

Local Hamiltonians in Quantum Computation

by

Daniel Nagaj

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2008

© Daniel Nagaj, MMVIII. All rights reserved.

The author hereby grants to MIT permission to reproduce and distribute publicly
paper and electronic copies of this thesis document in whole or in part.

Author
Department of Physics
May 16, 2008

Certified by
Edward H. Farhi
Professor of Physics; Director, Center for Theoretical Physics
Thesis Supervisor

Accepted by
Thomas J. Greytak
Lester Wolfe Professor of Physics
Associate Department Head for Education

Local Hamiltonians in Quantum Computation

by
Daniel Nagaj

Submitted to the Department of Physics
on May 16, 2008, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Physics

Abstract

In this thesis, I investigate aspects of local Hamiltonians in quantum computing. First, I focus on the Adiabatic Quantum Computing model, based on evolution with a time-dependent Hamiltonian. I show that to succeed using AQC, the Hamiltonian involved must have local structure, which leads to a result about eigenvalue gaps from information theory. I also improve results about simulating quantum circuits with AQC. Second, I look at classically simulating time evolution with local Hamiltonians and finding their ground state properties. I give a numerical method for finding the ground state of translationally invariant Hamiltonians on an infinite tree. This method is based on imaginary time evolution within the Matrix Product State ansatz, and uses a new method for bringing the state back to the ansatz after each imaginary time step. I then use it to investigate the phase transition in the transverse field Ising model on the Bethe lattice. Third, I focus on locally constrained quantum problems Local Hamiltonian and Quantum Satisfiability and prove several new results about their complexity. Finally, I define a Hamiltonian Quantum Cellular Automaton, a continuous-time model of computation which doesn't require control during the computation process, only preparation of product initial states. I construct two of these, showing that time evolution with a simple, local, translationally invariant and time-independent Hamiltonian can be used to simulate quantum circuits.

Thesis Supervisor: Edward H. Farhi

Title: Professor of Physics; Director, Center for Theoretical Physics

Acknowledgments

I gratefully acknowledge financial support from the MIT Presidential Fellowship, the National Security Agency (NSA) and Advanced Research and Development Activity (ARDA) under Army Research Office (ARO) contract W911NF-04-1-0216, and the W. M. Keck Foundation Center for Extreme Quantum Information Theory during my five years at MIT.

I wasn't alone on my journey to finishing this thesis. First, my thanks go to God for making this world a joy to explore. My parents gave me their love, support and freedom. My undergraduate advisor, Vlado Bužek, introduced me to exciting research. Eddie Farhi, my Ph.D. thesis advisor, taught me to be patient and to ask more questions. He gave me not only good direction, but also plenty of freedom in research. During my time at MIT, I have been honored to coauthor with Jeffrey Goldstone, Sam Gutman, Peter Shor, Iordanis Kerenidis, Shay Mozes, Pawel Wocjan and Igor Sylvester. Besides them, I learned much from discussions about research with Seth Lloyd, Ike Chuang, Guifre Vidal, Peter Love, Oded Regev, Stephen Jordan, Jake Taylor and an anonymous referee. I also benefited from being a teaching assistant for Gunther Roland, Young Lee, John Negele and Bruce Knuteson. Ashdown was a great place to live in because of Noel, Abhi, Jon, Ji-Eun, Marcus, Junsik and the housemasters, Terry and Ann Orlando. Problem sets were do-able with the help of Ambar, Mark, Stephen, Pouyan and Seungeun. Jon taught me to throw a frisbee and juggle, Dave taught me to drive and Ambar taught me to cook Indian food. Helen, Kim and Rachel helped me to grow in photography. Through the MIT Graduate Christian Fellowship, I found friends in Jon, Colin, Jeff, Wouter, Mary, Kevin, John, Ken, Bethany, Cynthia, Eric, Kyle, Stephen, Stephen and others. Looking back at all this today, I am thankful and my joy is great, for I share it with Pavla, the best of wives.

Contents

0	Introduction	17
0.1	Outline and Summary of Results	18
1	Preliminaries	19
1.1	Classical Computation	19
1.1.1	Introduction to Computation	19
1.1.2	Classical Complexity	22
1.1.3	Classical locally constrained problems	24
1.2	Quantum Computing	28
1.2.1	The Quantum Circuit model	29
1.2.2	Quantum Complexity	31
1.2.3	The Strengths and Limits of Quantum Computing	32
1.3	Local Hamiltonians in Quantum Computation	33
1.3.1	Feynman’s Hamiltonian Quantum Computer	34
1.3.2	Kitaev’s Local Hamiltonian (LH) Problem	36
2	Adiabatic Quantum Computation	41
2.1	Introduction to AQC	42
2.1.1	The power of AQC	43
2.2	AQC is equivalent to BQP	44
2.2.1	The Hamiltonian Computer model and universality	46
2.3	How to make the Quantum Adiabatic Algorithm Fail	49
2.3.1	General search starting with a one-dimensional projector	50
2.3.2	Search with a scrambled problem hamiltonian	55
2.3.3	Summary	59
3	Matrix Product States on Infinite Trees	61
3.1	Introduction	62
3.2	Matrix Product States	63
3.2.1	MPS for a spin chain	63
3.2.2	MPS on Trees	65
3.3	Simulating Quantum Systems with MPS	66
3.3.1	Unitary Update Rules	66
3.3.2	Imaginary Time Evolution	68
3.4	MPS and Translationally Invariant Systems	69

3.4.1	An Infinite Line	69
3.4.2	An Infinite Tree	70
3.4.3	Expectation Values	71
3.5	Quantum Transverse Field Ising Model	73
3.5.1	The Infinite Line.	73
3.5.2	The Infinite Tree (Bethe Lattice).	77
3.6	The Not 00 Model	78
3.6.1	Infinite Line	80
3.6.2	Infinite Tree	82
3.7	Stability and Correlation Lengths on the Bethe Lattice	82
4	Quantum Satisfiability	91
4.1	Introduction	92
4.1.1	Quantum k-SAT	92
4.1.2	Constructing the Hamiltonians	93
4.1.3	Constructing Clocks	95
4.2	A new 3-local QMA complete Local Hamiltonian (Quantum 5-SAT from 3-local terms)	100
4.2.1	Bravyi's Quantum 4-SAT	101
4.2.2	A qutrit clock implementation	102
4.2.3	Clock register construction	102
4.2.4	Checking correct application of gates and clock propagation	104
4.2.5	The new 3-local QMA complete construction (for qubits)	105
4.2.6	Checking gate application with 3-local terms	107
4.2.7	Clock propagation	107
4.2.8	Discussion and further directions	109
4.3	Quantum 2-SAT on a line	111
4.3.1	A Line of $d = 12$ Dimensional Particles (Quantum (12,12)-SAT)	111
4.3.2	Q-(11,11)-SAT on a line is QMA ₁ complete.	116
4.4	Quantum 2-SAT in general geometry (for qudits)	118
4.4.1	The Triangle Clock and Q-(5,3)-SAT	118
4.4.2	A QMA ₁ -complete 3-local Hamiltonian with restricted terms	122
4.5	Train Switch Construction: Universal Quantum Computation using a Quantum 3-SAT Hamiltonian	124
5	Hamiltonian Quantum Cellular Automata	131
5.1	Introduction	131
5.2	The HQCA in 1D with cell size $d = 10$	133
5.2.1	The Construction	133
5.2.2	Required Time Analysis	136
5.3	The HQCA in 1D with cell size $d = 20$	139
5.3.1	The Construction	139
5.3.2	Required Evolution Time Analysis	143
6	Conclusions	145
A	Kitaev's propagation Hamiltonian	149

B	Dyson Series for an AQC Hamiltonian	151
C	Continuous-time Quantum Walks in 1D	155
C.1	Quantum Walk on a Line	155
C.2	Quantum Walk on a Circle	158
D	Diffusion of Fermions on a Line	163
E	Transitions in a two level system	167

List of Figures

1-1	The Turing Machine.	20
1-2	A binary addition circuit for two two-bit numbers, $z_2 z_1 z_0 = x_1 x_0 \oplus y_1 y_0$. .	21
1-3	The implication graph for a 2-SAT instance $(x \vee y') \wedge (x' \vee y') \wedge (y \vee w) \wedge (y \vee z') \wedge (z \vee w')$. There is a directed loop containing both y and y' , so the instance is not satisfiable.	25
1-4	Illustration for the Cook-Levin theorem. Transforming a verifier circuit into a 3-SAT instance.	26
1-5	The GSAT and WALKSAT algorithms.	27
1-6	A reduction of Graph 3-colorability to 2-Satisfiability for trit-bit pairs, also called (3,2)-SAT. The trits with possible values $\{A,B,C\}$ are denoted by triangles, while the bits are denoted by circles. Each clause is represented by the pair of symbols which it rules out. For Graph 3-colorability, the excluded assignments are AA, BB and CC, ruling out the same color for two neighboring vertices.	27
1-7	The components of a quantum circuit.	30
1-8	A quantum circuit with measurements performed during the computation and its version with measurements deferred until the end of the computation. .	31
1-9	Feynman's computer with a hopping pointer particle (a single spin up). When the active site moves from c_2 to c_3 , gate U_3 gets applied to work qubits q_5 and q_6 . When the active spin jumps back from c_2 to c_1 , gate U_2 is uncomputed on qubits q_2 and q_3	35
2-1	Padding the circuit U with L gates by $2L$ extra identity gates. The state of the system after L steps stays unchanged.	47
2-2	Median required run time T versus bit number. At each bit number there are 50 random instances of Exact Cover with a single satisfying assignment. I choose the required run time to be the value of T for which quantum adiabatic algorithm has success probability between 0.2 and 0.21. For the projector beginning Hamiltonian I use (2.24) with $E = n/2$. The plot is log-linear. The error bars show the 95% confidence interval for the true medians. . . .	54
2-3	The scaled ground state energy E/n for a quantum adiabatic algorithm Hamiltonian of a decoupled problem. The lowest curve corresponds to the original decoupled problem. The upper "triangular" curves correspond to single instances of the n -bit decoupled problem, where the problem Hamiltonian was scrambled.	58
3-1	The Bethe lattice (infinite Cayley tree).	62

3-2	Two successive Schmidt decompositions on a line allow me to find the Γ tensor for the marked site and the two λ vectors for the bonds coming out of it.	64
3-3	The three Schmidt decompositions on a tree required to obtain the Γ tensor for the marked site and the three λ vectors for the bonds emanating from it.	65
3-4	The parametrization and update rules for the infinite line.	70
3-5	The two-layer, directed labeling of the tree.	70
3-6	The two-step interactions for the infinite tree.	71
3-7	Transverse Ising model on an infinite line. Fractional difference of the ground state energy obtained using MPS and the exact ground state, near the phase transition at $s_L = \frac{2}{3}$. The energy scale is logarithmic.	74
3-8	Transverse Ising model on an infinite line. The first and second derivative with respect to s of the ground state energy obtained via MPS compared with the exact result.	74
3-9	Transverse Ising model on an infinite line. Magnetization obtained using MPS vs s , with $B_z = 10^{-8}$	75
3-10	Transverse Ising model on an infinite line. Log-log plot of magnetization vs. $x - x_L$. I also plot a line with slope $\beta_L = 0.125$	76
3-11	Transverse Ising model on an infinite line. Log-log plot of the correlation length vs. $ x - x_L $. I also plot a line with slope $-\nu_L = -1$. For each χ in the plot, I choose a numerical value of the critical point s_L as the point at which the correlation length is maximal.	76
3-12	Transverse Ising model on an infinite tree. The first and second derivative with respect to s of the ground state energy obtained via MPS.	77
3-13	Transverse Ising model on an infinite tree. Magnetization vs. s , with $B_z = 10^{-8}$	78
3-14	Transverse Ising model on an infinite tree. Log-log plot of magnetization vs. $x - x_T$, with $B_z = 10^{-8}$. I also plot a line with slope 0.41.	79
3-15	Transverse Ising model on an infinite tree. A linear plot of the correlation length vs. s	79
3-16	The NOT 00 model on an infinite line. The ground state energy and its first two derivatives with respect to s	80
3-17	The NOT 00 model on an infinite line. Magnetization as a function of s and correlation length as a function of s	81
3-18	Ground state of the NOT 00 model on a line at $s = 1$. Comparison of the exact Schmidt coefficients for a division across the middle of a finite chain and of the MPS values ($\chi = 32$) for an infinite line.	82
3-19	The NOT 00 model on an infinite tree. The ground state energy and its first two derivatives with respect to s	83
3-20	The NOT 00 model on an infinite tree. Magnetization as a function of s and correlation length as a function of s	83
3-21	A system of classical spins on a lattice whose layers are Cayley trees. A and B denote the Boltzmann factors.	84
3-22	The $M = 2$ Cayley tree.	85
3-23	Computing the probability distribution of \vec{s}_0 and \vec{s}_d	88
4-1	The ‘triangle’ clock transition rules allowing to apply a controlled gate to two work qubits using interactions that involve one work qubit at a time.	98

4-2	The ‘pass the hat forward and then jump over it’ technique for a geometric clock (qubit transport) on a line from [5]. The progression of 20 states is shown.	100
4-3	Clock register consisting of $2L$ qubits and L qutrits.	103
4-4	Illustration of the two-step clock pointer propagation.	104
4-5	Clock register construction with $2L + 2L$ qubits.	106
4-6	Illustration of the two-step clock pointer propagation.	108
4-7	A quantum circuit for $n = 3$ qubits with $K = 4$ rounds of nearest neighbor gates $U_{k,s}$	112
4-8	The ‘triangle’ clock transition rules allowing to apply a controlled gate to two work qubits using interactions that involve only one work qubit at a time.	119
4-9	The ‘pulse’ clock construction consists of $T + 1$ clock qubits on a line, where only one of them is in the state $ 1\rangle$ at a time.	125
4-10	The ‘train switch’ clock construction for the application of a CNOT gate using only 3-local terms. There are four extra qubits inserted between c_t and c_{t+1} . The transitions from the original clock qubits to the gadget qubits are controlled by one of the work qubits. During the transition from u_1 to u_2 , the target work qubit is flipped. This gadget checks the application of a CNOT gate using 3-local projectors on qubits.	125
4-11	a) Computation on a line of states $ \Psi_t\rangle$. b) The line corresponding to a circuit U padded with extra identity gates. The region where the computation is done is marked. c) Computation on a cycle of states $ \Psi_t\rangle$ coming from two copies of the computation on a line. d) The cycle corresponding to a padded circuit U	127
5-1	a) A quantum circuit consisting of two sequences of gates acting on nearest neighbors. b) The previous circuit with a third sequence of identity gates added.	134
C-1	A contour plot of $ \lambda_j - \lambda_k ^{-1}$ with λ_j from (C.20). The sum in (C.31) is over the marked region.	159

List of Tables

1.1	Known complexity for classical satisfiability problems. $(3, 2)$ -SAT in the general strings column is the problem where each clause involves one trit (letter with three values) and one bit.	28
4.1	Known complexity of Quantum Satisfiability and Local Hamiltonian problems. My own results, presented later in this Chapter, are marked with *. .	93

Chapter 0

Introduction

Today's computers, working in binary, churn terabytes of classical information for us every day, solving computational problems. Nevertheless, for some of these problems, the required resources for even the best algorithms available today scale unfavorably with the problem size. Such computational difficulty can be turned into a resource of its own, motivating cryptographic schemes. Today's favorite, RSA, is based on the hardness of factoring large numbers. It can withstand attacks by even bigger and faster computers, as making the input size of the problem larger increases the required resources exponentially. Is there then a chance of ever solving hard problems? For many optimization problems like 3-SAT, the expected answer is no. There are problems thought to be easier than 3-SAT, such as graph isomorphism, the problem to determine whether two graphs are isomorphic or not. Many scientists have devoted many hours to obtain an efficient algorithm for it and did not succeed, but the jury is still out on whether one exists. Another problem thought hard classically is factoring. In 1994, Shor found an algorithm which runs in time polynomial, rather than exponential, in the number of digits of the number to be factored. However, Shor's algorithm requires a *quantum computer*, the target of considerable effort by many research groups, with nuclear spins, optical lattices, quantum dots, superconducting circuits or trapped ions as their building blocks. In fact, a small, single-purpose, NMR-based quantum computer was built by Isaac Chuang's group, ran Shor's algorithm, and factored 15. Needing to go much farther than this proof-of-principle quantum computer, today's effort focuses on developing and building scalable and fault-tolerant quantum computers.

The world is quantum-mechanical. It is interesting to investigate what the laws of physics would allow us to compute, if we could control a quantum system and use it as a computer. We would like to know the power and the limits of computers that would use quantum instead of classical information. From understanding the details of chemical reactions to finding ground state properties of quantum spin glass models, there are many problems whose nature is quantum mechanical. It is possible that a quantum computer could help solve these. Not only that, there are many classical problems where algorithmic speedups for quantum computers over their classical counterparts have been found, and there is hope that many are yet to be discovered.

Quantum mechanical systems behave according to the Schrödinger equation, governed by the Hamiltonian operator of the system, which describes its interactions. In our experience, the interactions between particles involve only a few particles at a time. A Hamiltonian describing such a system is called *local*. More specifically, if the particles are also spatially close, I will call the Hamiltonian *spatially local*, and emphasize this distinction throughout

the text. In this thesis, I ask—and answer—questions about the power of local Hamiltonians, whether how to simulate them, what could we do with them if we had them, or how hard it is to find their properties in the first place.

0.1 Outline and Summary of Results

After reviewing classical and quantum computing and introducing local Hamiltonians, illustrated by Feynman’s Hamiltonian Computer and Kitaev’s Local Hamiltonian problem in Chapter 1, I present the following new results in the next four chapters:

Chapter 2: Adiabatic Quantum Computation investigates the AQC model of quantum computation, which is based on time-evolving a system with a time-dependent Hamiltonian from a specific initial ground state. I define a general Hamiltonian Computer model, show that it is universal for quantum computing, and show that it also implies the universality of AQC for quantum computation. I improve the estimates on the runtime of AQC simulations of the quantum circuit model using the results I obtain later in Chapter 4. Then, based on [34], I prove two theorems about what not to do when designing quantum adiabatic algorithms, in essence saying that the Hamiltonians one uses must be local, if one hopes to succeed.

Chapter 3: Matrix Product States on Infinite Trees is based on [69] and contains a new numerical method for investigating the ground state properties of translationally invariant Hamiltonians on an infinite tree (Bethe lattice). This method is based on imaginary time evolution within the Matrix Product State ansatz. I use it to analyze two specific interaction models, first of which is the Ising model in transverse field on the Bethe lattice.

Chapter 4: Quantum Satisfiability analyzes the complexity of the quantum analogue of classical Satisfiability. I give several novel clock constructions useful for proofs of QMA completeness of Hamiltonian problems and for proofs of universality of Hamiltonian computers. Based on [70], I show that the 3-local Hamiltonian problem stays QMA-complete even without high-norm penalty terms required in previous constructions and that Quantum (3,2,2)-SAT is QMA_1 complete. I then present two previously unpublished results. First, Quantum 2-SAT on a line with 11-dimensional particles is QMA_1 complete. Second, the Quantum 3-SAT Hamiltonian is universal for quantum computing, when used in my Hamiltonian Computer model.

Chapter 5: Hamiltonian Quantum Cellular Automata is based on [71]. I define HQCA, a model of quantum computation with a local, time-independent, translationally invariant and problem-independent Hamiltonian. I present two such automata in 1D geometry, one built from particles with dimension $d = 10$, and the other from particles with dimension $d = 20$. These results are then related back to the universality of adiabatic quantum computation.

Finally, I summarize the thesis in Chapter 6, and discuss further research directions. Additional material can be found at the end of the thesis. Appendix A investigates Kitaev’s propagation Hamiltonian, Appendix B works out the details for a Trotter-like formula for classically simulating an adiabatic algorithm, Appendix C is an analysis of the continuous time quantum walk on a line and cycle, Appendix D concerns the diffusion of free fermions in 1D, and Appendix E looks at transitions in a two-level system.

Chapter 1

Preliminaries

This preliminary Chapter is aimed at a reader with little background in classical computing and complexity, and only basic knowledge of quantum mechanics. Kitaev's Classical and Quantum Computation [56] and Nielsen & Chuang's Quantum Information and Computation [73] are great starting points to investigate the topics mentioned herein in much more detail. The advanced reader should start straight with Section 1.3 on Local Hamiltonians, where I review the basic building blocks for many results in this thesis: Feynman's Hamiltonian computer and Kitaev's Local Hamiltonian problem.

This Chapter is organized as follows. I start with the classical model of computation in Section 1.1.1, review classical complexity in Section 1.1.2, and then look at classical locally constrained problems, k -Satisfiability and Max- k -Satisfiability, in Section 1.1.3.

In Section 1.2, I introduce quantum computation, discussing the quantum circuit model and universality in Section 1.2.1, quantum complexity in Section 1.2.2, and the strengths and limits of quantum computers in 1.2.3.

Finally, in Section 1.3 I turn my attention to the main topic of this thesis, local Hamiltonians in quantum computing, introducing Feynman's Hamiltonian computer model in Section 1.3.1 and Kitaev's QMA-complete Local Hamiltonian problem in Section 1.3.2.

1.1 Classical Computation

1.1.1 Introduction to Computation

I have a computer that I am typing this on. What can I do (compute) with it? I can surely add and multiply numbers really fast, but could I also use it to factor large numbers? It can effortlessly beat me in chess, but I am up to the challenge in go¹. Could I have programmed it to write this thesis? How big would the program have to be? How long would it take to execute? The question of computability and effectiveness of computation needs to be looked at in a particular model of computing. In 1937, Alan Turing introduced his remarkable Turing machine (TM) model [99, 100]. Consider a semi-infinite tape with symbols $s_i \in \{0, 1, \sqcup\}$ on it (see Figure 1-1), a read/write head with a finite memory (internal state space) positioned on the left end of the tape, and a finite set of instructions. The head is initialized in the state q_0 and positioned on the left of the tape, reading the symbol s_1 . The computation is then performed in sequential steps according to the instruction set.

¹though not against the very recent Monte Carlo programs like MoGo

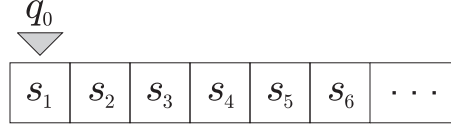


Figure 1-1: The Turing Machine.

These instructions have the form

$$[current\ state], [symbol\ being\ read] \longrightarrow [next\ state], [symbol\ to\ write], [where\ to\ move].$$

Depending on the current state of the machine and the symbol that is under the read/write head, the machine can change its state, write a symbol to the tape and move one step to the left (denoted by \triangleleft), one step to the right (\triangleright) or stay at its current position (\cdot). It is straightforward to run the machine according to the instruction set. The execution stops only when the internal state of the head changes to a special state **HALT**. The output of the TM is then what is written on the tape (or just the very last symbol written, right under the head).

Here is a simple example of a Turing Machine, with the head starting in the state q_0 on the left of the tape:

$$q_0, 0 \longrightarrow q_0, 0, \triangleright, \quad (1.1)$$

$$q_1, 0 \longrightarrow q_1, 0, \triangleright, \quad (1.2)$$

$$q_0, 1 \longrightarrow q_1, 1, \triangleright, \quad (1.3)$$

$$q_1, 1 \longrightarrow q_0, 1, \triangleright, \quad (1.4)$$

$$q_0, \sqcup \longrightarrow \text{HALT}, 0, \cdot, \quad (1.5)$$

$$q_1, \sqcup \longrightarrow \text{HALT}, 1, \cdot. \quad (1.6)$$

The first two rules (1.1)-(1.2) say that if the head reads a zero, it should keep its state and move to the right. Rules (1.3)-(1.4) apply when the head reads a 1. The internal state of the machine then flips ($q_0 \leftrightarrow q_1$), and the head again moves to the right. Finally, when the head reads \sqcup , it assumes it has reached the end of the string, writes down 0 or 1 according to its internal state, and halts. This TM computes the parity of the string such as

$$\underbrace{10010100110}_n \sqcup \sqcup \sqcup \sqcup \sqcup \dots \quad (1.7)$$

initially written on the tape. The *input string size* n is the number of symbols on the tape after which the rest of the tape is blank. Observe that this TM takes at most n steps before it halts on any input string of size n .

When a TM halts, it produces an output. In the above example, the output was the parity of the input string. However, a general TM does not have to halt on each possible input, leaving its output undefined. A function on a set S whose domain is not the whole S is called a *partial* function. A TM is thus computing some partial function $\phi_M : S \rightarrow S$ on S , the set of strings over the alphabet $\{0, 1, \sqcup\}$. Taking it from the other end, does a TM computing a partial function F exist for any $F : S \rightarrow S$? There is an uncountable

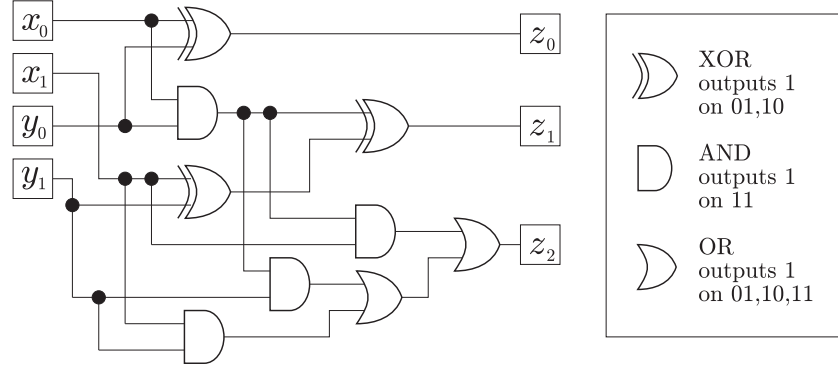


Figure 1-2: A binary addition circuit for two two-bit numbers, $z_2z_1z_0 = x_1x_0 \oplus y_1y_0$.

number of such functions, while the set of possible Turing Machines is countable. This means not all functions can be computable, motivating the following definition: A partial function $F : S \rightarrow S$ is *computable*, if there exists a Turing machine M such that $\phi_M = F$. An example of an uncomputable function is the *halting problem*: Does a given TM halt on all inputs? [91] A hypothetical halting problem-solving TM would then have to take the description of the TM whose halting properties it wants to determine as input.

