


Simulation of Strongly Correlated Quantum Many-Body Systems

Thesis by
Ersen Bilgin

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

2011
(Defended April 6, 2011)

Copyright notice and joint work:

Chapters 1, 2, and 3 contain material from [1], [2] and [3], and are joint work with David Poulin. Chapter 4 is from [3] and is joint work with Robert König. Chapter 5 contains material from [4] and is joint work with Sergio Boixo. All articles referenced above are copyrighted by the American Physical Society.

© 2011

Ersen Bilgin

All Rights Reserved

Acknowledgements

First of all, I would like to thank my advisor John Preskill for always keeping his door wide open when I needed advice and allowing me to follow my interests from the beginning. I was very fortunate to have him as my advisor. Under his invaluable direction, IQI has been a great place to work as a graduate student. I met many amazing people who spent time here either as visitors, postdocs, or students without whose influence, this thesis would have been impossible.

I also want to thank David Poulin for guiding me in my first year of research and for his invaluable input especially in the first two chapters of this thesis. He introduced me to belief propagation, and we continued collaborating on various projects after he left IQI for *Université de Sherbrooke*.

My collaborator for Anyonic Entanglement Renormalization in chapter 4, Robert König, has been a great colleague and friend. He introduced me to anyons, and was a pleasure to work with.

I worked with Sergio Boixo for the final chapter of this thesis. I had a great time exploring various ideas on quantum thermalization with him and working out all the details of our dimension reduction algorithm.

I also express my gratitude to everybody at IQI who have made graduate school a great experience for me: Ann Harvey, Panos Aliferis, Salman Beigi, Robin Blume-Kohout, Peter Brooks, Darrick Chang, Bill Fefferman, Lukasz Fidkowski, Steve Flammia, Alexey Gorshkov, Kovid Goyal, Jeongwan Haah, Ann Harvey, Liang Jiang, Stephen Jordan, Isaac Kim, Alexei Kitaev, Netanel Lindner, Yi-Kai Liu, Prabha Mandayam, Spiros Michalakis, Hui Khoon Ng, Norbert Schuch, Sujeet Shukla, Ben Reichardt, Gil Refael, Greg Ver Steeg, Stephanie Wehner, Jon Yard, Shengyu Zhang, and Michael Zwolak.

All my teachers and professors before Caltech, especially Simon Butterworth, David Tucker-Smith and William Wootters have made learning physics and research a great source

of enjoyment for me. Without their extraordinary teaching and encouragement, I might not even have started my Ph.D.

I have to thank all the wonderful people I met outside of work at Caltech, especially Caglar Tanrikulu, Omer Durak, Ercan Gurses, Ozgun Konca, Bahar Bingol, Selim Hanay, Muruvvet Buyukboyaci, Aycan Yurtsever, and Cevat Ustun. I will definitely miss all the lunch conversations under the warm California sun (Thank you Southern California for the great weather!) and all the great times we had together.

I would also like to thank my parents for their encouragement and their unconditional support in all my decisions.

Of course, I cannot even imagine this experience without my girlfriend Hang Song. It is a great comfort to go back home every day and have her there to share highs and lows of graduate student life.

Finally, I would like to thank DoE and NSF for their financial support.

Abstract

In this thesis, we address the problem of solving for the properties of interacting quantum many-body systems in thermal equilibrium. The complexity of this problem increases exponentially with system size, limiting exact numerical simulations to very small systems. To tackle more complex systems, one needs to use heuristic algorithms that approximate solutions to these systems. Belief propagation is one such algorithm that we discuss in chapters 2 and 3. Using belief propagation, we demonstrate that it is possible to solve for static properties of highly correlated quantum many-body systems for certain geometries at all temperatures. In chapter 4, we generalize the multiscale renormalization ansatz to the anyonic setting to solve for the ground state properties of anyonic quantum many-body systems. The algorithms we present in chapters 2, 3, and 4 are very successful in certain settings, but they are not applicable to the most general quantum mechanical systems. For this, we propose using quantum computers as we discuss in chapter 5. The dimension reduction algorithm we consider in chapter 5 enables us to prepare thermal states of any quantum many-body system on a quantum computer faster than any previously known algorithm. Using these thermal states as the initialization of a quantum computer, one can study both static and dynamic properties of quantum systems without any memory overhead.

Contents

Acknowledgements	iii
Abstract	v
1 Introduction	1
1.1 Motivation and Overview	1
1.2 The Quantum Many-Body Problem	3
1.2.1 Exponential Scaling of Quantum Simulations	3
1.2.2 Density Matrix	5
1.2.3 Thermal Equilibrium	6
2 Belief Propagation	10
2.1 Introduction	10
2.2 Classical Belief Propagation	11
2.2.1 Sum-Product Algorithm	11
2.2.2 Classical Hamiltonians	12
2.3 Graphical models	14
2.4 Quantum Belief Propagation Algorithm	14
2.4.1 Convergence	15
2.5 Quantum Belief Propagation for Quantum Many-Body	16
2.5.1 Replica	17
2.5.2 Sliding Window	18
2.6 Numerical Results	20
2.7 Conclusion	23

3	Coarse-Grained Belief Propagation	25
3.1	Introduction	25
3.2	Errors in Belief Propagation Results	26
3.3	Disordered System Revisited	29
3.4	Multiscale Entanglement Renormalization	31
3.5	Coarse-Grained Belief Propagation	37
3.6	Discussion	41
4	Anyonic Entanglement Renormalization	43
4.1	Introduction	43
4.2	Anyonic States and Operators	46
4.2.1	A Unified Treatment of Topological Order	46
4.2.2	The Anyonic Hilbert Space and Anyon Diagrams	47
4.2.2.1	Anyons on a disc	49
4.2.2.2	Anyons on a torus	53
4.2.3	Anyonic Hamiltonians: Long-range Effective Theories	56
4.3	Anyonic Entanglement Renormalization	57
4.3.1	The Setting	57
4.3.2	The Ansatz	58
4.3.3	Efficient Evaluation of Physical Quantities	59
4.3.4	Computational Cost and Refinements of the Ansatz	60
4.3.5	Example: Fibonacci Anyons	61
4.3.6	Distillable States for Composite Anyon Coding	63
4.4	Application to the Golden Chain and the Majumdar-Ghosh Chain	64
4.4.1	The Model	64
4.4.2	Exact Renormalization Group Fixed Point at the Majumdar-Ghosh Point	66
4.4.3	Numerical Variation over Ansatz States	68
4.5	Braiding and More General Arrangements of Anyons	76
4.6	Conclusions	79
5	Dimension Reduction	81
5.1	Introduction	81

5.2	Thermalization Using Phase Estimation	82
5.2.1	Dimension Reduction Overview	84
5.3	Perturbative Hamiltonian Update	85
5.3.1	Perturbative Update with Perfect Operations	89
5.3.2	Concatenation	93
5.3.3	Errors and Cost of Operations with Finite Precision	94
5.4	Time Requirements	98
5.5	Conclusion	101
6	Conclusion	102
	Bibliography	103

Chapter 1

Introduction

1.1 Motivation and Overview

The main purpose of physics is to model and analyze nature to gain a deeper understanding of the universe we live in. For mathematical models of physics to be useful in practice, predictions about nature need to be extracted from the model's mathematical statements. One example of a very successful physical model is quantum mechanics. Quantum mechanics has so far has survived many rigorous tests and is believed to describe the behavior of particles at the quantum realm accurately. Information gathered about nature through quantum mechanics have led to many technological advances such as transistors and lasers.

Whether we want to test quantum mechanics or use quantum mechanics to model and understand physical systems, we need to extract information about the system of interest from the equations of quantum mechanics. Starting from some initial conditions and information about the model, we can attempt to solve these equations analytically. While this may be possible for some simple systems, for more complex systems analytical solutions are not known to exist.

To solve for the properties of general quantum systems, we need to resort to numerical simulations. However, numerical simulations of quantum systems is not an easy task, as the resources required for the simulations grow exponentially with system size. This requirement limits us to very small system sizes for brute-force calculations. One way around this problem is developing algorithms for approximating solutions that have better scaling with system size. To achieve a better scaling, we usually constrain the number of available states in our simulations.

In fact, most physical systems of interest in nature have local interactions that severely

constrain the number of states required in their description. Their underlying geometric structure allows us to devise algorithms specific for locally interacting models that are much more efficient than algorithms for the general case. Over the years, many methods such as White’s density matrix renormalization group (DMRG) method [5] have been developed for this purpose. These methods mostly focus on the low energy sectors of locally interacting quantum systems [6–8]. They all owe their success to the fact that the constrained set of states for locally interacting models can be described very accurately with a numerically feasible number of parameters. For example, matrix product states (MPS) accurately approximate ground states of one-dimensional gapped Hamiltonians [9, 10], and projected entangled pair states (PEPS) are accurate in higher dimensions at finite temperatures [11, 12].

Even though these states require few parameters, finding the right set of parameters for a given system remains a formidable task [13]. Only in special cases, such as simulating imaginary time evolution within a variational set of states [14], or using iterative minimization procedures [8], have successful heuristics been devised that solve this problem.

In this thesis, we focus on approximation algorithms that solve for thermodynamic properties of quantum many-body systems on both classical and quantum computers. After briefly discussing the problem in section 1.2, we generalize the classical belief propagation algorithm to the quantum setting in chapter 2. Belief propagation (BP) is an algorithm to solve inference problems defined on graphical models, which had previously been used in many areas such as image recognition, artificial intelligence and statistical physics. To apply this algorithm to the quantum many-body systems, we first solve for properties of small clusters in our system, and use belief propagation to exchange classical information among the neighboring clusters. After describing the algorithm, we present some numerical results to show that the quantum belief propagation method works well at high temperatures as long as the underlying graph does not contain many small loops.

In chapter 3, we combine quantum belief propagation with multiscale entanglement renormalization ansatz (MERA) to expand the applicability of belief propagation to all temperatures. MERA is an ansatz that yields an accurate coarse-graining of quantum many-body systems by efficiently discarding high energy excitations using local operations. When belief propagation stops being effective as we lower the temperature, we coarse-grain our lattice using MERA and run belief propagation on the coarse-grained lattice. In the

coarse-grained lattice, the effective size of clusters increases, making BP accurate for lower temperatures. We continue this procedure of using BP on coarser lattices until we reach zero temperature. We also demonstrate the effectiveness of this method using numerical examples.

In chapter 4, we explore the anyonic entanglement renormalization algorithm, which is the anyonic generalization of the bosonic multiscale renormalization ansatz. The Hilbert space of anyonic systems is much smaller than that of bosonic systems for the same number of particles due to the constraints arising from the fusion rules for anyons. Therefore, simulation of anyons through embedding on a bosonic Hilbert space wastes scarce resources. Motivated by the success of MERA on bosonic systems, we discuss anyonic entanglement renormalization, which projects anyonic systems to their ground state by recursive coarse-graining.

Finally, in chapter 5, we address the problem of preparing quantum many-body states in thermal equilibrium on quantum computers. This task is important for both preparing initial states for quantum computers and for studying static properties of thermal quantum systems. Instead of preparing the entire thermal state using a single projective measurement as was previously done, we break up the projections. We first thermalize small clusters and recursively merge these clusters using projective measurements until we have thermalized the full system. This method achieves an exponential speedup over previously known methods.

1.2 The Quantum Many-Body Problem

1.2.1 Exponential Scaling of Quantum Simulations

The main difficulty in solving for the properties of quantum many-body systems is the exponential scaling of resource requirements with the number of subsystems. This is due to the exponential size of the Hilbert space the composite system is defined in. The Hilbert space is the complex vector space with inner product that is the state space of the associated physical system. If the system is in a pure state, then it can be described by its state vector, which is a unit vector in the system's Hilbert space.

All interactions in the system can be fully described by a Hamiltonian defined on this Hilbert space. This Hamiltonian coupled with the initial state for the system $|\psi(t = 0)\rangle$

can be used to solve for the evolution of the system by plugging them into Schrödinger's equation:

$$i\hbar \frac{d|\psi(t)\rangle}{dt} = H|\psi(t)\rangle. \quad (1.1)$$

The solution is simply given by

$$|\psi(t)\rangle = e^{-iHt/\hbar}|\psi(0)\rangle. \quad (1.2)$$

In the above equations, the state $|\psi\rangle$ is a unit vector in the Hilbert space of the system, and H is a hermitian matrix, whose eigenvectors span the same Hilbert space. In particular, in the diagonal form, the Hamiltonian of the system can be written as

$$H = \sum_k E_k |\psi_k\rangle\langle\psi_k|, \quad (1.3)$$

where $|\psi_k\rangle$ are the eigenvectors of H corresponding to the eigenvalue E_k , and $\langle\psi_k|$ is the vector dual to $|\psi_k\rangle$.

The exponentiation operation in (1.2) can be written using the definition of the exponential operator on matrices as

$$e^{-iHt/\hbar} = \sum_k e^{-iE_k t/\hbar} |\psi_k\rangle\langle\psi_k|. \quad (1.4)$$

As the exponentiation operation is defined in terms of the eigenvalues and eigenstates of the Hamiltonian, diagonalization of the Hamiltonian is required for any numerical simulations of the system. The computational cost of this exponentiation operation scales polynomially with the size of the matrix H . For a single particle system, this cost can be manageable, but the size of the Hilbert space along with the size of the matrix H increases exponentially as we add more particles.

Given two systems, A and B with Hilbert spaces \mathcal{H}_A and \mathcal{H}_B , we can define a joint Hilbert space \mathcal{H}_{AB} to be the tensor product of the Hilbert spaces of the two subsystems, $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$. The size of the joint Hilbert space \mathcal{H}_{AB} is now the product of the sizes of \mathcal{H}_A and \mathcal{H}_B by the definition of the tensor product.

For N -particle systems, the Hilbert space is therefore d^N dimensional, where d is the dimension of the Hilbert space of a single particle. Any simulation of the composite system

requires the diagonalization of $d^N \times d^N$ matrices, which takes exponential time in system size.

1.2.2 Density Matrix

In this thesis, we will be focusing on calculating properties of quantum many-body systems in thermal equilibrium. For this, we first need to introduce the formalism of ensembles of states, or *density matrices*. Density matrices provide full descriptions of quantum systems whose states are not known with full accuracy. For example, if our information about a given quantum system is limited to knowing with probability p_i that the system is in the pure state $|\psi_i\rangle$ for any i , then we can write the density matrix corresponding to the system as

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|. \quad (1.5)$$

To calculate the unitary evolution of ρ , we need to unitarily evolve each of the eigenstates using (1.2). This corresponds to

$$\rho(t) = \sum_i p_i e^{-iHt/\hbar} \rho(0) e^{iHt/\hbar}. \quad (1.6)$$

Similarly, measurements on density matrices can be expressed in terms of measurements on pure states. For example, probability of getting the measurement outcome m from the density matrix ρ can be written in terms of probabilities of outcome m from the $|\psi_i\rangle$'s.

$$P(m) = \sum_i P(m|i) p_i = \sum_i p_i \langle\psi_i|E_m|\psi_i\rangle, \quad (1.7)$$

where $P(m|i) = \langle\psi_i|E_m|\psi_i\rangle$ is the probability of outcome m given that we are in state ψ_i and E_m is the measurement operator associated with outcome m . We can further use the fact that $\text{Tr}(AB) = \text{Tr}(BA)$ to write $P(m)$ directly in terms of ρ instead of the form above,

which requires the diagonalization of ρ .

$$P(m) = \sum_i p_i \langle \psi_i | E_m | \psi_i \rangle \quad (1.8)$$

$$= \sum_i p_i \text{Tr}(\langle \psi_i | E_m | \psi_i \rangle) \quad (1.9)$$

$$= \sum_i p_i \text{Tr}(E_m | \psi_i \rangle \langle \psi_i |) \quad (1.10)$$

$$= \text{Tr}(E_m \sum_i p_i | \psi_i \rangle \langle \psi_i |) \quad (1.11)$$

$$= \text{Tr}(E_m \rho) \quad (1.12)$$

Here, we have used the commutativity of the trace operation with the sum. In this form, we can calculate all properties of the density matrix, ρ , using only multiplication and trace operations.

1.2.3 Thermal Equilibrium

Now that we have all the tools we need, we proceed to calculate the eigenstates and the associated probabilities of the density matrices for quantum systems in thermal equilibrium (with a very large reservoir). We can obtain the ratios of probabilities of occupying two eigenstates ψ_A and ψ_B from the corresponding multiplicities of the reservoir, Ω_A and Ω_B respectively. Here, multiplicity is a measure of the number of accessible states in the reservoir.

The fundamental assumption of statistical mechanics is that for an isolated system all accessible microstates of the system are equally likely to be occupied. Even though our system is not isolated, the system and the reservoir together are isolated from the rest of the universe. To compare the probabilities of two eigenstates of the system, we need to compare the number of microstates accessible when the system is in each of these states. Even though we are fixing the state of the system, there are many accessible states for the reservoir, all of which are equally likely. Intuitively, when the system is in a lower (higher) energy eigenstate, the bath will have a higher (lower) energy leading to a greater (fewer) number of accessible states. As all of these states are equally likely, we expect the system to occupy lower energy eigenstates with a greater probability.

The ratio of the probabilities of the two states ψ_A and ψ_B is given by the ratio of the

multiplicities of the reservoir when the system is in the corresponding state:

$$\frac{\mathcal{P}(\psi_A)}{\mathcal{P}(\psi_B)} = \frac{\Omega_A}{\Omega_B}. \quad (1.13)$$

By definition, the entropy is given by $S = k \ln \Omega$, where k is the Boltzmann's constant. Using this, we can write

$$\frac{\mathcal{P}(\psi_A)}{\mathcal{P}(\psi_B)} = \frac{e^{S_A/k}}{e^{S_B/k}} = e^{(S_A - S_B)/k}. \quad (1.14)$$

In the simplified case, where the volume and particle number in the reservoir are constant, the change in entropy of a microstate is only related to the change in energy. In fact, temperature of a system is defined to be the reciprocal of its entropy vs. energy curve:

$$1/T \equiv \frac{\partial S}{\partial E}. \quad (1.15)$$

This is because temperature can also be defined as the property that is the same for both systems when two systems are in thermal equilibrium. When two systems are in thermal equilibrium, we are most likely to find the joint system in the macrostate that maximizes their combined entropy (natural logarithm of available microstates in the macrostate). In fact, in the limit when the size of both systems is infinite, the likelihood of finding them in a low entropy macrostate goes to zero. This follows from a simple combinatorics argument and is also another way to state the second law of thermodynamics. When entropy is maximized, any infinitesimal energy transfer between the two systems should not change the total entropy ($\partial S_{total}/\partial E = 0$). So, the entropy gain by one system must equal the entropy loss of the other one. That is, in thermal equilibrium, the slope of the entropy vs. energy curve for both systems must be the same. As we also defined temperature to be the property that is the same for both systems in thermal equilibrium, it is natural to define temperature as in (1.15).

Now, let us return to the quantum system in thermal equilibrium. We already related the probability of occupying the two states ψ_A and ψ_B to the change in the entropy of the reservoir, $S_A - S_B$, in (1.14). Assuming that the quantum system is much smaller than the reservoir, the change in the energy of the system will only cause a small change in the entropy of the reservoir. Therefore, we can use the definition of temperature from (1.15) to

relate the change in energy of the reservoir to the change in its entropy.

$$S_A - S_B = (E_A^R - E_B^R)/T, \quad (1.16)$$

where the superscript R refers to the reservoir. Using the fact that the system and reservoir are isolated from the rest of the universe, the energy change of the reservoir is opposite of the energy change of the system, $E_A^R - E_B^R = -(E_A - E_B)$. With this, we finally have the ratio

$$\frac{\mathcal{P}(\psi_A)}{\mathcal{P}(\psi_B)} = \frac{e^{-E_A/kT}}{e^{-E_B/kT}}. \quad (1.17)$$

So far, we have only derived the ratio between two eigenstates of our system in terms of their energies and the temperature. To get the probabilities associated with each eigenstate, we rewrite the above equation as

$$\frac{\mathcal{P}(\psi_A)}{e^{-E_A/kT}} = \frac{\mathcal{P}(\psi_B)}{e^{-E_B/kT}}. \quad (1.18)$$

The left hand side of this equation depends only on A (is independent of B). Similarly, the right hand side is only a function of B and is independent of A . But both sides are equal to each other, making the ratio a constant. This constant is written as $1/Z$, where Z is referred to as the partition function. Since we picked the two states A and B arbitrarily, for any state ψ_i , we have

$$\mathcal{P}(\psi_i) = \frac{1}{Z} e^{-E_i/kT}. \quad (1.19)$$

For the probabilities to be normalized, we need $Z = \sum_i e^{-E_i/kT}$, where the sum is over all the eigenstates of our system Hamiltonian.

We now have all the tools needed to calculate properties of a quantum many-body system at a thermal equilibrium. We first need to diagonalize the Hamiltonian of the system to extract information about the energies associated with each eigenstate (the eigenvalues of the Hamiltonian). Then, the probabilities in the density matrix ρ are given by (1.19). This is equivalent to exponentiating the matrix $-H/kT$:

$$\rho(T) = \sum_i p_i |\psi_i\rangle\langle\psi_i| = \frac{1}{Z} \sum_i e^{-E_i/kT} |\psi_i\rangle\langle\psi_i| = \frac{1}{Z} e^{-H/kT}, \quad (1.20)$$

where the partition function $Z = \text{Tr}(e^{-H/kT})$.

Once we have the density matrix, we can use (1.12) to calculate any property of the system. As was mentioned earlier, the exponential scaling of the size of the matrix H with the number of subsystems makes this calculation intractable for all but very small systems. The rest of this thesis will be devoted to extracting approximate information about large systems without brute-force diagonalization of the full Hamiltonian. Even though most of the thesis will be focusing on simulations at finite temperatures, taking the limit of zero temperature in all algorithms allows us to explore ground state properties.

Chapter 2

Belief Propagation

2.1 Introduction

Belief propagation (BP) is a powerful algorithm designed to solve inference problems involving a large number of random variables. It operates on graphical models, where variables are located at the vertices of a graph and edges encode dependence relations between the variables. The algorithm is exact when the underlying graph is a tree. On graphs with small loops, generalizations of BP can also provide reliable approximations [15, 16]. It is also highly parallelizable in the sense that each random variable can be associated with a different processor, and messages can be exchanged between processors that are joined by an edge [17–20].

These features have made belief propagation an important tool in numerous scientific and technological fields ranging from information theory to image recognition, and from artificial intelligence to statistical physics. Indeed, it is one of the most powerful heuristic algorithms to solve problems such as decoding of low-density and turbo error correction codes [21–23], determining the phase diagram of quenched disordered systems [20, 24], and random satisfiability problems [25, 26].

When used to solve thermal properties of classical systems, the algorithm essentially reduces to a transfer matrix solution. On more general graphs, it can be described as performing a constrained minimization of the Bethe free energy of the system [15]. Therefore, it is often a very good approximation on graphs containing no small loops. On graphs with small loops, generalizations of BP can also provide reliable approximations [15, 16].

Recently, belief propagation and graphical models were generalized to the quantum setting [27, 28]. In this chapter, we characterize the performance of belief propagation




Figure 2.1. The graph that $f(x_1, \dots, x_N) = \tilde{f}(x_1, x_2)\tilde{f}(x_2, x_3)\dots\tilde{f}(x_{N-1}, x_N)$ is defined on. Each vertex corresponds to a variable x_i and the edges correspond to the functions $\tilde{f}(x_i, x_{i+1})$ of the vertices they connect to.

when used as a heuristic algorithm to solve inference problems—e.g., compute correlation functions—in the context of finite-temperature quantum many-body physics.

2.2 Classical Belief Propagation

2.2.1 Sum-Product Algorithm

Before we generalize the belief propagation algorithm to the quantum setting, we will first discuss classical belief propagation with the aid of a simple example. Belief propagation is a message passing algorithm that solves inference problems defined on graphs. One example of such problem could be the calculation of the marginal of a function. The method we discuss is also referred to as the sum-product algorithm.

Marginal of a function $f(x_1, x_2, \dots, x_N)$ on N th of its N variables is defined as the sum over all variables except x_N :

$$g(x_N) = \sum_{x_1, x_2, \dots, x_N} f(x_1, x_2, \dots, x_N). \quad (2.1)$$

If the function f does not have any internal structure, this sum has $\mathcal{O}(e^N)$ terms, and takes $\mathcal{O}(e^N)$ time to compute. However, if the function is factorizable in the form

$$f(x_1, x_2, \dots, x_N) = \tilde{f}(x_1, x_2)\tilde{f}(x_2, x_3)\dots\tilde{f}(x_{N-1}, x_N), \quad (2.2)$$

then we say this function is defined on a graph, and we can use belief propagation to solve for its marginal $g(x_N)$. The graph the function is defined on is a chain with vertices on figure 2.1.

To solve for the marginal $g(x_N)$, we can pass messages between vertices connected to each other by an edge until all messages converge. At each iteration, all messages between every pair of vertices get updated using the messages from the previous iteration. The right-

and left-moving messages are

$$m_{i \rightarrow i+1}(x_{i+1}, t) = \sum_{x_i} \tilde{f}(x_i, x_{i+1}) m_{i-1 \rightarrow i}(x_i, t-1), \text{ and} \quad (2.3)$$

$$m_{i \rightarrow i-1}(x_{i-1}, t) = \sum_{x_i} \tilde{f}(x_{i-1}, x_i) m_{i+1 \rightarrow i}(x_i, t-1), \quad (2.4)$$

and all messages will be initialized to $m_{i \rightarrow j}(x_j, 0) = 1$. Note that all messages are functions of a single variable, and operations on at most two variables is used for each message update.

At each step of the algorithm, we can also calculate our belief about the $g(x_N)$ as

$$b(g(x_N), t) = \sum_{x_{N-1}} \tilde{f}(x_{N-1}, x_N) m_{N-2 \rightarrow N-1}(x_{N-1}, t). \quad (2.5)$$

The beliefs correspond to the best estimate of $g(x_N)$ with the available information. As we are passing messages between neighboring vertices, only local information is available before $t < N$. For our example, at time t , each vertex has received information about the states of $2t$ nearest vertices. When every vertex has received information from all other vertices in the chain (at $t = N$ in this particular example), the messages converge to their final values. At this time, our belief about the marginal equals is the correct expression, i.e., $b(g(x_N), t) = g(x_N)$.

2.2.2 Classical Hamiltonians

One direct application of the classical belief propagation algorithm is the calculation of various properties of classical spin chains. Consider a system composed of N classical spins on a chain with Hamiltonian $H(x_1, \dots, x_N) = \sum_{\langle i, j \rangle} h(x_i, x_j)$, where $\langle i, j \rangle$ denotes neighboring sites. The partition function for this system is defined as

$$\mathcal{Z}(\beta) = \sum_{x_1, \dots, x_N} e^{-\beta H(x_1, \dots, x_N)} \quad (2.6)$$

$$= \sum_{x_1, \dots, x_N} e^{-\beta h_{N, N-1}} \dots e^{-\beta h_{3,2}} e^{-\beta h_{2,1}}, \quad (2.7)$$

where we use the shorthand $h_{i,j} = h(x_i, x_j)$. A brute-force calculation of this quantity requires summing over an exponential (in N) number of terms. However, taking advantage

of the local structure of the Hamiltonian, the sum can be rearranged as follows:

$$\sum_{x_N} e^{-\beta h_{N-1,N}} \left(\dots \sum_{x_2} \left(e^{-\beta h_{2,3}} \sum_{x_1} e^{-\beta h_{1,2}} \right) \dots \right).$$

Now, the sums can be performed sequentially with each sum only involving a small number of terms,¹ allowing the computation of \mathcal{Z} in a time proportional to N . This leads to an iterative rule where messages $m_{i \rightarrow j}$ are exchanged along the edge of the chain connecting two sites i, j with the update rule

$$m_{i \rightarrow i+1}(x_{i+1}) = \sum_{x_i} e^{-\beta h(x_i, x_{i+1})} m_{i-1 \rightarrow i}(x_i). \quad (2.8)$$

With the initialization $m_{0 \rightarrow 1}(x_1) = 1$, we see that $\mathcal{Z}(\beta) = \sum_{x_N} m_{N-1 \rightarrow N}(x_N)$.

