-site algorithm for fault-tolerant ground state preparation, which we extend to the case of degenerate Hamiltonians. Using matrix product states we evaluate the non-stabilizer quantum resources needed to represent thermalized subregions of a chaotic Potts model, which we use to address the feasibility for classical simulation of quantum hydrodynamics.

VARIATIONAL ALGORITHMS AND RESOURCES FOR NEAR-TERM QUANTUM SIMULATION

by

Troy Jacob Sewell

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2023

Advisory Committee:
Professor Jay Sau, Chair
Professor Stephen P. Jordan, Advisor/Co-Chair
Professor Maissam Barkeshli
Professor Brian Swingle
Professor Andrew Childs, Dean's Representative

© Copyright by
Troy Jacob Sewell
2023

Acknowledgments

It has been a long and winding road to reach this stage of my formal education, and I am immensely indebted to the many professors, teachers, mentors, friends, and family who have made it possible to continue this far. I am honored to have had the opportunity to pursue my Ph.D. at the University of Maryland, where I learned much from the many students and faculty in the Department of Physics and the Joint Center for Quantum Computing and Computer Science (QuICS), which have provided an environment so bountiful with interesting research in quantum computing. I am also thankful for the Department of Energy whose funding has made my research possible.

I am incredibly grateful for the guidance of my advisor Stephen Jordan as well as my senior collaborators Brian Swingle and Ning Bao. I am also grateful for Zohreh Davoudi and her research group for the inviting and engaging group meetings on quantum simulation, to Brian's research group which offered a virtual office during times of remote work, and my collaborator Christopher White whose enthusiasm for physics kept me engaged during the pandemic. I am thankful for my office mate and collaborator Aniruddha Bapat and our engaging conversations.

I want to thank both my parents for raising me in a comfortable and loving home, supporting my education and musical pursuits, and providing an example of creating the life you want to live. I thank my mother Charlene Jacob, for always entertaining my wild speculations and encouraging my curiosity, and my father James Sewell, who is always demonstrating that where there's a will there is a way.

I am thankful to my partner Natasha Jacobs, who has been there for me and is an inspiring example of someone who relentlessly follows their dreams. I also want to thank the late Van Jacobs, who always took an interest in my work and was an abundant source of helpful comments.

Lastly, I would be remiss not to mention the support of my all my friends outside of physics that have been there for me and helped to foster a broad spectrum of experience in life outside of numbers and books and computers.

Table of Contents

Acknowledgements	ii
Table of Contents	iii
List of Figures	vi
List of Abbreviations	viii
1 Introduction	1
2 Ground State Preparation	5
2.1 Degenerate Case	8
2.2 Resource Estimates	10
2.2.1 Kitaev Chain	10
2.3 Recursive state preparation	14
2.4 Overlap Bounds and the Schmidt spectrum	16
3 Variational Multi-scale Circuit Ansatz	22
3.1 Circuit Construction	26
3.2 Model and Variational Optimization	28
3.3 Numerical Benchmarking	31
3.4 Averaging Broken Symmetries	33
3.5 Comparison to Quantum Alternating Operator Ansatz	38
4 Fixed-Point State Preparation	41
4.1 Introduction	41
4.2 Preliminaries	46
4.3 Noise Models	53
4.4 Error Extrapolation	61
4.5 Experimental Data	63

5	TMERA	70
5.1	Background, intuition, and circuit ansatz	74
5.1.1	Background and intuition	74
5.1.2	Circuit ansatz	78
5.2	Model	82
5.3	Free energy minimization and properties of the TMERA	85
5.3.1	Entropy	86
5.3.2	Energy	87
5.4	Benchmarking	89
5.4.1	Energy and Entropy	89
5.4.2	State Fidelity	90
5.4.3	Two-point correlations	92
5.5	Discussion	95
5.6	Extension to systems away from criticality	96
6	Mana & Thermalization	97
6.1	Model, initial state, and methods	100
6.1.1	Model	100
6.1.2	Initial state	102
6.1.3	Method	105
6.1.4	Mana	105
6.2	Mana at infinite temperature	106
6.3	Mana at finite temperature	110
6.4	Discussion	113
7	Discussion	117
A	Gaussian Fermions	122
B	TMERA Appendix	130
B.1	Renormalization of temperature	130
B.2	Hamiltonian of Approximate Gibbs State	131
B.3	TFD purification	132
C	Mana Appendix	134
C.1	Initial state sampling	134
C.2	Defining and calculating mana	139
C.3	Convergence in Bond Dimension and Trotter Step	141
C.4	Subsystem entropy	143
C.5	Thermalization, typicality, and the eigenstate thermalization hypothesis: intuition	143
C.6	Gibbs state mana	146
C.7	Bounding the difference in mana between nearby states	147

C.8	Mana and canonical typicality	149
C.9	Krylov subspace evolution	151
	Bibliography	156

List of Figures

2.1	Kitaev chain spectral gap and overlap finite size scaling	11
2.2	Kitaev chain spectral gap phase diagram	12
2.3	Kitaev chain site-by-site overlap phase diagram	12
2.4	Energy gap above degenerate ground space	13
2.5	Principal angle phase diagram	14
2.6	Kitaev chain two halves ground state overlap	16
2.7	Largest Schmidt weight, site-by-site	20
2.8	Largest Schmidt weight, half system	21
3.1	DMERA circuit	23
3.2	Subsystem entanglement entropy	25
3.3	Ground state ansatz energy density and state fidelity	26
3.4	Subsystem infidelity	32
3.5	Symmetry averaged observable error	33
3.6	Relative improvement of symmetry averaging	38
3.7	Symmetry averaged correlation functions	40
4.1	Channel scaling dimensions	50
4.2	Past causal cone circuit diagram	54
4.3	Noisy channel tensor network	57
4.4	Subsystem fidelity	58
4.5	Channel noise response	58
4.6	Noisy fixed point energy density	59
4.7	Error susceptibility with channel depth	60
4.8	Error mitigation of noisy fixed point energy	64
4.9	Experimental convergence of energy density	67
4.10	Experimental convergence of local observables	69
5.1	TMERA circuit diagram	71
5.2	Energy versus entropy and state fidelity of TMERA	71
5.3	Alternate TMERA circuit diagram	76

5.4	Energy mode support	82
5.5	Relative error in free energy	85
5.6	Entropy by scale	88
5.7	Characteristic energy scale of circuit layers	89
5.8	Energy and entropy versus tempearture for TMERA states	91
5.9	Normalized infidelity for temperature and system size	93
5.10	Two-point correlations	94
6.1	Subsystem mana evolution, $T = \infty$	98
6.2	Energy gap ratio averages	101
6.3	Energy of stabilizer pairs	103
6.4	Subsystem mana versus entropy deficit, $T = \infty$	107
6.5	Subsystem mana evolution, finite temperature	109
6.6	Subsystem mana versus entropy deficit, finite temperature	111
6.7	Gibbs state subsystem mana	113
6.8	Long time Krylov evolution of mana for $L = 11$ subsystems	114
C.1	Initial state support on energy eigenstates	135
C.2	Half-system entropy numerical convergence, finite temperature	136
C.3	Subsystem entropy numerical convergence, finite temperature	137
C.4	Subsystem mana numerical convergence	140
C.5	Subsystem entropy evolution	144
C.6	Finite temperature Gibbs state subsystem mana scaling with system size . .	146
C.7	Krylov evolution of subsystem entropy, finite temperature	153
C.8	Late time subsystem mana versus energy deficit from Krylov evolution . . .	154

List of Abbreviations

AKLT	Affleck-Lieb-Kennedy-Tasaki
BPP	Bounded-error Probabilistic Polynomial time
BQP	Bounded-error Quantum Polynomial time
CCD	Charge-Coupled Device
CFT	Conformal Field Theory
ETH	Eigenstate Thermalization Hypothesis
DAOE	Dissipation Assisted Operator Evolution
DMERA	Deep Multi-scale Entanglement Renormalization Ansatz
DMRG	Density Matrix Renormalization Group
DMT	Density Matrix Truncation
GOE	Gaussian Orthogonal Ensemble
IR	Ifra-Red
JW	Jordan-Wigner
KW	Kramers-Wannier
L-BFGS	Limited memory Broyden-Fletcher-Goldfarb-Shanno
MERA	Multiscale Entanglement Renormalization Ansatz
MPO	Matrix Product Operator
MPS	Matrix Product State
MS	Mølmer-Sørensen
NISQ	Noisy Intermediate-Scale Quantum
NP	Nondeterministic Polynomial-time
PEPS	Projected Entangled Pair States
QAOA	Quantum Approximate Optimization Algorithm
QMA	Quantum Merlin Arthur
TEBD	Time-Evolving Block Decimation
TFD	Thermofield Double
TMERA	Thermal Multi-scale Etanglement Renormalization Ansatz
UV	Ultra-Violet
VQA	Variational Quantum Algorithm
VQE	Variational Quantum Eigensolver

Chapter 1: Introduction

The promise of practical quantum computing is to solve useful computational problems faster than possible by classical computers by using the unique affordances of information processing on quantum hardware [1]. Grover’s quantum search algorithm [2] is able to provide a polynomial time speedup over classical algorithms for unstructured search, but the most dramatic improvements offer an exponentially faster asymptotic scaling of resources compared to any known classical algorithms, as is the case for Shor’s quantum algorithm for factoring large numbers [3]. The physics of these quantum systems must be engineered to sustain some model of computation which is universal for quantum computing, i.e. some experiment with the system can be used to solve problems in the bounded-error quantum polynomial time (BQP) computational complexity class [4]. Problems of particular interest are those which are BQP-complete, meaning that quantum algorithms exist to solve the problem and any quantum algorithm may be reduced to an instance of it [5]. A few notable examples are computing topological invariants [6, 7], or simulating scattering problems in quantum field theory [8]. These problems in turn can serve as the basis for computing paradigms such as adiabatic quantum[9], topological quantum computing [10], or measurement based quantum computing [11].

While the aim of practical quantum computing is not yet achieved, many varieties of quantum hardware are in development to overcome the challenges of realizing precise control of a large scale quantum systems and claims of quantum advantage have already been made[12–14]. This milestone is meant to signal that something has been done by a quan-

tum computer which would be sufficiently difficult for any classical computer to replicate, although all existing quantum technologies are still too limited by size and noise to handle algorithms of much generality, which will require fault-tolerant quantum computation [15]. Although it is impossible to know how long it will take to develop the fully scalable fault-tolerant quantum computer needed to solve the most general quantum algorithms, some practical advantage may be accessible to noisy intermediate-scale quantum (NISQ) hardware of the near future [16].

Simulating quantum physics was one of the earliest motivations for quantum computing [17], and may also be one of the first accessible routes to realizing practical quantum advantage [18]. Quantum algorithms have been developed to tackle various simulation tasks of quantum systems, such as in quantum chemistry [19, 20] or quantum field theory [21, 22].

For special cases of quantum systems which are sufficiently constrained, classical algorithms exist to simulate them using resources scaling only polynomially in the system size, such as stabilizer or Gaussian fermion states. Clifford or matchgate circuits consist of unitary operations which respectively preserve the set of Gaussian fermion or stabilizer states, and so they may be simulated classically for large system sizes and are not able to perform universal quantum computation [23–28]. For any sufficiently general family of quantum states or dynamics, all known classical algorithms will scale exponentially in some computational resource, unless $BQP=BPP$. These classical simulation algorithms may try to generalize from some restricted cases, allowing better scaling of resources for large system sizes, but will inevitably scale exponentially in some computational resource if made sufficiently general.

In a similar vein, Clifford circuits and matchgate circuits may be enriched to universal models of quantum computation by adding additional operations which break the symmetry of the circuit class, or by supplying additional quantum computational resources in the form of non-stabilizer or non-Gaussian initial states, so-called “magic states” [29, 30].

Fault-tolerant quantum computation with stabilizer error-correcting codes usually involves Clifford gate circuits with magic states as a computational resource [31]. In measurement based quantum computation cluster states serve a very similar purpose to achieve universal computation [32].

Entanglement is another form of quantum resource which must be utilized in any algorithm to offer quantum advantage and can serve as a bottleneck to classical simulation algorithms [33]. Efficiently contractable tensor networks can represent quantum states with low entanglement using few resources and have proven useful in the simulation of many-body physics, such as the matrix product states (MPS) found using the density matrix renormalization group (DMRG), a classical variational method for finding ground state approximations [34–36].

In order for a quantum computer to do something truly difficult to replicate classically, as has already been claimed by some quantum devices, the problem and its implementation on hardware must evade the lack of any of these quantum resources which some classical algorithm may be able to leverage to its advantage.

Given limitations on size and the number of gates in current experiments, only relatively simple algorithms to prepare states and implement Hamiltonian evolution are accessible.

For state preparation on a NISQ device, a well-motivated family of parameterized circuits which are implementable on hardware may be used, with a hybrid quantum-classical feedback loop to optimize the circuit parameters and prepare an approximation of the target state [37]. For the case of ground state preparation this is known as a variational quantum eigensolver (VQE) [38, 39], and the circuit parameters are optimized to minimize the energy of the state. For thermal state preparation a similar procedure can be done to instead minimize the free energy for the target temperature.

The material in the following chapters is based on several papers by the author and collaborators which were posted online and published in academic journals prior to the

submission of this dissertation, which we cite along with the descriptions of each chapter.

In Chapter 2 we estimate the gate complexity of the site-by-site algorithm for fault-tolerant ground state preparation and extend the analysis to the case of degenerate Hamiltonians, based on material from Ref. [40]. In Chapter 3 we evaluate the use of a variational multi-scale circuit ansatz inspired by entanglement renormalization tensor networks to prepare the ground state of the critical Ising spin chain, based on material from Ref. [41]. We impose approximate symmetries of the critical ground state and use symmetry averaging to improve systematic error of these broken symmetries and compare to the performance of purely local variational circuit approximations. In Chapter 4 we show how subregion preparation of this circuit ansatz may be implemented on NISQ hardware and viewed as convergence to a noise-robust fixed point state by iterating a non-unitary quantum channel, based on material from Ref. [42]. We numerically simulate a local gate error to evaluate the noise resilience of local observables and customized error mitigation techniques, and demonstrate the convergence to a noisy fixed-point state with experimental data from ion-trap hardware. In Chapter 5 we then show how the separation of scales induced in the ground state optimization of the multi-scale circuit ansatz allows us to extend the circuit ansatz to thermal state preparation by using mixed product states to introduce additional energy and entropy to the state, based on material from Ref. [43]. Finally, in Chapter 6 we evaluate a quantity called mana which measures the non-stabilizer resources of a quantum state for local subregions of a thermalized state and discuss the implications of this for the difficulty of simulating quantum hydrodynamics, based on material from Ref. [44].

In chapters 2, 3, 4, and 5 we utilize numerical techniques for representing Gaussian fermion states to aid in the classical simulation of large system sizes. In chapter 6 we similarly use matrix product states to simulate states of one dimensional systems with short range entanglement.

Chapter 2: Ground State Preparation

The preparation of non-trivial quantum states is an important subroutine in quantum simulation. The canonical example is the (static) problem of preparing a quantum state that approximates the ground state of a Hamiltonian of interest, such as the electronic structure Hamiltonian in quantum chemistry or the Hamiltonian of a quantum field theory. In many cases, state preparation is also a necessary first step in the simulation of quantum dynamics, where one initializes a state that is expected to evolve in time with interesting dynamics. The algorithm proposed in [45] for simulating scattering in a scalar quantum field theory first involves a preparation of the vacuum of the interacting field theory, a subroutine that dominates the runtime cost of the full algorithm. It is therefore a well-motivated problem to design schemes for preparing the ground state of a known Hamiltonian. While such an algorithm cannot hope to be efficient in general due to the QMA-hardness of approximating the ground states of k -local Hamiltonians [46], it is nonetheless useful to understand when an efficient approximation is possible.

Classical methods for simulating real-time dynamics are often limited by their inability to accurately simulate interference effects (also known as the “sign problem”). The quantum approach is more natural: for a properly discretized field theory, one can map the local degrees of freedom directly to qubits of the quantum device, then prepare an initial state that corresponds to the state of interest (say a scattering state), and engineer interactions mimicking the true form of interaction seen in nature (analog) or provide a quantum circuit description of the time evolution (digital).

For theories in 1+1D spacetime, a fast state initialization algorithm was proposed [47] using ideas from tensor networks. A similarly motivated proposal used projected entangled pair states (PEPS) to prepare the ground state of a lattice in two spatial dimensions [48]. However, both proposals rely on knowing the tensor network description of the ground state and carry a (possibly large) classical computational overhead.

Recently in [49], an algorithm was introduced for the preparation of vacuum states of lattice Hamiltonians digitally on a quantum computer. Starting from the ground state $|g_{L_0}\rangle$ of the target Hamiltonian on a constant number of sites L_0 , the algorithm prepares successively larger ground states by adding one site at a time to the boundary of the lattice, until the number of sites equals the target size, L . Each site addition step involves a fixed-point Grover search [50] routine whose oracle calls invoke quantum phase estimation on the underlying ground states. The quantum phase estimation subroutine, in turn, uses time evolution under the target Hamiltonian. Therefore this “site by site” state preparation algorithm requires the ability to simulate the target Hamiltonian but is otherwise oblivious to the details of the ground state. Two additional conditions are assumed, for each $n \in \{L_0, \dots, L-1, L\}$:

1. The spectral gap m_n , or the difference between the energies of the first excited state and ground state, is non-zero: $m_n > 0$.
2. The successive ground state overlaps $\eta_n := |(\langle g_n | \otimes \langle q |) | g_{n+1} \rangle|$ are non-zero.

The second condition assumes the freedom to choose the state $|q\rangle$ on the added site(s) to maximize overlap.

On Hamiltonians that are geometrically local with each term bounded in spectral norm, the gate complexity of the site-by-site algorithm is $O\left(\frac{n}{m_n \eta_n} \log(1/\varepsilon)\right)$ for the iteration $|g_n\rangle \mapsto |g_{n+1}\rangle$. Success therefore crucially depends on how m_n and η_n scale with the number of sites n . When the spectral gap is constant, as is the case for massive field theories,

we have $\lim_{n \rightarrow \infty} m_n = O(1)$, giving a trivial contribution to the runtime.

A result of this work is analyzing the case where the Hamiltonian is degenerate, in which case the phase estimation oracle may be used to mark all states in the ground space using the gap m' between the ground space and the rest of the energy spectrum. In this case, there is no longer a single pair of ground states whose inner product we can use to upper bound the algorithm's runtime. However, we find that the quantity $\cos(\theta)$ may play the same role as the overlap η for the degenerate case, where θ is the largest principal angle between ground spaces of two successive Hamiltonians.

Reference [49] uses tensor network techniques to study the scaling of overlap for the massive Gross-Neveu model. The numerical calculations presented there on lattices up to $L = 50$ suggest that, for various settings of the coupling and bare mass, the overlap converges to a constant in the limit of large L . Similar results are obtained for the AKLT model. More generally, for models of interacting field theories, it is not known how the state overlap behaves as a function of system size and other model parameters.

Non-interacting models are useful for studying the behavior of physical properties over a broad range of system parameters, since they exhibit rich physics but are often computationally tractable. In the case of free fermions, the Hamiltonian, which is quadratic in the mode creation and annihilation operators, is diagonalizable in time $O(n^3)$ where n is the number of sites. The ground states of quadratic fermionic Hamiltonians are Gaussian, i.e. fully characterized by the expectation values of all degree-two monomials in the mode operators. As shown in [51], the magnitude of the inner product of any two fermionic Gaussian states can be exactly computed classically in time $O(n^3)$, while the gap is obtained directly as the smallest mode energy of the diagonalized Hamiltonian. Here, we use this feature to study the dependence of the spectral gap and the overlap for a range of free fermionic Hamiltonians, and provide concrete estimates for the runtime of the site-by-site algorithm on these models.

Although free fermion models are exactly solvable and therefore do not pose a computational challenge requiring quantum computers, they provide a tractable test case to investigate the general behavior of the ground state overlaps determining the performance of site-by-site quantum state preparation. Extrapolation to weakly coupled fermionic systems is possible through rigorous methods, and conjectures about behavior beyond weak coupling can be generated from the patterns observable in numerical data from the free case.

In Appendix A we describe details of the the free fermion formalism along with details of the model we analyze in this work, the Kitaev chain.

In this chapter we numerically estimate the spectral gap m of the Kitaev chain phase diagram, as well as the ground state overlaps η in the non-degenerate case and the quantity $\cos(\theta)$ for the degenerate case. In Section 2.3 we introduce a possible improvement to the site-by-site algorithm for one-dimensional systems by choosing a sequence of Hamiltonians which recursively joins together two ground states to double the system size. In Section 2.4 we also discuss some connections of the state overlap with the Schmidt spectrum.

2.1 Degenerate Case

In the case of a degenerate set of ground states (or some subspace of low-energy states) with an energy gap separating these states from the rest of the Hilbert space, the phase estimation oracle will mark each of these states and the Grover reflection operator will reflect about this multi-dimensional subspace. We will consider it a success to prepare any state within the ground space, and so for each step of the algorithm we may be in any state within the ground space of the Hamiltonian at that stage of the algorithm.

Jordan's lemma tells us that given a pair of projection operators Π_1 and Π_2 which

project onto the ground space of each Hamiltonian, the Hilbert space can be decomposed into one- and two-dimensional subspaces which are invariant under each of these projection operators. These same subspaces are invariant under the Grover iterate operator, which is a product of reflections about the subspace for each projector.

In the case of standard Grover search, one of these projectors must be rank-one and project onto the initial state, while the other projector may be higher dimensional in the case for multiple marked target states. This ensures that there is only one two-dimensional subspace where the Grover iterate has nontrivial action, which rotates the initial state into the target subspace. In this case where we are transitioning between degenerate ground spaces of two Hamiltonians, however, each projector has rank greater than one, and the Grover iterate generically has multiple two-dimensional subspaces where nontrivial rotations can occur. These rotation angles are given by the principal angles that characterize the geometric relation between the two ground spaces.

For each two-dimensional subspace i where the Grover iterate has nontrivial action, we can choose a pair of bases $\{|v_i\rangle, |v_i^\perp\rangle\}$ and $\{|w_i\rangle, |w_i^\perp\rangle\}$ such that

$$|v_i\rangle = \cos(\theta_i)|w_i\rangle + \sin(\theta_i)|w_i^\perp\rangle, \quad (2.1)$$

where $\theta_i \leq \pi/2$ is the principal angle for this subspace [52].

Using the standard Grover iterate, we would face the daunting problem of synchronizing the rotation frequencies of each subspace so that they are simultaneously in the ground space of the target Hamiltonian after a certain number of iterations, which could dramatically increase the number of iterations required to find a good approximate solution. This problem is solved, however, by using the fixed-point quantum search instead of standard Grover search. Each nontrivial subspace of the standard Grover iterate acts as an independent quantum search happening in parallel, and the fixed-point search ensures that each

subspace will converge as long as enough queries are used.

The number of queries required for each subspace to achieve success probability greater than $1 - \delta^2$ using the fixed-point quantum search is lower bounded by $\log(2/\delta)/\cos(\theta_i)$, hence the largest principal angle will give us the greatest lower bound. If this largest angle is $\pi/2$ then there is a possibility the algorithm will fail, depending on whether the initial state has support in this subspace.

The principal angles can be found by constructing a matrix of inner-products between an orthonormal basis of states for each ground space and computing the singular values of this matrix. These singular values will correspond to $\cos(\theta_i)$ for each of the principal angles, with the smallest eigenvalue corresponding to the largest principal angle [53]. Since the runtime of our algorithm is proportional to $1/\cos(\theta)$ this will give us an upper bound to the complexity.

2.2 Resource Estimates

2.2.1 Kitaev Chain

We begin with the Kitaev chain, a model of spinless fermions on a lattice in one spatial dimension, growing the open boundary condition ground state site-by-site. The Hamiltonian has single-site terms with coefficient $-\mu$ and nearest-neighbor coupling terms with coefficients t for hopping and Δ for squeezing.

$$H_{nn} = -\mu \sum_{j=1}^N c_j^\dagger c_j - \sum_{j=1}^{N-1} \left(t c_j^\dagger c_{j+1} + \Delta c_j^\dagger c_{j+1}^\dagger \right) + h.c. \quad (2.2)$$

Fixing $t = 1$, the phase diagram has critical lines at $|\mu| = 1$, and $\Delta = 0, |\mu| < 1$. These critical lines feature an inverse polynomially closing energy gap for finite system sizes. The regions with $|\mu| > 1$ are the topologically trivial gapped phases, while the regions with




Figure 2.1: Scaling of the energy gap (left) and overlap (right) with system size in the Kitaev chain for three values of μ along the line $\Delta = 1$ near the critical point at $\mu = 1$. The three points chosen are representative of the two phases on either side of the phase transition at $\mu = 1$ and the critical model at the phase transition itself. Best fit curves are drawn in dashed lines on the log-log plot for the energy gap, showing behavior consistent with asymptotically constant, polynomial, and exponential scaling. The overlaps for these regions show convergence to large constant values with large system size.

$|\mu| < 1$ are the topological phases. The trivial phase has an asymptotically constant gap, while the topological phases with open boundary conditions support Majorana zero modes with a gap closing exponentially in the system size. The scaling of the gap with system size for these three phases is shown in Fig. 2.1.

The finite system size scaling for the ground state overlaps at these same points in the phase diagram are also shown in Fig. 2.1. Despite the energy gap closing with system size for the critical point and point in the topological region, the ground state overlaps appear to converge to large constant values for each of these points at large system sizes. This can be seen by comparing Fig. 2.2 and Fig. 2.3, where the critical line at $\mu = 1$ marks where the gap begins to close in Fig. 2.2, while nothing special appears to happen along this line in Fig. 2.3.

The overlap of the target ground state with the ground state of one smaller system, plus one additional fermion, is relatively large everywhere except the critical line along $\Delta = 0$. Considering that we can choose the single site state optimally, this amounts to the fidelity of the $N - 1$ site ground state with the subsystem of the N site ground state. We expect these




Figure 2.2: Energy gap m for 1D spinless fermions, Kitaev chain open boundary conditions, $L = 72$, with logarithmic color scale taken base e .




Figure 2.3: Ground state overlap η when adding a single site for spinless fermion Kitaev chain with open boundary conditions, $L = 72$.




Figure 2.4: Energy gap m' between the lowest-two states (degenerate ground states when $|\mu| < 1$), and the rest of the spectrum, $L = 72$.

to have high overlap where there is small correlation length, due to edge effects having negligible effect on most of the state. However the overlaps appear to be large elsewhere, even along the critical lines.

The fidelity for mixed Gaussian states can be computed using a more complicated formula in terms of their respective covariance matrices [54]. For numerical stability we add a small boundary term to break the degeneracy.

Since our Hamiltonian is degenerate for large system sizes when $|\mu| < 1$, in this phase we should instead consider the gap between the degenerate ground space and the rest of the spectrum, which we denote m' . In the gapless region of phase space this secondary gap is large as seen in Fig.2.4 for system size $L = 72$.

To calculate the lower bound given by $\cos(\theta_{max})$ we calculate the inner products between the two ground states of each Hamiltonian, which we numerically separate by adding




Figure 2.5: Largest principal angle $\cos(\theta)$ when adding a single site for spinless fermion Kitaev chain with open boundary conditions, $L = 72$.

a small boundary term coupling the two ends of the Kitaev chain. The basis of states for the L -site Hamiltonian we call $\{|g_L\rangle, |e_L\rangle\}$. For the $(L - 1)$ -site Hamiltonian we use two basis sets, $\{|g_L\rangle|0\rangle, |e_L\rangle|0\rangle\}$ and $\{|g_L\rangle|1\rangle, |e_L\rangle|1\rangle\}$. We compute the matrix of inner products between the L - and $(L - 1)$ -site ground spaces for each basis set depending on whether the newly introduced fermion mode is occupied or not, and then take the supremum of these two choices for each point in the phase diagram which will serve as a lower bound to the actual value of $\cos(\theta)$, seen in Fig.2.5.

2.3 Recursive state preparation

In the previous sections we have looked at preparing a state by adding one site at a time. This can equivalently be thought of as a process where one first prepares two disjoint

ground states, one of size $n - 1$ and the other of size 1, and then “joins” them to form the n -site ground state with the coupling between the two sub-chains turned on. Nothing in the algorithm restricts us from choosing a different division. For instance, one can prepare two copies of ground states on $n/2$ sites, and then “join” them in the middle to form an n -site ground state. The runtime of this procedure will now depend on the overlap between the decoupled half-chain ground states and the full ground state. However, since we still only turn on a coupling term between two sites, it is not clear that this overlap should be worse than the overlap between an $n - 1$ -site and n -site ground state. On the other hand, by applying recursion and running disjoint couplings in parallel, one can get from a ground state on a constant number of sites to the full ground state on n sites using $\log n$ steps instead of n steps.

In the single-site scheme one must make a choice of starting state on the ancilla qubit, which affects the overlap. In the recursive scheme, there is no such choice to be made. Instead, the complexity of the algorithm is determined solely by the energy gaps of the Hamiltonians and the overlap between the tensor product between the ground state of the full Hamiltonian and the tensor product of the ground states on the half-systems. By considering these inner products we can compare the complexity of this recursive scheme to the single-site method.

As discussed in Appendix 2.4, this implies that in these cases there is not too much entanglement across a central cut and the ground state on half the system size has significant overlap with the largest Schmidt state.

The potential benefit of this recursive construction is that, while the overlaps at each step may be reduced, building up many copies of smaller systems to be joined together may be done in parallel to reduce total runtime.




Figure 2.6: Ground state overlap η when gluing together two Kitaev chains of equal length.

2.4 Overlap Bounds and the Schmidt spectrum

The ground state of the full lattice is in general entangled across any cut. Crucial to the performance of the site-by-site algorithm is the overlap between the full (entangled) ground state and a state that decomposes as a product across some cut. In the special case where we add one site at a time, the initial state is a product on the $(N - 1)$ -site lattice and the last site.

We establish bounds on the largest possible overlap between two states where one is a product across a given cut that partitions the system into two parts that we call L and R . Consider a Schmidt decomposition of the state, which expresses the entangled state as a positive linear combination of product states,

$$|\psi\rangle = \sum_{i=1}^r \sqrt{\lambda_i} |\psi_i^L\rangle |\psi_i^R\rangle, \quad (2.3)$$

where r denotes the Schmidt rank, and the coefficients $\{\lambda_1, \dots, \lambda_r\}$ are in descending order.

Suppose we evaluate the inner product of this state with the product state $|\phi\rangle = |\phi^L\rangle|\phi^R\rangle$. Let $a_i = \langle\phi^L|\psi_i^L\rangle$ and $b_i = \langle\phi^R|\psi_i^R\rangle$, which satisfy $\sum_i |a_i|^2 \leq 1$ and $\sum_i |b_i|^2 \leq 1$. Then, the overlap between the two states is

$$|\langle\phi_0|\psi\rangle|^2 = \left| \sum_{i=1}^r \sqrt{\lambda_i} a_i b_i \right|^2 \leq \sum_{i=1}^r \lambda_i |a_i|^2 \leq \lambda_1, \quad (2.4)$$

where we first used Cauchy-Schwarz inequality, followed by the fact that λ_1 is the largest Schmidt coefficient. Therefore, the squared overlap is bounded above by the largest Schmidt coefficient.

A product state with overlap λ_1 always exists, namely the maximal weight Schmidt state $|\psi_1^L\rangle|\psi_1^R\rangle$, which has $a_1 = b_1 = 1$. The maximal Schmidt coefficient λ_1 is also always at least $1/r$, with the minimal value obtained for a flat spectrum. The Schmidt decomposition is of maximal rank if it is equal to the Hilbert space dimension of the smaller subsystem, i.e. $r = \min(d_L, d_R)$ for d_L and d_R the Hilbert space dimensions of the left and right subsystems.

The largest Schmidt coefficient is also related to the entanglement entropy S as follows:

$$e^{-S} = \prod_{i=1}^r e^{\lambda_i \ln(\lambda_i)} = \prod_{i=1}^r \lambda_i^{\lambda_i} \leq \prod_{i=1}^r \lambda_1^{\lambda_i} = \lambda_1. \quad (2.5)$$

Thus, the entanglement entropy of a state across a bipartition gives a lower bound on the maximal overlap possible with a product state across the same bipartition which is tighter than the bounds given by the Hilbert space dimensions or the Schmidt rank.

In fact, in terms of Renyi entropies S_k , the Schmidt rank is equal to e^{S_0} , while the largest Schmidt coefficient is equal to e^{-S_∞} . Each successive Renyi entropy is bounded by the preceding one, $S_k \geq S_{k+1}$. So, the Renyi entropies give a sequence of tighter lower

bounds to the maximal Schmidt rank.

$$\frac{1}{r} = e^{-S_0} \leq e^{-S_k} \leq e^{-S_\infty} = \lambda_1 \quad (2.6)$$

For the case where one subsystem is a single qubit, an entangled state will have Schmidt rank 2, so there must always exist a product state with overlap at least $1/2$, with higher overlap product states existing when the qubit is not maximally entangled.

In practice, the state of the larger subsystem will be the ground state of the previous target Hamiltonian, while the smaller subsystem is introduced in a state we are free to choose.

Let us write the ground state of the system of size N as a Schmidt decomposition of product states on the first $N - k$ sites and the last k sites, with Schmidt rank r ,

$$|\psi_0^N\rangle = \sum_{i=1}^r \sqrt{\lambda_i} |\psi_i^{N-k}\rangle |\psi_i^k\rangle \quad (2.7)$$

To achieve high overlap it is necessary for the ground state of the previous Hamiltonian, $|\psi_0^{N-k}\rangle$, to have large support over the subspace spanned by the r Schmidt states $|\psi_i^{N-k}\rangle$. Any support outside of this subspace will be a loss in potential overlap, and support on the largest weight Schmidt states allows for the highest potential overlap. Let us also choose $|\psi_{r+1}^{N-k}\rangle$ so that $|\psi_0^{N-k}\rangle$ has support entirely on this subspace of Schmidt vectors extended by one additional state orthogonal to the r -dimensional subspace spanned by the Schmidt vectors. Such a choice always exists and can be constructed by, e.g., a Gram-Schmidt process.

$$|\psi_0^{N-k}\rangle = \sum_{i=1}^{r+1} a_i |\psi_i^{N-k}\rangle \quad (2.8)$$

If we prepare the ground state of the $N - k$ site system with some state on the smaller

system $|\phi^k\rangle = \sum_i b_i |\psi_i^k\rangle$, the overlap with the N site ground state is upper bounded by

$$\left| \langle \psi_0^N | \psi_0^{N-k} | \phi^k \rangle \right|^2 = \left| \sum_{i=1}^r a_i b_i \sqrt{\lambda_i} \right|^2 \leq \sum_{i=1}^r \lambda_i |a_i|^2. \quad (2.9)$$

This upper bound from the Cauchy-Schwarz inequality is achievable by setting $b_i = a_i^* \sqrt{\lambda_i} / \sqrt{\sum_i |a_i|^2 \lambda_i}$. This quantity is equivalent to the fidelity between the ground state of the $(N-k)$ -site system and the reduced density matrix σ of the N -site ground state on the first $N-k$ sites:

$$\mathcal{F}(|\psi_0^{N-k}\rangle, \sigma) = \langle \psi_0^{N-k} | \sigma | \psi_0^{N-k} \rangle = \sum_{i=1}^r \lambda_i |a_i|^2. \quad (2.10)$$

Therefore, while the largest possible overlap is the largest Schmidt weight λ_1 , the best achievable overlap with $|\psi_0^{N-k}\rangle |\phi^k\rangle$ over all possible choices of k is the fidelity of the $(N-k)$ -site ground state with the reduced density matrix of the N -site ground state on the $N-k$ sites.