On the other hand, for functions that are computable, it is interesting to ask how hard it is to compute them. For this, let me first ask how long does a specific Turing Machine run before it stops. This question must be related to the input string size n . The *running time* $T(n)$ of a TM is then defined as the largest number of steps it takes for the TM to stop on any input string of size n . There can be several different Turing Machines computing the same function, corresponding to different approaches to computing F . Some of them can be more efficient than others, with a better running time $T(n)$. An efficient TM is one whose running time $T(n)$ is small. The most common notion of an *efficient* TM is one whose running time scales only like a polynomial in the input size n (i.e. like $T(n) \propto n^3$ and not like $T(n) \propto 2^n$), and I adapt it in this thesis. On the other hand, functions F can be classified into *complexity classes* depending on what would the running time be for the best possible Turing Machine. I return to this topic in Section 1.1.2.

The notion of an *algorithm*, a prescription of what to do to compute something, is dependent on the computational model one uses. Turing himself defined what an algorithm is by his *Turing thesis*:

“Any algorithm can be realized by a Turing machine.”

Suprisingly, the alternative models of computation proposed in the early days of computer science turned out to be equivalent to the TM model [91, 56]. Today’s computers do not have long tapes with read/write heads inside them. The conventional *circuit model* of digital computers (see Figure 1-2), equivalent to the TM model, consists of wires carrying bits of information and logical gates like

x	NOT x
0	1
1	0

(1.8)

and two-bit gates

x	y	$x \text{ AND } y$	$x \text{ OR } y$	$x \text{ NAND } y$	$x \text{ XOR } y$
0	0	0	0	1	0
0	1	0	1	1	1
1	0	0	1	1	1
1	1	1	1	0	0

(1.9)

that are applied to them. What is the smallest set of gates one needs to simulate a Turing Machine (compute a computable function F)? It turns out that the NAND gate is *universal* in the sense that any circuit can be transformed into another built only from wires and NAND gates.

The single-bit gate NOT is reversible, as its input can be recovered from its output, while the gates AND, OR, XOR and NAND are all non-reversible. It turns out that reversible circuits are as powerful as their irreversible counterparts. The reversible 3-bit Toffoli gate, which flips the third bit if the first two bits are both 1, is universal [97]. This is important, because quantum computation is a generalization of classical reversible computation (see Section 1.2.1).

In the 1930's, analog computers like the *differential analyzer* for solving differential equations had their programs built in, as their components had to be physically rearranged for solving different problems. Similarly, for the TM mentioned so far, a different TM was required for computing each partial function F . However, Turing made a big breakthrough, paving the way for digital computation, by proving that it is possible to have a single *universal* TM, which can produce the output of any other TM. This is done by having its 'program' (the description of which particular computation to perform) given as part of the input. The universal TM then processes the 'data' part of the tape according to the 'program'. It is interesting to look for a simplest possible universal TM, and people have dedicated quite some effort to it [58, 112]. In Chapter 5, I will return to the idea of a classical universal reversible TM, and bring it to the world of quantum computing. I then investigate simple universal translationally and time-independent Hamiltonians which could be used to simulate any quantum computation in a 1D system.

1.1.2 Classical Complexity

Classical complexity examines how hard it is to compute a function F on strings S , or how hard it is to answer yes/no questions (decision problems) about strings S . A simple yes/no question about a string is: "Is the input string S a palindrome²?"

For comparing the hardness of problems, it is important to have a notion of reducibility. A *reduction* is a way of converting one problem to another in such a way that the solution to the second problem can be used to solve the first problem [91]. That multiplication can be reduced to addition, using addition as a subroutine polynomially many times, is a simple example of a polynomial reduction. I will frequently use reductions in this thesis, building on previous results about hardness of certain problems.

The complexity of computing a function or a decision problem can be classified by

- reducing the problem into another that is already classified,

²A palindrome is a string that reads the same backwards as it does normally, such as "livedondecalfacednodevil".

- showing that an already classified problem reduces to the one at hand,
- showing a lower bound $l(n) \leq T(n)$,
- finding an algorithm for it that runs in some $T(n)$.

To prove a lower bound $l(n)$ on the complexity of computing a function F , one must show that no algorithm computing F can possibly have running time less than $l(n)$. On the other hand, finding efficient algorithms is hard. We can solve many problems efficiently, but there are plenty others for which today's best known algorithm scales exponentially with the problem size (see Section 1.1.3). However, we do not know what tomorrow's algorithm designers will come up with. The best situation comes when a problem has a lower bound on its complexity, and an optimal algorithm that matches it. A trivial example is the problem of computing the *parity* of a string. The Turing Machine described by (1.1) is in fact optimal, matching the lower bound of n steps for a string of length n . A more interesting example is the *Ordered search* problem, where one needs to determine whether a specific number appears in an ordered list of n numbers. The list of numbers can be written on the one-dimensional tape of a Turing Machine which in the worst case would have to go through it all. On the other hand, I can view this problem differently. Consider the list of numbers to be hidden in a black box which the TM can query by writing a number m on a separate 'query' tape. Upon a query, the black box subroutine writes the m -th number from the list on the tape. The TM can then compare this number to the one it searches for, do some computation, and run another query. The complexity of the algorithm is then determined by the number of queries to the black box the TM needs to ask before it determines whether a number is in the black box list or not. The complexity of such black-box problems is called *query complexity*. The best classical algorithm for ordered search matches the corresponding lower bound, requiring to query $\log_2 n$ numbers in the list. In Section 1.2.3, I look at *Unstructured search*, another problem for which we have both a lower bound on the query complexity and an optimal algorithm, in the classical and in the quantum case as well.

I now introduce a few important classical complexity classes, all related to *decision problems*, which are simply yes/no questions.

Definition 1 (Complexity class P). *The class of decision problems solvable in polynomial time by a Turing machine.*

This is the basic class of *efficiently solvable* problems. It contains easy problems like determining whether two numbers are coprime, but also the problem of determining whether a number is prime [3]. On the other hand, it is not known whether it contains *factoring*, the problem of determining whether a number p has a divisor d in some range $a \leq d \leq b$. Today's best known classical algorithm for factoring has runtime $2^{\frac{n}{3}}$, where n is the number of bits of p . On the other hand, in 1994, Shor [90] discovered a breakthrough *quantum* algorithm with running time $T(n) \propto n^3$, whose use of the Fourier transform resulted in many of today's quantum algorithms. I refer the reader to Section 1.2.3 for more details about efficient quantum algorithms.

All the problems in P are contained in another class, NP, described by Scott Aaronson in the Complexity ZOO [1] as 'the class of dashed hopes and idle dreams'. Determining whether NP is in fact bigger than P or equal to it is a long standing open question in Computer Science with a million dollar prize from the Clay foundation awaiting the solver. The definition of NP is

Definition 2 (Complexity class NP (Non-deterministic, Polynomial time)). *NP is the class of decision problems such that, if the answer is ‘yes’, then there is a proof of this fact, of length polynomial in the size of the input, that can be verified by a deterministic polynomial-time algorithm. On the other hand, if the answer is ‘no’, then the algorithm must declare invalid any purported proof that the answer is ‘yes’.*

It is possible to expand the original Turing Machine model by allowing the TM to use an additional random generator, making it a *probabilistic Turing Machine*. This can be done by giving the TM access to a second tape with a random string written on it. The question whether this model is equivalent to the original, deterministic TM model is open, but believed to be true. The probabilistic Turing Machine model gives rise to the class BPP:

Definition 3 (Complexity class BPP (Bounded-error, Probabilistic, Polynomial time)). *The class of decision problems solvable in polynomial time by a randomized Turing machine.*

Great progress has been made in randomized algorithms recently, and it is a challenge for researchers to derandomize them. I give an example of a randomized algorithm in Section 1.1.3 (the WALKSAT algorithm for solving 3-SAT). Moreover, quantum computing, the topic of this thesis, is inherently randomized.

Analogously to generalizing P to BPP, the class NP motivates the class MA. Consider an interactive protocol between a powerful wizard Merlin and a logic, rational verifier Arthur, whose role is to verify purported proofs prepared by Merlin. The difference from NP is that Arthur can use a randomized algorithm, rather than verify the proof deterministically³:

Definition 4 (Complexity class MA (Merlin-Arthur)). *The class of decision problems solvable by a Merlin-Arthur protocol, which goes as follows. Merlin, who has unbounded computational resources, sends Arthur a polynomial-size purported proof that the answer to the problem is yes. Arthur must verify the proof in BPP (i.e. probabilistic polynomial-time), so that if the answer to the decision problem is*

1. *‘yes’, then there exists a proof which Arthur accepts with probability at least $2/3$.*
2. *‘no’, then Arthur accepts any “proof” with probability at most $1/3$.*

An alternative definition requires that if the answer is ‘yes’, then there exists a proof such that Arthur accepts with certainty. However, the definitions with one-sided and two-sided error are equivalent [91]. Moreover, the probabilities $2/3$ and $1/3$ in the definition of MA can be equivalently substituted by $1 - \epsilon$ and ϵ with arbitrarily small ϵ .

I introduce the quantum analogues of these classes in Section 1.2.2.

1.1.3 Classical locally constrained problems

Whether deciding on the best class schedule, finding a good route to spend a day touring a city or pairing people up at a dance class, our lives are full of local optimization problems. Some of them are simple, such as picking a single box of cereal with the lowest price/taste ratio from the shelf in the supermarket, some of them harder, like deciding which lenses to pack into my camera bag for the next trip. The latter is a question concerning a cost function involving only a few elements at a time (how useful will a particular lens in my bag

³If all random algorithms can be made deterministic (BPP=P), MA equals NP [1].

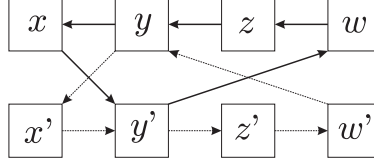


Figure 1-3: The implication graph for a 2-SAT instance $(x \vee y') \wedge (x' \vee y') \wedge (y \vee w) \wedge (y \vee z') \wedge (z \vee w')$. There is a directed loop containing both y and y' , so the instance is not satisfiable.

be, considering that I already bagged some others) and a global constraint (the size of the bag and my willingness to carry it all). In fact, this is an example of the *Knapsack* problem. Another locally constrained problem, *Satisfiability*, and its quantum analogue, are the focus of a large part of this thesis.

A boolean *clause* is an expression of the form $s_1 \vee s_2 \vee \dots \vee s_k$, where s_i are boolean variables or their negations. This clause is true for all bit strings $s_1 s_2 \dots s_k$ except for $\{00 \dots 0\}$. Consequently, the clause $s_1 \vee s'_2$, where s'_2 is the negation of bit s_2 , is true for all assignments $s_1 s_2$ except for $\{01\}$. Satisfiability is then defined as

Definition 5 (*k*-SATISFIABILITY, *k*-SAT). *Consider a collection $C = \{c_1, c_2, \dots, c_m\}$ of boolean clauses for a finite binary string $S = s_1 s_2 \dots s_n$, such that each c_i involves k of the string variables. Does a string S satisfying all the clauses in C exist?*

This is an instance of 2-SAT on four bits:

$$\begin{aligned}
 c_1 &= x \vee y', \\
 c_2 &= x' \vee y', \\
 c_3 &= y \vee w, \\
 c_4 &= y \vee z', \\
 c_5 &= z \vee w'.
 \end{aligned} \tag{1.10}$$

Here y' labels the negation of variable y , and the first clause, c_1 , is false if $x = 0$ and $y = 1$. The complexity of *k*-SAT varies with k . For 2-SAT, we have an algorithm with complexity scaling as $O(n^2)$. Draw a directed graph of implications between the bits and their negations as in Figure 1-3, where for example clause c_4 implies the connections $z \implies y$ and $y' \implies z'$. If a loop containing both a variable and its negation exists (it is easy to check this), the instance of 2-SAT is unsatisfiable. 2-SAT is thus in P. However, no known polynomial algorithm for 3-SAT exists. The Cook-Levin theorem says that 3-SAT is NP-complete. A problem Π is *complete* for a complexity class, when

- the problem Π is contained in the class, and
- every other problem in the class can be reduced to Π (solved by using the solution to Π as a subroutine polynomially many times).

It is straightforward to see that 3-SAT is in NP. The basic idea for the second direction of the NP-completeness proof is to encode the verification circuit in the definition of NP using a 3-SAT instance. It is instructive to see how this type of proof works in the classical case,

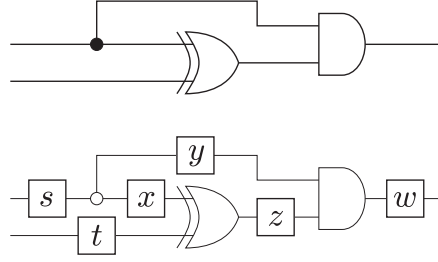


Figure 1-4: Illustration for the Cook-Levin theorem. Transforming a verifier circuit into a 3-SAT instance.

before I introduce its quantum analogue in Section 1.3.2. An instance x of a problem Π in NP is a yes/no question, denoted as $x \in (L_{yes} \cup L_{no})$, where L_{yes} is the set of instances with the answer ‘yes’. Each problem instance has a verifier circuit V associated with it. When V checks a purported proof that $x \in L_{yes}$, it accepts only a valid proof, never outputting 1 if the “proof” is wrong. This verifying procedure can be reduced to finding a satisfying assignment for a certain 3-SAT instance. Without loss of generality, circuit V consists of wires and gates involving at most 3 bits (two inputs, one output). Assign a boolean variable to each wire of the circuit as in Figure 1-4. The 3-SAT clauses then check the proper evaluation of the circuit. For example, the boolean clauses for y, z, w in Figure 1-4 (where w is the output of the AND gate on y and z) read

$$c_1 = y \vee z \vee w', \quad (1.11)$$

$$c_2 = y \vee z' \vee w', \quad (1.12)$$

$$c_3 = y' \vee z \vee w', \quad (1.13)$$

$$c_4 = y' \vee z' \vee w, \quad (1.14)$$

ruling out the assignments $yzw \in \{001, 011, 101, 110\}$. The last ingredient in the proof is a clause verifying that the circuit outputs 1:

$$c_{out} = w. \quad (1.15)$$

If a proof of $x \in L_{yes}$ which the circuit V accepts exists, an assignments of the bits corresponding to the computation satisfying the 3-SAT instance exists as well. On the other hand, if there is no “proof” that V accepts, there is no satisfying assignment to all the clauses in the 3-SAT instance. Therefore, if one could solve 3-SAT, one could solve any problem in NP. This is also true for k -SAT with $k \geq 3$.

To illustrate the idea of randomized algorithms mentioned in Section 1.1.2, let me present two randomized algorithms for the SAT problem. The procedures are given in Figure 1-5. The greedy algorithm GSAT [86] starts with a random bit string and proceeds to flip bits which give the greatest decrease in the number of broken clauses. It works great for 2-SAT, but not for 3-SAT, as it gets stuck in local minima. The second algorithm, WALKSAT [85], goes around this problem by randomly choosing between performing a step of GSAT or flipping a random bit in some unsatisfied clause. This randomized algorithm performs surprisingly well on many real-life instances of 3-SAT.

GSAT

```

for  $i = 1$  to  $max\_tries$ 
   $S$  = a randomly generated string
  for  $j = 1$  to  $max\_flips$ 
    if  $S$  satisfies  $C$ , return  $S$ 
    else flip any variable in  $S$  that
      results in the greatest decrease in
      the number of unsatisfied clauses
  end
end
return "No satisfying assignment found."

```

WALKSAT

With probability p :
 follow the standard GSAT
 scheme, making the best
 possible greedy move.

With probability $1 - p$:
 instead of making a GSAT step,
 randomly pick a variable
 occurring in some unsatisfied
 clause and flip it.

Figure 1-5: The GSAT and WALKSAT algorithms.

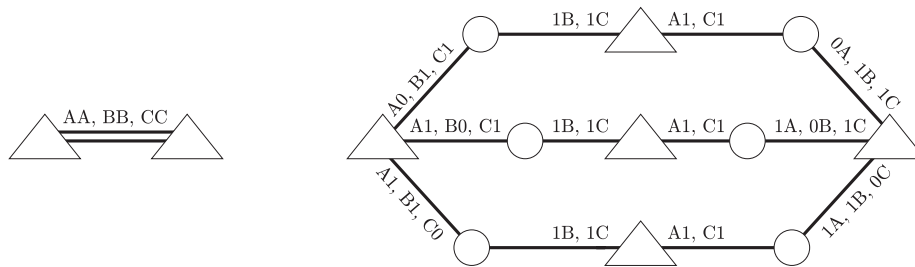


Figure 1-6: A reduction of Graph 3-colorability to 2-Satisfiability for trit-bit pairs, also called (3,2)-SAT. The trits with possible values $\{A,B,C\}$ are denoted by triangles, while the bits are denoted by circles. Each clause is represented by the pair of symbols which it rules out. For Graph 3-colorability, the excluded assignments are AA, BB and CC, ruling out the same color for two neighboring vertices.

SAT for non-bit strings

SAT can be generalized to non-bit strings. The simplest variant of this is (3,2)-SAT, with each clause involving a trit (a three-valued letter) and a bit. Although it doesn't seem much more complicated than regular 2-SAT, already this variant is NP-complete. It belongs to NP, because it can be reduced to 3-SAT by encoding the trit into two bits. On the other hand, it is also NP-hard, as the NP-complete *Graph 3-colorability* problem can be reduced to (3,2)-SAT as depicted in Figure 1-6.

MAX- k -SAT

The last variant of Satisfiability I introduce is MAX- k -SAT. In regular k -SAT, the question is whether an assignment satisfying all clauses exists or not. Sometimes, showing that no such assignment exists can be easy. However, the MAX- k -SAT problem poses a harder question: does an assignment breaking fewer than a clauses exist? The clauses in MAX- k -SAT define a cost function, and solving MAX- k -SAT for different a would determine its minimum value. Although regular 2-SAT is in P, MAX-2-SAT is NP-complete, so the existence of a polynomial algorithm for it is highly unlikely.

Classical	bit strings	general strings
k -SAT	$k = 2$: in P $k \geq 3$: NP-complete	(3, 2)-SAT : NP-complete
MAX- k -SAT	$k \geq 2$: NP-complete	

Table 1.1: Known complexity for classical satisfiability problems. (3, 2)-SAT in the general strings column is the problem where each clause involves one trit (letter with three values) and one bit.

In Table 1.1, I summarize the known results about classical satisfiability problems. The goal of Chapter 4 is to classify the complexity of the quantum analogues of these. I refer the reader interested in NP-complete problems and other variants of SAT to the book of Garey and Johnson [42].

1.2 Quantum Computing

Up until now, I have considered the model of computing governed by the laws of classical physics. One needs to ask what is the computing power of systems governed by quantum mechanics. Alternatively, how hard is it to simulate quantum-mechanical systems with classical computers? Already Feynman stated these questions in [38]. I assume the reader is familiar with quantum mechanics, so the following exposition will be brief, introducing only those topics which I later use.

Classical data can be easily measured, copied and erased. If someone gave me a binary string like 1001001, I could easily read it (measure it) and make many copies for myself. However, the world is quantum-mechanical. The state space of the simplest quantum system – a spin- $\frac{1}{2}$ particle, a *qubit*, is much bigger than the two possible values $\{0, 1\}$ of a classical bit. The first postulate of quantum mechanics is

- Associated to any isolated physical system is a complex vector space with inner product (i.e. a Hilbert space) known as the *state space* of the system. The system is completely described by its *state vector*, which is a unit vector in the system's state space.

The state space of a qubit is a 2-dimensional complex vector space $\mathcal{H} = \mathbb{C}^2$. The state of a qubit is given by

$$|\psi\rangle = c_0 \begin{bmatrix} 1 \\ 0 \end{bmatrix} + c_1 \begin{bmatrix} 0 \\ 1 \end{bmatrix} = c_0 |0\rangle + c_1 |1\rangle, \quad (1.16)$$

where $|0\rangle$ and $|1\rangle$ are basis vectors in the usual Dirac notation and c_0, c_1 are complex numbers for which $|c_0|^2 + |c_1|^2 = 1$.

How can a quantum computer be built from a system of qubits? In a classical computer, one can apply irreversible gates like NAND or reversible operations like the Toffoli gate. Besides measurements which I discuss below, the available operations for a quantum system are unitary transformations. Quantum mechanical systems obey the time-dependent Schrödinger equation. The next postulate of quantum mechanics is

- The time evolution of the state of a quantum system is described by the time-dependent *Schrödinger equation*

$$i\hbar \frac{d}{dt} |\psi(t)\rangle = H(t) |\psi(t)\rangle, \quad (1.17)$$

where $H(t)$ is a Hermitian operator called the Hamiltonian. Note that the value of the constant \hbar can be absorbed into the definition of H .

Because the Hamiltonian is Hermitian, the time evolved state $|\psi(t)\rangle$ and the initial state $|\psi(0)\rangle$ are related by a unitary transformation $U_{t,0}$:

$$|\psi(t)\rangle = U_{t,0} |\psi(0)\rangle. \quad (1.18)$$

When the Hamiltonian H is time-independent, this unitary transformation is

$$U_{t,0} = e^{-iHt}. \quad (1.19)$$

Unitarity of U means that $U_{t,0}^\dagger = U_{t,0}^{-1}$, and implies that time evolution according to the Schrödinger equation is reversible. The transformations one can perform on a quantum computer are thus unitary.

The final ingredient for quantum computation are measurements, used to read out the result. They could be used throughout the computation, but in the next Section I argue why it is sufficient to use measurements only at its end. Quantum measurements work in the following fashion:

- Consider a quantum observable \hat{M} , a Hermitian operator acting on the state space of the system being measured, expressed in terms of its eigenvectors as

$$\hat{M} = \sum_i \lambda_i |\phi_i^M\rangle \langle \phi_i^M|. \quad (1.20)$$

If the state of the quantum system is $|\psi\rangle$ immediately before the measurement then the probability to obtain the result λ_i is given by

$$p(i) = |\langle \phi_i^M | \psi \rangle|^2, \quad (1.21)$$

and the state of the system after the measurement is the corresponding eigenstate $|\phi_i^M\rangle$.

Imagine now that I want to determine what state the qubit (1.16) is in. As an example, let me choose to measure $\hat{\sigma}_z = |0\rangle\langle 0| - |1\rangle\langle 1|$, the z -component of the spin. I will get the result $+1$ with probability $|c_0|^2$, and the result -1 with probability $|c_1|^2$. It takes many copies of $|\psi\rangle$ and many measurements to determine c_0 and c_1 in general. Even if I made different measurements, distinguishing nonorthogonal states perfectly is impossible [73].

1.2.1 The Quantum Circuit model

Quantum circuits are a direct generalization of classical reversible circuits, using unitary transformations instead of classical reversible gates. A few basic single-qubit unitary gates,

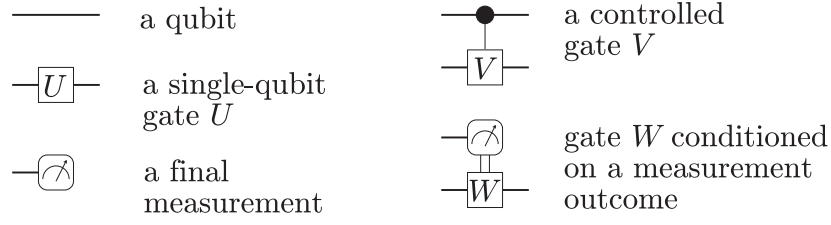


Figure 1-7: The components of a quantum circuit.

expressed as 2×2 matrices acting on column vectors in \mathbb{C}^2 are the Pauli matrices

$$\sigma_x = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad \sigma_y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \quad \sigma_z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \quad (1.22)$$

and the Hadamard gate H , the phase gate S and the $\frac{\pi}{8}$ gate T :

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \quad S = \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}, \quad T = \begin{bmatrix} e^{-i\frac{\pi}{8}} & 0 \\ 0 & e^{i\frac{\pi}{8}} \end{bmatrix}. \quad (1.23)$$

A general single qubit unitary gate can be expressed using the Pauli matrices as

$$U = e^{-i\frac{\phi}{2}\hat{n}\cdot\vec{\sigma}}, \quad (1.24)$$

where ϕ is a real number $\hat{n} = (n_x, n_y, n_z)$ is a unit vector and $\hat{n} \cdot \vec{\sigma} = n_x\sigma_x + n_y\sigma_y + n_z\sigma_z$

At least one two-qubit gate is required for quantum computation. The following two gates apply σ_x or σ_z to the target qubit, if the control qubit is in the state $|1\rangle$:

$$\text{CNOT} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \quad \text{C-Z} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}. \quad (1.25)$$

The elements of a quantum circuit are usually depicted as in Figure 1-7. At the end of the computation, the output of a quantum circuit is read out by performing a few final measurements. Although I could perform measurements throughout the computation and apply gates conditioned on the outcomes, it is possible to defer measurements until the very end of the computation by using controlled gates, as shown in Figure 1-8.

Universality

The NAND gate is universal for classical circuits and the Toffoli gate is universal for classical reversible circuits as mentioned in Section 1.1.1. On the other hand, can every unitary transformation on n qubits be decomposed into many applications of a small finite set of gates? Indeed, universal gate sets for quantum computing exist in the sense that they can approximate any unitary operation arbitrarily well. A universal gate set must be dense in the group $SU(n)$, and the approximations must converge fast. I point the interested reader to Appendix 3 of [73] for a thorough discussion of the Solovay-Kitaev theorem concerning

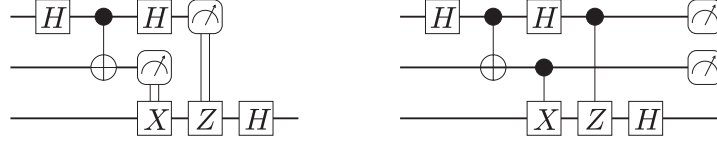


Figure 1-8: A quantum circuit with measurements performed during the computation and its version with measurements deferred until the end of the computation.

the convergence rate. Some well known gate sets universal for quantum computing are:

- single qubit gates, CNOT [73],
- single qubit gates, C–Z [73],
- H , S , T , CNOT [73].