This procedure can be generalized to arbitrary graphs by defining the update rule for the message $m_{i \rightarrow j}(x_j)$, passed from site i to a neighboring site j , to be

$$m_{i \rightarrow j}(x_j) = \sum_{x_i} e^{-\beta h(x_i, x_j)} \prod_{k \in \mathcal{N}(i) \setminus j} m_{k \rightarrow i}(x_i), \quad (2.9)$$

where $\mathcal{N}(i) \setminus j$ denotes the set of neighbors of site i other than j . On a tree, these messages will converge to their final value after a time equal to the tree's diameter. The one- and two-body beliefs:

$$b_j(x_j) = \frac{1}{\mathcal{Z}(\beta)} \prod_{k \in \mathcal{N}(j)} m_{k \rightarrow j}(x_j), \quad \text{and} \quad (2.10)$$

$$b_{i,j}(x_i, x_j) = \frac{1}{\mathcal{Z}(\beta)} \prod_{k \in \mathcal{N}(i)} m_{k \rightarrow i}(x_i) \prod_{k' \in \mathcal{N}(j)} m_{k' \rightarrow j}(x_j) e^{-\beta h(x_i, x_j)} \quad (2.11)$$

are equal to the reduced one- and two-body distribution respectively, and the partition function can be evaluated from any one of their normalization. When the underlying graph contains loops, BP is no longer exact but often provides accurate approximation to the true marginal states and partition function.

¹The procedure we are describing for a chain is usually referred to as the transfer matrix method in statistical physics.

2.3 Graphical models

To generalize the classical belief propagation algorithm to the quantum setting, we first need to introduce quantum graphical models. We consider quantum graphical models (G, ρ) that consist of a graph G and an n -bifactor state ρ . The graph $G = (V, E)$ has a set of vertices V and a set of edges E . Each $v \in V$ is a quantum system, with Hilbert space \mathcal{H}_v . A n -bifactor state ρ is a positive operator on $\mathcal{H} = \bigotimes_{v \in V} \mathcal{H}_v$, that can be expressed as

$$\rho = \frac{1}{Z} \left(\prod_{v \in V} \mu_v \right) \star^{(n)} \left(\bigodot_{(u,v) \in E} \nu_{u:v} \right), \quad (2.12)$$

where Z is some normalization factor, and μ_v and $\nu_{u:v}$ are positive operators on \mathcal{H}_u and $\mathcal{H}_u \otimes \mathcal{H}_v$, respectively. The operators $\nu_{u:v}$ are required to mutually commute when n is finite. The product $\star^{(n)}$ is defined as $X \star^{(n)} Y \equiv [X^{\frac{1}{2n}} Y^{\frac{1}{n}} X^{\frac{1}{2n}}]^n$, and has the property of producing a positive operator when both X and Y are positive. This product is noncommutative except in the limit $n \rightarrow \infty$, which defines the \odot product:

$$X \odot Y \equiv \lim_{n \rightarrow \infty} X \star^{(n)} Y = e^{(\log X + \log Y)}. \quad (2.13)$$

Both products $\star^{(n)}$ and \odot reduce to normal matrix product when X and Y commute.

A generic inference problem on a graphical model is to compute the reduced density operator on a subset $W \subset V$ of the quantum systems conditioned on the fact that a measurement was performed on a disjoint subset $U \subset V$, where both W and U are of constant size. This problem turns out to be equivalent to the seemingly simpler problem of computing the reduced state on any subset W of constant size, i.e., $\rho_W = \text{Tr}_{V-W} \{\rho\}$, where Tr_X denotes the partial trace over a systems in set X . Without additional assumptions on the structure of ρ , solving this problem requires resources that grow exponentially with the number of quantum systems $|V|$. However, the solution can sometimes be obtained or approximated by QBP in a time polynomial in $|V|$.

2.4 Quantum Belief Propagation Algorithm

Given a graphical model (G, ρ) , quantum belief propagation (QBP) consists of a sequence of exchanges of operator-valued messages between neighboring vertices, which carry infor-

mation about the state at other locations in the graph. More precisely, for $(u, v) \in E$, the message passed from vertex u to vertex v at time t is an operator on \mathcal{H}_v given by

$$m_{u \rightarrow v}(t) \propto \text{Tr}_u \left\{ \mu_u \star^{(n)} \left(\nu_{u:v} \odot \bigcirc_{v' \in n(u)-v} m_{v' \rightarrow u}(t-1) \right) \right\}, \quad (2.14)$$

where $n(u)$ denotes the neighbors of u . The proportionality factor can be chosen so that $\text{Tr}\{m_{u \rightarrow v}\} = 1$, and the messages are initialized $m_{u \rightarrow v}(0) = I$. At time t , the belief $b_{uv}(t)$ —which is meant to represent some approximation of the state $\rho_{uv} = \text{Tr}_{V-uv}\{\rho\}$ for $(u, v) \in E$ —is given by

$$b_{uv}(t) \propto (\mu_u \mu_v) \star^{(n)} \left(\nu_{u:v} \odot \bigcirc_{w \in n(u)-v} m_{w \rightarrow u} \bigcirc_{y \in n(v)-u} m_{y \rightarrow v} \right), \quad (2.15)$$

where all messages are taken at time t . When all operators defining the bifactor state commute, QBP reduces to the standard belief propagation algorithm [17–20].

Since the message update rule (2.14) at vertex u depends only on the incoming messages at that vertex, the algorithm can be operated in a highly parallel fashion where each quantum system u is associated with a processor, and messages are exchanged between processors u and v iff $(u, v) \in E$. Similarly, the beliefs on the pair (u, v) , (2.15), can be computed by combining the messages received at those vertices.

2.4.1 Convergence

In [27], it was shown that when G is a tree and (G, ρ) is either (i) a 1-bifactor state [cf. (2.12) with $n = 1$] or (ii) a quantum Markov network, QBP yields the exact solution in a time proportional to the graph's diameter—i.e., $b_{uv}(t) = \rho_{uv}$ for $t \geq \text{diameter}(G)$. Intuitively, this means that the algorithm must run for a time sufficiently long to allow messages to travel between any pair of vertices. When operated on loopy graphs, the beliefs do not necessarily converge to the correct density operators. A good heuristic in that case is to halt the algorithm when $b_{uv}(t)$ become almost time independent, which also happens in a time roughly equal to the graph's diameter in all the models we have investigated.

A graphical model (G, ρ) is a quantum Markov network when the conditional independence conditions $I(U : (V - n(U) - U) | n(U)) = 0$ are met for all $U \subset V$. The quantity $I(A : B|C) = S(AC) + S(BC) - S(C) - S(ABC)$ is the quantum conditional mutual

information [29, 30], and $S(A) = \text{Tr}\{\rho_A \log_2 \rho_A\}$ is the von Neumann entropy. As explained in [27, 31], the vanishing of $I(A : C|B)$ is equivalent to the condition $\rho_{ABC} = \rho_B^{-1} \odot \rho_{AB} \odot \rho_{BC}$. This equality is not verified in general, and the Kullback-Leibler distance between the right- and left-hand side is precisely the conditional mutual information $D(\rho_{ABC} || \rho_B^{-1} \odot \rho_{AB} \odot \rho_{BC}) = I(A : C|B)$.

To understand the workings of QBP, consider a bifactor state ρ_{uvw} on the line $u - v - w$. The reduced state on w is $\rho_w \propto \text{Tr}_{uv}\{(\mu_u \mu_v \mu_w) \star^{(n)} (\nu_{u:v} \odot \nu_{v:w})\}$. When $n = 1$, basic algebra implies that $\rho_w \propto \text{Tr}_v\{(\mu_v \mu_w) \star^{(1)} (\text{Tr}_u\{\mu_u \star^{(1)} \nu_{u:v}\} \odot \nu_{v:w})\}$; the operations Tr_u and $\star^{(1)}$ commute so to say. The computation of ρ_w can thus be broken into two steps: (i) Compute $m_{u \rightarrow v} = \text{Tr}_u\{\mu_u \star^{(1)} \nu_{u:v}\}$; (ii) compute $\rho_w \propto \text{Tr}_v\{(\mu_v \mu_w) \star^{(1)} (m_{u \rightarrow v} \odot \nu_{v:w})\}$. When $n \rightarrow \infty$ on the other hand, the operations Tr_u and \odot do not commute in general, but they do precisely when $I(u : w|v) = 0$ [27]. QBP is based on a generalization of these observations to arbitrary graphs.

QBP does not rely on the vanishing of the normalized connected correlation functions $\mathcal{C}(\sigma_A, \sigma_C) = \langle \sigma_A \sigma_C \rangle - \langle \sigma_A \rangle \langle \sigma_C \rangle$ [32], or equivalently [33, 34] on the vanishing of the mutual information $I(A : C) = S(A) + S(C) - S(AC)$. In many systems, the mutual information is not *a priori* short range. For instance in the $T \rightarrow 0$ limit, the thermal state of the one-dimensional Ising model in zero transverse field is an equal mixture of all spins up and all spins down, which has $I(A : B) = 1$ between any two disjoint regions, whereas $I(A : C|B) = 0$ for any three disjoint regions. To compute thermodynamical quantities, one generally introduces a symmetry-breaking field that randomly singles out either the all-up or all-down state, which both have $I(A : B) = 0$. Symmetry-breaking can be a delicate issue—for instance, on Cayley trees where a constant fraction of vertices live on the boundary [24]—and is circumvented by QBP.

2.5 Quantum Belief Propagation for Quantum Many-Body

In the context of quantum many-body physics, the inference problem consists of computing correlation functions for the thermal state of a system of interacting particles. Given a graph $G = (V, E)$, we consider the generic Hamiltonian

$$H = \sum_{v \in V} h_v + \sum_{(u,v) \in E} h_{uv} . \quad (2.16)$$

The thermal state at inverse temperature $\beta = 1/T$ is given by $\rho = \frac{1}{Z}e^{-\beta H}$ where $Z = \text{Tr}\{e^{-\beta H}\}$ is the partition function. Defining $\mu_v = e^{-\beta h_v}$ and $\nu_{u:v} = e^{-\beta h_{uv}}$ enables us to express any such thermal state as an ∞ -bifactor state, cf. (2.12).

Despite the fact that thermal states are bifactor states, the result from [27] cited above does not imply that correlation functions can be evaluated exactly and efficiently with QBP. This is primarily because G is not necessarily a tree, but also because thermal states are neither 1-bifactor nor quantum Markov networks in general. There is no general remedy to the first hurdle, unless the loops happen to be very small and can be eliminated by merging some vertices. Thus, QBP will need to be executed on a loopy graph and it is the primary goal of this chapter to determine the effects of such loops on the performance of QBP. Two pragmatic solutions, named the replica method and sliding window QBP, have been proposed to overcome the second set of obstacles [27].

2.5.1 Replica

The general idea of the replica method is to approximate the thermal state by a 1-bifactor state on which QBP can be executed directly and is guaranteed to converge in the absence of loops. In a first step, a Trotter-Suzuki (TS) decomposition is used to approximate a thermal state by an N_τ -bifactor state with finite N_τ . This produces a systematic error that scales as β/N_τ . Then, in a fashion reminiscent of the replica trick used in the study of spin glass, the N_τ -bifactor state is replaced by a 1-bifactor state at the expense of substituting the quantum system at each vertex by N_τ replicas:

$$\mu_v \rightarrow \left(\mu_v^{\frac{1}{N_\tau}}\right)^{\otimes N_\tau} T_v^{(N_\tau)} \quad \text{and} \quad \nu_{u:v} \rightarrow \left(\nu_{u:v}^{\frac{1}{N_\tau}}\right)^{\otimes N_\tau}, \quad (2.17)$$

where $T_v^{(N_\tau)}$ is the operator that cyclicly permutes the N_τ replicas of v . The operators $\nu_{u:v} = e^{-\beta h_{uv}}$ do not commute in general, but this can be fixed in practice on sparse graphs by merging some vertices. On a tree, the TS decomposition is the only source of error, so accuracy ϵ can be achieved at a computational cost that is exponential in β/ϵ . This method is particularly useful as it allows for a direct computation of correlation functions at arbitrary distances, see [27].

2.5.2 Sliding Window

While all quantum Markov networks are thermal states of some local Hamiltonian on G [27], the converse is not true in general. Sliding window QBP is motivated by the fact that quantum Markov networks are fixed points of coarse-graining procedures. Thermal states, regarded as ∞ -bifactor states, are used directly to implement the message passing rule in (2.14) with $n = \infty$, except that messages are computed not just using the nearest neighbors but with all vertices within a distance $\leq \ell$. On a line, for instance, vertex j receives a message from $(j-1, j-2, \dots, j-\ell)$ and one from $(j+1, j+2, \dots, j+\ell)$. In that case, sliding window QBP produces the exact solution efficiently if the conditional mutual information dies off at a finite distance.

As an example, consider a chain of qubits with Hamiltonian $H = \sum_{\langle i,j \rangle} h_{i,j}$ where $h_{i,j}$ are Hermitian operators acting on site i and j of a chain. Using the \odot -product between positive operators $A \odot B = e^{\log(A) + \log(B)}$, we can write the partition function in a form very similar to (2.7)

$$\mathcal{Z}(\beta) = \text{Tr}(e^{-\beta H}) \tag{2.18}$$

$$= \text{Tr}(e^{-\beta h_{N,N-1}} \odot \dots \odot e^{-\beta h_{3,2}} \odot e^{-\beta h_{2,1}}), \tag{2.19}$$

with the sums replaced by traces and products by \odot -products. This expression cannot be reorganized like its classical counterpart because—unlike ordinary products—the \odot -product does not obey a distributive law in the sense that $\text{Tr}_a(e^{-\beta h_{c,b}} \odot e^{-\beta h_{b,a}}) \neq e^{-\beta h_{c,b}} \odot \text{Tr}_a(e^{-\beta h_{b,a}})$.

However, the distributive law holds when the chain forms a quantum Markov network [27], i.e., $I(a : c|b) = 0$ where the quantity $I(a : c|b) = S(a, c) + S(b, c) - S(b) - S(a, b, c)$ is known as the conditional mutual information and S is the von Neumann entropy. Intuitively, this condition means that all correlations between a and c are mediated through b . This does not prevent a and c from being correlated, but implies that all information learned about a by measuring c can equivalently be obtained by measuring b instead. While this condition does not hold for generic quantum interactions, it becomes approximately true when the “Markov shield” b is sufficiently thick.

When the Markov shield is sufficiently thick, we can reorganize the sum of (2.19) as in the classical case, but keeping an l -site Markov blanket between the traced-out site and the




Figure 2.2. Calculating $m_{4 \rightarrow 5}$ from $m_{3 \rightarrow 4}$ in an iteration of BP algorithm with $l = 4$. In the first step $e^{h_{4,5}}$ is added to $m_{3 \rightarrow 4}$ using the \odot -product. Then, the first spin is traced out yielding $m_{4 \rightarrow 5}$.

end of the cluster:

$$\mathcal{Z}(\beta) = \text{Tr}_{1,\dots,N} \left(e^{-\beta h_{N-1,N}} \dots \odot e^{-\beta h_{1,2}} \right) \quad (2.20)$$

$$\approx \text{Tr}_{2,\dots,N} \left(e^{-\beta h_{N-1,N}} \dots \odot e^{-\beta h_{l+1,l+2}} \odot \text{Tr}_1 \left(e^{-\beta h_{l,l+1}} \odot \dots \odot e^{-\beta h_{1,2}} \right) \right) \quad (2.21)$$

$$\approx \text{Tr}_{3,\dots,N} \left(e^{-\beta h_{N-1,N}} \dots \odot e^{-\beta h_{l+2,l+3}} \odot \text{Tr}_2 \left(e^{-\beta h_{l+1,l+2}} \odot m_{l \rightarrow l+1} \right) \right) \quad (2.22)$$

$$\approx \dots \quad (2.23)$$

$$\approx \text{Tr}_{N-l,\dots,N} \left(e^{-\beta h_{N-1,N}} \odot m_{N-1 \rightarrow N} \right), \quad (2.24)$$

where we have defined $m_{i \rightarrow i+1} = \text{Tr}_{i-l+1} (e^{-\beta h_{i,i+1}} \odot m_{i-1 \rightarrow i})$, which is an operator acting on sites $i - l + 2$ to $i + 1$ (see figure 2.2). These equations require manipulating operators on up to $l + 1$ spins, so their complexity increases exponentially with the window size l .

Like in the classical setting (2.9), these message passing rules can be generalized to arbitrary graphs, enabling the computation of one- and two-body beliefs from which various quantities of interest such as energy can be computed. The method can also be adapted to estimate all correlation functions.

2.6 Numerical Results

We have numerically implemented the QBP algorithm on various graphs for the Ising and Heisenberg model whose Hamiltonians are

$$H_I = \sum_{v \in V} \vec{g} \cdot \vec{\sigma}_v + \sum_{(u,v) \in E} J_{uv} \sigma_u^z \sigma_v^z, \quad \text{and} \quad (2.25)$$

$$H_H = \sum_{(u,v) \in E} \vec{\sigma}_u \cdot \vec{\sigma}_v \quad (2.26)$$

respectively, and $\vec{\sigma} = (\sigma^x, \sigma^y, \sigma^z)$ are the usual Pauli matrices normalized so that $\sigma^2 = \mathbb{1}/4$. On a line, the homogeneous ($J_{uv} = 1$) Ising model has a zero temperature phase transition at the critical transverse field $\vec{g} = (\frac{1}{2}, 0, 0)$. Most of our simulations were performed at this critical value, as it is expected to represent the “hardest case.” Unless otherwise specified, it is henceforth assumed that $\vec{g} = (\frac{1}{2}, 0, 0)$ and $J_{uv} = 1$.

We used QBP to compute the energy density of the Ising model on an infinite line. This model can be solved exactly by means of a Jordan-Wigner transform that maps the interacting spin chain to a collection of free fermions [35]. Figure 2.3 shows the difference between the energy density computed with QBP and its exact value as a function of inverse temperature. Also shown are results obtained from a superoperator version of time-evolving block decimation (TEBD), which combines ideas from [14, 36]. Since a line is a tree, the error in the results obtained from the replica method is entirely caused by the TS decomposition. The results obtained for sliding window QBP are in remarkably good agreement with the exact value, and can be systematically improved by increasing ℓ . This reflects the fact that correlations are short-ranged in finite-temperature one-dimensional models. As expected the agreement improves for noncritical g . Results obtained for the Heisenberg model on the infinite line (not shown) are similar in all aspects.

To characterize the performance of QBP on more general graphs, we restrict our attention to systems with less than 12 spins, allowing comparison to direct brute-force numerical solutions. Figure 2.4 shows the correlation function $C(0, j) = \text{Tr}\{\sigma_0^z \sigma_j^z \rho\}$ for H_I on a frustrated 11-site circle. We assess the quality of the approximation \tilde{C} to the exact correlation




Figure 2.3. Critical Ising model on infinite line. Energy density estimate \tilde{E} using the method of replicas with $N_\tau = 10$, sliding window QBP with $\ell = 6$, and TEBD with $\chi = 150$, compared to the exact energy density E obtained from fermionization.

C by the average relative error,

$$\text{Error} = \frac{\sum_j |C(0, j) - \tilde{C}(0, j)|}{\sum_j |C(0, j)|}. \quad (2.27)$$

Sliding window is again in very good agreement with the exact value for a relatively small window size. For the values of $N_\tau \leq 10$ accessible with modest computational resources, the replica QBP reproduces the exact correlation function within a few percent at sufficiently high temperatures $\beta \lesssim 6$, which is consistent with the systematic error due to the TS decomposition.

Indeed, both the TS decomposition and the loopy QBP contribute to the total error, (2.27). By brute-force computation, it is possible to determine exactly what fraction of the error is caused by each of these approximations, and in almost all cases we have studied at critical g , both contributions were comparable. Figure 2.5a shows each contribution to the total error as a function of the transverse field $\vec{g} = (g, 0, 0)$.

The most successful applications of classical belief propagation algorithm are on graphs whose typical loop size is very large. This is the case for instance of low density parity check codes [21, 23] and spin glasses on Bethe lattices [24]. Intuitively, one expects a local algorithm like belief propagation to be relatively insensitive to the large-scale structure of the graph. We expect QBP to share this feature, and Figure 2.5b illustrates the effect of the loop size on the average relative error of the correlation function. The oscillatory




Figure 2.4. (Color online) Correlations for H_I on an 11-site circle at $\beta = 6$. Exact numerical solution (dash), sliding window with $\ell = 5$ (dash-dot), and the replica method for various values of N_τ (full). Left inset: Error (2.27) vs. the N_τ for different β .

behavior of the error is explained by the frustration present in odd-size circles. Save from these oscillations, the results show a global improvement as the loop size increases. Errors obtained from sliding window (not shown) also show a clear improvement as the loop size increases, but tend to have higher errors on even-size loops.




Figure 2.5. (Color online) Replica method with $N_\tau = 10$. a) Different contributions to the error (2.27) vs. transverse field for H_I on an 11-site circle at $\beta = 6$. b) Error vs. loop size, for $\beta = 1, 2, \dots, 10$.

We have tested QBP on a variety of graphs depicted in figure 2.6 a)-d). The resulting errors in the correlation functions are shown in figure 2.6. The computational cost is slightly




Figure 2.6. (Color online) Error (2.27) for Ising and Heisenberg models on various loopy graphs using the replica method. For the Ising model $N_\tau = 10$ and for the Heisenberg model $N_\tau = 4$. [d) is a torus.]

higher for the Heisenberg model because h_{uv} do not mutually commute. This restricts the computation to lower values of N_τ and consequently yields larger errors. Modulo this difference, the error is most prominent for graphs c) and d) which contain loops of size 3. In those cases, we found that the QBP algorithm was not converging: the magnitude of the errors is consistent with the magnitude of the time fluctuations of $b_{uv}(t)$, cf. (2.15). As expected, the predicted correlation function is in much better agreement with its exact value on graphs a) and b) that have only relatively large loops.

2.7 Conclusion

We have numerically characterized the performance of the recently proposed QBP algorithm. In the high temperature phase, both the replica and the sliding window QBP algorithms perform remarkably well on a tree with modest computational resources, cf. figure 2.3, and offer performances similar to TEBD. On loopy graphs, we found that the algorithm gives reliable approximations when the loop size is large. Most importantly, when the results deviated from the exact value, e.g., in the presence of small loops, the algorithm did not reach a steady state, i.e., the beliefs (2.15) were highly fluctuating as a function of time. This provides an indirect way of assessing the validity of the results.

In [37], a technique similar to what we have called the replica method was used to

investigate the phase diagram of quantum spin-glasses on Cayley trees. Based on the results we have presented, QBP should be suitable to study this phase diagram for more general Bethe lattices whose typical loop size scales as $\log|V|$. In the classical setting, it has been argued that the physics of random Bethe lattices and Cayley trees is greatly different [24]. We note that the randomness in quenched disordered systems should not affect the performances of QBP. In fact, our results obtained for random couplings J_{uv} and random local fields \vec{g} are typically in better agreement than the ones we have presented.

Chapter 3

Coarse-Grained Belief Propagation

3.1 Introduction

In chapter 2, we discussed the quantum belief propagation algorithm, which is a generalization of the belief propagation algorithm to the quantum setting [1, 27, 28, 37]. The gist of the method consists of solving the system exactly on a small cluster and using this solution to compute effective thermal Hamiltonians on the neighboring clusters. This procedure is repeated iteratively until it produces a correction to the bare Hamiltonian that accurately mimics a system of infinite size. Numerical tests indicate that the method is very accurate at high temperatures and on trees or graphs with only large loops [1]. However, BP becomes unreliable at low temperatures because the clusters must be larger or equal to the range of the effective Hamiltonian, which grows like the inverse of the temperature.

In this chapter, we describe how to quantify the errors accumulated at low temperatures. Using this error measure, we extend the applicability of belief propagation to all temperatures by introducing *coarse-grained belief propagation* (CGBP). This algorithm combines the strength of belief propagation from chapter 2 with *entanglement renormalization* (ER) [38, 39], which is a refinement of real space renormalization [40] devised by Vidal. The strength of entanglement renormalization comes from the realization that entanglement is organized on different length-scales in the ground state of many systems of interest, including critical systems [38, 41, 42] and systems with topological order [43, 44]. As a consequence of this organization, entanglement can be efficiently removed from the state by a sequence of local coarse-graining transformations. This leads to a very efficient scheme for finding low-energy states of local Hamiltonians, making ER a very effective method for low temperatures.

Coarse-grained belief propagation combines the features of BP and ER. Starting at high temperatures where plain BP is accurate, temperature is lowered until coarse-graining the lattice by eliminating its shortest length-scale degrees of freedom becomes favorable. Coarse-graining discards some high energy states, which results in a systematic error in the thermal state. On the other hand, it increases the effective size of the clusters, making BP more accurate. The coarse-graining procedure is continued as temperature is lowered to zero where plain ER is accurate. An estimate of the error caused by BP can be used to determine the temperatures at which each coarse-graining procedure should be performed.

In section 3.2, we review the BP method and describe how to estimate its accuracy. Based on this technique, Section 3.3 presents results obtained for the transverse field quantum Ising spin glass studied in [37]. Finally, in section 3.5 we introduce the CGBP algorithm, and benchmark it using an infinite quantum Ising chain.

3.2 Errors in Belief Propagation Results

In this section, and in the rest of the chapter, we will be focusing on the version of the belief propagation algorithm discussed in section 2.5.2 called sliding window belief propagation. To understand the error behavior of the algorithm physically, we will consider a one-dimensional chain with nearest neighbor interactions $H = \sum_i h_{i,i+1}$. The thermal Gibbs state at inverse temperature β is given by $\rho = e^{-\beta H} / \mathcal{Z}(\beta)$. The reduced state of spins $2, 3, \dots, N$ is obtained by tracing out the first spin of the chain, i.e.,

$$\rho_{2,\dots,N} = \text{Tr}_1(\rho). \quad (3.1)$$

We can formally define the *effective thermal Hamiltonian* H_{eff} acting on sites $2, \dots, N$ up to normalization by

$$\rho_{2,\dots,N} \propto e^{-\beta H_{\text{eff}}}. \quad (3.2)$$

In other words, H_{eff} is the traceless part of $-\frac{1}{\beta} \text{Log}(\rho_{2,\dots,N})$ where Log is the principal matrix logarithm. Similar effective Hamiltonians have also been studied in [45] in the context of reduced density matrices of fermionic and bosonic ground states. We also define the *effective*




Figure 3.1. (Color online) Upper bound on $\|V_{\text{eff}}^j\|$ (left) and correlations (right) of the Ising chain with critical transverse field. These exact values are obtained from a Jordan-Wigner transform. Note that $\|V_{\text{eff}}^j\|$ decays much faster than the correlations (notice the different length-scales). The values for $\beta = 10$ are in a different color as an aid to the eye.

thermal potential,

$$V = H_{\text{eff}} - \sum_{i=2}^{N-1} h_{i,i+1}, \quad (3.3)$$

as the term added to the bare Hamiltonian on sites 2 to N due to the presence of site 1.

At high temperatures, the effective thermal potential is short ranged. Consider for instance the one-dimensional Ising model with transverse field on an infinite chain

$$H = \sum_{i=-\infty}^{\infty} \sigma_i^z \sigma_{i+1}^z + B \sigma_i^x. \quad (3.4)$$

At zero temperature, this model exhibits a phase transition at the critical transverse field value $B = 1$. Figure 3.1 shows the value of the effective potential,

$$V_{\text{eff}} = -\frac{1}{\beta} \text{LogTr}_{-\infty, \dots, 0}(e^{-\beta H}) - H_{\text{bare}}, \quad (3.5)$$

obtained from cutting the critical Ising chain in half, i.e., tracing out spins $-\infty$ to 0 from the thermal state of an infinite chain. What is plotted is (an upper bound to) the operator norm of the cumulants of V_{eff} :

$$V_{\text{eff}}^j = \text{Tr}_{j+1, \dots, \infty} (V_{\text{eff}} - \sum_{k=1}^{j-1} V_{\text{eff}}^k), \quad (3.6)$$

with $V_{\text{eff}}^1 = \text{Tr}_{2, \dots, \infty} (V_{\text{eff}})$. We see that V_{eff} has a very short range, in fact much shorter than the correlation length in the system.

When the window size in the belief propagation algorithm is sufficiently large to support the effective thermal potential, the messages contain all the information required to calculate the states of the rest of the system. In this case, we can use the sliding window belief propagation algorithm as discussed in section 2.5.2.

However, the effective thermal potential V_{eff} cannot be computed exactly because of computational limitations. Instead, we can only estimate its value on a cluster of finite size. Thus, the main source of error in our method is due to the truncation of V_{eff} . We can assess the error by evaluating the portion of V_{eff} we discard. On a chain for instance, the error caused on the estimate of the beliefs is

$$\frac{1}{\mathcal{Z}(\beta)} (e^{-\beta H_{\text{eff}}} - e^{-\beta(H_{\text{eff}} - \sum_{j>l} V_{\text{eff}}^j)}) \approx \beta \sum_{j>l} \langle V_{\text{eff}}^j \rangle \quad (3.7)$$

for $\beta V_{\text{eff}}^j \ll 1$. Making the assumption (see figure 3.1) that $\|V_{\text{eff}}^j\|$ decreases exponentially with j , we estimate this quantity by

$$\beta \langle V_{\text{eff}}^{l+1} \rangle \approx \beta \langle V_{\text{eff}}^l \rangle \left(\frac{\langle V_{\text{eff}}^l \rangle}{\langle V_{\text{eff}}^{l-1} \rangle} \right), \quad (3.8)$$

which is our final error estimate. With the same reasoning, similar estimates can be derived for the error associated with different observables.