Uhlmann's theorem states that the fidelity between two mixed states is the maximal fidelity of any purification. We may always fix the purification of one state, however, because the fidelity is not changed by a unitary transformation acting only on the purifying subsystem. In our case, we may fix the purification of the mixed state to be the ground state of the larger system, and the maximization occurs over the unentangled degrees of freedom added to the already pure state to match the purified Hilbert space dimension. The new degrees of freedom are introduced in an optimal state to reproduce the fidelity of the smaller mixed state.

Next, we numerically investigate the tightness of the upper bound on overlap which the Schmidt coefficients provide by comparing these to the overlaps themselves from section 2.2. As shown in figures 2.7 and 2.8 this bound is frequently far from tight and provides only a loosely informative guide to the regions in phase space in which the algorithm is




Figure 2.7: Largest Schmidt coefficient λ , upper bounding the ground state overlap, computed for the case of adding a single site for spinless fermion Kitaev chain with open boundary conditions, $L = 72$.

efficient.




Figure 2.8: Largest Schmidt coefficient compared against ground state overlap for the case of gluing together two Kitaev chains of equal length.

Chapter 3: Variational Multi-scale Circuit Ansatz

The multi-scale entanglement renormalization ansatz (MERA) tensor networks represent quantum states on a d dimensional lattice using a $d + 1$ dimensional representation. The extra dimension can be interpreted as scale, with slices of the network along this dimension corresponding to successively coarse grained states. In one dimension, the tree-like MERA tensor network can represent ground states of critical systems, reproducing polynomially decaying correlation functions and logarithmic scaling of subsystem entanglement entropy [55–57]. In the case where the local tensors are fermionic Gaussian unitaries, the networks can be viewed as a wavelet transform on fermion operators and can be rigorously shown to support good approximations of local free-fermion ground states [58–60].

The local unitary structure of a MERA tensor network may be interpreted as a quantum circuit which introduces further UV degrees of freedom scale-by-scale to entangle qubits of the target ground state. However, on a quantum computer it is more natural to allow increased circuit depth rather than local bond dimension as a means of increasing expressivity of an ansatz, giving rise to a class of quantum circuits known as DMERA [61] which could be used as a variational circuit for ground state preparation. Of particular interest for near-term application is that local observables and correlation functions of ground states prepared by DMERA circuits feature an inherent resilience to local noise stemming from circuit topology [42, 61], which will be investigated further in Chapter 4. It is also possible to prepare subregions of DMERA states on quantum computers much smaller than the total




Figure 3.1: State preparation by successive application of $D = 4$ scale transformation circuits $U_\ell(\theta)$ with new qubits initialized in the state $|0\rangle$. Colors show the identity of local gate parameters imposed to replicate approximate translation and scaling symmetry.

system size, which may help to leverage the capabilities of small quantum devices [16]. This has been demonstrated on an ion-trap quantum computer, where the critical ground state subregion can be prepared as a robust fixed-point state of a local quantum channel derived from a DMERA quantum circuit [42].

MERA tensor networks can be contracted in polynomial time classically. However, the polynomial scaling with bond dimension is quite severe, *e.g.* $O(\chi^8)$ for 1D and $O(\chi^{16})$ for 2D using the schemes proposed in [56]. Thus direct execution of MERA on quantum computers can yield large polynomial speedups, and can additionally serve as a method for preparing initial states in the context of a quantum algorithm for simulating quantum dynamics.

This work investigates the feasibility of multi-scale circuits for variational ground state preparation on quantum computers in the fashion of a variational quantum eigensolver (VQE) [38, 39], or more generally an ansatz circuit for some variational quantum algo-

rithm (VQA) [62], where some family of parameterized circuits are minimized according to some Hamiltonian energy in hopes that a good approximation of the ground state can be found. Another MERA-inspired quantum circuit ansatz for variational ground state preparation was studied in Ref. [63], however, the circuit ansatz studied in the present work more directly follows the constructions of Refs. [58, 59, 61].

Conformal field theories and the use of renormalization are of central importance to the study of quantum field theories, which may generally be viewed as deformations from critical fixed points of a renormalization group flow. The prospect to simulate conformal field theories with critical lattice models and including the use of renormalization theory is seen as an important step in the ability to handle more general study of field theory simulations on quantum computers [64]. Ground states of critical systems feature large correlation lengths which are challenging to reproduce using a more locally entangled ansatz such as matrix product states or short depth quantum circuits. Scale invariance, however, makes these states somewhat simpler to describe using a multi-scale ansatz which may incorporate this symmetry directly into the circuit parameterization. Gapped states may be better able to be represented by an ansatz with local correlations, but may also be prepared using a multi-scale circuit with parameters that differ between scales, and using fewer scales overall due to the lack of long range correlations.

We benchmark the viability of the DMERA ansatz for variational state preparation by numerically optimizing for approximate ground states of the critical Ising spin chain in one dimension. We are able to find high fidelity ground state approximations using relatively low circuit depth D of each scaling transformation. The ansatz states and local energy density converge to their exact values exponentially with D , with relative error in the energy density below 10^{-8} for $D = 6$. Key features of critical ground states in one dimension such as polynomially decaying correlation functions and logarithmic scaling subsystem entanglement entropy are found on average. See Fig. 3.1 for a representation of




Figure 3.2: Relative error in average entanglement entropy of N qubit subsystems for an $L = 256$ qubit state, which is exponentially small in D . The subsystem entropy scales logarithmically with N for critical ground states in one dimension, with DMERA circuits featuring an excess of entropy for $D < 3$ and a small entropy deficit for $D > 3$.

the multi-scale variational circuit with $D = 4$.

Symmetry-averaging is also used to improve the systematic error of local observable expectation values. This can be implemented in hybrid variational schemes by using classical post processing to average observables which are related by symmetries which are explicitly broken in the circuit ansatz. In the case of the critical Ising model spatial translation and Kramers-Wannier (KW) symmetry are used, although we also note that these symmetries are still approximately reproduced in the ansatz states themselves, with smaller variance over these symmetry groups for increasing D . Similar ideas have been studied in the context of Refs. [65, 66], where symmetry projection is used in post-processing to improve broken symmetries of ansatz states.




Figure 3.3:

Left: Relative error of variational fixed-point ground state energy for channels of depth D , using the infinite volume energy density of $-4/\pi$. Trend lines showing the exponential scaling of energy density error for our ansatz (solid) as well as those from the analytic wavelet construction of [59](dashed) and the numerically optimized non-Gaussian MERA in [57] (dotted).

Right: Normalized state infidelity $1 - \mathcal{F}^{1/L}$ of states prepared by depth D scaling transformation for ground states of various system sizes L .

Both quantities appear to decay exponentially with D with a coefficient of approximately -4.89 .

3.1 Circuit Construction

We consider states which double the number of qubits with each scaling transformation, interleaving new qubits in the zero state between qubits from each of the previous layers. After ℓ scaling circuits we have a state on $L = 2^\ell$ qubits. Our goal is to prepare an approximation to the ground state of a corresponding critical Ising Hamiltonian. Each transformation is a quantum circuit $U(\theta)_\ell$ of depth D , parameterized in terms of angles θ .

$$|\psi_{\ell+1}\rangle = U_\ell(\theta) \left(|\psi_\ell\rangle \otimes |0\rangle^{\otimes 2^\ell} \right). \quad (3.1)$$

Approximate translation and scale invariance of the state are imposed by having each scale transformation be a periodic brickwork circuit invariant under translation by two sites, and constraining each scaling transformation circuit at different layers to be made up of the same gates, albeit acting on different numbers of qubits. The approximate translation sym-

metry reduces the total number of variational parameters from scaling with the total number of gates, $O(DL)$, to only with total circuit depth $O(D \log L)$. The approximate scale symmetry further reduces the number of circuit parameters down to only $O(D)$, the depth of a single scaling transformation circuit. (See fig 3.1).

These symmetries cannot be imposed exactly due to the discrete nature of the gates and scaling transformations, so the exact symmetries are broken by the ansatz state, yet retained as approximate symmetries. Deviations in local observables are exponentially small in D but constitute a major source of error for experimentally realistic values of D , *e.g.* $D = 1, 2, \dots, 6$ as studied here. However, as shown in Fig. 3.5, the error in two-point correlation functions related by KW symmetry are nearly out of phase with each other. Consequently, this source of error can be reduced by averaging expectation values over this known symmetries. This symmetry averaging lowers error in the correlation functions by approximately two orders of magnitude, yielding relative error below 10^{-7} for correlation functions using the $D = 6$ state. This improvement is consistent across both of the example Hamiltonians considered in this work (eq. 3.3 and 3.6), as shown in Figs. 3.5 and 3.6.

In [58–60], wavelet transformations on fermionic modes are used to construct approximations to free fermion ground states. These constructions use knowledge of the underlying physics to analytically construct the desired wavelet and achieve accuracy exponential in the order of the wavelet (corresponding to circuit depth D). In this work we instead determine the variational parameters in our DMERA ansatz by numerical optimization. Inspired by the free-fermion MERA constructions of [58] we take

$$u(x, y) = \begin{bmatrix} \cos(x) & 0 & 0 & \sin(x) \\ 0 & \cos(y) & \sin(y) & 0 \\ 0 & -\sin(y) & \cos(y) & 0 \\ -\sin(x) & 0 & 0 & \cos(x) \end{bmatrix} \quad (3.2)$$

as the local gates in our circuit, where the two variational parameters specify the rotation on the odd and even parity subspaces of the pair of qubits. Because we use real-valued parity-conserving gates, the ansatz states will always be parity-even states with time-reflection symmetry.

Similarly to [59], we find that variationally optimized parameters achieve substantially more accurate energy densities than analytically constructed parameters, even though the latter become exact in the limit of DMERA of infinite depth or MERA of infinite bond dimension.

3.2 Model and Variational Optimization

We focus on the transverse-field Ising model at criticality, namely

$$H_I = - \sum_{j=1}^L X_j X_{j+1} + Z_j. \quad (3.3)$$

This spin chain is known to be well described by a conformal field theory with $c = 1/2$ at low energies, but is also integrable due to the Jordan Wigner duality relating it to the following free fermion model

$$H = i \sum_{j=1}^{2L} \gamma_j \gamma_{j+1}, \quad (3.4)$$

where γ_j are Majorana operators. In this convention there are $2L$ Majorana operators for L spatial sites, with operators γ_{2j} and γ_{2j-1} together comprising the fermion at spatial site j . At the critical point the coupling between Majorana operators at the same spatial site and neighboring sites are equal, which is sometimes referred to as the "half-shift symmetry", and allows for the simple description of the Hamiltonian in Eqn. (3.4). For the

spin Hamiltonian on an even number of sites with periodic boundary conditions, the ground state has even parity and is equivalent to the ground state of the Majorana Hamiltonian with anti-periodic boundary conditions, so $\gamma_{2L+1} = -\gamma_1$.

To achieve efficient numerical simulations we exploit two properties. First, the gates of eq. 3.2 are matchgates, which allows us to simulate their action on Gaussian fermion states efficiently as linear transformations on Majorana operators [26, 67]. By constraining to matchgate tensors, tensor networks may be used to represent Gaussian fermion operations and states [27, 28]. We use these techniques to compute two-point correlation functions out to large distances such as 512 lattice spacings. Secondly, the expectation values of local operators in MERA and DMERA states can be accelerated by exploiting the constant-width causal cones imposed by these circuits. Local reduced density matrices of the global pure state defined by the ansatz can be obtained as output from iterated quantum channels on a constant number of qubits. These local density matrices are all that is needed to compute the energy of the variational state under a local Hamiltonian. As the number of layers in the MERA ansatz is increased the density matrix converges exponentially to the fixed point of this channel, which can be easily computed, as described in [42]. We tune variational parameters by minimizing the energy density of this fixed-point state.

For sufficiently shallow circuits, we find well-optimized variational parameters using gradient descent and L-BFGS with restarts. For higher depth circuits we obtain variational parameters by starting from an optimized circuit for a given value of D with an additional row of gates inserted at the end of the circuit with parameters initialized near the identity, and then re-optimizing to obtain a solution for $D + 1$. Similarly, we may alternatively insert two rows of near-identity gates anywhere within the scaling circuit to obtain a solution for $D + 2$. For the larger depth circuits, the optimization by adding additional parameters to a smaller solution was better at finding good minima. Random initialization of all parameters was less likely to find an energy minimum any lower than found previously

with shorter depth circuits. We therefore relied on this technique of bootstrapping a better $D + 1$ minimum from a depth D minimum in order to find circuits which continued to show exponentially small error in the energy with circuit depth.

By using the total energy to estimate the average energy density, we are implicitly averaging over symmetries of the Hamiltonian, since the full symmetry group we consider for the critical Ising model relates each local term in the Hamiltonian to every other. By not preparing the whole state and only measuring some terms of the Hamiltonian, we run the risk of minimizing the energy in a non-variational way, i.e. minimizing an estimate of the global energy below what is possible for the ground state, while inadvertently raising the energy for terms not measured because the global energy still cannot be below that of the ground state. While this could happen in principle, the spatial distribution of expectation values does not have a large variance due to the approximate translation symmetry still found in the circuit ansatz, and so this likely would not happen in a catastrophic way in practice.

Imposing swap-symmetry on the local gates would constrain the odd-parity rotation angle $y = 0$. For the isometry gates only the output legs are swapped, so instead this angle would be set to $y = -\pi/4$, resulting in an ansatz state which is symmetric with respect to spatial reflections. This choice of parameters is optimal for the $D = 1$ case. However, for $D > 1$ we find that by relaxing this constraint, reflection-symmetry broken states can be prepared using the same circuit depth which have lower energy and higher fidelity. By transforming the parameters $y \rightarrow -y$ for non-isometry gates and $y \rightarrow \pi/2 - y$, a reflected ansatz state with the same energy density and fidelity is prepared.

The circuit parameters we use throughout this paper for $D \leq 6$ can be found at the end of Appendix A. The corresponding error in energy density appears to be exponentially small in D with a scaling coefficient of -4.89 , seen in Fig. 3.3. We compare this fit line to two similar multi-scale ground state representations for the Ising model, the wavelet based

construction found in Ref. [59] which uses a very similar circuit ansatz but constructs the ideal parameters analytically rather than through energy minimization, and the more traditional MERA tensor network states in Ref. [56] found through numerical optimization of the energy. In order to compare these different methods, we note that local gates of the circuit based methods can be grouped together to form a MERA tensor network of bond dimension $\chi = 2^{D-1}$. So, while the MERA from [56] does not have a circuit depth D , we will use this equivalence in order to compare ansatz states of similar complexity. That being said, the circuit based constructions of this work and [59] are much more constrained ansatz states compared to a generic MERA of bond dimension $\chi = 2^{D-1}$, due to the local circuit structure.

We find that while maintaining much of the simplicity of the wavelet construction of [59], our numerically optimized states are able to outperform those constructions at equal circuit depth and appear to scale with a larger scaling coefficient. Our ansatz states even outperform the general MERA states of [56] for short depth circuits, but are beaten at larger bond dimensions, which should be expected due circuits being much more constrained than a MERA state of equivalent bond dimension.

3.3 Numerical Benchmarking

Having minimized the energy density for circuits with $D \leq 6$, we can compare our ansatz states with the true ground state for a variety of other measures. We compare the state fidelity and subsystem fidelities for different (sub)system sizes and circuit depths as well as logarithmic scaling of subsystem entropy. Finally we look at the two-point correlation functions and show how averaging local observables over symmetry groups broken by the circuit ansatz can greatly improve the estimation accuracy of these observables.

We prepare finite sized ground states using the fixed-point optimized circuit parameters




Figure 3.4: Normalized infidelity of N qubit subsystems of a $L = 256$ qubit state, averaged over different subsystem locations.

and applying $\mathcal{L} = \log_2 L$ scaling transformation layers to prepare the state on L qubits.

Interestingly, the infidelity of our ansatz states seems to be exponentially small in D , with the same scaling coefficient of -4.89 as the energy density, see Fig. 3.3. We compare different sized systems by normalizing the state infidelity as $1 - \mathcal{F}^{1/L}$, which we find clusters together ansatz states of the same D but different system sizes L .

In Fig. 3.4 we compare the infidelity of N qubit subsystems of a $L = 256$ qubit system, normalized by subsystem size. This behaves similarly to that of global state fidelity with identical exponential scaling in D and clustering of subsystem infidelity from the normalization. The smallest subsystems, however, have an even smaller normalized infidelity than that of the global state, and so local observables should be even more accurate than global state properties.

The relative error in average subsystem entropy is shown in Fig. 3.2, showing exponential accuracy with D . Critical ground states in one dimension feature more long-range entanglement than area-law states, which would have constant entropy in one dimension.




Figure 3.5:

Left: Relative error in the translation-averaged expectation value of the quadratic Majorana operators $\gamma_{2i-1}\gamma_{2j}$ (solid) and $\gamma_{2i}\gamma_{2j+1}$ (dashed) at fixed distance $|i-j|$. These two cases are related by a KW transformation, and have systematic errors which appear oscillatory on a logarithmic scale and nearly out of phase with each other.

Right: Further averaging over the KW symmetry cancels much of the error of quadratic Majorana observables, reducing the error in the expectation value estimate by orders of magnitude.

Instead, subsystem entropy scales logarithmically with the subsystem size which becomes a major challenge of strictly local circuits which can only entangle qubits within a finite range. Multi-scale circuits introduce entanglement at all scales and can therefore prepare states with this scaling of subsystem entropy. In fact, the ansatz states for $D \leq 3$ have on average an excess of entropy rather than a deficit.

3.4 Averaging Broken Symmetries

Although we impose a quasi-translation invariance in the local gates of our ansatz, the actual circuit breaks this translation invariance within each circuit layer down to translations by an even number of sites. The resulting state is only exactly invariant under translations by $L/2$ sites due to this breaking of symmetry at each circuit layer. Despite this, an approximate symmetry is still maintained, so local observables related by translations have a small variance which decreases exponentially with D .

In the Majorana representation there are two operators γ_{2i-1} and γ_{2i} at site i , so a transla-

tion by n sites amounts to shifting the index of all Majorana operators by $2n$, i.e. $\gamma_j \rightarrow \gamma_{j+2n}$.

At criticality, the Ising model possesses an additional symmetry which constrains the expectation value of local Z and XX operators to be identical [68]. The KW transformation maps between these sets of operators as

$$Z_i \rightarrow X_i X_{i+1}, \quad X_i X_{i+1} \rightarrow Z_{i+1} \quad (3.5)$$

which is a symmetry of the Hamiltonian. This symmetry becomes more apparent in the Majorana representation, since the Jordan-Wigner transformation maps Z_i to $\gamma_{2i-1} \gamma_{2i}$ and $X_i X_{i+1}$ to $\gamma_{2i} \gamma_{2i+1}$, so the KW transformation amounts to shifting the Majorana index by one and is sometimes known as a “half-shift” symmetry of the Majorana operators since two of these transformations amount to a spatial translation by one site. This transformation is clearly a symmetry of the Hamiltonian in Eq. (3.4), but is broken away from criticality when the Z and XX operators have different coefficients.

On the left panel of Fig. 3.5 we look at the relative error in the expectation values of the two sets of quadratic Majorana operators $\gamma_{2i-1} \gamma_{2j}$ and $\gamma_{2i} \gamma_{2j+1}$. Translation symmetry dictates that the ground state expectation values are the same for equal $i - j$, so we can average over the group of symmetries, the cyclic group on L elements, and we plot the relative error after translation-averaging in the left panel of Fig. 3.5. These two sets of observables decay inverse polynomially in the distance between the two Majorana operators as $|i - j|^{-1}$, and are related to each other by the KW transformation.

Interestingly, the relative error in these two sets of correlation functions appear oscillatory when plotted on a logarithmic scale in the distance $|i - j|$, which we attribute to the breaking of continuous scaling symmetry down to an approximate symmetry due to the discrete scaling transformations. The amplitude of these deviations decays with D , but also appears to be nearly out of phase with each other for the two sets of operators related by the

KW transformation. Averaging over this additional $\mathbb{Z}_2$ symmetry we can cancel much of this remaining systematic error due to the respective deviations being nearly out of phase, resulting in a much more accurate estimation of the observables after averaging of the KW symmetry, seen in the right panel of Fig. 3.5.

We can assess how well this cancellation by KW symmetry works by comparing the average amount of error of correlators at a fixed distance, denoted $|\overline{\delta\langle\mathcal{O}_{ij}\rangle}|$, with the amount of error in the symmetry-averaged expectation value, denoted $|\delta\langle\overline{\mathcal{O}_{ij}}\rangle|$. In Fig. 3.6 we see that this cancellation works remarkably well, with the symmetry averaged mean showing multiple orders of magnitude better accuracy than the typical local observable, particularly for large D . This shows that states of larger circuit depth feature both a lower amount of variance in their oscillatory expectation values as well as a higher degree of error cancellation, which compound to produce highly accurate estimators of local observables after symmetry averaging.

In Fig. 3.6 we also include values from circuits approximating the ground state of an alternate Ising model

$$H_{I'} = \sum_{i=1}^L -X_i X_{i+1} + X_{i-1} Z_i X_{i+1}. \quad (3.6)$$

This Hamiltonian is related to the Ising model by a constant depth unitary $\prod_j \frac{1}{\sqrt{2}}(-I + iX_j X_{j+1})$, which smears the single Pauli Z term into a 3-body operator, but is still equivalent to a quadratic fermion model by Jordan-Wigner duality. This other set of parameters for the $H_{I'}$ ground state feature similar degrees of error cancellation.

While these two Hamiltonians share much of the same physics due to their equivalence up to a short-depth local transformation, the variational circuits we find to approximate their respective ground states are not made equivalent by any short depth local transformation. Both sets of circuits preparing ground state approximations to H_I or $H_{I'}$ are scale-invariant variational circuits with different sets of parameters. It would be possible to prepare ap-

proximate ground states for H_I by using the circuits for $H_{I'}$ and then applying the short-depth transformation that relates the two Hamiltonians, or vice versa, but the resulting circuits are not scale-invariant due to the final circuit layer differing from all the previous ones.

This breaking of scale symmetry is explicitly used in some applications of MERA, where an initial and final layer of tensors which differ from the scale-invariant layers are used and can be referred to as "transition layers". These may be beneficial to include explicitly because the finite volume and lattice spacing of the model already break scale-invariance at the largest and smallest scales, and so a circuit with complete scale symmetry may not be the most ideal representation of the ground state. We restrict our attention to circuits with identical variational parameters for all scaling transformations within the quantum circuit in this work, and to this extent we consider the circuits found for ground states of H_I and $H_{I'}$ to give independent confirmation to the efficacy of these circuits.

Subsystem entropy also varies with location due to lack of exact translation invariance. However, entropy is a much more coarse grained property of a state and did not display a great degree of error canceling from translation averaging. That is, the subsystem entropies were typically biased in a particular direction for all subsystems, so the average entropy was not much more accurate than that of a typical subsystem.

Here, we have specifically considered the translational invariance and KW symmetries of two-point functions on a one-dimensional N -site translationally-invariant lattice with periodic boundary conditions. This symmetry averaging procedure can be directly generalized to arbitrary groups of symmetries, *e.g.* S_n symmetries on permutation-symmetric systems, the translation and rotation symmetries in a fixed number of spatial dimensions, or internal global or gauge symmetries, which could be either discrete or continuous.

Averaging over translations by a symmetry group G when measuring an observable comes at no cost in qubit count or circuit depth. Rather, the requirement is for repeated

executions of the circuit and measurement. Such repeated measurements are already necessary, as $1/\varepsilon^2$ measurements are needed to measure an observable with statistical error of order ε . Instead of repeating $1/\varepsilon^2$ measurements on an individual choice of G -translation of a given observable, one can instead choose a uniformly random G -translation of the observable for each measurement. In the absence of any *a priori* information to suggest some particular G -translation as preferable, this G -averaging will in the worst case yield equivalent accuracy to a randomly chosen but fixed G -translation, and in the best case can result in greatly improved accuracy due to cancellation of errors of opposite sign. The latter situation is observed in the two examples H_I and $H_{I'}$ considered here. It remains for future work to investigate which models exhibit such strong benefit from cancellation and which do not.

While we do not need to increase the circuit size we may still have to compile multiple circuits in order to measure the full set of symmetry related observables supported on different subregions of the state. Each circuit will remain easy to implement, but it is often more costly on NISQ hardware to draw measurement samples from multiple circuits compared to the same number of samples from the same circuit. Despite this, it is still less resource intensive than doing a full projection onto the trivial representation of the symmetry group as in Ref. [65] and therefore well suited to NISQ applications.

Since Pauli strings only have eigenvalues ± 1 , their variance is $1 - \langle \mathcal{O} \rangle^2$ for an observable $\mathcal{O}$. The symmetry averaged expectation value will still only have measurement outcomes ± 1 , so aggregating the measurements from different observables will result in sample with variance $1 - \langle \overline{\mathcal{O}} \rangle^2$.

This technique is complementary to the widely-used Richardson extrapolation and randomized resampling methods introduced in [69], as it mitigates systematic error induced by the ansatz rather than random error introduced by the hardware. These observations also suggest that imposing strict symmetries could be overly restrictive for a circuit ansatz.




Figure 3.6: Improvement in accuracy from symmetry averaging also increases with D . The ratio is plotted between the error magnitude in the symmetry-averaged correlation function, denoted by $|\delta\langle\mathcal{O}_{ij}\rangle|$, and the average error in particular expectation values with fixed $|i - j|$, denoted by $|\delta\langle\mathcal{O}_{ij}\rangle|$. Also included are the ratios for approximate ground states of the modified Ising model H_I , which show similar improvement from symmetry averaging.

Instead, it could be more advantageous to consider ansatz states which can break these symmetries while retaining accurate symmetry-averaged expectation values. With more quantum resources available, one could consider a coherent analogue of this error mitigation scheme in which states with approximate symmetries are projected onto the proper symmetric subspace after preparation to achieve higher fidelity states. Specifically, one can project onto the G -invariant subspace preparing a control register in the uniform superposition, executing a controlled G translation on the target state, and then measuring the control register in the Hadamard basis and postselecting on the zero outcome of this measurement.

3.5 Comparison to Quantum Alternating Operator Ansatz

Another natural approach to variational state preparation is using a strictly local circuit by evolving commuting terms in the Hamiltonian in sequence to form a quantum alternating

operator ansatz [70, 71]. Using a short depth ansatz makes it feasible for NISQ devices, and these circuits may be constrained to conserve symmetries of the Hamiltonian, such as preparing only translation invariant states. In this case, they do not benefit from symmetry averaging, unlike DMERA. This approach was used to find exact variational circuits for ground states of the critical Ising model on $2p$ spins using p rounds of a QAOA ansatz [72], and has been implemented experimentally on an ion-trap device [73].

Exact variational circuits are not possible with fewer than $L/2$ rounds, as this is the minimum circuit depth needed for the ansatz to “see the whole graph” [74]. That is, the past causal cone of some subregion of a local circuit ansatz contains the entire initial state only when the circuit depth is L or greater, making global entanglement possible. Before a local circuit can see the whole system it also is unaware of system size. So, even though the ground state of a finite size Ising chain with periodic boundary conditions has energy density lower than the infinite volume value of $-4/\pi$, a local circuit with depth less than L could not prepare a state with energy below this threshold. This also means the energy density for a local circuit ansatz will be identical for all system sizes greater than the circuit depth.

The energy density and normalized ground state infidelity for a local QAOA circuit ansatz are shown in Fig. 3.7. The circuits are initialized using parameters provided in [72] for preparing exact ground states for system sizes $L = 2p$, and then the energy is optimized for systems with $L > 2p$ so that the circuit does not see the whole system, and therefore cannot prepare the exact ground state. The resulting states give us an idea of how short depth local circuits can perform when preparing approximate ground states. Both errors in energy density and state fidelity appear consistent with polynomial decay in the circuit depth. This contrasts with the DMERA ansatz which achieves exponentially improving accuracy with increasing D .




Figure 3.7:

Left: Error of variational energy density from the infinite volume value of $-4/\pi$ against total circuit depth / number of rounds p for systems $L > 2p$.

Right: Normalized pure state infidelity versus number of rounds p , for various system sizes $L > 2p$. Error in both the energy density and normalized infidelity for short depth ansatz states appear consistent with polynomial decay in p for large system sizes.

Chapter 4: Fixed-Point State Preparation

4.1 Introduction

The current period of experimentation on digital quantum devices that are still quite limited in terms of size and error rates has been dubbed the Noisy Intermediate-Scale Quantum (NISQ) era [16], and it is an ongoing question within the field of quantum computation to find what techniques, if any, may be leveraged to implement practical quantum algorithms on these early devices. Prior to the development of fully fault-tolerant quantum computers, error mitigation procedures will be needed to achieve robust quantum computation in the presence of noise. In [61], Kim and Swingle proposed that the stable nature of renormalization group fixed points could provide intrinsic robustness to a class of quantum state-preparation algorithms based on a variant of the Multi-scale Entanglement Renormalization Ansatz (MERA) that they named Deep MERA, or DMERA. Similar circuits were also considered for their noise-resilience in [75]. MERA are a well-known variational tensor network ansatz for describing ground states of scale-invariant quantum systems, such as at critical points [57, 76]. The expectation values of local observables can be computed in polynomial time classically using MERA. Unfortunately, the polynomial scaling with bond dimension χ is of very high degree (e.g. $O(\chi^8)$ for 1D and $O(\chi^{16})$ for 2D using the schemes proposed in [56]), making MERA impractical for many classical computations. However, a MERA tensor network may be interpreted directly as a quantum circuit for preparing the state described. Thus MERA circuits can be used to calculate local ob-

servables on a quantum computer with a polynomial but large scaling advantage relative to classical computers. DMERA is a variant of MERA more directly tailored to quantum state preparation applications by dealing with local quantum circuits of variable depth rather than general unitary tensors of variable bond dimension.

Given a quantum circuit for preparing a quantum state and a subset of the qubits of the state to be prepared, one can define the causal cone of the subset by discarding all qubits and gates at the earliest stage possible such that the reduced density matrix on the chosen subset remains unchanged. A noteworthy property of DMERA circuits is that geometrically local patches selected from one-dimensional DMERA states have causal cones of constant width, independently of the volume of the one dimensional state from which they are taken. Viewed operationally, the causal cone can be regarded as a quantum channel for preparing the reduced density matrix for the local patch of the overarching state. This channel is composed by iterating a procedure in which new qubits are introduced in the $|0\rangle$ state, unitary gates are applied, and some qubits are discarded. From a renormalization group point of view, this procedure can be regarded as zooming in on a smaller region of the underlying scale-invariant system by fine-graining the lattice discretization, and then discarding the qubits that move out of the field of view under this zooming process [61]. The ability to measure local observables and correlation functions by preparing the causal cone of a small subset of qubits from a large state can dramatically reduce the resources needed and effects of noise [77], both of which are crucial for developing realistic applications for NISQ devices.

The quantum channel and renormalization group flow points of view are related according to $\lambda = 2^{-\Delta}$, where λ is an eigenvalue of the superoperator defining the local quantum channel being applied at each iteration that scales the state by a factor of two, while Δ is the scaling dimension of the associated operator in the underlying conformal field theory [78]. If the scaling dimensions are all positive, as is generically the case, then the

quantum channel has a unique attractive fixed point toward which any initial state will converge exponentially at a rate bounded by the smallest scaling dimension. This provides the state-preparation procedure with a degree of intrinsic self-correction. Dissipative quantum state preparation and computation has a long history [79–87] and within this context the DMERA perspective adds a systematic way of deriving concrete quantum circuit implementations of the needed quantum channels for a class of state preparation problems as well as an immediate way to predict convergence rates based on known scaling dimensions of conformal field theories.

Here, we experimentally and numerically investigate the application of DMERA for the preparation of the ground state of the critical Ising model. Via the Jordan-Wigner transformation, this model can be solved using techniques for free-fermion models. Exploiting the structure provided by this exact solution, Evenbly and White have derived wavelet-based analytic expressions for the tensors by which a local free-fermion MERA can describe the ground state [58]. Due to the exactly solvable nature of this model, the circuit ansatz is an example of a matchgate circuit which can be simulated efficiently classically [26]. We thus do not propose this model for a demonstration of quantum speedup. Rather, by choosing a model with analytically determined parameters in the tensors, we can make a clean investigation of convergence and error-robustness without confounding issues coming from optimization of variational parameters. A modified version of the circuit with additional gates dual to non-quadratic fermion operators would break the integrability of the circuit dynamics. This would lead to structurally analogous circuits that are challenging to simulate classically, and likewise able to prepare non-Gaussian fermionic states which arise in interacting fermion models and nonintegrable spin chains.

A noteworthy feature of this state preparation method is that it requires the repeated introduction of fresh ancilla qubits in the $|0\rangle$ state and subsequent discarding of qubits that are no longer needed. It is due to this open-system nature of the procedure that entropy

can be removed and self-correction can be attained. This is a necessary feature for all dissipative quantum algorithms, as well as for the “holographic” methods of [88–90], but is not possible on most present-day quantum computing test beds. However, measurement and resetting of individual qubits during the execution of a quantum circuit is possible on the quantum CCD architecture used by the Honeywell (now Quantinuum) HØ device on which our experiments were performed [91]. In principle, the quantum channel derived from DMERA can be repeated indefinitely on a quantum CCD device, such that the attractive nature of the fixed point constitutes an active form of error correction which perpetually steers the quantum state back toward its target. This is in stark contrast with most NISQ algorithms which rely on limited circuit depth to limit the accumulation of error.

While we focus our numerical benchmarking and experimental data on the 1D quantum Ising model at criticality, multi-scale ansatz states which prepare states via renormalization-inspired circuits are expected to be applicable to many other systems of interest. Some nonintegrable models which feature scale symmetry and have been previously studied with traditional MERA techniques would be natural models to consider for future study of a multi-scale quantum circuits and their local channel fixed-points. These include the critical points of the 1D lattice models studied in Ref. [57], such as the Potts, XX, and Heisenberg spin chains. In two spatial dimensions, MERA tensor networks are able to represent states with topological order [92], such as logical states of the toric code [93], and so multi-scale circuits could also be used to prepare ground states of more sophisticated systems with topological order which may also be useful for quantum error-correction [94]. Separation of length-scales and energy scales is ubiquitous in many physical models of interest and naturally arises from locality of physical Hamiltonians, and so a multi-scale ansatz may be applicable to preparing states of interest for these models. Strictly local circuits make it difficult to introduce entanglement between distant regions of a state, which multi-scale circuits circumvent by introducing long range entanglement through the initial layers of

coarse-grained state preparation.