Another universal gate set uses the quantum Toffoli gate:

- H , Toffoli [88]

It is intriguing, because it uses only real matrices. It is possible to do quantum mechanics with real numbers, as superpositions and minus signs, not imaginary numbers are essential. Consider the state (1.16) of a qubit, and add to the system another qubit labeled r . This qubit can be used to work around the requirement of imaginary amplitudes.

$$|\Psi\rangle = \text{Re}(c_0) |0\rangle |0\rangle_r + \text{Im}(c_0) |0\rangle |1\rangle_r \quad (1.26)$$

$$+ \text{Re}(c_1) |1\rangle |0\rangle_r + \text{Im}(c_1) |1\rangle |1\rangle_r \quad (1.27)$$

is a state of the larger system with only real amplitudes, and has all the information contained in the original state $|\psi\rangle$ (1.16). Arbitrary unitary transformations on $|\psi\rangle$ can be then mapped onto real unitary transformations on the larger system [88]. Finally, a single two-qubit gate

- W ,

the controlled $\frac{\pi}{2}$ -rotation about the y -axis is universal for quantum computation [88]. In Chapter 5, I use this single universal gate in my Hamiltonian Quantum Cellular Automaton constructions.

1.2.2 Quantum Complexity

In recent years, quantum complexity classes have been defined and studied in an attempt to understand the capacity and limitations of quantum computers and quantum algorithms and their relation to classical complexity classes. The basic quantum complexity class is BQP, the analogue of BPP.

Definition 6 (Complexity class BQP (Bounded error, Quantum, Polynomial time). *Consider a decision problem $\Pi = (L_{yes} \cup L_{no})$, where L_{yes} is the set of all instances x of the problem Π with the answer ‘yes’. The problem Π is in BQP, if there exists a uniform family*

of circuits⁴ U , with the following properties. For a problem instance x with the answer ‘yes’ ($x \in L_{yes}$), the probability of the circuit U outputting ‘yes’ is greater than $\frac{2}{3}$. On the other hand, when $x \in L_{no}$, the probability of the circuit outputting ‘yes’ is less than $\frac{1}{3}$.

In short, BQP is the class of problems one can solve efficiently using a quantum computer. The complexity class QMA, also known as BQNP, studied and defined in [57] and [56], is the quantum analogue of the classical complexity class MA, the probabilistic setting of NP.

Definition 7 (Complexity class QMA). *A decision problem $\Pi = (L_{yes} \cup L_{no})$ of size n is in the class QMA if there exists a polynomial time quantum verifier circuit V such that for every instance x of the problem*

1. $\forall x \in L_{yes} : \text{there exists a witness state } |\varphi(x)\rangle \text{ such that the computation } V|x\rangle \otimes |\varphi(x)\rangle \otimes |0 \dots 0\rangle_{\text{ancilla}} \text{ yields the answer 1 with probability at least } p ;$
2. $\forall x \in L_{no} : \text{for any witness state } |\varphi(x)\rangle, \text{ the computation } V|x\rangle \otimes |\varphi(x)\rangle \otimes |0 \dots 0\rangle_{\text{ancilla}} \text{ yields the answer 1 with probability at most } p - \epsilon,$

where $p > 0$ and $\epsilon = \Omega(1/\text{poly}(n))$.

The first known QMA-complete problem, local Hamiltonian, is a quantum analogue of classical MAX-k-SAT. Building on ideas that go back to Feynman [38], Kitaev [56] has shown that 5-local Hamiltonian is QMA complete and I present this result in Section 1.3.2. A special case of Local Hamiltonian, Quantum k -SAT, is the topic of Chapter 4. There I show that several of its variants are complete for QMA₁, the class QMA with one-sided error:

Definition 8 (Complexity class QMA₁). *The class QMA₁ is the class QMA with single sided error, i.e. with $p = 1$ in the above definition.*

In the classical case, MA with one-sided error is equivalent to regular MA. It is not known whether QMA=QMA₁.

1.2.3 The Strengths and Limits of Quantum Computing

There are quite a few problems on which quantum computers would algorithmically outperform classical ones. In quantum computing, we can use resources unknown in classical computing: superpositions, interference and entanglement. Much has been discovered since Shor’s factoring algorithm in 1994 [90]. Stephen Jordan has a thorough and up-to-date summary of the algorithmic successes of quantum computing in his thesis [51]. However, the speedups are rarely exponential. Not only that, solving problems such as *Parity*⁵ cannot be sped up on a quantum computer at all [35]. I refer the reader interested in the limitations of quantum computers to [8, 2, 10, 15, 17].

Algorithm designers bound the complexity of a problem from above by finding better algorithms, while lower bound enthusiasts limit it from below. The quantum variant of the black-box problem *Unstructured search* is an example of a problem for which these two approaches meet. The lower bound on its query complexity is known, and an optimal quantum algorithm, Grover’s search, exists.

⁴Here *uniform family of circuits* means that there is one circuit for each problem size (and not for each problem instance), with a number of gates polynomial in the problem size.

⁵*Parity*: determine the sum modulo 2 of n bits.

Definition 9 (Unstructured Search). *Consider a quantum black-box (oracle)*

$$\hat{O}|s\rangle = \mathbb{I} - 2|w\rangle\langle w|. \quad (1.28)$$

which flips the phase of a single, unknown to you, canonical basis quantum state $|w\rangle$, and acts trivially on any other canonical basis state. Find the state $|w\rangle$.

Looking at a single state of n bits unsuccessfully does not give me any information about where to look next, making this problem unstructured. Classically, the required number of queries scales like the number of possible n -bit strings, 2^n , and is matched by the simplest randomized algorithm one can think of—check the strings randomly until you find the marked one. The query complexity of a quantum algorithm for unstructured search is the number of necessary calls to \hat{O} in the worst case. In 1996, Grover discovered his famous algorithm requiring only $\sqrt{2^n}$ queries, getting a square root speedup over the best possible classical algorithm. Also around that time, Bennett, Bernstein, Brassard and Vazirani found a matching lower bound $l(n) \propto \sqrt{2^n}$ [17]. Therefore, Grover’s algorithm is the best possible quantum algorithm for unstructured search. Unstructured search can be also recast in a continuous setting, where quantum computation is done by continuous-time Hamiltonian evolution instead of application of discrete unitary gates. Instead of a quantum black box, in [37] Farhi and Gutmann use an oracle Hamiltonian

$$H_w = -E|w\rangle\langle w|, \quad (1.29)$$

where E determines an energy scale. This oracle Hamiltonian is always on in the system. What one can do to find $|w\rangle$ in this Hamiltonian setting is to prepare the system in some initial state, and let the system evolve according to a Hamiltonian

$$H(t) = H_D(t) + H_w, \quad (1.30)$$

where $H_D(t)$ is an additional arbitrary time-dependent Hamiltonian. If measuring the system after some time produces $|w\rangle$, the algorithm is successful. Farhi and Gutmann show that the required running time for any such Hamiltonian algorithm (any $H_D(t)$ one uses) necessarily scales like $\sqrt{2^n}/E$. In Chapter 2, I use this lower bound to show what not to do when designing quantum adiabatic algorithms.

1.3 Local Hamiltonians in Quantum Computation

Finding the properties of relatively simple quantum systems is the topic of much condensed matter physics. Even though all the interactions are local, involving only a few neighboring particles, the resulting possibilities are amazingly rich. In this thesis I focus on two properties of systems governed by Local Hamiltonians. First, a (quantum) computer can be constructed from a system governed by a simple Hamiltonian. Second, the ground states of local Hamiltonians can be hard to find, as they can encode solutions to interesting optimization problems.

Hamiltonian time evolution according to the Schrödinger equation is the underlying principle of all quantum computing, producing the unitary operations required the quantum circuit model. Controlling the Hamiltonian of a quantum system means changing the

interaction of the particles with each other and their surroundings. For example, in addition to the natural interaction of the particles in a quantum system, we can apply magnetic field to spins on a lattice, laser pulses to ions in an ion trap or electric fields to electrons in quantum dots. External control can allow us to perform the desired unitary transformations in the system, making it into a computer. Feynman [38] thought of how to make a Hamiltonian Computer for simulating other quantum systems long before people invented quantum circuits. I introduce his idea in Section 1.3.1, and build on it throughout this thesis, especially in Chapters 4 and 5.

Besides investigating the time evolution of a quantum system, analyzing the spectrum of its eigenvalues is another source of interesting problems. It is long known that finding ground states of certain Hamiltonians is hard. If we could find ground states of a system governed by the Ising model 3D, we could solve NP-complete problems [13]. Finding minima of locally constrained optimization problems can be recast into finding the ground state energy of a local Hamiltonian. I am interested in showing how hard it is to find it for some Hamiltonians. In one case, in Chapter 3, I give a numerical method for finding the ground state energy of the translationally invariant Ising model in transverse field on a Bethe lattice, based on the Matrix Product State ansatz. On the other hand, building on Kitaev's Local Hamiltonian problem I introduce in Section 1.3.2, I classify the complexity of finding the ground state of several Hamiltonians in Chapter 4.

1.3.1 Feynman's Hamiltonian Quantum Computer

Already in 1985, Feynman [38] proposed the following interpretation/implementation of a computer using a quantum-mechanical systems. What he originally had in mind was to simulate a sequence of reversible classical gates in a quantum mechanical system. However, his construction applies equally well to implementing a quantum circuit. Take a sequence of L unitary transformations U_t on n work qubits

$$U = U_L U_{L-1} \dots U_2 U_1. \quad (1.31)$$

Consider now a system with two registers,

$$\mathcal{H} = \mathcal{H}_{work} \otimes \mathcal{H}_{clock}. \quad (1.32)$$

The first one holds the n work qubits, which I label q_1, \dots, q_n throughout the thesis, while the second register holds a pointer particle hopping on a line $0, 1, \dots, L$, serving as a *clock* for the computation. I can realize the clock register using $L + 1$ qubits labeled c_0, \dots, c_L , and initialize it in the state

$$|0\rangle_c = |1_{c_0} 0_{c_1} 0_{c_2} \dots 0_{c_L}\rangle. \quad (1.33)$$

The state of the clock register corresponding to “time” t is

$$|t\rangle_c = |0_{c_0} \dots 0_{c_{t-1}} 1_{c_t} 0_{c_{t+1}} \dots 0_{c_L}\rangle. \quad (1.34)$$

A simple time-independent Hamiltonian can facilitate the evaluation of U on the work register of this system. As the pointer particle hops on the line $0, 1, \dots, L$, performing a quantum walk, the corresponding unitary operations are applied (or uncomputed) on the

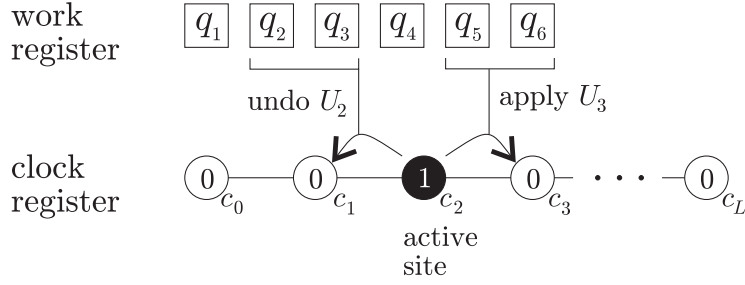


Figure 1-9: Feynman's computer with a hopping pointer particle (a single spin up). When the active site moves from c_2 to c_3 , gate U_3 gets applied to work qubits q_5 and q_6 . When the active spin jumps back from c_2 to c_1 , gate U_2 is uncomputed on qubits q_2 and q_3 .

work qubits as depicted in Figure 1-9. The Hamiltonian with the desired dynamics is

$$H_{Fnmn} = \sum_{t=1}^L U_t \otimes \sigma_{(t)}^+ \sigma_{(t-1)}^- + U_t^\dagger \otimes \sigma_{(t)}^- \sigma_{(t-1)}^-, \quad (1.35)$$

where $\sigma_{(t)}^+ = |1\rangle \langle 0|_{c_t}$ and $\sigma_{(t)}^- = |0\rangle \langle 1|_{c_t}$ are the raising and lowering operators for the clock register spin c_t and U_t acts on the corresponding work qubits. When the system is initialized in the state

$$|\Psi_0\rangle = |\psi_0\rangle \otimes |0\rangle_c = |\psi_0\rangle \otimes |10 \dots 0\rangle_c, \quad (1.36)$$

where $|\psi_0\rangle$ is some initial state of the work qubits, time evolution with H_{Fnmn} according to the Schrödinger equation brings the state into a superposition of states

$$|\Psi_t\rangle = |\psi_t\rangle \otimes |t\rangle_c, \quad (1.37)$$

where $|\psi_t\rangle$ is the state of the work qubits after the first t gates of the circuit:

$$|\psi_t\rangle = U_t U_{t-1} \dots U_1 |\psi_0\rangle. \quad (1.38)$$

After some time τ , when I measure the clock register of $|\Psi\rangle$ and obtain L , the state of the system after the measurement becomes $|\psi_L\rangle \otimes |L\rangle_c$. The work register of this state contains the desired output of the circuit U .

I can obtain $|\psi_L\rangle$ with high probability using a slightly modified system. First, I pad the sequence U (1.31) with $2L$ extra identity operations

$$U' = \mathbb{I}_{3L} \mathbb{I}_{3L-1} \dots \mathbb{I}_{L+1} U, \quad (1.39)$$

and then expand the clock register to $3L + 1$ qubits accordingly. All the states $|\psi_t\rangle$ (1.38) for $t \geq L$ now are the same, equal to the desired output state $|\psi_L\rangle$. Starting from (1.36), I let the system evolve for a time τ chosen uniformly at random between 0 and $O(L \log L)$. Because the quantum walk on a line is rapidly mixing, the probability to measure a state with the clock register $|t \geq L\rangle_c$ at time τ is close to $\frac{2}{3}$, as shown in Appendix C. This yields the desired state $|\psi_L\rangle$ of the work qubits. If a measurement of the clock register results in

$|\psi_t\rangle \otimes |t < L\rangle_c$, repeat the experiment.

The clock register representation of the states $|t\rangle_c$ given in (1.34) is only one of the possibilities. I describe a different one in the next Section, and several others in Chapter 4. Considering this, Feynman's Hamiltonian can be rewritten in a general form, encompassing different implementations of $|t\rangle_c$, as

$$H_{Fnmn} = \sum_{t=1}^T H_F^t, \quad (1.40)$$

$$H_F^t = U_t \otimes X_{t,t-1} + (U_t \otimes X_{t,t-1})^\dagger, \quad (1.41)$$

where U_t acts on the corresponding work register qubits, and the operator $X_{t,t-1}$ acts on the clock register as

$$X_{t,t-1} = |t\rangle \langle t-1|_c, \quad (1.42)$$

increasing the time from $t-1$ to t . In Feynman's original formulation, this corresponds to moving the position of the single up spin to the right as $|\dots 0100 \dots\rangle_c \rightarrow |\dots 0010 \dots\rangle_c$, and the operator $X_{t,t-1}$ can be written as $X_{t,t-1} = |01\rangle \langle 10|_{c_{t-1}, c_t}$, where c_{t-1} and c_t are two of the clock register qubits. Analogously, the Hermitian conjugate of $X_{t,t-1}$ moves the time register backwards as

$$(X_{t,t-1})^\dagger = |t-1\rangle \langle t|_c. \quad (1.43)$$

This is my preferred formulation of Feynman's Hamiltonian throughout this thesis. It shows up in Chapter 2 where I use it to prove results about the universality of adiabatic quantum computing, while in Chapter 4 it serves a basic building block for my results about the complexity of Quantum Satisfiability. Finally, in Chapter 5 it inspires my Hamiltonian Quantum Cellular Automaton model.

1.3.2 Kitaev's Local Hamiltonian (LH) Problem

In this section I introduce Kitaev's *Local Hamiltonian* and sketch the proof that it is QMA complete. In [56] Kitaev defined the following promise problem:

Definition 10 (Local Hamiltonian). *Take a k -local Hamiltonian $H = \sum_{j=1}^r H_j$ acting on n qubits, composed of $r = \text{poly}(n)$ terms, each of those acting nontrivially on a constant number k of qubits. Given two numbers a, b , where $b > a$ and the separation $b - a$ is greater than $n^{-\alpha}$ for some constant α , determine whether*

1. *H has an eigenvalue not exceeding a , or*
2. *all eigenvalues of H are greater than b .*

The two numbers b and a are separated by at least an inverse polynomial in the number of qubits in the system, n . LH is a promise problem, therefore we know that either 1 or 2 are true. Kitaev proved that

Theorem 1 (LH is QMA-complete). *5-Local Hamiltonian is QMA-complete.*

One of the directions of the proof is to show that LH is in QMA. The verifier circuit U , saying whether Arthur accepts a given state or not, is constructed as follows. Add an ‘answer’ qubit to the system and for each term $H_j = \sum \lambda_s |\psi_s\rangle \langle \psi_s|$, construct a measurement operator

$$W_j : |\psi_s\rangle \otimes |0\rangle_{ans} \rightarrow |\psi_s\rangle \otimes \left(\sqrt{\lambda_s} |0\rangle + \sqrt{1 - \lambda_s} |1\rangle \right)_{ans}. \quad (1.44)$$

Taking a general state $|\eta\rangle = \sum_s y_s |\psi_s\rangle$, the probability the circuit accepts it (i.e. that I measure 1 on the output of U) is

$$p_j(1) = \langle \eta | \otimes \langle 0 |_{ans} W_j^\dagger (\mathbb{I} \otimes |1\rangle \langle 1|_{ans}) W_j | \eta \rangle \otimes |0\rangle_{ans} \quad (1.45)$$

$$= \sum_s (1 - \lambda_s) y_s^* y_s = 1 - \langle \eta | H_j | \eta \rangle. \quad (1.46)$$

Add now another ‘operator’ register to the system. The general circuit for checking all of the terms in the Hamiltonian combines the checking operators W_j into

$$W = \sum_{j=1}^r |j\rangle \langle j|_{op} \otimes W_j, \quad (1.47)$$

where r is the number of terms in the Hamiltonian. After applying W to the system in the state

$$\left(\sum_{j=1}^r |j\rangle_{op} \right) \otimes |\eta\rangle \otimes |0\rangle_{ans}, \quad (1.48)$$

one measures the answer qubit. The probability of getting the outcome 1 is then $p_1 = 1 - \frac{\langle \eta | H | \eta \rangle}{r}$. This is enough to show that LH is in QMA.

For the other direction, one needs to construct a corresponding instance of the Hamiltonian problem for every instance of a problem in QMA. Each of those has a verifier circuit U with the required accept/decline properties associated with it. Given a quantum circuit U , Kitaev showed how to construct H with the properties required in the definition of LH (Definition 10). If the circuit U accepts the state $|\psi\rangle$, this Hamiltonian’s ground state can be constructed from it, and the corresponding expectation value in this state will be smaller than a . On the other hand, if the circuit does not accept any state with high probability, the ground state of H will have energy higher than b .

Kitaev’s Hamiltonian

Consider a quantum system consisting of a work register with n qubits q_1, \dots, q_n , and a clock register.

$$\mathcal{H} = \mathcal{H}_{work} \otimes \mathcal{H}_{clock}. \quad (1.49)$$

Kitaev’s Hamiltonian is a sum of three terms,

$$H_{Kitaev} = H_{prop} + H_{out} + H_{input}. \quad (1.50)$$

The first term checks the proper progression of the circuit U . Taking the idea from

Feynman's Hamiltonian computer (see Section 1.3.1), Kitaev used (1.41) to construct a Hamiltonian whose ground state is the uniform superposition over the computation called the *history state*:

$$|\phi\rangle_{\text{history}} = \frac{1}{\sqrt{L+1}} \sum_{t=0}^L |\psi_t\rangle \otimes |t\rangle_c, \quad (1.51)$$

where

$$|\psi_t\rangle = U_t U_{t-1} \dots U_2 U_1 |\psi_0\rangle, \quad (1.52)$$

and $|\psi_0\rangle$ is any initial state of the work qubits. This Hamiltonian is

$$H_{\text{prop}} = \frac{1}{2} \sum_{t=1}^L H_{\text{prop}}^t, \quad (1.53)$$

where each term H_{prop}^t is a projector constructed from Feynman's Hamiltonian (1.41) as

$$H_{\text{prop}}^t = \mathbb{I} \otimes (P_{t-1} + P_t) - \underbrace{\left(U_t \otimes X_{t,t-1} + (U_t \otimes X_{t,t-1})^\dagger \right)}_{H_F^t}, \quad (1.54)$$

with $P_t = |t\rangle \langle t|_c$ the projector onto the state $|t\rangle_c$ of the clock register, and the time-increasing operator $X_{t,t-1}$ changing the clock from $t-1$ to t given by (1.42). The only states that are not obviously annihilated by H_{prop}^t live in the subspace of states with the clock register in the states $|t-1\rangle_c$ or $|t\rangle_c$. Within that subspace, the states with zero energy have the form

$$|\phi\rangle = |\alpha\rangle \otimes |t-1\rangle_c + U_t |\alpha\rangle \otimes |t\rangle_c, \quad (1.55)$$

making it easy to check that the history state (1.51) is annihilated by the non-negative Hamiltonian H_{prop} . It can be thought of as checking the correct propagation of the computation, as the expectation value of H_{prop} in any state that is not a history state is greater than zero. In fact, the second lowest eigenvalue can be bounded by $\lambda_1 \geq \frac{c_1}{L^2}$ for some constant c_1 (see Appendix A for a detailed analysis of the Hamiltonian (1.53)).

The last of the work register qubits, q_n , is the designated output qubit for the circuit U . The second term H_{out} in (1.50) adds an energy penalty to states whose q_n at time L is $|0\rangle$, meaning that the circuit did not accept them.

$$H_{\text{out}} = |0\rangle \langle 0|_{q_n} \otimes |L\rangle \langle L|_c. \quad (1.56)$$

This term ensures that H_{Kitaev} (1.50) can have a low eigenvalue only if the circuit U outputs 1 on some input state $|\psi_0\rangle$ with high probability.

The work register has n qubits q_1, \dots, q_n . Out of these, at time $t=0$, some are initialized in a purported proof state, while others are ancilla qubits required for the verification circuit. The last part of H_{Kitaev} ,

$$H_{\text{input}} = \sum_{k \in \text{ancilla}} |1\rangle \langle 1|_{q_k} \otimes |0\rangle \langle 0|_c, \quad (1.57)$$

checks the correct initialization of these ancilla qubits by adding an energy penalty to states whose ancillae at time $t = 0$ (when the clock register is in the state $|0\rangle_c$) are not in the state $|0\rangle$.

These three terms together make up Kitaev's Hamiltonian (1.50). Let me now sketch the highlights of the proof that (1.50) has the properties described in Definition 10. First, the expectation value of H_{Kitaev} in the history state (1.51) corresponding to an initial state $|\psi_{yes}\rangle$ that the circuit U accepts can be shown to be upper bounded by $\lambda_0^{yes} \leq \frac{\epsilon}{L}$ for some ϵ . Second, one needs to lower bound the lowest eigenvalue of H_{Kitaev} in the case the circuit U doesn't accept any state with high probability, and show that it is separated from λ_0^{yes} at least as an inverse polynomial in L . All three terms in (1.50) are non-negative. The proof uses a geometrical argument, showing that the angle between the null spaces of H_{prop} and $(H_{input} + H_{out})$ is lower bounded by θ , an inverse polynomial in L . The smallest eigenvalue of H_{Kitaev} can then be bounded from below by the smallest nonzero eigenvalue of H_{prop} multiplied by a geometric factor coming from θ . The resulting lower bound on the lowest eigenvalue of (1.50) is then $\lambda_0^{no} \geq c(1 - \sqrt{\epsilon})L^{-3}$. H_{Kitaev} thus has the desired properties.

The final step in the proof that 5-local Hamiltonian is QMA complete is to implement the clock register. I describe several clock register constructions in detail in Section 4.1.3. One of these is Kitaev's *domain wall* clock. Similarly to Feynman (1.34), he encoded his clock register in unary using $L + 1$ clock qubits c_t . However, the states $|t\rangle_c$ are now encoded as

$$|t\rangle_c = |1_{c_0}1_{c_1} \dots 1_{c_t}0_{c_{t+1}} \dots 0_{c_{L-1}}0_{c_L}\rangle \quad (1.58)$$

The Hilbert space of the $L + 1$ clock qubits c_0, \dots, c_L is much bigger than the space spanned by the proper clock states $|t\rangle_c$ given by (1.58), which I call the *legal clock subspace*, \mathcal{H}_{legal} . By adding a clock-checking Hamiltonian

$$H_{clock} = \sum_{t=0}^{L-1} |01\rangle\langle 01|_{c_t, c_{t+1}} \quad (1.59)$$

to H_{Kitaev} (1.50), one can ensure that the low-energy eigenstates of H are close to the subspace of states with a proper unary-encoded clock, i.e. whose clock register states are close to the legal clock subspace \mathcal{H}_{legal} . The final form of the Hamiltonian is then

$$H_{Kitaev5} = H_{prop} + H_{out} + H_{input} + H_{clock}. \quad (1.60)$$

In this clock implementation, P_t , the projector onto the state $|t\rangle$, is only a 2-local operator

$$P_t^c = |t\rangle\langle t|_c = |10\rangle\langle 10|_{c_t, c_{t+1}}, \quad (1.61)$$

with a special case $P_L^c = |1\rangle\langle 1|_{c_L}$, acting trivially on the rest of the qubits. On the other hand, the operator $X_{t,t-1}$ (1.42) in this encoding is 3-local (acting nontrivially on three clock qubits):

$$X_{t,t-1} = |t\rangle\langle t-1|_c = |110\rangle\langle 100|_{c_{t-1}, c_t, c_{t+1}}, \quad (1.62)$$

with a special case at the right end of the clock register,

$$X_{L,L-1} = |L\rangle\langle L-1|_c = |11\rangle\langle 10|_{c_{L-1}, c_L}. \quad (1.63)$$

Recalling that 2-qubit gates are universal for quantum computation, the terms in H_{prop} (1.53), and thus also H_{Kitaev} , become at most 5-local. It remains to show that in the case the circuit U rarely outputs 1, the lowest eigenvalue of (1.60) is also bounded from below by an inverse polynomial in n . This proof follows from the geometric arguments outlined above without much complication. Therefore, 5-local Hamiltonian is QMA-complete.