Figure 3.2 compares this error estimate to the true error produced by BP for the evaluation of the energy density of the critical one-dimensional Ising chain. Clearly, the error estimate accurately bounds the true error. This figure also illustrates the power of BP by comparing the accuracy with which BP estimates the energy density of an infinite chain to what is achieved by brute-force diagonalization with equivalent computational resources.




Figure 3.2. Error on the energy density as a function of temperature for the critical Ising chain. For QBP with window size $l = 10$, we show the error estimate of (3.8) and the true error obtained by comparison with analytical solution. Also shown is the finite size error for exact diagonalization of a 11-site chain and error caused by iTEBD [14] with parameters $\delta T = 0.001$ and $\chi = 150$, both of which require equivalent computational resources.

BP largely outperforms diagonalization down to temperatures of order 0.1. This behavior is expected because, as seen on figure 3.1 (red line), the range of the effective thermal potential becomes larger than the window size ($l = 10$) at this temperature.

3.3 Disordered System Revisited

The use of belief propagation in physics originates in the study of disordered systems—spin glasses—where it is more often referred to as the “cavity method” [24]. Along with the sliding window algorithm outlined in the previous section, we presented in [1, 27] a second distinct way to generalize BP to the quantum setting. This method, which we named “replica BP,” maps the quantum system to a classical system with one additional spatial dimension of length equal to the inverse temperature β . The edges of a graph become ribbons on which classical BP can be employed for sufficiently low β . This technique was independently introduced by Laumann et al. [37] for the study of the transverse field quantum Ising spin glass. The continuous imaginary-time limit of this procedure was later studied in [46].

The numerical results obtained in [1] suggest that, for a given amount of computational power, sliding window BP is much more accurate than replica BP. In this section, we revisit




Figure 3.3. Schematic illustration of the procedure to calculate the message from a node to its parent. First, two messages from its children are merged using the \odot -product. The bare Hamiltonian term relating the current node to its parent is also added. Tracing out the leaves of this 8-spin message gives the message to the node’s parent.

the spin-glass model of [37] using sliding window BP, and apply the method outlined in the previous section to estimate the accuracy of our results.

The spins are located at the vertices of a degree-3 Cayley tree. The Hamiltonian has Ising coupling between neighboring spins and a transverse field B :

$$H = \sum_{\langle ij \rangle} \sigma_i^z \sigma_j^z + B \sum_i \sigma_i^x + \sum_{i \in \text{Boundary}} r_i \sigma_i^z. \quad (3.9)$$

The last term is a random parallel boundary field introduced to create frustration in the system. The strength of the boundary fields r_i are chosen at random uniformly in $[-1, 1]$. The quantity of interest in this setting is the Edwards-Anderson order parameter $q_{\text{EA}} = \langle \sum_j \langle \sigma_j^z \rangle^2 \rangle_Q$ on the lattice at the thermodynamic limit, where $\langle \cdot \rangle_Q$ refers to the quench average over the random boundary field configurations. Note that this order parameter is defined along the axis perpendicular to the applied external field B , but parallel to the random boundary field. Hence, q_{EA} is zero in the paramagnetic phase. It becomes nonzero on the onset of the glassy phase where the system settles into one of many metastable randomly polarized state.

Our numerical simulations are performed on a tree of depth 12. To further reduce the finite-size effects, the EA order parameter is only evaluated on the central spin of the lattice, away from the boundary. The order parameter can be computed from the single-body belief b_j obtained from BP. On a chain, we would start from one end and propagate messages to the other end of the chain. On the tree, we start from the leaves, and propagate messages towards the middle. The basic message passing step is demonstrated in figure 3.3. First, two messages from the children of a given node are combined using the \odot -product. The

bare Hamiltonian term relating the node to its parent is then added. Tracing out the leaves of this 8-spin message gives the message to the node's parent. This procedure is repeated until the central site is reached where three messages are joined to produce the belief. The quench average is obtained by repeating this procedure 100 times with different boundary field configurations.

Because the graph contains no loops, the error estimate presented in the previous section is reliable. However, note that the statistical fluctuations of the quench average are not included in this error estimate. The statistical fluctuations of the average of q_{EA} over the many instances of boundary fields range from 2 percent at low temperatures to 14 percent at high temperatures. Therefore, the main source of error in parts of figure 3.4 is the statistical fluctuations, which are not shown in the plot, and can be systematically reduced by increasing the sample size.

Fig 3.4 shows the EA order parameter q_{EA} in the transverse field-temperature diagram, along with the estimated BP error. The glassy/paramagnetic phase transition line agrees with results of [37] up to the statistical fluctuations above $T \gtrsim 0.3$. At low temperatures ($T \lesssim 0.3$), our results indicate a phase transition line with a decreasing value of B as T is lowered. We could not think of any physical mechanism that could explain this behavior. Moreover, this happens in a region of the phase diagram where the BP error is high. Hence, we suspect that the true phase transition line has a monotonous behavior in temperature and that the glassy phase persists all the way to zero temperature for a transverse field $B \lesssim 1.65$. This conclusion and, more generally, our entire phase diagram is in very good agreement with that of [37].

3.4 Multiscale Entanglement Renormalization

We have seen that BP provides reliable estimates of thermal expectation values as long as the effective thermal potential V_{eff} is sufficiently short ranged to be tracked numerically. As seen, e.g., in figure 3.1, the range of V_{eff} grows linearly with β , so BP becomes unreliable at low temperatures. To probe lower temperatures, one needs to increase the window size l , which is not feasible because resources scale as $O(2^l)$. On the other hand, as we lower the temperature, high energy excitations become increasingly irrelevant. This fact leads to efficient algorithms for zero temperature simulations such as entanglement renormalization




Figure 3.4. The Edwards-Anderson order parameter for transverse field quantum Ising spin glass on a degree-3 Cayley tree with random parallel boundary fields. The total depth of the tree is 12, and the plots show the average of 100 instances of random boundary fields. Values of q_{EA} range from 0 to 1 and contours are equally spaced. The error on the order parameter is estimated using the procedure outlined at the end of Section 3.2.

and DMRG. These algorithms become rapidly inaccurate at finite temperature because they are only able to keep track of a small number of eigenstates. In this section, we will describe a method that interpolates between BP at high T to ER at $T = 0$. Before we do so, we briefly review ER, see [38, 39] for a detailed description.

Entanglement renormalization [38, 39] builds on the multiscale renormalization ansatz (MERA), which asserts that certain degrees of freedom can be decoupled from the ground state of local Hamiltonians by unitary transformations acting on small spatial regions. A concrete example of this scheme is illustrated in figure 3.5 in the case of a one-dimensional lattice. The lattice is first partitioned into clusters each containing 3 consecutive sites. A disentangling transformation u (a unitary transformation on $(\mathbb{C}^d)^{\otimes 2}$) is applied on the boundary of each cluster in order to minimize the correlations between neighboring clusters. Finally, local degrees of freedom are discarded from each cluster by means of an isometry v mapping $(\mathbb{C}^d)^{\otimes 3}$ to \mathbb{C}^d .

This procedure is applied repeatedly. At each iteration, the disentanglers and isometries transform the Hamiltonian H_i to a new Hamiltonian H_{i+1} acting on a smaller lattice and retaining only the lowest eigenstates of the previous Hamiltonian. ER is halted when only a few sites remain in the lattice so it can be handled exactly numerically. The disentanglers




Figure 3.5. Schematics of entanglement renormalization for a ternary MERA of a one-dimensional lattice.

and isometries are chosen in such a way that the final state (or more generally subspace) minimizes the energy of the initial Hamiltonian H_0 . This minimization problem is in general hard, but good heuristics have been devised for it [38, 39].

It is convenient to use the diagrammatic formalism of tensor networks to describe the unitaries and the local isometries constituting a MERA (the latter correspond to the elimination of local degrees of freedom, but the division into “disentangling” unitaries and isometries is often arbitrary). In this formalism, the identity on \mathbb{C}^d is represented by a single directed edge, and operators $\mathbf{O} : (\mathbb{C}^d)^{\otimes n} \rightarrow (\mathbb{C}^d)^{\otimes m}$ are represented by shaded boxes with n ingoing and m outgoing (ordered) edges. Labels on the edges from $1, \dots, d$ correspond to the elements of a fixed orthonormal basis of \mathbb{C}^d . Products of operators are taken by connecting outgoing with ingoing edges, and summing over labels of edges with no free ends. The trace of an operator on $(\mathbb{C}^d)^{\otimes n}$ is taken by connecting each outgoing strand with the corresponding ingoing strand, and then contracting the tensor network (i.e., summing over all edge labelings). Partial traces are computed analogously by connecting up subsets of edges. Tensor products are obtained by placing diagrams next to each other.

With these conventions, the property $\mathbf{W}^\dagger \mathbf{W} = \mathbb{I}_{(\mathbb{C}^d)^{\otimes n}}$ of an isometry $\mathbf{W} : (\mathbb{C}^d)^{\otimes n} \rightarrow (\mathbb{C}^d)^{\otimes m}$ (for $m \geq n$) takes the simple form,

$$\begin{array}{c}
 \begin{array}{c}
 \uparrow \uparrow \uparrow \uparrow \\
 \text{W}^\dagger \\
 \uparrow \uparrow \uparrow \uparrow \\
 \text{W} \\
 \uparrow \uparrow \uparrow \uparrow
 \end{array}
 \begin{array}{l}
 (n \text{ strands}) \\
 (m \text{ strands}) \\
 (n \text{ strands})
 \end{array}
 =
 \begin{array}{c}
 | | | | \\
 | | | | \\
 | | | | \\
 | | | |
 \end{array}
 \cdot
 \end{array}
 \quad (3.10)$$

Identity in (3.10) is crucial for the definition of the MERA ansatz.

Consider a system of N qudits arranged on a line with periodic boundary conditions. For




Figure 3.6. Two examples of MERA structures: For a periodic qudit chain (at the bottom of the figure), a variational family of states is obtained by varying over the contents of the boxes. In (a), these are (adjoints of) isometries $\mathbf{W} : (\mathbb{C}^d)^{\otimes 2} \rightarrow \mathbb{C}^d$, unitaries $\mathbf{U} : (\mathbb{C}^d)^{\otimes 2} \rightarrow (\mathbb{C}^d)^{\otimes 2}$ and (the adjoint of) a state $|\varphi\rangle \in \mathbb{C}^d$ for the strip at the top. Tree-like structures may also be considered, and lead to particularly simple expressions in the homogenous case [47]. Note that we choose to represent the top tensor by a (periodic) strip; the reason for this will become obvious once we move to the anyonic setting.

any tensor network as shown in figure 3.6, one obtains a variational family of states $|\Psi\rangle = |\Psi_{|\varphi\rangle, \{\mathbf{W}\}}\rangle \in (\mathbb{C}^d)^{\otimes N}$ parametrized by the isometries $\{\mathbf{W}\}$ and the state $|\varphi\rangle \in \mathbb{C}^d$ at the top of the structure. This is the MERA ansatz. Figure 3.6 represents a map $(\mathbb{C}^d)^{\otimes N} \rightarrow \mathbb{C}$, which can be understood as the bra $\langle\Psi|$ of the represented state. The fact that the recipe specifies $\langle\Psi|$ instead of $|\Psi\rangle$ is a matter of preference, but the chosen convention has a natural operational interpretation: One may think of figure 3.6 as a renormalization prescription by decomposing the tensor network along different horizontal cuts L_i . The strips $[L_i, L_{i+1}[$ represent coarse-graining maps, which reduce the number of degrees of freedom at each level.

Importantly, local expectation values and pair correlation functions can be computed efficiently for such a state. This is because the expectation value of an observable \mathbf{O} is given by the contraction of the tensor network obtained by sandwiching \mathbf{O} between the MERA network and its adjoint as in figure 3.7 (b). (3.10) then allows to simplify the network resulting in a significantly smaller network corresponding to the “causal cone” of the operator, see figure 3.7 (c). This network can be efficiently contracted, as the number of tensors scales logarithmically with the number of sites N . For example, for the case of




Figure 3.7. Tensor networks associated with the MERA of figure 3.6 (a). The network (a) represents the projection $|\Psi\rangle\langle\Psi|$ onto the state described by the MERA. The contraction of the network (b) gives the expectation value $\langle\Psi|\mathbf{O}|\Psi\rangle$ of a local operator \mathbf{O} . The tensor network (c) gives the same value as (b), and is obtained from it by using (3.10) repeatedly.

figure 3.7 (b), we get

$$\langle \Psi | \mathbf{O} | \Psi \rangle = \langle \varphi | \mathcal{E}_{L_{\max}} \circ \mathcal{F}_{L_{\max}} \cdots \mathcal{E}_{L_0} \circ \mathcal{F}_{L_0}(\mathbf{O}) | \varphi \rangle, \quad (3.11)$$

where $L_{\max} = \log_2 N - 1$, and the superoperators $\{\mathcal{E}_L, \mathcal{F}_L\}$ associated to a level L are

$$\mathcal{E}_L \left(\begin{array}{c} \uparrow \uparrow \\ \square \mathbf{O} \\ \uparrow \uparrow \end{array} \right) = \begin{array}{c} \uparrow \uparrow \\ \square \mathbf{W}^\dagger \\ \uparrow \uparrow \\ \square \mathbf{O} \\ \uparrow \uparrow \\ \square \mathbf{W} \\ \uparrow \uparrow \end{array}, \quad \mathcal{F}_L \left(\begin{array}{c} \uparrow \uparrow \\ \square \mathbf{O} \\ \uparrow \uparrow \end{array} \right) = \begin{array}{c} \uparrow \uparrow \\ \square \mathbf{U}^\dagger \\ \uparrow \uparrow \\ \square \mathbf{O} \\ \uparrow \uparrow \\ \square \mathbf{U} \\ \uparrow \uparrow \end{array} \quad (3.12)$$

for $L < L_{\max}$, and

$$\mathcal{E}_{L_{\max}} \left(\begin{array}{c} \uparrow \uparrow \\ \square \mathbf{O} \\ \uparrow \uparrow \end{array} \right) = \begin{array}{c} \uparrow \uparrow \\ \square \mathbf{W}^\dagger \\ \uparrow \uparrow \\ \square \mathbf{O} \\ \uparrow \uparrow \\ \square \mathbf{W} \\ \uparrow \uparrow \end{array}, \quad \mathcal{F}_{L_{\max}} \left(\begin{array}{c} \uparrow \uparrow \\ \square \mathbf{O} \\ \uparrow \uparrow \end{array} \right) = \begin{array}{c} \uparrow \uparrow \\ \square \mathbf{U}^\dagger \\ \uparrow \uparrow \\ \square \mathbf{O} \\ \uparrow \uparrow \\ \square \mathbf{U} \\ \uparrow \uparrow \end{array}. \quad (3.13)$$

In particular, the reduced two-site density operator of the state $|\Psi\rangle$ can be computed using the adjoint superoperators as

$$\mathrm{tr}_{n-2} |\Psi\rangle\langle\Psi| = \mathcal{F}_1^\dagger \circ \mathcal{E}_1^\dagger \cdots \mathcal{F}_{L_{\max}}^\dagger \circ \mathcal{E}_{L_{\max}}^\dagger (|\varphi\rangle\langle\varphi|).$$

Similar expressions can be obtained for translates as well as for pair correlation functions evaluated at certain specific distances [47–49].

Given a MERA-structure as in figure 3.6, we have described a recipe giving a variational family of states for a chain of N d -dimensional qudits parametrized by isometries of the form $\mathbf{W} : (\mathbb{C}^d)^{\otimes n} \rightarrow (\mathbb{C}^d)^{\otimes m}$ (the constants $n \geq m$ depend on the MERA-structure). There is a natural way of enlarging this family: the edges (in higher levels of the network) may be interpreted as corresponding to \mathbb{C}^χ , where the refinement parameter $\chi > d$ is larger than the dimension d of the physical qudits. Increasing this so-called bond dimension χ amounts to using isometries $\mathbf{W} : (\mathbb{C}^\chi)^{\otimes n} \rightarrow (\mathbb{C}^\chi)^{\otimes m}$ (and correspondingly a state $|\varphi\rangle \in \mathbb{C}^\chi$ at the top). To motivate a similar refinement in the anyonic setting, we point out that if $\chi = d^s$ is an integer power of the physical qudit dimension d , then this enlargement of the variational

family is equivalent to replacing each edge in the MERA-structure by s edges. Note also that, since the number of isometries is of order $O(N\log N)$ and each isometry is described by fewer than χ^{n+m} parameters, the total number of parameters describing a MERA is of order $O(\text{poly}(\chi)N\log N)$.

We conclude this short exposition by mentioning two subclasses of MERA-ansatz states, which are of particular interest: a *translation-invariant* MERA has identical isometries within every level. As a consequence, the number of parameters of such a MERA scales as $O(\text{poly}(\chi)\log N)$. A *scale-invariant* MERA is one where all isometries (with identical domain and range) are chosen to be the same: Here the number of parameters is only $O(\text{poly}(\chi))$ independent of N (see [41] for a more detailed discussion of the complexity of MERA and [39] for concrete examples). For a scale-invariant MERA, it is possible to write down two-point and three-point correlators explicitly in terms of eigenvalues of a certain coarse-graining superoperator defined in terms of the isometries as in (3.12). This allows to numerically extract critical exponents and parameters from the associated CFT, as described, e.g., in [47–49].

3.5 Coarse-Grained Belief Propagation

Both ER and BP revolve around the idea that some correlations are short range in the state of interest. BP becomes exact when the conditional mutual information $I([-\infty, j] : [j+l, \infty] | [j+1, j+l-1])$ vanishes for a sufficiently large window size l . In other words, BP can work in the presence of arbitrary long range classical correlations but the purely quantum correlations must be short ranged. These quantum correlations tend to increase like the inverse temperature β (cf. figure 3.1), which limits BP to relatively high temperatures, unless the window size l can somehow be increased while keeping computational cost low. In contrast, ER makes use of disentanglers to eliminate short-range quantum correlations in the system and coarse-grain the lattice. Because it only keeps a few low-energy states, it is limited to very low temperatures.

Coarse-grained belief propagation interpolates between these two methods and provides accurate thermal expectation values over a very large range of temperatures. At high temperature, CGBP reduces to ordinary BP. As the temperature is lowered, the error attributed to BP increases. At some temperature T_1 , it becomes favorable to coarse-grain




Figure 3.8. (Color online) Energy vs. temperature for the one-dimensional critical quantum Ising model. Presented are the plain BP results for $l = 10$ (dark blue), CGBP results for various levels of coarse-graining with $\chi = 4$ and $l = 5$ (various colors), the plain ER result $\chi = 4$ (light green) along with the exact result obtained from Jordan-Wigner transformations (black). The labels on the temperature axis correspond to the switching temperatures between various levels of coarse-graining in the CGBP algorithm. E.g., plain BP is very accurate down to $T_1 \approx 0.22$, where BP combined with one level of coarse-graining (red) becomes more accurate. The inset shows the absolute error in the combined CGBP result with respect to the exact solution (blue), and the error estimate calculated using the procedure given in section 3.5. While plain BP and plain ER are both very inaccurate for temperature range $0.005 \lesssim T \lesssim 0.22$, combining the two using CGBP yields very accurate results.

the lattice using one step of ER. This procedure discards some high energy states, leading to a systematic error, but it effectively increases the BP window length l by a constant factor (3 in the ternary ER scheme illustrated in figure 3.5). This increase in the window length improves the accuracy of BP and compensates for the loss of high energy states. As the temperature is lowered, the lattice is further coarse-grained at T_2, T_3 , etc. until the CGBP reduces to ordinary ER. Figure 3.8 illustrates this behavior for the one-dimensional critical quantum Ising model. Each coarse-graining level provides a reliable estimate only for a small temperature range, but the union of these ranges cover the entire temperature domain. Thus, CGBP provides accurate estimates of thermodynamical observables in temperature ranges that are accessible to neither BP nor ER.

To calculate the optimal coarse-graining temperatures T_i , we would need independent error assessments for ER and BP. Then, we could switch to a coarser lattice whenever the

increase in window size compensates for the coarse-graining of the Hamiltonian. However, we are not aware of a reliable method to estimate the error caused by ER. Instead of using error estimates, we simultaneously perform BP on two different coarse-grained levels and determine the switching temperatures by comparing the results.

More precisely, let $\bar{x}_i(T)$ be the expectation value of some observable X at temperature T obtained by BP on the i th level of coarse-graining. To determine the switching temperature T_{i+1} from the i th level of coarse-graining to the $(i+1)$ th, we calculate both \bar{x}_i and \bar{x}_{i+1} as we slowly lower the temperature. At high temperatures, the dominating error is attributed to the discarded high energy states, so \bar{x}_i is more accurate than \bar{x}_{i+1} . On the other hand, at low temperatures, as the range of the effective thermal potential gets larger than the BP window, \bar{x}_{i+1} becomes more accurate than \bar{x}_i . The two sources of errors are balanced when $|\bar{x}_i - \bar{x}_{i+1}|$ reaches a minimum (see figure 3.9), so T_i should be chosen at the position of this minimum. There can be exceptions to this rule that result from accidental crossings of \bar{x}_{i+1} and \bar{x}_i . In that case, the error estimate for BP presented in section 3.2 can be used to discriminate between the multiple minima. Indeed, the switching should occur when the value of $|\bar{x}_i - \bar{x}_{i+1}|$ is the closest to the error attributed to BP on the i th level because both numbers are estimates of the BP error. The switching temperatures for figure 3.8 were chosen following this method.

We can use the same reasoning to estimate the total error $\delta x(t)$ on our final estimate $\bar{x}(T)$ obtained by joining the $\bar{x}_i(T)$ over their respective range. We define $\delta x_i^{\text{BP}}(T)$ to be the error attributed to BP on the i th level of coarse-graining. This quantity can be estimated as described in section 3.2. For $T > T_1$, we have $\delta x(T) = \delta x_0^{\text{BP}}(T)$ since BP is the only source of error. Between T_1 and T_2 , there are two contributions to the error: the error $\delta x_1^{\text{BP}}(T)$ attributed to BP on the 1st coarse-grained level and the error attributed to ER caused by discarding high energy states. This second error decreases as temperature is lowered, and at $T = T_1$ it is equal to the BP error (this is how T_1 was defined). Thus, we obtain for $T_2 \leq T < T_1$ the bound $\delta x(T) \leq \delta x_0^{\text{BP}}(T_1) + \delta x_1^{\text{BP}}(T)$. More generally, for $T_{i+1} \leq T < T_i$ we get $\delta x(T) \leq \sum_{j \leq i} \delta x_{j-1}^{\text{BP}}(T_j) + \delta x_i^{\text{BP}}(T)$. See the inset of figure 3.8 for the error estimate of CGBP with $\chi = 4$ and $l = 5$.

In addition to CGBP with $\chi = 4$ and $l = 5$, we have also investigated various other combinations of χ and l . At equal computational costs—which are of $O(\chi^{3l})$ —the results are qualitatively equivalent in the sense that they have roughly equivalent maximum error.




Figure 3.9. The exact error in the energies calculated at the first and second levels of coarse-graining, $|E_1 - E_{exact}|$ and $|E_2 - E_{exact}|$ are plotted along with $|E_1 - E_2|$. Note that for both $|E_1 - E_{exact}|$ and $|E_2 - E_{exact}|$, MERA error dominates at high temperatures, and BP error dominates at low temperatures. The minimum of their difference $|E_1 - E_2|$ occurs when the BP error at level 1 of coarse-graining is equal to the MERA error at level 2 of coarse-graining. Furthermore, the high temperature portion of $|E_1 - E_2|$ is dominated by the MERA error at level 2 and the low temperature portion of $|E_1 - E_2|$ is dominated by the BP error at level 1.

However, varying χ and l at fixed computational cost can improve the results for a given temperature. For instance, at very low temperatures, the computations with a higher χ yield results with better accuracy. On the other hand, larger l and smaller χ perform better at high temperatures. Thus, the values of χ and l could also be varied dynamically in the simulation, but we leave out this possibility for the moment.

The results we obtain with CGBP compare favorably with results obtained by other methods using equivalent computational resources. This can be seen by comparing the inset of figure 3.8 to the various curves shown on figure 3.2. The CGBP result with $\chi = 4$, $l = 5$ has at least four digits of accuracy for the entire temperature range. The complexity of this simulation is equivalent to exact diagonalization of an 11-site chain, that produces a result accurate to roughly three digits. Many methods rely on a Trotter-Suzuki decomposition by discretizing imaginary time in intervals δT , which creates a systematic bias $O(\delta T^3)$ in the result. This is the case for instance of imaginary time dependent block decimation [14]. This algorithm has complexity $(\chi d^2)^3 T_{\min}/\delta T$ where $d = 2$ for the Ising model is the dimension of the spins. The parameters $\chi = 150$ and $\delta T = 0.001$ yield the

same complexity as the $\chi = 4$ $l = 5$ CGBP algorithm, and produce an error of roughly 10^{-2} at sufficiently low temperatures.

3.6 Discussion

The concept of an effective thermal potential, obtained by tracing out sites from the thermal Gibbs state, gives a clear physical picture of the workings of belief propagation. Adding a site to a spin chain can modify the thermal state even far away from the added site, as far as the correlation length of the system. However, this effect can be mimicked by adding a short range thermal potential to the Hamiltonian of the original spin chain. At finite temperature, we have seen that the range of this thermal potential can be much shorter than the correlation length of the system, making it more suitable for numerical simulations. This underlies the success of belief propagation and provides a means to assess its accuracy. We have illustrated these concepts and methods on the critical quantum Ising chain and the transverse field quantum Ising spin glass. For this last system, our findings are, within estimated error bars, in agreement with those of [37] obtained from a different belief propagation implementation.

At lower temperatures however, the range of the effective thermal potential becomes too large to handle numerically. For these temperature ranges, we have introduced the coarse-grained belief propagation algorithm by combining belief propagation with entanglement renormalization. Coarse-graining discards some high-energy states, which leads to a systematic error in the thermal states. On the other hand, it increases the accuracy of BP by shortening the range of the effective thermal potential by a constant factor. CGBP seeks an optimal compromise between these two effects in order to accurately probe temperature regimes where neither ER nor BP are reliable. Thus, CGBP truly extends the domain of applicability of the two underlying approaches. Moreover, results obtained by CGBP compare favorably to other methods using equivalent computational resources.

The drawback of CGBP is that it inherits some intrinsic limitations of the underlying approaches. For instance, ER is applicable only when there exists a coarse-graining method, which preserves the locality of the Hamiltonian. For graphs with exponential spreading structure, such as the Cayley tree studied in section 3.3, we are not aware of such coarse-graining procedures. This is the reason we have not implemented CGBP on that system.

Another limitation comes from the shortcoming of BP on graphs containing many small loops such as two-dimensional lattices. Classically, this limitation can be alleviated using generalized BP [15, 16]. Versions of generalized BP that incorporates the smallest loops into messages can be used to reduce errors for loopy quantum many-body systems. However, only very high temperatures can be probed with such methods, and it is difficult to analyze the errors in their results.

Chapter 4

Anyonic Entanglement Renormalization

4.1 Introduction

Strongly correlated quantum many-body systems with topological order have been proposed as a substrate for building fault-tolerant quantum computers [50–53]. Under this proposal, logical information is stored within a subspace of a fixed number of quasi-particles (anyons). Computation is performed by exploiting the nonabelian statistics obeyed by these anyons under braiding. Compared to more conventional implementations of quantum computers, this offers an intrinsic resilience to noise: local perturbations cannot decohere the stored information because of the nonlocal nature of the encoding.

At present, perhaps the most promising candidate systems exhibiting nonabelian statistics are fractional quantum Hall systems [54–56]. Some amount of experimental evidence is already available in this setting [57–59]. Other proposed systems include topological insulators [60–63] and lattice spin systems [50, 64, 65] whose local interactions could be engineered artificially [66, 67]. Independently of the proposed physical realization, the stability of the topologically ordered phase with respect to (local) perturbations is of great interest for topological quantum computation. This presents a formidable theoretical challenge. Possible approaches range from the study of concrete physical models (see, e.g., [68–71]) or bounds on the gap for general families of models [72–74] to the investigation of effective Hamiltonians describing the interanyon interactions (see e.g., [75]). In the context of the latter, it is natural to study paradigmatic models such as one-dimensional anyonic chains. These are the natural counterpart of spin chains. The study of such systems has led, e.g.,

to an exactly solved model called the golden chain [76], and novel realizations of infinite randomness critical phases [77–79].