Scaling symmetry found in critical models reduces the ansatz complexity by identifying gates from different layers, but noncritical states may also be prepared by optimizing parameters for different layers of the multi-scale circuit independently, resulting in a renormalization-like state preparation for non-scale invariant models. Non-critical models with an energy gap, and therefore a finite correlation length ξ , will also only require $\log(\xi)$ layers to prepare because regions separated further than this length do not require long-range entanglement. In quantum chemistry, using Gausslet basis functions has been a topic of recent study which adds multi-scale structure to traditional numerical techniques to provide higher accuracy simulations in a non-scale-invariant setting [95, 96], and so multi-scale representations may provide analogous benefits to simulation of quantum chemistry on quantum computers as well.

While other circuit ansatz architectures may also be suitable to prepare a given model, the multi-scale structure offers the robustness of local observables that is not found in other variational ansatz circuits, and may be a preferable way to prepare states on NISQ devices for that reason. We see this state preparation procedure as one way that the use of physics-inspired methods may help to inform ansatz circuits and simulation methods on NISQ devices, and hope that further influence from physics may be used to develop NISQ applications.

In addition to quantum simulation, MERA-inspired approaches to quantum machine learning have been proposed, such as multi-scale tensor network classifiers [97] and quantum convolutional neural networks [98], and may also benefit from properties of multi-scale quantum circuits studied in this work.

4.2 Preliminaries

The DMERA circuits we consider consist of a depth D layer with D variational parameters which is iterated to drive the system into an approximate ground state supported on a limited number of qubits. The circuit appears to be local within one depth D scale transformation. However, if we fix the geometry of the qubits in the final state of the circuit, earlier layers require interaction of pairs of geometrically distant qubits. This fact favors implementation on near-term ion-trap computers where nonlocal gates can be implemented natively without swap gates. Larger-depth implementations of the quantum channel achieve smaller systematic error arising from the limitations of the ansatz, but incur greater amount of error from imperfect implementation because of the increased number of gates in the past causal cone. We explore numerically the tension between approximation accuracy and noise susceptibility to find the optimal channel depth for a given error rate. This suggests that while the overall depth can be quite large with fixed points stabilized by local dissipative dynamics, NISQ devices will still need to focus on channels with relatively short depth to limit the effect of noise. We find numerically that the approximation accuracy is exponential in D while the noise susceptibility scales linearly with D , suggesting an optimal depth such that $\varepsilon \sim e^{-D}/D$ for gate error rate ε . For large enough D this is just $D \sim \log(1/\varepsilon)$, however, for the larger error rates of NISQ devices we will be working in the regime of small depth where subleading terms and constant factors are crucial for choosing the optimal circuit for particular task at hand.

Each layer of the circuit implements one scale transformation by a spatially homogeneous depth D brickwork circuit of two-qubit gates. Within a single scale the local gates are spatially homogeneous, built by one initial layer of "isometry" gates $W(\theta_1)$ and $D - 1$ layers of gates $U(\theta_i)$ for depth i within the layer, as seen in Fig. 3.1.

The two-qubit gates W and U are matchgates and therefore preserve parity. They are

each parameterized in terms of a single angle, the rotation angle of the even-parity subspace, which gives us D variational parameters for the angle at each depth within the scaling transformation. The odd-parity subspace is fixed by permutation symmetry of the qubits involved in each gate. The unitary gate U is symmetric under the simultaneous permutation of incoming and outgoing qubits, which fixes the action on the odd-parity subspace to be the identity, up to a phase. The W gates are in spirit isometries with only one incoming qubit, since one of the two qubits is always introduced as a $|0\rangle$ state, so they are symmetric under permutations of only the two outgoing qubits. This has the effect of propagating an incoming $|1\rangle$ to the symmetric pair $|01\rangle + |10\rangle$, while an incoming $|0\rangle$ state becomes an arbitrary even-parity state which is always permutation-symmetric.

$$W(\theta) = \begin{pmatrix} \cos(\theta - \frac{\pi}{4}) & 0 & 0 & \sin(\theta - \frac{\pi}{4}) \\ 0 & \frac{1}{\sqrt{2}} & \frac{-1}{\sqrt{2}} & 0 \\ 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 0 \\ -\sin(\theta - \frac{\pi}{4}) & 0 & 0 & \cos(\theta - \frac{\pi}{4}) \end{pmatrix} \quad U(\theta) = \begin{pmatrix} \cos(\theta) & 0 & 0 & \sin(\theta) \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ -\sin(\theta) & 0 & 0 & \cos(\theta) \end{pmatrix} \quad (4.1)$$

Since our state is scale-invariant we apply the same scale transformation repeatedly, implementing identical steps of fine-graining to prepare the renormalization UV fixed point defined by our local scaling operation. Since the past causal cone of a local region is constant width, we can use this to define a quantum channel on a constant number of qubits that implements the local scaling operation on a fixed subsystem size. Since each scaling layer introduces new qubits in order to fine-grain the state by zooming in on the virtual lattice, the constant qubit channel also needs to discard qubits at the edges which end up outside of the causal cone. This combination of preparing new ancilla and discarding edge qubits is built into the definition of the local scaling channel along with the unitary

scale transformation, making the channel non-unitary but physically implementable on a quantum computer.

Since the unitary circuit implementing the channel is only invariant under translation by an even number of sites, there are two different ways the subregion of retained qubits can align with the circuit. The two choices of subregion are akin to choosing to follow a subregion moving to the left or right when descending down the state preparation circuit of a large spin chain. Sampling over infinite sequences of these left and right channels would correspond to an averaging of the location of our n qubit subregion over the infinite sized spin chain. We implicitly perform this averaging of subregion location by studying the channel which is the statistical mixture of the left and right channels with equal probability, which ensures that the fixed point energies measured are still variational for the infinite sized system.

There is a single eigenvalue of the quantum channel with $\lambda = 1$, corresponding to the fixed point operator of the channel. All other operators have eigenvalues with magnitude less than one, meaning those operators will decay exponentially with successive iterations of the channel. The scaling dimensions of these operators are given by $\Delta = -\log_2 \lambda$. A local n -qubit channel has 2^{2n} eigenvalues, the largest of which correspond to the slowly decaying low energy modes described by the CFT with small scaling dimension. We compare the scaling dimensions Δ of the primary operators of the CFT describing the critical Ising model with those derived from the local quantum channels, seen in Table 4.1. In this table we also compare the fixed point energy density of these channels E^*/L to the exact energy density in the thermodynamic limit $E_0/L = -4/\pi$. Additional scaling dimensions are shown in Fig. 4.1, where the smallest 32 scaling dimensions in the CFT are compared with those derived from the largest 32 eigenvalues of the $D = 2$ channel on three qubits. The largest eigenvalues match well with the expected scaling dimensions, suggesting that the low energy sector of the spin-chain is indeed well described as a CFT, while the smallest

	exact	$D = 2$	$D = 4$
E^*/L	-1.27323..	-1.23948..	-1.26757..
Δ_I	0	0	0
Δ_σ	0.125	0.136..	0.123..
Δ_ε	1	1	1

Table 4.1: Critical data extracted from quantum channels preparing approximations of the critical Ising ground state, using wavelet-derived angles of [58] compared with exact values. The exact ground state energy density in the thermodynamic limit $E_0/L = -4/\pi$ is compared to the fixed point energy density E^*/L of local quantum channels. The scaling dimensions Δ of the three primary operators listed above are derived from the three largest eigenvalues of the quantum channel. Additional dimensions compared in Fig. 4.1.

eigenvalues that correspond to rapidly decaying high energy modes that are not universal and do not correspond to quantities of the CFT.

The positive scaling dimensions and unique fixed point imply that local subsystems converge exponentially and only $\log(1/\varepsilon)$ layers are needed to prepare the target state with precision ε . However, even the exact fixed point is still just an approximation of the true ground state, limited in precision by the depth D of the ansatz used, and the actual fixed point state prepared by a NISQ device will have perturbations from the ideal fixed point introduced by noise.

We can understand why the DMERA circuit based on wavelets is able to prepare the critical Ising state by considering the action of the circuit in terms of fermion modes via a Jordan-Wigner transformation. The critical Ising model spin Hamiltonian $H_I = -\sum_i X_i X_{i+1} + Z_i$ is transformed into the Majorana Hamiltonian $H_M = -i \sum_i \gamma_i \gamma_{i+1}$ ¹, both of which have ground state energy density $E_0/L = -4/\pi$. Being translation invariant, the eigenmodes of the massless free-fermion Hamiltonian dual to the critical Ising model are plane-waves. Matchgate circuits implement linear transformations on fermion modes, which would require depth scaling linearly with the system size to implement an exact

¹The Ising model with periodic boundary conditions is equivalent to a free-fermion model with anti-periodic boundary conditions in the even parity subspace and periodic boundary conditions in the odd parity subspace, but for local quantities boundary effects may be ignored.




Figure 4.1: Scaling dimensions Δ derived from the eigenvalues λ of the quantum channel for $D = 2$. Exact values for the 32 smallest scaling dimensions of the Ising CFT compared with those derived from the largest 32 eigenvalues of the noiseless channel. Perturbations to these scaling dimensions in noisy versions of the channel are also included.

transformation mapping local modes to the proper plane-wave modes via a local quantum circuit. This can be cast as a QAOA circuit with depth $L/2$ with $O(L^2)$ gates to prepare the exact critical Ising ground state for a size L system [72]. The wavelet transformation offers an alternate basis of local wave-packets, which can be mapped to from local modes using a $O(D \log(L))$ depth circuit with $O(DL)$ gates by breaking the strict locality of the QAOA circuit. Since the ground state is invariant under transformations that do not mix the occupied and unoccupied modes, it is conceivable that these wave-packet modes could be made from only occupied modes and therefore prepare the exact ground state. In practice the quasi-local modes found through variational optimization of a low-depth DMERA circuit architecture only prepare an approximate ground state, albeit one whose accuracy increases exponentially with the layer depth D . The analysis of [58] considered the tight-binding model of 1D lattice of spinless fermions at half-filling, which can be decoupled into two copies of the Hamiltonian H_M with local transformations. They designed

the circuit and set of angles to specify a wavelet transformation on the fermion modes that occupies the Dirac sea of momentum modes with $|k| < \pi/2$, with vanishing support at the Fermi points where $|k| = \pi/2$. Larger-depth variations of the circuit then correspond to higher-order wavelets which are able to prepare a better approximation by canceling higher derivatives of the mode occupation around the Fermi points. These circuits can be directly adapted for the critical Ising model as outlined in the Appendix of [58], which serves as the basis of our analysis. Additional work on wavelet circuits has been done exploring application to both free-fermion and free-boson systems [59, 60, 99].

The relation of different wavelet basis states by both translations and scaling transformations makes them especially suited for critical models, which have these operations as symmetries. Both the wavelet circuits and MERA tensor networks reflect these symmetries in the spatial homogeneity of tensors at a given depth in the circuit, and the homogeneity of the different circuit layers in the hierarchical structure. This discrete scaling symmetry allows us to interpret the past causal cone of a local region as the iteration of a single quantum channel, and the translation invariance helps to reduce the number of variational parameters needed to find good circuits.

Due to the inherently discrete nature of qubits, the continuum of a field theory and its symmetry group must be discretized and broken to be represented on digital quantum hardware. Additionally the staggered structure of the individual quantum gates used to prepare the state and the fixed scaling by two for each circuit layer break translation and scale invariance of the unitary used to prepare a pure quantum state. These result in small deviations from exact translation invariance and power-law scaling of the correlation functions for the resulting ground state. These deviations can be improved by averaging over different samples of the state transformed by these different symmetries. This can be done implicitly by equally sampling over the past causal cone from different locations of the global state, or symmetrizing the measurement data by reflections or the global parity symmetry.

We find that the wavelet-based gates derived in [58] admit particularly efficient decomposition in terms of Mølmer-Sørensen (MS) gates, the native two-qubit gate set on ion trap devices, which we will sometimes denote $XX(\theta)$. This conveniently allows the implementation of these circuits on NISQ ion-trap computers with little gate overhead. MS gates can typically be applied to any pair of qubits in an ion-trap and so the qubit architecture within a trap is thought of as globally connected. Ion-trap computers engineer this all-to-all connectivity in different ways, for instance by moving individual ions around to bring pairs into interaction sites in the quantum CCD architecture [91].

$$XX(\theta) = e^{-i\frac{\theta}{2}X_iX_j} \quad (4.2)$$

The XX native interaction affords a simple implementation of our state preparation circuit, with each gate W and U being decomposable in terms of two MS gates and additional single qubit Z rotations $Z(\phi) = \exp(i\frac{\phi}{2}Z)$. All of our gates are matchgates and thus cannot form a universal gate set in a geometrically local circuit. The addition of local X or Y rotations into the circuit would expand the gate set to become universal. The non-local interaction of these gates is also crucial for implementing a hierarchical algorithm like DMERA, where new qubits need to be interleaved at each layer and would be costly to achieve with swap gates. Importantly, the non-local gates are always performed over strings of qubits in the all zero state, which means they can be performed with local matchgates and fermionic swap gates, so they can still be considered local matchgate circuits that are classically simulable. Local Z rotations are often implemented virtually via the internal clock of the hardware Hamiltonian, making them practically noise free. We considered two possible native gate decompositions of W and U , denoted $\mathcal{C}_1$ and $\mathcal{C}_2$, drawn out in Fig. 4.2.

$$\begin{aligned}
W_{\mathcal{C}_1}(\theta) &= XX(\pi/2) \cdot Z_1(\theta) \cdot Z_2(\theta - \pi/2) \cdot XX(-\pi/2) \\
U_{\mathcal{C}_1}(\theta) &= XX(\pi/2) \cdot Z_1(\theta) \cdot Z_2(\theta) \cdot XX(-\pi/2) \\
W_{\mathcal{C}_2}(\theta) &= Z_1(\pi/2) \cdot XX(-\theta) \cdot Z_1(-\pi/2) \cdot Z_2(\pi/2) \cdot XX(\theta - \pi/2) \cdot Z_2(-\pi/2) \\
U_{\mathcal{C}_2}(\theta) &= Z_1(\pi/2) \cdot XX(-\theta) \cdot Z_1(-\pi/2) \cdot Z_2(\pi/2) \cdot XX(\theta) \cdot Z_2(-\pi/2)
\end{aligned} \tag{4.3}$$

To truly prepare an approximation of the critical Ising ground state, we would need to apply the transitional layer of $XX(\pi/2)$ gates to every pair of nearest-neighbor qubits after the layers of DMERA circuit have been performed [58]. Instead, we directly target the ground state of the Hamiltonian by applying the transition layer to the usual Ising Hamiltonian H_I , resulting in $H_{I^*} = \sum_i -X_i X_{i+1} + X_{i-1} Z_i X_{i+1}$.

4.3 Noise Models

The largest single source of error is two-qubit gate infidelity, which we model by imprecise rotation angle of MS gates given by a Gaussian distribution with standard deviation σ . This type of two-qubit gate error model is commonly used for modeling noise in ion-trap devices, as the two qubit gates are typically the dominant source of noise and fluctuations in laser intensity during the execution of a MS gate can result in overrotation error [100]. This noisy version of the MS gate, denoted $\widetilde{XX}$, is equivalent to a quantum channel with Kraus operators $K_0 = \sqrt{\frac{1+e^{-\sigma^2/2}}{2}} XX(\theta)$ and $K_1 = \sqrt{\frac{1-e^{-\sigma^2/2}}{2}} XX(\theta - \pi)$, as shown in Eq. 4.4. This is equivalent to a probability of $(1 - e^{-\sigma^2/2})/2$ to flip both bits after the ideal MS gate has been applied, which for small error is $\approx \sigma^2/4$ to leading order.

This kind of noise is correlated between qubits, since the bit flips only occur in pairs, and biased towards bit flip errors rather than other two-qubit errors due to its source in gate imprecision.




Figure 4.2: Quantum circuit for causal cone of a three qubit observable using three layers of the $D = 2$ circuit ansatz starting from the all zeros state, presented in two alternate gate decompositions, $\mathcal{C}_1$ (left) and $\mathcal{C}_2$ (right), both constructed with two-qubit Mølmer-Sørensen gates and single qubit Z rotations.

$$\begin{aligned}
\widetilde{XX}_{\theta,\sigma}(\rho) &= \int_{-\infty}^{\infty} d\phi \frac{e^{-(\phi-\theta)^2/2\sigma^2}}{\sqrt{2\pi\sigma^2}} XX(\phi)\rho XX(\phi)^\dagger \\
&= \left(\frac{1+e^{-\sigma^2/2}}{2}\right) XX(\theta)\rho XX(\theta)^\dagger + \left(\frac{1-e^{-\sigma^2/2}}{2}\right) XX(\theta-\pi)\rho XX(\theta-\pi)^\dagger
\end{aligned} \tag{4.4}$$

$$\begin{aligned}
\mathcal{D}_p(\rho) &= (1-p)\rho + \frac{p}{3}(X\rho X + Y\rho Y + Z\rho Z) \\
&= \left(1 - \frac{4p}{3}\right)\rho + \frac{2p}{3}\mathbb{I}_2
\end{aligned} \tag{4.5}$$

Spontaneous logical errors from environmental noise can be modeled by periodically applying a single qubit depolarizing channel $\mathcal{D}_p$ independently to each qubit, equivalent to applying a Pauli operator randomly with small probability. This noise process is uncorrelated between the qubits and invariant under single qubit rotations. The Kraus operators of single qubit depolarizing noise, shown in Eq. 4.5, are $K_0 = \sqrt{1-p}I$, $K_i = \sqrt{p/3}\sigma_i$, for Pauli matrices σ_i . For the depolarizing noise model, a single qubit depolarizing channel with error probability $p/2$ is applied to each qubit before and after the MS gate is applied.

The depolarizing channel can be made anisotropic by biasing the probabilities for each Pauli operator. This may arise from a particular device coupling to the environment in a way that favors dephasing or bit-flip errors, while still assuming no correlation of the noise process between qubits. Ion-trap devices have relatively long coherence times and so environmental noise of this kind is expected to be smaller than gate implementation error. With these local noise models we ignore state preparation and measurement errors, and make many assumptions including that errors are homogeneous in space and time, lack long-range correlations beyond the local gate associated to the error, and are state-

independent.

The leading term in the average gate infidelity $1 - \overline{\mathcal{F}}$, where $\overline{\mathcal{F}}$ is the average gate fidelity, scales as σ^2 and p respectively for the angle imprecision and depolarizing models. These will be taken as the relevant noise parameters ε for their respective noise models. We also expect the fixed point infidelity and local observable expectation values to scale linearly in these parameters to leading order, which is confirmed by our numerics. The coefficient of this scaling is determined by specific details of the noise models and circuit implementation, including native gate sets and choices of gate decomposition. The tensor network for the channel acting on the density matrix of three qubits is given in Fig. 4.3 with the $\mathcal{C}_1$ gate decomposition in terms of native XX and Z gates.

We see in Fig. 4.4 that errors from a prior layer tend to dilute evenly among the qubits in subsequent layers. Here we apply one noisy scale transformation to a global state on four bits, followed by multiple noiseless circuit layers which each double the system size of the ground state on a circle. The global state fidelity with the reference pure state prepared by the noiseless circuit is constant as more layers are applied, and the subsystem entropy stays roughly constant. Since prior errors mix and stabilize across the system as the total size grows, the error on a constant size subsystem decreases exponentially as more layers are applied. An individual gate error introduces a perturbation to the fixed point state, which may be decomposed into the channel eigenoperators. Each operator perturbing the fixed point state decays exponentially according to its eigenvalue, which means the true convergence of an expectation value to the asymptotic value is a sum of exponential decays. In Fig. 4.5 we see that this can be fitted well by a single exponential with a rate of convergence back to the fixed point value is given by a scaling dimension of $\Delta \approx 1.89$. This convergence rate is consistent across different noise strengths in the gate imprecision model.

The stability of this channel against noise can be thought of in terms of an equilibrium reached by the noisy channel by discarding error-laden qubits and introducing new high




Figure 4.3: Tensor network diagram for the $\mathcal{C}_1$ circuit with $D = 2$ on three qubit density matrices. Two copies of the unitary circuit are given to act on each set of indices of the density matrix tensor, with initialized zero states given by one-index tensors, and the discarded qubits given by connecting the dangling legs from each circuit copy. Small red two-index tensors are also given for each MS gate which stores the weights of the Kraus operators which are summed over to implement a given local noise model for the MS gate at that position in the circuit. The colors of all the single and two-qubit gates are chosen to be identical for identical rotation angles. The choice of open indices at the bottom of the diagram make this an instance of the left channel. The left channel is averaged with the right channel, constructed by choosing right three qubits as output, to form the variational channel that implicitly averages over subregion location in the spin chain.




Figure 4.4: Average fidelity $\mathcal{F}_\ell$ of size ℓ subsystems between ideal DMERA circuit preparing L spin ground state and states where the initial layers have imprecise angles. As successive layers dilute early circuit errors subsystem fidelities remain roughly constant as a fraction of total system size, i.e. ℓ/L .




Figure 4.5: Response of the noiseless $D = 2$ channel to noisy gates, the channel with σ^2 noise is applied once to the noiseless fixed point state. The difference between the noisy energy density $\tilde{E}$ after successive applications of the noiseless channel and the fixed point energy density E^* is plotted on a log scale, showing an exponential convergence back to the fixed point energy density.




Figure 4.6: Percent error of fixed point energy densities for $D = 2, 3, 4$ channels using the angle imprecision noise model in the $\mathcal{C}_1$ (solid) and $\mathcal{C}_2$ (dashed) gate decompositions. A modified model where angle variance scales with the target rotation angle was also tried for $\mathcal{C}_2$ (dotted), but would have no effect on $\mathcal{C}_1$ due to fixed MS rotation angles.

fidelity qubits each layer. By introducing new qubits, the average number of gates a given qubit has seen remains low. By discarding qubits, a prior error may end up entirely on the discarded qubits, or in the phase of the entanglement between the discarded qubits, both of which would have no effect on observables or fidelities of the subregion we keep.

Given that an n -qubit channel has $O(nD)$ gates that each have a gate error ϵ , and the errors decay with average eigenvalue λ under action of the channel, we can expect the total error of an n -qubit fixed point to be $\delta \sim nD\epsilon + nD\epsilon\lambda + nD\epsilon\lambda^2 + \dots = nD\epsilon/(1 - \lambda)$.

We do find numerical support that the energy density does scale to leading order in σ^2 and p for our respective noise models. In Fig. 4.6, which shows the scaling with gate imprecision, we also see that the different choices of gate decomposition can greatly affect the scaling of the noise, with $\mathcal{C}_1$ scaling better than $\mathcal{C}_2$. Interestingly, we also see that the $D = 4$ channel unexpectedly scales better than the $D = 3$ for the $\mathcal{C}_1$ gate decomposition, but not for $\mathcal{C}_2$. These behaviors were not reflected in the local depolarizing noise model,




Figure 4.7: Rate of error growth in the fixed point energy density, $d\tilde{E}/d\epsilon$, versus channel depth D , using $\mathcal{C}_1$ and $\mathcal{C}_2$ circuit decompositions and the MS gate imprecision model with $\epsilon = \sigma^2$. The error growth rate for each channel depth averaged over many sets of angles with low energy fixed point, suggests a linear growth in error susceptibility with channel depth.

which had very similar scaling for different decompositions and increasing error susceptibility with depth. This could be due to the uncorrelated and local unitary invariance of the depolarizing noise, which may allow it to propagate in a more regular manner within circuits and avoid the irregularities seen in the gate imprecision model.

In [73] an alternate angle imprecision model is suggested for MS gates, where the Gaussian variance is proportional to the angle of rotation. This angular dependence of the variance would alter the noise model of $\mathcal{C}_2$ gate decomposition, but not that of $\mathcal{C}_1$ decomposition used on the ion-trap experiment where the MS gate angles are fixed. Shown dotted in Fig. 4.6, this model has effectively lower noise since the MS gates have a small rotation angle compared to the fixed angle of $\pi/2$ in the $\mathcal{C}_1$ decomposition.

Since the linear scaling of the energy density in the error parameter tends to increase for larger-depth channels, we compare $d\tilde{E}/d\epsilon$ for channels of different depth to find dependence on D of the fixed point error. To avoid potential peculiarities arising from particular

sets of angles having an atypical noise susceptibility, as seen with $\mathcal{C}_1$, we averaged $d\tilde{E}/d\epsilon$ over many sets of angles for fixed depth D in order to assess the typical behavior of channels of a given depth. The different circuits sampled, however, were chosen so that they have a fixed point energy density less than -1.2 . This constraint is needed because we are only interested in the behavior of circuits that provide reasonable ground state approximations, and $d\tilde{E}/d\epsilon$ can be negative near $\epsilon = 0$ for channels with higher energy fixed points. After averaging the noise susceptibility $d\tilde{E}/d\epsilon$ over different channels of the same depth with low energy fixed point states, we find in Fig. 4.7 that the circuit decomposition $\mathcal{C}_2$ does typically have worse scaling than $\mathcal{C}_1$ under the gate imprecision model, and the fixed point energy of both circuits seem to scale linearly with channel depth D .

4.4 Error Extrapolation

Another kind of error modeling that is important for error mitigation is a functional form to model the effect that error processes have on measurement expectation values. The error mitigation technique of zero-noise extrapolation works by assuming a functional form with certain free parameters, one crucially being the noise-free expectation value, and solving for these parameters by measuring the expectation value at different values of the noise parameters [69]. In practice the noise parameters are not able to be freely tuned, so what is often done is that certain noise sources are amplified by inverting and repeating parts of the circuit n times, $U(U^\dagger U)^n$, so that the source of noise has been applied $2n + 1$ times to a logically equivalent circuit. The repeated unitary elements could be any unitary subsection of the circuit, from individual gates to large sections of the circuit. Since there are many methods of experimentally amplifying error processes and functional forms to model the effects of errors on observables, many zero-noise extrapolation schemes are possible with varying success.

In our numerical simulations of zero-noise extrapolation, gate repetitions are implemented on individual MS gates. This could be performed everywhere in the circuit, but would triple the total MS gate count for the lowest order extrapolation. Since the DMERA channel acts to dissipate perturbations from the fixed point state, errors from earlier layers of the circuit are exponentially suppressed as seen in Fig. 4.5. To reduce the overhead of inserting extra gates into the entire circuit getting the most of error extrapolation we can focus on amplifying and subtracting noise from the final layers of the circuit which have a greater impact on the measured expectation value.

One way to model noise on an observable is as a polynomial perturbation. The measured energy density $\tilde{E}$ is the noiseless asymptotic value E^* , plus a polynomial perturbation in the unknown error parameters ε_i for some independent sources of error, which we assume contribute independently without interaction to the measured energy $\tilde{E}$. To fit a polynomial with m parameters, a collection of at least m appropriate noise amplification measurements are needed in addition to the unamplified measurement to solve for all the coefficients.

Aggregating the error from each channel application into a single source, we expect the error contribution ε from each application to be identical, but with error from earlier stages of the circuit reduced by the scaling parameter λ . This suggests a polynomial model of the error with three free parameters:

$$\tilde{E} = E^* + \varepsilon + \varepsilon\lambda + \varepsilon\lambda^2 + \dots = E^* + \frac{\varepsilon}{1-\lambda} \quad (4.6)$$

An alternate functional form for the error is a product of exponentials applied to the target energy. Following the same logic, we get an analogous functional form for the noisy energy as an exponential with three free parameters.

$$\tilde{E} = E^* e^{-\varepsilon} e^{-\varepsilon\lambda} e^{-\varepsilon\lambda^2} \dots = E^* e^{\frac{-\varepsilon}{1-\lambda}} \quad (4.7)$$

This multiplicative series model is equivalent to the additive series ansatz when expanding the exponential to first order, since $E^*(1 - \frac{\epsilon}{1-\lambda} + \dots) \approx E^* - \frac{E^*\epsilon}{1-\lambda} = E^* + \frac{\epsilon'}{1-\lambda}$, and also has the correct behavior for large ϵ where the energy density should approach zero.

These model parameters can be fit by amplifying errors in the last and second to last layers of the circuit independently. While this error mitigation technique requires measuring three different versions of the circuit, the gate count overhead for the amplified circuits is not large compared to amplifying the entire circuit. Numerically the scheme also performs much better than full linear extrapolation applied by amplifying individual gates in the entire circuit (Fig. 4.8). The values of observables obtained through zero-noise extrapolation is not guaranteed to lie between the noisy measurements and the true noiseless value, so the extrapolated energies could lie below the noiseless fixed point values and may not be a variational upper bound to the ground state energy. In the examples considered in our numerics, however, no extrapolation schemes resulted in energies below the noiseless fixed point values.

4.5 Experimental Data

Comparison to other experiments: Previous experiments using NISQ era ion-traps have prepared critical ground states and thermofield double states for the quantum transverse-field Ising model for small systems using the QAOA variational ansatz [73]. It is known that the critical ground state on L qubits can be prepared exactly with a depth $L/2$ QAOA circuit [72]. Since the critical model has polynomially decaying correlations, it is expected that a geometrically local QAOA ansatz would need depth scaling linearly with the system size in order to prepare a good approximation to the ground state. However, with non-local gates it could be possible to lower the required depth to prepare a high fidelity ground state approximation. Using the DMERA ansatz it is possible to prepare polynomially scaling




Figure 4.8: Percent error of fixed point energy densities for $D = 2$ and $D = 3$ channels using the angle imprecision error model and the $\mathcal{C}_1$ gate decomposition. Unmitigated fixed point energies (solid) are considerably improved through zero-noise extrapolation of MS gates using a geometric series error ansatz with additive(dot-dash) or multiplicative(dash) error, both of which outperform standard linear extrapolation of gates in the entire circuit (dotted) that requires much larger gate overhead.

correlations with only a logarithmic-depth circuit by utilizing non-local gates. While the size of our experiment is similarly limited, the DMERA state preparation is better thought of as a finite subsystem of the thermodynamic ground state, while the QOAO protocol prepares the exact ground state of a small sized system.

We run our experiment on the Honeywell HØ ion-trap quantum computer, a predecessor to the model H1 and H2 devices now available through Quantinuum. This device has six qubits with global connectivity achieved by bussing the pair of ions to an interaction site when applying the two-qubit gate. Each application of the n -qubit channel requires twice as many to implement, so we are limited to three qubit channels but we can apply the channel arbitrary number of times with mid-circuit qubit reset.

The native MS gates of the ion-trap are also given in terms of ZZ rotations, rather than XX rotations as used earlier in the paper. Rather than sandwich each MS gate with explicit Hadamard gates to transform them to the other convention, we will implicitly transform the circuits and states to the convention native to the ion-trap by interchanging X and Z operators in the experimental implementation. This means our new Hamiltonian will be $H_{I^*} = \sum_i -Z_i Z_{i+1} + Z_{i-1} X_i Z_{i+1}$ and new ancilla will be prepared in the state $|+\rangle$ via a Hadamard transformation on the reset qubits.

To see the effectiveness of the state preparation algorithm and stability against noise, we measure the expectation value of local observables as a function of the number of iterations of our three qubit channel. With limited resources, we prioritized sampling as many different numbers of channel iterations as possible to see convergence and stability, rather than repeating individual measurements many times to get better measurement precision. The longest circuits measured had twelve layers of the DMERA channel, each having 120 MS gates, 156 single qubit gates, and 33 qubits recycled mid-circuit, and the average number of samples per measurement was 200.

For initial states we choose products of single qubit states $|\psi_1\rangle$ and $|\psi_2\rangle$ since they can

be prepared with high fidelity. The product of the first state $|\psi_1\rangle$ is one of the two mean-field ground states, i.e. a pure product state with the lowest energy, which will actually be below the energy of our noisy fixed point. The second state $|\psi_2\rangle$ is orthogonal to $|\psi_1\rangle$ and serves as a higher energy initial state which will decrease in energy under applications of the noisy channel. These initial states also have opposite signs for the expectations of local X and Z observables which we see converging towards their fixed point values experimentally.

Each iteration of the channel follows the circuit / tensor network diagram of Fig. 4.3. For the initial layer, qubits 1,3, and 5 are all initialized to either $|\psi_1\rangle$ or $|\psi_2\rangle$ while qubits 2, 4, and 6 act as the new ancilla and are initialized in the $|+\rangle$ state. The native ion-trap gates are applied in a brickwork-like circuit according to Fig. 4.3. Qubits 1 and 6 are reset to the $|+\rangle$ state to be used as the ancilla for the next channel iteration. We must also choose whether to reset qubit 2 or qubit 5 to be used as the third ancilla qubit, which correspond with choosing the right or left subregions to continue to the next iteration. So, if we choose the left subregion we will have reset qubits 1, 5, 6, which will become the new ancilla qubits in positions 2, 4, and 6, and the unreset qubits 2, 3, and 4 will take the position of qubits 1, 3, and 5. If we instead choose to propagate the right subregion then we will reset qubits 1, 2, and 6, and the qubits 3-5 will take the position of qubits 1, 3, 5 for the next layer.

This process is iterated for however many circuit layers we wish to apply before measurement. After all iterations are completed, qubits 2-5 are left as the local subregion on four qubits of our approximate ground state. We can then measure these four qubits in the Z basis in order to estimate expectation values of local Z_i observables and the terms $Z_i Z_{i+1}$ that appear in the Hamiltonian H_I^* . Alternatively we can apply the Hadamard to 3 or 4 in order to measure the expectation value of these local observables X_i and the terms $Z_{i-1} X_i Z_{i+1}$ that appear in the Hamiltonian.

For an exhaustive experimental implementation we would want to randomize the choice




Figure 4.9: Experimental convergence of energy density from initial product states $|\psi_1\rangle$ and $|\psi_2\rangle$. Data from the HØ ion-trap computer plotted against simulation curves of noiseless(dashed) and best fit noisy(solid) circuits.

of left or right subregions to propagate to the next layer, which corresponds to choosing various local subregions of the global ansatz state. However, each of these choices would result in a different quantum circuit to be compiled and submitted to hardware, which is costly. Instead, we choose to alternate between left and right subregions as we iterate the channels and take data from this particular circuit choice.

$$\begin{aligned}
 |\psi_1\rangle &= \sqrt{\frac{3+2\sqrt{2}}{6}}|0\rangle - \sqrt{\frac{3-2\sqrt{2}}{6}}|1\rangle & E_1/L &= -\frac{32}{27} \\
 |\psi_2\rangle &= \sqrt{\frac{3-2\sqrt{2}}{6}}|0\rangle + \sqrt{\frac{3+2\sqrt{2}}{6}}|1\rangle & E_2/L &= -\frac{16}{27}
 \end{aligned} \tag{4.8}$$

While there are some outlying data points we find general agreement with the convergence of local observables to somewhat stable values after multiple iterations of our quantum channel. We also see that the convergence rate can vary greatly depending on the

exact observable, with the Z expectation converging much slower than either X or Hamiltonian terms ZZ and ZXZ . We include in the experimental data for the expectation values of energy density (Fig. 4.9) and local observables (Fig.4.10) the simulated noiseless channel (dashed) and the best fit value of the noise parameter σ^2 for the angle imprecision model, which we found to be $\sigma^2 = 0.112$. This noise model corresponds to a fixed point energy density of -1.08 , 15.2% higher than the true ground state energy density of $-4/\pi$.