In Chapter 4, I return to this problem, focusing on a special case called Quantum Satisfiability, defined by Bravyi in [20]. When all the terms in H are projectors (non-negative), one can ask whether there exists a state with energy exactly 0, or whether the ground state energy of H is greater than b . I investigate the complexity of several variants of Q-SAT and utilize the techniques used in Kitaev's proof I just outlined.

Chapter 2

Adiabatic Quantum Computation

Farhi et al. [36] introduced Adiabatic Quantum Computing (AQC) as a method to find a minimum of a classical cost function by using a quantum system with a time-dependent, local Hamiltonian. I describe this model in Section 2.1, and then review the proof of Aharonov et al. [7], who showed that AQC is equivalent to the quantum circuit model, in the light a previously unpublished result of Seth Lloyd [63] in Section 2.2. This motivates my definition of a Hamiltonian Computer model in Section 2.2.1, which I later use to prove universality of the Quantum 3-SAT Hamiltonian in Section 4.5 and of two Hamiltonian Quantum Cellular Automata in Chapter 5.

There is hope that there may be combinatorial search problems, defined on n bits so that $N = 2^n$, where for certain “interesting” subsets of the instances the run time of the Quantum Adiabatic Algorithm grows sub-exponentially in n . A positive result of this kind would greatly expand the known power of quantum computers. At the same time it is worthwhile to understand the circumstances under which the algorithm is doomed to fail. Section 2.3, based on the paper [34],

How to make the Quantum Adiabatic Algorithm Fail

Edward Farhi, Jeffrey Goldstone, Sam Gutmann and Daniel Nagaj

The quantum adiabatic algorithm is a Hamiltonian based quantum algorithm designed to find the minimum of a classical cost function whose domain has size N . We show that poor choices for the Hamiltonian can guarantee that the algorithm will not find the minimum if the run time grows more slowly than square root of N . These poor choices are nonlocal and wash out any structure in the cost function to be minimized and the best that can be hoped for is Grover speedup. These failures tell us what not to do when designing quantum adiabatic algorithms.

contains two results about the necessity of local structure for successful AQC algorithms. First, in Section 2.3.1 I show a connection between a bad choice of the starting Hamiltonian in an adiabatic algorithm and an information theoretical lower bound on the running time of any quantum algorithm for unstructured search (see Section 1.2.3). Second, I show in Section 2.3.2 that knowing the spectrum of the final Hamiltonian is not enough to determine whether finding its ground state via an AQC algorithm is feasible. The local structure determining how the final spectrum arises is essential.

2.1 Introduction to AQC

Consider a classical optimization problem with cost function $h(z)$ on n bits, with $z = 0, \dots, 2^n - 1$. I can construct a Hamiltonian for n qubits, diagonal in the computational (z) basis, according to this cost function as

$$H_P = \sum_{z=0}^{N-1} h(z) |z\rangle. \quad (2.1)$$

The goal is now to find the ground state of H_P , which encodes the solution to the classical optimization problem with cost function $h(z)$. If I could prepare the ground state of H_P , a measurement of the spins in the z -basis would produce the desired optimal assignment z (or one of them, if there are several).

To prepare the ground state of H_P , a ‘beginning’ Hamiltonian H_B is introduced with a known and easy to construct ground state $|g_B\rangle$. The quantum computer is a system governed by the time dependent Hamiltonian

$$H(t) = \left(1 - \frac{t}{T}\right) H_B + \left(\frac{t}{T}\right) H_P = (1 - s)H_B + sH_P, \quad (2.2)$$

where the time t runs from 0 to T . It is also sometimes useful to denote $s = t/T$, with $0 \leq s \leq 1$. This defines a path in Hamiltonian space between H_B and H_P , with the rate of change of $H(t)$ controlled by T . The Schrödinger equation for a state $|\psi(t)\rangle$ reads

$$i \frac{d}{dt} |\psi(t)\rangle = H(t) |\psi(t)\rangle. \quad (2.3)$$

I now choose the state at $t = 0$ to be the ground state of H_B

$$|\psi(0)\rangle = |g_B\rangle \quad (2.4)$$

and run the algorithm for time T (let the system evolve according to $H(t)$). The power of the algorithm comes from the Adiabatic theorem. Suited for our investigation, in this work I choose to refer to its formulation due to Jeffrey Goldstone [43]:

Theorem 2 (Adiabatic Theorem). *Consider a time-dependent Hamiltonian $H(t)$, and the time evolution of the state $|\phi_0\rangle$, the ground state of $H(0)$, from time $t = 0$ to $t = T$. Then*

$$\epsilon(T) = \|U(T, 0)\phi_0(0)\rangle - |\phi_0(T)\rangle\| = O\left(\frac{1}{T} \max_t \frac{M(t)^2}{g(t)^3}\right), \quad (2.5)$$

where $U(T, 0)$ is the time evolution operator from time 0 to T corresponding to $H(t)$, while $M(t) = T \left\| \frac{dH}{dt} \right\|$ and $g(t) = E_1(t) - E_0(t)$ is the energy gap of $H(t)$.

Thus, by the adiabatic theorem, if the gap is large enough, T is large enough, and the norm of the derivative of $H(t)$ is bounded, $|\psi(T)\rangle$ will have a large component in the ground state subspace of H_P . A measurement of z can then be used to find the minimum of the classical cost function $h(z)$ I wanted to find. The algorithm is useful if the required run time T does not grow exponentially with the number of qubits n .

Rescaling the Hamiltonian by a factor E would result in a shorter required running time. However, this is a known tradeoff between energy and time. It is thus usual to think

the required resources of an AQC algorithm as $T \cdot \|H\|$, or as the required running time for a rescaled Hamiltonian whose norm is $\|H\| = 1$.

Note that if the eigenvalue $\lambda_0(t)$ is degenerate, i.e. when we have a subspace $\mathcal{L}_0(t)$ instead of a single state $|\phi_0(t)\rangle$ corresponding to the eigenvalue $\lambda_0(s)$, the formulation of the theorem can be generalized to involve the projection onto the subspace $\mathcal{L}_0(T)$ instead of the overlap with the single eigenstate $|\phi_0(T)\rangle$.

It used to be common to state the adiabatic theorem with $\epsilon(T) = \text{const.}$ for $T \propto \Delta^{-2}$, where $\Delta = \min_t g(t)$ is the minimum energy gap. However, several authors have recently rigorously investigated the sufficient conditions for the adiabatic theorem and only proved results in which the scaling of T with the gap was worse than Δ^{-2} [26, 46, 9]. In this work, I choose to use Goldstone's formulation which for the particular Hamiltonian I investigate (2.2) implies constant overlap of the time evolved state with the ground state of H_P if the control time depends on the gap like $T \propto \Delta^{-3}$.

Also note that other time-dependent paths in Hamiltonian space from H_B to H_P are possible, each of them defining a specific Quantum Adiabatic Algorithm. In fact, it has been shown [31], that a path change in the Hamiltonian space can be used to beat certain counterexamples in which the standard, linear-path QAC requires exponential time to run.

2.1.1 The power of AQC

Although AQC was introduced with the hope of solving local optimization problems, the verdict on the practical use of AQC for this purpose is so far 'not discouraging'. Several numerical studies provided results that did not show exponential behavior in the required running time [33, 44]. Nevertheless, because of computer resource limitations, these could not be done for bit numbers greater than $n = 24$, while behavior change might be expected for much larger values of n [19]. Only very recently, Young et al. [113] utilized a different numerical method, investigating the gap in a QAC algorithm for the NP-complete *Exact Cover* problem using a Quantum Monte Carlo method, and found polynomial behavior of the gap for up to $n = 128$ qubits.

On the other hand, soon after AQC was introduced, examples of cases in which a simple linear adiabatic algorithm failed were constructed by van Dam et al. [101]. However, all of these problems were circumvented in Farhi et al. [31] by changing the path in Hamiltonian space by adding a random term H_C to the Hamiltonian as

$$H(t) = (1 - s)H_B + sH_P + s(1 - s)H_C. \quad (2.6)$$

In the cleverly constructed example of [101], the gap above the ground state is exponentially small only at a single point on the interpolation path in s . Adding the term H_C avoids the bad point on the Hamiltonian path where energy levels cross, making the scaling of the required runtime favorable again.

Recently, more papers investigating cases where certain adiabatic algorithms fail appeared [108, 109]. Nevertheless, they share a common design flaw, a bad choice of the starting Hamiltonian. In Section 2.3, I show that poor choices for the Hamiltonian can guarantee that the algorithm will not find the minimum if the run time grows more slowly than \sqrt{N} . These poor choices are nonlocal and wash out any structure in the cost function to be minimized and the best that can be hoped for is Grover speedup. These failures tell us what not to do when designing quantum adiabatic algorithms.

AQC is interesting on its own as an implementation of a quantum computer, as Aharonov

et al. in [7] (see Section 2.2) proved that its computational power is equivalent to that of the conventional quantum circuit model. Because of the energy gap, the computation in the AQC model is inherently protected against noise. The robustness and fault tolerance of the AQC has been recently investigated by Childs et al. [23], Jordan et al. [92], Lidar [61] and Lloyd [63]. The challenge today is to find more new, entirely adiabatic algorithms, such as the adiabatic state preparation of Aharonov and Ta-Shma [6].

2.2 AQC is equivalent to BQP

Aharonov et al. proved the universality of the AQC model in [7]:

Theorem 3 (Equivalence of AQC and the quantum circuit model). *The model of adiabatic computation is polynomially equivalent to the standard model of quantum computation.*

One of the proof directions is straightforward. One can simulate time evolution with an AQC Hamiltonian on a standard quantum computer, dividing the time into small intervals and approximating the time evolution in each slice. The usual procedure utilizes the finite n approximation of the Trotter-Suzuki formula [98, 93, 94, 96, 95]

$$e^{i(A+B)\Delta t} = \lim_{n \rightarrow \infty} \left(e^{iA\Delta t/n} e^{iB\Delta t/n} \right)^n, \quad (2.7)$$

and the resulting unitary transformations are applied using a circuit-based quantum computer. The precision of this approximation can be improved, as the AQC Hamiltonian is changing only linearly in time by integrating the Dyson series for the time-dependent Hamiltonian $H(t)$. Assume that $H(t)$ consists of two types of terms, the first containing only terms built from σ_i^x and the second containing only terms built from σ_i^z . The Hamiltonian can then be written as

$$H(t) = \left(1 - \frac{t}{T}\right) H_x + \left(\frac{t}{T}\right) H_z, \quad (2.8)$$

where H_x and H_z contain only terms built from the respective type of operator. As shown in Appendix B, the approximation of the time evolution operator $U(t + \Delta t, t)$ in the form of a product of exponentials, correct to order $(\Delta t)^2$, is

$$W_2 = e^{-ia_2 H_x} e^{-ib_2 H_z} e^{-ic_2 H_x}, \quad (2.9)$$

with

$$a_2 = c_2 = \left(1 - \frac{t}{T} - \frac{\Delta t}{2T}\right) \frac{\Delta t}{2}, \quad (2.10)$$

$$b_2 = \left(\frac{t}{T} + \frac{\Delta t}{2T}\right) \Delta t. \quad (2.11)$$

On the other hand, one can encode the evaluation of any quantum circuit U into time evolution with a sufficiently slowly changing time-dependent AQC Hamiltonian. The proof in [7] uses ideas from Kitaev's construction I described in Section 1.3.2. Consider a system with two registers, work and clock (1.32), with the work register holding n qubits q_1, \dots, q_n and with the clock register consisting of $L + 1$ qubits c_0, \dots, c_L , where L is the number of

gates in the circuit U . Encode the clock states $|t\rangle_c$ in unary as in (1.58). Aharonov et. al. choose the initial Hamiltonian

$$H_{init} = H_{input} + H_{clock} + H_{clockinit}. \quad (2.12)$$

The term H_{input} (1.57) checks that when the clock is in the state $|0\rangle_c = |1_{c_0} 0_{c_1} \dots 0_{c_L}\rangle$, the state of the work qubits is initialized in $|00\dots 0\rangle$, by adding an energy penalty to all states for which this is not true. Second, H_{clock} (1.59) checks whether the clock register states are proper unary-clock states. Finally, in the unary clock encoding, the term

$$H_{clockinit} = |1\rangle\langle 1|_{c_1} \quad (2.13)$$

is a projector onto all states whose clock register is not in the state $|0\rangle_c$, the only legal clock state whose second clock qubit (c_1) is in the state $|0\rangle$. Because the three terms H_{init} are non-negative, the ground state of H_{init} is the state annihilated by all of them:

$$|\phi_{init}\rangle = |\phi_0(0)\rangle = |00\dots 0\rangle \otimes |1_{c_1} 0_{c_2} \dots 0_{c_T}\rangle. \quad (2.14)$$

The final Hamiltonian is chosen to be a modification of (1.60)

$$H_{final} = H_{prop} + H_{input} + H_{clock}, \quad (2.15)$$

omitting the term H_{out} . The term H_{prop} (1.53) has a degenerate ground state subspace containing all the possible history states (1.51) for a given quantum circuit U . The ground state of the complete H_{final} is the specific history state (1.51) for the quantum circuit U starting in the initial state (2.14). Using classical Markov chain methods, the authors then show that the gap above the ground state of the AQC Hamiltonian

$$H(\tau) = \left(1 - \frac{\tau}{T}\right) H_{init} + \frac{\tau}{T} H_{final} \quad (2.16)$$

is bounded from below by an inverse polynomial in L (the length of the computation U). This is possible, because it is enough to look for this gap in the subspace of legal unary clock states. Moreover, we know that the eigenvalue gap between the ground state subspace and the first excited state of H_{prop} is polynomially large (see Appendix A for details).

The authors then encode the clock register in unary, using the 3-local construction of [54], and prove that 3-local AQC is universal for quantum computation. However, the cost of this translation from a quantum circuits to AQC is steep. First, the required runtime for the adiabatic algorithm is proportional to at least the inverse gap squared. Second, the gap of the final Hamiltonian is small, because of the penalty terms necessary for the 3-local clock construction of [54]. The required runtime, for a rescaled Hamiltonian with $\|H\| = O(1)$, is then claimed to be $T \propto L^{14}$. However, the authors in [7] obtained this using a formulation of the Adiabatic theorem which was later found to be wrong, with a Δ^{-2} dependence of the runtime. If, in fact, I use the adiabatic theorem with $\Delta^{-2+\delta}$, their result becomes $T \propto L^{14+3\delta}$. The proper adiabatic theorem I stated above has $\delta = 1$, estimating the running time for the simulation of quantum circuits by AQC as $T \propto L^{17}$. Deift, Ruskai and Spitzer [26] later showed that one of the gap estimates in the proof (not in the adiabatic theorem) of [7] can be made better, resulting in the total required runtime scaling $T \propto L^{12+3\delta}$ (or, with the adiabatic theorem I use, L^{15}). It is possible that the adiabatic theorem with $\delta < 1$ in the condition $T \propto \Delta^{-(2+\delta)}$ is provable [46]. However,

for now, I stay on the safe side and use the adiabatic theorem with $\delta = 1$, as proved by Goldstone [43].

Later, Mizel et al. [65] showed that the required running time for simulating quantum circuits using AQC scales like $n^2 L^2$, where L is the number of gates in the circuit, and n is the number of work qubits. However, this running time is for a Hamiltonian whose norm is $\|H\| = L$, so that the real running time for a rescaled Hamiltonian scales like $n^2 L^3$. Their construction is built on 2-particle, 4-site interactions of electrons hopping through an array of quantum dots. If we think of their model as implemented using qubits instead of hopping particles, it becomes 4-local. Also, they use the knowledge about where in their algorithm the gap is small, running the algorithm slower around that spot, allowing them to have their runtime scale as good as $T \propto \Delta^{-1}$.

When I want to keep the interactions 3-local (for qubits), much improvement is necessary. For this, I suggest using my new 3-local Hamiltonian construction of Section 4.2 instead of [54]. The required running time for the AQC algorithm then becomes $T \propto L^{5+2\delta} = L^7$. This scales worse than Mizel's required runtime, but my Hamiltonian is only 3-local for qubits. Also, when implementing my static-qubit construction using real hopping particles like Mizel [65], the required interactions couple 2 particles, one of which hops between two sites, giving a 2-particle, 3-site interaction. The improvement over the previous 3-local constructions is possible because my clock register encoding uses only constant norm penalty terms for illegal clock states, as opposed to penalty terms scaling like high polynomials in L as in [54, 53]. Moreover, the Hamiltonians used can be further restricted using the results of Sections 4.4.2 and 4.5, as explained therein.

This, however, is not the strongest result I want to present here. Using an unpublished result of Seth Lloyd [63], in the next Section I show that the resources necessary to perform universal quantum computation using 3-local AQC scale only as $\tau \cdot \|H\| = O(L^2 \log^2 L)$, where τ is the algorithm runtime. However, in this case it is not the adiabatic theorem that gives the computational model its power, it is rather the underlying structure of the final Hamiltonian whose dynamics is a quantum walk on a line.

2.2.1 The Hamiltonian Computer model and universality

In this section, I look at the dynamics of the Hamiltonian (2.16) in more detail. The result presented here was pointed out to me by Seth Lloyd during my Part III doctoral exam¹ and I thank him for the access to the write-up of his unpublished work [63]. He showed that even when the control time T for the change in the Hamiltonian is short, the result of the quantum circuit U can still be obtained from the state of the system with high probability after waiting for time τ upper bounded by a polynomial in L , the number of gates in the circuit U .

What can happen to the state $|\phi_{init}\rangle$ (2.14) when we change the Hamiltonian from H_{init} to H_{final} faster than the required conditions in the statement of the adiabatic theorem? Let me consider an extreme scenario, jumping from H_{init} to H_{final} instantly. The system will not be in the ground state of H_{final} , but in a superposition of some of its eigenstates, those which it could transition to under the action of the Hamiltonian (2.16). To see what these are, let me first analyze the time evolution of $|\phi_{init}\rangle$ directly with H_{final} (2.15). The term determining the dynamics of the system is H_{prop} , as the terms H_{input} and H_{clock} annihilate the initial state and all the states H_{final} brings it to. Consider now the subspace of states

¹I was made to work it out on the board, and was quite puzzled by it at the time, thinking I must have gotten it wrong.

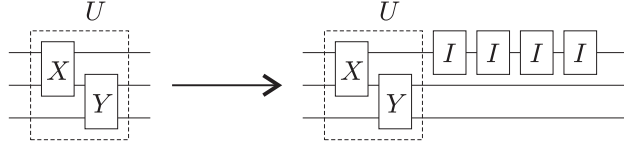


Figure 2-1: Padding the circuit U with L gates by $2L$ extra identity gates. The state of the system after L steps stays unchanged.

which one can transition to from $|s\rangle \otimes |0\rangle_c$ using arbitrary powers of H_{prop} . There is one such subspace for each work register basis state $|s\rangle$, and these subspaces are not connected by H_{prop} . As described in much more detail in Appendix A, after a unitary basis change

$$|s\rangle \otimes |t\rangle_c \rightarrow |s_t\rangle \otimes |t\rangle_c = (U_t U_{t-1} \dots U_2 U_1) |s\rangle \otimes |t\rangle_c, \quad (2.17)$$

the Hamiltonian H_{prop} becomes block diagonal. The matrix form of H_{prop} in each of these blocks is

$$H'_{prop} = \mathbb{I} - \frac{1}{2}B - \frac{1}{2}P_0 - \frac{1}{2}P_L, \quad (2.18)$$

where the matrix B (see Appendix C) is the adjacency matrix for a line of length $L + 1$ and P_0 and P_L are projectors onto the endpoints of this line. This gives the system the dynamics of a quantum walk on a line.

Because the blocks for different work register basis states $|s\rangle$ are decoupled, the specific initial state $|\phi_{init}\rangle = |0\rangle \otimes |0\rangle_c$ time evolved for some time τ with H_{prop} can thus end up only in some superposition of the states

$$|\Psi_t\rangle = |\psi_t\rangle \otimes |t\rangle_c = (U_t U_{t-1} \dots U_2 U_1) |00 \dots 0\rangle \otimes |t\rangle_c, \quad (2.19)$$

where $|\psi_t\rangle$ is the state of the work qubits after t gates of the quantum circuit U have been applied to an initial all-zero state $|00 \dots 0\rangle$. The span of all these states $|\Psi_t\rangle$ defines the subspace \mathcal{H}_0 . The state of the system $|\phi(\tau)\rangle$ after time τ has no overlap with states outside of \mathcal{H}_0 , because the Hamiltonian H_{prop} is constructed so that it doesn't couple \mathcal{H}_0 to the subspaces generated by H_{prop} from some different state $|s \neq 0\rangle \otimes |0\rangle_c$.

Note that the state $|\psi_L\rangle$ contains the result of the quantum circuit U in its work register. I would like to find the system with the computation already done, so I boost the probability by padding the initial circuit as in Figure 2-1. After the original L gates of the circuit, let me add $2L$ extra identity gates, taking $L \rightarrow 3L$. Now, all the states $|\Psi_t\rangle$ with $t > L/3$ contain the output of the circuit U in their work register, as the work register of $|\Psi_t\rangle$ is $U|00 \dots 0\rangle$ and stays unchanged for all $t > L/3$ because I padded the circuit with identity gates.

The time evolution of $|\phi_{init}\rangle$ with H_{prop} corresponds to a continuous-time quantum walk on a line with specific boundary conditions. I can get rid of the boundary conditions by wrapping the walk around on a circle, which simplifies the analysis, given in detail in Appendix C. There I show that this walk mixes so fast, that when I choose a time τ uniformly at random between 0 and $\tau_1 \propto L \log^2 L$ the probability to find the system in a state $|t\rangle_c$ with $t > L/3$ is close to $2/3$. Therefore, when I let the state $|\phi_{init}\rangle$ evolve for a

time chosen uniformly at random between 0 and $O(L \log^2 L)$ and measure the clock qubit $c_{L/3}$, I will obtain the result 1 with probability close to $2/3$. This signifies ‘success’, as it implies that the clock register of the state I measured was one of the states with $t > L/3$. The work qubits then contain the result of the quantum computation U .

To conclude, when I change H_{init} to H_{final} quickly, the resulting state of the system will be some superposition of states $|\Psi_t\rangle$, close to $|\Psi_0\rangle$. When I then let the system evolve for extra time $\tau < O(L \log^2 L)$ with the Hamiltonian H_{final} , I can measure the clock register, and with high probability I will obtain a result that tells me that the work register now contains the result of the quantum circuit U . The overall resources necessary for this procedure then scale like $\tau \cdot \|H\| = O(L^2 \log^2 L)$, as the Hamiltonian I use has norm $O(L)$.

Let me now formalize the model of computation I just described.

Definition 11 (Hamiltonian Computer (HC) model). *A Hamiltonian Computer aiming to simulate a quantum circuit U with L gates is defined by a system consisting of two registers, work and clock, a k -local Hamiltonian H , a simple to prepare initial state $|\psi_0\rangle$, and a bound on the required evolution time. The clock register states must be constructed in such a way that a measurement distinguishing $|t \leq L/3\rangle_c$ (failure) and $|t > L/3\rangle_c$ (success) is simple. The computation procedure goes as follows:*

1. *Prepare the system in the initial state $|\psi_0\rangle$, the ground state of a k -local Hamiltonian H_0 .*
2. *Turn off the initial Hamiltonian H_0 , and let the system evolve with a time independent Hamiltonian H for a time τ chosen uniformly at random between 0 and $\text{poly}(L)$.*
3. *Make a measurement determining whether the clock register is in a state $|t > L/3\rangle_c$, which signifies success. The probability to obtain such state is greater than $p = \frac{1}{2}$. The result of the quantum circuit U is then in the work register of the system. Restart otherwise.*

Feynman’s Hamiltonian Computer with a unary clock $|t\rangle_c = |0 \dots 010 \dots 0\rangle_c$ (see Section 1.3.1) is one of the possible implementations of a HC. The initial Hamiltonian H_0 is one whose ground state is $|\psi_0\rangle = |0_{q_1} \dots 0_{q_n}\rangle \otimes |0_{c_0} \dots 0_{c_L}\rangle$, and the time independent Hamiltonian H is given by (1.35). To determine success/failure, I need to measure the first third of the clock qubits. When I don’t find the spin up there, it means I have succeeded. In the case of Feynman’s Hamiltonian, the Hamiltonian itself is 4-local, and Lloyd [63] bounds the required runtime from above by L^2 . As seen in the previous Section, Aharonov et al. [7] use the 3-local Hamiltonian H_{final} of Kempe and Regev [54] to simulate quantum circuits using AQC, and require resources of the order $L^{12+3\delta}$. However, when I encode the clock register and the clock state transitions as in Section 4.5, the gap of H_{final} whose norm is $O(L)$ scales like L^{-2} . The mixing of the quantum walk is then such that the required running time becomes only $O(L \log^2 L)$. When I rescale my Hamiltonian to $\|H\| = 1$, the resources necessary for this HC model are $O(L^2 \log^2 L)$. For more detail on this construction, see Section 4.5.