A peculiar aspect of anyonic systems is the structure of the Hilbert space of N anyons. In contrast to the space of N qudits, this space does not decompose into an N -fold tensor product. Instead, its dimension scales as d^N , where the quantum dimension d is generally not an integer. For analytical studies, it is sometimes convenient to embed this space into a larger tensor product space. However, this approach is inconvenient when using variational methods: with the most straightforward encoding, it may be unclear how to vary (numerically) over the subset of physical states. Furthermore, this leads to a significant increase in computational complexity, particularly because locality is only approximately preserved by such an embedding. Even for one-dimensional anyonic systems, these issues complicate the direct application numerical methods such as DMRG [5]. While such methods may presently be the most powerful and successful tools available, these difficulties motivate the search for alternative approaches.

There is another more profound reason why the naïve application of numerical methods for qudits may be suboptimal: local anyonic operators preserve the total topological charge in their support. Since a realistic Hamiltonian consists of locally acting terms, topological charge conservation severely constrains the action of the Hamiltonian both at the local level as well as on larger scales. Incorporating this fact into the method of study should therefore be highly beneficial. With a qudit embedding, the meaning of this constraint is obscured and may be hard to make use of.

In this chapter, we introduce a variational method for anyonic systems, which avoids using unphysical additional states and optimally exploits the special structure of the Hilbert space.¹ This is motivated in part by the goal of facilitating numerical studies. Perhaps more importantly, variational families of ansatz states can provide significant insight into the nature of quantum correlations in a given system. Even in this regard, a qudit embedding of anyons is generally undesirable, as a physically motivated ansatz for qudit systems may lose its significance when applied directly to the anyonic setting.

Our scheme is inspired by renormalization group studies of anyonic chains [77–79], composite anyon distillation [81] and the multiscale entanglement renormalization ansatz

¹Shortly after posting this work to the arXiv, Pfeifer et al. have also presented work on the anyonic MERA [80]. They use the scale-invariant MERA to numerically extract CFT data in the thermodynamic limit, studying the Fibonacci golden chain at the AFM point.

(MERA) [38, 41] for (nonanyonic) spin chains, which was discussed in section 3.4. In fact, it can be understood as the natural anyonic counterpart of the latter and shares many of its properties. In particular, it can be seen as a renormalization group scheme and is thus especially suited for describing scale-invariant systems. The scheme allows—in principle—to extract, e.g., critical exponents when considering fixed points of the renormalization group flow. More generally, it provides procedures for computing parameters of the corresponding conformal field theory (CFT) in the continuum limit. While such procedures are also available for other variational methods (e.g., matrix product states [9, 82–85], by transfer matrix methods), they are particularly natural in the present context due to the scale-invariant form of the ansatz.

We formulate our ansatz for one-dimensional (periodic) chains of anyons. This allows us to numerically test its validity for the golden chain [76] and its relatives [86]. Ultimately, though, it is desirable to find methods for two-dimensional systems. Our work makes some progress toward this goal: the nonanyonic MERA extends naturally to two dimensions, and this is also the case for its anyonic counterpart. We briefly comment on such generalizations.

Entanglement renormalization, while originally defined for qudit systems, has been extended to free bosons [87] and interacting fermions [88]. Anyonic statistics are peculiar to two dimensions, encompassing both fermions and bosons (which are abelian) in addition to interesting nonabelian generalizations. Our ansatz applies to all such models, but its use is especially suggestive in the nonabelian case. This is due to the special form of the Hilbert space mentioned earlier. Interestingly, the exact details of the exchange statistics play essentially no role in the definition of the ansatz, though they become important when evaluating expectation values of local operators in the case of, e.g., two-dimensional arrangements of anyons.

In a wider context, entanglement renormalization is a special instance of the class of tensor network states, which also includes, e.g., MPS [9, 82–85], PEPS [89] as well as the closely related TERG-states [90]. These have wide applicability beyond qudit systems. In particular, a general framework for fermions has been developed [91–95]. Our focus is on entanglement renormalization because of its unique operational interpretation, as well as the possibility of efficiently computing expectation values without approximations, even in two dimensions (in contrast to, e.g., PEPS [96]). Furthermore, anyonic entanglement renormalization is conceptually related to previous analytical studies for anyonic chains [77–

79]. However, our work suggests that anyonic generalizations of other tensor network states should also be possible along similar lines, though further work may be needed to evaluate their descriptive power.

The structure of this chapter is as follows. In section 4.2, we give some background on anyonic systems and their description in terms of fusion diagrams. We then present the anyonic entanglement renormalization ansatz in section 4.3 and show how to efficiently evaluate expectation values of local observables and correlation functions. We also discuss an operational interpretation in the context of composite anyon distillation. In section 4.4, we apply the ansatz to the golden chain and identify a renormalization group fixed point. We conclude in section 4.6.

4.2 Anyonic States and Operators

In this section, we provide a short introduction to anyons, emphasizing the aspects relevant to the definition of the anyonic entanglement renormalization ansatz: the anyonic Hilbert space and the isotopy-invariant formalism for describing states and operators. We also discuss the origin of the dynamics, that is, the definition of Hamiltonians for anyons. For a thorough and accessible introduction to anyons and topological quantum computation, we recommend [51] (see also [53] for a recent review).

4.2.1 A Unified Treatment of Topological Order

Anyons arise as localized quasi-particle excitations in what can roughly be referred to as two-dimensional topologically ordered quantum media, e.g., qudit lattice systems with certain Hamiltonians [50, 64, 65], quantum Hall systems [54–56] or topological insulators [60–63]. Independently of their physical realization, their state space and exchange statistics are described by the axioms of a topological quantum field theory (TQFT) (see, e.g., [97]). This formalism is extraordinarily useful for studying low-energy processes [75] as well as for the application to quantum computation [51], as it abstracts out the relevant physics: it specifies the particle content, i.e., what particle types occur, describes what their internal degrees of freedom are, and how they are affected by braiding (exchanges) and fusion (which corresponds to bringing particles together). The relation between states of (and operations on) the physical system and the abstract anyonic state space is discussed extensively in the

literature; see, e.g., [50, 51, 55, 64, 65, 98].

The algebraic object underlying such an anyonic theory is a modular category. Roughly, this consists of (i) a finite set of particle types Ω equipped with an involution $*$: $\Omega \rightarrow \Omega$ and containing a distinguished trivial particle $1 \in \Omega$, (ii) fusion rules, i.e., a set of allowed triples of particles, (iii) a quantum dimension $d_a > 0$ associated to every particle a , (iv) a tensor F indexed by 6 particles, and (v) a 3-index tensor R . These are required to satisfy a number of consistency conditions (see, e.g., [51]) the most important of which express associativity of fusion and compatibility of fusion with braiding.

In more physical terms, the involution associates an antiparticle a^* to every particle a , with $1^* = 1$ corresponding to the absence of a particle. The fusion rule summarizes the possible outcomes when bringing two particles together. The quantum dimensions give a rough measure of the growth of the anyonic state space when adding particles, and the F -tensor relates different bases of this space (as explained below). Finally, the R -tensor encodes braiding of pairs of particles. This will be mostly irrelevant for our discussion, but becomes relevant for two-dimensional arrangements of anyons as explained in section 4.5.

4.2.2 The Anyonic Hilbert Space and Anyon Diagrams

We proceed by explaining the construction of the anyonic Hilbert space using the data (i)–(iv) of a modular category. We will mostly follow the detailed exposition in [99], but will require slightly more general definitions when dealing with operators (see, e.g., appendix of [64] for more details).

The state space of a set of anyons depends on their types and on the surface the quantum medium is embedded in. We will discuss two cases in detail below: anyons pinned to fixed locations on a disc and on a torus. Usually, we assume that these are arranged on a chain, though one may also consider, e.g., regular two-dimensional lattices; their particular geometric arrangement only becomes important when considering Hamiltonians, but does not affect the definition of the Hilbert space.

Starting point are certain trivalent graphs with directed edges. These correspond to pants decomposition of the surface with punctures inserted at the anyons' positions. Fixing such a graph, a basis of the anyonic Hilbert space is given by all labelings of the edges with particle labels from Ω satisfying the fusion rules at every vertex. We will give explicit examples for the punctured sphere and the torus below.

Labeled graphs related by reversing the direction of an edge and simultaneously replacing its particle label by the associated antiparticle represent the same vectors. This is somewhat analogous to the formalism of Feynman diagrams. Indeed, anyon diagrams may, to some extent, be interpreted as particle world lines, though this analogy has its limitations. Note also that in many anyon theories of interest, such as the Fibonacci category considered below, every particle is its own antiparticle, $a^* = a$, and it is sufficient to work with undirected graphs.

Dividing up trivalent graphs into neighborhoods of their vertices, one arrives at the following alternative description: the total anyonic Hilbert space is the direct sum of tensor products of two-anyon fusion spaces V_{ab}^c (respectively their dual splitting space V_c^{ab}) corresponding to every vertex, where $a, b, c \in \Omega$, and where the sum is taken over all fusion-consistent labelings (see below). The space V_c^{ab} can be thought of as the internal degrees of freedom of two anyons of type a and b whose combined topological charge is c . Equivalently, it is the space of two anyons a and b on a disc with total topological charge c at the boundary. The latter is—in principle—a measurable quantity. We assume for simplicity that fusion is multiplicity-free, i.e., $\dim V_{ab}^c \in \{0, 1\}$, but our techniques directly generalize to models with fusion multiplicities (see [99] for the necessary adaptations in the diagrammatic calculus).

We pick a normalized vector $|ab; c\rangle \in V_c^{ab}$ in each splitting space and represent it using the isotopy-invariant formalism [99] as

$$|ab; c\rangle = \left(\frac{d_c}{d_a d_b}\right)^{1/4} \begin{array}{c} a \quad b \\ \diagdown \quad \diagup \\ \quad c \\ \uparrow \end{array} \in V_c^{ab},$$

$$\langle ab; c| = \left(\frac{d_c}{d_a d_b}\right)^{1/4} \begin{array}{c} \quad c \\ \uparrow \\ a \quad b \\ \diagup \quad \diagdown \end{array} \in V_{ab}^c,$$

where a scalar d_a called quantum dimension is associated to every label a . More generally, going from a vector to its dual vector corresponds to flipping the diagram along the horizontal axis and reversing all the arrows. Isotopy invariance means that diagrams may be continuously deformed as long as endpoints are held fixed and edges are not passed through each other or around open endpoints.

4.2.2.1 Anyons on a disc

Consider the space $V_c^{\vec{a}}$ of n anyons of types $\vec{a} \in \Omega^n$ pinned to fixed locations on a disc with total charge $c \in \Omega$ at the boundary. For this space, we will often use the standard basis given by the vectors

$$|\vec{a}, \vec{b}; c\rangle_n = \left(\frac{d_c}{\prod_i d_{a_i}} \right)^{1/4} \begin{array}{c} a_1 \quad a_2 \quad a_3 \quad \dots \quad a_n \\ \swarrow \quad \searrow \quad \swarrow \quad \searrow \quad \swarrow \quad \searrow \\ b_1 \quad \dots \quad b_{n-2} \\ \swarrow \quad \searrow \\ c \end{array} , \quad (4.1)$$

where $\vec{b} \in \Omega^{n-1}$ is such that this diagram is fusion consistent. This corresponds to the decomposition

$$V_c^{\vec{a}} \cong \bigoplus_{\vec{b} \in \Omega^{n-1}} V_{b_1}^{a_1 a_2} \otimes V_{b_2}^{b_1 a_3} \otimes \dots \otimes V_{b_{n-2}}^{b_{n-3} a_{n-1}} \otimes V_c^{b_{n-2} a_n} .$$

One may switch between different bases using the (unitary) F move, i.e., the isomorphism between the two decompositions of

$$V_d^{abc} \cong \bigoplus_e V_e^{ab} \otimes V_d^{ec} \cong \bigoplus_f V_f^{bc} \otimes V_d^{af} ,$$

which is specified by the coefficients $[F_d^{abc}]_{e,f}$ in

$$\begin{array}{c} a \quad b \quad c \\ \swarrow \quad \searrow \quad \swarrow \\ e \quad \searrow \\ \uparrow d \end{array} = \sum_f [F_d^{abc}]_{e,f} \begin{array}{c} a \quad b \quad c \\ \swarrow \quad \searrow \quad \swarrow \\ \quad \searrow \quad \swarrow \\ f \\ \uparrow d \end{array} . \quad (4.2)$$

The matrices $[F_d^{abc}]$ collectively constitute the F -tensor of the modular category mentioned earlier. Matrix elements of related isomorphisms such as

$$V_{cd}^{ab} \cong \bigoplus_e V_{ce}^a \otimes V_d^{eb} \cong \bigoplus_f V_{cd}^f \otimes V_f^{ab}$$

the total charge, i.e., if it has block-diagonal form

$$U_{\vec{a}}^{\vec{a}'} \cong \bigoplus_c \left(U_{\vec{a}}^{\vec{a}'}(c) : V_c^{\vec{a}} \rightarrow V_c^{\vec{a}'} \right) . \quad (4.7)$$

Such operators act locally on subsets of n anyons of specified types $\vec{a} \in \Omega^n$ and map them to a state of anyons of specified types $\vec{a}' \in \Omega^m$. In fact, (4.6) (respectively (4.7)) represent the most general locally acting operator, their single most important property being charge preservation. Additional properties such as unitarity impose further conditions on the form of the map/matrix $U_{\vec{a}}^{\vec{a}'}(c)$ for each c .

More generally, we are interested in operators that act between spaces of the form $\bigoplus_c \left(\bigoplus_{\vec{a} \in \Gamma_n} V_c^{\vec{a}} \right)$, where $\Gamma_n \subset \Omega^n$ specifies a set of n -tuples of available (spatial) particle configurations on a chain. A total charge-preserving operator of this kind takes the form

$$U_{\Gamma_m}^{\Gamma_n} = \bigoplus_c U_{\Gamma_m}^{\Gamma_n}(c) , \quad (4.8)$$

where $U_{\Gamma_m}^{\Gamma_n}(c) \in \text{End} \left(\bigoplus_{\vec{a} \in \Gamma_m} V_c^{\vec{a}} , \bigoplus_{\vec{a}' \in \Gamma_n} V_c^{\vec{a}'} \right)$. This can be understood as the projection of a general charge-preserving operator onto inputs and outputs from Γ_m and Γ_n , respectively.

More precisely, let

$$\mathbb{I}_{\Gamma_n} : \bigoplus_c \left(\bigoplus_{\vec{a} \in \Omega^n} V_c^{\vec{a}} \right) \rightarrow \bigoplus_c \left(\bigoplus_{\vec{a} \in \Omega^n} V_c^{\vec{a}} \right)$$

denote the projection onto the subset of states with anyon labels from Γ_n defined by

$$\begin{aligned} \mathbb{I}_{\Gamma_n} &:= \sum_{\vec{a} \in \Gamma_n} \sum_{\substack{c \in \Omega \\ \vec{b} \in \Omega^{n-1}}} |\vec{a}, \vec{b}; c\rangle_{nn} \langle \vec{a}, \vec{b}; c| \\ &= \sum_{\vec{a} \in \Gamma_n} \begin{array}{c} | \\ \uparrow \\ a_1 \end{array} \begin{array}{c} | \\ \uparrow \\ a_2 \end{array} \begin{array}{c} | \\ \uparrow \\ a_3 \end{array} \dots \begin{array}{c} | \\ \uparrow \\ a_n \end{array} . \end{aligned}$$

Then

$$U_{\Gamma_n}^{\Gamma_m} = \mathbb{I}_{\Gamma_m} U_{\Gamma_n}^{\Gamma_m} \mathbb{I}_{\Gamma_n} , \quad (4.9)$$

which expresses the fact that the domain and range of the operator is restricted to states

with certain particle configurations.

Similar to tensor network states, we represent charge-preserving operators of the form (4.8) taking the space of m anyons to the space of n anyons by shaded boxes with m ingoing and n outgoing (unlabeled) edges; summation over these edge labels is left implicit. It is important to note, however, that such a box represents a different object compared to the case of tensor networks: it is defined by a family of maps $\{U_{\Gamma_n}^{\Gamma_m}(c)\}_{c \in \Omega}$, which specify a weighted superposition of certain labeled trivalent graphs embedded in the box, with m in- and n outgoing edges as in (4.6).

The diagrammatic representation of anyonic operators and vectors satisfies simple rules with respect to composition, tensor products and (partial) tracing. Adjoint operators are obtained by flipping the diagram, reversing the arrows and complex conjugating all coefficients. Operators are multiplied or applied to vectors by stacking their representations on top of each other and connecting up out- with ingoing edges. Tensor products are obtained by placing diagrams next to each other.

Traces and partial traces corresponds to connecting up in- and outgoing strands of an operator, while simultaneously inserting the operator

$$\begin{array}{c} \uparrow \\ \textcircled{\mathbb{D}} \\ \uparrow \end{array} = \begin{array}{c} \uparrow \\ \textcircled{\mathbb{D}} \\ | \end{array} = \begin{array}{c} | \\ \textcircled{\mathbb{D}} \\ \uparrow \end{array} = \sum_c d_c^{-1} |c\rangle \langle c| \quad (4.10)$$

into each line. (Here we use the convention that either input or output of an operator that acts diagonally on anyon labels may be represented by an undirected edge.) For example, the partial trace over the n th anyon for an operator U acting on n anyons takes the form

$$\text{tr}_n \begin{array}{c} \uparrow \dots \uparrow \\ \text{O} \\ \uparrow \dots \uparrow \end{array} = \begin{array}{c} \uparrow \dots \uparrow \\ \text{O} \\ \uparrow \dots \uparrow \end{array} \begin{array}{c} \textcircled{\mathbb{D}} \\ \curvearrowright \end{array} .$$

The (complete) trace is obtained by connecting up all the strands in this way and computing

the vacuum coefficient, i.e.,

The diagram on the right hand side. in (4.11) is in fact proportional to the vacuum graph. Using (4.3) repeatedly, it is straightforward to check that the states in (4.1) are orthonormal.

The rules for the diagrammatic representation of anyonic operators are remarkably similar to the contraction of tensor networks, although their origin is distinct. In particular, the notion of evaluating of a diagram is very different: A tensor network associates a scalar quantity to every labeling of the edges. This means that contraction, i.e., summing over all labelings, results in a scalar. In contrast, anyonic diagrams associate a trivalent graph to every labeling, and the contraction results in a formal linear combination of equivalence classes of labeled trivalent graphs. Equivalence is defined by isotopy and the local rules from (4.2), (4.3) and (4.4).

4.2.2.2 Anyons on a torus

Here we are interested in periodic chains of anyons arranged on a line, and a few modifications of the above formalism are necessary. As in [76, 86, 100], we consider a chain of anyons arranged on a topologically nontrivial cycle wrapping around a torus. We denote the space of n anyons of types $\vec{a} = (a_1, \dots, a_n)$ arranged on such a chain with periodic boundary conditions by $V_{\text{periodic}}^{\vec{a}}$. This space does not naturally decompose into subspaces with specified total charge. An orthonormal basis is given by the basis vectors

$$|\vec{a}, \vec{b}\rangle_n = \frac{1}{(\prod_i d_{a_i})^{1/4}} \begin{array}{c} \uparrow a_1 \quad \uparrow a_2 \quad \dots \quad \uparrow a_{n-1} \quad \uparrow a_n \\ \rightarrow b_n \quad \rightarrow b_1 \quad \rightarrow b_2 \quad \dots \quad \rightarrow b_{n-2} \quad \rightarrow b_{n-1} \quad \rightarrow b_n \end{array}, \quad (4.12)$$


where $\vec{b} \in \Omega^n$ is such that each vertex satisfies the fusion rules. The basis specified by (4.12) corresponds to a decomposition of the Hilbert space as

$$V_{\text{periodic}}^{\vec{a}} \cong \bigoplus_{\vec{b}=(b_1, \dots, b_N)} V_{b_N}^{a_1 b_1} \otimes V_{b_1}^{a_1 b_2} \otimes \dots \otimes V_{b_{N-1}}^{a_N b_N} .$$

The representation of the bra $\langle \vec{a}, \vec{b} |_n$ of the vector in (4.12) is again obtained by flipping the diagram and reversing the arrows, i.e.,

$$\langle \vec{a}, \vec{b} |_n = \frac{1}{(\prod_i d_{a_i})^{1/4}} \begin{array}{c} b_n \quad b_1 \quad b_2 \quad \dots \quad b_{n-2} \quad b_{n-1} \quad b_n \\ \leftarrow \quad \leftarrow \quad \leftarrow \quad \dots \quad \leftarrow \quad \leftarrow \quad \leftarrow \\ \uparrow \quad \uparrow \quad \uparrow \quad \dots \quad \uparrow \quad \uparrow \\ a_1 \quad a_2 \quad \dots \quad a_{n-1} \quad a_n \end{array} . \quad (4.13)$$

Local operators acting on a subset of the n anyons are represented as before by shaded boxes, but global operators can not be represented as superpositions of graphs embedded in such a planar surface. Instead, it is convenient to embed the chain and associated anyonic diagrams along one of the fundamental nontrivial cycles of the torus. (This is already implicit in (4.12).) A global operator $\mathbf{O} : \bigoplus_{\vec{a} \in \Omega^m} V_{\text{periodic}}^{\vec{a}} \rightarrow \bigoplus_{\vec{a} \in \Omega^n} V_{\text{periodic}}^{\vec{a}}$ mapping between periodic chains of (possibly different) lengths m and n is then represented by a shaded strip parallel to the chain on the torus, with m ingoing and n outgoing edges, i.e.,



Application of such an operator is again equivalent to attaching it to a diagram. Finally, (partial) traces are computed simply by connecting up edges. For example, for an operator \mathbf{O} acting on $\bigoplus_{\vec{a}} V_{\text{periodic}}^{\vec{a}}$, the partial trace over the n th anyon is

$$\text{tr}_n \begin{array}{c} \uparrow \quad \uparrow \quad \dots \quad \uparrow \quad \uparrow \\ \text{---} \quad \mathbf{O} \quad \text{---} \\ \uparrow \quad \uparrow \quad \dots \quad \uparrow \quad \uparrow \end{array} = \begin{array}{c} \text{---} \\ \uparrow \quad \uparrow \quad \dots \quad \uparrow \quad \uparrow \\ \text{---} \quad \mathbf{O} \quad \text{---} \\ \uparrow \quad \uparrow \quad \dots \quad \uparrow \quad \uparrow \end{array} .$$

Here we used isotopy on the torus to get a convenient expression on the right hand side. The (complete) trace is the result of connecting up all strands, and then computing the

coefficient of the empty graph, i.e.,

$$\text{tr} \left[\begin{array}{c} \uparrow \uparrow \dots \uparrow \uparrow \\ \text{---} \text{---} \text{---} \text{---} \text{---} \\ \uparrow \uparrow \dots \uparrow \uparrow \end{array} \right] = \left[\begin{array}{c} \text{---} \text{---} \text{---} \text{---} \text{---} \\ \uparrow \uparrow \dots \uparrow \uparrow \\ \text{---} \text{---} \text{---} \text{---} \text{---} \\ \text{O} \end{array} \right]_{\text{vac}} . \quad (4.14)$$

Here we denote the coefficient of the empty graph in a formal superposition X by $[X]_{\text{vac}}$. That is, it is obtained by writing X as a superposition of states with flux a , i.e., with a line labeled a going around the torus, and taking the coefficient for $a = 1$.

We can now show that the (partial) trace defined in this fashion is equivalent to the orthonormality of the set of vectors from (4.12). That is, we now show that the basis states $\{|\vec{a}, \vec{b}\rangle\}$ of $\bigoplus_{\vec{a}} V_{\text{periodic}}^{\vec{a}}$ (cf. (4.12)) are indeed orthonormal with respect to the Hilbert-Schmidt inner product defined by the trace from (4.14). By definition, we have

$$|\langle \vec{a}', \vec{b}' | \vec{a}, \vec{b} \rangle|^2 = \prod_i \frac{\delta_{a_i, a'_i}}{d_{a_i}^{1/2}} \left[\begin{array}{c} \leftarrow b'_n \leftarrow b'_1 \leftarrow b'_2 \dots \leftarrow b'_{n-2} \leftarrow b'_{n-1} \leftarrow b'_n \\ \uparrow a_1 \uparrow a_2 \dots \uparrow a_{n-1} \uparrow a_n \\ \rightarrow b_n \rightarrow b_1 \rightarrow b_2 \dots \rightarrow b_{n-2} \rightarrow b_{n-1} \rightarrow b_n \\ \text{---} \text{---} \text{---} \text{---} \text{---} \\ \text{O} \end{array} \right]_{\text{vac}} .$$

Inserting the projection $\mathbb{I}_{\{(e_1, e_2)\}}$ onto a pair of anyons (e_1, e_2) decomposed into total charge c as

$$\mathbb{I}_{\{(e_1, e_2)\}} = \begin{array}{c} | \\ \uparrow e_1 \\ | \end{array} \begin{array}{c} | \\ \uparrow e_2 \\ | \end{array} = \sum_c \sqrt{\frac{d_c}{d_{e_1} d_{e_2}}} \begin{array}{c} e_1 \quad e_2 \\ \diagdown \quad \diagup \\ \quad \uparrow c \\ \diagup \quad \diagdown \\ e_1 \quad e_2 \end{array}$$

into the horizontal lines leads to

$$|\langle \vec{a}', \vec{b}' | \vec{a}, \vec{b} \rangle|^2 = \delta_{\vec{a}, \vec{a}'} \sum_{\vec{c}} \left(\prod_i \frac{1}{\sqrt{d_{a_i}}} \right) [X]_{\text{vac}} , \quad (4.15)$$

where

$$X = \sum_{\vec{c}} \sqrt{\prod_j \frac{d_{c_j}}{d_{b_j} d_{b'_j}}} \begin{array}{c} \begin{array}{c} \leftarrow c_n \leftarrow c_1 \leftarrow c_2 \dots \leftarrow c_3 \leftarrow c_2 \dots \leftarrow c_n \\ \uparrow a_1 \uparrow a_2 \dots \uparrow a_3 \uparrow a_2 \dots \uparrow a_n \\ \rightarrow b_n \rightarrow b_1 \rightarrow b_2 \dots \rightarrow b_3 \rightarrow b_2 \dots \rightarrow b_n \end{array} \end{array} .$$

Since we are interested in the vacuum coefficient of X , we can set all $c_i = 1$, getting

$$[X]_{\text{vac}} = \prod_j \frac{\delta_{b_j, b'_j}}{d_{b_j}} \left[\begin{array}{c} \text{Diagram: A sequence of } \theta \text{-like graphs. The first graph has vertices } b_n \text{ and } b_1 \text{ with an internal vertex } a_1. \text{ The second graph has vertices } b_1 \text{ and } b_2 \text{ with an internal vertex } a_2. \text{ This is followed by an ellipsis.} \\ \text{vac} \end{array} \right]. \quad (4.16)$$

But each θ -like graph on the right hand side is proportional to the empty graph, with scalar

$$\left[\begin{array}{c} \text{Diagram: A } \theta \text{-like graph with vertices } c \text{ and } b \text{ and an internal vertex } a. \\ \text{vac} \end{array} \right] = \sqrt{d_a d_b d_c}. \quad (4.17)$$

This can be verified by applying (4.3) twice. Inserting (4.17) into (4.16) gives

$$[X]_{\text{vac}} = \prod_j \delta_{b_j, b'_j} \sqrt{d_{a_j}}.$$

When combined with (4.15), this implies the claim.

4.2.3 Anyonic Hamiltonians: Long-range Effective Theories

A remarkable feature of the state space of anyons is its topological degeneracy: the Hamiltonian of the quantum medium assigns equal energy to each state. Furthermore, this degeneracy is stable under local perturbations, a feature that makes anyons particularly suited for encoding and processing quantum information. These properties hold up to exponentially small corrections in the interanyon distances.

If the interanyon separation falls below a certain length scale, the microscopic details of the system become relevant and the topological degeneracy is generally lifted. Such a degeneracy lifting has been examined in various quantum media [101–103]. In the system-independent anyonic formalism, Bonderson [75] has shown that a general interaction between two anyons can be interpreted as tunneling of topological charge, and that generic tunneling fully lifts the topological degeneracy.

Hamiltonian terms responsible for such tunneling and more generally arbitrary multi-anyon interactions take the form of Hermitian operators as in (4.6) with $m = n$. The exact form of the effective Hamiltonian governing the energy splitting depends on the geometric arrangement of the anyons. In a regular lattice, nearest neighbor ($m = 2$) and next-to-nearest neighbor ($m = 3$) interactions are most relevant physically as the interaction

strength decays exponentially with distance. Latticelike arrangements of anyons arise when certain spatial distributions are energetically favored, e.g., by inserting defects into the quantum medium, which couple to additional quantum numbers (such as electric charge) of the anyons.