The local observables considered are Pauli strings, therefore they are traceless and would have an expectation value of zero for the maximally mixed state. Strong levels of noise in the error models considered drive the fixed-point expectation values towards zero for these observables, but for moderate levels of noise the expectations are stabilized at nonzero values. Experimentally we can clearly see the effects of noise in real quantum hardware but also see that after many iterations of the local channel these expectation values are robustly kept away from zero for those observables which have nonzero expectation in the ground state. This convergence to a robust value is expected to improve as hardware capabilities are refined, and together with more advanced and specially tailored error mitigation techniques can provide a more reliable way of measuring local observables on NISQ devices.




Figure 4.10: Convergence of local expectation values, for single qubit X and Z measurements (top) and multi-spin Hamiltonian terms ZZX and ZZ (bottom). Experimental datapoints from the H $\emptyset$ ion-trap computer plotted against noiseless and best fit noisy simulation curves, for initial product states $|\psi_1\rangle$ and $|\psi_2\rangle$.

Chapter 5: TMERA

Quantum computers promise efficient simulation of quantum physics [101]. Nearly any quantum simulation task requires the preparation of some physically motivated state to be measured or used as the initial state of a dynamical simulation. States in thermal equilibrium, although dynamically trivial, are of general importance to many physical problems.

But preparing Gibbs states on a quantum computer is in general not easy. The problem is NP complete, because computational problems can be encoded in statistical mechanics models, and QMA complete because in the low-temperature limit the Gibbs state becomes the ground state [102–104]. Indeed any algorithm preparing Gibbs states [37, 73, 105–128] is only efficient for a restrictive subclass of problems (e.g. Hamiltonians with local commuting terms [112] or short-range correlated states [113]), and many face challenges like barren plateaus, complicated energy or entropy measurement schemes, or reliance on oracles that themselves require expensive implementations.

We seek to prepare Gibbs states on current and near-future quantum computers. These computers will be NISQ (noisy intermediate-scale quantum [16]) devices, limited in the number of qubits available and the fidelity with which quantum operations can be performed on those qubits. One can avoid accumulation of errors on these devices by working within a family of short-depth circuits—an ansatz for the state to be prepared. Parameters for a circuit preparing the target state can then be found through a process of hybrid quantum-classical variational optimization, in which quantum measurements are used by a classical optimization algorithm to determine the optimal circuit parameters. For ex-




Figure 5.1: The TMERA circuit: a unitary DMERA quantum circuit optimized for ground state preparation (green) is applied to a product of mixed qubit states (colored dots). The circuit ansatz separates energies by scale, with low energy qubits introduced in very mixed/hot states (red) and high energy qubits introduced in relatively pure/cold states (blue).




Figure 5.2: The global state energy and entropy (left) and state fidelity (right) for an $L = 512$ spin system show that the ansatz is able to prepare high fidelity Gibbs state approximations for the critical Ising chain at a variety of temperatures. The ansatz improves with the circuit depth D of the scale transformations.

ample, to prepare the ground state of a Hamiltonian one would measure the energy for different circuit parameters in order to classically estimate energy gradients and perform gradient descent on the circuit parameters. For ground states this approach is known as a *variational quantum eigensolver* (VQE) [38, 39, 62] and has been demonstrated for small systems on existing quantum computing devices [38, 72, 73]. More generally, *variational quantum algorithms* (VQA) can be used for other types of near-term applications, where a small ansatz circuit is optimized to prepare a desired state [37]. Another example is the quantum approximate optimization algorithm (QAOA) [71, 129] which has been applied to combinatorial optimization problems like MaxCut. For Gibbs states we can instead use the free energy as the scalar quantity to be optimized [116], however, this can prove challenging due to the need to know the entropy of the state, which is not a directly measurable observable.

We suggest that for many practical applications, the Gibbs state preparation problem becomes tractable when we take advantage of *scale separation*—the fact that in physical Hamiltonians length scales are associated with energy scales. The separation of degrees of freedom at different length scales is made manifest in multi-scale quantum circuits. This motivates our variational ansatz for Gibbs states, which we call the *thermal multi-scale renormalization ansatz* (TMERA).

Multiscale circuits like MERAs or DMERAs progressively enlarge the Hilbert space by tensoring on new qubits in the $|0\rangle$ state. This progressive enlargement implements a fine-graining scaling transformation. The TMERA instead progressively enlarges the Hilbert space by tensoring on new qubits in mixed states with entropy controlled by the scale at which the new qubits are introduced, seen in Fig. 5.1. The mixed state qubits can be prepared by independently sampling computational basis states on each qubit, with a scale-dependent sampling bias that sets the entropy at a given energy range. Since each qubit is introduced at a given length scale, and therefore energy scale, a probability can

be associated to it for a target temperature. Short length scale (UV) degrees of freedom have high energy, hence lower excitation probability, while long length scale (IR) modes have low energy, hence higher excitation probability. As an example of a product spectrum ansatz for preparing Gibbs states [117], our work further refines the ansatz to the case of multi-scale circuits, where the separation of energy scales and approximate translation invariance can be incorporated to improve expressability of the ansatz beyond what would be possible with a strictly local one dimensional circuit.

A product spectrum ansatz allows us to know the entropy of the final state a priori, since the unitary quantum circuit does not change the entropy of the initial product state. This avoids the pitfall of having to estimate the entropy via quantum measurements which can be practically prohibitive.

Despite the limitations of using a low-depth variational circuit with a product state spectrum, we find that high fidelity approximations to Gibbs states can be found for large systems at various temperatures using few variational parameters. We benchmark our ansatz with the transverse-field Ising model in one dimension at its critical point. Scaling and translation symmetry of this model strongly constrains the system’s Hamiltonian and correlations, and therefore can greatly simplify the underlying quantum circuit ansatz, requiring only a constant number of variational parameters for the DMERA ground state circuit. Scale-invariance of critical models also constrains the energy density to scale according to a fixed critical exponent. This reduces the ansatz for excitation probabilities to two variational parameters which set the critical exponent and the temperature. However, the energy associated to each scale of the DMERA may also be measured experimentally, which effectively solves the Gibbs state problem in the context of our ansatz by fixing a single parameter family of states which approximate Gibbs states of all temperatures. For the critical Ising model we find that the TMERA ansatz gives a global fidelity > 0.5 on 256 qubits using a log depth quantum circuit.

We structure the paper as follows. In Sec. 5.1 we describe our guiding intuition and present the TMERA circuit ansatz. In Sec. 5.2 we give the model we use for benchmarking (the transverse field Ising model), in Sec. 5.3 we discuss properties of the TMERA, including the energy and entropy measurements required for variational optimization of the free energy, and in Sec. 5.4 we benchmark our model by measuring its fidelity with the true thermal state, its entropy as a function of energy, and other physical quantities.

5.1 Background, intuition, and circuit ansatz

5.1.1 Background and intuition

The art of a successful variational algorithm is finding a circuit ansatz that is expressive enough to represent a state close to the target state, but simple enough to be effectively optimized and implemented given the limitations of existing quantum technology. Physical intuition can guide us in designing a circuit ansatz for physical problems—for example, a low-depth local brickwork circuit will leave subsystems separated by more than the circuit depth unentangled, and so the circuit would be best suited for states without long-range entanglement such as short-range correlated states which obey an area law for subsystem entanglement entropy. If the state is translation invariant, the circuit may likewise be constrained to be spatially homogeneous or periodic in order to reduce the ansatz complexity.

If the model is also scale invariant, perhaps because it sits at a quantum critical point, we should construct an ansatz that is likewise scale-invariant. The *multi-scale entanglement renormalization ansatz* (MERA) is such a scale-invariant ansatz, initially developed to target critical ground states using tensor networks [56, 130]. The most basic MERA prepares the state in a scale-by-scale manner, preparing a sequence of states $\{|\psi_\ell\rangle\}_\ell$, each on 2^ℓ qubits. The $|\psi_\ell\rangle$ can be viewed as a coarse-grained version of the target state. The MERA, then, consists of a series of identical “fine-graining” layers, known as a *scaling*

transformation. Each layer maps an incoming Hilbert space of $2^{\ell-1}$ qubits to an outgoing Hilbert space of 2^ℓ qubits; it does this by applying an isometry to each qubit that embeds that qubit's Hilbert space in a larger Hilbert space, and then applying unitaries between neighboring qubits to tune local correlations.

A more expressive variant of the MERA can be constructed by replacing the unitaries and isometries with low-depth quantum circuits. This circuit ansatz is known as a *deep MERA* (DMERA) [61]. Like a MERA, a DMERA is built from successive scaling transformations. But now the scaling transformation has two steps: first, it adds ℓ new qubits interleaved between the ℓ qubits resulting from the previous layer and initializes the new qubits in the $|0\rangle$ state; second, it applies a brickwork unitary circuit on the 2^ℓ qubits, entangling new with old. Formally the scale transformation is

$$|\psi_{\ell+1}\rangle = U_\ell(\theta) \left(|\psi_\ell\rangle \otimes |0\rangle^{\otimes 2^\ell} \right). \quad (5.1)$$

Fig. 5.3 illustrates a closely related circuit (our TMERA, to be introduced below); were that circuit a DMERA, all of the colored balls would represent the state $|0\rangle$.

The DMERA is approximately scale invariant because we require each of the scaling transformations to have identical parameters. The total number of two-qubit gates for preparing an L qubit DMERA state is $O(DL)$, with D the depth of the brickwork unitaries U_ℓ , and so the number of variational parameters could scale at most as $O(DL)$. Each two-qubit gate is an element of $SU(4)$ and therefore has at most 15 needed to parameterize the local gates. Imposing translation symmetry within each layer reduces the number of independent gates down to $D \log_2(L)$, i.e. the total circuit depth, which constrains the number of variational parameters to be at most $15D \log_2(L)$. Further imposing scaling symmetry by identifying the gates within different scaling transformations U_ℓ reduces the number of independent gates to D , with at most $15D$ variational parameters. We further restrict our




Figure 5.3: Circuit diagram of the TMERA showing three scale transformations with $D = 4$ on a subregion of qubits in a one dimensional lattice. Time moves from top to bottom as the scale goes from IR to UV. The gates are chosen to be identical in the different scale transformations to impose a discrete scale symmetry, and by translations within each circuit U_{ℓ} to impose approximate translation symmetry. Individual qubits are introduced in mixed states $\sigma_{\beta,\ell}$ (colored dots), with the degree of mixture determined by the target temperature and energy scale E_{ℓ} of the given circuit layer.

gates to be real-valued and parity-preserving, which reduces the number of parameters for each gate to two.

Turn now to variational preparation of a thermal state. Our *thermal multi-scale entanglement renormalization ansatz* (TMERA) prepares a mixed state by replacing the input states $|0\rangle$ in the DMERA scaling transformation with mixed product states $\sigma_{\beta,\ell}$; the states $\sigma_{\beta,\ell}$ depending on system's inverse temperature β and the scale ℓ at which the qubit is introduced.

To understand how this changes the state, imagine flipping a single one of the input states $|0\rangle$ of Eq. 5.1 at a single layer ℓ to $|1\rangle$. Once this change propagates to the physical qubits, the result is an wavelet excitation at a length scale $L/2^{\ell}$, where L is the total length

of the system. The excitation introduced at layer ℓ therefore has an energy

$$E_\ell \sim (2^\ell/L)^z \quad (5.2)$$

where z is the model's dynamical critical exponent.

The DMERA circuit propagates excitations local on the input qubits to wavepacket excitations in the final state, which are localized in both space and energy. The most energy-localized basis for a translation invariant system would be a plane-wave basis, i.e. the momentum modes of a free field theory, but this transformation from a local to plane-wave basis would require circuit depth linear in the system size to transform local modes to completely delocalized momentum modes. This is seen in Ref. [72], where the exact ground state can be prepared by a local circuit only once the circuit depth is at least $L/2$. This may also be expected because a local circuit must have depth $L/2$ for the future causal cone of any initial qubit to span all of the final qubits, which would be needed for the circuit to "see" the whole state and find a global minimum of the energy [74]. The DMERA circuit achieves a happy medium by being somewhat localized in energy, but still localized enough in space to be of short depth. The DMERA circuit can be viewed as implementing a kind of many-body quantum wavelet transform analogous to the Fourier transform between local and plane-wave bases. Wavelet transforms are used in classical signal processing for time-frequency analysis where localization in both time and frequency are required.

The energy localization of the wavelet basis allows us to use the local input qubits as a good basis to use for approximately filling out the distribution of our target Gibbs state and motivates the use of a product spectrum ansatz. Since initial scales are associated with low energy excitations, they have the largest probability of being excited in a thermal state, with successive layers having smaller probabilities dictated by the energy localization of that scale. All qubits introduced at a given scale will have the same excitation probability

because they are all related by translations which leave the circuit invariant below that scale.

Fig. 5.4 is meant to illustrate that degrees of freedom introduced at different scales within the circuit are supported locally in the space of energy eigenmodes when using a one-particle description. To compute the quantity we plot, we find the energy modes e_α for our quadratic fermion Hamiltonian and take the squared inner product with the local modes introduced at scale ℓ which become wavepacket modes for our final state. For each energy mode we sum over all local modes i within a fixed scale ℓ to find how much each energy eigenmode is supported in that scale. The resulting plot shows that modes from different scales ℓ are each supported on a compact range of energy eigenmodes.

Finding the characteristic energy E_ℓ for modes introduced at each scale then gives us a prescription for how to introduce entropy at different scales via the mixed states $\sigma_\ell(\beta)$. The scale dependence of $\sigma_\ell(\beta)$ is also the only way that our ansatz breaks scale symmetry due to the energy scale set by the target temperature.

5.1.2 Circuit ansatz

Precisely, then, the thermal multi-scale entanglement renormalization ansatz (TMERA) prepares an approximate thermal state $\rho_{\beta,N}$ on $L = 2^N$ qubits at inverse temperature β in a scale-by-scale manner using a sequence of scale transformations $\mathcal{S}_\ell$. Each scale transformation $\mathcal{S}_\ell$ prepares the state $\rho_{\beta,\ell}$ from $\rho_{\beta,\ell-1}$ by interleaving $2^{\ell-1}$ new qubits, each in the mixed state $\sigma_{\beta,\ell}$, and then applying the unitary circuit $U_\ell(\theta)$. The scale transformation is therefore

$$\begin{aligned} \rho_{\beta,\ell} &\equiv \mathcal{S}_{\beta,\ell}[\rho_{\beta,\ell-1}] \\ &= U_\ell(\theta)(\rho_{\beta,\ell-1} \otimes \sigma_{\beta,\ell-1}^{\otimes 2^{\ell-1}})U_\ell^\dagger(\theta). \end{aligned} \tag{5.3}$$

We illustrate three applications of the scale transformation on subregion of the lattice in Fig. 5.3. The balls represent $\sigma_{\beta,\ell}$, and the rounded rectangles the gates of the brickwork circuit U_ℓ .

The mixed state $\sigma_{\beta,\ell}$ is a Gibbs state for a Hamiltonian Pauli matrix Z , chosen so that we have a mixture between the occupied and unoccupied state of the spin or fermion. The resulting input state when taking all initialized qubits into account will be a classical state, a probability distribution over computational basis states.

$$\sigma_{\beta,\ell} = \frac{\text{sech}(\beta E_\ell)}{2} e^{\beta E_\ell Z}. \quad (5.4)$$

These single qubit density matrices may be prepared stochastically by taking an ensemble over initial computational basis states, or by taking one qubit from a pair which are entangled to produce $\sigma_{\beta,\ell}$ as the reduced state of that qubit. While the classical ensemble may seem more simple, it would require sampling many different circuits which only differ by some bit flips applied to initialized qubits. Depending on the exact ensemble of computational states it may be unlikely for any particular circuit to be sampled more than once, which would become costly when having to compile many different circuits. For this reason it may prove advantageous to instead prepare mixed qubit states via pairwise entanglement method, which would only require a single circuit to be compiled and sampled. With mid-circuit reset the unused qubit of the entangled pair may be reset and reused to initialize further mixed qubit states.

We imagine several different variational approaches to minimize the free energy for a particular target temperature depending on what is already known about the model. We also assume here that the unitary circuit has already been optimized for ground state preparation, which induces the appropriate separation of energy scales. The parameters E_ℓ are meant to represent the average energies of excitations introduced at scale ℓ , which would propagate

through the circuit to form wave packets localized to $O(L/2^\ell)$ spatial sites. The parameters E_ℓ also represent the window of energy modes that these wave packets are supported on when the one-particle description is available, as shown in Fig. 5.4. For a critical system, we expect E_ℓ to scale with a critical z as

$$E_\ell = E_{UV}(2^\ell/L)^z. \quad (5.5)$$

If z is not already known then we may treat it as a variational parameter in addition to E_{UV} which will fix the UV energy scale. Alternatively, and certainly if the underlying system is not at a critical point, the energy scales E_ℓ can each be taken as independent variational parameters. For the Ising model we know that the scaling exponent should be $z = 1$, but further verify that this scaling is indeed the optimal one by conducting an independent optimization of the energies E_ℓ which recovers the expected scaling. One may also allow for further optimization of the circuit parameters when minimizing the free energy, but we have not explored that option in the current work.

As we expect the energies associated with short length scales to be larger than those associated at long length scales, it should always be true that $E_{\ell+1} > E_\ell$. At very low temperatures we can expect most qubits to be almost pure, except for those at the longest length scales with the lowest associated energy. Similarly, at very high temperatures we can expect that most qubits would be nearly maximally mixed, except for those at short length scales with the largest associated energy. So, for these extreme temperatures only a few of these energies will be practically relevant, while intermediate temperatures will rely more on the full range of energy scales. Since the state will remain near maximally mixed until the final layers of the circuit for high temperature target states, any prior layers would be unnecessary as we can instead apply those final layers to the maximally mixed state for a nearly identical outcome.

In fact, we find that for any temperature at most five different circuit layers will introduce qubits with entropy more than 0.01 away from either zero or $\log(2)$ with a TMERA ansatz when using energies scaling as in Eqn. (5.5) with $z = 1$. This means if 1% error in the individual qubit entropy is acceptable, then all except five circuit layers can be prepared as pure states or maximally mixed states. An especially crude version of our ansatz could take every qubit to be either pure or maximally mixed. In this scenario, we would simply implement the ground state preparation circuit but starting at some intermediate layer ℓ using the maximally mixed state on $2^{\ell-1}$ qubits as our initial state. This version of the ansatz would essentially be "cutting off" the ground state circuit above some temperature induced length scale.

Our ansatz also naturally gives an algorithm for preparing approximations to the purification of the thermal state known as the thermofield double (TFD) state. The TFD state is a pure state on two copies of the system, where the reduced density matrix on either copy looks like a Gibbs state at a given temperature. In our case, the qubits introduce in mixed product states for one copy would instead become pairs of entangled qubits for preparing an approximate TFD state. Identical unitary circuits would then be applied to each collection of entangled qubits so that two copies of the approximate Gibbs state are prepared which purify each other. The circuit diagram in Fig. 5.1 can be doubly interpreted as applying a unitary circuit to a product state density matrix as is the case for Gibbs state preparation, or applying two unitary circuits to each half of a product of entangled pairs as in the case for TFD state preparation.




Figure 5.4: The support of the first-quantized wavepacket modes $|\ell, i\rangle$ introduced at layer ℓ with the energy eigenmodes $|e_\alpha\rangle$. The inner product is summed over all 2^ℓ from each layer, showing the support for each scale is localized in energy, calculated for system size $L = 512$.

5.2 Model

We benchmark the TMERA on the one dimensional critical transverse-field Ising model

$$H_I = -\sum_j^L \sigma_j^x \sigma_{j+1}^x + \sigma_j^z. \quad (5.6)$$

At low energy this spin chain is known to be well described by a conformal field theory with central charge equal to $1/2$ [131].

The transverse field Ising model is free fermion integrable. The Jordan-Wigner transformation maps the Hamiltonian (5.6) to the quadratic Majorana fermion Hamiltonian

$$H_f = i \sum_{j=1}^{2L} \gamma_j \gamma_{j+1}. \quad (5.7)$$

Although spin chain models like (5.6) are natural for digital quantum computers, free

fermion models like (5.7) are more natural for classical computers. The ground state of (5.7) is well approximated by a fermionic version of a DMERA circuit, as shown in Refs. [58, 59]. We will likewise be considering a fermionic version of the TMERA algorithm in this work in order to classically simulate using free-fermion numerical methods (see App. A for further details). Many states of interest—including the Gibbs states with which we are concerned—are Gaussian fermion states. They are therefore (at least in principle) preparable by matchgate circuits, which are equivalent to unitary Gaussian operations. Matchgate circuits consist of local two-qubit gates in one dimension, where the gates are restricted to be independent $SU(2)$ operations on the even and odd parity subspaces. These gates are equivalent to rotations generated by the operators $X_i X_{i+1}$, $Y_i Y_{i+1}$, or Z_i , each of which is equivalent to a quadratic fermion operator via Jordan-Wigner duality. Consequently, quantum spin chains with Hamiltonians that are a sum of only these operators are integrable by the equivalence to a quadratic fermion Hamiltonian, and have energy eigenstates which are Gaussian fermion states.

Covariance matrix techniques can simulate large matchgate circuits efficiently [132]. For models that are not free fermion, the target Gibbs state will no longer be Gaussian, and the ansatz circuit must include gates outside the set of matchgates. This is not an obstacle for a quantum computer, but it is an obstacle for a classical computer—which is why we restrict ourselves to the free-fermion Hamiltonian (5.7).

Matchgates map single-fermion states to single-fermion states. Since the DMERA (and TMERA) for a free fermion model like (5.7) consists entirely of matchgates, we can rephrase the wavelet excitation picture of Sec. 5.1 even more explicitly. The matchgate DMERA circuits transform local fermion modes introduced at scale ℓ into wavepacket modes with support over 2^ℓ lattice sites in the final state. This linear transformation on the modes is equivalent to a discrete wavelet transform. A Gaussian fermion state may be specified by a set of fermion operators which annihilate it, and any linear combination of

these operators also annihilate the ground state. A Gaussian DMERA circuit ansatz may be understood as preparing the state annihilated by these wave packet modes introduced at each scale, which span approximately the same subspace of operators as the actual ground state annihilation operators. Analogously to how periodic functions may have a simple representation in a Fourier basis, the success of a DMERA circuit ansatz lies in the ability to find a good wavelet basis that admits a simple approximation of the target ground state with acceptable fidelity.

While this clear picture in terms of free fermion modes breaks down for non-integrable models, we may still think of the DMERA circuit ansatz with a universal gate set as a more general kind of many-body wavelet transformation. This ansatz will still have a product spectrum, but may still be successful as long as the target state has an entanglement structure that admits this scale-by-scale coupling of degrees of freedom.

The underlying DMERA circuits for ground state preparation that we use are from [41], where circuit parameters were optimized using the expectation value of the energy density. The energy density was optimized using the fixed-point density matrix of the local quantum channel, which implements a DMERA scaling circuit on a constant number of qubits, effectively optimizing the average energy density of the infinite volume ansatz state. These parameters are then used to prepare the ground state circuit of a finite sized system using a finite number of circuit layers.

This parameterized family of circuits is a modest generalization of the circuits proposed in [58] to prepare the Ising model ground state, with each local two-qubit gate being the real-valued matchgate $u(x,y)$ that implements independent rotations of the even- and odd-




Figure 5.5: Relative error in the free energy $\delta F/|F|$ for TMERA states for temperatures T and circuit depth D , system size $L = 512$.

parity subspaces

$$u(x,y) = \begin{bmatrix} \cos(x) & 0 & 0 & \sin(x) \\ 0 & \cos(y) & \sin(y) & 0 \\ 0 & -\sin(y) & \cos(y) & 0 \\ -\sin(x) & 0 & 0 & \cos(x) \end{bmatrix}. \quad (5.8)$$

So, the DMERA circuits have $2D$ parameters in total after imposing constraints for approximate scale and translation invariance expected in a critical ground state.

5.3 Free energy minimization and properties of the TMERA

Variational optimization of the Gibbs state at temperature T requires minimizing the free energy

$$F = E - TS \quad (5.9)$$

over our variational parameters. This in turn requires measuring energy and entropy of each Gibbs state. While the energy of the ansatz state must be measured from the expectation values of local terms of the Hamiltonian, the entropy can be inferred directly from the excitation probabilities of wavelet modes.

5.3.1 Entropy

A TMERA prepares a mixed state by repeated application of the scaling transformation

$$\mathcal{S}_{\beta,\ell}[\rho_{\beta,\ell-1}] = U_\ell(\boldsymbol{\theta})(\rho_{\beta,\ell-1} \otimes \sigma_{\beta,\ell-1}^{\otimes 2^{\ell-1}})U_\ell^\dagger(\boldsymbol{\theta}). \quad (5.10)$$

Consequently the final mixed state is—up to a unitary—a tensor product of the mode mixed states $\sigma_{\beta,\ell-1}$, and the entropy is the sum of the mode entropies. Because each scale introduces $2^{\ell-1}$ modes, the total entropy is

$$S = S_1 + \sum_{\ell=1}^N 2^{\ell-1} S_\ell, \quad (5.11)$$

where

$$\begin{aligned} S_\ell &= -p_\ell \log(p_\ell) - (1 - p_\ell) \log(1 - p_\ell) \\ &= \log(\cosh(\beta E_\ell)) - \beta E_\ell \tanh(\beta E_\ell) + \log 2, \end{aligned} \quad (5.12)$$

is the entropy of a single qubit in the state $\rho = Z^{-1} e^{-\beta E_\ell}$. The extra S_1 term in the total entropy S is required because the first layer introduces two qubits rather than just one, for a total of $L = 2^N$ contributions to the entropy from single qubit mixed states.

In variationally optimizing the free energy (5.9), one would explicitly evaluate the full expressions (5.11) and (5.12). But we can get physical insight into the origin of the state's entropy by considering the UV and IR limits of (5.11) and (5.12). In those limits, the

entropy of each mode is

$$S_\ell \approx \log(2) - \frac{1}{2}(\beta E_\ell)^2 \quad \beta E_\ell \ll 1 \text{ (IR)} \quad (5.13a)$$

$$S_\ell \approx 2\beta E_\ell e^{-2\beta E_\ell} \quad \beta E_\ell \gg 1 \text{ (UV)}; \quad (5.13b)$$

and the energy is $E_\ell \propto 2^{\ell z}$. Comparing these entropies with the total entropy of Eq. (5.11), we see a competition between the number of modes and the entropy per mode. In the UV limit the entropy per mode may be exponentially small, but there are exponentially many modes. In the IR limit, by contrast the individual mode entropies approach their maximal value of $\log(2)$, but there are $O(1)$ modes. In Fig. 5.6 we show the total entropy contribution from each scale at a variety of temperatures. For high temperatures all modes are close to maximally mixed, so the entropy contribution follows the exponential decay of the number of modes at each scale. For low temperatures, however, the UV modes are relatively unexcited and do not contribute much entropy despite the large number of them, and the bulk of the total entropy comes from scales with energy low enough to have non-negligible excitation probability. A given scale ℓ will contribute more entropy than the shorter length scale $\ell + 1$ when $S_\ell/2 > S_{\ell+1}$, which occurs when $\beta E_\ell \approx 1.26$.

5.3.2 Energy

Since the system is translation-invariant, the energy can be computed by measuring the energy density on a single bond. Such a measurement does not require the full DMERA: rather, it requires the past causal cone of the single bond. The past causal cone of a subsystem converges to a width independent of subsystem size or total system size, determined only by the scaling transformation depth D . With the ability to reset and reuse qubits, at least $4D - 2$ qubits are needed to implement the this local depth D channel that prepares a subregion of the state.




Figure 5.6: Entropy introduced at each scale. At high temperature most of the entropy comes from the many UV modes near infinite temperature. At low temperature a non-monotonicity results from competition between the number modes at each scale and the entropy per mode, for system size $L = 512$.

We are benchmarking on a free-fermion integrable model with a matchgate DMERA, so the entropy-maximizing energy is just the mode expectation value of the physical Hamiltonian. In Fig. 5.7 we plot those mode energies, and we see that the mode energies broadly follow the expected scaling $E_\ell \propto 2^{\ell z}$ with $z = 1$. They depart from that scaling at the lowest ℓ (IR-most modes), where finite size effects become important, and at the highest ℓ (UV-most modes), where lattice effects become important.

The energy introduced by flipping one of the initialized qubits, in effect exciting some wavepacket mode localized to a length scale of $L/2^\ell$, exactly coincides with the optimal values of E_ℓ when minimizing the free energy. This implies that it would be possible to measure the optimal values of E_ℓ experimentally for each scale, in effect requiring zero additional variational parameters after finding a suitable ground state circuit.




Figure 5.7: Average energy E_ℓ of modes excited at the ℓ^{th} scale, with dashed line showing an exponential scaling proportional to $2^{z\ell}$ with $z = 1$ for reference, system size $L = 512$.

5.4 Benchmarking

In Sec. 5.3 we described how to variationally optimize a TMERA. In this section we measure how well the resulting state approximates the Gibbs state.

5.4.1 Energy and Entropy

One measure of accuracy of our TMERA states is observing the relationships between energy, entropy, and temperature. One way of defining a state in thermal equilibrium is that it is the density matrix of highest entropy with a given expectation value of the energy, or equivalently the lowest energy state with a given entropy. The solid black line in Fig. 5.2 traces this extremal relation between energy and entropy which is occupied by positive temperature Gibbs states. All other quantum states must lie in the region below this extremal line, so the degree to which our thermal ansatz states can hug this line is a measure of their success.

For a state with $L = 512$, the largest difference in entropy for the TMERA state with $D = 6$ is around 0.014 bits per qubit and occurs near temperature $T = 0.28$, while the $D = 1$ state has the largest entropy difference of around 0.06 bits per qubit near $T = 0.19$. The largest energy differences are 0.011 per qubit for $D = 6$ near $T = 1.3$ and 0.096 per qubit for $D = 1$ near temperature $T = 0.87$.

At low temperatures we observe polynomial scaling of entropy and energy with temperature, as seen in Fig. 5.8. The log-log plots seem consistent with linear and quadratic scaling for the entropy and energy respectively. This is exactly what we expect from the long-wavelength description in terms of a conformal field theory. In any such theory, the scaling symmetry implies that the temperature sets a thermal length scale $\xi(T) \sim 1/T$ (with $k_B, \hbar$, and the “speed of light” set to unity). As the only dimensionful quantity in the problem, this length scale determines the temperature dependence of the entropy density, $s \sim \xi^{-d}$, and the energy density, $\varepsilon \sim \xi^{-(d+1)} \sim T \xi^{-d}$, where d is the dimension of space. Setting $d = 1$ gives $s \sim T$ and $\varepsilon \sim T^2$, consistent with the data. Note that the shorter depth ansatz states deviate at low temperature when they cannot lower the energy beyond the ground state ansatz for depth D .

5.4.2 State Fidelity

The state fidelity $\mathcal{F}$ between two states is given by

$$\mathcal{F}(\rho, \sigma) = \text{tr} \sqrt{\sqrt{\sigma} \rho \sqrt{\sigma}}, \quad (5.14)$$

or simply $\langle \psi | \phi \rangle$ when the states are pure. The fidelity between two Gaussian fermion states can be computed using covariance matrices according to Eq.A.16. We compute the state fidelity between the thermal state at $e^{\beta H} / \mathcal{Z}$ using the Hamiltonian from Eqn.5.7 and




Figure 5.8: **Top:** Log-log plot of the energy above the ground state energy per qubit vs temperature, with dashed line showing a quadratic scaling for reference. **Bottom:** Log-log plot of the entropy per qubit vs temperature, with dashed line showing a linear scaling for reference. System size $L = 512$

the TMERA ansatz state targeting the same inverse temperature β . Despite the simplicity of the TMERA ansatz, our approximate state still has fidelity $\mathcal{F} > 0.4$ for all temperatures using a $D = 6$ circuit and system size $L = 512$, see Fig. 5.2. Even $D = 3$ ansatz states have fidelity higher than 0.2 for all temperatures at this system size.

The system-size normalized infidelity scales as

$$1 - \mathcal{F}^{1/L} = s(T), \quad s(T) < 10^{-2}. \quad (5.15)$$

(Fig. 5.9) For high temperatures the fidelity displays this scaling even for the smallest system sizes we treat ($L = 16$); for low temperatures the fidelity approaches this scaling for $L \gtrsim 128$.

While ground state infidelity becomes exponentially small with increasing D , the normalized infidelity, roughly the infidelity per site, seems unable to decrease much less than 10^{-3} for some $O(1)$ temperatures. This may be improved somewhat by allowing further optimization of the unitary circuit parameters or allowing higher depth circuits, however, the greatly-degenerate product spectrum of any TMERA density matrix is an inherent limitation to the expressability of the ansatz at any circuit depth.

5.4.3 Two-point correlations

Scale symmetry of the ground state constrains the correlation function $\langle \gamma_i \gamma_j \rangle$ to decay polynomially with distance as $\sim |i - j|^{-1}$, up to effects from finite size and finite lattice spacing. The thermal energy scale breaks scale symmetry at finite temperature and introduces a correlation length, leading to exponential decay of this correlation function, with $\log(\langle \gamma_i \gamma_j \rangle_\beta) \sim -\beta |i - j|$.

The TMERA states, however, maintain polynomially decaying quadratic fermion operators, albeit with $\langle \gamma_i \gamma_j \rangle \sim |i - j|^{-2}$, see Fig. 5.10. Despite this very different behavior of the




Figure 5.9: Normalized infidelity for $D = 6$ TMERA states for various temperatures T and system sizes L .

correlation function over long distances, our approximate Gibbs states still have relatively good fidelity. These values of these long distance observables are exponentially small, so despite greatly overestimating them proportionally by using a polynomially decaying function, the error is still within a small additive precision and does not impact the overall fidelity too greatly.

One interpretation of our ansatz is that the DMERA circuit transforms a Hamiltonian consisting of only single qubit Z operators to an effective Hamiltonian with long range couplings that decay polynomially. The TMERA ansatz then prepares the exact Gibbs state for this effective Hamiltonian, which also features polynomially decaying correlation functions. Despite these long range couplings, the effective Hamiltonian is still close enough to our target Hamiltonian that the TMERA states can still serve as Gibbs state approximations. See App. B.2 for further details on this interpretation.




Figure 5.10: Quadratic fermion correlation functions, featuring exponential decay with distance for exact Gibbs states at finite temperature (solid) and polynomial decay for TMERA states (dashed). The lowest temperature of $T = 1/L$ converges to the highly accurate ground state correlation function due to finite size, while higher temperature correlation functions only match well for larger expectation values, i.e. those within a correlation length set by the temperature. Computed for system size $L = 512$ and depth $D = 6$.