What remains is the question how is one protected against noise in such Hamiltonian Computer models. In AQC, it was the energy gap. However, as I have shown, the energy gap computed by Aharonov et al. [7] is not relevant, because the transitions to higher momenta eigenstates of the quantum walk still contain the result of the computation. The relevant energy barrier is the energy cost to transition from the subspace \mathcal{H}_{legal} to $\mathcal{H}_{legal}^\perp$, in which case we would lose the computation. For this purpose, I use a clock-checking

Hamiltonian like (1.59), whose norm is $O(L)$. Lloyd [63] showed, that the energy gap for such Hamiltonian scales like L^{-1} . In summary, when simulating quantum circuits in a HC model, one requires resources $\tau \cdot \|H\| = O(L^2 \log^2 L)$, and is protected by an energy gap scaling like L^{-1} , where L is the number of gates in the circuit. This energy barrier to decoherence can be increased, at the cost of increasing the locality of interactions, by using error correcting codes like the ones developed for AQC by Jordan et al. [92].

2.3 How to make the Quantum Adiabatic Algorithm Fail

After looking at the power of AQC in general and showing how it can be used to simulate any quantum circuit effectively, let me now return to its original intended purpose: solving locally constrained optimization problems.

Recently, a paper announcing the failure of AQC in certain cases appeared [108], putting considerable effort into showing that the gap in a specific AQC algorithm is exponentially small. However, the authors make a bad choice of the starting Hamiltonian H_B in (2.2). In this section I prove some general results which show that with certain choices of H_B or H_P the algorithm will not succeed if T is $o(\sqrt{N})$, that is $T/\sqrt{N} \rightarrow 0$ as $N \rightarrow \infty$, so that improvement beyond Grover speedup is impossible. I view these failures as due to poor choices for H_B and H_P , which teach us what not to do when looking for good algorithms. I guarantee failure by removing any structure which might exist in $h(z)$ from either H_B or H_P . By structure I mean that z is written as a bit string and both H_B and H_P are sums of terms involving only a few of the corresponding qubits. After the preprint [34] was posted on the quant-ph archive, two similar papers [109, 45] concerning this topic also appeared.

In Section 2.3.1 I show that regardless of the form of $h(z)$ if H_B is a one dimensional projector onto the uniform superposition of all the basis states $|z\rangle$, then the quantum adiabatic algorithm fails. Here all the $|z\rangle$ states are treated identically by H_B so any structure contained in $h(z)$ is lost in H_B . In Section 2.3.2 I consider a scrambled H_P that I get by replacing the cost function $h(z)$ by $h(\pi(z))$ where π is a permutation of 0 to $N-1$. Here the values of $h(z)$ and $h(\pi(z))$ are the same but the relationship between input and output is scrambled by the permutation. This effectively destroys any structure in $h(z)$ and typically results in algorithmic failure.

The quantum adiabatic algorithm is a special case of Hamiltonian based continuous time quantum algorithms, where the quantum state obeys (2.3) and the algorithm consists of specifying $H(t)$, the initial state $|\psi(0)\rangle$, a run time T and the operators to be measured at the end of the run. In the Hamiltonian language, the Grover problem can be recast as the problem of finding the ground state of

$$H_w = E(\mathbb{I} - |w\rangle\langle w|), \quad (2.20)$$

where w lies between 0 and $N-1$. The algorithm designer can apply H_w , but in this oracular setting, w is not known. In [37] the following result was proved. Let

$$H(t) = H_D(t) + H_w, \quad (2.21)$$

where H_D is any time dependent “driver” Hamiltonian independent of w . Assume also that the initial state $|\psi(0)\rangle$ is independent of w . For each w we want the algorithm to be

successful, that is $|\psi(T)\rangle = |w\rangle$. It then follows that

$$T \geq \frac{\sqrt{N}}{2E}. \quad (2.22)$$

The proof of this result is a continuous-time version of the BBBV oracular proof [17]. My proof techniques in this paper are similar to the methods used to prove the result just stated.

2.3.1 General search starting with a one-dimensional projector

In this section I consider a completely general cost function $h(z)$ with $z = 0, \dots, N-1$. The goal is to use the quantum adiabatic algorithm to find the ground state of H_P given by (2.1) with $H(t)$ given by (2.2). Let

$$|s\rangle = \frac{1}{\sqrt{N}} \sum_{z=0}^{N-1} |z\rangle \quad (2.23)$$

be the uniform superposition over all possible values z . If I pick

$$H_B = E(\mathbb{I} - |s\rangle\langle s|) \quad (2.24)$$

and $|\psi(0)\rangle = |s\rangle$, then the adiabatic algorithm fails in the following sense:

Theorem 4. *Let H_P be diagonal in the z basis with a ground state subspace of dimension k . Let*

$$H(t) = (1 - t/T)E(\mathbb{I} - |s\rangle\langle s|) + (t/T)H_P.$$

Let P be the projector onto the ground state subspace of H_P and let $b > 0$ be the success probability, that is, $b = \langle \psi(T) | P | \psi(T) \rangle$. Then

$$T \geq \frac{b}{E} \sqrt{\frac{N}{k}} - \frac{2\sqrt{b}}{E}.$$

Proof. Keeping H_P fixed, introduce $N-1$ additional beginning Hamiltonians as follows. For $x = 0, \dots, N-1$ let V_x be a unitary operator diagonal in the z basis with

$$\langle z | V_x | z \rangle = e^{2\pi i z x / N}$$

and let

$$|x\rangle = V_x |s\rangle = \frac{1}{\sqrt{N}} \sum_{z=0}^{N-1} e^{2\pi i z x / N} |z\rangle$$

so that the $\{|x\rangle\}$ form an orthonormal basis. Note also that

$$|x=0\rangle = |s\rangle.$$

We now define

$$H_x(t) = (1 - t/T)E(\mathbb{I} - |x\rangle\langle x|) + (t/T)H_P,$$

with corresponding evolution operator $U_x(t_2, t_1)$. Note that $H(t)$ above is $H_0(t)$ with the corresponding evolution operator U_0 . For each x let the system evolve with $H_x(t)$ from

the ground state of $H_x(0)$, which is $|x\rangle$. Note that $H_x = V_x H_0 V_x^\dagger$ and $U_x = V_x U_0 V_x^\dagger$. Let $|f_x\rangle = U_x(T, 0) |x\rangle$. For each x the success probability is $\langle f_x | P | f_x \rangle$, which is equal to b since P commutes with V_x . The key point is that if I run the Hamiltonian evolution with H_x backwards in time, I would then be finding x , that is, solving the Grover problem. However, this should not be possible unless the run time T is of order \sqrt{N} .

Let U_R be the evolution operator corresponding to an x -independent reference Hamiltonian

$$H_R(t) = (1 - t/T)E + (t/T)H_P.$$

Let $|g_x\rangle = \frac{1}{\sqrt{b}}P |f_x\rangle$ be the normalized component of $|f_x\rangle$ in the ground state subspace of H_P . Consider the difference in backward evolution from $|g_x\rangle$ with Hamiltonians H_x and H_R , and sum on x ,

$$S(t) = \sum_x \left\| U_x^\dagger(T, t) |g_x\rangle - U_R^\dagger(T, t) |g_x\rangle \right\|^2.$$

Clearly $S(T) = 0$, and

$$S(0) = \sum_x \left\| U_x^\dagger(T, 0) |g_x\rangle - U_R^\dagger(T, 0) |g_x\rangle \right\|^2.$$

Now $|g_x\rangle = \sqrt{b}|f_x\rangle + \sqrt{1-b}|f_x^\perp\rangle$ where $|f_x^\perp\rangle$ is orthogonal to $|f_x\rangle$. Since $U_x^\dagger(T, 0) |f_x\rangle = |x\rangle$,

$$S(0) = \sum_x \left\| \sqrt{b}|x\rangle + \sqrt{1-b}|x^\perp\rangle - |i_x\rangle \right\|^2,$$

where for each x , $|x^\perp\rangle$ and $|i_x\rangle$ are normalized states with $|x^\perp\rangle$ orthogonal to $|x\rangle$. Since H_R commutes with H_P , $|i_x\rangle = U_R^\dagger(T, 0) |g_x\rangle$ is an element of the k -dimensional ground state subspace of H_P . Then

$$\begin{aligned} S(0) &= 2N - \sum_x \left[\sqrt{b}\langle x | i_x \rangle + \sqrt{1-b}\langle x^\perp | i_x \rangle + c.c. \right] \\ &\geq 2N - 2\sqrt{b} \sum_x \left| \langle x | i_x \rangle \right| - 2N\sqrt{1-b}. \end{aligned}$$

Choosing a basis $\{|G_j\rangle\}$ for the k dimensional ground state subspace of H_P and writing $|i_x\rangle = a_{x1}|G_1\rangle + \dots + a_{xk}|G_k\rangle$ gives

$$\begin{aligned} \sum_x \left| \langle x | i_x \rangle \right| &\leq \sum_{x,j} |a_{xj}| \cdot \left| \langle x | G_j \rangle \right| \\ &\leq \sqrt{\sum_{x,j} |a_{xj}|^2 \sum_{x',j'} \left| \langle x' | G_{j'} \rangle \right|^2} = \sqrt{Nk}. \end{aligned} \tag{2.25}$$

Thus

$$S(0) \geq 2N(1 - \sqrt{1-b}) - 2\sqrt{b}\sqrt{Nk}. \tag{2.26}$$

Using the Schrödinger equation, one can find the time derivative of $S(t)$:

$$\begin{aligned}
\frac{d}{dt}S(t) &= -\sum_x \frac{d}{dt} \left[\langle g_x | U_x(T, t) U_R^\dagger(T, t) | g_x \rangle + c.c. \right] \\
&= -i \sum_x \langle g_x | U_x(T, t) [H_x(t) - H_R(t)] U_R^\dagger(T, t) | g_x \rangle + c.c. \\
&= -2 \operatorname{Im} \sum_x (1 - t/T) E \langle g_x | U_x(T, t) | x \rangle \langle x | U_R^\dagger(T, t) | g_x \rangle.
\end{aligned}$$

Now

$$\begin{aligned}
\left| \frac{d}{dt}S(t) \right| &\leq 2E(1 - t/T) \sum_x \left| \langle g_x | U_x(T, t) | x \rangle \langle x | U_R^\dagger(T, t) | g_x \rangle \right| \\
&\leq 2E(1 - t/T) \sum_x \left| \langle x | U_R^\dagger(T, t) | g_x \rangle \right|.
\end{aligned}$$

The same technique as in (2.25) gives

$$\left| \frac{d}{dt}S(t) \right| \leq 2E(1 - t/T) \sqrt{Nk}.$$

Therefore

$$\int_0^T \left| \frac{d}{dt}S(t) \right| dt \leq ET \sqrt{Nk}.$$

Now $S(0) \leq S(T) + \int_0^T \left| \frac{d}{dt}S(t) \right| dt$ and $S(T) = 0$ so

$$S(0) \leq ET \sqrt{Nk}.$$

Combining this with (2.26) gives

$$ET \sqrt{Nk} \geq 2N(1 - \sqrt{1 - b}) - 2\sqrt{b} \sqrt{Nk},$$

which implies what I wanted to prove:

$$T \geq \frac{b}{E} \sqrt{\frac{N}{k}} - \frac{2\sqrt{b}}{E}.$$

□

How to interpret Theorem 4? The goal is to find the minimum of the cost function $h(z)$ using the quantum adiabatic algorithm. It is natural to pick for H_B a Hamiltonian whose ground state is $|s\rangle$, the uniform superposition of all $|z\rangle$ states. However if I pick H_B to be the one dimensional projector $E(\mathbb{I} - |s\rangle \langle s|)$ the algorithm will not find the ground state if T/\sqrt{N} goes to 0 as N goes to infinity. The problem is that H_B has no structure and makes no reference to $h(z)$. My hope is that the algorithm might be useful for interesting computational problems if H_B has structure that reflects the form of $h(z)$.

Note that Theorem 4 explains the algorithmic failure discovered by Žnidarič and Horvat [108] for a particular set of $h(z)$.

For a simple but convincing example of the importance of the choice of H_B , suppose I take a decoupled n bit problem which consists of n clauses each acting on one bit, say for each bit j

$$h_j(z) = \begin{cases} 0 & \text{if } z_j = 0, \\ 1 & \text{if } z_j = 1, \end{cases}$$

so

$$h(z) = z_1 + z_2 + \cdots + z_n. \quad (2.27)$$

Let me pick a beginning Hamiltonian reflecting the bit structure of the problem,

$$H_B = \sum_{j=1}^n \frac{1}{2} (1 - \sigma_x^{(j)}). \quad (2.28)$$

The ground state of H_B is $|s\rangle$, The quantum adiabatic algorithm acts on each bit independently, producing a success probability of

$$p = (1 - q(T))^n,$$

where $q(T) \rightarrow 0$ as $T \rightarrow \infty$ is the transition probability between the ground state and the excited state of a single qubit. As long as $nq(T) \rightarrow \text{const.}$ I have a constant probability of success. This can be achieved for T of order \sqrt{n} , because for a two level system with a nonzero gap, the probability of a transition is $q(T) = O(T^{-2})$. (For details, see Appendix E.) However, I know from Theorem 4 that a poor choice of H_B would make the quantum adiabatic algorithm fail on this simple decoupled n bit problem by destroying the bit structure.

Next, suppose the satisfiability problem I am trying to solve has clauses involving say 3 bits. If clause c involves bits i_c, j_c and k_c I may define the clause cost function

$$h_c(z) = \begin{cases} 0 & \text{if } z_{i_c}, z_{j_c}, z_{k_c} \text{ satisfy clause } c, \\ 1 & \text{otherwise.} \end{cases}$$

The total cost function is then

$$h(z) = \sum_c h_c(z).$$

To get H_B to reflect the bit and clause structure I may pick

$$H_{B,c} = \frac{1}{2} \left[(1 - \sigma_x^{(i_c)}) + (1 - \sigma_x^{(j_c)}) + (1 - \sigma_x^{(k_c)}) \right]$$

with

$$H_B = \sum_c H_{B,c}. \quad (2.29)$$

In this case the ground state of H_B is again $|s\rangle$. With this setup, Theorem 4 does not apply.

I did a numerical study of a particular satisfiability problem, Exact Cover. For this

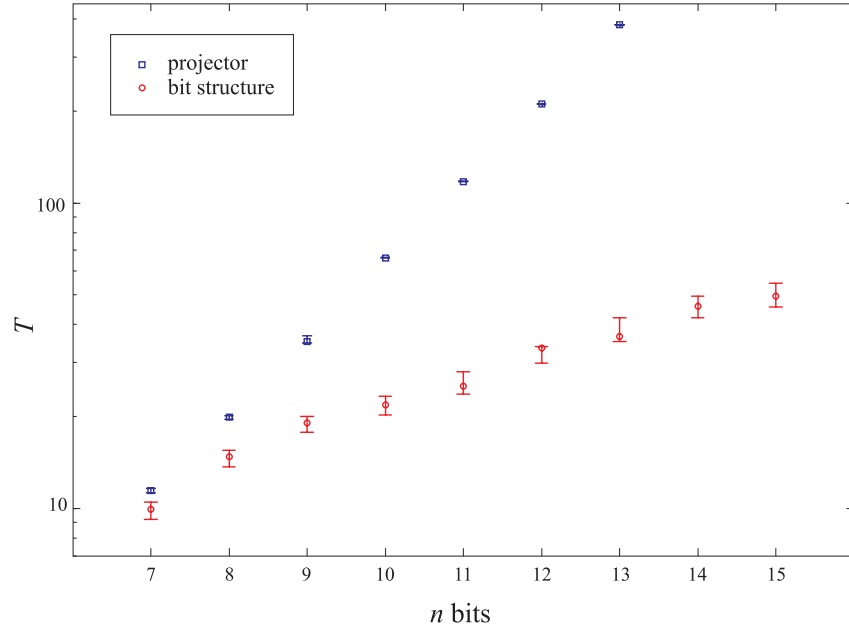


Figure 2-2: Median required run time T versus bit number. At each bit number there are 50 random instances of Exact Cover with a single satisfying assignment. I choose the required run time to be the value of T for which quantum adiabatic algorithm has success probability between 0.2 and 0.21. For the projector beginning Hamiltonian I use (2.24) with $E = n/2$. The plot is log-linear. The error bars show the 95% confidence interval for the true medians.

problem if clause c involves bits i_c, j_c and k_c , the cost function is

$$h_c(z) = \begin{cases} 0 & \text{if } z_{i_c} + z_{j_c} + z_{k_c} = 1, \\ 1 & \text{otherwise.} \end{cases}$$

Some data is presented in FIG. 1. Here we see that with a structured beginning Hamiltonian the required run times are substantially lower than with the projector H_B .

2.3.2 Search with a scrambled problem hamiltonian

In the previous section I showed that removing all structure from H_B dooms the quantum adiabatic algorithm to failure. In this section I remove structure from the problem to be solved (H_P) and show that this leads to algorithmic failure. Let $h(z)$ be a cost function whose minimum I seek. Let π be a permutation of $0, 1, \dots, N-1$ and let

$$h^{[\pi]}(z) = h(\pi^{-1}(z)).$$

I will show that no continuous time quantum algorithm (of a very general form) can find the minimum of $h^{[\pi]}$ for even a small fraction of all π if T is $o(\sqrt{N})$. Classically, this problem takes order N calls to an oracle.

Without loss of generality let $h(0) = 0$, and $h(1), h(2), \dots, h(N-1)$ all be positive. For any permutation π of $0, 1, \dots, N-1$, define a problem Hamiltonian $H_{P,\pi}$, diagonal in the z basis, as

$$H_{P,\pi} = \sum_{z=0}^{N-1} h^{[\pi]}(z) |z\rangle \langle z| = \sum_{z=0}^{N-1} h(z) |\pi(z)\rangle \langle \pi(z)|.$$

Now consider the Hamiltonian

$$H_\pi(t) = H_D(t) + c(t)H_{P,\pi} \tag{2.30}$$

for an arbitrary π -independent driving Hamiltonian $H_D(t)$ with $|c(t)| \leq 1$ for all t . Using this composite Hamiltonian, evolve the π -independent starting state $|\psi(0)\rangle$ for time T , reaching the state $|\psi_\pi(T)\rangle$. This setup is more general than the quantum adiabatic algorithm since I do not require $H_D(t)$ or $c(t)$ to be slowly varying. Success is achieved if the overlap of $|\psi_\pi(T)\rangle$ with $|\pi(0)\rangle$ is large.

I first show

Lemma 1.

$$\sum_{\pi, \pi'} \left\| |\psi_\pi(T)\rangle - |\psi_{\pi'}(T)\rangle \right\|^2 \leq 4h^*TN!\sqrt{N-1}, \tag{2.31}$$

where the sum is over all pairs of permutations π, π' that differ by a single transposition involving $\pi(0)$, and $h^* = \sqrt{\sum h(z)^2/(N-1)}$.

Proof. For two different permutations π and π' let $|\psi_\pi(t)\rangle$ be the state obtained by evolving from $|\psi(0)\rangle$ with H_π and let $|\psi_{\pi'}(t)\rangle$ be the state obtained by evolving from $|\psi(0)\rangle$ with $H_{\pi'}$.

Now

$$\begin{aligned}
\frac{d}{dt} \left\| |\psi_\pi(t)\rangle - |\psi_{\pi'}(t)\rangle \right\|^2 &= -\frac{d}{dt} \langle \psi_\pi(t) | \psi_{\pi'}(t) \rangle + c.c. \\
&= i \langle \psi_\pi(t) | (H_\pi(t) - H_{\pi'}(t)) | \psi_{\pi'}(t) \rangle + c.c. \\
&\leq 2 \left| \langle \psi_\pi(t) | (H_\pi(t) - H_{\pi'}(t)) | \psi_{\pi'}(t) \rangle \right|.
\end{aligned}$$

Consider the case when π and π' differ by a single transposition involving $\pi(0)$. Specifically, $\pi' = \pi \circ (a \leftrightarrow 0)$ for some a . Now if $\pi(0) = i$ and $\pi(a) = j$, we have $\pi'(0) = j$ and $\pi'(a) = i$. Therefore, since $h(0) = 0$,

$$H_{P,\pi} - H_{P,\pi'} = c(t)h(a) (|j\rangle \langle j| - |i\rangle \langle i|) = c(t)h(a) (|\pi(a)\rangle \langle \pi(a)| - |\pi'(a)\rangle \langle \pi'(a)|),$$

so that

$$\frac{d}{dt} \sum_{\pi, \pi'} \left\| |\psi_\pi(t)\rangle - |\psi_{\pi'}(t)\rangle \right\|^2 \leq 2|c(t)| \sum_{\pi, \pi'} h(a) \left| \langle \psi_\pi(t) | (|\pi(a)\rangle \langle \pi(a)| - |\pi'(a)\rangle \langle \pi'(a)|) | \psi_{\pi'}(t) \rangle \right|.$$

This further simplifies to

$$\begin{aligned}
\frac{d}{dt} \sum_{\pi, \pi'} \left\| |\psi_\pi(t)\rangle - |\psi_{\pi'}(t)\rangle \right\|^2 &\leq 2 \sum_{\pi, \pi'} h(a) \left(\left| \langle \psi_\pi(t) | \pi(a) \rangle \right| + \left| \langle \pi'(a) | \psi_{\pi'}(t) \rangle \right| \right) \\
&= 2 \sum_{\pi} \sum_{a \neq 0} h(a) \left| \langle \psi_\pi(t) | \pi(a) \rangle \right| + 2 \sum_{\pi'} \sum_{a \neq 0} h(a) \left| \langle \pi'(a) | \psi_{\pi'}(t) \rangle \right| \\
&= 4 \sum_{\pi} \sum_{a \neq 0} h(a) \left| \langle \psi_\pi(t) | \pi(a) \rangle \right| \\
&= 4 \sum_{\pi} \sum_a h(a) \left| \langle \psi_\pi(t) | \pi(a) \rangle \right| \\
&\leq 4 \sum_{\pi} \sqrt{\sum_a h(a)^2} = 4h^* N! \sqrt{N-1}.
\end{aligned}$$

where I used the Cauchy-Schwartz inequality to obtain the last line. Integrating this inequality for time T , I obtain the result I wanted to prove,

$$\sum_{\pi, \pi'} \left\| |\psi_\pi(T)\rangle - |\psi_{\pi'}(T)\rangle \right\|^2 \leq 4h^* T N! \sqrt{N-1},$$

where the sum is over π and π' differing by a single transposition involving $\pi(0)$. \square

Next I establish

Lemma 2. Suppose $|1\rangle, |2\rangle, |L\rangle$ are orthonormal vectors and $\left| \langle \psi_i | i \rangle \right|^2 \geq b$ for normalized vectors $|\psi_i\rangle$, where $i = 1, \dots, L$. Then for any normalized $|\varphi\rangle$,

$$\sum_{i=1}^L \left\| |\psi_i\rangle - |\varphi\rangle \right\|^2 \geq bL - 2\sqrt{L}. \quad (2.32)$$

Proof. Write

$$\begin{aligned} \sum_i \left\| |\psi_i\rangle - |\varphi\rangle \right\|^2 &\geq \sum_i \left| \langle i|\psi_i\rangle - \langle i|\varphi\rangle \right|^2 \\ &\geq \sum_i \left| \langle i|\psi_i\rangle \right|^2 - 2 \sum_i \left| \langle i|\psi_i\rangle \right| \left| \langle i|\varphi\rangle \right| \end{aligned}$$

and use the Cauchy-Schwartz inequality to obtain

$$\begin{aligned} \sum_i \left\| |\psi_i\rangle - |\varphi\rangle \right\|^2 &\geq bL - 2 \sqrt{\sum_i \left| \langle i|\psi_i\rangle \right|^2} \sqrt{\sum_i \left| \langle i|\varphi\rangle \right|^2} \\ &\geq bL - 2\sqrt{L}. \end{aligned}$$

□

Let me now state the main result of this Section.

Theorem 5. *Suppose that a continuous time algorithm of the form (2.30) succeeds with probability at least b , i.e. $\left| \langle \psi_\pi(T) | \pi(0) \rangle \right|^2 \geq b$, for a set of $\epsilon N!$ permutations. Then*

$$T \geq \frac{\epsilon^2 b}{16h^*} \sqrt{N-1} - \frac{\epsilon \sqrt{\epsilon/2}}{4h^*}. \quad (2.33)$$

Proof. For any permutation π , there are $N-1$ permutations π'_a obtained from π by first transposing 0 and a . For each π let \mathcal{S}_π be the subset of those $N-1$ permutations on which the algorithm succeeds with probability at least b . Any such permutation appears in exactly $N-1$ of the sets \mathcal{S}_π , therefore

$$\sum_\pi |\mathcal{S}_\pi| = (N-1)\epsilon N!.$$

Let M be the number of sets \mathcal{S}_π with $|\mathcal{S}_\pi| \geq \frac{\epsilon}{2}(N-1)$. Now

$$\begin{aligned} \sum_\pi |\mathcal{S}_\pi| &= \sum_{|\mathcal{S}_\pi| \geq \frac{\epsilon}{2}(N-1)} |\mathcal{S}_\pi| + \sum_{|\mathcal{S}_\pi| < \frac{\epsilon}{2}(N-1)} |\mathcal{S}_\pi| \\ \sum_\pi |\mathcal{S}_\pi| &\leq M(N-1) + (N! - M) \frac{\epsilon}{2}(N-1), \\ (N-1)\epsilon N! &\leq M(N-1) + N! \frac{\epsilon}{2}(N-1), \end{aligned}$$

so $M \geq \frac{\epsilon}{2}N!$, i.e. at least $\frac{\epsilon}{2}N!$ of the sets \mathcal{S}_π must contain at least $\frac{\epsilon}{2}(N-1)$ permutations on which the algorithm succeeds with probability at least b . For the corresponding π , I have

$$\sum_{\pi'_a} \left\| |\psi_\pi(T)\rangle - |\psi_{\pi'_a}(T)\rangle \right\|^2 \geq b \frac{\epsilon}{2}(N-1) - 2 \sqrt{\frac{\epsilon}{2}(N-1)}.$$

by Lemma 2. (Note that the algorithm is not assumed to succeed with probability b on π .)