Paradigmatic models of such effective Hamiltonians have been considered extensively in the literature. They can be thought of as describing the dynamics of the internal degrees of freedom of anyons pinned to fixed sites. We refer to [86] for an introductory discussion of such models. We discuss explicit examples for Fibonacci anyon chains in section 4.4.

4.3 Anyonic Entanglement Renormalization

4.3.1 The Setting

We have described the origin of anyonic Hamiltonians as long-range effective descriptions of quantum media in section 4.2.3. We now turn to the problem of defining a variational ansatz for such systems. For concreteness, we restrict our attention to one-dimensional chains of anyons, arranged in a periodic fashion along a topologically nontrivial cycle on the torus. We discuss more general two-dimensional arrangements in section 4.5.

One may consider different spatial distributions of anyons on the chain. For example, in a setting with several nontrivial anyon types, one may be interested in the effective behavior of a staggered, i.e., alternating arrangement of anyons. While our ansatz could in principle be adapted to such cases, here we consider the simplest nontrivial setting. We assume that a subset $\Omega_{\text{eff}} \subset \Omega$ of anyons is allowed in each site. The Hilbert space of a (periodic) chain of length n on the torus is then given by

$$\mathcal{H}_{\text{chain},n} \cong \bigoplus_{\vec{a} \in \Omega_{\text{eff}}^n} V_{\text{periodic}}^{\vec{a}} . \quad (4.18)$$

The most commonly considered case (e.g., the golden chain [76]) is when $\Omega_{\text{eff}} = \{a\}$ consists of a single anyon a , that is, each site is occupied by a particle of type a . Our formulation is slightly more general, as it allows to consider Hamiltonians that can create and destroy particles on sites of the chain (resp. change particle types) when $\Omega_{\text{eff}} = \Omega$. This is important when the quantum medium assigns nearly degenerate energies to different distributions of anyons. Further intermediate cases could be considered.

4.3.2 The Ansatz

Consider a MERA-structure for a periodic chain with N sites as in figure 3.6. We associate to this structure a family of states in the anyon Hilbert space $\mathcal{H}_{\text{chain},N}$ (cf. (4.18)) as follows:

1. To the strip at the top, associate the bra $\langle\varphi|$ of a normalized state

$$|\varphi\rangle \in \mathcal{H}_{\text{chain},n}, \quad (4.19)$$

that is,

$$\begin{array}{c} \text{---} \\ \uparrow \uparrow \dots \uparrow \uparrow \end{array} \mapsto \langle\varphi|, \quad (4.20)$$

where n is the number of ingoing edges.

2. To every intermediate box with m ingoing and n outgoing strands ($m \geq n$), associate the adjoint \mathbf{W}^\dagger of an isometric charge-conserving map

$$\mathbf{W} \in \text{End} \left(\bigoplus_{\vec{a} \in \Omega_{\text{eff}}^n} V^{\vec{a}}, \bigoplus_{\vec{a}' \in \Omega_{\text{eff}}^m} V^{\vec{a}'} \right), \quad (4.21)$$

that is,

$$\begin{array}{c} \uparrow \dots \uparrow (n \text{ strands}) \\ \text{---} \\ \uparrow \dots \uparrow (m \text{ strands}) \end{array} \mapsto \mathbf{W}^\dagger.$$

3. Regard the state $|\varphi\rangle$ ((4.20)) and the family of maps $\{\mathbf{W}\}$ ((4.21)) as variational parameters specifying a state $|\Psi\rangle = |\Psi_{|\varphi\rangle, \{\mathbf{W}\}}\rangle \in \mathcal{H}_{\text{chain},N}$ of the chain. The state is determined by the following recipe: in the MERA-structure of figure 3.6, replace every box by the superposition of trivalent labeled graphs representing the associated object (i.e., (4.20) and (4.21)). The result is a superposition of labeled graphs, each with N ingoing edges. This superposition represents $\langle\Psi|$.

More explicitly, the maps in (4.21) are of the form

$$\mathbf{W} = \bigoplus_c \mathbf{W}(c) , \quad (4.22)$$

with $\mathbf{W}(c) \in \text{End}(\bigoplus_{\vec{a} \in \Omega_{\text{eff}}^n} V_c^{\vec{a}}, \bigoplus_{\vec{a}' \in \Omega_{\text{eff}}^m} V_c^{\vec{a}'})$ satisfying

$$\mathbf{W}(c)^\dagger \mathbf{W}(c) = \mathbb{I}_{\bigoplus_{\vec{a} \in \Omega_{\text{eff}}^n} V_c^{\vec{a}}} . \quad (4.23)$$

Eqs. 4.22 and 4.23 severely constrain the set of allowed maps \mathbf{W} for certain (m, n) and Ω_{eff} .

This ansatz is clearly motivated by entanglement renormalization for qudits. In fact, it has the same operational interpretation: the MERA-structure of figure 3.6, after replacing each box by the superposition of graphs specified by $\{\mathbf{W}\}$, is a procedure for successively mapping the chain to a coarse-grained chain by local gates and isometries. Indeed, due to charge conservation, boxes with the same number of in- and outgoing edges correspond to local unitaries on the anyons, while boxes with fewer outputs than inputs correspond to local isometries reducing the number of degrees of freedom. Such reductions preserve the total charge in their support. Importantly, since the range of each operator \mathbf{W}^\dagger is contained in $\bigoplus_{\vec{a} \in \Omega_{\text{eff}}^m} V_c^{\vec{a}}$, states supported on $\mathcal{H}_{\text{chain}, N}$ are mapped to coarse-grained chains of the same type at each level in Fig 3.6, i.e., with particles from the subset $\Omega_{\text{eff}} \subset \Omega$ on all sites.

4.3.3 Efficient Evaluation of Physical Quantities

Having introduced a set of variational ansatz states parametrized by $(|\varphi\rangle, \{\mathbf{W}\})$, we argue that quantities of physical interest such as expectation values of local observables and correlation functions can be efficiently computed from these parameters. Here we use the formal equivalence of the manipulation rules of anyonic states and operators with usual tensor contractions.

The anyonic analog of (3.10) is

$$\begin{array}{c} \begin{array}{|c|} \hline \uparrow \\ \hline \end{array} \begin{array}{|c|} \hline \uparrow \\ \hline \end{array} \begin{array}{|c|} \hline \uparrow \\ \hline \end{array} \\ \hline \square \square \square \end{array} = \begin{array}{c} \begin{array}{|c|} \hline \uparrow \\ \hline \end{array} \begin{array}{|c|} \hline \uparrow \\ \hline \end{array} \begin{array}{|c|} \hline \uparrow \\ \hline \end{array} \\ \hline \square \text{W}^\dagger \\ \hline \begin{array}{|c|} \hline \uparrow \\ \hline \end{array} \begin{array}{|c|} \hline \uparrow \\ \hline \end{array} \begin{array}{|c|} \hline \uparrow \\ \hline \end{array} \\ \hline \square \text{W} \\ \hline \begin{array}{|c|} \hline \uparrow \\ \hline \end{array} \begin{array}{|c|} \hline \uparrow \\ \hline \end{array} \begin{array}{|c|} \hline \uparrow \\ \hline \end{array} \end{array} \quad (4.24)$$

for any operator $\mathbf{W} \in \text{End} \left(\bigoplus_{\vec{a} \in \Omega_{\text{eff}}^n} V^{\vec{a}}, \bigoplus_{\vec{a}' \in \Omega_{\text{eff}}^m} V^{\vec{a}'} \right)$ with the required isometry property (where $m \geq n$). Each of the small dark boxes represents the projection $\mathbb{I}_{\Omega_{\text{eff}}}$ onto the subset of allowed anyon labels. Note that, if the edges on the left-hand side of (4.24) are connected to an operator \mathbf{W}' as specified in the ansatz, these may be omitted. This is because both the support and range of \mathbf{W}' are already restricted to the set of allowed anyons Ω_{eff} (cf. (4.21)). In particular, we formally recover the rule (3.10) in this case.

The second important ingredient is the formula in (4.14) for the trace of an operator on the chain. (4.24) and (4.14) immediately imply that expectation values of local operators can be efficiently evaluated for a state $|\Psi\rangle = |\Psi_{|\varphi\rangle, \{\mathbf{W}\}}\rangle$ in essentially the same manner as for MERA states of qudits. The same is true for two-point correlation functions for certain distances of the points related to the MERA structure.

Consider for example a local observable \mathbf{O} acting on two sites of the chain. The expectation value of this operator, given a density matrix ρ , is equal to the diagrammatic expression

$$\text{tr}(\mathbf{O}\rho) = \left[\begin{array}{c} \text{---} \\ \uparrow \uparrow \uparrow \uparrow \uparrow \\ \dots \uparrow \uparrow \uparrow \uparrow \uparrow \\ \uparrow \uparrow \uparrow \uparrow \uparrow \\ \text{---} \\ \rho \\ \uparrow \uparrow \uparrow \uparrow \uparrow \\ \text{---} \\ \text{vac} \end{array} \right] . \quad (4.25)$$

If $\rho = |\Psi\rangle\langle\Psi|$ is an anyonic MERA-state corresponding, e.g., to the structure of figure 3.6 (a), we conclude from (4.25) that

$$\langle\Psi|\mathbf{O}|\Psi\rangle = [X]_{\text{vac}} ,$$

where X is the superposition of trivalent labeled graphs specified by the diagram in figure 3.7 (b). Using (4.24), this immediately reduces to $[X']_{\text{vac}}$, where X' is the superposition in figure 3.7 (c). This can be efficiently evaluated using the superoperators defined in (3.12) and (3.13) (again interpreted as anyonic expressions). Two-point correlations functions can be computed analogously.

4.3.4 Computational Cost and Refinements of the Ansatz

To count the number of parameters needed to describe the anyonic ansatz states, let $\mathcal{D} = \max_{a \in \Omega_{\text{eff}}} d_a$ be the maximal quantum dimension of the particles used. Since the number

of states of the form in (4.1) can be upper bounded by $O(\mathcal{D}^n)$, a charge-conserving map as in (4.21) is described by fewer than $O(\mathcal{D}^{n+m})$ parameters. Similarly, a state $|\varphi\rangle$ as in (4.19) is described by $O(\mathcal{D}^n)$ parameters.

As with the MERA for spin chains, the family of ansatz states may be enlarged by replacing an edge by $s > 1$ edges; this is analogous to increasing the bond dimension. In this case, isometries are described by $O(\mathcal{D}^{s(n+m)})$ parameters as opposed to $O(d^{s(n+m)})$ in the qudit case. (Note that, by definition, $\mathcal{D} < d$ for any embedding of anyonic states into qudits: for example, for the Fibonacci anyons considered below, $\mathcal{D} \approx 1.618$.) In summary, a general anyonic MERA is described by $O(\text{poly}(\mathcal{D}^s)N\log N)$ parameters, and translation-invariant and scale-invariant MERAs by $O(\text{poly}(\mathcal{D}^s)\log N)$ and $O(\text{poly}(\mathcal{D}^s))$ parameters, respectively.

The remainder of this story is the same as that of the MERA for spin chains, and we refer to the extensive literature (e.g., [39]) on this subject. For example, the computational cost of computing local expectation values is roughly the same as for qudit chains (with χ replaced by \mathcal{D}^s), and methods used, e.g., for varying over the isometries in figure 3.6 can directly be adapted to the anyonic framework. Compared to the qudit chain setting, an additional advantage stems from the fact that the isometries are charge preserving and thus take a block-diagonal form. Therefore, matrix multiplication and singular value decompositions can be performed on matrices whose dimensions are a constant factor smaller than with a naïve qudit ansatz. Similarly, methods for extracting critical exponents from scale-invariant MERA states [49, 104] may be applied in the anyonic setting.

4.3.5 Example: Fibonacci Anyons

The Fibonacci theory has one nontrivial particle τ with quantum dimension $d_\tau = \phi = \frac{\sqrt{5}+1}{2}$ equal to the golden ratio and fusion rule $\tau \times \tau = 1 + \tau$. The F -matrix is given by

$$\begin{aligned} \begin{array}{c} \diagup \\ \cdot \\ \diagdown \end{array} &= \frac{1}{\phi} \begin{array}{c} \diagup \\ \cdot \\ \cdot \\ \diagdown \end{array} + \frac{1}{\sqrt{\phi}} \begin{array}{c} \diagup \\ \cdot \\ \diagup \\ \cdot \\ \diagdown \end{array}, \\ \begin{array}{c} \diagdown \\ \cdot \\ \diagup \end{array} &= \frac{1}{\sqrt{\phi}} \begin{array}{c} \diagup \\ \cdot \\ \cdot \\ \diagdown \end{array} - \frac{1}{\phi} \begin{array}{c} \diagup \\ \cdot \\ \diagup \\ \cdot \\ \diagdown \end{array}, \end{aligned} \tag{4.26}$$

where we use the convention that a solid line represents the τ -label, while a dotted line represents the trivial label 1. Consider a periodic chain of τ -anyons. To get a corresponding family of ansatz states, we set $\Omega_{\text{eff}} = \{\tau\}$. Using the convention that a solid line represents

an edge with label τ and a dotted line represents an edge with label 1, the Hilbert spaces $\mathcal{H}_{\text{chain},n}$ of a periodic chains with $n \in \{1, 2\}$ particles are spanned by (cf. (4.12))

$$\begin{aligned} \mathcal{H}_{\text{chain},1} &= \mathbb{C} \begin{array}{c} | \\ \hline \end{array} =: \mathbb{C}|\tau\rangle, \\ \mathcal{H}_{\text{chain},2} &= \text{span} \left\{ \begin{array}{c} | \quad | \\ \hline \end{array}, \begin{array}{c} | \quad | \quad | \\ \hline \end{array}, \begin{array}{c} | \quad | \quad | \quad | \\ \hline \end{array} \right\} \\ &=: \text{span} \{|\tau\tau\rangle, |1\tau\rangle, |\tau 1\rangle\}. \end{aligned} \quad (4.27)$$

In general, the space $\mathcal{H}_{\text{chain},n}$ is spanned by vectors corresponding to (periodic) sequences of 1s and τ s, with the fusion constraint forbidding neighboring 1s (this defines an embedding into a subspace of $(\mathbb{C}^2)^{\otimes n}$). This determines the form of the state $|\varphi_n\rangle \in \mathcal{H}_{\text{chain},n}$ corresponding to the top box in structures as in figure 3.6 (a) and (b), respectively. Next, we consider the constraints on the maps $\mathbf{W} = \mathbf{W}_{(n,m)}$ in (4.22) with n input and m output anyons, for small (n, m) . The standard form of charge preserving isometries/unitaries is

$$\mathbf{W}_{(1,2)} = \frac{e^{i\theta}}{\phi^{1/4}} \begin{array}{c} \vee \\ | \\ \wedge \end{array}, \quad (4.28)$$

$$\mathbf{W}_{(2,2)} = \frac{e^{i\theta_1}}{\phi^{1/2}} \begin{array}{c} \vee \\ | \\ \wedge \end{array} + \frac{e^{i\theta_2}}{\phi} \begin{array}{c} \vee \\ \vdots \\ \wedge \end{array}, \quad (4.29)$$

$$\mathbf{W}_{(1,3)} = \frac{\alpha}{\phi^{1/2}} \begin{array}{c} \vee \\ | \\ \wedge \end{array} + \frac{\beta}{\phi^{1/2}} \begin{array}{c} \vee \\ \vdots \\ \wedge \end{array}, \quad (4.30)$$

$$\mathbf{W}_{(2,3)} = \frac{e^{i\theta}}{\phi^{5/4}} \begin{array}{c} \vee \\ | \\ \wedge \end{array} + \frac{1}{\phi^{3/4}} \left(\alpha \begin{array}{c} \vee \\ | \\ \wedge \end{array} + \beta \begin{array}{c} \vee \\ \vdots \\ \wedge \end{array} \right), \quad (4.31)$$

where $\theta \in [0, 2\pi[$ and $|\alpha|^2 + |\beta|^2 = 1$. In all these cases, the map is completely specified by a phase and/or a qubit state. A less trivial case is $\mathbf{W}_{(3,3)}$, which is equal to a phase times the projection onto the span of $\begin{array}{c} \vee \\ | \\ \wedge \end{array}$ plus a two-by-two unitary $\mathbf{W}(\tau)$ on the span of $\left\{ \begin{array}{c} \vee \\ \vdots \\ \wedge \end{array}, \begin{array}{c} \vee \\ | \\ \wedge \end{array} \right\}$. For later reference, we also state the most general form of an isometry with

two input and four output strands:

$$\begin{aligned}
 \mathbf{W}_{(2,4)} &= \frac{1}{\phi^{3/2}} \left(\alpha \begin{array}{c} \text{---} \text{---} \\ \text{---} \text{---} \\ \text{---} \text{---} \\ \text{---} \text{---} \end{array} + \beta \begin{array}{c} \text{---} \text{---} \\ \text{---} \text{---} \\ \text{---} \text{---} \\ \text{---} \text{---} \end{array} \right) \\
 &+ \frac{1}{\phi} \left(\gamma \begin{array}{c} \text{---} \text{---} \\ \text{---} \text{---} \\ \text{---} \text{---} \\ \text{---} \text{---} \end{array} + \delta \begin{array}{c} \text{---} \text{---} \\ \text{---} \text{---} \\ \text{---} \text{---} \\ \text{---} \text{---} \end{array} + \varepsilon \begin{array}{c} \text{---} \text{---} \\ \text{---} \text{---} \\ \text{---} \text{---} \\ \text{---} \text{---} \end{array} \right), \tag{4.32}
 \end{aligned}$$

where $|\alpha|^2 + |\beta|^2 + |\gamma|^2 + |\delta|^2 + |\varepsilon|^2 = 1$.

From (4.27), (4.28) and (4.29), we conclude that the family of states associated with figure 3.6 (a) is rather uninteresting as the variational parameters $(|\varphi\rangle, \{\mathbf{W}\})$ are merely a set of phases. In contrast, the structure in figure 3.6 (b) gives rise to a less trivial family of states due to (4.30).

We will give additional nontrivial explicit examples in section 4.4.

4.3.6 Distillable States for Composite Anyon Coding

The MERA ansatz for qudits is motivated by quantum circuits. Indeed, a MERA-description of a state provides a circuit preparing the state starting from the top-level state $|\varphi\rangle$. This is achieved by realizing isometries using ancillae prepared in pure states; it corresponds to running the coarse-graining procedure in reverse. This use of the MERA has been proposed for example as a way of efficiently preparing topologically ordered states [43].

The one-to-one correspondence between preparation circuits of a certain form and states described by the entanglement renormalization ansatz clearly extends to anyons. However, there is an additional relation in the anyonic setting, which corresponds to running the coarse-graining forwards (instead of backwards as in the qudit case): a subclass of the anyonic ansatz is in one-to-one correspondence with distillation procedures preparing a logical state $|\varphi\rangle$ of composite anyons. The anyonic renormalization ansatz therefore provides an alternative characterization of “distillable” states in the framework of composite anyon coding [81].

The goal of composite anyon coding is to prepare a suitable state for computation starting from some unknown initial state, but without using measurements. In the terminology of anyonic entanglement renormalization, composite anyons are anyons at higher levels in

the coarse-graining scheme. Initial states that allow to prepare a given target state $|\varphi\rangle$ of composite anyons can be characterized as follows. They can be represented by an anyonic renormalization ansatz (corresponding to the preparation scheme) with the following properties: the state at the top is fixed to $|\varphi\rangle$, and all adjoints of isometries \mathbf{W}^\dagger are implementable by braiding and fusion. The latter condition means that the renormalization scheme only consists of unitaries effected by braiding, and coarse-graining, that is, any reduction in the number of anyons, is achieved by bunching together some particles. Such fusion-based coarse-graining is given by a product of the isometry

$$\sum_c |a, b; c\rangle \langle c| = \sum_c \left(\frac{d_c}{d_a d_b} \right)^{1/4} \begin{array}{c} a \quad b \\ \diagdown \quad / \\ \quad c \end{array} \quad a = a(c), b = b(c)$$

whose adjoint describes the fusion of a pair of particles. An anyonic renormalization ansatz with these properties can directly be implemented using the operations commonly envisioned to be available for manipulating anyons. Thus we can regard such schemes as state preparation circuits for topological quantum computers.

4.4 Application to the Golden Chain and the Majumdar-Ghosh Chain

4.4.1 The Model

In this section, we consider the use of the anyonic MERA ansatz in the context of the golden chain [76] and its relatives. An introduction to these models can be found in [86]. Specifically, we consider the Fibonacci-anyonic analog of the Heisenberg and the Majumdar-Ghosh (MG) spin chains. The former was introduced in [76] and consists of a uniform chain of N Fibonacci anyons with Hamiltonian terms favoring one of either possible total charge of two neighboring τ -particles. Concretely, the Hamiltonian is given by

$$J_2 \cdot H^{\text{golden}} = -J_2 \sum_{i=1}^N H_2^i, \quad (4.33)$$

where each term H_2^i is a projection onto trivial charge of the anyons i and $i + 1$, i.e., it has the diagrammatic representation $\frac{1}{\phi} \begin{array}{c} \vee \\ \vdots \\ \wedge \end{array}$. In analogy with the familiar $SU(2)$ spin chains, the




Figure 4.1. Phase diagram of the model in (4.35) as obtained in [106]. Phases I and II are gapped, with the exact ground states known at the Majumdar-Ghosh (MG) point $\theta = 3\pi/2$ (see Section 4.4.2) and at $\tan \theta = \phi/2$. There are two extended critical phases for which an exact mapping was established to the 3-state Potts and the tricritical Ising model at the FM- and AFM-golden chain points, respectively [76]. A small sliver of an incommensurate phase is found near $\theta = 1.075\pi$ next to a phase with \mathbb{Z}_4 -symmetry; both these phases are believed to be critical. See [106] for a detailed discussion.

case $J_2 > 0$ energetically favoring trivial total charge is referred to as “antiferromagnetic” (AFM) coupling, whereas $J_2 < 0$ is called “ferromagnetic” (FM) coupling.

The Majumdar-Ghosh (MG) chain [105] is a model of $SU(2)$ spin-1/2 particles arranged on a chain, with three-particle interactions favoring either total spin 3/2 (called ferromagnetic/FM) or 1/2 (called antiferromagnetic/AFM). Its anyonic analog [106] takes the form

$$J_3 \cdot H^{\text{MG}} = -J_3 \sum_{i=1}^N H_3^i, \quad (4.34)$$

where H_3^i is the projection onto trivial charge of three neighboring anyons. Using the F-matrix from (4.26), it is straightforward to rewrite the terms H_2^i and H_3^i in the standard basis from (4.12). Corresponding expressions can be found in [106]. Expressed in the standard embedding into $(\mathbb{C}^2)^{\otimes n}$, this leads to 3- and 4-qubit terms, respectively.

The Hamiltonian in (4.33) was studied in detail in [76]. Criticality and a two-dimensional CFT description were established numerically. Furthermore, an exact mapping to a standard integrable two-dimensional classical lattice model [107] known as the RSOS model was given. These studies were extended to the one-parameter family of Hamiltonians

$$H_\theta := \cos \theta \cdot H^{\text{golden}} + \sin \theta \cdot H^{\text{MG}}, \quad (4.35)$$

the chain.

To specify the scale-invariant MERA-ansatz describing the ground space from (4.36), consider the MERA-structure of Figure 3.6(a) with refinement parameter $s = 2$, i.e., with every strand replaced by two. Due to this doubling and scale invariance, the corresponding ansatz then is described by an isometry $\mathbf{W} = \mathbf{W}_{(2,4)}$ of the form in (4.32) and a unitary \mathbf{U} acting on four τ -anyons. We set the unitary equal to the identity, and the isometry equal to

$$\mathbf{W} := \frac{1}{\sqrt{\phi}} \left| \cup \right\rangle, \quad (4.37)$$

which corresponds to the parameters $(\alpha, \beta, \gamma, \delta, \epsilon) = (\frac{1}{\phi}, \frac{1}{\sqrt{\phi}}, 0, \frac{1}{\phi}, \frac{1}{\sqrt{\phi}})$ in (4.32). This completes the specification of the MERA up to translation (which we fix later), as the top-level state $|\varphi\rangle$ depends on the actual state considered. Note that the MERA-structure of Figure 3.6(a) with unitaries equal to the identity is known as a tree-tensor network [109–111] in the nonanyonic setting.

Let us argue that the renormalization group scheme has the ground states from (4.36) as fixed points. This is most easily seen for the state $|1\tau 1\tau 1\tau \cdots\rangle$ using the diagrammatic calculus: applying a layer of (adjoints of) isometries corresponds to stacking $N/4$ parallel copies of (4.37) on top of the state. We assume that the isometries are aligned in such a way that this takes the form

for the state $|1\tau 1\tau 1\tau \cdots\rangle$, where we suppressed scalar factors. Here we used isotopy invariance and (4.3). With (4.12), it is easy to verify that the operation $(\mathbf{W}^\dagger)^{\otimes N/4}$ also preserves the norm of the state. This establishes the claimed fixed-point property for $|1\tau 1\tau 1\tau \cdots\rangle$. To verify the claim for the second state $|\tau\tau_x\tau\tau_x\tau\tau_x \cdots\rangle$, it suffices to observe that this state is essentially equivalent to the former but with τ -flux, i.e.,

because of the identity

$$\text{---} \text{---} \text{---} = \text{---} \text{---} \text{---} .$$

This immediately implies that this state is fixed by the same renormalization group scheme.


4.4.3 Numerical Variation over Ansatz States

To assess the suitability of the anyonic entanglement renormalization ansatz as a numerical method, we have implemented an algorithm for numerically minimizing the energy by varying over the parameters of a translation-invariant ansatz. The algorithm is based on iterative optimization of the (identical) isometries at each level and the top-level state. It is described in detail in [39] for nonanyonic spin chains. For a fixed isometry, it proceeds by computing its *environment*, that is, the contraction of the MERA-network with the isometry omitted. The resulting tensor can be interpreted as a linear map whose singular value decomposition dictates how the isometry is updated. Adapting this to the anyonic setting is straightforward: here the environment always has a block-diagonal form with respect to total charge. Compared to the algorithm of [39], the only significant difference lies in the implementation of the ascending and descending superoperators (see, e.g., (3.12) and (3.13)) used to compute the environments. As with all anyonic operations, they require applying basis changes into compatible treelike bases (cf. (4.2)). These basis changes can be precomputed.


This randomized optimization algorithm is susceptible to local minima, and its convergence depends on the choice of initial points. In practice, these issues appear to be minor and can be addressed by starting with a large number of initial points and postselecting after a few iterations.

Ground state energy and correlation functions

We have applied the variational algorithm to periodic chains of 12 and 16 Fibonacci anyons governed by the Hamiltonian in (4.35). These system sizes were chosen to allow for comparison with exact diagonalization data and to test the suitability of different MERA structures.




(a) Ternary MERA structure applied to 12 anyons



(b) Binary MERA structure applied to 16 anyons

Figure 4.2. Ground state energy approximated variationally by an anyonic entanglement renormalization ansatz. The inset in each figure shows the relative error $\Delta E = (E_{\text{MERA}} - E_0)/(E_{\text{max}} - E_0)$. Panel (a) is based on a ternary anyonic MERA for a system of 12 anyons, while panel (b) shows the result of a binary anyonic MERA for a chain of 16 anyons. At the Majumdar-Ghosh point ($\theta = 3\pi/2$), the variational procedure recovers the exact fixed-point discussed in section 4.4.2 for the binary MERA. In (b), the approximation is best around around the golden chain-point with AFM-couplings $\theta = 0$. This is consistent with the fact that the ground state has a \mathbb{Z}_2 -sublattice ordering [76], which is compatible with the coarse-graining structure of a binary MERA. For $\theta = \pi$, the ground state has a \mathbb{Z}_3 -sublattice ordering for which the binary MERA structure is less suited. However, the approximation is still better than in intermediate regions where the Hamiltonian has both 2-local and 3-local terms. Similarly, panel (a) shows that a ternary structure appears to be suitable for capturing the \mathbb{Z}_3 -sublattice ordering at the golden chain-FM-point $\theta = \pi$.

For the 12- and 16-anyon chains, we use the “ternary” and “binary” MERA-structures



with $s = 2$ (i.e., every strand is doubled). The former consists of a single level of coarse-graining isometries, while the latter has two such levels. The variationally obtained ground state energy is compared to its exact value in figure 4.2. As shown, we find good agreement between the variationally estimated ground state energy and its exact value, over a wide range of values of the parameter θ .