5.5 Discussion

We propose a variational quantum circuit ansatz for preparing approximate Gibbs states called a *thermal multiscale renormalization ansatz*, or TMERA. Our ansatz targets states in which scale invariance is broken only by temperature. The TMERA builds on a deep multiscale renormalization ansatz (DMERA) for the ground state, which we presume is already known (perhaps via an independent variational optimization). For a 1d system of length L , the TMERA has $O(\log L)$ parameters if the system’s dynamical critical exponent is not known; if that exponent is known, it has a single parameter. The TMERA requires $O(DL)$ gates, $D \sim 1 - 10$ the depth of a brickwork unitary circuit used in scaling.

In a TMERA both measurements required for optimizing free-energy—energy and entropy—are natural. The entropy follows directly from the structure of the TMERA and the variational parameters. The energy can be measured using only the past causal cone of a two-site region; this past causal cone is a subset of the full TMERA involving only $O(D \log L)$ gates.

We benchmarked our TMERA on the free-fermion integrable transverse-field Ising model at its critical point. We found that with a $D = 6$ circuit the ansatz produced fidelity $\mathcal{F} \gtrsim 0.4$ on a 512-site system, per-site infidelity $1 - \mathcal{F}^{1/L} < 10^{-2}$ for all system sizes and temperatures, and nearly maximal entropy at fixed energy.

Neither the basic inputs nor the physical intuition behind the TMERA are specific to the transverse field Ising model we used for benchmarking. The inputs are a scale-invariant DMERA network describing the ground state and a scaling ansatz for the characteristic energy as a function of scale; these inputs reify the intuition that ground states of models at critical points are scale-invariant, and thermal states are governed by the model’s dynamical scaling. A critical point is described at long wavelengths by a conformal field theory (or a more general non-relativistic analog) and the scaling symmetry of the long-

wavelength theory is expected to imply the existence of analogous structures independent of the details of the theory. In other words, our construction can plausibly apply to a wide variety of quantum critical points in diverse dimensions. It may of course be challenging to classically simulate the resulting tensor networks, especially in higher dimensions, but our architecture would remain a potentially powerful quantum variational ansatz that leverages our significant a priori expectations for the physics.

5.6 Extension to systems away from criticality

For a critical ground state, the underlying ground state circuit is scale-invariant. The temperature imposes an energy scale E_{UV} that breaks scale symmetry in the TMERA via the different mode energies at each scale $E_\ell \approx E_{UV}(2^\ell/L)^z$. If we tune away from the critical point, perhaps by changing a parameter h to $h \neq h_c$, we introduce another length scale $\xi \sim (h - h_c)^{-\nu}$ and energy scale ξ^{-z} . The ground state is no longer scale invariant, so one might worry that the TMERA is useless, or at least no improvement on a brickwork circuit. But the system is locally scale invariant for lengths $l \lesssim (h - h_c)^{-\nu}$. So the brickwork circuit will be redundant, perhaps highly so.

We believe a TMERA modified to take into account the scale $\xi \sim (h - h_c)^{-\nu}$ will eliminate the redundancy. The correct TMERA will be cut off at

$$n(h - h_c) \sim \log \xi \sim -\nu \log(h - h_c) \quad (5.16)$$

layers. This number of layers is not a priori known—but learning it is a matter for ground-state DMERA optimization, not TMERA optimization. Given the ground state DMERA, the TMERA construction and optimization will proceed as we have described here with fewer gates and parameters than a brickwork circuit product spectrum ansatz.

Chapter 6: Mana & Thermalization

Quantum hydrodynamics—the long-wavelength, long-time dynamics governing transport of conserved quantities—is believed to be efficiently simulable on classical computers, even for strongly-interacting systems. If a system’s Hamiltonian satisfies the *eigenstate thermalization hypothesis* (ETH) [133–135] it will rapidly reach local thermodynamic equilibrium. After that time local observables are well described by a Gibbs state with spatially varying thermodynamic potentials. Since hydrodynamics is presumptively local [136, 137] one expects a local approximation to be enough to compute long-time dynamics. References [138, 139] offer quantitative arguments that this is the case.

Recent work has built on these conceptual insights to create workable numerical methods in one dimension. The “generalized relaxation time approximation” [140] treats integrable models perturbed by small integrability-breaking terms; it replaces the detailed effect of the perturbation by a local phenomenological collision integral with a single parameter, a relaxation time. Near-equilibrium transport properties are accessible in non-equilibrium steady-state setups [141–145]. For unitary quench dynamics far from integrability there is a new generation of matrix product operator (MPO) methods, *density matrix truncation* (DMT) [146] and *dissipation assisted operator evolution* (DAOE) [147]. These methods assume that non-local information is unimportant and can be discarded, if one proceeds carefully—but so far the only practical use of either of these methods has been for a model close to free-fermion integrability [148]. All these approaches share two key assumptions: that local approximations to the full state “simple”, in some sense; and that




Figure 6.1: **Subsystem mana in real-time evolution** of infinite-temperature starting states on a chain of 50 sites. Bold lines show the average over 20 random starting states (cf Sec. 6.1.2), faint lines show trajectories of individual states, and dot mark points at which we measure the mana with time step $dt = 1/8$. We use TEBD (cf Sec. 6.1.3). In each case we see a rapid rise to approximately the average mana of a Haar state (black line in bottom plot), due to local thermalization of the subsystem, followed by a slow decay as the subsystem entangles with the rest of the system.

capturing those local properties is enough to simulate the system’s dynamics (at least in some hydrodynamical regime). These MPO methods are restricted to one-dimensional systems.

But the insights that led to these methods are not limited to one-dimensional physics. In two (or more) spatial dimensions, ETH Hamiltonians will still locally thermalize, and one expects that Gibbs states will still have compact representations. The questions, then, are—what data structures and algorithms are suitable for higher-dimensional mixed-state dynamics? How computationally intensive is the early-time complexity “hump”? And how computationally intensive is the long-time dynamics?

We propose that near-Clifford algorithms and data structures offer a promising avenue for simulations of hydrodynamics in more than one dimension. Clifford circuits are efficiently simulable on classical computers [23–25], because they map each Pauli operator to

a single other Pauli operator. Pure Clifford circuits have been used to construct analytically tractable ETH Hamiltonians [149], and related quantum cellular automaton models have been used to study hydrodynamics [150, 151]. Circuits with few non-Clifford gates (or many gates that are nearly Clifford) are also classically simulable [152–159]. But we wish to do more than build model systems: we wish to simulate any given (ETH) Hamiltonian, in any dimension.

At first sight, Clifford or near-Clifford circuits are ill-suited to simulating hydrodynamics. They can map any particular input state to at most a finite number of states. (In particular, they map the computational basis state only to stabilizer states.) By contrast the orbit of an initial state under evolution generated by an ETH Hamiltonian traces out a continuous manifold. Moreover capturing chaotic growth in out-of-time-order correlations requires many non-Clifford gates [160]. But if we satisfy ourselves with non-unitary local approximations to the system’s dynamics—as we do in using DMT, DAOE, or the relaxation-time approximation—a path opens up. Stochastic Clifford circuits (that is, averages over an ensemble of Clifford circuits) can simulate many mixed states, in fact every state in the convex hull of the stabilizer states; one might construct such an ensemble by using randomized Trotter decompositions or similar techniques [161–165].

Each algorithm for near-Clifford simulation has classical computational complexity exponential in some measure of the distance of the circuit from a pure Clifford circuit. One can estimate these circuit measures by computing so-called *magic monotones* [166] for the states produced by the circuits. A magic monotone is a function on a quantum state that is non-decreasing under Clifford gates and certain other reasonable operations; consequently, it lower-bounds the number of non-Clifford operations required to produce the state. Many magic monotones exist [156, 166–170].

We ask: under what circumstances are local approximations to hydrodynamics accessible to near-Clifford simulations? We use the *mana* [166] of subsystem reduced density

matrices as a proxy for that accessibility. We choose it in part because it is closely related to the quantity controlling the difficulty of the Monte Carlo method of Ref. [152], and in part because it is computable without solving a minimization problem. We consider time evolution of a stabilizer state, and measure the mana of local reduced density matrices as a function of time. We find that local reduced density matrices display a clear “complexity hump” (Fig. 6.1): for times $t \lesssim \varepsilon^{-1}$ the local energy scale¹ these local subsystems rise to nearly maximal mana while for $t \gtrsim \varepsilon^{-1}$ the mana decreases, broadly following the mana of a Haar state with the appropriate entropy. For finite-temperature states, we additionally notice that the subsystem mana deviates from the Gibbs state mana (which is zero for sufficiently small inverse temperature [171]). We attribute this to mana’s sensitivity to small deviations from canonical typicality.

The chapter is organized as follows. In Sec. 6.1 we describe our model (a variant of the $q = 3$ Potts model), our procedure for choosing initial states, and our numerical methods, and we briefly describe mana. In Sec. 6.2 we treat the evolution of mana for infinite-temperature states, while in Sec. 6.3 we treat the evolution of mana for finite-temperature states.

6.1 Model, initial state, and methods

6.1.1 Model

Studies of thermalization and hydrodynamics typically use a transverse-field Ising model with additional longitudinal field. Because mana is only defined for qudits of odd dimen-

¹In the model (6.1) an appropriate local energy scale is

$$\varepsilon = \sqrt{2(J^2 + h_x^2 + h_z^2)} \approx 2.4$$

for our parameter values $J = h_x = h_z = 1$.




Figure 6.2: **Gap ratio** of Eq. (6.2) for the model Eq. (6.1) with $h_x = J = 1$, as a function of h_z , on a system of $L = 9$ sites. For $0.1 \lesssim h_z \lesssim 1.0$ the model has the GOE gap ratio $r \approx 0.53$ (upper dotted line); for $h_z \ll 0.1$ it drops below the Poisson gap ratio $r \approx 0.39$ (lower dotted line)

sion,² we cannot use that model. Instead we use the analogous qutrit model

$$\begin{aligned}
 H = & -J \sum_j [Z_j^\dagger Z_j + h.c.] - h_x \sum_j [X_j + X_j^\dagger] \\
 & - h_z \sum_j [Z_j + Z_j^\dagger]
 \end{aligned} \tag{6.1}$$

where $X = \sum_{m=0}^2 |m+1 \bmod 3\rangle\langle m|$ and $Z = \sum_{m=0}^2 e^{-2\pi i m/3} |m\rangle\langle m|$ are the clock and shift operators. We take $J = h_x = h_z = 1$ and system size $L = 50$, except where otherwise specified.

This model is ETH for most parameter values. Even with $h_z = 0$, where this model becomes the $\mathbb{Z}_3$ Potts model, it is only integrable at the critical point $J = h_x$ [172–175].

We add the longitudinal-field term to robustly break integrability, even at $J = h_x$, and the $\mathbb{Z}_3$ onsite rotation symmetry, which will lead to more complicated hydrodynamics.

²We performed much of this work before the Rényi entropy of magic [170] was defined for mixed states. Going forward we suggest others use that monotone.

To check that this Hamiltonian is in fact ETH, we measure the eigenstate gap ratio. The eigenstate gap ratio is

$$r := \left\langle \frac{\min(\delta_\alpha, \delta_{\alpha+1})}{\max(\delta_\alpha, \delta_{\alpha+1})} \right\rangle, \quad (6.2)$$

with $\delta_\alpha := E_{\alpha+1} - E_\alpha$ and α indexing the energy eigenvalues in increasing order. The average is then taken over all values of α . We additionally average over symmetry sectors. We plot r as a function of h_z for a system of $L = 9$ sites in Fig. 6.2; we find that the system has level-spacing statistics of a Gaussian orthogonal ensemble (GOE) for $0.1 \lesssim h_z \lesssim 1.0$, indicating that it satisfies the ETH. (In the limit $h_z \gg 1$ the uniform field term dominates; this will lead to a gap ratio $r = 0$ due to degeneracies. Similarly, in the limit $h_z \ll 1$ the model regains its $\mathbb{Z}_3$ symmetry, again leading to degeneracies and a gap ratio $r = 0$. In each case the model presumably remains ETH for sufficiently large system sizes.)

Despite the fact that we ultimately seek two-dimensional algorithms, we use the one-dimensional model (6.1). We do so precisely because there already exist data structures and algorithms—matrix product states (MPS) and time-evolving block decimation (TEBD)—for one-dimensional systems. To use an effective model like [138] risks assuming our conclusion.

6.1.2 Initial state

We wish to study the effect of local thermalization on subsystem mana. To do so cleanly, we choose each of our initial states to be the product of onsite stabilizer states (eigenstates of the single-site generalized Pauli matrices) picked to give constant energy density. We use a product state because we expect that the trajectory of a subsystem's mana will be intimately tied to the way it entangles with the rest of the system. (Additionally, choosing our initial state to be a product state keeps our matrix product state bond dimensions tractable for slightly longer.) We use stabilizer states, which have zero mana,




Figure 6.3: **Energy of pairs of single-qutrit stabilizer states** in the Hamiltonian (6.1) with $J = h_x = h_z = 1$. The energy for the stabilizer product pair is defined as the bond term for the pair plus half the single site terms from each qutrit. X_α and Z_α label eigenstates of the shift operator X and the clock operator Z , respectively, while S_α are eigenstates of the other onsite generalized Pauli operators ZX, ZX^2 .

so we can watch the initial growth of the mana, as well as its transfer from short-range degrees of freedom to long-range degrees of freedom. We choose constant initial energy density because we do not wish to confuse the effects of local thermalization with the those of long-time hydrodynamical relaxation (which will drive the system's dynamics after the initial thermalization).

Our requirement that the energy density be spatially homogeneous strongly constrains our initial state, and indeed the energy densities we can choose. Fig. 6.3 illustrates the energy densities we can achieve with the product of two stabilizer states. The energy density is a two-site operator—call the energy density on sites $j, j+1$ by

$$\begin{aligned}\varepsilon_{j,j+1} \equiv & -J[Z_j^\dagger Z_{j+1} + h.c.] \\ & - \frac{1}{2}h_x[(X_j + X_{j+1}) + h.c.] \\ & + \frac{1}{2}h_z[(Z_j + Z_{j+1}) + h.c.] \end{aligned} \quad (6.3)$$

for $1 < j < L-1$. (One must take care at the ends of an open chain, i.e. $j = 1, L-1$.) We can construct a state at a particular constant energy density by choosing the states $|\psi_1\rangle, |\psi_2\rangle$ on sites 1 and 2 to have that energy density, and then choosing the site on each successive site $j > 2$ such that $\varepsilon_{j-1,j} = \varepsilon_{1,2}$. This state selection fixes the energy density on the $L-1$ pairs of neighboring sites, but fails to account for the extra contribution of half the on-site field terms for the first and last site of the chain. The initial state is only accepted if the total energy is L times the specified local energy density. This allows for some fluctuations in energy density at the ends of the chain arising from the extra on-site field contributions as long as these fluctuations do not alter the desired total energy. Only a discrete set of energy densities is possible, and many possible energy densities admit only one state. If there is only one state in a manifold, we cannot average to avoid non-generic effects. We therefore restrict ourselves to energy density manifolds with at least two states per bond;

this gives a manifold with $\mathcal{N} > 2^{L-1}$ possible states. Fig. 6.3 illustrates the possibilities for $J = h_t = h_l = -1$.

We give some further details of the initial state configuration in Appendix C.1

6.1.3 Method

The majority of our simulations use time evolution with TEBD [176, 177] with a second-order Trotter decomposition. We find that Trotter step $dt = 1/16$ and bond dimension $\chi = 512$ give good convergence for $t \lesssim 6$ as measured by the half-chain entanglement entropy; the subsystem mana is converged for $t \lesssim 8$. We indicate the regime $6 \leq t \leq 8$ by dotted lines and lower color saturation. App. C.3 gives details of our convergence testing.

6.1.4 Mana

We measure *mana* $\mathcal{M}$ [166]. In this section we give a very brief précis of the relevant properties of mana. In App. C.2 we describe how to calculate it. For slightly less brief précis from a similar point of view see Refs. [178, 179]; for more details see Refs. [166, 180, 181].

Mana is a *magic monotone*, meaning that it is non-increasing under Clifford unitaries, partial traces, and (on average) stabilizer measurements. For a pure state, the mana is zero if and only if the state is a stabilizer state [180, 181]. Classical statistical mixtures of stabilizer states likewise have zero mana, but some zero-mana mixed states are not statistical mixtures of stabilizer states.

Mana is *multiplicative*, in the sense that for two density matrices ρ_1, ρ_2

$$\mathcal{M}(\rho_1 \otimes \rho_2) = \mathcal{M}(\rho_1) + \mathcal{M}(\rho_2) . \quad (6.4)$$

One therefore expects it to be extensive for states with short range correlations. In fact on

a system of ℓ qudits each with dimension d one can bound

$$\mathcal{M}(\rho) \leq \frac{1}{2}(\ell \ln d - S_2) , \quad (6.5)$$

S_2 is the second Rényi entropy of the state ρ , and the mana of a Haar-random state is extensive with subextensive corrections [178].

Mana is computed in terms of the *Wigner norm* $\mathcal{W}$, defined in App. C.2:

$$\mathcal{M}(\rho) = \ln \mathcal{W}(\rho) .$$

Sometimes it is convenient to work in terms of this Wigner norm, which shares the properties of mana, suitably translated.

The Wigner function for bosons is a quasiprobability distribution which acts as a phase space representation of the bosonic quantum state. It is not a true probability distribution because the Wigner function can take on negative values, although it still integrates to one. The one norm will exceed one whenever there are negative values in the Wigner function, which in a bosonic system indicates non-Gaussianity of the state [182]. In our case, we use a discrete version of the Wigner function for spin systems, and the negativity of this representation will analogously signal when a state fails to be a stabilizer state.

6.2 Mana at infinite temperature

Consider a length- ℓ subsystem of our chain (we consider the ℓ most central sites). The evolution of the subsystem mana is governed by the competition between two effects. The subsystem's internal dynamics locally randomize the state, for a fast rise in the mana density. At the same time the coupling between the subsystem and its complement steadily increases the subsystem's entropy: since the mana is bounded by Eq. (6.5) this ongoing in-




Figure 6.4: **Subsystem mana in real-time evolution of infinite-temperature starting states** as a function of entropy deficit $\Delta(t) = \ell \ln 3 - S_2(t)$, with S_2 the subsystem second Rényi entropy (cf Fig. 6.1). Bold lines show the average over 20 random starting states (cf Sec. 6.1.2), faint lines show trajectories of individual states, and dot mark points at which we measure the mana with time step $dt = 1/8$. In each case we see a rapid rise to approximately the average mana of a Haar state, due to local thermalization of the subsystem, followed by a slow decay as the subsystem entangles with the rest of the system.

crease of entanglement must decrease the mana. For short times the first effect dominates and the mana rapidly rises; for longer times the second is dominates and the mana must slowly decay. In the long-time limit $\ell \ln 3 - S_2 = 0$ (up to a Page correction, which will be small for $\ell \ll L/2$), so the subsystem mana must be

$$\lim_{t \rightarrow \infty} \mathcal{M} = 0 .$$

The system as a whole, by contrast, is in a pure state. One expects this pure state to be essentially a random state in the microcanonical ensemble, and so to have Haar-like mana

$$\mathcal{M} \approx \frac{1}{2}(\ell \ln 3 - \ln \pi/2) . \quad (6.6)$$

This whole-system saturation was observed for a related model in Ref. [183].

In Fig. 6.1 we show the subsystem mana as a function of time. The subsystems consist of the ℓ central sites of the MPS state. We clearly see both effects—rapid initial rise due to local randomization, followed by decay to the infinite temperature value $\mathcal{M} = 0$.

The bound Eq. (6.5) depends solely on the “entropy deficit”

$$\Delta = \ell \ln d - S_2 ; \quad (6.7)$$

one is entitled to ask how close the subsystem mana comes to saturating that bound. Moreover one might expect the decay of the subsystem mana to track the mana of a subsystem of a Haar-random state. For large Hilbert space dimension the Wigner norm of a subsystem of a Haar state is controlled solely by the entropy deficit, and given by [178]

$$\langle \mathcal{W} \rangle_{\text{Haar}} = \sqrt{2/\pi} (\sigma/\mu) e^{-\mu^2/2\sigma^2} + \text{erf}(\mu/\sigma\sqrt{2}) \quad (6.8)$$




Figure 6.5: **Subsystem mana in real-time evolution of finite-temperature starting states** on a chain of 50 sites. Bold lines show the average over 20 random starting states (cf Sec. 6.1.2), faint lines show trajectories of individual states, and dot mark points at which we measure the mana with time step $dt = 1/8$. As in the infinite-temperature case we see a rapid rise, due to local thermalization of the subsystem, followed by a slow decay as the subsystem entangles with the rest of the system—but the decay no longer continues to mana $\mathcal{M}_\ell = 0$ at accessible times.

where

$$\frac{\sigma^2}{\mu^2} = e^\Delta - 1 .$$

For $\Delta \gtrsim \ln \pi/2$ the mana resulting from this expression becomes [178]

$$\mathcal{M} \approx \frac{1}{2} [\Delta - \ln \pi/2] . \quad (6.9)$$

In Fig. 6.4 we plot the subsystem mana against the entropy deficit Δ . Each subsystem starts at zero entropy S_2 , hence large entropy deficit Δ , and moves right to left to smaller entropy deficit. (We plot the entropy as a function of time in App. C.4.) Dot mark points where we measure the mana, every time step $dt = 1/8$. We see again a fast early rise and a long-time decay, matching the Haar prediction based on (6.8).

At intermediate times the state mana undershoots the Haar prediction. This is because different parts of the subsystem are not fully entangled. Heuristically, the subsystem behaves like a tensor product of Haar states on smaller subsystems. The mana is the sum of the manas of these smaller subsystems (cf (6.4)), each of which comes with a $\ln \pi/2$ correction from (6.9). Once the system is fully (internally) entangled one can think of it as a single Haar state with a single $\ln \pi/2$ correction: this gives the black line in Fig. 6.4.

6.3 Mana at finite temperature

How does this picture change when the initial state has nonzero energy? Fig. 6.5 shows subsystem mana as a function of time, and Fig. 6.6 shows it as a function of the subsystem's entropy deficit. In each case we show a variety of energies, labeled by equilibrium temperature.³

³The model has one conserved quantity, the energy. Every energy E_0 corresponds to an equilibrium temperature $\beta(E_0)$ such that

$$E_0 = Z^{-1} \text{Tr} H e^{-\beta(E_0)H} ;$$




Figure 6.6: **Subsystem mana in real-time evolution of finite-temperature starting states** as a function of entropy deficit $\Delta(t) = \ell \ln 3 - S_2(t)$, S_2 the subsystem second Rényi entropy (cf Fig. 6.1). Stars mark Gibbs state values. Time increases as entropy deficit decreases, i.e. from right to left. In each case the long-time endpoint ($t = 8$, leftmost on each curve) has mana and entropy noticeably larger than the Gibbs value.

We see an initial rise followed by a slow decay, broadly following the Haar value, as in the infinite-temperature case. In both Fig. 6.5 and Fig. 6.6, the $\beta = 0.233$ averages display much more variation than other temperatures. We believe that the variation is due to the tight constraints on initial states at this energy (cf App. C.1). Although there are exponentially many suitable initial states, they are locally similar.

But careful examination of Fig. 6.6 presents a mystery. We extract the subsystem mana and entropy of a Gibbs state from exact-diagonalization simulations on small systems (cf Fig. 6.7 and App. C.6); we mark those values in Fig. 6.6 with a dot.

We attribute the discrepancy to a subtle finite-time effect related to deviations from so-
if the state after time evolution locally approximates any Gibbs state, it must be this Gibbs state.

called *canonical typicality* [184–190]. Essentially, the discrepancy is controlled by the size of the Hilbert space of the region with which our subsystem is entangled. As that region grows the mana will approach its thermal value—but our MPS simulations are limited to times for which the entangled region is small.

To understand this, consider first subsystem mana in finite systems at long times. Heuristically, one expects long-time states to behave much like Haar-random states on a microcanonical subspace (we rehearse the standard intuition behind this statement in App. C.5). Almost all such states have subsystem reduced density matrices near in trace norm to the subsystem reduced density matrices of the microcanonical density matrix [186], hence to that of the Gibbs state.

But this is not enough: two density matrices close in trace norm can have widely divergent Wigner norms. We show in App. C.7 that if density matrices ρ, σ on a dimension- d subspace have

$$\|\rho - \sigma\|_1 \leq \eta ,$$

then

$$\left| \|\rho\|_W - \|\sigma\|_W \right| \leq \min(d^2\eta, d^{1/2}\sqrt{\eta}) . \quad (6.10)$$

In App. C.8 we argue that the Hilbert space dimension factors in (6.10) mean that it can only give very wide bounds on mana when combined with the result of Popescu. Moreover a Gibbs state can have zero mana but be near the boundary of the zero-mana region; a nearby thermal state may then have small but nonzero mana.

In Fig. 6.8 we plot subsystem mana as a function of time in small (Krylov-accessible) systems at much longer times than are accessible to matrix product states; we see a steady-state deviation between the mana of the time-evolved state and the Gibbs distribution mana.

So much for finite systems at large times. What about large systems at finite times? We have considered fixed-sized subsystems; as the surrounding system becomes large, even




Figure 6.7: **Subsystem mana for Gibbs states** of the Hamiltonian (6.1) on 9 qutrits with periodic boundary conditions. We take $J = h_z = h_z = 1$.

bounds based on (6.10) and the typicality result of [186] will strongly constrain the mana. But a finite-time state will be very different from most random microcanonical states: those states have nearly maximal entanglement, while the state at time t has entanglement entropy $S \approx 2ct$. That finite-time state, then, can be crudely modeled by a Haar state from the microcanonical subspace on a system of size $l + 2ct$. (To construct a less crude model one might use random matrix product states. [191–193] A random MPS would mimic not only the entanglement of the subsystem with its surroundings, but also its internal entanglement structure.) This crude model is broadly consistent with Fig. 6.5, in that the decay times increase linearly with subsystem time. We leave a more careful comparison of numerics with predictions from canonical typicality to future work.

6.4 Discussion

We would be remiss not to outline some limitations of our work. Most seriously, comparisons between the mana complexity hump of Fig. 6.1 or 6.5 and the entropy complexity




Figure 6.8: **Long-time subsystem mana per site in Krylov evolution** of an $L = 11$ site chain with thermal values marked with horizontal dashed lines. The average value of five samples is given by a dark solid line with each sample plotted using a light line. Thermal values for all samples except $\beta = -0.227$ are too small to be seen on the plot scale.

hump of e.g. [147] Fig. 2 can be importantly misleading. The cost of an MPO calculation is, broadly speaking,

$$[\text{MPO cost}] \sim \sum_{\text{bonds}} \exp[\alpha S_{\text{bond}}] \propto L, \quad (6.11)$$

where S_{bond} is the entanglement entropy across a bond and α is some power. The simulation cost may be dominated by the peak bond entropy, but it is still polynomial in system size. Straightforward near-Clifford simulations, by contrast, will have cost

$$[\text{near-Clifford cost}] \sim \exp \left[\gamma \sum_{\text{subsystems}} \mathcal{M}_{\text{subsystem}} \right] \quad (6.12)$$

$$\sim e^{\gamma m L}$$

for some finite mana density m . That is, even if we take advantage of the insensitivity of hydrodynamics to details of long-range correlations, the cost is still exponential in system size. Worse, since the peak mana is close to the maximum mana, this suggests that there are no cost savings to be had from short-range approximations.

We believe this obstacle is superable. Because the peak mana occurs at short times, when the system displays only short-range entanglement, one should be able to decouple simulation of different subsystems, effectively exchanging sum and exponential in (6.12) and giving cost polynomial (indeed linear) in system size. Concretely, imagine dividing the system into subsystems of some length l , separating those subsystems separately for $t \propto l$ and then re-introducing the couplings. While this dramatically changes the early- to intermediate-time dynamics, it should not change the long-time hydrodynamics, and it suggests that more sophisticated schemes are possible.

The nonzero finite-temperature long-time mana is another limitation of our work. Absent decoupling tricks like those required for early-time simulations, this nonzero subsystem mana density means that the simulation cost is exponential in system size, albeit with

small exponent.

And looming behind these limitations is the fact that while mana controls the difficulty of some classical simulation methods and lower-bounds the non-Clifford resources required to create a state, these bounds are not constructive: knowing that a state has low mana does not give one a recipe for constructing it. We leave algorithms—the analogues of DMT or DAOE—to future work.

We have framed this result in terms of classical simulation, because we expect that to be its immediate application, but in many ways it is more naturally understood in the context of future error-corrected quantum computers. In many error-correction architectures Clifford gates are “easy”. Non-Clifford gates, by contrast, must be performed by costly magic-state distillation and injection schemes: magic comes dear. Our results suggest that good local approximations to long-time states will dramatically reduce requirements for this key resource.

Chapter 7: Discussion

In this dissertation we have studied the quantum resources required for quantum simulation tasks on classical and quantum hardware. We have also considered in detail multi-scale circuit ansatz for state preparation on NISQ devices, including noise-resilience, error mitigation, and experiments on an ion-trap quantum computer.

In Chapter 2 we have studied the performance of site-by-site ground state preparation on free-fermion systems. Using computationally efficient formulas, we can access large system sizes and explore the phase space thoroughly at system sizes of up to $L = 72$. The quantities of interest are two system parameters, the energy gap and the overlap between subsequent ground states, which together determine the gate complexity of the algorithm. We find that the state overlaps correlate well with the energy gap of the system, sharing many features in their respective phase diagrams. The state overlaps are substantial for large regions of phase space and it is mainly the closing energy gap of critical systems or topological edge modes that impedes the performance of this algorithm. We also extend the algorithm to handle cases of ground state degeneracy, where the gap between degenerate ground states and the rest of the spectrum may be used, in addition to the quantity $\cos(\theta)$, where θ is the largest principal angle between a pair of ground state subspaces.

It is interesting to know how well these results extend to broader classes of interacting systems. While it is possible that a favorable scaling of the overlap with system size depends on the underlying Gaussian structure, this relation is not apparent in the numerical results. Indeed, as long as the system is gapped, any perturbative interaction of bounded

norm would change the state overlap by a term quadratic in the perturbation strength. In fact, the features present in the gap and overlap phase diagrams can be understood qualitatively via general properties of geometrically local many-body systems such as the correlation length and area laws. Based on this, we expect the state overlaps to behave similarly beyond the free fermion setting as well.

In Chapters 3, 4, and 5 we have considered the variational preparation of ground states and thermal states of the critical Ising model and its fermionic counterpart through the Jordan-Wigner transformation. The critical free-fermion model is special in many ways, but we expect similar circuit architectures to be successful more broadly for states of local Hamiltonians. In these situations, locality and separation of length scales have been invaluable tools in physical analysis which can also be leveraged in designing good quantum circuits for preparing physically relevant states. Adapting to systems which are non-integrable could be done by allowing a more general set of local gates which goes beyond Gaussian fermion operations (i.e. universal rather than matchgate circuits).

Additionally, states with finite correlation length ξ could be prepared using $O(\log \xi)$ scaling layers optimized independently, which may be interpreted as a nontrivial renormalization group flow which disentangles the state at scales beyond the correlation length. The variational ansatz states we find are able to outperform similar analytic constructions while maintaining much of their simplicity, compared to a more general variational ansatz such as a large bond dimension MERA state. Finding a way to more precisely formalize variational multi-scale circuits as genuine scaling transformations beyond state preparation may also find use in algorithms for simulating conformal field theories and quantum field theories using closely ideas related from renormalization and wavelet theory [64, 99, 194–196].

For NISQ applications, mitigation of error under tight hardware constraints is a key consideration. The symmetry averaging methods for mitigating ansatz-induced error complement the intrinsic error resilience features of the multi-scale circuit architecture.

We were able to numerically and experimentally demonstrate the convergence and noise-resilience of local subregion preparation, which manifests as a dissipative circuit targeting a renormalization fixed point state. Numerical noise models of gate imprecision and local depolarizing noise support a linear scaling of the fixed point observables in the local error parameters σ^2 and p , and the channel depth D . We also showed evidence that different choices of gate decomposition may greatly affect the noise susceptibility for some kinds of noise.

Error mitigation techniques such as zero-noise extrapolation are also seen to dramatically improve the mitigated error rate of observables with relatively little gate overhead when using extrapolation schemes designed using knowledge about the error propagation for a dissipative channel. These effectively lower noise levels achieved from classical post-processing can be leveraged to obtain the approximation accuracy of larger-depth circuits even in regimes where their raw output is noisier than shallower-depth circuits.

We are interested in how similar DMERA circuits which prepare states in a noise-resilient way could be adapted for models away from a critical point, for thermal states, or for non-integrable models. It has also yet to be seen how the scaling with channel depth and benefits of this error mitigation scheme hold up for real experimental data, but which may all be tested relatively soon on NISQ devices.

There may also be interesting connections with this work and holographic models of quantum gravity and error correction. This relation could arise from the interpretation of MERA-like circuits as a description of the AdS/CFT correspondence [197] and the recognition of error correcting properties in the structure of the duality [198–200]. Alternatively, the robustness of these circuits may be interpreted holographically in terms of the dilution of noise in an expanding universe [61, 201].

In chapter 5 we propose a variational quantum circuit ansatz for preparing approximate Gibbs states called a *thermal multiscale renormalization ansatz*, or TMERA. Our ansatz

targets states in which scale invariance is broken only by temperature, as in a quantum critical fan. The TMERA builds on a deep multiscale renormalization ansatz (DMERA) for the ground state, which we presume is already known (perhaps via an independent variational optimization). For a 1d system of length L , the TMERA has $O(\log L)$ parameters if the system’s dynamical critical exponent is not known; if that exponent is known, it has a single parameter. The TMERA requires $O(DL)$ gates, $D \sim 1 - 10$ the depth of a brickwork unitary circuit used in scaling.

In a TMERA both measurements required for optimizing free-energy—energy and entropy—are natural. The entropy follows directly from the structure of the TMERA and the variational parameters. The energy can be measured using only the past causal cone of a two-site region; this past causal cone is a subset of the full TMERA involving only $O(D \log L)$ gates.

We benchmarked our TMERA on the free-fermion integrable transverse-field Ising model at its critical point. We found that with a $D = 6$ circuit the ansatz produced fidelity $\mathcal{F} \gtrsim 0.4$ on a 512-site system, per-site infidelity $1 - \mathcal{F}^{1/L} < 10^{-2}$ for all system sizes and temperatures, and nearly maximal entropy at fixed energy.

Neither the basic inputs nor the physical intuition behind the TMERA are specific to the transverse field Ising model we used for benchmarking. The inputs are a scale-invariant DMERA network describing the ground state and a scaling ansatz for the characteristic energy as a function of scale; these inputs reify the intuition that ground states of models at critical points are scale-invariant, and thermal states are governed by the model’s dynamical scaling. A critical point is described at long wavelengths by a conformal field theory (or a more general non-relativistic analog) and the scaling symmetry of the long-wavelength theory is expected to imply the existence of analogous structures independent of the details of the theory. In other words, our construction can plausibly apply to a wide variety of quantum critical points in diverse dimensions. It may of course be challenging to

classically simulate the resulting tensor networks, especially in higher dimensions, but our architecture would remain a potentially powerful quantum variational ansatz that leverages our significant a priori expectations for the physics.