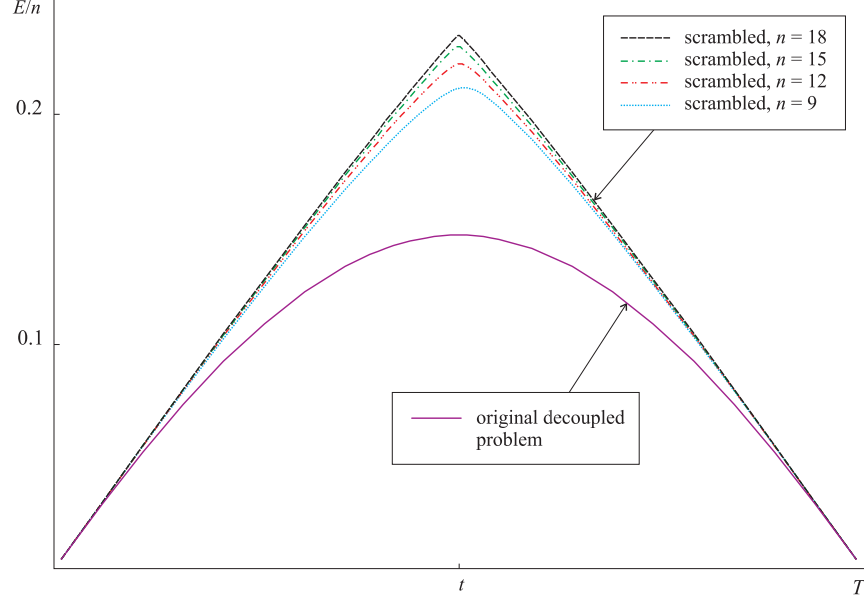


Figure 2-3: The scaled ground state energy E/n for a quantum adiabatic algorithm Hamiltonian of a decoupled problem. The lowest curve corresponds to the original decoupled problem. The upper “triangular” curves correspond to single instances of the n -bit decoupled problem, where the problem Hamiltonian was scrambled.

Since there are at least $\frac{\epsilon}{2}N!$ such π ,

$$\sum_{\pi, \pi'} \left\| |\psi_{\pi}(T)\rangle - |\psi_{\pi'}(T)\rangle \right\|^2 \geq \frac{\epsilon}{2}N! \left(b\frac{\epsilon}{2}(N-1) - 2\sqrt{\frac{\epsilon}{2}(N-1)} \right),$$

where the sum is over all permutations π and π' which differ by a single transposition involving $\pi(0)$. Combining this with Lemma 1 I obtain

$$T \geq \frac{\epsilon^2 b}{16h^*} \sqrt{N-1} - \frac{\epsilon\sqrt{\epsilon/2}}{4h^*},$$

which is what I wanted to prove. \square

What I have just shown is that no continuous time algorithm of the form (2.30) can find the minimum of $H_{P,\pi}$ with a constant success probability for even a fraction $\epsilon N!$ of all permutations π if T is $o(\sqrt{N})$. A typical permutation π yields an $H_{P,\pi}$ with no structure relevant to any fixed H_D and the algorithm cannot find the ground state of $H_{P,\pi}$ efficiently.

To illustrate the nature of this failure for the quantum adiabatic algorithm for a typical permutation, consider again the decoupled n bit problem with $h(z)$ given by (2.27) and H_B given by (2.28). The lowest curve in FIG. 2 shows the ground state energy divided by n as a function of t . (Since the system is decoupled this is actually the ground state energy of a single qubit.) I then consider the n bit scrambled problem for different values of n . At each n I pick a single random permutation π of $0, \dots, (2^n - 1)$ and apply it to obtain a cost function $h(\pi^{-1}(z))$ while keeping H_B fixed. The ground state energy divided by n is

now plotted for $n = 9, 12, 15$ and 18 . From these scrambled problems it is clear that if I let n get large the typical curves will approach a triangle with a discontinuous first derivative at $t = T/2$. For large n , the ground state changes dramatically as t passes through $T/2$. In order to keep the quantum system in the ground state we need to go very slowly near $t = T/2$ and this results in a long required run time.

2.3.3 Summary

I have shown two main results about the performance of the quantum adiabatic algorithm when used to find the minimum of a classical cost function $h(z)$ with $z = 0, \dots, N - 1$. Theorem 4 says that for any cost function $h(z)$, if the beginning Hamiltonian is a one dimensional projector onto the uniform superposition of all the $|z\rangle$ basis states, the algorithm will not find the minimum of h if T is less then of order \sqrt{N} . This is true regardless of how simple it is to classically find the minimum of $h(z)$.

In Theorem 5 I start with any beginning Hamiltonian and classical cost function h . Replacing $h(z)$ by a scrambled version, i.e. $h^{[\pi]}(z) = h(\pi(z))$ with π a permutation of 0 to $N - 1$, will make it impossible for the algorithm to find the minimum of $h^{[\pi]}$ in time less than order \sqrt{N} for a typical permutation π . For example suppose we have a cost function $h(z)$ and have chosen H_B so that the quantum algorithm finds the minimum in time of order $\log N$. Still scrambling the cost function results in algorithmic failure.

These results do not imply anything about the more interesting case where H_B and H_P are structured, i.e., sums of terms each operating only on several qubits.

Chapter 3

Matrix Product States on Infinite Trees

When can one actually simulate local Hamiltonians or find their ground states efficiently? It is possible for some Hamiltonians in 1D and in tree geometry, by using an approximate method based on Matrix Product States. This is what I show in this Chapter, based on the paper [69]

The Quantum Transverse Field Ising Model on an Infinite Tree from Matrix Product States

Daniel Nagaj, Edward Farhi, Jeffrey Goldstone, Peter Shor and Igor Sylvester

We give a generalization to an infinite tree geometry of Vidal's infinite time-evolving block decimation (iTEBD) algorithm [106] for simulating an infinite line of quantum spins. We numerically investigate the quantum Ising model in a transverse field on the Bethe lattice using the Matrix Product State ansatz. We observe a second order phase transition, with certain key differences from the transverse field Ising model on an infinite spin chain. We also investigate a transverse field Ising model with a specific longitudinal field. When the transverse field is turned off, this model has a highly degenerate ground state as opposed to the pure Ising model whose ground state is only doubly degenerate.

This Chapter is organized as follows. Section 3.2 is a review of the MPS ansatz and contains its generalization to the tree geometry. In Section 3.3, I review the numerical procedure for unitary updates and give a recipe for applying imaginary time evolution within the MPS ansatz. In Section 3.4, I adapt Vidal's iTEBD method for simulating translationally invariant one-dimensional systems to systems with tree geometry. Section 3.5 contains my numerical results for the quantum Ising model in a transverse field for translationally invariant systems. In Section 3.5.1, I first test my method for the infinite line, and then in Section 3.5.2 I present new results for the infinite tree. I turn to the NOT model in Section 3.6 and show that my numerics work well for this system even when there is a high ground state degeneracy. In Section 3.7 I investigate the stability of my tree results and conjecture that they may be good approximations to a local description far from the boundary of a large finite tree system.

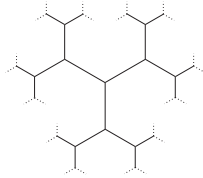


Figure 3-1: The Bethe lattice (infinite Cayley tree).

3.1 Introduction

The matrix product state (MPS) description [29][75] has brought a new way of approaching many-body quantum systems. Several methods of investigating spin systems have been developed recently combining state of the art many-body techniques such as White's Density Matrix Renormalization Group [110][83] (DMRG) with quantum information motivated insights. Vidal's Time Evolving Block Decimation (TEBD) algorithm [105] [104] uses MPS and emphasizes entanglement (as measured by the Schmidt number), directing the computational resources into that bottleneck of the simulation. It provides the ability to simulate time evolution and it was shown that MPS-inspired methods handle periodic boundary conditions well in one dimension [103], areas where the previous use of DMRG was limited. TEBD has been recast into the language of DMRG in [111][25] and adapted to finite systems with tree geometry in [89]. DMRG is especially successful in describing the properties of quantum spin chains, the application of basic DMRG-like methods is limited for quantum systems with higher dimensional geometry. New methods like PEPS [102][68] generalize MPS to higher dimensions, opening ways to numerically investigate systems that were previously inaccessible.

Here I am interested in investigating infinite translationally invariant systems. Several numerical methods to investigate these were developed recently. The iTEBD algorithm [106] (see also Sec.3.4) is a generalization of TEBD to infinite one-dimensional systems. A combination of PEPS with iTEBD called iPEPS [50] provides a possibility of investigating infinite translationally invariant systems in higher dimensions. My contribution is a method to investigate the ground state properties of infinite translationally invariant quantum systems on the Bethe lattice using imaginary time evolution with Matrix Product States. The Bethe lattice is an infinite tree with each node having three neighbors, as depicted in Fig.3-1. It is translationally invariant in that it looks the same at every vertex. This geometry is interesting, because of the following connection to large random graphs with fixed valence. Moving out from any vertex in such a random graph, you need to go a distance of order $\log n$, where n is the number of vertices in the graph, before you detect that you are not on the Bethe lattice, that is, before you see a loop.

I choose to investigate the quantum transverse field Ising model on the Bethe lattice. Note that I work directly on the infinite system, never taking a limit. First I test the iTEBD method on a system with a known exact solution, the infinite line. Then I turn to the Bethe lattice with the new method I provide. In both cases, the Hamiltonian is given by

$$H = \frac{J}{2} \sum_{\langle i,j \rangle} (1 - \sigma_z^i \sigma_z^j) + \frac{h}{2} \sum_i (1 - \sigma_x^i), \quad (3.1)$$

where the sum over i is over all sites, and the sum over $\langle i, j \rangle$ is over all bonds (nearest neighbors). I show that imaginary time evolution within the MPS ansatz provides a very good approximation for the exact ground state on an infinite line, resulting in nearly correct critical exponents for the magnetization and correlation length as one approaches the phase transition. I obtain new results for the quantum Ising model in transverse field on the infinite tree. Similarly to the infinite line, I observe a second order phase transition and obtain the critical exponent for the magnetization, $\beta_T \approx 0.41$ (different than the mean-field result). However, the correlation length does not diverge at the phase transition for this system and I conjecture that it has the value $1/\ln 2$.

I also investigate a model where besides an antiferromagnetic interaction of spins I add a specific longitudinal field $\frac{1}{4}\sigma_z^i$ for each spin:

$$H_{\text{NOT } 00} = J \sum_{\langle i, j \rangle} \frac{1}{4} (1 + \sigma_z^i + \sigma_z^j + \sigma_z^i \sigma_z^j) + \frac{h}{2} \sum_i (1 - \sigma_x^i). \quad (3.2)$$

I choose the longitudinal field in such a way that the interaction term in the computational basis takes a simple form, $|00\rangle \langle 00|_{ij}$, giving an energy penalty to the $|00\rangle$ state of neighboring spins. (I follow the usual convention that spin up in the z -direction is called 0.) I call it the NOT 00 model accordingly. This model is interesting from a computational viewpoint. The degeneracy of the ground state of $H_{\text{NOT } 00}$ at $h = 0$ is high for both infinite line and infinite tree geometry of interactions. I am interested in how my numerical method deals with this case, as opposed to the double degeneracy of the ground state of (3.1) at $h = 0$. I do not see a phase transition in this system as I vary J and h .

3.2 Matrix Product States

If one's goal is to numerically investigate a system governed by a local Hamiltonian, it is convenient to find a local description and update rules for the system. A Matrix Product State description is particularly suited to spin systems for which the connections do not form any loops. Given a state of this system, I will first show how to obtain its MPS description, and then how to utilize this description in a numerical method for obtaining the time evolution and approximating the ground state (using imaginary time evolution). I begin with matrix product states on a line (a spin chain), and then generalize the description to a tree geometry. In 3.3, I give a numerical method of updating the MPS description for both real and imaginary time simulations.

3.2.1 MPS for a spin chain

Given a state $|\psi\rangle$ of a chain of n spins

$$|\psi\rangle = \sum_{\dots s_i s_{i+1} \dots} c_{\dots, s_i, s_{i+1}, \dots} |s_1\rangle_1 \dots |s_i\rangle_i |s_{i+1}\rangle_{i+1} \dots |s_n\rangle_n, \quad (3.3)$$

I wish to rewrite the coefficients c_{s_1, \dots, s_n} as a matrix product (see [77] for a review of MPS)

$$c_{\dots, s_i, s_{i+1}, \dots} = \sum_{\dots abc \dots} \dots \lambda_a^{(i-1)} \Gamma_{a,b}^{(i), s_i} \lambda_b^{(i)} \Gamma_{b,c}^{(i+1), s_{i+1}} \lambda_c^{(i+1)} \dots \quad (3.4)$$

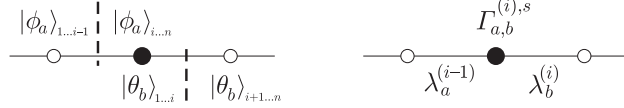


Figure 3-2: Two successive Schmidt decompositions on a line allow me to find the Γ tensor for the marked site and the two λ vectors for the bonds coming out of it.

using n tensors $\Gamma^{(i)}$ and $n-1$ vectors $\lambda^{(i)}$. The range of the indices a, b, \dots will be addressed later. After decomposing the chain into two subsystems, one can rewrite the state of the whole system in terms of orthonormal bases of the subsystems. $\lambda^{(i)}$ is the vector of Schmidt coefficients for the decomposition of the state of the chain onto the subsystems $1 \dots i$ and $i+1 \dots n$.

In order to obtain the λ 's and the Γ 's for a given state $|\psi\rangle$, one has to perform the following steps. First, perform the Schmidt decomposition of the chain between sites $i-1$ and i as

$$|\psi\rangle = \sum_{a=1}^{\chi_{i-1}} |\phi_a\rangle_{1,\dots,i-1} \lambda_a^{(i-1)} |\phi_a\rangle_{i,\dots,n}, \quad (3.5)$$

where the states on the left and on the right of the division form orthonormal bases required to describe the respective subsystems of the state $|\psi\rangle$. The number χ_{i-1} (the Schmidt number) is the minimum number of terms required in this decomposition. The Schmidt decomposition for a split between sites i and $i+1$ gives

$$|\psi\rangle = \sum_{b=1}^{\chi_i} |\theta_b\rangle_{1,\dots,i} \lambda_b^{(i)} |\theta_b\rangle_{i+1,\dots,n}. \quad (3.6)$$

These two decompositions (see FIG.3-2) describe the same state, allowing me to combine them to express the basis of the subsystem i, \dots, n using the spin at site i and the basis of the subsystem $i+1, \dots, n$ as

$$|\phi_a\rangle_{i,\dots,n} = \sum_{s=0,1} \sum_{b=1}^{\chi_i} \Gamma_{a,b}^{(i),s} \lambda_b^{(i)} |s\rangle_i |\theta_b\rangle_{i+1,\dots,n}, \quad (3.7)$$

where I inserted the $\lambda_b^{(i)}$ for convenience. This gives me the tensor $\Gamma^{(i)}$. It carries an index s corresponding to the state $|s\rangle$ of the i -th spin, and indices a and b , corresponding to the two consecutive divisions of the system (see FIG.3-2). Because $|\phi_a\rangle$ (and $|\theta_b\rangle$) are orthonormal states, the vectors λ and tensors Γ obey the following normalization conditions. From (3.6) I have

$$\sum_{b=1}^{\chi_i} \lambda_b^{(i)2} = 1, \quad (3.8)$$

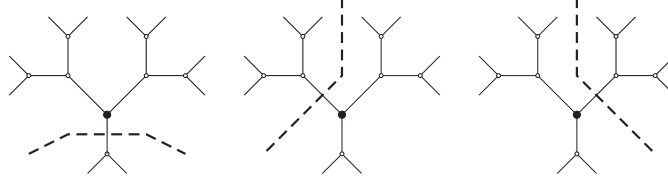


Figure 3-3: The three Schmidt decompositions on a tree required to obtain the Γ tensor for the marked site and the three λ vectors for the bonds emanating from it.

while (3.7) implies

$$\langle \phi_{a'} | \phi_a \rangle_{i, \dots, n} = \sum_{s=0,1} \sum_{b=1}^{\chi_i} \Gamma_{a',b}^{(i),s*} \lambda_b^{(i)} \Gamma_{a,b}^{(i),s} \lambda_b^{(i)} = \delta_{a,a'} , \quad (3.9)$$

and

$$\langle \theta_{b'} | \theta_b \rangle_{1, \dots, i} = \sum_{s=0,1} \sum_{a=1}^{\chi_{i-1}} \lambda_a^{(i-1)} \Gamma_{a,b'}^{(i),s*} \lambda_a^{(i-1)} \Gamma_{a,b}^{(i),s} = \delta_{b,b'} . \quad (3.10)$$

3.2.2 MPS on Trees

Matrix Product States are natural not just on chains, but also on trees, because these can also be split into two subsystems by cutting a single bond, allowing for the Schmidt-decomposition interpretation as described in the previous section. The Matrix Product State description of a state of a spin system on a tree, i.e. such that the bonds do not form loops, is a generalization of the above procedure. Tree-tensor-network descriptions such as the one given here have been previously described in [89].

Specifically, for the Bethe lattice with 3 neighbors per spin, I introduce a vector $\lambda_{a_k}^{(k)}$ for each bond k and a four-index (one for spin, three for bonds) tensor $\Gamma_{a_k, a_l, a_m}^{(i), s_i}$ for each site i . I can then rewrite the state $|\psi\rangle$ analogously to (3.3), (3.4) as

$$|\psi\rangle = \left(\prod_{k \in \text{bonds}} \sum_{a_k=1}^{\chi_k} \lambda_{a_k}^{(k)} \right) \left(\prod_{i \in \text{sites}} \sum_{s_i} \Gamma_{a_l, a_m, a_n}^{(i), s_i} \right) |\dots\rangle |s_i\rangle |\dots\rangle , \quad (3.11)$$

where a_l, a_m, a_n are indices corresponding to the three bonds l, m and n coming out of site i . Each index a_l appears in two Γ tensors and one λ vector. To obtain this description, one needs to perform a Schmidt decomposition across each bond. This produces the vectors $\lambda^{(l)}$. To obtain the tensor $\Gamma^{(i)}$ for site i , one needs to combine the three decompositions corresponding to the bonds of site i as depicted in Fig.3-3. Analogously to (3.7), expressing the orthonormal basis for the first subsystem marked in Fig.3-3 in terms of the state of the spin $|s_i\rangle$ and the orthonormal bases for the latter two subsystems in Fig.3-3, one obtains the tensor $\Gamma_{a_l, a_m, a_n}^{(i), s}$ for site i .

The normalization conditions for a MPS description of a state on a tree are analogous

to (3.8)-(3.10). I have

$$\sum_{a_k} \lambda_{a_k}^{(k)2} = 1, \quad (3.12)$$

$$\sum_{s=0,1} \sum_{a_k=1}^{\chi_k} \sum_{a_l=1}^{\chi_l} \Gamma_{a_k, a_l, a_{m'}}^{(i), s*} \lambda_{a_k}^{(k)2} \lambda_{a_l}^{(l)2} \Gamma_{a_k, a_l, a_m}^{(i), s} = \delta_{a_m, a_{m'}}, \quad (3.13)$$

and two other variations of (3.13) with k, l and m interchanged.

3.3 Simulating Quantum Systems with MPS

I choose to first describe the numerical procedures for a chain of spins. Then, at the end of the respective subsections, I note how to generalize these to tree geometry.

3.3.1 Unitary Update Rules

The strength of the MPS description of the state lies in the efficient application of local unitary update rules such as $U = e^{-iA\Delta t}$ (where A is an operator acting only on a few qubits). First, I describe the numerical procedure in some detail, and then, in the next Section, discuss how to modify the procedure to also implement imaginary time evolution.

Given a state $|\psi\rangle$ as a Matrix Product State, I want to know what happens after an application of a local unitary. In particular, for a 1-local U acting on the i -th spin, it suffices to update the local tensor

$$\Gamma_{a,b}^{(i),s} \xrightarrow{U} U_{s'}^s \Gamma_{a,b}^{(i),s'}. \quad (3.14)$$

The update rule for an application of a 2-local unitary V acting on neighboring spins i and $i+1$, requires several steps. First, using a larger tensor

$$\Theta_{a,c}^{s,t} = \lambda_a^{(i-1)} \sum_b \left(\Gamma_{a,b}^{(i),s} \lambda_b^{(i)} \Gamma_{b,c}^{(i+1),t} \right) \lambda_c^{(i+1)}, \quad (3.15)$$

I rewrite the state $|\psi\rangle$ as

$$|\psi\rangle = \sum_{a,c} \sum_{s,t} \Theta_{a,c}^{s,t} |\phi_a\rangle_{1\dots i-1} |s\rangle_i |t\rangle_{i+1} |\phi_c\rangle_{i+1\dots n}. \quad (3.16)$$

After the application of V , the tensor Θ in the description of $|\psi\rangle$ changes as

$$\Theta_{a,c}^{s,t} \xrightarrow{V} \sum_{s',t'} V_{s',t'}^{s,t} \Theta_{a,c}^{s',t'}. \quad (3.17)$$

One now needs to decompose the updated tensor Θ to obtain the updated tensors $\Gamma^{(i)}$, $\Gamma^{(i+1)}$ and the vector $\lambda^{(i)}$. I use the indices a, s and c, t of Θ to introduce combined indices (as) and (ct) and form a matrix $T_{(as),(ct)}$ with dimensions $2\chi_{i-1} \times 2\chi_{i+1}$ as

$$T_{(as),(ct)} = \Theta_{a,c}^{s,t}. \quad (3.18)$$

Using the singular value decomposition (SVD), this matrix can be decomposed into $T = Q\Lambda W$, where Q and W are unitary and Λ is a diagonal matrix. In terms of matrix elements, this reads

$$T_{(as),(ct)} = \sum_b Q_{(as),b} D_{b,b} W_{b,(ct)}. \quad (3.19)$$

The diagonal matrix $D = \text{diag}(\lambda^{(i)})$ gives me the updated Schmidt vector $\lambda^{(i)}$. The updated tensors $\Gamma^{(i)}$ and $\Gamma^{(i+1)}$ can be obtained from the matrices Q , W and the definition of Θ (3.15) using the old vectors $\lambda^{(i-1)}$ and $\lambda^{(i+1)}$ which do not change with the application of the local unitary V . After these update procedures, the conditions (3.8)-(3.10) are maintained.

The usefulness/succintness of this description depends crucially on the amount of entanglement across the bipartite divisions of the system as measured by the Schmidt numbers χ_i . To exactly describe a general quantum state $|\psi\rangle$ of a chain of n spins, the Schmidt number for the split through the middle of the chain is necessarily $\chi_{n/2} = 2^{n/2}$. Suppose I start my numerical simulation in a state that is exactly described by a MPS with only low χ_i 's. The update step described above involves an interaction of two sites, and thus could generate more entanglement across the $i, i+1$ division. After the update, the index b in $\lambda_b^{(i)}$ would need to run from 1 to $2\chi_i$ to keep the description exact (unless χ_i already is at its maximum required value $\chi_i = 2^{\min\{i, n-i\}}$). This makes the number of parameters in the MPS description grow exponentially with the number of update steps.

So far, this description and update rules have been exact. Let me now make the description an approximate one (use a block-decimation step) instead. First, introduce the parameter χ , which is the maximum number of Schmidt terms I keep after each update step. If the amount of entanglement in the system is low, the Schmidt coefficients $\lambda_b^{(i)}$ decrease rapidly with b (I always take the elements of λ sorted in decreasing order). A MPS ansatz with restricted $\chi_i = \chi$ will hopefully be a good approximation to the exact state $|\psi\rangle$. However, I also need to keep the restricted χ throughout the simulation. After a two-local unitary update step, the vector $\lambda^{(i)}$ can have 2χ entries. However, if the $b > \chi$ entries in $\lambda_b^{(i)}$ after the update are small, I am justified to truncate $\lambda^{(i)}$ to have only χ entries and multiply it by a number so that it satisfies (3.8). I also truncate the Γ tensors so that they keep dimensions $2 \times \chi \times \chi$. The normalization condition (3.10) for $\Gamma^{(i)}$ will be still satisfied exactly, while the error in the normalization condition (3.9) will be small. This normalization error can be corrected as discussed in the next section. This procedure keeps the state within the MPS ansatz with restricted χ .

The procedure described above allows me to efficiently approximately implement local unitary evolution. To simulate time evolution

$$|\psi(t)\rangle = e^{-iHt} |\psi(0)\rangle, \quad (3.20)$$

with a local Hamiltonian like (3.1), I first divide the time t into small slices Δt and split the Hamiltonian into two groups of commuting terms $H_k^{(x)}$ and $H_m^{(z)}$. Each time evolution step $e^{-iH\Delta t}$ can then be implemented as a product of local unitaries using the second order Trotter-Suzuki formula

$$U_2 = \left(\prod_k e^{-iH_k^{(x)} \frac{\Delta t}{2}} \right) \left(\prod_m e^{-iH_m^{(z)} \Delta t} \right) \left(\prod_k e^{-iH_k^{(x)} \frac{\Delta t}{2}} \right). \quad (3.21)$$

The application of the product of the local unitaries within each group can be done almost in parallel (in two steps, as described in Section 3.4), as they commute with each other.

These update rules allow me to efficiently approximately simulate the real time evolution (3.20) with a local Hamiltonian H for a state $|\psi\rangle$ within the MPS ansatz with parameter χ . The number of parameters in this MPS description with restricted χ is then $n(2\chi^2)$ for the tensors $\Gamma^{(i)}$ and $(n-1)\chi$ for the vectors $\lambda^{(i)}$. The simulation cost of each local update step scales like $O(\chi^3)$, coming from the SVD decomposition of the matrix Θ . For a system of n spins, I thus need to store $O(2n\chi^2 + n\chi)$ numbers and each update will take $O(n\chi^3)$ steps.