Figure 4.3 shows that ground state energies computed using the binary MERA are sufficient to obtain a rough estimate for the location of the phase boundaries.

To study whether the anyonic MERA correctly reproduces correlations in the ground state, we have computed the (translation-averaged) two-point correlation functions

$$C(r) = \frac{1}{N} \sum_{i=1}^N (\langle H_2^i H_2^{i+r} \rangle - \langle H_2^i \rangle \cdot \langle H_2^{i+r} \rangle) \quad (4.39)$$

of the local topological charge density (as measured in terms of the local projection H_2^i onto trivial charge for a pair). The result of this computation for the AFM-point $\theta = 0$ and the FM-point $\theta = \pi$ are shown in figure 4.4 and figure 4.5, respectively. They exhibit a remarkably good agreement with the exact values.

In figure 4.6, we show the error when computing nearest-neighbor and long-range correlations as a function of the Hamiltonian parameter θ . We observe good accuracy in regions where the ground state energy is well approximated (compare figure 4.2). As expected, the considered MERA-structures are less suited for describing ground state correlations at intermediate values of θ , where the Hamiltonian has both 2-anyon and 3-anyon interactions.

We emphasize that the structures in (4.38) with $s = 2$ are two of the simplest possible leading to nontrivial families of ansatz states for Fibonacci anyons. The support of the unitaries/isometries (4 anyons) is only marginally larger than that of the terms in the Hamiltonian from (4.35). This suggests that the approximation by such states may be




Figure 4.3. The first and second (inset) derivative of the ground state energy with respect to the parameter θ approximately reveals the location of the phase boundaries (red vertical lines). This data was obtained from the binary MERA for 16 anyons used to produce figure 4.2 (b). For each of the 40 points $\theta_0 \in [0, 2\pi]$, the ground state energies $E_0(\theta_0 \pm \Delta\theta)$ at two neighboring points at distance $\Delta\theta = 10^{-3}$ were approximated using the MERA. The plots show the corresponding discrete approximation to the first and second derivative at each point. We stress that only limited information can be gained from this plot. In particular, it does not reveal the nature of the phase transition. We refer to the detailed discussion in [86], where, e.g., the CFT descriptions of the transitions out of the tricritical Ising phase have been identified.



(a) $C(r)$ with the ternary MERA at the AFM-point(b) $C(r)$ with the binary MERA at the AFM-point

Figure 4.4. Two-point correlation function $C(r)$ (cf. (4.39)) of the local topological charge density at the AFM point. They reveal a \mathbb{Z}_2 -sublattice ordering of the ground state wave function. The inset shows the absolute error $|\delta C(r)| = |C_{\text{MERA}}(r) - C_{\text{exact}}(r)|$.



(a) $C(r)$ with the ternary MERA at the FM-point(b) $C(r)$ with the binary MERA at the FM-point

Figure 4.5. Two-point correlation function $C(r)$ (cf. (4.39)) of the local topological charge density at the FM point. The inset shows the absolute error $|\delta C(r)| = |C_{\text{MERA}}(r) - C_{\text{exact}}(r)|$.



(a) $|\delta C(1)|$ and $|\delta C(6)|$ for the ternary MERA(b) $|\delta C(1)|$ and $|\delta C(8)|$ for the binary MERA

Figure 4.6. The absolute error $|\delta C(r)| = |C_{\text{MERA}}(r) - C_{\text{exact}}(r)|$ of the nearest-neighbor ($r = 1$) and long-range ($r = 6$ respectively $r = 8$) charge-charge correlation functions. Qualitatively, these deviations agree with the errors in the ground state energy (see figure 4.2). Note that at the Majumdar-Ghosh point ($\theta = 3\pi/2$), the plotted $|\delta C|$ is large even though the MERA accurately represents one of the ground states. This is because of the degeneracy and the fact that C_{exact} is computed from the completely mixed state on the ground space.

rather coarse. Nevertheless, the ansatz provides reasonable approximations to the ground state energies (for all θ), and correlation functions at the AFM- and FM-points where the Hamiltonian consists of nearest-neighbor-terms. This illustrates that anyonic entanglement renormalization successfully exploits the constraints imposed by conservation of topological charge. Future work may go beyond this proof of principle by considering refined families with parameters $s > 2$. This should lead to significant improvements as in the nonanyonic setting: Here accurate results for correlation functions are usually obtained only for high bond dimension (e.g., $\chi = 22$ for Ising and Potts chains at criticality [39]).

Larger systems

To test the scalability of the method, we have additionally computed the ground state energy of chains of length $N \in \{32, 64, 128\}$ using the binary MERA-structure (with $s = 2$) obtained by adding levels to (4.38). Since this is beyond the reach of exact diagonalization, we consider the AFM-point, which allows us to compare our results to the CFT-predictions of [76].

Explicitly, we use the fact that the low-lying spectrum of a periodic one-dimensional critical quantum systems of length N takes the form

$$E = \varepsilon N + \frac{2\pi v}{N} \left(h_L + h_R - \frac{c}{12} \right) . \quad (4.40)$$

Here ε and v are nonuniversal constants, c is the central charge of the CFT, and h_L and h_R are the conformal weights of the holomorphic and antiholomorphic part of the local field associated with the energy level. The latter parameters are defined in terms of a representation of the Virasoro algebra and are tabulated for unitary minimal CFTs. In [76], the CFT corresponding to the AFM-point was unambiguously identified as that describing the tricritical Ising model at its critical point, with central charge $c = 7/10$. The ground state energy E_0 corresponds to $h_L = h_R = 0$, whereas the first excited energy E_1 is determined by $h_L = h_R = 3/80$.

Using the exact values of E_0 and E_1 for $N = 16$, we determine the nonuniversal constants in (4.40) (approximately). The resulting prediction for the ground state energy density E_0/N for system sizes $N \leq 16$ differs by only $10^{-2}\%$ from that obtained in the same way using the exact spectrum at $N = 8$. This suggests that finite-size effects are negligible, and we use the prediction from (4.40) obtained in this fashion to study the anyonic MERA for larger systems.

Figure 4.7 shows the result of this computation. We find that the ground state energy density is well approximated by the anyonic MERA ansatz even for larger systems.

Throughout this chapter, we have considered finite-size systems. A modification of the optimization algorithm for scale-invariant qudit MERA allows to numerically estimate data of the associated CFT in the continuum limit [47–49]. Adapting this to the anyonic setting, this can provide an additional benchmark when comparing to the predictions of [76, 106]. Indeed, results of this kind were recently presented in [80] (see note below).




Figure 4.7. Ground state energy per site obtained for $N = 16, 32, 64$ and 128 anyons at the AFM-point using the anyonic MERA (blue crosses). The red line shows the CFT prediction for the ground state energy obtained by extrapolating from $N = 16$ (as discussed in the text). The dotted curves indicate the (modulus of the) deviation from this prediction when using $N = 8$ instead. The inset illustrates the relative error of the MERA with respect to the CFT prediction.




Figure 4.8. Two-chain ladder arrangement of Fibonacci anyons as discussed in [86], with nearest-neighbor interactions. We consider the ordering indicated in (b).

4.5 Braiding and More General Arrangements of Anyons

In the previous chapters, our focus has been on one-dimensional chains of anyons. Here we briefly comment on the generalization of the anyonic entanglement renormalization ansatz to systems of anyons arranged in a more general way (e.g., a regular lattice). Such systems have been studied before (see [108]), and the modifications necessary to define corresponding anyonic Hamiltonians are nicely explained in [86]. In fact, these modifications directly carry over to anyonic entanglement renormalization when a linear ordering of the anyons at every level is chosen. An analogous situation arises when considering fermionic tensor networks, and corresponding techniques [94, 95] can thus be extended to anyons.

For concreteness, we first consider the example of the two-leg ladder model discussed in [86], closely following that presentation. This is a system of two chains of m anyons each placed next to each other as shown in figure 4.8(a) along a nontrivial cycle on the

torus. The figure also indicates two-anyon nearest neighbor interactions, which are, e.g., projections onto trivial charge as in the golden chain.

The Hilbert space of this system is $V_{\text{periodic}}^{\vec{a}}$ with $\vec{a} = \underbrace{(\tau, \dots, \tau)}_{2m}$, as this space depends only on the topology of the surface (an $n = 2m$ -punctured torus) and the boundary labels (all equal to τ). However, the basis in (4.12) is now ambiguous, and we must choose a linear ordering of the anyons. A convenient choice of an ordering is shown in figure 4.8(b).

Given this ordering, some of the nearest-neighbor terms in the Hamiltonian now act on nonneighboring pairs $(i, i+2)$ of strands in the diagrammatic representation. To make sense of such terms, it is necessary to introduce an additional basis change corresponding to a transposition $(i, i+1) \mapsto (i+1, i)$. Physically, such basis changes correspond to movements of the anyons, and this is where the (nonabelian) braid group statistics appears: pairs of neighboring anyons may be exchanged in either clockwise or anticlockwise fashion. Thus there are two inequivalent ways of transposing neighboring pairs. It is natural to define the Hamiltonian as the result of transposing, applying the charge projection, and undoing the transposition, averaged over either version of exchanging.

In the diagrammatic formalism, clockwise- and anticlockwise exchanges of neighboring anyons are represented by over- and undercrossings, respectively. The following computational rules involving the R -matrix of the modular category are then added to the isotopy-invariant calculus:

$$\begin{array}{c} b \ a \\ \swarrow \ \searrow \\ \uparrow c \end{array} = R_c^{ab} \begin{array}{c} b \ a \\ \searrow \ \swarrow \\ \uparrow c \end{array} \quad , \quad \begin{array}{c} b \ a \\ \swarrow \ \searrow \\ \uparrow c \end{array} = \overline{R_{c^*}^{a^*b^*}} \begin{array}{c} b \ a \\ \searrow \ \swarrow \\ \uparrow c \end{array} .$$

With (4.2)-(4.4), these rules imply that over- and undercrossings may be resolved according to

$$\begin{array}{c} b \ a \\ \swarrow \ \searrow \\ \diagdown \ \diagup \end{array} = \sum_c \sqrt{\frac{d_c}{d_a d_b}} R_c^{ab} \begin{array}{c} b \ a \\ \swarrow \ \searrow \\ \uparrow c \\ \swarrow \ \searrow \\ b \ a \end{array} , \quad (4.41)$$

$$\begin{array}{c} b \ a \\ \swarrow \ \searrow \\ \diagup \ \diagdown \end{array} = \sum_c \sqrt{\frac{d_c}{d_a d_b}} \overline{R_{c^*}^{a^*b^*}} \begin{array}{c} b \ a \\ \swarrow \ \searrow \\ \uparrow c \\ \swarrow \ \searrow \\ b \ a \end{array} ,$$

where the sums are restricted to fusion-compatible labels c . (4.41) gives the matrix ele-

the linear maps in (4.41) to their inputs and outputs, respectively. This eventually reduces the evaluation to the expectation value of a local operator for a related anyonic MERA-state. The efficiency of this procedure depends on the number of crossings that need to be resolved, and therefore on a judicious choice of orderings. This feature is identical to fermionic tensor networks/MERAs, where it is necessary to keep track of the number of swaps.

4.6 Conclusions

We have proposed a variational ansatz for periodic chains of interacting anyons. It is the natural counterpart of entanglement renormalization for spin chains. Based on the empirical evidence for the numerical accuracy of the latter, it is reasonable to expect that the ansatz is a powerful tool for describing critical anyonic systems. Indeed, we have obtained numerical evidence for its suitability by comparing with exact diagonalization data in the case of (a variation of) the golden chain.

Our ansatz makes optimal use of the anyonic structure of the Hilbert space by incorporating conservation of topological charge at different scales. We expect this to lead to significant computational savings and improved accuracy compared to more conventional methods based on embedding the anyons into qubits. It may be more pronounced for more general anyon models than the Fibonacci anyon chains numerically studied here.

Beyond providing an efficient numerical tool, the proposed ansatz is a starting point for interesting generalizations. For example, along the lines of [39], one may define a Hamiltonian renormalization group flow based on the ansatz. This flow generalizes the perturbative renormalization prescription analytically considered in [77–79] in the context of random couplings. More importantly, the current work may serve as a stepping stone for the development of variational methods for two-dimensional systems of interacting anyons. We give a rough sketch of corresponding adaptations in section 4.5.

The transition from entanglement renormalization for qudits to anyons is remarkably simple on a conceptual level: it boils down to the replacement of isometries by topological charge-preserving isometries, and a reinterpretation of networks in terms of the isotopy-invariant calculus. This suggests that more generally, tensor network states such as MPS or PEPS may also be adapted to the anyonic setting using similar substitutions. This should

add significantly to the repertoire of variational methods for topologically ordered systems.

Chapter 5

Dimension Reduction

5.1 Introduction

Many open problems in condensed matter physics concern strongly correlated quantum many-body systems. These are typically not solvable analytically, and we have to resort to numerical simulations. We discussed several numerical methods in the previous three chapters of this thesis. These methods were successful under certain the limits and geometric restrictions. Unfortunately, there are no known numerical methods for general Hamiltonians with no restrictions, due to the exponential scaling of the dimension of the corresponding Hilbert space with system size. This problem is one of the main motivations for the quest of quantum computers. Indeed, quantum computers can efficiently simulate unitary evolutions of quantum many-body systems with local interactions [112, 113], because they can inherently deal with exponentially large Hilbert spaces.

Nevertheless, the preparation of the desired initial state on a quantum computer is still a difficult problem in general [13, 114–117]. In this chapter, we propose a new method for thermal Gibbs state preparation that achieves an exponential speedup over other known algorithms. Prior to our method, there have been several proposals to tackle this problem [118–121]. Some significant alternatives have worse complexity scaling than ours [118, 121], while others apply to a restricted set of systems [120]. The quantum metropolis algorithm [119], in particular, would often be faster, but lacks complexity bounds. For one-dimensional systems, the classical algorithm proposed in [11] can be turned into a quantum one with only a polynomial time complexity overhead with respect to our method, but it requires an exponentially higher number of bits than the number of qubits we use, and it does not extend to higher-dimensional systems.

The time complexity of our method for one-dimensional systems is dominated by the quantity $N^{\|h\|/T}$, where N is the number of subsystems, T is the temperature, and $\|h\|$ is a bound on the operator norm of the local terms of the Hamiltonian, the interaction strength. Note that this scaling is polynomial in N . The memory of the quantum computer scales simply with N , an exponential improvement over general classical algorithms. Our algorithm can also be massively parallelized, and when run in a cellular automaton architecture the memory scales as $N^{\|h\|/T}$, but the total time would be linear in N (the total number of steps would still be the same). In two and higher dimensions, our method lowers the number of effective dimensions by one. This results in an exponential speedup, but the exponential scaling with N remains.

The overall scaling appears to be optimal: the known complexity of thermalizing one-dimensional quantum systems makes a guaranteed polynomial scaling with temperature extremely unlikely [13, 117]. We also expect the grouping of $\|h\|/T$ in the exponent by dimensional analysis. In other words, the relevant temperature scale is set by the Hamiltonian.

5.2 Thermalization Using Phase Estimation

It is easier to introduce this method by explaining the proposal in [121] first. In order to prepare a thermal state of a given Hamiltonian, the probability of each eigenstate needs to be set to the correct Gibbs probability. This can be done by first rotating an ancilla conditioned on the energies of the eigenstates, then by projecting this ancilla to its original state.

The energy of any given eigenstate, $|\psi_a\rangle$, can be measured using the quantum phase estimation (QPE) procedure (followed by a quantum Fourier transform). Given a Hamiltonian $H = \sum_a E_a |\psi_a\rangle\langle\psi_a|$, the QPE circuit measures and writes down the energy of any input state $|\psi_a\rangle$ onto energy registers initialized to $|0\rangle$. That is, $\text{QPE}(H)$ is the transformation:

$$\text{QPE}(H) \equiv \sum_a |\psi_a, E_a\rangle\langle\psi_a, 0|. \quad (5.1)$$

After QPE, we can rotate an ancilla initialized to $|0\rangle$ by $\arccos(e^{-\beta E_a/2})$ using a condi-




Figure 5.1. QPE followed by a unitary rotation of an ancilla qubit by $\arccos(e^{-\beta E_a/2})$, where ψ_a is an eigenstate of the Hamiltonian H and E_a is the corresponding energy. This is the basic procedure for preparing a thermal state $\rho \propto e^{-\beta H}$.

tional rotation as in figure 5.2. The resulting state from this circuit can be written as

$$e^{-\beta E_a} |\psi_a\rangle\langle\psi_a| \otimes |E_a\rangle\langle E_a| \otimes |0\rangle\langle 0| + \dots, \quad (5.2)$$

where the ellipsis corresponds to the other terms where the rotated ancilla is not in the state $|0\rangle\langle 0|$.

Now, instead of a single eigenstate, $|\psi_a\rangle$, as input, if we use the maximally mixed state, $\mathbb{1} = \sum_a |\psi_a\rangle\langle\psi_a|$, the resulting density matrix after the circuit becomes

$$\sum_a e^{-\beta E_a} |\psi_a\rangle\langle\psi_a| \otimes |E_a\rangle\langle E_a| \otimes |0\rangle\langle 0| + \dots. \quad (5.3)$$

Projecting the ancilla to state $|0\rangle$ results in the density matrix:

$$(1/Z) \sum_a e^{-\beta E_a} |\psi_a\rangle\langle\psi_a| \otimes |E_a\rangle\langle E_a| \otimes |0\rangle\langle 0|, \quad (5.4)$$

where Z is the partition function. At this point, we can simply trace out the energy registers and the ancilla to obtain the thermalized state, $\rho = (1/Z) \sum_a e^{-\beta E_a} |\psi_a\rangle\langle\psi_a|$.

The projection discussed above can be done by measuring the ancilla in the $|0\rangle$ - $|1\rangle$ basis. If the measurement outcome is the state $|0\rangle$, we have successfully performed the projection, and we fail if we obtain the state $|1\rangle$. Before the projection, all eigenstates of the system, $|\psi_a\rangle$, are equally likely. The unitary rotation conditioned on the energy registers ensures that the probability of obtaining the outcome $|0\rangle$ is higher for lower-energy states. Because of the precise rotation we chose, the probabilities of the eigenstates are updated to their corresponding Gibbs weights conditioned on the successful projection.

The probability of a successful projection depends on the eigenvalues of the Hamiltonian, H . Precisely, this probability is Z/d^N , where d is the dimension of each local subsystem.

If the projective measurement fails, i.e., we obtain the outcome $|1\rangle$, we need to restart the whole procedure. As a result, the number of trials scales exponentially with the system size. It is possible to obtain a quadratic speedup over this scaling using Grover's amplitude amplification [121], but the algorithm still scales exponentially with system size.

5.2.1 Dimension Reduction Overview

We overcome the problem of exponential time cost by dividing the overall procedure into a sequence of projections and arranging them so that we only need to rebuild a small section after most failures (see figure 5.2). We first thermalize small regions independently. Each of these thermalization procedures require some projective measurements to be successful. When these measurements fail, they only affect the state of the associated small region. This forces a restart of the thermalization procedure of that region alone, which is not as costly as the thermalization of the whole system. Once we have thermalized small regions, we recursively merge them until we have the fully thermalized system as in figure 5.2.

There are very few projective measurements that require a restart of thermalization of large portions of the system in this procedure. Most failed projective measurements destroy the state of very small regions. This causes the running time to be only polynomial in system size for one-dimensional systems. The method also trivially generalizes to higher dimensions and reduces the scaling of the cost with the system dimension by one compared to a direct projection.




Figure 5.2. The procedure to thermalize an 8-qubit chain. After thermalizing individual qubits at level $k = 0$, we pair them up and merge them by adding the Hamiltonian that connects the two qubits. This procedure is then repeated recursively as we merge two already thermalized regions of size 2^k at level k to obtain a thermalized chain of size 2^{k+1} at level $k + 1$.

5.3 Perturbative Hamiltonian Update

The key technical ingredient for our method is a general procedure that, given sufficient copies to a thermal state $\rho^{(0)} \propto e^{-\beta H}$ (from previous steps), constructs the state $\rho^{(1)} \propto e^{-\beta(H+h)}$ with high fidelity. This is depicted in figure 5.3, where Hamiltonian H corresponds to the halves to be merged and h to the link between the two halves, but the procedure is more general. This method will use controlled evolutions with the Hamiltonians H and $H + h$. The final cost, quantified by the total evolution required, will be shown to be $\mathcal{O}(e^{\beta\|h\|})$. This is a probabilistic method and, when it fails, we need to restart with a new copy of $\rho^{-\beta H}$. The number of initial copies required also scales like $\mathcal{O}(e^{\beta\|h\|})$.




Figure 5.3. Merging separate thermal states is done with the general method that constructs $e^{-\beta(H+h)}$ from copies of $e^{-\beta H}$.

We implement the merging perturbatively. We will see how to generate (with high probability) the state $\rho^{(\epsilon)} \propto e^{-\beta(H+\epsilon h)}$ from sufficient copies of $\rho^{(0)}$. We then repeat the same process to produce the sequence

$$\rho = \rho^{(0)} \rightarrow \rho^{(\epsilon)} \rightarrow \rho^{(2\epsilon)} \rightarrow \dots \rightarrow \rho^{(1)}. \quad (5.5)$$

Every transformation in the sequence has some probability of failure, in which case we restart. If all of the steps succeed, we approximate the state $\rho^{(1)}$ with an error of $\mathcal{O}(\epsilon)$. It is important to remark that all the errors in this chapter are in the trace norm. That is, in this case, for input ρ , we build a state σ such that $\|\sigma - \rho^{(1)}\|_{\text{Tr}} \in \mathcal{O}(\epsilon)$.

We update the state $\rho \propto e^{-\beta H}$ to $\rho^{(\epsilon)} \propto e^{-\beta(H+\epsilon h)}$, to first order in ϵ , in two steps. The first step is probabilistic and updates the probabilities of the Gibbs state through postselection. If it fails, we will be forced to restart, and this will be the dominant part for the cost of the algorithm. In the second step we update the eigenbasis to the eigenbasis of $\rho^{(\epsilon)}$. We now first perfect phase estimation and perfect dephasing, and later account for the

cost and errors of these operations.

We first derive a perturbative expression for $\rho^{(\epsilon)}$, which holds also for the degenerate parts of the spectrum. For this, we introduce a parameter λ , formally just an inverse temperature and use Dyson series in imaginary time.

Lemma 1. *Let ρ denote the Gibbs state of H at temperature β , and $\rho^{(\epsilon)}$ the Gibbs of $H + \epsilon h$ at the same temperature. Define $J = \|h\|$ to be the coupling strength. Then,*

$$\rho^{(\epsilon)} = \left(\rho - \epsilon\beta \sum_{lr} \frac{p_l - p_r}{\beta(E_r - E_l)} h_{lr} \right) (1 - \epsilon\beta \text{Tr}(\rho h))^{-1} .$$

with an error $\mathcal{O}((\epsilon\beta J)^2)$ in the trace norm. We define $(p_l - p_r)/(\beta(E_r - E_l)) = p_l$ when $E_l = E_r$ (this also makes the function continuous). This expression is valid when there are degeneracies.

Proof. We derive a perturbative expansion for $\rho^{(\epsilon)}$, which holds also for degenerate parts of the spectrum. For this, we introduce a parameter λ , formally just an inverse temperature and use Dyson series expansion in imaginary time. We write

$$\partial_\lambda e^{-\beta(H+\epsilon h)\lambda} = -\beta(H + \epsilon h)e^{-\beta(H+\epsilon h)\lambda} , \quad (5.6)$$

which has the solution

$$\begin{aligned} (1/Z)e^{-\beta(H+\epsilon h)\lambda} &= (1/Z)e^{-\beta H\lambda} - \epsilon\beta/Z \int_0^\lambda d\lambda_1 e^{-\beta H(\lambda-\lambda_1)} h e^{-\beta(H+\epsilon h)\lambda_1} \\ &= (1/Z)e^{-\beta H\lambda} - \epsilon\beta/Z \int_0^\lambda d\lambda_1 e^{-\beta H(\lambda-\lambda_1)} h e^{-\beta H\lambda_1} \\ &\quad + (\epsilon\beta)^2/Z \int_0^\lambda d\lambda_1 \int_0^{\lambda_1} d\lambda_2 e^{-\beta H(\lambda-\lambda_1)} h e^{-\beta H(\lambda_1-\lambda_2)} h e^{-\beta(H+\epsilon h)\lambda_2} . \end{aligned} \quad (5.7)$$

Note that this expansion is well behaved. Theorem 5.4.7 in [122] states that for all unitarily invariant norms,

$$\left\| \int_0^1 A^t X B^{1-t} dt \right\| \leq 1/2 \|AX + XB\| . \quad (5.9)$$

Using this, we can bound the trace-norm of the first order term in the expansion directly,

and get

$$(1/Z) \left\| \epsilon\beta \int_0^1 d\lambda_1 e^{-\beta H(1-\lambda_1)} h e^{-\beta H \lambda_1} \right\|_{\text{Tr}} \leq (1/Z) \epsilon\beta/2 \|e^{-\beta H} h + h e^{-\beta H}\|_{\text{Tr}} \quad (5.10)$$

$$\leq \epsilon\beta \|h\| \|e^{-\beta H}\|_{\text{Tr}}/Z \leq \epsilon\beta \|h\|, \quad (5.11)$$

where $\|h\|$ is the operator norm of h .

In fact, the sum of all perturbation terms (first-order and higher) can be bounded by $\epsilon\beta\|h\|$ by applying (5.9) to (5.7), and using the bound $\text{Tr}(e^{-\beta(H+\epsilon h)}) = \text{Tr}(e^{-\beta H})(1 + \mathcal{O}(\epsilon))$ from Ref. [121]. That is,

$$\frac{1}{Z} \left\| \epsilon\beta \int_0^1 d\lambda_1 e^{-\beta H(1-\lambda_1)} h e^{-\beta(H+\epsilon h)\lambda_1} \right\|_{\text{Tr}} \leq \frac{1}{Z} \epsilon\beta/2 \|e^{-\beta H} h + h e^{-\beta(H+\epsilon h)}\|_{\text{Tr}} \quad (5.12)$$

$$\leq \epsilon\beta/2 \|h\| (\|e^{-\beta H}\|_{\text{Tr}} + \|e^{-\beta(H+\epsilon h)}\|_{\text{Tr}})/Z \quad (5.13)$$

$$\leq \epsilon\beta \|h\| (1 + \mathcal{O}(\epsilon\beta\|h\|)). \quad (5.14)$$

For sufficiently small $\epsilon\beta\|h\|$, the first two terms of the Dyson expansion are the major contribution, and we can use the bound on the whole series to write

$$\|\epsilon\beta\|h\| A_1(\tilde{H}, \tilde{h}) + (\epsilon\beta\|h\|)^2 A_2(\tilde{H}, \tilde{h})\|_{\text{Tr}} \leq \kappa(\epsilon\beta\|h\|), \quad (5.15)$$

where κ is a small constant, $\tilde{h} = h/\|h\|$ and $\tilde{H} = \beta H$ are dimensionless matrices, and

$$A_1(\tilde{H}, \tilde{h}) = \frac{1}{Z} \int_0^1 d\lambda_1 e^{-\tilde{H}(1-\lambda_1)} \tilde{h} e^{-\tilde{H}\lambda_1} \quad (5.16)$$

$$A_2(\tilde{H}, \tilde{h}) = \frac{1}{Z} \int_0^1 d\lambda_1 \int_0^{\lambda_1} d\lambda_2 e^{-\tilde{H}(1-\lambda_1)} \tilde{h} e^{-\tilde{H}(\lambda_1-\lambda_2)} \tilde{h} e^{-\tilde{H}\lambda_2}. \quad (5.17)$$

We get

$$\epsilon\beta\|h\| \|A_2(\tilde{H}, \tilde{h})\|_{\text{Tr}} \leq \kappa + \|A_1(\tilde{H}, \tilde{h})\|_{\text{Tr}} \quad (5.18)$$

But we have already shown that $\|A_1(\tilde{H}, \tilde{h})\|_{\text{Tr}} \leq 1$, and so

$$\epsilon\beta\|h\| \|A_2(\tilde{H}, \tilde{h})\|_{\text{Tr}} \leq \kappa + 1. \quad (5.19)$$

Because this bound holds for all \tilde{H} and \tilde{h} , we get

$$\|A_2(\tilde{H}, \tilde{h})\|_{\text{Tr}} \leq \kappa_2 . \quad (5.20)$$

That is, if such a bound didn't exist (for small, but fixed, $\epsilon\beta\|h\|$), then the expression on the left could not be bounded by the constant $\kappa + 1$. Therefore, we get a bound for the second order contribution

$$(\epsilon\beta\|h\|)^2 \frac{1}{Z} \left\| \int_0^1 d\lambda_1 \int_0^{\lambda_1} d\lambda_2 e^{-\tilde{H}(1-\lambda_1)} \tilde{h} e^{-\tilde{H}(\lambda_1-\lambda_2)} \tilde{h} e^{-\tilde{H}\lambda_2} \right\|_{\text{Tr}} = \mathcal{O}((\epsilon\beta\|h\|)^2) . \quad (5.21)$$

We now rewrite the first order perturbation in terms of projectors onto the eigenstates of ρ .

$$\sum_{l,r} \int_0^1 e^{-\beta E_l(1-\lambda_1) - \beta E_r \lambda_1} d\lambda_1 P_l h P_r / Z = \sum_{l \neq r} \frac{p_l - p_r}{\beta(E_r - E_l)} P_l h P_r + \sum_k p_k P_k h P_k , \quad (5.22)$$

where $p_l = e^{-\beta E_l} / Z$. Notice that $\sum_l p_l = 1$. Also notice that if $E_l = E_r + \delta$, then

$$\frac{p_l - p_r}{\beta(E_r - E_l)} = p_l(1 + \mathcal{O}(\beta\delta)) . \quad (5.23)$$

That is, if the energies are sufficiently close, then their probabilities are similar, canceling the gap dependence. Furthermore, writing $p = \max\{p_l, p_r\}$ and $\delta = |E_l - E_r|$ we have

$$\frac{p_l - p_r}{\beta(E_r - E_l)} = p \frac{1 - e^{-\beta\delta}}{\beta\delta} \leq p . \quad (5.24)$$

We have so far an unnormalized first order approximation:

$$\rho^{(\epsilon)} \propto \rho - \epsilon\beta \left(\sum_{l \neq r} \frac{p_l - p_r}{\beta(E_r - E_l)} P_l h P_r + \sum_k p_k P_k h P_k \right) \cdots . \quad (5.25)$$

To simplify the notation, we will just write

$$\rho^{(\epsilon)} \propto \rho - \epsilon\beta \sum_{l,r} \frac{p_l - p_r}{\beta(E_r - E_l)} P_l h P_r , \quad (5.26)$$

with the understanding that the fraction goes to p_l when $r = l$.