Finally, in chapter 6 we analyze the non-stabilizer resource of mana and how it manifests in local subsystems under thermalizing dynamics. We find that local subsystems of zero-energy initial states have zero mana after a short local thermalization time proportional to the subsystem size, consistent with a characterization as an infinite temperature Gibbs state. This suggests that infinite-temperature hydrodynamics may be simulated classically methods with low overhead. In the context of quantum simulation, there may be effective dynamics using mixtures of stabilizer states and Clifford circuits, or circuits with few non-Clifford gates, that also effectively reproduce the infinite-temperature hydrodynamics. At finite temperature the landscape is more complicated: for sufficiently high (but still finite) temperature the Gibbs state, hence the long-time thermal state, has zero mana; while for somewhat lower temperatures, the subsystem mana is small but nonzero. Additionally, regardless of temperature, the subsystem mana takes a long time to relax to the thermal value, because mana is sensitive to small deviations from canonical typicality.

Because mana is sensitive to small deviations from the Gibbs state, approximate simulations may be able to achieve some desired precision using states with much lower mana than that of the target state. This could allow approximate simulations to use fewer non-Clifford resources required than would be for exact simulation, which would broaden the scope of physics accessible using near-Clifford simulation techniques and reduce the cost of quantum simulation.

We have explored in this dissertation the capabilities and limitations of near term quantum simulation. We hope that this work may prove useful to the greater scientific community on the long road toward realizing practical quantum computers as a computational tool afforded to humanity.

Appendix A: Gaussian Fermions

Quadratic fermionic Hamiltonians are a widely used and useful class of toy models that often allow closed-form analytical expressions or numerically efficient computation procedures for physical properties such as ground state energies and multi-point correlation functions. The eigenstates of these Hamiltonians are known as Gaussian fermion states and the unitary transformations generated by them preserve this set of states. Gaussian fermions may be represented using polynomial resources and are therefore a classically simulable set of states. Some basics of the representation of Gaussian fermionic states and operations are outlined here following [51]. See also the Grassmann representation in [67].

A system of non-interacting, spinless fermions on L sites can be modeled by a Hamiltonian which is quadratic in the mode creation and annihilation operators $c_j^\dagger, c_j$ on site $j = 1, \dots, L$, is given by

$$H = \sum_{i,j} A_{ij} c_i^\dagger c_j + \frac{1}{2} \sum_{i,j} (B_{ij} c_i^\dagger c_j^\dagger + h.c.). \quad (\text{A.1})$$

The mode operators satisfy anti-commutation relations $\{c_i, c_j\} = 0, \{c_i, c_j^\dagger\} = \delta_{ij}$, which allows us to restrict the form of the coupling matrices A, B . Namely, we assume that $A_{ij} = A_{ji}^*$ (Hermitian), and $B_{ij} = -B_{ji}$ (anti-symmetric). The Hamiltonian can then be written

more compactly in a bilinear form,

$$H = \frac{1}{2} \begin{bmatrix} \vec{c}^\dagger & \vec{c} \end{bmatrix} \begin{bmatrix} A & B \\ -B^* & -A^* \end{bmatrix} \begin{bmatrix} \vec{c} \\ \vec{c}^\dagger \end{bmatrix}, \quad (\text{A.2})$$

where we omit an overall additive constant $\frac{1}{2}\text{Tr}(A)$. The same Hamiltonian can also be expressed in terms of $2L$ Majorana operators which we define for a given site j as

$$\gamma_{2j-1} = \frac{(c_j + c_j^\dagger)}{\sqrt{2}}, \quad \gamma_{2j} = \frac{(c_j - c_j^\dagger)}{\sqrt{2}i}. \quad (\text{A.3})$$

Like their Dirac counterparts, the Majorana operators satisfy canonical anti-commutation relations $\{\gamma_i, \gamma_j\} = \delta_{ij}$.

The ground states (and all excited states as well for fermions) of a quadratic Hamiltonian are Gaussian, and therefore are fully characterized by all linear and quadratic expectation values $\langle \gamma_i \rangle, \langle \gamma_i \gamma_j \rangle$. Parity conservation (which follows from H having a quadratic form) implies that the first moments $\langle \gamma_i \rangle$ vanish. The expectation values of all quadratic products of the Majorana modes can be arranged in the form of the *covariance matrix* given by

$$\Gamma_{jk} = i\langle [\gamma_j, \gamma_k] \rangle \equiv i\langle \gamma_j \gamma_k - \gamma_k \gamma_j \rangle. \quad (\text{A.4})$$

Γ is real, Hermitian, and anti-symmetric by construction, and fully characterizes the state. It can also be shown that the eigenvalues of Γ are $\lambda_i = \pm 1$.

Consider an arbitrary basis change that transforms γ_j to $\tilde{\gamma}_j$ via the linear map $\lambda_j = O_{jk} \gamma_k$.

Since the $\tilde{\gamma}_j$ must satisfy anti-commutation relations, it follows that

$$\delta_{ij} = \{\tilde{\gamma}_i, \tilde{\gamma}_j\} \quad (\text{A.5})$$

$$= \sum_{k,l} O_{ik} O_{jl} \{\gamma_k, \gamma_l\} = \sum_{k,l} O_{ik} O_{jl} \delta_{kl} \quad (\text{A.6})$$

$$= \sum_k O_{ik} O_{kj}^T = [OO^T]_{ij}, \quad (\text{A.7})$$

which implies that O is an element of the orthogonal group $SO(2n)$ (and conversely, any orthogonal matrix corresponds to a valid basis transformation). The covariance matrix transforms under O as

$$\Lambda_{ij} = i\langle[\tilde{\gamma}_i, \tilde{\gamma}_j]\rangle \quad (\text{A.8})$$

$$= i \sum_{k,l} O_{ik} O_{jl} [\gamma_k, \gamma_l] = i \sum_{k,l} O_{ik} O_{jl} \Gamma_{kl} \quad (\text{A.9})$$

$$= i \sum_{k,l} O_{ik} \Gamma_{kl} O_{lj}^T = [O\Gamma O^T]_{ij}. \quad (\text{A.10})$$

In other words, $\Lambda = O\Gamma O^T$. Any real antisymmetric matrix M can be expressed in a canonical form $M = P^T \Omega_M P$ for some orthogonal matrix P , where Ω_M is given by

$$\Omega_M = \begin{bmatrix} O & E_M \\ -E_M & O \end{bmatrix}, \quad (\text{A.11})$$

where O is the null matrix. The matrix E_M is a diagonal matrix containing the eigenvalues of M times i . For the covariance matrix Γ , we can write $\Gamma = P^T \Omega P$, where $\Omega = \begin{bmatrix} O & I \\ -I & O \end{bmatrix}$ is known as the *symplectic form*. The modes corresponding to this transformation, $\tilde{\gamma}_j = P_{jk} \gamma_k$, are the eigenmodes of the system.

We can write the (real antisymmetric) coupling matrix of the quadratic Majorana Hamil-

tonian as $P^T \begin{bmatrix} O & E \\ -E & O \end{bmatrix} P$ where P is diagonalizing transformation. This brings H into decoupled form

$$H = \frac{i}{2} \sum E_j \tilde{\gamma}_j \tilde{\gamma}_j', \quad (\text{A.12})$$

where E_i are the energies of the eigenmodes of H . The difference between the energies of the vacuum (no modes occupied) and the first excited state (lowest energy mode occupied), or the spectral gap, is simply given by the lowest mode energy, $m := \min_j E_j$.

The transformation P that brings the ground state covariance matrix Γ into canonical form also diagonalizes H . This provides a way to compute Γ from the coupling matrix of the Hamiltonian. Namely, one finds the matrix P via the diagonalization of the coupling matrix, and then computes $\Gamma = P^T \Omega P$.

Subsystems of a Gaussian state are defined by restricting the covariance matrix to a subset of the Majorana operators, eliminating the rows and columns associated with the discarded operators. Although we can relate the fermionic system to a system of qubits using a Jordan-Wigner duality, restricting to taking fermionic subsystems is a different restriction of the Hilbert space from the tracing out of qubit degrees of freedom. However, all even parity operators on the reduced fermionic and qubit subsystems will agree. Odd parity operators have zero expectation value in Gaussian states.

Any mixed Gaussian state can be considered a subsystem of some pure Gaussian state. The covariance matrix for a mixed state can be purified by adding additional Majorana modes, introducing the new rows and columns with entries so that the new covariance matrix satisfies $\Gamma^2 = -I$.

Any mixed Gaussian state may also be thought of as a thermal Gibbs state for some quadratic Hamiltonian. The covariance matrix for the Gibbs state of this Hamiltonian at

inverse temperature β is then

$$\Gamma_\beta = i \tanh(i\beta h) \quad (\text{A.13})$$

For some mixed covariance matrix Γ with eigenvalues $|\lambda_i| < 1$ then the Hamiltonian for which it is a Gibbs state is then given by $\tilde{h} = -i/\beta \text{atanh}(-i\Gamma)$.

Because all Gaussian states have a product spectrum, the entropy of any mixed Gaussian state is the sum of entropies S_j from the states eigenmodes, each of which can be written in terms of the respective eigenvalue of the covariance matrix.

$$\begin{aligned} S_j &= -\frac{1+\lambda_j}{2} \log \frac{1+\lambda_j}{2} - \frac{1-\lambda_j}{2} \log \frac{1-\lambda_j}{2} \\ &= \log 2 - \log \sqrt{1-\lambda_j^2} - \lambda_j \log \sqrt{\frac{1+\lambda_j}{1-\lambda_j}} \end{aligned} \quad (\text{A.14})$$

The covariance matrix organizes information about a Gaussian state in a useful form. Given two pure Gaussian states $|\psi_1\rangle, |\psi_2\rangle$ and corresponding covariance matrices Γ_1, Γ_2 , the magnitude of overlap between them, $|\langle\psi_1|\psi_2\rangle|$, has a particularly simple form

$$|\langle\psi_1|\psi_2\rangle| = \det\left(\frac{\Gamma_1 + \Gamma_2}{2}\right)^{1/4}. \quad (\text{A.15})$$

The fidelity between two mixed Gaussian states can also be computed in terms of their covariance matrices using the formula from[54],

$$\mathcal{F}_{\rho,\sigma} = 2^{-\frac{n}{2}} \det(I - \Gamma_\rho \Gamma_\sigma)^{\frac{1}{4}} \det\left(I - \sqrt{I + \Gamma_{\rho,\sigma}^2}\right)^{\frac{1}{4}} \quad (\text{A.16})$$

with $\Gamma_{\rho,\sigma} = \frac{\Gamma_\rho + \Gamma_\sigma}{I - \Gamma_\rho \Gamma_\sigma}$.

Local matchgate circuits are equivalent to unitary Gaussian operations and can be im-

plemented using covariance matrix methods. Local two-qubit gates become orthogonal transformations on four consecutive Majorana operators.

The quadratic fermionic Hamiltonian we work with which is dual to the critical transverse-field Ising model is shown in Eq. (5.7).

For finite sized spin chain with periodic boundary conditions, the dual quadratic fermion Hamiltonian must have anti-periodic boundary conditions for the even parity states, and periodic boundary conditions for odd parity states. Thus, with periodic boundary conditions the full Ising Hamiltonian is not strictly dual to a single quadratic fermion Hamiltonian, but all of its eigenstates are dual to a Gaussian fermion states.

To avoid these issues we will focus just on the Majorana fermion model with anti-periodic boundary conditions for our numerics, and in this case our circuits may be interpreted as local operations on fermionic degrees of freedom. In fact, because we are working with a free-fermion model, we will also restrict our gates to be unitary Gaussian fermion operations, equivalent to Bogliubov transformations. In terms of spins these operations are dual to matchgate circuits, and can be simulated classically via Gaussian fermion methods. Although we will restrict our attention to noninteracting fermions, the methods can be straightforwardly applied to qubits on a quantum computer for preparing states of spin systems, and the local gates generalized beyond matchgates to target interacting models.

Because local matchgate circuits are equivalent to linear transformations on the fermionic modes, this subgroup of quantum states and operations can be simulated by classical polynomial time algorithms [26], and may be implemented numerically using covariance matrix methods or via fermionic linear optics [67]. Under Jordan-Wigner duality, local two-qubit gates become orthogonal transformations on four consecutive Majorana operators. These techniques can be extended to tensor networks as matchgate tensor networks [27, 28].

The two-qubit gates $u(x,y)$ used in Chapters 4, 3, and 5 can be written implemented as

an orthogonal transformations $\tilde{u}(x', y')$ on four Majorana operators, using the parameters $x' = x + y$ and $y' = x - y$ as in Eq.(A.17).

$$\tilde{u}(x', y') \equiv \begin{bmatrix} \cos(x') & 0 & -\sin(x') & 0 \\ 0 & \cos(y') & 0 & \sin(y') \\ \sin(x') & 0 & \cos(x') & 0 \\ 0 & -\sin(y') & 0 & \cos(y') \end{bmatrix} \quad (\text{A.17})$$

The variational parameters found and used for the Ising model ground state are reported as θ_D for a depth D ansatz with $2D$ parameters, see eq. (A.18). Here, the $(2i - 1)^{th}$ parameter is the x' parameter for the i^{th} layer of gates within the scaling transformation and the $(2i - 1)^{th}$ parameter is the y' parameter for the same layer of gates.

We also found alternate sets of parameters for the modified Ising model in Eq. (3.6), listed in (A.19). These circuits have similar features and the symmetry-averaging improvement is displayed in Fig. 3.6 along with those of the original Ising model states.

$$\begin{aligned} \theta_1 &= [0.43188, -1.13891] \\ \theta_2 &= [0.1379, -0.56374, -0.53456, 0.18071] \\ \theta_3 &= [-1.79716, -1.51891, 0.64486, 2.0904, 0.10994, -0.31494] \\ \theta_4 &= [-1.66528, -1.55101, 1.05114, 1.82904, \\ &\quad 0.43426, -0.74951, 0.07349, -0.20764] \\ \theta_5 &= [-1.61046, -1.56358, 1.29107, 1.69499, 0.80132, \\ &\quad -1.06696, 0.30863, -0.54539, 0.05158, -0.14651] \\ \theta_6 &= [-1.58682, -1.56831, 1.42666, 1.6278, 1.08403, -1.28163, \\ &\quad 0.62053, -0.85436, 0.2323, -0.4140, 0.03862, -0.11017] \end{aligned} \quad (\text{A.18})$$

$$\begin{aligned}
\tilde{\theta}_1 &= [-0.22107 \quad -1.79187] \\
\tilde{\theta}_2 &= [0.37921, \quad -1.58672, \quad -0.23588, \quad -0.8906] \\
\tilde{\theta}_3 &= [0.09744, \quad -1.62352, \quad -0.86247, \quad -0.29472, \quad -0.18693, \quad 0.51259] \\
\tilde{\theta}_4 &= [-0.0058, \quad -0.21063, \quad 1.6268, \quad 0.78319, \\
&\quad -1.08715, \quad 0.1295, \quad 2.77886, \quad -0.08755] \\
\tilde{\theta}_5 &= [-0.08357, \quad -1.73696, \quad -0.09269, \quad -0.14014, \quad 0.14162, \\
&\quad 1.32826, \quad -1.05885, \quad 0.11338, \quad -0.42156, \quad -0.40656] \\
\tilde{\theta}_6 &= [0.03614, \quad -1.56255, \quad -1.11474, \quad -0.28192, \quad 0.55476, \quad 2.59916, \\
&\quad -0.08826, \quad 0.10361, \quad -1.05898, \quad -1.74665, \quad -0.05094, \quad 0.25421]
\end{aligned} \tag{A.19}$$

The nearest-neighbor coupled quadratic fermion Hamiltonian in one spatial dimension known as the Kitaev chain can be written in terms of Majorana fermions as:

$$H = i \sum_{j=1}^{2L} \mu \gamma_{2i-1} \gamma_{2i} + \frac{(t+\Delta)}{2} \gamma_{2i} \gamma_{2i+1} + \frac{(t-\Delta)}{2} \gamma_{2i-1} \gamma_{2i+1}. \tag{A.20}$$

Taking anti-periodic boundary conditions, $\gamma_{2L+1} = -\gamma_1$. This directly gives us the excitation spectrum $\varepsilon_p = \sqrt{(\mu + t \cos p)^2 + \Delta^2 \sin^2 p}$.

This model may be related to the one dimensional XY model with transverse field and periodic boundary conditions via the Jordan-Wigner transformation. In this dual spin model the coupling constants are $J_x = t + \Delta$, $J_y = t - \Delta$, and $h = \mu$.

$$H_{XY} = - \sum_{j=1}^L h Z_j + \frac{1}{2} (J_x X_j X_{j+1} + J_y Y_j Y_{j+1}). \tag{A.21}$$

When $\mu = \Delta = t = 1$ the model simplifies to that of Eq. (5.7) or Eq. (5.6) for fermions.

Appendix B: TMERA Appendix

B.1 Renormalization of temperature

In preparing a state with a DMERA circuit, an intermediate state of smaller system size is prepared after each scale transformation, which can be thought of as a course grained version of the final state. For ground state preparation of a critical these intermediate states have nearly identical local expectation values and subsystem entropies to the final state due to the scale invariance. However, for a Gibbs state the finite temperature breaks scale invariance, as we see from the different excitation probabilities at each scale. This is also reflected in the different energy density and correlation lengths of the intermediate states prepared during the course of the approximate Gibbs state preparation. These intermediate states may also be thought of as course grained versions of the final Gibbs state. Are these intermediate states also well described as thermal states at a higher temperature and smaller system size? Since we optimize the probabilities for a particular final system size its not necessary that the intermediate states would also be optimal. However, due to the scaling of the mode energies at different scales according roughly to $E_\ell \sim 2^\ell$, it's plausible that the state before the last scale transformation may look like a thermal state at twice the temperature.

B.2 Hamiltonian of Approximate Gibbs State

We can find a Hamiltonian for which our approximate Gibbs state is an exact thermal state by taking the logarithm of our density matrix, which is unambiguous up to an overall shift in the energies, so we can fix the trace of the Hamiltonian to be zero. For a high temperature state, the Hamiltonian is the leading order perturbation to the normalized identity matrix.

$$\rho_\beta = \frac{e^{-\beta H}}{\mathcal{Z}_\odot} \approx \frac{1}{L}(I - \beta H) \quad (\text{B.1})$$

In terms of our covariance matrix formalism, an analogous procedure can be used to get a quadratic fermion Hamiltonian from one of our approximate thermal states. The coupling matrix of this Hamiltonian proportional to a high temperature thermal state approximation, so the features of its two-point correlation functions become features of the Hamiltonian for which it is a thermal state. This means the couplings of this Hamiltonian will be nonlocal but decaying roughly inverse polynomial with distance. The couplings at a fixed distance will also not be strictly equal due to broken translation invariance but instead vary spatially around mean values.

$$H_\beta = \frac{-i}{\beta} \text{atanh}(-i\tilde{\Gamma}_\beta) \quad (\text{B.2})$$

The spectrum of the modes is flat due to the degeneracy of excitation probabilities for all modes at a given scale. The wavepacket modes are themselves the eigenmodes of this Hamiltonian, and so they have compact spatial support. In the Heisenberg picture, the ground state DMERA circuit can be seen as transforming an initial Hamiltonian of only single-body terms to this long range Hamiltonian. Time evolution according to this

Hamiltonian is therefore equivalent to time evolution of the trivial decoupled Hamiltonian conjugated by the DMERA circuit. This evolution will quickly deviate from evolution by the true local translational invariant Hamiltonian because the eigenmodes being evolved by the long range Hamiltonian are the wavepacketmodes which have compact support over varying lengthscales, and have a flat spectrum compared to the true Hamiltonian.

We also derive this Hamiltonian from a particular thermal state approximation, so in principle the approximate Gibbs states at different temperatures could give rise to different Hamiltonians. In principle it may be possible that for any given induced Hamiltonian, only its Gibbs state at the particular temperature of the Gibbs state approximation is a good analogue to the true thermal state of our target Hamiltonian. However, if these Hamiltonians induced at different temperatures do not deviate much, then choosing any one of them should give good approximations to our target thermal states for the whole range of thermal states.

By the construction of the Hamiltonian, it is equivalent to a product state Hamiltonian conjugated by a logarithmic depth circuit. This mean time evolution by this Hamiltonian can be efficiently fast-forwarded [202]. This is true even in the case where the circuit is not a matchgate circuit and therefore is not equivalent to some quadratic fermion model.

B.3 TFD purification

A mixed state can be purified by adding additional degrees of freedom that the state is entangled with, such that the reduced state on the original subsystem is the same mixed state. with the additional Hilbert space dimension equal at least to the rank of the mixed state. The purified state may be written as a Schmidt decomposition in the energy eigenbasis with the coefficients equal to the square root of the normalized Boltzman weights. A natural choice of purifying system is a second copy of the same thermal state, since it has

the minimal Hilbert space size and has the same eigenvalue spectrum.

$$\psi_\beta = \frac{1}{\sqrt{\mathcal{Z}_\odot}} \sum_i e^{-\beta E_i/2} |E_i\rangle_L |\tilde{E}_i\rangle_R$$

We can prepare an ansatz state for the TFD using an extension of our TMERA circuit, where each mixed qubit is half of an entangled pair of qubits. Pairs of qubits destined for the two halves of the TFD state will be entangled to the degree required that their entanglement entropy matches the desired entropy for qubits at a given scale ℓ for the target temperature T . Two copies of the DMERA circuit are then applied to each collection of appropriately entangled qubits to produce an ansatz state for the TFD at temperature T . By the product-state nature of the TMERA ansatz, these TFD states will only have Bell-pair like entanglement between the two halves.

Appendix C: Mana Appendix

C.1 Initial state sampling

For our MPS time evolution, we choose pure initial states with no mana or entanglement, ie tensor products of the twelve single qutrit stabilizer states. We also want to sample from initial states with the same energy density for local subsystems to thermalize to a consistent temperature without hydrodynamic energy transport. Looking at the energy for pairs of single qutrit stabilizers we can see what energy densities and therefore temperatures will be accessible to this choice of initial states. The energy densities for pairs of stabilizer states are grouped into the X eigenstates (1, 11, 12), the Z eigenstates (2, 9, 10), and other stabilizer states (3-8). The energies with highest degeneracy are any of the non X or Z eigenstates, which have zero energy density when paired together, or a slightly positive or negative energy density when paired with one of the X or Z eigenstates. When $h_l = 0$ states 3-8 only have nonzero energy density when paired with an X eigenstate. This means our initial state sequences will be random among states 3-8 for infinite temperature, or alternating between an X or Z eigenstate and one of the other eigenstates. The colored tables and graphs show the energy densities for different pairs of these single qutrit stabilizer states, colored so negative energies are red and positive energies are blue. Variation of the on-site longitudinal field only affects pairs with a Z eigenstate since all other single qutrit stabilizers have zero expectation value for this term.

The infinite-temperature states—those with $\epsilon_{j,j+1} = 0$ —deserve special consideration,




Figure C.1: **Top:** Average support of initial states over energy eigenstates in 8 qutrit system. The quantity $|\langle \psi_0 | E_k \rangle|$, which measures the support of an initial state $|\psi_0\rangle$ with each energy eigenstate (labeled by their energy density E_k), is averaged over stabilizer product states with the proper energy for each temperature. The relative support of initial states on different energy eigenstates roughly scales as the square root of the Boltzman weights $\sim \exp(-\beta E_k/2)$. **Bottom:** Normalized density of states with 0.2 width bins for energy density.




Figure C.2: 50 qutrit MPS, half-system entropy convergence with bond dimension and trotter step.




Figure C.3: 50 qutrit MPS, subsystem entropy convergence with bond dimension and Trotter step for a single initial state at each temperature.

both for their simplicity and their importance. Consider the six eigenstates of ZX and ZX^2 . call them $|\phi_\alpha\rangle$, and write

$$|\phi_\alpha\phi_\beta\rangle = |\phi_\alpha\rangle \otimes |\phi_\beta\rangle . \quad (\text{C.1})$$

Then term-by-term

$$\langle\phi_\alpha\phi_\beta|\varepsilon_{j,j+1}|\phi_\alpha\phi_\beta\rangle = 0 , \quad (\text{C.2})$$

for every choice of ϕ_α, ϕ_β , and any state of the form

$$|\phi_{\mathbf{ff}}\rangle = |\phi_{\alpha_1}\rangle \otimes |\phi_{\alpha_2}\rangle \otimes \dots \otimes |\phi_{\alpha_L}\rangle \quad (\text{C.3})$$

has zero energy density everywhere. Our infinite-temperature states, then, are $|\phi_{\mathbf{ff}}\rangle$ for random strings α .

The initial states are chosen to be stabilizer states, so that they have zero initial mana, and product states, so that there is no initial entanglement and all subsystems are pure. The energy of these initial states are determined by all the nearest neighbor pairs of stabilizer states. The only stabilizer states with nonzero expectation value of Hamiltonian terms are the eigenstates of the generalized Z and X operators. The transverse and longitudinal field terms each have one low energy eigenvstate, denoted X_0 and Z_0 respectively. These terms also have two degenerate positive energy eigenstates denoted X_1 and X_2 for the transverse field and Z_1 and Z_2 for the longitudinal field. The states Z_i also have nonzero expectation value for the bond terms, with negative energy when the nearest neighbor pairs are the same and positive when they are different. The other six stabilizer states are denoted S_i and have zero expectation with all Hamiltonian terms.

We consider in Fig. C.1 the magnitude of the inner product of initial states with energy eigenstates in an 8-qutrit system, averaged over initial states for a given temperature. We find that initial states have support over the full range of energy eigenstates, with the aver-

age support scaling roughly as the square root of the appropriate Boltzmann weights at that temperature, i.e. $|\langle \psi_0 | E_k \rangle| \sim \exp(-\beta E_k/2)$.

The exception is for the case where $\beta = 0.233$ ($E_0/L = -2$), where there are many energy eigenstates orthogonal to all of the stabilizer product states at that temperature. This makes sense because this temperature has the fewest number of initial states and significant correlations in those initial states imposed by the energy density selection. Unlike at other temperatures these constraints for $\beta = 0.233$ has restricted the initial states to have support over only a subspace of the energy eigenstates. The the average support within this subspace, however, still scales according to the appropriate Boltzmann weights.

C.2 Defining and calculating mana

Consider a system of ℓ qudits each of dimension d . (NB this notation is different from that of Sec. 6.3 and Appendices C.7, C.8. There we write d for the dimension of a whole Hilbert space; here we write d for the dimension of the component qubits.)

To calculate the mana, first recall the generalized Pauli strings

$$T_{\mathbf{p}\mathbf{q}} = \omega^{2^{-1}\mathbf{p}\cdot\mathbf{q}} \prod_j Z_j^{p_j} X_j^{q_j} \quad (\text{C.4})$$

where $\omega = e^{2\pi i/d}$, $2^{-1} = (d+1)/2$ is the multiplicative inverse of 2 in the field $\mathbb{Z}_d$, and Z_j and X_j are the clock and shift operators on site j . Just as the Pauli strings generalize continuous single-particle Weyl operators to a product of discrete rings, we can generalize the Wigner function to a discrete Wigner function

$$W_{\mathbf{p}\mathbf{q}}(\rho) = d^{-\ell} \sum_{\mathbf{p}'\mathbf{q}'} \omega^{\mathbf{p}\cdot\mathbf{q}' - \mathbf{p}'\cdot\mathbf{q}} \text{Tr} \rho T_{\mathbf{p}'\mathbf{q}'} \quad (\text{C.5})$$




Figure C.4: 50 qutrit MPS, subsystem mana convergence with bond dimension and Trotter step for the same single state at each temperature as in Fig. C.3

This may also be written in terms of the *phase-space point operators* $A_{\mathbf{pq}}$ as

$$W_{\mathbf{pq}}(\rho) = d^{-\ell} \text{Tr}[\rho A_{\mathbf{pq}}], \quad (\text{C.6})$$

where

$$A_{\mathbf{pq}} = T_{\mathbf{pq}} A_0 T_{\mathbf{pq}}^\dagger \quad \text{with} \quad A_0 = d^{-\ell} \sum_{\mathbf{pq}} T_{\mathbf{pq}}. \quad (\text{C.7})$$

The discrete Wigner function has $\sum_{\mathbf{pq}} W_{\mathbf{pq}}(\rho) = \text{Tr} \rho = 1$, and—for classical mixtures of stabilizer states σ_1, σ_2 —

$$W_{\mathbf{pq}}(\alpha \sigma_1 + (1 - \alpha) \sigma_2) \geq 0. \quad (\text{C.8})$$

(Among pure states *only* stabilizer states have $W_{\mathbf{p},\mathbf{q}}(\rho) \geq 0$ for all $\mathbf{p}, \mathbf{q}$. This property is called the *discrete Hudson's theorem*. [180]) The Wigner norm

$$\|\rho\|_W = \sum_{\mathbf{pq}} |W_{\mathbf{pq}}(\rho)| \quad (\text{C.9})$$

measures the size of the negative part of the Wigner function, and the mana is

$$\mathcal{M}(\rho) = \ln \|\rho\|_W. \quad (\text{C.10})$$

C.3 Convergence in Bond Dimension and Trotter Step

The accuracy of simulations of time evolution using matrix product states is limited by both the bond dimension of the MPS and the Trotter step size of the TEBD evolution. Trotterization of the time-evolution operator to second order with time step ηt incurs an error of order δt^3 , but a truncation also occurs at each time step; this truncation projects the

state back into the space of MPS of the given bond dimension and causes additional error.

As the initial product state evolves in time entanglement grows in the system but is ultimately limited by the bond dimension of the MPS. The largest entanglement in the system would occur across a central bipartition, so we compare the entanglement entropy for this bipartition with several different bond dimensions and Trotter step sizes for a particular unentangled zero energy initial state of a 50 site MPS. We see in Fig. C.2 that the half system entanglement entropy for MPS evolution of bond dimension $\chi = 512$ and $\chi = 1024$ diverge at a time $t \sim 6$, while Trotter error is not significant (at either bond dimension) until later times $t \sim 8$. We therefore consider our simulations reliable

While the half system entropy suggests global properties of the MPS begin to diverge around time $t = 6$, entropies and mana for subsystem of size $l \leq 6$ are much better converged as seen in Figures C.3, C.4. No noticeable difference in these quantities is seen for the entire simulated duration up to $t = 12$ for the smallest subsystems, with slight differences noted at later times for the largest subsystems. The spread in entropy of the six qutrit subsystem between these different samples at time $t = 12$ is 2.7%.

We also see from these examples looking at a single state that while mana of larger systems is always greater or equal to mana of subsystems, the entropy of subsystems can grow at different rates for early times. In our examples at finite energy density the one and six qutrit subsystems have a slower entropy growth than the intermediate subsystems. This is due to the initial state on the edges of our subsystem and how this affects the mixing for initial dynamics. In the case of our zero energy states all sites have single qutrit stabilizers which are neither X nor Z eigenstates and mix at identical rates. For the finite energy density cases our initial states are different eigenstates of the X or Z operators, and so the initial mixing rate depends on the exact pairing of stabilizer states on each edge of the subsystem. For later times once the systems have been sufficiently mixed the larger subsystems have larger entropies as expected.

C.4 Subsystem entropy

We often make reference to the entropy deficit Δ of our subsystems, which is defined as the maximal entropy of the subsystem minus the second Rényi entropy of that subsystem. In Fig. C.5 we plot the average Rényi entropy of the central subsystems of our 50 qutrit MPS over the course of our TEBD evolution which is used to determine the entropy deficit. We also plot in dashed lines the subsystem entropies of Gibbs states, which subsystems should converge to in the process of thermalization. Note that these thermal entropies are not the maximal subsystem entropies, so the asymptotic entropy deficit is still nonzero even after subsystems have converged to thermal values of the entropy. The infinite temperature case is an exception, since the thermal entropies are maximal and the entropy deficit is zero.

We see that small subsystems quickly converge to their thermal values and stabilize at this entropy. However, at later times we begin to see decreases in subsystem entropies which arise from truncation errors in the MPS evolution and are not reflective of accurate dynamics of the subsystem entropies. Unfortunately, some of the larger subsystems fail to reach their thermal entropy values before these truncation errors become relevant and start to decrease subsystem entropy.

C.5 Thermalization, typicality, and the eigenstate thermalization hypothesis: intuition

Imagine evolving a tensor product state $|\psi\rangle$ by a Hamiltonian H . Since $|\psi\rangle$ is the eigenstate of some local Hamiltonian, it has energy uncertainty

$$[\langle\psi|H^2|\psi\rangle - \langle\psi|H|\psi\rangle^2]^{1/2} = \Delta\sqrt{L}, \quad (\text{C.11})$$




Figure C.5: 50 qutrit MPS, subsystem entropy measured in trits, infinite and finite temperature initial states, $J = h_x = h_z = 1$.

where Δ is some $O(1)$ constant with dimensions energy. The time-evolved state is

$$|\psi(t)\rangle = \sum_j e^{-iE_j t} \langle E_j | \psi \rangle |E_j\rangle$$

only has weight near the initial energy $E = \langle \psi | H | \psi \rangle$ of $|\psi\rangle$. The phase factors $e^{-iE_j t}$ break the delicate conspiracy between overlaps $\langle E_j | \psi \rangle$ that results in the state ψ at the initial time $t = 0$, and we can think of the state as randomly chosen from a distribution on energy eigenstates with width given by (C.11). That distribution, in turn, is similar to a microcanonical distribution with the same width.

This cartoon assumes that the eigenstates have no structure. If they do, the overlaps $\langle E_j | \psi \rangle$ will be weighted towards eigenstates that resemble the initial state, leading to a failure of thermalization.

Suppose the system has a classical limit. Berry's conjecture [134, 203] is that the eigenstates are appropriately structureless if the corresponding classical system is chaotic. (More precisely the energy eigenfunctions are random in such a way that the Wigner function, averaged over a small phase space volume, matches the microcanonical ensemble.)

But Berry's conjecture does not contemplate the discrete Wigner functions we work with! In order to take the average over a small phase space volume, we would have to work with large onsite Hilbert space dimension d . More broadly the small- d model does not have an obvious classical limit, chaotic or integrable.

If the system does not have a classical limit, as ours does not, Berry's conjecture becomes the eigenstate thermalization hypothesis (ETH) [133]—that local expectation values in eigenstates match thermal expectation values. One expects the system to thermalize if its Hamiltonian satisfies the ETH, and not otherwise. The ETH is a hypothesis, not a conjecture like Berry's conjecture: a particular Hamiltonian may or may not satisfy the ETH, so it must be checked (generally numerically.) The message of our Fig. 6.2 is that our




Figure C.6: **Subsystem mana of Gibbs states** with inverse temperature $\beta = -0.25$ for different system sizes. The subsystem mana converges rapidly with system size which allows us to use these small system estimates for our large scale MPS simulations.