The update procedure generalizes to tree geometry by taking the tensors Γ with dimensions $2 \times \chi \times \chi \times \chi$ as in Section 3.2.2. For a local update (on two neighboring spins i and $i+1$ with bonds labeled by l, m, n and n, o, p) I rewrite the state $|\psi\rangle$ analogously to (3.16) as

$$|\psi\rangle = \sum_{a_k, a_l, a_o, a_p} \sum_{s, t} \Theta_{(a_k a_l), (a_o a_p)}^{s, t} |\phi_{a_k}\rangle |\phi_{a_l}\rangle |s\rangle_i |t\rangle_{i+1} |\phi_{a_o}\rangle |\phi_{a_p}\rangle. \quad (3.22)$$

using the tensor

$$\Theta_{(a_k a_l), (a_o a_p)}^{s, t} = \lambda_{a_k}^{(k)} \lambda_{a_l}^{(l)} \sum_{a_m} \left(\Gamma_{a_k, a_l, a_m}^{(A), s} \lambda_{a_m}^{(m)} \Gamma_{a_m, a_o, a_p}^{(i+1), t} \right) \lambda_{a_o}^{(o)} \lambda_{a_p}^{(p)}, \quad (3.23)$$

with combined indices $(a_k a_l)$ and $(a_o a_p)$. One then needs to update the tensor Θ as described above (3.17)-(3.19). The decomposition procedure to get the updated vector $\lambda^{(m)}$ and the new tensors $\Gamma^{(i)}$ and $\Gamma^{(i+1)}$ now requires $O(\chi^6)$ computational steps. The cost of a simulation on n spins thus scales like $O(n\chi^6)$.

3.3.2 Imaginary Time Evolution

Using the MPS ansatz, I can also use imaginary time evolution with e^{-Ht} instead of (3.20) to look for the ground state of systems governed by local Hamiltonians. One needs to replace each unitary term $e^{-iA\Delta t}$ in the Trotter expansion (3.21) of the time evolution with $e^{-A\Delta t}$ followed by a normalization procedure. However, the usual normalization procedure for imaginary time evolution (multiplying the state by a number to keep $\langle\psi|\psi\rangle = 1$) is now not enough to satisfy the MPS normalization conditions (3.8)-(3.10) for the tensors Γ and vectors λ I use to describe the state $|\psi\rangle$.

The unitarity of the real time evolution automatically implied that the normalization conditions (3.8),(3.10) were satisfied after an exact unitary update. While there already was an error in (3.9) introduced by the truncation of the $\chi + 1 \dots 2\chi$ entries in $\Gamma^{(i)}$, the non-unitarity of imaginary time evolution update steps introduces further normalization errors. It is thus important to properly normalize the state after every application of terms like $e^{-A\Delta t}$ to keep it within the MPS ansatz.

In [106], Vidal dealt with this problem by taking progressively shorter and shorter steps Δt during the imaginary time evolution. This procedure results in a properly normalized state only at the end of the evolution, after the time step decreases to zero (and not necessarily during the evolution). I propose a different scheme in which I follow each local update $e^{-A\Delta t}$ by a normalization procedure (based on Vidal's observation) to bring the state back to the MPS ansatz at all times. The simulation I run (evolution for time t) thus consists of many short time step updates $e^{-H\Delta t}$, each of which is implemented using

a Trotter expansion as a product of local updates $e^{-A\Delta t}$. Each of these local updates is followed by my normalization procedure.

I now describe the iterative normalization procedure in detail for the case of an infinite chain, where it can be applied efficiently, as the description of the state $|\psi\rangle$ requires only two different tensors Γ (see Section 3.4.1). One needs to apply the following steps over and over, until the normalization conditions are met with chosen accuracy.

First, for each nearest neighbor pair $i, i+1$ with even i , combine the MPS description of these two spins (3.15)-(3.16), forming the matrix T (3.18). Do a SVD decomposition of T (3.19) to obtain a new vector $\lambda^{(i)}$. The decomposition does not increase the number of nonzero elements of $\lambda^{(i)}$, as the rank of the $2\chi \times 2\chi$ matrix T (3.19) was only χ (coming from (3.15)). I thus take only the first χ values of $\lambda^{(i)}$ and rescale the vector to obey $\sum_{a=1}^{\chi} \lambda_a^{(i)2} = 1$. using this new $\lambda^{(i)}$, I obtain tensors $\Gamma^{(i)}$, $\Gamma^{(i+1)}$ from (3.19), and truncate them to have dimensions $\chi \times \chi \times 2$. Second, I repeat the previous steps for all nearest neighbor pairs of spins $i, i+1$ with i odd.

I observe that repeating the above steps over and over results in exponential decrease in the error in the normalization of the Γ tensors. Note though, that the rate of decrease in normalization errors becomes much slower near the phase transition for the transverse field Ising model on an infinite line (see Section 3.5.1).

In practice, I apply this normalization procedure by using the same subroutine for the local updates $e^{-A\Delta t}$, except that I skip the step (3.17), which is equivalent to applying the local update with $\Delta t = 0$. The normalization procedure is thus equivalent to evolving the state repeatedly with zero time step (composing two tensors Γ and decomposing them again) and imposing the normalization condition on the vectors λ . Note though, following from the definition of the SVD, that each decomposition assures that one of the conditions (3.9),(3.10) is retained exactly for the updated tensors Γ . The errors in the other normalization condition for the Γ tensors are decreased in each iteration step.

The numerical update rules for a system with tree geometry are a simple analogue of the update rules for MPS on spin chains. Every interaction couples two sites, with tensors $\Gamma_{a,b,c}^{(A),s}$ and $\Gamma_{c,d,e}^{(B),t}$, with the three lower indices corresponding to the bonds emanating from the sites. One only needs to reshape the tensors into $\Gamma_{(ab),c}^{(A),s}$ and $\Gamma_{c,(de)}^{(B),t}$ and proceed as described in (3.15) and below.

3.4 MPS and Translationally Invariant Systems

3.4.1 An Infinite Line

For systems with translational symmetry such as an infinite line all the sites are equivalent. I assume that the ground state is translationally invariant, and furthermore pick the tensors $\Gamma^{(i)}$ and vectors $\lambda^{(i)}$ to be site independent. For fixed χ the number of complex parameters in the translationally invariant MPS ansatz on the infinite line scales as $2\chi^2$.

When using imaginary time evolution to look for the ground state of this system, within this ansatz, it is technically hard to keep the translational symmetry and the normalization conditions after each update. Numerical instabilities plagued my efforts to impose the symmetry in the procedures described above. In [106], Vidal devised a method to deal with this problem. Let us break the translational symmetry of the ansatz by labeling the sites A and B as in FIG.3-4. This doubles the number of parameters in the ansatz. The state update now proceeds in two steps. Let the site pairs AB interact and update the tensors

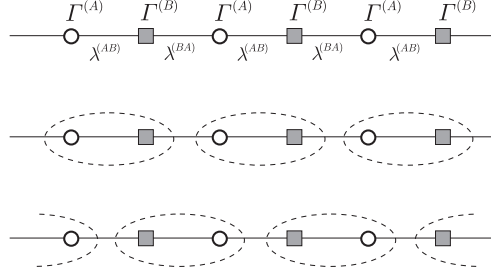


Figure 3-4: The parametrization and update rules for the infinite line.

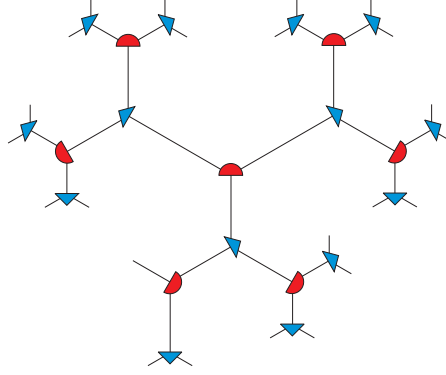


Figure 3-5: The two-layer, directed labeling of the tree.

$\Gamma^{(A)}$, $\Gamma^{(B)}$ and the vector $\lambda^{(AB)}$. Then let the neighbor pairs BA interact, after which I update the tensors $\Gamma^{(B)}$, $\Gamma^{(A)}$ and the vector $\lambda^{(BA)}$. What I observe is that after many state updates the elements of the resulting $\Gamma^{(A)}$ and the $\Gamma^{(B)}$ tensors differ at a level which is way below my numerical accuracy (governed by the normalization errors) and I am indeed obtaining a translationally invariant description of the system.

One of the systems easily investigated with this method (iTEBD) is the Ising model in a transverse field (5.37) on an infinite line. Vidal’s numerical results for the real time evolution and imaginary time evolution [106] of this system show remarkable agreement with the exact solution. I take a step further and also numerically obtain the critical exponents for this system. Further details can be found in Section 3.5, where I compare these results for the infinite line to the results I obtain for the Ising model in transverse field on the Bethe lattice.

3.4.2 An Infinite Tree

For the infinite Bethe lattice, My approach is a modification of the above procedure introduced by Vidal. In order to avoid the numerical instabilities associated with imposing site-independent Γ and λ after the update steps, I break the translational symmetry by labeling the “layers” of the tree A and B (denoted by half-circles and triangles), as in FIG.3-5.

The Bethe lattice is also symmetric under the permutation of directions. Tensors Γ with

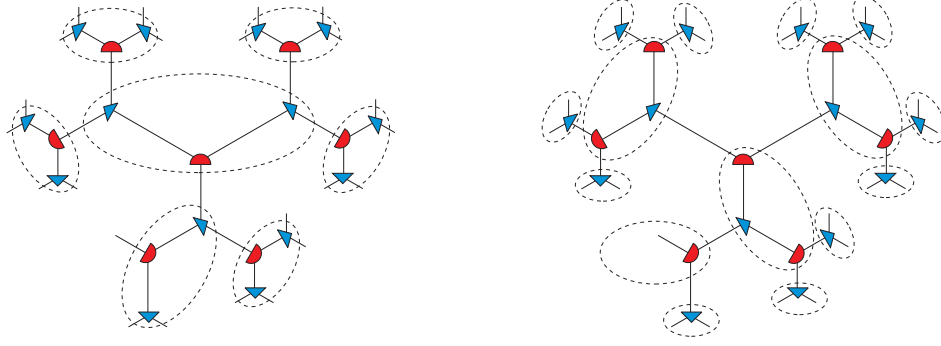


Figure 3-6: The two-step interactions for the infinite tree.

full directional symmetry obey $\Gamma_{a,b,c} = \Gamma_{b,c,a} = \Gamma_{c,a,b} = \Gamma_{c,b,a} = \Gamma_{b,a,c} = \Gamma_{a,c,b}$. However, for the purpose of simple organization of interactions, I will also partially break this symmetry by consistently labeling an ‘inward’ bond for each node, as denoted by the flat sides of the semi-circles and the longer edges of the triangles in Fig.3-5. This makes the first of the three indices of $\Gamma_{a,b,c}$ special. However, I keep the residual symmetry $\Gamma_{a,b,c} = \Gamma_{a,c,b}$. I can enforce this by interacting a spin with both of the spins from the next layer at the same time. The update procedure for the interaction between the spins now splits into two steps, interacting the layers in the AB order first, and then in the BA order as in Fig.3-6. Similarly to what I discovered for the line, the differences in the elements of the final $\Gamma^{(A)}$ and $\Gamma^{(B)}$ are well below the numerical accuracy of my procedure.

The scaling of this procedure is more demanding than the $O(\chi^3)$ simulation for a line. The number of entries in the matrix Θ used in each update step is $2\chi^2 \times 4\chi^4$, therefore the SVD decomposition requires $O(\chi^8)$ steps. The scaling of my numerical method is thus $O(\chi^8)$ for each update step.

3.4.3 Expectation Values

A nice property of the MPS state description is that it allows efficient computation of expectation values of local operators. First, for a translationally invariant system on a line (with only one tensor Γ and one vector λ), I have for an operator $O^{(i)}$ acting only on the i -th spin

$$\langle \psi | O^{(i)} | \psi \rangle = \sum_{s_i, s'_i=0,1} O_{s_i, s'_i}^{(i)} \sum_{a=1}^{\chi} \sum_{b=1}^{\chi} (\lambda_a \Gamma_{a,b}^{s'_i*} \lambda_b) (\lambda_a \Gamma_{a,b}^{s_i} \lambda_b), \quad (3.24)$$

where $O_{s_i, s'_i}^{(i)} = \langle s'_i | O^{(i)} | s_i \rangle$. Similarly, for the expectation values of $O^{(i)} O^{(j)}$ (assuming $j > i$),

$$\begin{aligned} \langle \psi | O^{(i)} O^{(j)} | \psi \rangle &= \sum_{s_i, s'_i, \dots, s_j, s'_j} \sum_{a,e} \sum_{b,b',\dots} O_{s_i, s'_i}^{(i)} O_{s_j, s'_j}^{(j)} \\ &\times (\lambda_a \Gamma_{a,b'}^{s'_i*} \lambda_{b'} \Gamma_{b',c'}^{s_i+1*} \lambda_{c'} \dots \lambda_{d'} \Gamma_{d',e}^{s'_j*} \lambda_e) \\ &\times (\lambda_a \Gamma_{a,b}^{s_i} \lambda_b \Gamma_{b,c}^{s_i+1} \lambda_c \dots \lambda_d \Gamma_{d,e}^{s_j} \lambda_e). \end{aligned} \quad (3.25)$$

Defining a $\chi^2 \times \chi^2$ matrix B (where one should think of (bb') as one combined index ranging from 1 to χ^2) as

$$B_{(bb'),(cc')} = \sum_s \Gamma_{b,c}^s \Gamma_{b',c'}^{s*} \lambda_c \lambda_{c'}, \quad (3.26)$$

and vectors v and w with elements again denoted by a combined index $(bb') = 1 \dots \chi^2$ as

$$v_{(bb')} = \sum_{s_i, s'_i} O_{s_i, s'_i}^{(i)} \sum_a (\lambda_a)^2 \Gamma_{a,b}^{s_i} \Gamma_{a,b'}^{s'_i*} \lambda_b \lambda_{b'}, \quad (3.27)$$

$$w_{(dd')} = \sum_{s_j, s'_j} O_{s_j, s'_j}^{(j)} \sum_e \Gamma_{d,e}^{s_j} \Gamma_{d',e}^{s'_j*} (\lambda_e)^2, \quad (3.28)$$

I can rewrite (3.25) as

$$\langle \psi | O^{(i)} O^{(j)} | \psi \rangle = v^T \underbrace{B B \dots B}_{j-i-1} w, \quad (3.29)$$

There is a relationship between the eigenvalues of the matrix B and the correlation function $\langle O^{(i)} O^{(j)} \rangle - \langle O^{(i)} \rangle \langle O^{(j)} \rangle$. One of the eigenvalues of B is $\mu_1 = 1$, with the corresponding right eigenvector

$$\beta_{(cc')}^{(1R)} = \sum_s \sum_c \Gamma_{b,c}^s \Gamma_{b',c}^{s*} (\lambda_c)^2, \quad (3.30)$$

and left eigenvector

$$\beta_{(bb')}^{(1L)} = \delta_{b,b'} \lambda_b^2, \quad (3.31)$$

which can be verified using the normalization conditions (3.9) and (3.10). I numerically observe that $\mu_1 = 1$ is also the largest eigenvalue. (Note that $|\mu_k| > 1$ would result in correlations unphysically growing with distance.) Denote the second largest eigenvalue of B as μ_2 . Using the eigenvectors of B , I can express B^{j-i-1} in (3.29), as

$$B^{j-i-1} = \beta^{(1L)} \beta^{(1R)T} + \mu_2^{j-i-1} \beta^{(2L)} \beta^{(2R)T} + \dots \quad (3.32)$$

When computing the correlation function, the term that gets subtracted exactly cancels the leading term involving $\mu_1 = 1$. Therefore, if $|\mu_2|$ is less than 1, (3.32) implies

$$\langle O^{(i)} O^{(j)} \rangle - \langle O^{(i)} \rangle \langle O^{(j)} \rangle \propto \mu_2^{|j-i|}. \quad (3.33)$$

The correlation function necessarily falls off exponentially in this case, and the correlation length ξ is related to μ_2 as $\xi = -1/\ln \mu_2$.

The computation of expectation values for a MPS state on a system with a tree geometry can be again done efficiently. For single-site operators $O^{(i)}$, the formula is an analogue of (3.24) with three λ vectors for each Γ tensors which now have three lower indices. For

two-site operators, the terms in (3.29) now become

$$B_{(cc'),(dd')} = \sum_s \sum_e \Gamma_{c,e,d}^s \Gamma_{c',e,d'}^{s*} (\lambda_e)^2 \lambda_d \lambda_{d'}, \quad (3.34)$$

$$v_{(cc')} = \sum_{s_i, s'_i} O_{s_i, s'_i}^{(i)} \sum_{a,b} (\lambda_a)^2 (\lambda_b)^2 \Gamma_{a,b,c}^{s_i} \Gamma_{a,b,c'}^{s'_i*} \lambda_c \lambda_{c'}, \quad (3.35)$$

$$w_{(dd')} = \sum_{s_j, s'_j} O_{s_j, s'_j}^{(j)} \sum_{e,f} \Gamma_{d,e,f}^{s_j} \Gamma_{d',e,f}^{s'_j*} (\lambda_e)^2 (\lambda_f)^2. \quad (3.36)$$

The correlation length is again related to the second eigenvalue of the B matrix as in (3.33).

3.5 Quantum Transverse Field Ising Model

My goal is to investigate the phase transition for the Ising model in transverse magnetic field (3.1) on the infinite line and on the Bethe lattice. I choose to parametrize the Hamiltonian as

$$H = \frac{s}{2} \sum_{\langle i,j \rangle} (1 - \sigma_z^i \sigma_z^j) + \frac{b(1-s)}{2} \sum_i (1 - \sigma_x^i), \quad (3.37)$$

where $0 \leq s \leq 1$ and b is the number of bonds for each site ($b = 2$ for the line, $b = 3$ for the tree). I will investigate the ground state properties of (3.37) as we vary s . The point $s = 0$ corresponds to a spin system in transverse magnetic field, while $s = 1$ corresponds to a purely ferromagnetic interaction between the spins.

3.5.1 The Infinite Line.

I present the results for the case of an infinite line and compare them to exact results obtained via fermionization (see e.g. [82], Ch.4). Vidal has shown [106] that imaginary time evolution within the MPS ansatz is capable of providing a very accurate approximation for the ground state energy and correlation function. I show that even using χ smaller than used in [106], I obtain the essential information about the nature of the phase transition in the infinite one-dimensional system. I also obtain the critical exponents for the magnetization and the correlation length.

In FIG.3-7, I show the how the ground state energy obtained using imaginary time evolution with MPS converges to the exact energy as χ increases. The exact solution for a line has a second order phase transition at the critical value of s , $s_L = \frac{2}{3}$, and the ground state energy and its first derivative are continuous, while the second derivative diverges at $s = s_L$. I plot the first and second derivative of E with respect to s obtained numerically and compare them to the exact values in FIG. 3-8, observing the expected behavior already for low χ .

The derivative of the exact magnetization $M = \langle \sigma_z \rangle$ is discontinuous at s_L , with the magnetization starting to rise steeply from zero as

$$M \propto (x - x_L)^\beta, \quad (3.38)$$

with the critical exponent $\beta_L = \frac{1}{8}$. Here x is the ratio of the ferromagnetic interaction

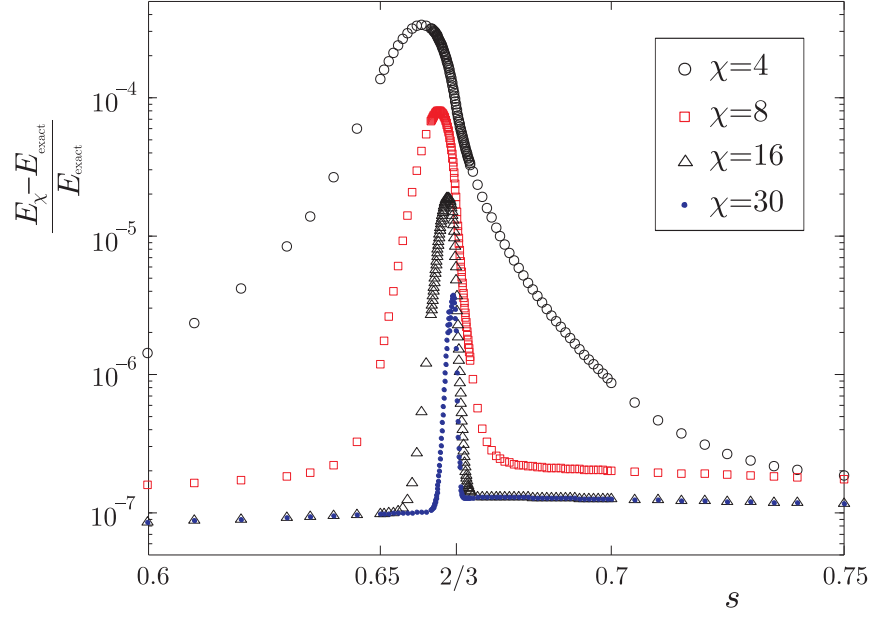


Figure 3-7: Transverse Ising model on an infinite line. Fractional difference of the ground state energy obtained using MPS and the exact ground state, near the phase transition at $s_L = \frac{2}{3}$. The energy scale is logarithmic.

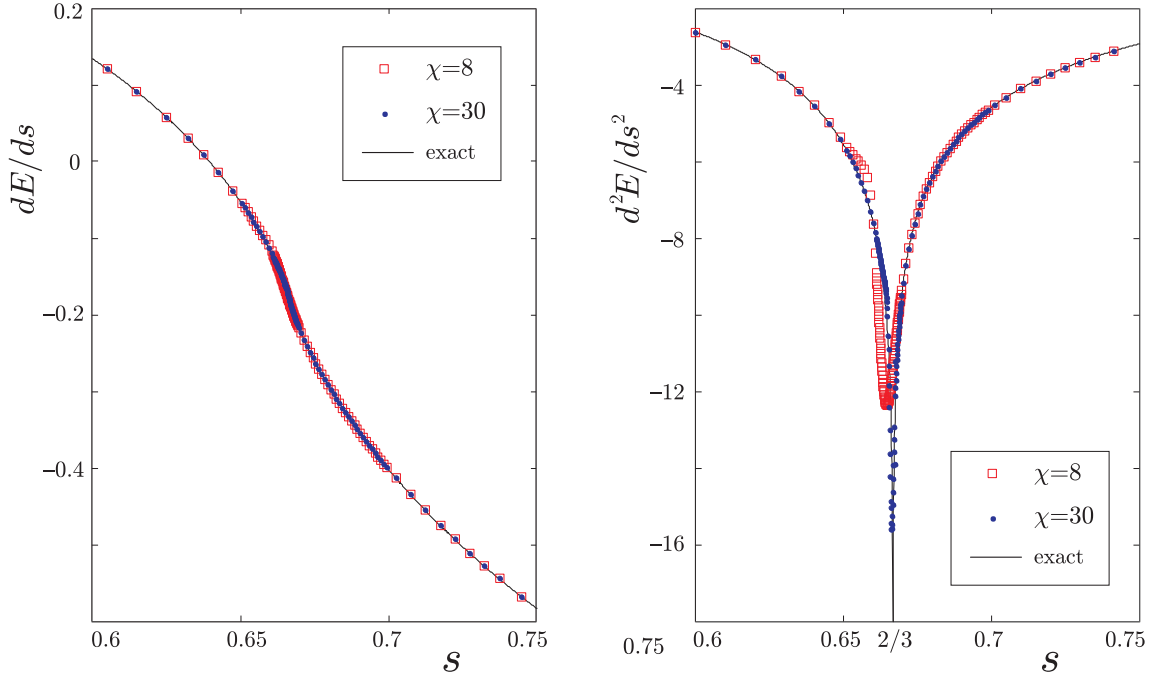


Figure 3-8: Transverse Ising model on an infinite line. The first and second derivative with respect to s of the ground state energy obtained via MPS compared with the exact result.

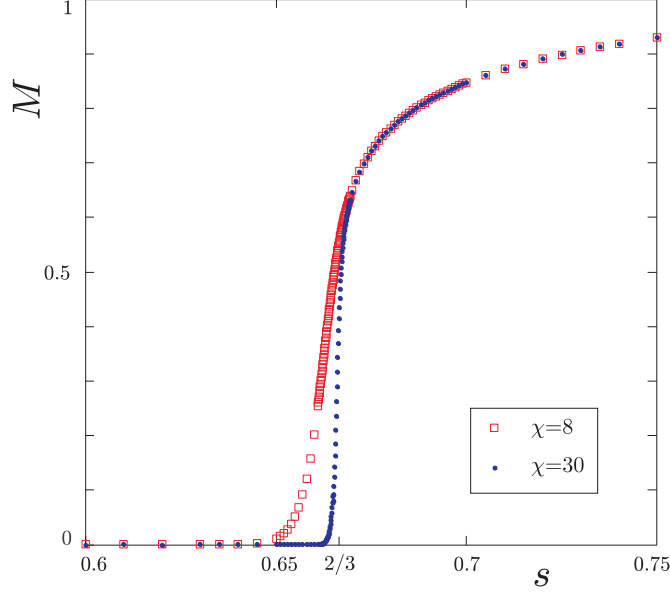


Figure 3-9: Transverse Ising model on an infinite line. Magnetization obtained using MPS vs s , with $B_z = 10^{-8}$.

strength to the transverse field strength in (3.37) and is

$$x = \frac{s}{2(1-s)}, \quad (3.39)$$

with the value $x = x_L = 1$ at the phase transition ($s_L = \frac{2}{3}$). I plot the magnetization obtained with my method in FIG.3-9. To obtain the magnetization depicted in the plot, I used a small symmetry breaking longitudinal field with magnitude $B_z = 10^{-8}$. In FIG.3-10, I plot M vs. $x - x_L$ on a log-log scale. I also plot a line with slope 0.125. Observe that as χ increases, the data is better represented by a line down to smaller values of $x - x_L$. For the largest χ displayed, a straight line fit of the data between $4 \times 10^{-3} \leq x - x_L \leq 10^{-1}$ gives me a slope of 0.120. Note that the mean field value of the critical exponent for magnetization is 0.5.