□

At this point it is useful to introduce some notation. We are interested in approximation in the trace norm. This is justified because the trace norm bounds the probability of error, over all measurements, when distinguishing two quantum states. With that in mind, we denote $\|\sigma_1 - \sigma_2\|_{\text{Tr}} \in \mathcal{O}(f)$, for some function f , by $\sigma_1 = \sigma_2 + \mathcal{O}(f)$. We also say that σ_1 is σ_2 up to an error of order f .

To calculate the normalization, we take the trace of the expression on the previous lemma to get

$$\text{Tr} \left(\rho - \epsilon\beta \frac{p_l - p_r}{\beta(E_r - E_l)} P_l h P_r \right) = 1 - \epsilon\beta \text{Tr} \left(\sum_l p_l P_l h P_l \right) = 1 - \epsilon\beta \text{Tr} \rho h . \quad (5.27)$$

And then

$$\rho^{(\epsilon)} = \left(\rho - \epsilon\beta \frac{p_l - p_r}{\beta(E_r - E_l)} P_l h P_r \right) / (1 - \epsilon\beta \text{Tr} \rho h) + \mathcal{O}((\epsilon\beta J)^2) . \quad (5.28)$$

5.3.1 Perturbative Update with Perfect Operations

We now give an algorithm that, assuming perfect phase estimation and perfect dephasing, performs the update of the state $\rho \propto e^{-\beta H}$, to $\rho^{(\epsilon)} \propto e^{-\beta(H+\epsilon h)}$ to first order in $\epsilon\beta\|h\|$. We present this subroutine in two steps. In the first step we implement, probabilistically, a transformation, which is the exponential of $\epsilon\beta h$, to first order. We assume that $h \geq 0$. This step is probabilistic and updates the probabilities through postselection: if it succeeds, we are projecting on the correct states with the correct probability. If it fails, we will be forced to restart, and this will be the dominant part for the cost of the algorithm, which we account for in the next section. In the second step we dephase in the eigenbasis of $\rho^{(\epsilon)}$, which corresponds to updating the eigenbasis.

The intuitive meaning of the next subroutine is that it implements, probabilistically, an approximate conjugation $\sigma \rightarrow e^{-\epsilon\beta h/2} \sigma e^{-\epsilon\beta h/2}$.

Lemma 2. *Assume perfect phase estimation of $e^{2\pi i \beta h t}$ with $1/t > \beta h \geq 0$. Then we can implement the map*

$$\mathcal{E}_{\beta h, \epsilon}(\sigma) \propto (1 - \epsilon\beta h/2) \sigma (1 - \epsilon\beta h/2) ,$$

with probability $1 - \epsilon\beta \text{Tr} \sigma h$. The cost is that of two phase estimation calls.

Proof. We use phase estimation and postselection. The phase estimation with $e^{2\pi i\beta ht}$ gives the map

$$\sum_a P_a |E_a\rangle\langle 0|, \quad (5.29)$$

where P_a is a projection onto the eigenspace of h with energy E_a . The correction for t and β on the phase estimation is trivial. This energy gets written in an ancilla register, which is initialized to $|0\rangle$. Now rotate a second ancilla to $(1 - \epsilon\beta E_a/2)|0\rangle + \dots |1\rangle$, conditional on the value of the previous ancilla register, to obtain

$$\sum_a (1 - \epsilon\beta E_a/2) P_a |E_a 0\rangle\langle 00| + \dots. \quad (5.30)$$

We now undo the phase estimation to obtain

$$\sum_a (1 - \epsilon\beta E_a/2) P_a |00\rangle\langle 00| + \dots = (1 - \epsilon\beta h/2) |00\rangle\langle 00| + \dots. \quad (5.31)$$

Finally we measure the ancillae, and fail unless we obtain 0 for the second ancilla. \square

Let us consider what is the first order approximation to the updated probabilities of $\rho^{(\epsilon)}$. Denote the first order correction to the eigenvalues by $-\epsilon\beta E_k^{(1)} = -\epsilon\beta \text{Tr} P_k h$. Then the corrected probabilities are, up to normalization,

$$p_k e^{-\epsilon\beta E_k^{(1)}} \approx p_k (1 - \epsilon\beta \text{Tr} P_k h). \quad (5.32)$$

These are the probabilities of the result of the operation of the previous lemma with input ρ , which we denote by ρ_{prob} ,

$$\rho_{\text{prob}} = \mathcal{E}_{h,\epsilon}(\rho) \propto \rho - \epsilon\beta/2(h\rho + \rho h). \quad (5.33)$$

In the previous expression, the error is bounded, in the trace norm, by $\mathcal{O}((\epsilon\beta\|h\|)^2)$. That is, the map of Lemma 2 updates the probability correctly to first order.

We also need to change the eigenbasis of ρ , and not only the probabilities. The standard procedure to evolve from the eigenbasis of some operator to the eigenbasis of a related operator is the use of the adiabatic approximation. The adiabatic approximation can also

be seen as a consequence of the Zeno effect, which can be achieved through measurements or dephasing [123, 124]. Here we implement the Zeno effect directly, or, more specifically, we assume access to pure dephasing in the eigenbasis of $H + \epsilon h$ (see [124]). Notice that this could be implemented if we had access to perfect phase estimation.

Lemma 3. *Assume access to a perfect dephasing operation in the eigenbasis of $H + \epsilon h$. Then we can transform ρ_{prob} to $\rho^{(\epsilon)}$ with an error $\mathcal{O}((\epsilon\beta J)^2)$ in the trace norm, where $J = \|h\|$. The cost is that of one dephasing operation.*

Proof. For this proof, it is more convenient to work in the eigenbasis of the new Hamiltonian, $H + \epsilon h$. Denote by $\{P_k^{(\epsilon)}\}$ the projectors on the eigenstates of $H + \epsilon h$, and define $\{p_k^{(\epsilon)}\}$ and $\{E_k^{(\epsilon)}\}$ to be the corresponding probabilities and energies in $\rho^{(\epsilon)}$, such that $\rho^{(\epsilon)} = \sum_k p_k^{(\epsilon)} P_k^{(\epsilon)}$. Using these projectors and the first order approximation above we can write

$$\rho = \rho^{(\epsilon)} + \epsilon\beta \sum_{lr} \frac{(p_l^{(\epsilon)} - p_r^{(\epsilon)})P_l^{(\epsilon)}hP_r^{(\epsilon)}}{\beta(E_r^{(\epsilon)} - E_l^{(\epsilon)})} + \mathcal{O}((\epsilon\beta J)^2). \quad (5.34)$$

After the phase estimation subroutine, we obtained $\rho_{\text{prob}} = (1 - \epsilon\beta h/2)\rho(1 - \epsilon\beta h/2)$. Rewriting this in the new eigenbasis, we have

$$\begin{aligned} \rho_{\text{prob}} &= \rho^{(\epsilon)} - \sum_k p_k^{(\epsilon)} \epsilon\beta h P_k^{(\epsilon)} / 2 - \sum_k p_k^{(\epsilon)} \epsilon\beta P_k^{(\epsilon)} h / 2 \\ &+ \epsilon\beta \sum_{lr} \frac{(p_l^{(\epsilon)} - p_r^{(\epsilon)})P_l^{(\epsilon)}hP_r^{(\epsilon)}}{\beta(E_r^{(\epsilon)} - E_l^{(\epsilon)})} \\ &- \frac{\epsilon^2\beta^2}{2} \sum_{lr} \frac{(p_l^{(\epsilon)} - p_r^{(\epsilon)})hP_l^{(\epsilon)}hP_r^{(\epsilon)}}{\beta(E_r^{(\epsilon)} - E_l^{(\epsilon)})} \\ &- \frac{\epsilon^2\beta^2}{2} \sum_{lr} \frac{(p_l^{(\epsilon)} - p_r^{(\epsilon)})P_l^{(\epsilon)}hP_r^{(\epsilon)}h}{\beta(E_r^{(\epsilon)} - E_l^{(\epsilon)})} + \mathcal{O}((\epsilon\beta J)^2). \end{aligned} \quad (5.35)$$

Dephasing ρ_{prob} in the new basis, i.e., applying the map $\sum_{k,l:E_k=E_l} P_k^{(\epsilon)} \cdot P_l^{(\epsilon)}$, we get

$$\sum_{k,l:E_k=E_l} P_k^{(\epsilon)} \rho_{\text{prob}} P_l^{(\epsilon)} = \sum_k p_k^{(\epsilon)} P_k^{(\epsilon)} \quad (5.36)$$

$$- \sum_{k,l:E_k=E_l} \frac{\epsilon\beta}{2} p_k^{(\epsilon)} P_k^{(\epsilon)} h P_l^{(\epsilon)} - \sum_{k,l:E_k=E_l} \frac{\epsilon\beta}{2} p_k^{(\epsilon)} P_k^{(\epsilon)} h P_l^{(\epsilon)} \quad (5.37)$$

$$+ \epsilon\beta \sum_{k,l:E_k=E_l} p_k^{(\epsilon)} P_k^{(\epsilon)} h P_l^{(\epsilon)} \quad (5.38)$$

$$- \frac{\epsilon^2 \beta^2}{2} \sum_{klm:E_k=E_m} \frac{(p_l^{(\epsilon)} - p_k^{(\epsilon)}) P_k^{(\epsilon)} h P_l^{(\epsilon)} h P_m^{(\epsilon)}}{\beta(E_k^{(\epsilon)} - E_l^{(\epsilon)})} \quad (5.39)$$

$$- \frac{\epsilon^2 \beta^2}{2} \sum_{klm:E_k=E_m} \frac{(p_k^{(\epsilon)} - p_l^{(\epsilon)}) P_k^{(\epsilon)} h P_l^{(\epsilon)} h P_m^{(\epsilon)}}{\beta(E_l^{(\epsilon)} - E_k^{(\epsilon)})} \quad (5.40)$$

$$+ \mathcal{O}((\epsilon\beta J)^2). \quad (5.41)$$

Simplifying the above expression, we get

$$\sum_k P_k^{(\epsilon)} \rho_{\text{prob}} P_k^{(\epsilon)} = \sum_k p_k^{(\epsilon)} P_k^{(\epsilon)} \quad (5.42)$$

$$- \epsilon^2 \beta^2 \sum_{klm:E_k=E_m} \frac{(p_l^{(\epsilon)} - p_k^{(\epsilon)}) P_k^{(\epsilon)} h P_l^{(\epsilon)} h P_m^{(\epsilon)}}{\beta(E_k^{(\epsilon)} - E_l^{(\epsilon)})}. \quad (5.43)$$

The $\mathcal{O}(\epsilon^2)$ term can be rewritten as

$$\epsilon^2 \beta^2 / 2 (1/Z) \sum_{kl:E_k=E_l} P_k \left(h \int_0^1 d\lambda_1 e^{-\beta H(1-\lambda_1)} h e^{-\beta H \lambda_1} \right. \quad (5.44)$$

$$\left. + \int_0^1 d\lambda_1 e^{-\beta H(1-\lambda_1)} h e^{-\beta H \lambda_1} h \right) P_l. \quad (5.45)$$

The norm of the expression before dephasing can be bounded using 5.11 as

$$(1/Z)(1/2) \left\| h \int_0^1 d\lambda_1 e^{-\beta H(1-\lambda_1)} h e^{-\beta H \lambda_1} \right\| \quad (5.46)$$

$$+ \left\| \int_0^1 d\lambda_1 e^{-\beta H(1-\lambda_1)} h e^{-\beta H \lambda_1} h \right\|_{\text{Tr}} \quad (5.47)$$

$$\leq (1/Z)(1/2) \|h\| \left\| \int_0^1 d\lambda_1 e^{-\beta H(1-\lambda_1)} h e^{-\beta H \lambda_1} \right\|_{\text{Tr}} \leq \|h\|^2. \quad (5.48)$$

Since dephasing does not increase the trace-norm (dephasing can also be implemented using phase randomization), we can bound the trace-norm of the higher-order terms by

$\mathcal{O}((\epsilon\beta J)^2)$.

□

5.3.2 Concatenation

Putting it all together, and concatenating, we obtain the following corollary.

Lemma 4. *Assume access to perfect phase estimation and dephasing operations, and to copies of $\rho \propto e^{-\beta H}$. Then we can transform $e^{-\beta H}/\text{Tre}^{-\beta H}$ to $e^{-\beta(H+h)}/\text{Tre}^{-\beta(H+h)}$. For a target error in the trace norm $\mathcal{O}(\epsilon\beta^2\|h\|^2)$, let p denote a lower bound for $1-\epsilon\beta\text{Tr}\rho^{m\epsilon}h$, where $\rho^{m\epsilon} \propto e^{-\beta(H+mh)}$. An obvious choice is $p = 1-\epsilon\beta\|h\|$. Define $s = \lceil 1/\epsilon \rceil$. The number of operations used is a random variable with mean μ bounded by*

$$\mu \leq \frac{1}{(1-p)p^s} \in \mathcal{O}(p^{-s}).$$

The probability that we are not successful after n trials can be bounded by $e^{-p^s n/s}/(1-p^s)$ (that is, the tail has an exponential decay rate). Note that $p^{-s} \in e^{\beta\|h\|}(1 + \mathcal{O}(\epsilon\beta^2\|h\|^2))$. We can use the same bounds for the number of initial states ρ used.

Proof. Fix ϵ according to the target error in the trace norm, $\mathcal{O}(\epsilon\beta^2\|h\|^2)$. We concatenate the procedures of Lemmas 2 and 3. We prepare an approximation to the target state if, starting with ρ , we obtain a *success run* of length s , that is, all consecutive s concatenations of operations are successful. The final error, in this case, is bounded as $\mathcal{O}(\epsilon\beta^2\|h\|^2)$ by the errors of Lemmas 2 and 3 and the triangle inequality of the trace norm.

The number of operations is random. We need a success run of size s , and we restart with the state ρ after every failure. The distribution of the number of operations is known in the theory of success runs. For the mean μ we refer to the literature (see, for instance, [125] Eq. (7.7)). The tail probability is easy to bound using the worst case bound p [126]. For a sequence of random variables with Bernoulli distribution of probability p , define the random variable Z_j to denote the a success run of length s starting at j , that is, all the s trails starting at j are successful. Denote by $F(n)$ the probability that there is success run in the first n trails. We have

$$F(n) = P(Z_j = 0, j = 1, 2, \dots, n - s + 1) \leq P(Z_1 = 0, Z_{s+1} = 0, \dots, Z_{\lfloor n/s \rfloor} = 0). \quad (5.49)$$

The random variables $Z_1, Z_{s+1}, \dots, Z_{\lfloor n/s \rfloor}$ are independent, so

$$F(n) \leq \prod_{k=0}^{\lfloor n/s \rfloor - 1} P(Z_{ks+1}) = (1 - p^s)^{\lfloor n/s \rfloor} \leq \frac{e^{-p^s n/s}}{1 - p^s}. \quad (5.50)$$

□

5.3.3 Errors and Cost of Operations with Finite Precision

In the previous section we were considering the behavior of our algorithms assuming perfect phase estimation, which is not realistic. We now assume access to the Hamiltonian H and account for the effects of the errors inherent in the phase estimation algorithm.

Lemma 5. *If we can perform controlled evolutions with h , we can implement the map*

$$\mathcal{E}_{h,\epsilon}(\sigma) \propto (1 - \epsilon\beta h/2)\sigma(1 - \epsilon\beta h/2) \quad (5.51)$$

with probability $1 - \epsilon\beta \text{Tr}\sigma h$ with error $\mathcal{O}(\epsilon^2\beta^2\|h\|^2)$ and cost (evolution time) $\mathcal{O}(\epsilon^{-1}\beta^{-1}\|h\|^{-2}\log(1/(\epsilon\beta\|h\|)))$.

Proof. We use high precision phase estimation [121, 127–129]. The cost (evolution time with h), for precision δ and error ϵ , scales as $\mathcal{O}(\log(1/\epsilon)/\delta)$. We write, as in Lemma 2, $h = \sum E_a P_a$. High precision phase estimation implements the transformation

$$\sum_a P_a |0\rangle \rightarrow \sum_a P_a \left(\sum_{\pm} c_a^{\pm} |E_a^{\pm}\rangle + q_a |\xi_a\rangle \right) \quad (5.52)$$

with $|E_a - E_a^{\pm}| \leq \delta$ and $q_a \leq \epsilon$. We will follow the same steps as in Lemma 2.

First consider the effect of the error term $\sum_a q_a P_a |\xi_a\rangle$. All the manipulations conserve the projectors P_a and do not increase the norm of $|\xi_a\rangle$, so the final error due to these terms, on input σ , can be bounded with terms like $\|\sum_a q_a P_a \sigma\|_{\text{tr}} \leq \epsilon$. The final error can increase as a result of the normalization when projecting onto the post-selected state, if the preparation is successful. We will see that this effect is negligible.

For the other term we get, before undoing the phase estimation

$$\begin{aligned} & \sum_{a,\pm} c_a^\pm (1 - \epsilon\beta E_a^\pm/2) P_a |E_a^\pm 0\rangle + \dots \\ &= \sum_{a,\pm} (1 - \epsilon\beta E_a/2) c_a^\pm P_a |E_a^\pm 0\rangle + \sum_{a,\pm} \epsilon\beta/2 (E_a - E_a^\pm) c_a^\pm P_a |E_a^\pm 0\rangle + \dots . \end{aligned} \quad (5.53)$$

The last term is again an error, which can be bounded by $\mathcal{O}(\epsilon\beta\delta)$ in the trace norm.

We undo the phase estimation on $\sum_{a,\pm} (1 - \epsilon\beta E_a/2) c_a^\pm P_a |E_a^\pm 0\rangle$. This gives

$$\sum_a (1 - \epsilon\beta E_a/2) P_a |00\rangle + \dots = (1 - \epsilon\beta h/2) |00\rangle , \quad (5.54)$$

with corrections in $\mathcal{O}(\epsilon)$, which we are not writing anymore. Finally we project to the same as in Lemma 2. We see that the probability that the projection is successful (and the normalization) is $1 - \epsilon\beta \text{Tr}\sigma h + \mathcal{O}(\epsilon^2\beta^2\|h\|^2 + \epsilon + \epsilon\beta\delta)$. If we choose $\epsilon, \epsilon\beta\delta \in \mathcal{O}(\epsilon^2\beta^2\|h\|^2)$, we obtain the lemma. \square

Finally, we deal with the errors related to imperfect dephasing. We can try to give some intuition behind this method. There are two main cases: similar or dissimilar energies. For similar energies, the previous map already does the proper update, including the update of the eigenbasis. For dissimilar energies, the previous map updates the probabilities correctly, to first order, but not the eigenbasis. If we now perform an imperfect phase estimation, the problem will be corrected as long as the energies are dissimilar enough.

Theorem 1. *If we can perform controlled evolutions with h , and evolutions with $H + \epsilon h$, we can transform ρ to $\rho^{(\epsilon)}$ with probability $1 - \epsilon\beta \text{Tr}\rho h$, error $\mathcal{O}(\epsilon^2\beta^2\|h\|^2)$, and cost (evolution time) $\mathcal{O}(\log(1/(\epsilon\beta\|h\|)))/(\epsilon\|h\|)$.*

Proof. We can implement dephasing using high-precision QPE with accuracy δ and precision ϵ in time $\mathcal{O}(\log(1/\epsilon)/\delta)$. We measure the energy using the Hamiltonian $H + \epsilon h$ and forget the result. Using high precision phase estimation, as in Eq. (5.52), we get

$$\sigma \rightarrow \sum_{ab} P_a^{(\epsilon)} \sigma P_b^{(\epsilon)} \left(\sum_{\pm} c_a^\pm (c_b^\pm)^* |E_a^{(\epsilon)\pm}\rangle \langle E_b^{(\epsilon)\pm}| \right) + \mathcal{O}(\epsilon) . \quad (5.55)$$

The states $|E_a^{(\epsilon)\pm}\rangle$ are binary discretizations of the energy, as is standard on phase estimation. If we now forget the result of the phase estimation, which amounts to a partial trace

of the ancillae, we obtain

$$\sigma \rightarrow \sum_a P_a^{(\epsilon)} c_a^+ \sigma \left(\sum_{b: E_a^{(\epsilon)+} = E_b^{(\epsilon)+}} P_b^{(\epsilon)} (c_b^+)^* + \sum_{b: E_a^{(\epsilon)+} = E_b^{(\epsilon)-}} P_b^{(\epsilon)} (c_b^-)^* \right) \quad (5.56)$$

$$+ \sum_a P_a^{(\epsilon)} c_a^- \sigma \left(\sum_{b: E_a^{(\epsilon)-} = E_b^{(\epsilon)+}} P_b^{(\epsilon)} (c_b^+)^* + \sum_{b: E_a^{(\epsilon)-} = E_b^{(\epsilon)-}} P_b^{(\epsilon)} (c_b^-)^* \right) + \mathcal{O}(\varepsilon) \quad (5.57)$$

$$(5.58)$$

$$= \sum_{a,b: |E_a^{(\epsilon)} - E_b^{(\epsilon)}| \leq 2\delta} w_{ab} P_a^{(\epsilon)} \sigma P_b^{(\epsilon)} + \mathcal{O}(\varepsilon), \quad (5.59)$$

where

$$w_{ab} = c_a^+ (c_b^+)^* \delta_{E_a^{(\epsilon)+}, E_b^{(\epsilon)+}} + c_a^+ (c_b^-)^* \delta_{E_a^{(\epsilon)+}, E_b^{(\epsilon)-}} + c_a^- (c_b^+)^* \delta_{E_a^{(\epsilon)-}, E_b^{(\epsilon)+}} + c_a^- (c_b^-)^* \delta_{E_a^{(\epsilon)-}, E_b^{(\epsilon)-}}. \quad (5.60)$$

Note that $|w_{ab}| \leq 1$ and $w_{aa} = 1 - \mathcal{O}(\varepsilon)$. From now on we drop the $\mathcal{O}(\varepsilon)$ corrections for simplicity of notation: they are made sufficiently small with respect to other errors with only a logarithmic overhead.

We now apply the above map to ρ_{prob} from (5.35) to obtain

$$\begin{aligned} \sum_{a,b: |E_a^{(\epsilon)} - E_b^{(\epsilon)}| \leq 2\delta} w_{ab} P_a^{(\epsilon)} \rho_{\text{prob}} P_b^{(\epsilon)} &= \sum_a w_{aa} P_a^{(\epsilon)} p_a \\ &- \epsilon \beta / 2 \sum_{a,b: |E_a^{(\epsilon)} - E_b^{(\epsilon)}| \leq 2\delta} w_{ab} (p_a^{(\epsilon)} + p_b^{(\epsilon)}) P_a^{(\epsilon)} h P_b^{(\epsilon)} \\ &+ \epsilon \beta \sum_{a,b: |E_a^{(\epsilon)} - E_b^{(\epsilon)}| \leq 2\delta} w_{ab} \frac{(p_a^{(\epsilon)} - p_b^{(\epsilon)})}{\beta(E_b^{(\epsilon)} - E_a^{(\epsilon)})} P_a^{(\epsilon)} h P_b^{(\epsilon)}. \end{aligned} \quad (5.61)$$

Now, we rewrite $\frac{(p_a^{(\epsilon)} - p_b^{(\epsilon)})}{\beta(E_b^{(\epsilon)} - E_a^{(\epsilon)})} = \frac{p(e^{\beta \Delta_{ab}} - 1)}{\beta \Delta_{ab}} = p(1 + \beta \Delta_{ab}/2 + \mathcal{O}(\beta^2 \Delta_{ab}^2))$, where $p = \min(p_a^{(\epsilon)}, p_b^{(\epsilon)})$ and $\Delta_{ab} = |E_a^{(\epsilon)} - E_b^{(\epsilon)}|$. From here on, we ignore the $\mathcal{O}(\beta^2 \Delta_{ab}^2)$ terms, but they can be treated in the same way as the $\mathcal{O}(\beta \Delta_{ab})$ terms to give a final $\mathcal{O}(\epsilon \beta^3 \|h\| \delta^2)$

correction. Using this, we can rewrite (5.61) as

$$\begin{aligned}
& \sum_a w_{aa} P_a^{(\epsilon)} p_a - \epsilon\beta/2 \sum_{a,b:|E_a^{(\epsilon)}-E_b^{(\epsilon)}|\leq 2\delta} w_{ab}(p_a^{(\epsilon)} + p_b^{(\epsilon)}) P_a^{(\epsilon)} h P_b^{(\epsilon)} \\
& + \epsilon\beta \sum_{a,b:|E_a^{(\epsilon)}-E_b^{(\epsilon)}|\leq 2\delta, E_a^{(\epsilon)} > E_b^{(\epsilon)}} w_{ab} p_a^{(\epsilon)} P_a^{(\epsilon)} h P_b^{(\epsilon)} \\
& + \epsilon\beta \sum_{a,b:|E_a^{(\epsilon)}-E_b^{(\epsilon)}|\leq 2\delta, E_a^{(\epsilon)} \leq E_b^{(\epsilon)}} w_{ab} p_b^{(\epsilon)} P_a^{(\epsilon)} h P_b^{(\epsilon)} \\
& + \epsilon\beta \sum_{a,b:|E_a^{(\epsilon)}-E_b^{(\epsilon)}|\leq 2\delta} w_{ab} \min(p_a, p_b) \beta \Delta_{ab} P_a^{(\epsilon)} h P_b^{(\epsilon)} + \mathcal{O}(\epsilon\beta^3 \|h\| \delta^2). \quad (5.62)
\end{aligned}$$

We want to rewrite w_{ab} as the inner product of two unit vectors. To this end, for each energy E_a define

$$|f_a\rangle = \sum_E \left(c_a^+ \delta_{E, E_a^{(\epsilon)+}} + c_a^- \delta_{E, E_a^{(\epsilon)-}} \right) |E\rangle. \quad (5.63)$$

The vectors $|E\rangle$ are the binary energy representations contained in the ancillae of the high precision phase estimation. Then $w_{ab} = \langle f_b | f_a \rangle$.

We can further rewrite $\Delta_{ab} = |E_a - E_b| = 2\delta - 2\delta(1 - |E_a - E_b|/(2\delta))$. We also know that the sum is over a, b such that $|E_a - E_b| \leq 2\delta$. We can absorb that condition into Δ_{ab} and write $\Delta_{ab} = 2\delta - 2\delta g_{ab}$, where g_{ab} is $1 - |E_a - E_b|/(2\delta)$. Using 5.2.13 from [122] and $|E_a - E_b| \leq 2\delta$, we know that the matrix g_{ab} is positive definite.