Hamiltonian at our parameters does satisfy the ETH.

For a thorough and accessible review see [135], especially Sec. 2.4 on the semiclassical limit and Berry’s conjecture and Sec. 4.1-4.2 on the eigenstate thermalization hypothesis. [204] offers a review of recent work on quantum chaos in the semiclassical limit from a very different point of view.

C.6 Gibbs state mana

We use exact diagonalization on a system of up to nine sites with periodic boundary conditions to determine the thermal properties of subsystems for the larger qutrit chain. We use the energy densities of these systems to map initial state energies to temperatures. Fig. C.6

We would like to compare the asymptotic mana of subsystems that have thermalized with nonzero initial energy density to the expected value of thermal subsystem mana. From

these plots we can extract the subsystem mana for a given inverse temperature as well as the energy density for a given inverse temperature. From these we can find the subsystem mana we expect for an initial state of a given energy density.

C.7 Bounding the difference in mana between nearby states

Proposition 1. *Suppose two density matrices ρ, σ on a Hilbert space of dimension d are nearby in trace distance:*

$$\|\rho - \sigma\|_1 \leq \eta .$$

Then

$$\|\rho - \sigma\|_W \leq d^2 \eta \quad [\text{linear bound}]$$

$$\|\rho - \sigma\|_W \leq \sqrt{d\eta} \quad [\text{square root bound}].$$

Applying the reverse triangle inequality to this result will give

Corollary C.1.

$$\left| \|\rho\|_W - \|\sigma\|_W \right| \leq \min(d^2 \eta, \sqrt{d\eta}) . \quad (\text{C.12})$$

and subadditivity of the logarithm will give

Corollary C.2.

$$\ln \|\rho\|_W \leq \ln \|\sigma\|_W + \min \left(2 \ln d + \ln \eta, \frac{1}{2} \ln d + \frac{1}{2} \ln \delta \right) . \quad (\text{C.13})$$

Now to prove 1.

Proof. Start with the linear bound. The difference in Wigner functions is

$$\begin{aligned} |W_{pq}(\rho) - W_{pq}(\sigma)| &= d^{-1} |\text{Tr} A_{pq}(\rho - \sigma)| \\ &\leq d^{-1} \|A_{pq}\|_1 \|\rho - \sigma\|_1 . \end{aligned} \quad (\text{C.14})$$

Since the phase-space point operators A_{pq} of (C.7) are unitary, $\|A_{pq}\|_1 = d$, so

$$|W_{pq}(\rho) - W_{pq}(\sigma)| \leq \eta \quad (\text{C.15})$$

and

$$\|\rho - \sigma\|_W = \sum_{pq} |W_{pq}(\rho) - W_{pq}(\sigma)| \leq d^2 \eta . \quad (\text{C.16})$$

as desired.

Turn to the square root bound. Write

$$\Delta := \rho - \sigma ; \quad (\text{C.17})$$

the hypothesis is that $\|\Delta\|_1 \leq \eta$. Then, using Cauchy-Schwarz,

$$\|\Delta\|_W = \sum_{pq} |W_{pq}(\Delta)| \leq d \left[\sum_{pq} W_{pq}^2(\Delta) \right]^{1/2} \quad (\text{C.18})$$

But

$$\sum_{pq} W_{pq}^2 = \frac{1}{d} \|\Delta\|_2^2 \leq \frac{1}{d} \|\Delta\|_1 \quad (\text{C.19})$$

since $-1 \preceq \Delta \preceq 1$, so

$$\|\Delta\|_W \leq \sqrt{d\eta} \quad (\text{C.20})$$

as desired. □

C.8 Mana and canonical typicality

Let us use the result of Popescu et al.[186] to be more precise about how what canonical typicality means for mana.

That result is as follows. The subspace of interest is the vector space spanned by the eigenvectors within $\Delta\sqrt{L}/2$ of E ; call it

$$R = \text{span}\{|E_j\rangle : |E_j - E| < \Delta\sqrt{L}/2\} ,$$

and its dimension d_R . The relevant microcanonical ensemble is

$$\mathcal{E} = \frac{1}{d_E} \sum_{|E_j - E| < \Delta\sqrt{L}/2} |E_j\rangle\langle E_j| . \quad (\text{C.21})$$

Call the subsystem of interest S , and its Hilbert space dimension d_S ; call the rest of the system $B = \bar{S}$. The microcanonical ensemble $\mathcal{E}$ traces down to

$$\Omega_S = \text{Tr}_B \mathcal{E} \quad (\text{C.22})$$

on S ; the effective accessible dimension on B is

$$d_E^{\text{eff}} = \frac{1}{\text{Tr} \Omega_B^2} , \quad (\text{C.23})$$

where $\Omega_B = \text{Tr}_S \mathcal{E}$. Finally, write μ for the Haar measure on the vector (sub)space R . Then

$$\mu\left(|\phi\rangle \in R : \|\rho_S(\phi) - \Omega_S\|_1 \geq \eta\right) \leq \eta' \quad (\text{C.24})$$

with

$$\begin{aligned}\eta &= \varepsilon + \frac{1}{2}3^{L/2}\sqrt{d_E^{\text{eff}}} \\ \eta' &= 4e^{-Cd_R^2\varepsilon}.\end{aligned}\tag{C.25}$$

Now use (C.24) and (C.25) together with Cor. C.1 to bound the difference in Wigner norms between $\rho_S(\phi)$ and $|\Omega_S\rangle$. Write

$$\mathcal{W}_\phi = \ln\|\rho_S(\phi)\|_W$$

$$\mathcal{W}_\Omega = \ln\|\Omega_S\|_W$$

for the subsystem mana of the randomly chosen state $|\phi\rangle$ and the microcanonical ensemble Ω , respectively. Applying Cor. C.1 from App. C.7, we find that

$$\mu\left(\phi \in R : \mathcal{W}_\phi - \mathcal{W}_\rho \geq \delta\right) \leq \eta' \tag{C.26}$$

with

$$\delta = \min(d_S^2\eta, \sqrt{d\eta}), \tag{C.27}$$

η as in Eq. (C.25).

To get some intuition for what this result means, let us look for the η (and hence δ) corresponding to $\eta' = 1/2$: that is, (an upper bound on) the median discrepancy in Wigner norm. This $\eta' = 1/2$ gives

$$\varepsilon = \frac{\ln 8}{Cd_R^2}$$

for

$$\begin{aligned}\eta &= \varepsilon + \frac{1}{2} \sqrt{\frac{d_S}{d_E^{\text{eff}}}} \\ &\leq \frac{\ln 8}{Cd_R^2} + \frac{d_S}{2\sqrt{d_R}}\end{aligned}$$

using that [bound on d_E^{eff} from Popescu]. The second term dominates the first, so

$$\eta \lesssim \frac{d_S}{2\sqrt{d_R}}.$$

(Were η' sufficiently small this would not be the case.) Now

$$\delta \approx \frac{d_S}{\sqrt{2\sqrt{d_R}}} \sim d_S d_R^{-1/4}. \quad (\text{C.28})$$

C.9 Krylov subspace evolution

We want to corroborate our stories about thermalization and subsystem mana in small enough systems where more exact methods than TEBD evolution of MPS states can be done for long times. We also need the system to be large enough to see thermalization of the smallest subsystems and finite size correction emerging for larger subsystems, which is challenging using exact diagonalization of a chain of qutrit sites. So, we turn to Krylov subspace methods to time evolve states of our Potts model for intermediate system sizes and high-precision long time dynamics. The Potts model is the same as in Eq. 6.1, but with periodic boundary conditions and the coefficients $J = h_x = h_z = 1$. We evolve the product stabilizer initial states for an $L = 11$ qutrit chain up to time $t = 100$ and measure the mana and entropy of subsystems less than half the system size.

The time evolved state which has thermalized is compared to the Gibbs state values

obtained through exact diagonalization of an 8 qutrit system. The subsystem entropies asymptotically approach the thermal values, with deficits visible in four and five qutrit systems due to finite size effects. See Fig. C.7.

The entropy of the $\beta = 0.135$ states are lower and have larger standard deviation between different initial states than other temperatures. The entropies of the $\beta = 0.233$ states have fluctuations which are the same across the initial state, which have more correlation due to their relatively lower statistics from energy constraints. These fluctuations are maintained in averaging and result in a low standard deviation.

For subsystem mana, subsystems smaller than four which have zero thermal mana match, with larger subsystems having finite size corrections. However, for $\beta = 0.233$ and $\beta = -0.227$, which have small but nonzero mana in most subsystems, have significantly more mana than thermal values even in smaller subsystems. See Fig. C.8 for comparison of late time subsystem mana and entropy deficit with thermal subsystem values.

What we also notice for the Gibbs states is that for $-0.12 \leq \beta \leq 0.2$ all subsystems have zero mana, thus the Gibbs state itself may be inside the stabilizer hull. For our zero energy initial states, most subsystems are maximally mixed with finite size effects reducing the entropy only for the largest subsystems. The maximally mixed state is in a sense at the center of the stabilizer hull and far from any states with mana, so all subsystems of the late time states for zero energy initial states have zero mana except for the $\ell = 5$ subsystem where finite size effects are seen.

Our Gibbs states with $\beta = 0.135, -0.96$ are also inside the stabilizer hull, but closer to the boundary. The late time subsystems for these initial states also have zero mana for most subsystems, though the $\ell = 4$ subsystems noticeably have mana where the same sized subsystems for zero energy initial states did not. This could still be from finite size effects, but being closer to the stabilizer hull boundary the introduction of mana does not require as large of a deviation from the true Gibbs state subsystems.




Figure C.7: 11 qutrit Krylov evolution, subsystem entropy in trits up to $t = 100$, thermal subsystem values of 8-qutrit system dashed.




Figure C.8: 11 qutrit Krylov evolution at $t = 100$ compared with finite temperature subsystems of an 8 qutrit chain. Entropy deficit in trits versus subsystem mana, with the relationship for Haar random states given by the solid black line.

For Gibbs states with $\beta = 0.233, -0.237$ subsystems do have a small amount of mana and lie just outside the stabilizer hull. We see however an excess of mana in subsystems over the Gibbs state values even for small subsystems. These states, although close to the Gibbs state subsystems in some sense, have enough 'wiggle room' in a small neighborhood to amount significantly more mana.

Our subsystem mana estimates for thermal states were carried out by numerically finding the exact Gibbs state for small systems and then computing the mana of various sized subsystems. We see in Fig. C.6 that the subsystem mana converge rapidly with overall system size and so we can accurately use these values as estimates of the subsystem mana for large system thermal states. Moreover, the finite size effects we do see increase the subsystem mana, so it is unlikely that these values would underestimate the mana of thermal subsystems. This shows further that the large excess of mana seen in long time simulations is due to canonical typicality effects of time evolved pure states.

Bibliography

- [1] Ashley Montanaro. Quantum algorithms: an overview. *npj Quantum Information*, 2(1):15023, January 2016. ISSN 2056-6387. doi: 10.1038/npjqi.2015.23. URL <https://doi.org/10.1038/npjqi.2015.23>.
- [2] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, page 212–219, New York, NY, USA, 1996. Association for Computing Machinery. ISBN 0897917855. doi: 10.1145/237814.237866. URL <https://doi.org/10.1145/237814.237866>.
- [3] P.W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, 1994. doi: 10.1109/SFCS.1994.365700.
- [4] Michael A. Nielsen, Isaac Chuang, and Lov K. Grover. Quantum Computation and Quantum Information. *American Journal of Physics*, 70(5):558–559, 05 2002. ISSN 0002-9505. doi: 10.1119/1.1463744. URL <https://doi.org/10.1119/1.1463744>.
- [5] Pawel Wocjan and Shengyu Zhang. Several natural bqp-complete problems, 2006.
- [6] Dorit Aharonov and Itai Arad. The BQP-hardness of approximating the jones polynomial. *New Journal of Physics*, 13(3):035019, mar 2011. doi: 10.1088/1367-2630/13/3/035019. URL <https://doi.org/10.1088/1367-2630/13/3/035019>.
- [7] Gorjan Alagic, Stephen P. Jordan, Robert König, and Ben W. Reichardt. Estimating turaev-viro three-manifold invariants is universal for quantum computation. *Physical Review A*, 82(4), oct 2010. doi: 10.1103/physreva.82.040302. URL <https://doi.org/10.1103/physreva.82.040302>.
- [8] Stephen P. Jordan, Hari Krovi, Keith S. M. Lee, and John Preskill. BQP-completeness of scattering in scalar quantum field theory. *Quantum*, 2:44, jan

2018. doi: 10.22331/q-2018-01-08-44. URL <https://doi.org/10.22331/q-2018-01-08-44>.
- [9] Tameem Albash and Daniel A. Lidar. Adiabatic quantum computation. *Reviews of Modern Physics*, 90(1), jan 2018. doi: 10.1103/revmodphys.90.015002. URL <https://doi.org/10.1103/revmodphys.90.015002>.
- [10] Michael H. Freedman, Alexei Kitaev, Michael J. Larsen, and Zhenghan Wang. Topological quantum computation, 2002.
- [11] H. J. Briegel, D. E. Browne, W. Dür, R. Raussendorf, and M. Van den Nest. Measurement-based quantum computation. *Nature Physics*, 5(1):19–26, January 2009. ISSN 1745-2481. doi: 10.1038/nphys1157. URL <https://doi.org/10.1038/nphys1157>.
- [12] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandão, Fernando G. S. L. Brandão, David A. Buell, B. Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Brooks Foxen, Austin G. Fowler, Craig Gidney, Marissa Giustina, R. Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Michael J. Hartmann, Alan Ho, Markus R. Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander N. Korotkov, Alexander N. Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, E. Lucero, Dmitry I. Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matt McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John Platt, Chris Quintana, Eleanor Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, John M. Martinis, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779): 505–510, 2019. doi: 10.1038/S41586-019-1666-5.
- [13] Han-Sen Zhong, Hui Wang, Yu-Hao Deng, Ming-Cheng Chen, Li-Chao Peng, Yi-Han Luo, Jian Qin, Dian Wu, Xing Ding, Yi Hu, Peng Hu, Xiao-Yan Yang, Wei-Jun Zhang, Hao Li, Yuxuan Li, Xiao Jiang, Lin Gan, Guangwen Yang, Lixing You, Zhen Wang, Li Li, Nai-Le Liu, Chao-Yang Lu, and Jian-Wei Pan. Quantum computational advantage using photons. *Science*, 370(6523):1460–1463, 2020. doi: 10.1126/science.abe8770. URL <https://www.science.org/doi/abs/10.1126/science.abe8770>.

- [14] Lars S. Madsen, Fabian Laudenbach, Mohsen Falamarzi. Askarani, Fabien Rortais, Trevor Vincent, Jacob F. F. Bulmer, Filippo M. Miatto, Leonhard Neuhaus, Lukas G. Helt, Matthew J. Collins, Adriana E. Lita, Thomas Gerrits, Sae Woo Nam, Varun D. Vaidya, Matteo Menotti, Ish Dhand, Zachary Vernon, Nicolás Quesada, and Jonathan Lavoie. Quantum computational advantage with a programmable photonic processor. *Nature*, 606(7912):75–81, June 2022. ISSN 1476-4687. doi: 10.1038/s41586-022-04725-x. URL <https://doi.org/10.1038/s41586-022-04725-x>.
- [15] Peter W. Shor. Fault-tolerant quantum computation. *Proceedings of 37th Conference on Foundations of Computer Science*, pages 56–65, 1996. URL <https://api.semanticscholar.org/CorpusID:7508572>.
- [16] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018. ISSN 2521-327X. doi: 10.22331/q-2018-08-06-79. URL <http://arxiv.org/abs/1801.00862>. arXiv: 1801.00862.
- [17] Richard Phillips Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21:467–488, 1999.
- [18] Andrew J. Daley, Immanuel Bloch, Christian Kokail, Stuart Flannigan, Natalie Pearson, Matthias Troyer, and Peter Zoller. Practical quantum advantage in quantum simulation. *Nature*, 607(7920):667–676, July 2022. ISSN 0028-0836. doi: 10.1038/s41586-022-04940-6.
- [19] Bela Bauer, Sergey Bravyi, Mario Motta, and Garnet Kin-Lic Chan. Quantum algorithms for quantum chemistry and quantum materials science. *Chemical Reviews*, 120(22):12685–12717, oct 2020. doi: 10.1021/acs.chemrev.9b00829. URL <https://doi.org/10.1021%2Facs.chemrev.9b00829>.
- [20] Ivan Kassal, Stephen P. Jordan, Peter J. Love, Masoud Mohseni, and Alán Aspuru-Guzik. Polynomial-time quantum algorithm for the simulation of chemical dynamics. *Proceedings of the National Academy of Sciences*, 105(48):18681–18686, dec 2008. doi: 10.1073/pnas.0808245105. URL <https://doi.org/10.1073/pnas.0808245105>.
- [21] Stephen P. Jordan, Keith S. M. Lee, and John Preskill. Quantum algorithms for quantum field theories. *Science*, 336(6085):1130–1133, jun 2012. doi: 10.1126/science.1217069. URL <https://doi.org/10.1126%2Fscience.1217069>.
- [22] John Preskill. Simulating quantum field theory with a quantum computer, 2018.
- [23] Daniel Gottesman. The Heisenberg Representation of Quantum Computers. *arXiv:quant-ph/9807006*, July 1998. URL <http://arxiv.org/abs/quant-ph/9807006>. arXiv: quant-ph/9807006.

- [24] Scott Aaronson and Daniel Gottesman. Improved Simulation of Stabilizer Circuits. *Physical Review A*, 70(5):052328, November 2004. ISSN 1050-2947, 1094-1622. doi: 10.1103/PhysRevA.70.052328. URL <http://arxiv.org/abs/quant-ph/0406196>. arXiv: quant-ph/0406196.
- [25] Simon Anders and Hans J. Briegel. Fast simulation of stabilizer circuits using a graph-state representation. *Physical Review A*, 73:022334, February 2006. ISSN 1050-2947. doi: 10.1103/PhysRevA.73.022334. URL <http://adsabs.harvard.edu/abs/2006PhRvA..73b2334A>.
- [26] Richard Jozsa and Akimasa Miyake. Matchgates and classical simulation of quantum circuits. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 464(2100):3089–3106, jul 2008. doi: 10.1098/rspa.2008.0189. URL <https://doi.org/10.1098/rspa.2008.0189>.
- [27] Norbert Schuch and Bela Bauer. Matrix product state algorithms for gaussian fermionic states. *Physical Review B*, 100(24), dec 2019. doi: 10.1103/physrevb.100.245121. URL <https://doi.org/10.1103/physrevb.100.245121>.
- [28] A. Jahn, M. Gluza, F. Pastawski, and J. Eisert. Holography and criticality in matchgate tensor networks. *Science Advances*, 5(8):eaaw0092, 2019. doi: 10.1126/sciadv.aaw0092. URL <https://www.science.org/doi/abs/10.1126/sciadv.aaw0092>.
- [29] Sergey Bravyi and Alexei Kitaev. Universal quantum computation with ideal clifford gates and noisy ancillas. *Physical Review A*, 71(2), feb 2005. doi: 10.1103/physreva.71.022316. URL <https://doi.org/10.1103/physreva.71.022316>.
- [30] M. Hebenstreit, R. Jozsa, B. Kraus, S. Strelchuk, and M. Yoganathan. All pure fermionic non-gaussian states are magic states for matchgate computations. *Physical Review Letters*, 123(8), aug 2019. doi: 10.1103/physrevlett.123.080503. URL <https://doi.org/10.1103/physrevlett.123.080503>.
- [31] E. Knill. Fault-tolerant postselected quantum computation: Schemes, 2004.
- [32] Robert Raussendorf, Daniel E. Browne, and Hans J. Briegel. Measurement-based quantum computation on cluster states. *Physical Review A*, 68(2), aug 2003. doi: 10.1103/physreva.68.022312. URL <https://doi.org/10.1103/physreva.68.022312>.
- [33] Guifré Vidal. Efficient classical simulation of slightly entangled quantum computations. *Physical Review Letters*, 91(14), oct 2003. doi: 10.1103/physrevlett.91.147902. URL <https://doi.org/10.1103/physrevlett.91.147902>.
- [34] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac. Matrix product state representations, 2007.

- [35] Steven R. White. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.*, 69:2863–2866, 1992. doi: 10.1103/PhysRevLett.69.2863.
- [36] Ulrich Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics*, 326(1):96–192, jan 2011. doi: 10.1016/j.aop.2010.09.012. URL <https://doi.org/10.1016%2Fj.aop.2010.09.012>.
- [37] M. Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles. Variational quantum algorithms. *Nature Reviews Physics*, 3(9): 625–644, September 2021. ISSN 2522-5820. doi: 10.1038/s42254-021-00348-9. URL <https://www.nature.com/articles/s42254-021-00348-9>. Number: 9 Publisher: Nature Publishing Group.
- [38] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O’Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5(1), jul 2014. doi: 10.1038/ncomms5213. URL <https://doi.org/10.1038%2Fncomms5213>.
- [39] Jules Tilly, Hongxiang Chen, Shuxiang Cao, Dario Picozzi, Kanav Setia, Ying Li, Edward Grant, Leonard Wossnig, Ivan Rungger, George H. Booth, and Jonathan Tennyson. The variational quantum eigensolver: a review of methods and best practices, 2021. URL <https://arxiv.org/abs/2111.05176>. arXiv:2111.05176.
- [40] Troy Sewell, Aniruddha Bapat, and Stephen Jordan. Estimating gate complexities for the site-by-site preparation of fermionic vacua, 2022.
- [41] Troy J. Sewell, Ning Bao, and Stephen P. Jordan. Variational quantum simulation of the critical Ising model with symmetry averaging. *Phys. Rev. A*, 107:042620, Apr 2023. doi: 10.1103/PhysRevA.107.042620. URL <https://link.aps.org/doi/10.1103/PhysRevA.107.042620>.
- [42] Troy J. Sewell and Stephen P. Jordan. Preparing renormalization group fixed points on nisq hardware, 2021. URL <https://arxiv.org/abs/2109.09787>. arXiv:2109.09787.
- [43] Troy J. Sewell, Christopher David White, and Brian Swingle. Thermal multi-scale entanglement renormalization ansatz for variational gibbs state preparation, 2022.
- [44] Troy J. Sewell and Christopher David White. Mana and thermalization: Probing the feasibility of near-clifford hamiltonian simulation. *Physical Review B*, 106(12), sep 2022. doi: 10.1103/physrevb.106.125130. URL <https://doi.org/10.1103%2Fphysrevb.106.125130>.
- [45] Stephen P Jordan, Keith SM Lee, and John Preskill. Quantum algorithms for fermionic quantum field theories. *arXiv preprint arXiv:1404.7115*, 2014.

- [46] A.Y. Kitaev, A. Shen, and M.N. Vyalyi. *Classical and Quantum Computation*. Graduate studies in mathematics. American Mathematical Society, 2002. ISBN 9780821832295. URL <https://books.google.com/books?id=08vZYhafYEAC>.
- [47] Ali Hamed Moosavian and Stephen Jordan. Faster quantum algorithm to simulate fermionic quantum field theory. *Physical Review A*, 98(1):012332, 2018.
- [48] Martin Schwarz, Kristan Temme, and Frank Verstraete. Preparing projected entangled pair states on a quantum computer. *Physical review letters*, 108(11):110502, 2012.
- [49] Ali Hamed Moosavian, James R Garrison, and Stephen P Jordan. Site-by-site quantum state preparation algorithm for preparing vacua of fermionic lattice field theories. *arXiv preprint arXiv:1911.03505*, 2019.
- [50] Theodore J Yoder, Guang Hao Low, and Isaac L Chuang. Fixed-point quantum search with an optimal number of queries. *Physical review letters*, 113(21):210501, 2014.
- [51] Sergey Bravyi and David Gosset. Complexity of quantum impurity problems. *Communications in Mathematical Physics*, 356(2):451–500, 2017.
- [52] Daniel Nagaj, Pawel Wocjan, and Yong Zhang. Fast amplification of qma, 2009.
- [53] Ake Björck and Gene Golub. Numerical methods for computing angles between linear subspaces. *Mathematics of Computation*, 27:123, 07 1973. doi: 10.2307/2005662.
- [54] Brian G. Swingle and Yixu Wang. Recovery map for fermionic gaussian channels. *Journal of Mathematical Physics*, 60(7):072202, jul 2019. doi: 10.1063/1.5093326. URL <https://doi.org/10.1063%2F1.5093326>.
- [55] G. Vidal. Class of quantum many-body states that can be efficiently simulated. *Phys. Rev. Lett.*, 101:110501, Sep 2008. doi: 10.1103/PhysRevLett.101.110501. URL <https://link.aps.org/doi/10.1103/PhysRevLett.101.110501>.
- [56] G. Evenbly and G. Vidal. Algorithms for entanglement renormalization. *Physical Review B*, 79:144108, 2009. arXiv:0707.1454.
- [57] Glen Evenbly and Guifre Vidal. Quantum criticality with the multi-scale entanglement renormalization ansatz, 2011. URL <https://arxiv.org/abs/1109.5334>.
- [58] Glen Evenbly and Steven R. White. Entanglement Renormalization and Wavelets. *Physical Review Letters*, 116:140403, Apr 2016. doi: 10.1103/PhysRevLett.116.140403. URL <https://link.aps.org/doi/10.1103/PhysRevLett.116.140403>.

- [59] Jutho Haegeman, Brian Swingle, Michael Walter, Jordan Cotler, Glen Evenbly, and Volkher B. Scholz. Rigorous free-fermion entanglement renormalization from wavelet theory. *Phys. Rev. X*, 8:011003, Jan 2018. doi: 10.1103/PhysRevX.8.011003. URL <https://link.aps.org/doi/10.1103/PhysRevX.8.011003>.
- [60] Freek Witteveen, Volkher Scholz, Brian Swingle, and Michael Walter. Quantum circuit approximations and entanglement renormalization for the dirac field in 1+1 dimensions, 2019.
- [61] Isaac H. Kim and Brian Swingle. Robust entanglement renormalization on a noisy quantum computer, 2017. arXiv:1711.07500.
- [62] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2):023023, feb 2016. doi: 10.1088/1367-2630/18/2/023023. URL <https://doi.org/10.1088%2F1367-2630%2F18%2F2%2F023023>.
- [63] Qiang Miao and Thomas Barthel. A quantum-classical eigensolver using multi-scale entanglement renormalization, 2021. URL <https://arxiv.org/abs/2108.13401>.
- [64] Tobias J. Osborne and Alexander Stottmeister. Quantum simulation of conformal field theory, 2021. URL <https://arxiv.org/abs/2109.14214>.
- [65] Kazuhiro Seki, Tomonori Shirakawa, and Seiji Yunoki. Symmetry-adapted variational quantum eigensolver. *Physical Review A*, 101(5), may 2020. doi: 10.1103/physreva.101.052340. URL <https://doi.org/10.1103%2Fphysreva.101.052340>.
- [66] Takashi Tsuchimochi, Masaki Taii, Taisei Nishimaki, and Seiichiro L. Ten-no. Adaptive construction of shallower quantum circuits with quantum spin projection for fermionic systems, 2022. URL <https://arxiv.org/abs/2205.07097>.
- [67] Sergey Bravyi. Lagrangian representation for fermionic linear optics. *Quantum Info. Comput.*, 5(3):216–238, may 2005. ISSN 1533-7146.
- [68] Djordje Radicevic. Spin structures and exact dualities in low dimensions, 2018. URL <https://arxiv.org/abs/1809.07757>.
- [69] Kristan Temme, Sergey Bravyi, and Jay Gambetta. Error Mitigation for Short-Depth Quantum Circuits. *Physical Review Letters*, 119:180509, Nov 2017. doi: 10.1103/PhysRevLett.119.180509. URL <https://link.aps.org/doi/10.1103/PhysRevLett.119.180509>.

- [70] Stuart Hadfield, Zhihui Wang, Bryan O’Gorman, Eleanor Rieffel, Davide Venturelli, and Rupak Biswas. From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms*, 12(2):34, feb 2019. doi: 10.3390/a12020034. URL <https://doi.org/10.3390%2Fa12020034>.
- [71] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm, 2014.
- [72] Wen Wei Ho and Timothy H. Hsieh. Efficient variational simulation of non-trivial quantum states. *SciPost Physics*, 6(3), Mar 2019. ISSN 2542-4653. doi: 10.21468/scipostphys.6.3.029. URL <http://dx.doi.org/10.21468/SciPostPhys.6.3.029>.
- [73] D. Zhu, S. Johri, N. M. Linke, K. A. Landsman, N. H. Nguyen, C. H. Alderete, A. Y. Matsuura, T. H. Hsieh, and C. Monroe. Generation of thermofield double states and critical ground states with a quantum computer, 2020.
- [74] Edward Farhi, David Gamarnik, and Sam Gutmann. The quantum approximate optimization algorithm needs to see the whole graph: Worst case examples, 2020. URL <https://arxiv.org/abs/2005.08747>.
- [75] Johannes Borregaard, Matthias Christandl, and Daniel Stilck Franca. Noise-robust exploration of many-body quantum states on near-term quantum devices. *NPJ Quantum Information*, 7:45, 2021. arXiv:1909.04786.
- [76] Robert N. C. Pfeifer, Glen Evenbly, and Guifré Vidal. Entanglement renormalization, scale invariance, and quantum criticality. *Phys. Rev. A*, 79:040301, Apr 2009. doi: 10.1103/PhysRevA.79.040301. URL <https://link.aps.org/doi/10.1103/PhysRevA.79.040301>.
- [77] Omar Shehab, Isaac H. Kim, Nhung H. Nguyen, Kevin Landsman, Cinthia H. Alderete, Daiwei Zhu, C. Monroe, and Norbert M. Linke. Noise reduction using past causal cones in variational quantum algorithms, 2019.
- [78] V. Giovannetti, S. Montangero, and R. Fazio. Quantum Multiscale Entanglement Renormalization Ansatz Channels. *Physical Review Letters*, 101(18), Oct 2008. ISSN 1079-7114. doi: 10.1103/physrevlett.101.180503. URL <http://dx.doi.org/10.1103/PhysRevLett.101.180503>.
- [79] Frank Verstraete, Michael M. Wolf, and J. Ignacio Cirac. Quantum computation and quantum-state engineering driven by dissipation. *Nature Physics*, 5:633–636, 2009. arXiv:0803.1447.
- [80] A. Beige. Quantum computing using dissipation. In *Institute of Physics Conference Series 173*, page 35, 2003. arXiv:quant-ph/0306125.

- [81] A. Beige, D. Braun, B. Tregenna, and P. L. Knight. Quantum computing using dissipation to remain in a decoherence-free subspace. *Physical Review Letters*, 85:1762, 2000. arXiv:quant-ph/0004043.
- [82] Paolo Zanardi, Jeffrey Marshall, and Lorenzo Campos Venuti. Dissipative universal Lindbladian simulation. *Physical Review A*, 93:022312, 2016. arXiv:1506.04311.
- [83] M. Kliesch, T. Barthel, C. Gogolin, M. Kastoryano, and J. Eisert. A dissipative quantum Church-Turing theorem. *Physical Review Letters*, 107:120501, 2011. arXiv:1105.3986.
- [84] B. Kraus, H. P. Buchler, S. Diehl, A. Kantian, A. Micheli, and P. Zoller. Preparation of entangled states by quantum Markov processes. *Physical Review A*, 78:042307, 2008. arXiv:0803.1463.
- [85] Jeffrey Marshall, Lorenzo Campos Venuti, and Paolo Zanardi. Modular quantum information processing by dissipation. *Physical Review A*, 94:052339, 2016.
- [86] Ilya Sinayskiy and Francesco Petruccione. Dissipative quantum computing with open quantum walks. In *AIP Conference Proceedings 1633*, page 186, 2014.
- [87] Ilya Sinayskiy and Francesco Petruccione. Efficiency of open quantum walk implementation of dissipative quantum computing algorithms. *Quantum Information Processing*, 11(5):1301–1309, 2012. arXiv:1401.6658.
- [88] Isaac H. Kim. Holographic quantum simulation, 2017.
- [89] Michael Foss-Feig, David Hayes, Joan M. Dreiling, Caroline Figgatt, John P. Gaebler, Steven A. Moses, Juan M. Pino, and Andrew C. Potter. Holographic quantum algorithms for simulating correlated spin systems. arXiv:2005.03023, 2020.
- [90] Galit Anikeeva, Isaac H. Kim, and Patrick Hayden. Recycling qubits in near-term quantum computers. arXiv:2012.01676, 2020.
- [91] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, and et al. Demonstration of the trapped-ion quantum CCD computer architecture. *Nature*, 592(7853):209–213, Apr 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-03318-4. URL <http://dx.doi.org/10.1038/s41586-021-03318-4>.
- [92] Miguel Aguado and Guifré Vidal. Entanglement renormalization and topological order. *Physical Review Letters*, 100(7), feb 2008. doi: 10.1103/physrevlett.100.070404. URL <https://doi.org/10.1103/physrevlett.100.070404>.
- [93] Johannes M. Oberreuter and Stefan Kehrein. Representation of excited states and topological order of the toric code in mera, 2015.

- [94] H. Bombin. An introduction to topological quantum codes, 2013.
- [95] Steven R. White and E. Miles Stoudenmire. Multisliced gausslet basis sets for electronic structure. *Physical Review B*, 99(8), feb 2019. doi: 10.1103/physrevb.99.081110.
- [96] Yiheng Qiu and Steven R. White. Hybrid gausslet/gaussian basis sets, 2021. URL <https://arxiv.org/abs/2103.02734>.
- [97] Justin Reyes and Miles Stoudenmire. A multi-scale tensor network architecture for classification and regression, 2020. URL <https://arxiv.org/abs/2001.08286>.
- [98] Iris Cong, Soonwon Choi, and Mikhail D. Lukin. Quantum convolutional neural networks. *Nature Physics*, 15(12):1273–1278, aug 2019. doi: 10.1038/s41567-019-0648-8.
- [99] Freek Witteveen and Michael Walter. Bosonic entanglement renormalization circuits from wavelet theory. *SciPost Physics*, 10(6), Jun 2021. ISSN 2542-4653. doi: 10.21468/scipostphys.10.6.143. URL <http://dx.doi.org/10.21468/SciPostPhys.10.6.143>.
- [100] Colin J Trout, Muyuan Li, Mauricio Gutiérrez, Yukai Wu, Sheng-Tao Wang, Lum-ing Duan, and Kenneth R Brown. Simulating the performance of a distance-3 surface code in a linear ion trap. *New Journal of Physics*, 20(4):043038, apr 2018. doi: 10.1088/1367-2630/aab341. URL <https://doi.org/10.1088%2F1367-2630%2Faab341>.
- [101] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6-7):467–488, June 1982. ISSN 0020-7748, 1572-9575. doi: 10.1007/BF02650179. URL <http://link.springer.com/10.1007/BF02650179>.
- [102] Julia Kempe, Alexei Kitaev, and Oded Regev. The Complexity of the Local Hamiltonian Problem. *SIAM Journal on Computing*, 35(5):1070–1097, January 2006. ISSN 0097-5397. doi: 10.1137/S0097539704445226. URL <https://epubs.siam.org/doi/abs/10.1137/S0097539704445226>. Publisher: Society for Industrial and Applied Mathematics.
- [103] Brielin Brown, Steven T. Flammia, and Norbert Schuch. Computational Difficulty of Computing the Density of States. *Physical Review Letters*, 107(4):040501, July 2011. doi: 10.1103/PhysRevLett.107.040501. URL <https://link.aps.org/doi/10.1103/PhysRevLett.107.040501>. Publisher: American Physical Society.
- [104] Sergey Bravyi, Anirban Chowdhury, David Gosset, and Pawel Wocjan. On the complexity of quantum partition functions, October 2021. URL <http://arxiv.org/abs/2110.15466>. Number: arXiv:2110.15466 arXiv:2110.15466 [quant-ph].

- [105] Barbara M. Terhal and David P. DiVincenzo. Problem of equilibration and the computation of correlation functions on a quantum computer. *Physical Review A*, 61(2):022301, January 2000. doi: 10.1103/PhysRevA.61.022301. URL <https://link.aps.org/doi/10.1103/PhysRevA.61.022301>. Publisher: American Physical Society.
- [106] David Poulin and Pawel Wocjan. Sampling from the Thermal Quantum Gibbs State and Evaluating Partition Functions with a Quantum Computer. *Physical Review Letters*, 103(22):220502, November 2009. doi: 10.1103/PhysRevLett.103.220502. URL <https://link.aps.org/doi/10.1103/PhysRevLett.103.220502>. Publisher: American Physical Society.
- [107] Ersen Bilgin and Sergio Boixo. Preparing Thermal States of Quantum Systems by Dimension Reduction. *Physical Review Letters*, 105(17):170405, October 2010. doi: 10.1103/PhysRevLett.105.170405. URL <https://link.aps.org/doi/10.1103/PhysRevLett.105.170405>. Publisher: American Physical Society.
- [108] K. Temme, T. J. Osborne, K. G. Vollbrecht, D. Poulin, and F. Verstraete. Quantum Metropolis sampling. *Nature*, 471(7336):87–90, March 2011. ISSN 1476-4687. doi: 10.1038/nature09770. URL <https://www.nature.com/articles/nature09770>. Number: 7336 Publisher: Nature Publishing Group.
- [109] Man-Hong Yung and Alán Aspuru-Guzik. A quantum–quantum Metropolis algorithm. *Proceedings of the National Academy of Sciences*, 109(3):754–759, January 2012. doi: 10.1073/pnas.1111758109. URL <https://www.pnas.org/doi/full/10.1073/pnas.1111758109>. Publisher: Proceedings of the National Academy of Sciences.
- [110] Arnau Riera, Christian Gogolin, and Jens Eisert. Thermalization in Nature and on a Quantum Computer. *Physical Review Letters*, 108(8):080402, February 2012. doi: 10.1103/PhysRevLett.108.080402. URL <https://link.aps.org/doi/10.1103/PhysRevLett.108.080402>. Publisher: American Physical Society.
- [111] Anirban Narayan Chowdhury and Rolando D. Somma. Quantum algorithms for Gibbs sampling and hitting-time estimation, March 2016. URL <http://arxiv.org/abs/1603.02940>. Number: arXiv:1603.02940 arXiv:1603.02940 [quant-ph].
- [112] Michael J. Kastoryano and Fernando G. S. L. Brandão. Quantum Gibbs Samplers: The Commuting Case. *Communications in Mathematical Physics*, 344(3):915–957, June 2016. ISSN 1432-0916. doi: 10.1007/s00220-016-2641-8. URL <https://doi.org/10.1007/s00220-016-2641-8>.
- [113] Fernando G. S. L. Brandão and Michael J. Kastoryano. Finite Correlation Length Implies Efficient Preparation of Quantum Thermal States. *Communications in Mathematical Physics*, 365(1):1–16, January 2019. ISSN 1432-0916. doi: 10.1007/s00220-018-3150-8. URL <https://doi.org/10.1007/s00220-018-3150-8>.

- [114] William Cottrell, Ben Freivogel, Diego M. Hofman, and Sagar F. Lokhande. How to build the thermofield double state. *Journal of High Energy Physics*, 2019(2):58, February 2019. ISSN 1029-8479. doi: 10.1007/JHEP02(2019)058. URL [https://doi.org/10.1007/JHEP02\(2019\)058](https://doi.org/10.1007/JHEP02(2019)058).
- [115] Guillaume Verdon, Jacob Marks, Sasha Nanda, Stefan Leichenauer, and Jack Hidary. Quantum Hamiltonian-Based Models and the Variational Quantum Thermalizer Algorithm. Technical Report arXiv:1910.02071, arXiv, October 2019. URL <http://arxiv.org/abs/1910.02071>. arXiv:1910.02071 [quant-ph] type: article.
- [116] Jingxiang Wu and Timothy H. Hsieh. Variational thermal quantum simulation via thermofield double states. *Physical Review Letters*, 123(22), Nov 2019. ISSN 1079-7114. doi: 10.1103/physrevlett.123.220502. URL <http://dx.doi.org/10.1103/PhysRevLett.123.220502>.
- [117] John Martyn and Brian Swingle. Product spectrum ansatz and the simplicity of thermal states. *Physical Review A*, 100(3), Sep 2019. ISSN 2469-9934. doi: 10.1103/physreva.100.032107. URL <http://dx.doi.org/10.1103/PhysRevA.100.032107>.
- [118] Anirban N. Chowdhury, Guang Hao Low, and Nathan Wiebe. A Variational Quantum Algorithm for Preparing Quantum Gibbs States. Technical Report arXiv:2002.00055, arXiv, January 2020. URL <http://arxiv.org/abs/2002.00055>. arXiv:2002.00055.
- [119] Mario Motta, Chong Sun, Adrian T. K. Tan, Matthew J. O’Rourke, Erika Ye, Austin J. Minnich, Fernando G. S. L. Brandão, and Garnet Kin-Lic Chan. Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution. *Nature Physics*, 16(2):205–210, February 2020. ISSN 1745-2481. doi: 10.1038/s41567-019-0704-4. URL <https://www.nature.com/articles/s41567-019-0704-4>. Number: 2 Publisher: Nature Publishing Group.
- [120] Bela Bauer, Sergey Bravyi, Mario Motta, and Garnet Kin-Lic Chan. Quantum Algorithms for Quantum Chemistry and Quantum Materials Science. *Chemical Reviews*, October 2020. doi: 10.1021/acs.chemrev.9b00829. URL <https://pubs.acs.org/doi/pdf/10.1021/acs.chemrev.9b00829>. Publisher: American Chemical Society.
- [121] R. Sagastizabal, S. P. Premaratne, B. A. Klaver, M. A. Rol, V. Negîrneac, M. Moreira, X. Zou, S. Johri, N. Muthusubramanian, M. Beekman, C. Zachariadis, V. P. Ostroukh, N. Haider, A. Bruno, A. Y. Matsuura, and L. DiCarlo. Variational preparation of finite-temperature states on a quantum computer. Technical report, arXiv, December 2020. URL <http://arxiv.org/abs/2012.03895>. arXiv:2012.03895.

- [122] Jin-Guo Liu, Liang Mao, Pan Zhang, and Lei Wang. Solving quantum statistical mechanics with variational autoregressive networks and quantum circuits. *Machine Learning: Science and Technology*, 2(2):025011, February 2021. ISSN 2632-2153. doi: 10.1088/2632-2153/aba19d. URL <https://doi.org/10.1088/2632-2153/aba19d>. Publisher: IOP Publishing.
- [123] Shi-Ning Sun, Mario Motta, Ruslan N. Tazhigulov, Adrian T.K. Tan, Garnet Kin-Lic Chan, and Austin J. Minnich. Quantum Computation of Finite-Temperature Static and Dynamical Properties of Spin Systems Using Quantum Imaginary Time Evolution. *PRX Quantum*, 2(1):010317, February 2021. doi: 10.1103/PRXQuantum.2.010317. URL <https://link.aps.org/doi/10.1103/PRXQuantum.2.010317>. Publisher: American Physical Society.
- [124] Xue-Yi Guo, Shang-Shu Li, Xiao Xiao, Zhong-Cheng Xiang, Zi-Yong Ge, He-Kang Li, Peng-Tao Song, Yi Peng, Kai Xu, Pan Zhang, Lei Wang, Dong-Ning Zheng, and Heng Fan. Thermal variational quantum simulation on a superconducting quantum processor, July 2021. URL <http://arxiv.org/abs/2107.06234>. arXiv:2107.06234.
- [125] Dominik S. Wild, Dries Sels, Hannes Pichler, Cristian Zanoci, and Mikhail D. Lukin. Quantum Sampling Algorithms, Phase Transitions, and Computational Complexity. *Physical Review A*, 104(3):032602, September 2021. ISSN 2469-9926, 2469-9934. doi: 10.1103/PhysRevA.104.032602. URL <http://arxiv.org/abs/2109.03007>. arXiv:2109.03007 [cond-mat, physics:quant-ph].
- [126] Youle Wang, Guangxi Li, and Xin Wang. Variational quantum Gibbs state preparation with a truncated Taylor series. *Physical Review Applied*, 16(5):054035, November 2021. ISSN 2331-7019. doi: 10.1103/PhysRevApplied.16.054035. URL <http://arxiv.org/abs/2005.08797>. arXiv:2005.08797.
- [127] Thais de Lima Silva, Márcio M. Taddei, Stefano Carrazza, and Leandro Aolita. Fragmented imaginary-time evolution for intermediate-scale quantum signal processors, January 2022. URL <http://arxiv.org/abs/2110.13180>. arXiv:2110.13180.
- [128] Zohreh Davoudi, Niklas Mueller, and Connor Powers. Toward Quantum Computing Phase Diagrams of Gauge Theories with Thermal Pure Quantum States, August 2022. URL <http://arxiv.org/abs/2208.13112>. arXiv:2208.13112.
- [129] Leo Zhou, Sheng-Tao Wang, Soonwon Choi, Hannes Pichler, and Mikhail D. Lukin. Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices. *Phys. Rev. X*, 10:021067, Jun 2020. doi: 10.1103/PhysRevX.10.021067. URL <https://link.aps.org/doi/10.1103/PhysRevX.10.021067>.

- [130] G. Vidal. Entanglement renormalization. *Physical Review Letters*, 99(22), nov 2007. doi: 10.1103/physrevlett.99.220405. URL <https://doi.org/10.1103%2Fphysrevlett.99.220405>.
- [131] Philippe Di Francesco, Pierre Mathieu, and David Sénéchal. *Conformal Field Theory*. Springer New York, 1997. doi: 10.1007/978-1-4612-2256-9. URL <https://doi.org/10.1007/978-1-4612-2256-9>.
- [132] E. Knill. Fermionic linear optics and matchgates, 2001. URL <https://arxiv.org/abs/quant-ph/0108033>.
- [133] J. M. Deutsch. Quantum statistical mechanics in a closed system. *Physical Review A*, 43(4):2046–2049, February 1991. doi: 10.1103/PhysRevA.43.2046. URL <https://link.aps.org/doi/10.1103/PhysRevA.43.2046>.
- [134] Mark Srednicki. Chaos and quantum thermalization. *Physical Review E*, 50(2):888–901, August 1994. doi: 10.1103/PhysRevE.50.888. URL <https://link.aps.org/doi/10.1103/PhysRevE.50.888>.
- [135] Luca D’Alessio, Yariv Kafri, Anatoli Polkovnikov, and Marcos Rigol. From Quantum Chaos and Eigenstate Thermalization to Statistical Mechanics and Thermodynamics. *Advances in Physics*, 65(3):239–362, May 2016. ISSN 0001-8732, 1460-6976. doi: 10.1080/00018732.2016.1198134. URL <http://arxiv.org/abs/1509.06411>. arXiv: 1509.06411.
- [136] P. M. Chaikin and T. C. Lubensky. *Principles of Condensed Matter Physics*. Cambridge University Press, Cambridge, October 2000. ISBN 978-0-521-79450-3.
- [137] Dieter Forster. *Hydrodynamic fluctuations, broken symmetry, and correlation functions*. Frontiers in physics ; 47. W.A. Benjamin, Advanced Book Program, Reading, Mass., 1975. ISBN 0-8053-2561-1 978-0-8053-2561-4 0-8053-2560-3 978-0-8053-2560-7. Section: xix, 326 pages : illustrations ; 25 cm.
- [138] Christopher David White. Effective dissipation rate in a Liouvillean graph picture of high-temperature quantum hydrodynamics. *arXiv:2108.00019 [cond-mat]*, July 2021. URL <http://arxiv.org/abs/2108.00019>. arXiv: 2108.00019.
- [139] C. W. von Keyserlingk, Frank Pollmann, and Tibor Rakovszky. Operator backflow and the classical simulation of quantum transport. *arXiv:2111.09904 [cond-mat]*, November 2021. URL <http://arxiv.org/abs/2111.09904>. arXiv: 2111.09904.
- [140] Javier Lopez-Piqueres, Brayden Ware, Sarang Gopalakrishnan, and Romain Vasseur. Hydrodynamics of non-integrable systems from relaxation-time approximation. *arXiv:2005.13546 [cond-mat, physics:quant-ph]*, July 2020. URL <http://arxiv.org/abs/2005.13546>. arXiv: 2005.13546.

- [141] Tomaz Prosen and Marko Znidarič. Matrix product simulations of non-equilibrium steady states of quantum spin chains. *Journal of Statistical Mechanics: Theory and Experiment*, 2009(02):P02035, February 2009. ISSN 1742-5468. doi: 10.1088/1742-5468/2009/02/P02035. URL <http://arxiv.org/abs/0811.4188>. arXiv: 0811.4188.
- [142] S. Langer, F. Heidrich-Meisner, J. Gemmer, I. P. McCulloch, and U. Schollwöck. Real-time study of diffusive and ballistic transport in spin- $\frac{1}{2}$ chains using the adaptive time-dependent density matrix renormalization group method. *Physical Review B*, 79(21):214409, June 2009. doi: 10.1103/PhysRevB.79.214409. URL <https://link.aps.org/doi/10.1103/PhysRevB.79.214409>. Publisher: American Physical Society.
- [143] Marko Žnidarič, Tomaž Prosen, Giuliano Benenti, Giulio Casati, and Davide Rossini. Thermalization and ergodicity in one-dimensional many-body open quantum systems. *Physical Review E*, 81(5):051135, May 2010. doi: 10.1103/PhysRevE.81.051135. URL <https://link.aps.org/doi/10.1103/PhysRevE.81.051135>. Publisher: American Physical Society.
- [144] Tomaz Prosen. Matrix product solutions of boundary driven quantum chains. *Journal of Physics A: Mathematical and Theoretical*, 48(37):373001, September 2015. ISSN 1751-8113, 1751-8121. doi: 10.1088/1751-8113/48/37/373001. URL <http://arxiv.org/abs/1504.00783>. arXiv: 1504.00783.
- [145] Hendrik Weimer, Augustine Kshetrimayum, and Román Orús. Simulation methods for open quantum many-body systems. *Reviews of Modern Physics*, 93(1):015008, March 2021. ISSN 0034-6861, 1539-0756. doi: 10.1103/RevModPhys.93.015008. URL <http://arxiv.org/abs/1907.07079>. arXiv: 1907.07079.
- [146] Christopher David White, Michael Zaletel, Roger S. K. Mong, and Gil Refael. Quantum dynamics of thermalizing systems. *Physical Review B*, 97(3), jan 2018. doi: 10.1103/physrevb.97.035127. URL <https://doi.org/10.1103/PhysRevB.97.035127>.
- [147] Tibor Rakovszky, C. W. von Keyserlingk, and Frank Pollmann. Dissipation-assisted operator evolution method for capturing hydrodynamic transport. *arXiv:2004.05177 [cond-mat]*, April 2020. URL <http://arxiv.org/abs/2004.05177>. arXiv: 2004.05177.
- [148] Bingtian Ye, Francisco Machado, Christopher David White, Roger S. K. Mong, and Norman Y. Yao. Emergent hydrodynamics in Floquet quantum systems. *arXiv:1902.01859 [cond-mat, physics:quant-ph]*, July 2019. URL <http://arxiv.org/abs/1902.01859>. arXiv: 1902.01859.

- [149] Sarang Gopalakrishnan and Bahti Zakirov. Facilitated quantum cellular automata as simple models with nonthermal eigenstates and dynamics. *arXiv:1802.07729 [cond-mat, physics:quant-ph]*, February 2018. URL <http://arxiv.org/abs/1802.07729>. arXiv: 1802.07729.
- [150] V. Alba, J. Dubail, and M. Medenjak. Operator Entanglement in Interacting Integrable Quantum Systems: The Case of the Rule 54 Chain. *Physical Review Letters*, 122(25):250603, June 2019. doi: 10.1103/PhysRevLett.122.250603. URL <https://link.aps.org/doi/10.1103/PhysRevLett.122.250603>. Publisher: American Physical Society.
- [151] Jason Iaconis, Andrew Lucas, and Rahul Nandkishore. Multipole conservation laws and subdiffusion in any dimension. *arXiv:2009.06507 [cond-mat]*, September 2020. URL <http://arxiv.org/abs/2009.06507>. arXiv: 2009.06507.
- [152] Hakop Pashayan, Joel J. Wallman, and Stephen D. Bartlett. Estimating Outcome Probabilities of Quantum Circuits Using Quasiprobabilities. *Physical Review Letters*, 115(7):070501, August 2015. doi: 10.1103/PhysRevLett.115.070501. URL <https://link.aps.org/doi/10.1103/PhysRevLett.115.070501>. Publisher: American Physical Society.
- [153] Sergey Bravyi, Graeme Smith, and John A. Smolin. Trading Classical and Quantum Computational Resources. *Physical Review X*, 6:021043, April 2016. doi: 10.1103/PhysRevX.6.021043. URL <http://adsabs.harvard.edu/abs/2016PhRvX...6b1043B>.
- [154] Sergey Bravyi and David Gosset. Improved Classical Simulation of Quantum Circuits Dominated by Clifford Gates. *Physical Review Letters*, 116:250501, June 2016. ISSN 0031-9007. doi: 10.1103/PhysRevLett.116.250501. URL <http://adsabs.harvard.edu/abs/2016PhRvL.116y0501B>.
- [155] Ryan S. Bennink, Erik M. Ferragut, Travis S. Humble, Jason A. Laska, James J. Nutaro, Mark G. Pleszkoch, and Raphael C. Pooser. Unbiased simulation of near-Clifford quantum circuits. *Physical Review A*, 95:062337, June 2017. ISSN 1050-2947. doi: 10.1103/PhysRevA.95.062337. URL <http://adsabs.harvard.edu/abs/2017PhRvA...95f2337B>.
- [156] Mark Howard and Earl T. Campbell. Application of a resource theory for magic states to fault-tolerant quantum computing. *Physical Review Letters*, 118(9):090501, March 2017. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.118.090501. URL <http://arxiv.org/abs/1609.07488>. arXiv: 1609.07488.
- [157] Sergey Bravyi, Dan Browne, Padraic Calpin, Earl Campbell, David Gosset, and Mark Howard. Simulation of quantum circuits by low-rank stabilizer decompositions. *arXiv e-prints*, 1808:arXiv:1808.00128, July 2018. URL <http://adsabs.harvard.edu/abs/2018arXiv180800128B>.

- [158] Yifei Huang and Peter Love. Approximate stabilizer rank and improved weak simulation of Clifford-dominated circuits for qudits. *Physical Review A*, 99:052307, May 2019. ISSN 1050-2947. doi: 10.1103/PhysRevA.99.052307. URL <http://adsabs.harvard.edu/abs/2019PhRvA..99e2307H>.
- [159] Kaifeng Bu and Dax Enshan Koh. Efficient classical simulation of Clifford circuits with nonstabilizer input states. *arXiv e-prints*, 1902:arXiv:1902.11257, February 2019. URL <http://adsabs.harvard.edu/abs/2019arXiv190211257B>.
- [160] Lorenzo Leone, Salvatore F. E. Oliviero, You Zhou, and Alioscia Hamma. Quantum Chaos is Quantum. *Quantum*, 5:453, May 2021. ISSN 2521-327X. doi: 10.22331/q-2021-05-04-453. URL <http://arxiv.org/abs/2102.08406>. arXiv: 2102.08406.
- [161] Earl Campbell. A random compiler for fast Hamiltonian simulation. *Physical Review Letters*, 123(7):070503, August 2019. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.123.070503. URL <http://arxiv.org/abs/1811.08017>. arXiv: 1811.08017.
- [162] Andrew M. Childs, Aaron Ostrander, and Yuan Su. Faster quantum simulation by randomization. *Quantum*, 3:182, September 2019. ISSN 2521-327X. doi: 10.22331/q-2019-09-02-182. URL <http://arxiv.org/abs/1805.08385>. arXiv: 1805.08385.
- [163] Andrew M. Childs, Yuan Su, Minh C. Tran, Nathan Wiebe, and Shuchen Zhu. A Theory of Trotter Error. *Physical Review X*, 11(1):011020, February 2021. ISSN 2160-3308. doi: 10.1103/PhysRevX.11.011020. URL <http://arxiv.org/abs/1912.08854>. arXiv: 1912.08854.
- [164] Minh C. Tran, Yuan Su, Daniel Carney, and Jacob M. Taylor. Faster Digital Quantum Simulation by Symmetry Protection. *PRX Quantum*, 2(1):010323, February 2021. ISSN 2691-3399. doi: 10.1103/PRXQuantum.2.010323. URL <http://arxiv.org/abs/2006.16248>. arXiv: 2006.16248.
- [165] Paul K. Faehrmann, Mark Steudtner, Richard Kueng, Maria Kieferova, and Jens Eisert. Randomizing multi-product formulas for improved Hamiltonian simulation. *arXiv:2101.07808 [quant-ph]*, October 2021. URL <http://arxiv.org/abs/2101.07808>. arXiv: 2101.07808.
- [166] Victor Veitch, Seyed Ali Hamed Mousavian, Daniel Gottesman, and Joseph Emerson. The Resource Theory of Stabilizer Computation. *New Journal of Physics*, 16(1):013009, January 2014. ISSN 1367-2630. doi: 10.1088/1367-2630/16/1/013009. URL <http://arxiv.org/abs/1307.7171>. arXiv: 1307.7171.

- [167] Xin Wang, Mark M. Wilde, and Yuan Su. Efficiently computable bounds for magic state distillation. *arXiv e-prints*, 1812:arXiv:1812.10145, December 2018. URL <http://adsabs.harvard.edu/abs/2018arXiv181210145W>.
- [168] Bartosz Regula. Convex geometry of quantum resource quantification. *Journal of Physics A: Mathematical and Theoretical*, 51(4):045303, January 2018. ISSN 1751-8113, 1751-8121. doi: 10.1088/1751-8121/aa9100. URL <http://arxiv.org/abs/1707.06298>. arXiv: 1707.06298.
- [169] Michael Beverland, Earl Campbell, Mark Howard, and Vadym Kliuchnikov. Lower bounds on the non-Clifford resources for quantum computations. *arXiv:1904.01124 [quant-ph]*, November 2019. URL <http://arxiv.org/abs/1904.01124>. arXiv: 1904.01124.
- [170] Lorenzo Leone, Salvatore F. E. Oliviero, and Alioscia Hamma. Stabilizer Rényi entropy. *arXiv:2106.12587 [quant-ph]*, January 2022. URL <http://arxiv.org/abs/2106.12587>. arXiv: 2106.12587.
- [171] Saubhik Sarkar, Chiranjib Mukhopadhyay, and Abolfazl Bayat. Characterization of an operational quantum resource in a critical many-body system. *New Journal of Physics*, 22(8):083077, August 2020. ISSN 1367-2630. doi: 10.1088/1367-2630/aba919. URL <http://arxiv.org/abs/1908.08058>. arXiv: 1908.08058.
- [172] V. A. Fateev and A. B. Zamolodchikov. Self-dual solutions of the star-triangle relations in ZN-models. *Physics Letters A*, 92(1):37–39, October 1982. ISSN 0375-9601. doi: 10.1016/0375-9601(82)90736-8. URL <http://www.sciencedirect.com/science/article/pii/0375960182907368>.
- [173] G. von Gehlen and V. Rittenberg. Zn-symmetric quantum chains with an infinite set of conserved charges and Zn zero modes. *Nuclear Physics B*, 257:351–370, January 1985. ISSN 0550-3213. doi: 10.1016/0550-3213(85)90350-5. URL <http://www.sciencedirect.com/science/article/pii/0550321385903505>.
- [174] Helen Au-Yang, Barry M. McCoy, Jacques H. H. Perk, Shuang Tang, and Mu-Lin Yan. Commuting transfer matrices in the chiral Potts models: Solutions of star-triangle equations with genus>1. *Physics Letters A*, 123(5):219–223, August 1987. ISSN 0375-9601. doi: 10.1016/0375-9601(87)90065-X. URL <http://www.sciencedirect.com/science/article/pii/037596018790065X>.
- [175] Paul Fendley. Parafermionic edge zero modes in Z_n -invariant spin chains. *Journal of Statistical Mechanics: Theory and Experiment*, 2012(11):P11020, November 2012. ISSN 1742-5468. doi: 10.1088/1742-5468/2012/11/P11020. URL <http://arxiv.org/abs/1209.0472>. arXiv: 1209.0472.

- [176] Guifre Vidal. Efficient classical simulation of slightly entangled quantum computations. *Physical Review Letters*, 91(14), October 2003. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.91.147902. URL <http://arxiv.org/abs/quant-ph/0301063>. arXiv: quant-ph/0301063.
- [177] G. Vidal. Efficient simulation of one-dimensional quantum many-body systems. *Physical Review Letters*, 93(4), July 2004. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.93.040502. URL <http://arxiv.org/abs/quant-ph/0310089>. arXiv: quant-ph/0310089.
- [178] Christopher David White and Justin H. Wilson. Mana in Haar-random states. *arXiv:2011.13937 [cond-mat, physics:quant-ph]*, November 2020. URL <http://arxiv.org/abs/2011.13937>. arXiv: 2011.13937.
- [179] Christopher David White, ChunJun Cao, and Brian Swingle. Conformal field theories are magical. *arXiv:2007.01303 [cond-mat, physics:hep-th, physics:quant-ph]*, July 2020. URL <http://arxiv.org/abs/2007.01303>. arXiv: 2007.01303.
- [180] D. Gross. Hudson’s Theorem for finite-dimensional quantum systems. *Journal of Mathematical Physics*, 47(12):122107, December 2006. ISSN 0022-2488, 1089-7658. doi: 10.1063/1.2393152. URL <http://arxiv.org/abs/quant-ph/0602001>. arXiv: quant-ph/0602001.
- [181] David Gross. Non-negative Wigner functions in prime dimensions. *Applied Physics B*, 86(3):367–370, February 2007. ISSN 0946-2171, 1432-0649. doi: 10.1007/s00340-006-2510-9. URL <http://arxiv.org/abs/quant-ph/0702004>. arXiv: quant-ph/0702004.
- [182] Anatole Kenfack and Karol yczkowski. Negativity of the wigner function as an indicator of non-classicality. *Journal of Optics B: Quantum and Semiclassical Optics*, 6(10):396–404, aug 2004. doi: 10.1088/1464-4266/6/10/003. URL <https://doi.org/10.1088%2F1464-4266%2F6%2F10%2F003>.
- [183] Kanato Goto, Tomoki Nosaka, and Masahiro Nozaki. Chaos by Magic. *arXiv:2112.14593 [cond-mat, physics:hep-th, physics:quant-ph]*, December 2021. URL <http://arxiv.org/abs/2112.14593>. arXiv: 2112.14593.
- [184] Hal Tasaki. From Quantum Dynamics to the Canonical Distribution: General Picture and a Rigorous Example. *Physical Review Letters*, 80(7):1373–1376, February 1998. doi: 10.1103/PhysRevLett.80.1373. URL <https://link.aps.org/doi/10.1103/PhysRevLett.80.1373>. Publisher: American Physical Society.
- [185] Sheldon Goldstein, Joel L. Lebowitz, Roderich Tumulka, and Nino Zanghi. Canonical Typicality. *Physical Review Letters*, 96(5):050403, February 2006. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.96.050403. URL <http://arxiv.org/abs/cond-mat/0511091>. arXiv: cond-mat/0511091.

- [186] Sandu Popescu, Anthony J. Short, and Andreas Winter. Entanglement and the foundations of statistical mechanics. *Nature Physics*, 2(11):754–758, November 2006. ISSN 1745-2481. doi: 10.1038/nphys444. URL <https://www.nature.com/articles/nphys444>.
- [187] Peter Reimann. Typicality for Generalized Microcanonical Ensembles. *Physical Review Letters*, 99(16):160404, October 2007. doi: 10.1103/PhysRevLett.99.160404. URL <https://link.aps.org/doi/10.1103/PhysRevLett.99.160404>. Publisher: American Physical Society.
- [188] Sheldon Goldstein, Joel L. Lebowitz, Christian Mastrodonato, Roderich Tumulka, and Nino Zanghi. On the Approach to Thermal Equilibrium of Macroscopic Quantum Systems. *Physical Review E*, 81(1):011109, January 2010. ISSN 1539-3755, 1550-2376. doi: 10.1103/PhysRevE.81.011109. URL <http://arxiv.org/abs/0911.1724>. arXiv: 0911.1724.
- [189] Sheldon Goldstein, Joel L. Lebowitz, Christian Mastrodonato, Roderich Tumulka, and Nino Zanghi. Normal Typicality and von Neumann’s Quantum Ergodic Theorem. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 466(2123):3203–3224, November 2010. ISSN 1364-5021, 1471-2946. doi: 10.1098/rspa.2009.0635. URL <http://arxiv.org/abs/0907.0108>. arXiv: 0907.0108.
- [190] Peter Reimann. Generalization of von Neumann’s Approach to Thermalization. *Physical Review Letters*, 115(1):010403, July 2015. doi: 10.1103/PhysRevLett.115.010403. URL <https://link.aps.org/doi/10.1103/PhysRevLett.115.010403>. Publisher: American Physical Society.
- [191] Silvano Garnerone, Thiago R. de Oliveira, and Paolo Zanardi. Typicality in random matrix product states. *Physical Review A*, 81(3):032336, March 2010. doi: 10.1103/PhysRevA.81.032336. URL <https://link.aps.org/doi/10.1103/PhysRevA.81.032336>. Publisher: American Physical Society.
- [192] Silvano Garnerone, Thiago R. de Oliveira, Stephan Haas, and Paolo Zanardi. Statistical properties of random matrix product states. *Physical Review A*, 82(5):052312, November 2010. doi: 10.1103/PhysRevA.82.052312. URL <https://link.aps.org/doi/10.1103/PhysRevA.82.052312>. Publisher: American Physical Society.
- [193] Benoit Collins, Carlos E. Gonzalez-Guillen, and David Perez-Garcia. Matrix Product States, Random Matrix Theory and the Principle of Maximum Entropy. *Communications in Mathematical Physics*, 320(3):663–677, June 2013. ISSN 0010-3616, 1432-0916. doi: 10.1007/s00220-013-1718-x. URL <http://arxiv.org/abs/1201.6324>. arXiv: 1201.6324.

- [194] Gavin K. Brennen, Peter Rohde, Barry C. Sanders, and Sukhwinder Singh. Multi-scale quantum simulation of quantum field theory using wavelets. *Physical Review A*, 92(3), sep 2015. doi: 10.1103/physreva.92.032315. URL <https://doi.org/10.1103%2Fphysreva.92.032315>.
- [195] Mohsen Bagherimehrab, Yuval R. Sanders, Dominic W. Berry, Gavin K. Brennen, and Barry C. Sanders. Nearly optimal quantum algorithm for generating the ground state of a free quantum field theory. *PRX Quantum*, 3(2), jun 2022. doi: 10.1103/prxquantum.3.020364. URL <https://doi.org/10.1103%2Fprxquantum.3.020364>.
- [196] Tobias J. Osborne and Alexander Stottmeister. Conformal field theory from lattice fermions, 2021. URL <https://arxiv.org/abs/2107.13834>.
- [197] Brian Swingle. Entanglement Renormalization and Holography. *Physical Review D*, 86(6), Sep 2012. ISSN 1550-2368. doi: 10.1103/physrevd.86.065007. URL <http://dx.doi.org/10.1103/PhysRevD.86.065007>.
- [198] Fernando Pastawski, Beni Yoshida, Daniel Harlow, and John Preskill. Holographic quantum error-correcting codes: Toy models for the bulk/boundary correspondence. *Journal of High Energy Physics*, 2015(6), Jun 2015. ISSN 1029-8479. doi: 10.1007/jhep06(2015)149. URL [http://dx.doi.org/10.1007/JHEP06\(2015\)149](http://dx.doi.org/10.1007/JHEP06(2015)149).
- [199] A. Almheiri, X. Dong, and D. Harlow. Bulk locality and quantum error correction in AdS/CFT. *Journal of High Energy Physics*, 2015(4), Apr 2015. ISSN 1029-8479. doi: 10.1007/jhep04(2015)163. URL [http://dx.doi.org/10.1007/JHEP04\(2015\)163](http://dx.doi.org/10.1007/JHEP04(2015)163).
- [200] Isaac H. Kim and Michael J. Kastoryano. Entanglement renormalization, quantum error correction, and bulk causality. *Journal of High Energy Physics*, 2017(4), Apr 2017. ISSN 1029-8479. doi: 10.1007/jhep04(2017)040. URL [http://dx.doi.org/10.1007/JHEP04\(2017\)040](http://dx.doi.org/10.1007/JHEP04(2017)040).
- [201] Ning Bao, ChunJun Cao, Sean M. Carroll, and Aidan Chatwin-Davies. De Sitter Space as a Tensor Network: Cosmic No-Hair, Complementarity, and Complexity. *Physical Review D*, 96(12), Dec 2017. ISSN 2470-0029. doi: 10.1103/physrevd.96.123536. URL <http://dx.doi.org/10.1103/PhysRevD.96.123536>.
- [202] Shouzhen Gu, Rolando D. Somma, and Burak Ş ahinoğlu. Fast-forwarding quantum evolution. *Quantum*, 5:577, Nov 2021. doi: 10.22331/q-2021-11-15-577. URL <https://doi.org/10.22331%2Fq-2021-11-15-577>.
- [203] M. V. Berry. Regular and irregular semiclassical wavefunctions. *Journal of Physics A: Mathematical and General*, 10(12):2083–2091, December 1977. ISSN 0305-4470. doi: 10.1088/0305-4470/10/12/016. URL <https://doi.org/10.1088/0305-4470/10/12/016>. Publisher: IOP Publishing.

- [204] Klaus Richter, Juan Diego Urbina, and Steven Tomsovic. Semiclassical roots of universality in many-body quantum chaos. Technical Report arXiv:2205.02867, arXiv, May 2022. URL <http://arxiv.org/abs/2205.02867>. arXiv:2205.02867 [cond-mat, physics:hep-th, physics:nlin, physics:quant-ph] type: article.