Finally, we have that $\min(p_a, p_b) = \sqrt{p_a p_b} \min(e^{-\beta(E_a - E_b)/2}, e^{-\beta(E_b - E_a)/2}) = \sqrt{p_a p_b} e^{-\beta\Delta_{ab}/2}$. From 5.2.17 in [122], we know that $e^{-\beta\Delta_{ab}/2}$ is a positive definite matrix.

Combining everything together, we can bound the trace norm of the the last line of (5.62):

$$\left\| 2\epsilon\beta^2\delta \sum_{a,b} \min(p_a, p_b) (1 - g_{ab}) w_{ab} h_{ab} \right\|_{\text{Tr}} \quad (5.64)$$

$$\leq 2\epsilon\beta^2\delta \left\| \sum_{a,b} \sqrt{p_a} e^{-\Delta_{ab}} (1 - g_{ab}) w_{ab} h_{ab} \sqrt{p_b} \right\|_{\text{Tr}} \quad (5.65)$$

$$\leq 2\epsilon\beta^2\delta \left\| \tilde{\rho} \circ h \circ w \circ e^{-\Delta} - \tilde{\rho} \circ h \circ w \circ g \circ e^{-\Delta} \right\|_{\text{Tr}}, \quad (5.66)$$

where $\tilde{\rho}$ is the density matrix corresponding to the pure state whose amplitudes are given

by $\{\sqrt{p_a}\}$. We can further replace h with $(h + \|h\|\mathbb{1}) - \|h\|\mathbb{1}$, where $\mathbb{1}$ is the identity matrix, and $(h + \|h\|\mathbb{1})$ is a positive definite matrix whose norm is less than $2\|h\|$. Using 3.4.3 and 2.7.12(iii) from [122] then gives us the bound

$$2\epsilon\beta^2\delta\left\|\tilde{\rho}\circ((h+\|h\|\mathbb{1})-\|h\|\mathbb{1})\circ w\circ e^{-\Delta}-\tilde{\rho}\circ((h+\|h\|\mathbb{1})-\|h\|\mathbb{1})\circ w\circ g\circ e^{-\Delta}\right\|_{\text{Tr}} \quad (5.67)$$

$$\leq 2\epsilon\beta^2\delta\left\|\tilde{\rho}\circ(h+\|h\|\mathbb{1})\circ w\circ e^{-\Delta}\right\|_{\text{Tr}}+\left\|\tilde{\rho}\circ(\|h\|\mathbb{1})\circ w\circ e^{-\Delta}\right\|_{\text{Tr}} \quad (5.68)$$

$$+\left\|\tilde{\rho}\circ(h+\|h\|\mathbb{1})\circ w\circ g\circ e^{-\Delta}\right\|_{\text{Tr}}+\left\|\tilde{\rho}\circ(\|h\|\mathbb{1})\circ w\circ g\circ e^{-\Delta}\right\|_{\text{Tr}} \quad (5.69)$$

$$\leq 12\epsilon\beta^2\delta\|h\|\|\tilde{\rho}\|_{\text{Tr}}\leq 12\epsilon\beta^2\delta\|h\|. \quad (5.70)$$

Simplifying the rest of the terms in (5.62) further gives:

$$\sum_a w_{aa}P_a^{(\epsilon)}p_a-\epsilon\beta/2\sum_{a,b:|E_a^{(\epsilon)}-E_b^{(\epsilon)}|\leq 2\delta} w_{ab}|p_b^{(\epsilon)}-p_a^{(\epsilon)}|P_a^{(\epsilon)}hP_b^{(\epsilon)}+\mathcal{O}(\epsilon\beta^2\delta\|h\|). \quad (5.71)$$

We can use the same trick as before and write $|p_a-p_b|=\sqrt{p_ap_b}|e^{-\beta(E_a-E_b)/2}-e^{-\beta(E_b-E_a)/2}|=\sqrt{p_ap_b}(\beta|E_a-E_b|+\mathcal{O}(\beta^2\delta^2))=\sqrt{p_ap_b}\beta(2\delta-2\delta g_{ab})+\mathcal{O}(\beta^2\delta^2)$, where g is defined above. Using the same procedure, the last error term is also bounded by $\mathcal{O}(\epsilon\beta^2\delta\|h\|)$.

Finally, choosing $\epsilon\beta^2\delta\|h\|=\varepsilon=\epsilon^2\beta^2\|h\|^2$ gives an overall error of $\mathcal{O}(\beta^2\epsilon^2\|h\|^2)$. \square

5.4 Time Requirements

We can merge two regions already thermalized into one large thermal region with the two subroutines just described using a sequence of small perturbative steps. Each step is successful with the probability $p\geq 1-\epsilon\beta\|h\|$. The average number of steps until we generate a complete sequence without failures is $\langle m\rangle\in\mathcal{O}(e^{\beta\|h\|})$.¹ Each time that we fail we need to produce two new thermal regions to be merged. The average number of failures is $\langle\alpha\rangle\in\mathcal{O}(e^{\beta\|h\|})$.

Now we analyze the average number of steps $\langle\tau(k)\rangle$ required to prepare a thermalized chain of length 2^k at level k in figure 5.2. Since α and $\tau(k-1)$ are independent random

¹This is also known from the theory of success runs. We give the average cost, but the tail has an exponential decay rate, so the worst case cost is similar (see, for instance [126]).

variables, we can calculate the expectation value of $\tau(k)$ as

$$\langle \tau(k) \rangle = 2\langle \alpha \rangle \langle \tau(k-1) \rangle + \langle m \rangle . \quad (5.72)$$

The cost of preparing the states at level 0 is smaller than the cost $\langle m \rangle$ of merging two sections, so we get

$$\langle \tau(k) \rangle \leq \langle m \rangle \sum_{j=0}^k (2\langle \alpha \rangle)^j = \langle m \rangle \frac{(2\langle \alpha \rangle)^{k+1} - 1}{2\langle \alpha \rangle - 1} = \mathcal{O} \left(\langle m \rangle (2\langle \alpha \rangle)^k \right) . \quad (5.73)$$

For the full chain,

$$\langle \tau(\log_2 N) \rangle = \mathcal{O} \left(e^{\beta \|h\|} \left(2e^{\beta \|h\|} \right)^{\log_2 N} \right) = \mathcal{O} \left(\exp(\log N (\beta \|h\| / \log 2 + 1)) \right) . \quad (5.74)$$

Now we bound the error. Each perturbative step as an error $\mathcal{O}(\epsilon^2 \beta^2 \|h\|^2)$, so adding a link gives a total error $\mathcal{O}(\epsilon \beta^2 \|h\|^2)$. The error at level k , $\epsilon(k)$, is bounded by

$$\epsilon(k) = 2\epsilon(k-1) + \mathcal{O}(\epsilon \beta^2 \|h\|^2) = \mathcal{O}(\epsilon \beta^2 \|h\|^2 2^{k+1}) . \quad (5.75)$$

The total error is $\mathcal{O}(N \epsilon \beta^2 \|h\|^2)$. If we choose $\epsilon = \bar{\epsilon} / (N \beta^2 \|h\|^2)$, we get a total error of $\mathcal{O}(\bar{\epsilon})$ in trace-norm. Finally, because the evolution time of each step is $\mathcal{O}(\log(1/(\epsilon \beta \|h\|)) / (\epsilon \|h\|))$, we obtain the dominant contribution to the total evolution time $\beta N^{\beta \|h\| / \log 2} / \bar{\epsilon}$.

Notice that $\langle \alpha \rangle = \langle m \rangle$, depending on the area that we are merging. For dimension D , at level k , we are going to merge lower dimensional hypercubes of edge size 2^{k-1} . To simplify the calculations (even though it does not give an optimal scaling of D in the exponent), we merge all the hypercubes at once. Then the area to merge is given by the surface of the face of the hypercube, $2^{(k-1)(D-1)}$, times the number of edges of the hypercube, $2^{D-1} D$. This gives

$$\langle \alpha(k) \rangle = e^{\beta \|h\| D 2^{k(D-1)}} . \quad (5.76)$$

After each failure we have to rebuild the 2^D nodes of the hypercube, each with a cost

$\langle \tau^D(k-1) \rangle$. This gives

$$\langle \tau^{(D)}(k) \rangle = \langle \alpha(k) \rangle (2^D \langle \tau^{(D)}(k-1) \rangle + 1) . \quad (5.77)$$

We can write $\langle \tau(-1) \rangle = 0$. A close expression for $\langle \tau(k) \rangle$ is

$$\langle \tau^{(D)}(k) \rangle = \sum_{j=0}^k \prod_{i=0}^j 2^D \langle \alpha(k-i) \rangle . \quad (5.78)$$

We can bound

$$\langle \tau(\log N) \rangle \leq \log_2 N \prod_{i=0}^{\log_2 N} 2^D \langle \alpha(\log_2 N - i) \rangle . \quad (5.79)$$

The exponent from the multiplication of the α 's is

$$\beta \|h\| \sum_{i=0}^{\log_2 N} 2^{(D-1)i} < \beta \|h\| \frac{2^{(\log_2 N + 1)(D-1)}}{2^{D-1} - 1} = \beta \|h\| 2^{\log_2 N (D-1)} \frac{2^{D-1}}{2^{D-1} - 1} \leq 2\beta \|h\| N^{D-1} , \quad (5.80)$$

which gives the bound

$$\langle \tau(\log N) \rangle < \log_2 N 2^{ND} e^{2\beta \|h\| DN^{D-1}} . \quad (5.81)$$

The dominant contribution is the term $e^{2\beta \|h\| DN^{D-1}}$.

The total time complexity and error can be analyzed by imagining that all the interactions between each site are added sequentially and successfully. The number of interactions is N^D , so the total error is $\epsilon DN^D \beta^2 \|h\|^2$. For a total error $\bar{\epsilon}$ we choose $\epsilon = \bar{\epsilon} / (DN^D \beta^2 \|h\|^2)$. We plug this into the running time of each elementary operation, as before, to obtain the dominant term

$$\beta \frac{e^{2\beta \|h\| DN^{D-1}}}{\bar{\epsilon}} . \quad (5.82)$$

5.5 Conclusion

We have presented an algorithm that prepares a thermal state of a one-dimensional quantum system in time polynomial in the system size and exponential in the inverse temperature (as required by the existence of QMA-complete ground state problems in one-dimension). This algorithm can be trivially generalized into D dimensions, by dividing up the D -dimensional system into many smaller pieces, and by combining these pieces recursively. This would result in an exponential speedup by reducing the time requirements of the simulation from $\exp(N^D)$ to $\exp(N^{D-1})$, where N is the linear dimension of the system. We do not get polynomial scaling with system size for $D > 1$ because the intersections of two neighboring regions scale with N^{D-1} . Note that this is to be expected because there exist two-dimensional ground states with constant gap that encode the solution to NP-complete problems.

There are also several possible improvements to the scaling of this algorithm. If one is interested in thermalizing a classical system with a small quantum perturbation one can first solve for the classical part of the Hamiltonian. Then, one would only need to use projections for the quantum perturbation. Also, if one is interested in thermalizing a quantum system with short-ranged quantum correlations, one can also use belief propagation [1, 2, 27, 28] to reduce the storage requirements from $\mathcal{O}(N)$ qubits to $\mathcal{O}(l \log(N))$, where l is a constant related to the quantum correlation length. This can be done by tracing out parts of the blocks that do not share any entanglement with the boundary to be merged.

Chapter 6

Conclusion

In this thesis, we discussed various classical and quantum algorithms to simulate static properties of quantum many-body systems. We observed in chapter 1 that analytical solutions for quantum many-body problems are very rare. If we want to extract any information about these systems from the equations of quantum mechanics, we need to resort to numerical calculations. As these calculations require exponential resources with system size, we stressed the need to find approximating algorithms. Throughout the rest of this thesis, we focused on alternative approximation schemes that improved the efficiency of the previously known methods.

In chapter 2, we generalized the classical belief propagation algorithm to the quantum setting. This algorithm allows us to solve for thermodynamic properties of certain quantum many-body systems in the high temperature limit. In chapter 3, we combine belief propagation with multiscale renormalization ansatz, which in return allows us to use belief propagation algorithm from chapter 2 at all temperatures. In chapter 4, we introduced anyonic multiscale renormalization ansatz, which solves for the ground state properties of anyonic systems. Finally, in chapter 5, we described a quantum algorithm, which prepares thermal states of any quantum many-body system on a quantum computer and achieves an exponential speedup over previously known methods.

The common point of the algorithms discussed in this thesis is the fact that they exploit the locality of physical Hamiltonians. Even though we need exponential resources to simulate the most general quantum many-body system, we can approximate solutions very well for systems defined on physically motivated geometries. In nature, most interactions are two-body and most Hamiltonians of interest have some local geometric structure. We have demonstrated in this thesis that quantum many-body problems can be solved efficiently by

exploiting this structure. We described both classical and quantum algorithms for these simulations, which have been more efficient than previously known methods.

Bibliography

- [1] D. Poulin and E. Bilgin, Phys. Rev. A **77**, 052318 (2008).
- [2] E. Bilgin and D. Poulin, Phys. Rev. B **81**, 054106 (2010).
- [3] R. König and E. Bilgin, Phys. Rev. B **82**, 125118 (2010).
- [4] E. Bilgin and S. Boixo, Phys. Rev. Lett. **105**, 170405 (2010).
- [5] S. R. White, Phys. Rev. Lett. **69**, 2863 (1992).
- [6] I. Affleck, T. Kennedy, E. H. Lieb, and H. Tasaki, Comm. Math. Phys. **115**, 477 (1988).
- [7] G. Vidal, Phys. Rev. Lett. **91**, 147902 (2003).
- [8] F. Verstraete and J. I. Cirac, “Renormalization algorithms for quantum-many body systems in two and higher dimensions,” (2004), arXiv:cond-mat/0407066 .
- [9] F. Verstraete and J. I. Cirac, Phys. Rev. B **73**, 094423 (2006).
- [10] M. B. Hastings, J. Stat. Mech.-Theory. E. , P08024 (2007).
- [11] M. B. Hastings, Phys. Rev. B **73**, 085115 (2006).
- [12] M. B. Hastings, Phys. Rev. B **76**, 035114 (2007).
- [13] N. Schuch, I. Cirac, and F. Verstraete, Phys. Rev. Lett. **100**, 250501 (2008).
- [14] G. Vidal, Phys. Rev. Lett. **98**, 070201 (2007).
- [15] J. S. Yedidia, W. T. Freeman, and Y. Weiss, *Mitsubishi Electric Research Laboratories Tech. Rep. No. TR-2001-22*, pp. 239–236 ((unpublished), 2002).
- [16] M. Wainwright, T. Jaakkola, and A. Willsky, IEEE T. Inform. Theory **49**, 1120 (2003).
- [17] S. Aji and R. McEliece, IEEE T. Inform. Theory **46**, 325 (2000).
- [18] J. S. Yedidia, *Advanced Mean Field Methods: Theory and Practice* (MIT Press, Cambridge, MA, 2001).
- [19] D. J. C. MacKay, *Information Theory, Inference and Learning Algorithms* (Cambridge University Press, Cambridge, UK, 2003).

- [20] M. Mézard and A. Montanari, *Constraint Satisfaction Networks in Physics and Computation* (Clarendon Press, Oxford, 2007).
- [21] R. G. Gallager, *Low Density Parity Check Codes* (MIT Press, Cambridge, MA, 1963).
- [22] C. Berrou, A. Glavieux, and P. Thitimajshima, in *ICC'93*, pp. 1064–1070 (IEEE, Genève, Switzerland, 1993).
- [23] T. Richardson and R. Urbanke, *Modern Coding Theory* (Cambridge University Press, Cambridge, England, 2008).
- [24] M. Mézard and G. Parisi, *Eur. Phys. J. B* **20**, 217 (2001).
- [25] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky, *Nature* **400**, 133 (1999).
- [26] M. Mézard, G. Parisi, and R. Zecchina, *Science* **297**, 812 (2002).
- [27] M. Leifer and D. Poulin, *Ann. Phys.* **323**, 1899 (2008).
- [28] M. B. Hastings, *Phys. Rev. B* **76**, 201102 (2007).
- [29] E. H. Lieb and M. B. Ruskai, *J. Math. Phys.* **14**, 1938 (1973).
- [30] P. Hayden, R. Jozsa, D. Petz, and A. Winter, *Comm. Math. Phys.* **246**, 359 (2004).
- [31] M. B. Ruskai, *J. Math. Phys.* **43**, 4358 (2002).
- [32] M. M. Wolf, F. Verstraete, M. B. Hastings, and J. I. Cirac, *Phys. Rev. Lett.* **100**, 070502 (2008).
- [33] M. Fannes, *Comm. Math. Phys.* **31**, 291 (1973).
- [34] C. Fuchs and J. van de Graaf, *IEEE T. Inform. Theory* **45**, 1216 (1999).
- [35] P. Pfeuty, *Ann. Phys.* **57**, 79 (1970).
- [36] M. Zwolak and G. Vidal, *Phys. Rev. Lett.* **93**, 207205 (2004).
- [37] C. Laumann, A. Scardicchio, and S. L. Sondhi, *Phys. Rev. B* **78**, 134424 (2008).
- [38] G. Vidal, *Phys. Rev. Lett.* **99**, 220405 (2007).
- [39] G. Evenbly and G. Vidal, *Phys. Rev. B* **79**, 144108 (2009).
- [40] K. G. Wilson, *Rev. Mod. Phys.* **47**, 773 (1975).
- [41] G. Vidal, *Phys. Rev. Lett.* **101**, 110501 (2008).
- [42] L. Cincio, J. Dziarmaga, and M. M. Rams, *Phys. Rev. Lett.* **100**, 240603 (2008).
- [43] M. Aguado and G. Vidal, *Phys. Rev. Lett.* **100**, 070404 (2008).
- [44] R. König, B. W. Reichardt, and G. Vidal, *Phys. Rev. B* **79**, 195123 (2009).
- [45] I. Peschel and V. Eisler, *J. Phys. A* **42**, 504003 (2009).

- [46] F. Krzakala, A. Rosso, G. Semerjian, and F. Zamponi, *Phys. Rev. B* **78**, 134428 (2008).
- [47] P. Silvi, V. Giovannetti, S. Montangero, M. Rizzi, J. I. Cirac, and R. Fazio, *Phys. Rev. A* **81**, 062335 (2010).
- [48] V. Giovannetti, S. Montangero, and R. Fazio, *Phys. Rev. Lett.* **101**, 180503 (2008).
- [49] R. N. C. Pfeifer, G. Evenbly, and G. Vidal, *Phys. Rev. A* **79**, 040301 (2009).
- [50] A. Y. Kitaev, *Ann. Phys.* **303**, 2 (2003).
- [51] J. Preskill, “Topological quantum computation,” *Lecture Notes (Chapter 9)* (2004).
- [52] M. H. Freedman, A. Kitaev, M. J. Larsen, and Z. Wang, *Bull. Amer. Math. Soc.* **40**, 31 (2003).
- [53] C. Nayak, S. H. Simon, A. Stern, M. Freedman, and S. Das Sarma, *Rev. Mod. Phys.* **80**, 1083 (2008).
- [54] G. Moore and N. Read, *Nuclear Physics B* **360**, 362 (1991).
- [55] N. Read and E. Rezayi, *Phys. Rev. B* **59**, 8084 (1999).
- [56] E. H. Rezayi and N. Read, *Phys. Rev. B* **79**, 075306 (2009).
- [57] M. Dolev, M. Heiblum, V. Umansky, A. Stern, and D. Mahalu, *Nature* **452**, 829 (2008).
- [58] I. P. Radu, J. B. Miller, C. M. Marcus, M. A. Kastner, L. N. Pfeiffer, and K. W. West, *Science* **320**, 899 (2008).
- [59] R. L. Willett, L. N. Pfeiffer, and K. W. West, *P. Natl. Acad. Sci.* **106**, 8853 (2009).
- [60] C. L. Kane and E. J. Mele, *Phys. Rev. Lett.* **95**, 226801 (2005).
- [61] J. E. Moore and L. Balents, *Phys. Rev. B* **75**, 121306 (2007).
- [62] L. Fu, C. L. Kane, and E. J. Mele, *Phys. Rev. Lett.* **98**, 106803 (2007).
- [63] R. Roy, *Phys. Rev. B* **79**, 195322 (2009).
- [64] A. Kitaev, *Ann. Phys.* **321**, 2 (2006).
- [65] M. A. Levin and X.-G. Wen, *Phys. Rev. B* **71**, 045110 (2005).
- [66] L.-M. Duan, E. Demler, and M. D. Lukin, *Phys. Rev. Lett.* **91**, 090402 (2003).
- [67] A. Micheli, G. K. Brennen, and P. Zoller, *Nat. Phys.* **2**, 341 (2006).
- [68] S. Trebst, P. Werner, M. Troyer, K. Shtengel, and C. Nayak, *Phys. Rev. Lett.* **98**, 070602 (2007).
- [69] I. S. Tupitsyn, A. Kitaev, N. V. Prokof'ev, and P. C. E. Stamp, *Phys. Rev. B* **82**, 085114 (2010).

- [70] J. Vidal, S. Dusuel, and K. P. Schmidt, *Phys. Rev. B* **79**, 033109 (2009).
- [71] J. Vidal, R. Thomale, K. P. Schmidt, and S. Dusuel, *Phys. Rev. B* **80**, 081104 (2009).
- [72] I. Klich, *Ann. Phys.* **325**, 2120 (2010).
- [73] S. Bravyi, M. B. Hastings, and S. Michalakis, *J. Math. Phys.* **51**, 093512 (2010).
- [74] S. Bravyi and M. Hastings, “A short proof of stability of topological order under local perturbations,” (2010), arXiv:1001.4363 .
- [75] P. Bonderson, *Phys. Rev. Lett.* **103**, 110403 (2009).
- [76] A. Feiguin, S. Trebst, A. W. W. Ludwig, M. Troyer, A. Kitaev, Z. Wang, and M. H. Freedman, *Phys. Rev. Lett.* **98**, 160409 (2007).
- [77] N. E. Bonesteel and K. Yang, *Phys. Rev. Lett.* **99**, 140405 (2007).
- [78] L. Fidkowski, G. Refael, N. E. Bonesteel, and J. E. Moore, *Phys. Rev. B* **78**, 224204 (2008).
- [79] L. Fidkowski, H.-H. Lin, P. Titum, and G. Refael, *Phys. Rev. B* **79**, 155120 (2009).
- [80] R. N. C. Pfeifer, P. Corboz, O. Buerschaper, M. Aguado, M. Troyer, and G. Vidal, *Phys. Rev. B* **82**, 115126 (2010).
- [81] R. König, *Phys. Rev. A* **81**, 052309 (2010).
- [82] S. Rommer and S. Östlund, *Phys. Rev. B* **55**, 2164 (1997).
- [83] G. Vidal, *Phys. Rev. Lett.* **93**, 040502 (2004).
- [84] F. Verstraete, D. Porras, and J. I. Cirac, *Phys. Rev. Lett.* **93**, 227205 (2004).
- [85] F. Verstraete, V. Murg, and J. I. Cirac, *Adv. Phys.* **57**, 143 (2008).
- [86] S. Trebst, M. Troyer, Z. Wang, and A. W. W. Ludwig, *Progr. of Theor. Phys. Suppl.* **176**, 384 (2008).
- [87] G. Evenbly and G. Vidal, *New J. Phys.* **12**, 025007 (2010).
- [88] P. Corboz and G. Vidal, *Phys. Rev. B* **80**, 165129 (2009).
- [89] F. Verstraete and J. I. Cirac, *Phys. Rev. A* **70**, 060302 (2004).
- [90] Z.-C. Gu, M. Levin, and X.-G. Wen, *Phys. Rev. B* **78**, 205116 (2008).
- [91] I. Pižorn and F. Verstraete, *Phys. Rev. B* **81**, 245110 (2010).
- [92] C. V. Kraus, N. Schuch, F. Verstraete, and J. I. Cirac, *Phys. Rev. A* **81**, 052338 (2010).
- [93] P. Corboz, R. Orús, B. Bauer, and G. Vidal, *Phys. Rev. B* **81**, 165104 (2010).
- [94] C. Pineda, T. Barthel, and J. Eisert, *Phys. Rev. A* **81**, 050303 (2010).

- [95] T. Barthel, C. Pineda, and J. Eisert, Phys. Rev. A **80**, 042333 (2009).
- [96] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Phys. Rev. Lett. **98**, 140506 (2007).
- [97] K. Walker, “On Witten’s 3-manifold invariants,” (1991), available at <http://canyon23.net/math/>.
- [98] R. Koenig, G. Kuperberg, and B. W. Reichardt, Ann. Phys. **325**, 2707 (2010).
- [99] P. Bonderson, M. Freedman, and C. Nayak, Ann. Phys. **324**, 787 (2009).
- [100] C. Gils, E. Ardonne, S. Trebst, A. W. W. Ludwig, M. Troyer, and Z. Wang, Phys. Rev. Lett. **103**, 070401 (2009).
- [101] V. Lahtinen, G. Kells, A. Carollo, T. Stitt, J. Vala, and J. K. Pachos, Ann. Phys. **323**, 2286 (2008).
- [102] M. Baraban, G. Zikos, N. Bonesteel, and S. H. Simon, Phys. Rev. Lett. **103**, 076801 (2009).
- [103] M. Cheng, R. M. Lutchyn, V. Galitski, and S. Das Sarma, Phys. Rev. Lett. **103**, 107001 (2009).
- [104] S. Montangero, M. Rizzi, V. Giovannetti, and R. Fazio, Phys. Rev. B **80**, 113103 (2009).
- [105] C. K. Majumdar and D. K. Gosh, J. Math. Phys. **10**, 1388 (1969).
- [106] S. Trebst, E. Ardonne, A. Feiguin, D. A. Huse, A. W. W. Ludwig, and M. Troyer, Phys. Rev. Lett. **101**, 050401 (2008).
- [107] G. E. Andrews, R. J. Baxter, and P. J. Forrester, J. Stat. Phys. **35**, 193 (1984).
- [108] A. W. W. Ludwig, D. Poilblanc, S. Trebst, and M. Troyer, “Two-dimensional quantum liquids from interacting non-abelian anyons,” arXiv:1003.3453 .
- [109] Y.-Y. Shi, L.-M. Duan, and G. Vidal, Phys. Rev. A **74**, 022320 (2006).
- [110] L. Tagliacozzo, G. Evenbly, and G. Vidal, Phys. Rev. B **80**, 235127 (2009).
- [111] V. Murg, F. Verstraete, O. Legeza, and R. M. Noack, Phys. Rev. B **82**, 205105 (2010).
- [112] R. P. Feynman, Int. J. Theor. Phys. **21**, 467 (1982).
- [113] S. Lloyd, Science **273**, 1073 (1996).
- [114] A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and Quantum Computation* (Amer. Math. Soc., 2002).
- [115] J. Kempe, A. Kitaev, and O. Regev, SIAM J. Comput. **35**, 1070 (2006).
- [116] R. Oliveira and B. M. Terhal, Quant. Inf. Comp. **8**, 0900 (2008).
- [117] D. Aharonov, D. Gottesman, S. Irani, and J. Kempe, Commun. Math. Phys. **287**, 41 (2009).

- [118] B. M. Terhal and D. P. DiVincenzo, Phys. Rev. A **61**, 022301 (2000).
- [119] K. Temme, T. J. Osborne, K. G. Vollbrecht, D. Poulin, and F. Verstraete, Nature **471**, 87 (2011).
- [120] M. Cramer and J. Eisert, New J. Phys. **12**, 055020 (2010).
- [121] D. Poulin and P. Wocjan, Phys. Rev. Lett. **103**, 220502 (2009).
- [122] R. Bhatia, *Positive Definite Matrices* (Princeton University Press, Princeton, New Jersey, 2007).
- [123] A. M. Childs, E. Deotto, E. Farhi, J. Goldstone, S. Gutmann, and A. J. Landahl, Phys. Rev. A **66**, 032314 (2002).
- [124] S. Boixo, E. Knill, and R. D. Somma, QIC **9**, 833 (2009).
- [125] W. Feller, *An Introduction to Probability Theory and Its Applications, Vol. 1, 3rd Edition*, 3rd ed. (Wiley, 1968).
- [126] N. Balakrishnan and M. V. Koutras, *Runs and Scans with Applications*, 1st ed. (Wiley-Interscience, 2001).
- [127] E. Knill, G. Ortiz, and R. D. Somma, Phys. Rev. A **75**, 012328 (2007).
- [128] R. D. Somma, S. Boixo, H. Barnum, and E. Knill, Phys. Rev. Lett. **101**, 130504 (2008).
- [129] C. Chiang and P. Wocjan, “Quantum algorithm for preparing thermal gibbs states - detailed analysis,” (2010), arXiv:1001.1130 .