The correlation function $\langle \sigma_z^{(i)} \sigma_z^{(j)} \rangle - \langle \sigma_z^{(i)} \rangle \langle \sigma_z^{(j)} \rangle$ can be computed efficiently using (3.25). Away from criticality, it falls off exponentially as $e^{-|i-j|/\xi}$. The falloff of the correlation function is necessarily exponential as long as μ_2 , the second eigenvalue of B , is less than 1. The exact solution for the correlation length ξ near the critical point has the form

$$\xi \propto |x - x_L|^{-1} \quad (3.40)$$

as x approaches x_L . Already at low χ the iTEBD method captures the divergence of the correlation length. In FIG. 3-11, I plot ξ vs. $|x - x_L|$ on a log-log plot, together with a line with slope -1 . Again, as χ increases, the data is better represented by a line closer to the phase transition. For the highest χ displayed, a straight line fit of the data between $2 \times 10^{-2} \leq x_L - x \leq 4 \times 10^{-1}$ gives me a slope of -0.92 . Note that the mean-field value of the critical exponent for the correlation length is 0.5 (corresponding to slope -0.5 in the

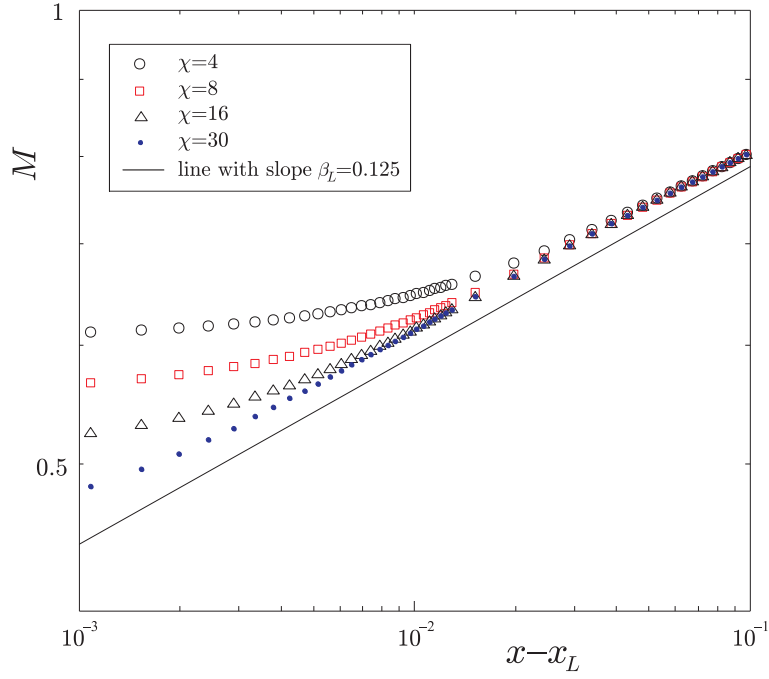


Figure 3-10: Transverse Ising model on an infinite line. Log-log plot of magnetization vs. $x - x_L$. I also plot a line with slope $\beta_L = 0.125$.

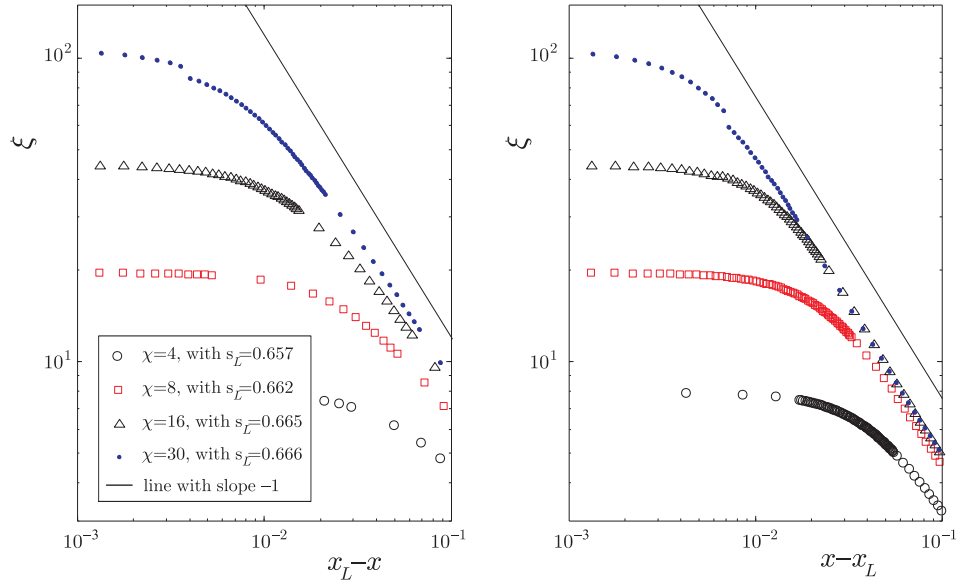


Figure 3-11: Transverse Ising model on an infinite line. Log-log plot of the correlation length vs. $|x - x_L|$. I also plot a line with slope $-\nu_L = -1$. For each χ in the plot, I choose a numerical value of the critical point s_L as the point at which the correlation length is maximal.

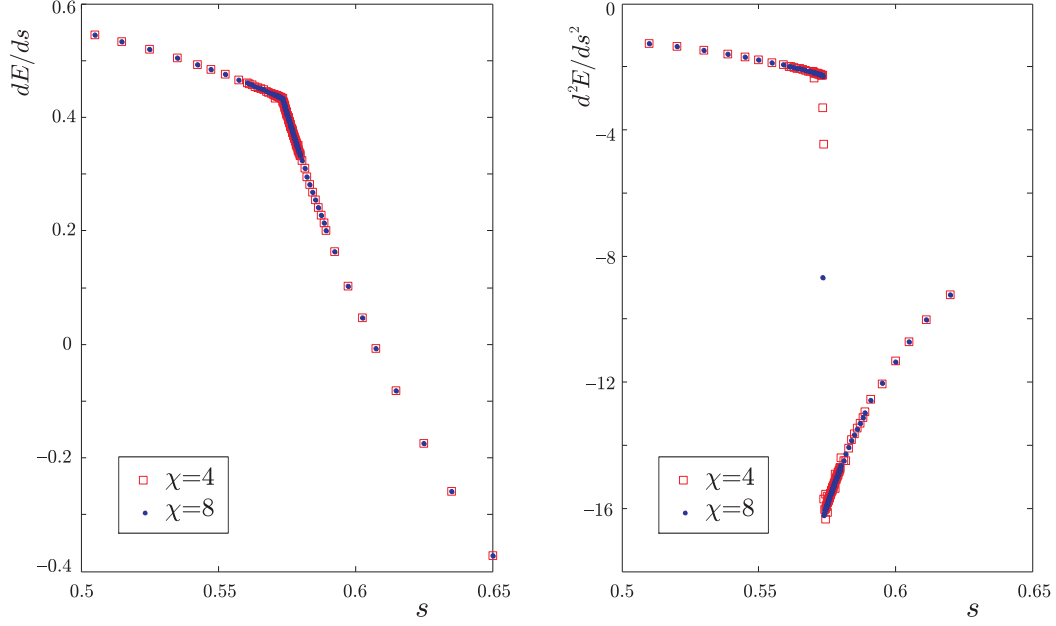


Figure 3-12: Transverse Ising model on an infinite tree. The first and second derivative with respect to s of the ground state energy obtained via MPS.

graph).

3.5.2 The Infinite Tree (Bethe Lattice).

The computational cost of the tree simulation is more expensive with growing χ than the line simulation, so I give my results only up to $\chi = 8$. I run the imaginary time evolution with 10000 iterations (each iteration followed by several normalization steps) for each point s , taking a lower χ result as the starting point for the procedure. I also add a small symmetry-breaking longitudinal field with magnitude $B_z = 10^{-8}$.

I see that the energy and its first derivative with respect to s are continuous. However, I now observe a finite discontinuity in the second derivative of the ground state energy (see FIG. 3-12), as opposed to the divergence on the infinite line. This happens near $s = s_T \approx 0.5733$.

Similarly to the one-dimensional case, the magnetization quickly grows for $s > s_T$, while it has (nearly) zero value for $s < s_T$ (see FIG. 3-13). In FIG.3-14, I plot the magnetization vs. $x - x_T$ on a log-log scale, where x is

$$x = \frac{s}{3(1-s)}, \quad (3.41)$$

with the value $x = x_T \approx 0.451$ at the phase transition (where $s = s_T \approx 0.5733$). I want to test whether the magnetization behaves like

$$M \propto (x - x_T)^\beta \quad (3.42)$$

for x close to x_T , which would appear as a line on the log-log plot. As χ grows, the data is

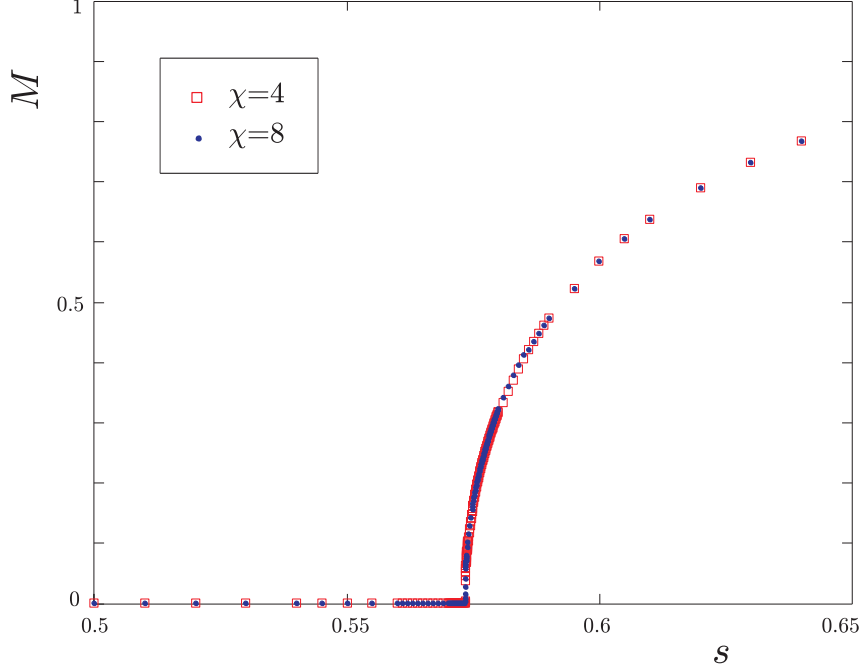


Figure 3-13: Transverse Ising model on an infinite tree. Magnetization vs. s , with $B_z = 10^{-8}$.

better represented by a straight line closer to the phase transition. If I fit the $\chi = 8$ data for $4 \times 10^{-4} \leq x - x_L \leq 10^{-3}$, I get $\beta = 0.41$. I add a line with this slope to the plot. Note that the mean-field value for the exponent β is 0.5, just as it is for the infinite line.

I observe that the correlation length now rises up only to a finite value (see FIG. 3-15). As I increase χ , the second eigenvalue of the B matrix, μ_2 , approaches a maximum value close to $\frac{1}{2}$. I conjecture that the limiting value of μ_2 is indeed $\frac{1}{2}$, which corresponds to a finite correlation length with value $(\ln 2)^{-1}$. Note that for the infinite line, the second eigenvalue of B approaches 1, and so the correlation length is seen to diverge at the phase transition.

3.6 The Not 00 Model

I now look at a model with a different interaction term. Starting with an antiferromagnetic interaction, I add a specific longitudinal field at each site. As in the previous section, I parametrize the Hamiltonian (3.2) with a single parameter s :

$$H_{\text{NOT00}} = s \sum_{\langle i,j \rangle} \underbrace{\frac{1}{4} (1 + \sigma_z^i + \sigma_z^j + \sigma_z^i \sigma_z^j)}_{H_{ij}} + \frac{b(1-s)}{2} \sum_i (1 - \sigma_x^i), \quad (3.43)$$

with $b = 2$ on the line and $b = 3$ on the tree. I choose the longitudinal field in such a way that the nearest-neighbor interaction term H_{ij} becomes a projector, expressed in the

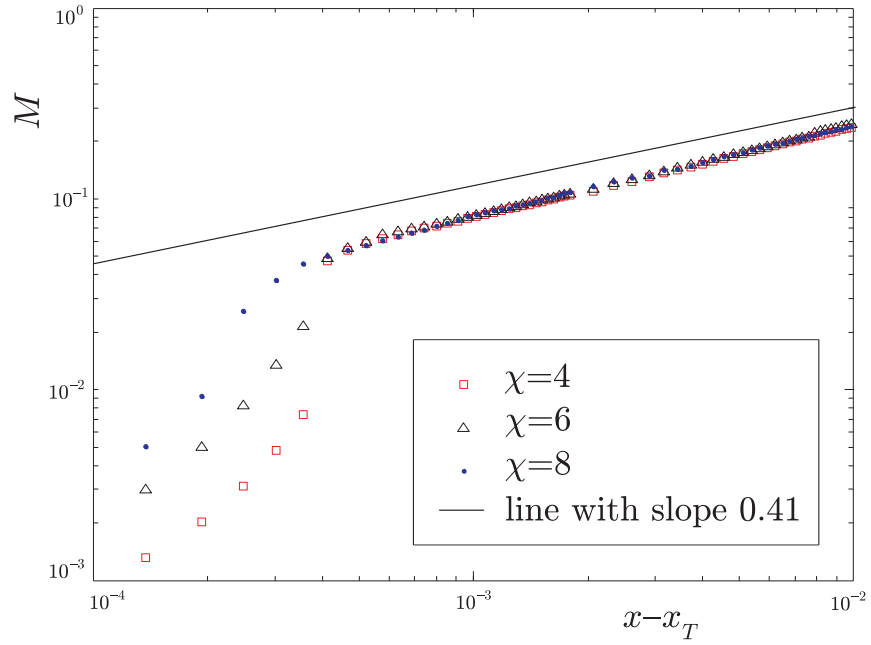


Figure 3-14: Transverse Ising model on an infinite tree. Log-log plot of magnetization vs. $x - x_T$, with $B_z = 10^{-8}$. I also plot a line with slope 0.41.

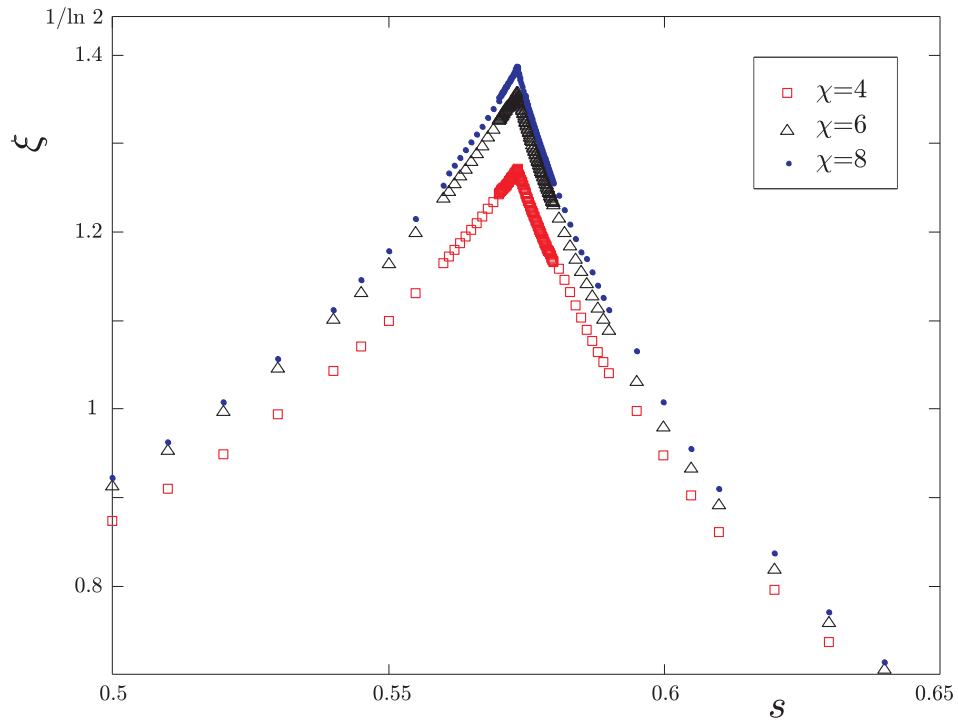


Figure 3-15: Transverse Ising model on an infinite tree. A linear plot of the correlation length vs. s .

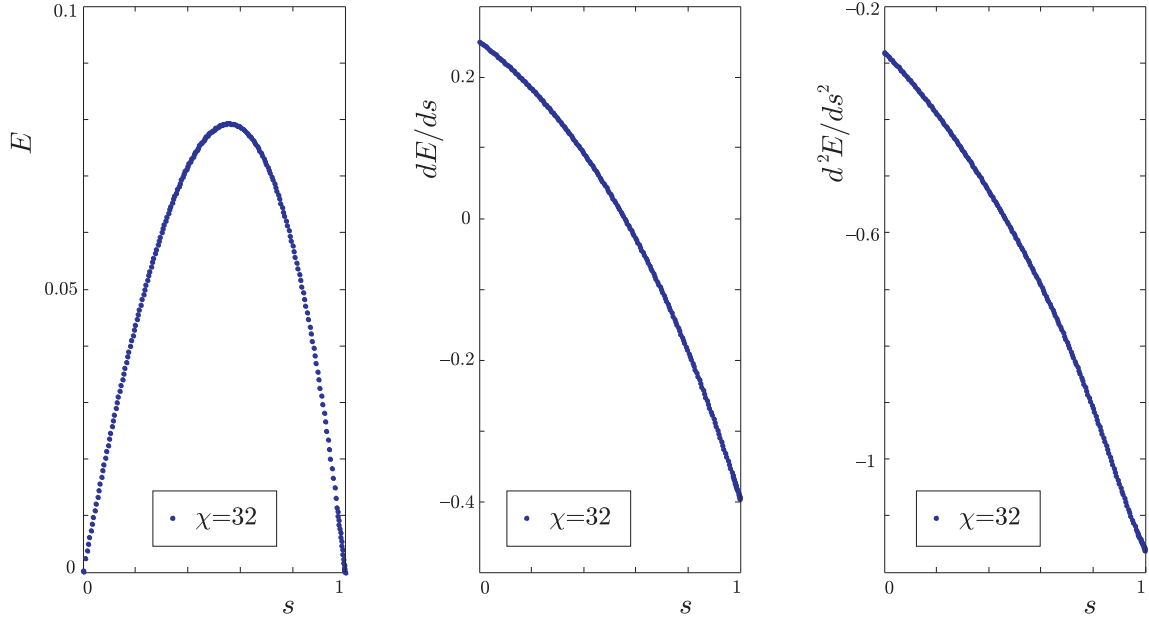


Figure 3-16: The NOT 00 model on an infinite line. The ground state energy and its first two derivatives with respect to s .

computational basis as

$$H_{ij} = |00\rangle \langle 00|_{ij}, \quad (3.44)$$

thus penalizing only the $|00\rangle$ configuration of neighboring spins. Accordingly, I call this model NOT 00. The ground state of the transverse Ising model (3.37) at $s = 1$ has degeneracy 2. For (3.43) on the infinite line or the Bethe lattice, the degeneracy of the ground state at $s = 1$ is infinite, as any state that does not have two neighboring spins in state $|0\rangle$ has zero energy.

3.6.1 Infinite Line

I use my numerics to investigate the properties of (3.43) on the infinite line as a function of s . My numerical results show continuous first and second derivatives of the energy with respect to s (see FIG.3-16). The magnetization $M = \langle \sigma_z \rangle$ decreases continuously and monotonically from 0 at $s = 0$ to a final value of -0.606 at $s = 1$ (see FIG.3-17). The second eigenvalue of the B matrix (3.33) rises continuously from 0 at $s = 0$, approaching 0.603 at $s = 1$ (see FIG.3-17). Because $\mu_2 < 1$, the correlation length ξ is finite for all values of s in this case. These results imply that there is no phase transition for this model as I vary s .

As a test of my results, I compute the magnetization at $s = 1$ exactly for this model on a finite chain (and ring) of up to $n = 17$ spins. I maximize the expectation value of $H_B = \sum_i \sigma_x^i$ within the subspace of all allowed states at $s = 1$ (with no two zeros on neighboring spins), thus minimizing the expectation value of the second term in (3.43) for s approaching 1. I compute the magnetization $M = \langle \sigma_z^i \rangle$ for the middle $i = \lfloor \frac{n}{2} \rfloor$ spin

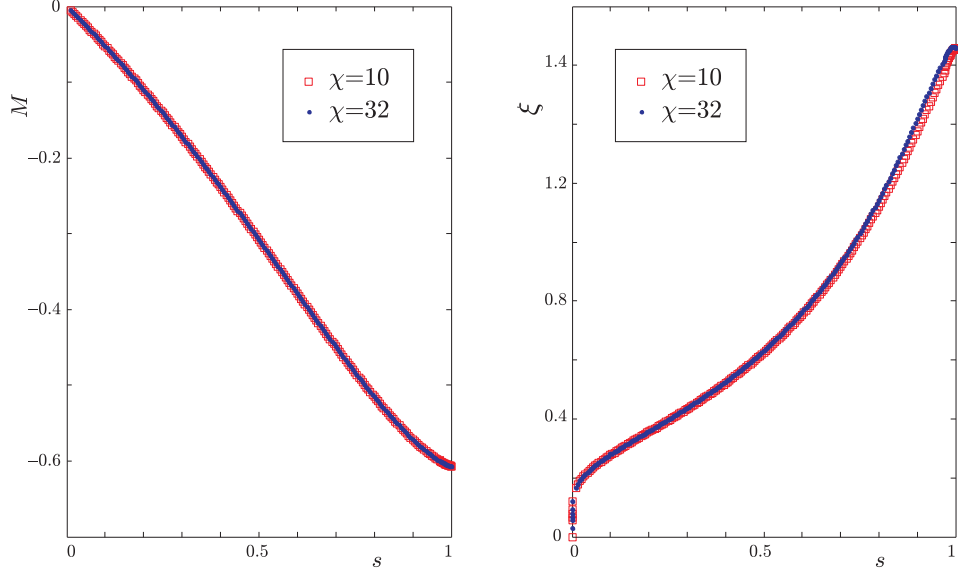


Figure 3-17: The NOT 00 model on an infinite line. Magnetization as a function of s and correlation length as a function of s .

for the ground state of the NOT 00 model exactly for a finite chain and ring of up to 17 spins at $s = 1$. As I increase n , the value of M converges to -0.603 (much faster for the ring, as the values of M for $n = 14, 17$ differ by less than 10^{-4}). Recall that I obtained $M = -0.606$ from my MPS numerics with $\chi = 16$ for the NOT 00 model on an infinite line. I also compare the values of the Schmidt coefficients across the central division of the finite chain ($n = 16$) to the elements of the λ vector obtained using my MPS numerics with $\chi = 32$. I observe very good agreement for the 11 largest values of λ_k , with the difference that my MPS values keep decreasing (exponentially), while the finite-chain values flatten out at around $\lambda_{k>14} \approx 10^{-9}$ (see FIG.3-18). The behavior of the components of λ from MPS doesn't change with increasing χ .

In the ground state of (3.43) at $s = 1$, the overlap with the $|00\rangle$ state of any two neighboring spins is exactly 0. If the Γ tensors are the same at every site, the component of the state $|\psi\rangle$ that has overlap with the state $|00\rangle$ on nearest neighbors can be expressed as

$$\sum_{a,b,c} (\lambda_a \Gamma_{ab}^0 \lambda_b \Gamma_{bc}^0 \lambda_c) |\phi_a\rangle |00\rangle |\phi_c\rangle. \quad (3.45)$$

Furthermore, when the Γ tensors are symmetric, the elements of the λ vectors must be allowed to take negative values to make this expression equal to zero. Note that until now, I used only positive λ vectors, knowing that they come from Schmidt decompositions, which give me the freedom to choose the components of λ to be positive and decreasing.

The negative signs in the λ vector can be absorbed into every other Γ tensor, resulting in a state with two different Γ (for the even and odd-numbered sites) and only positive λ 's. In fact, this is what I observe in my numerics, which assume positive λ , but allow two different Γ tensors (see 3.4.1). If I allow the elements of λ to take negative values, my numerically obtained Γ tensors are identical.

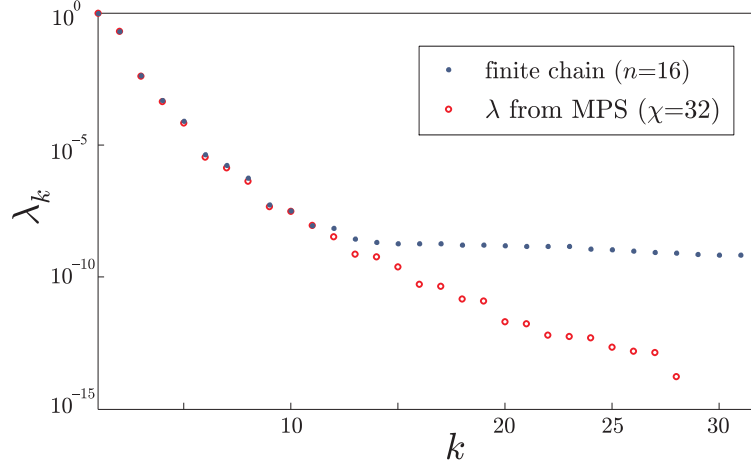


Figure 3-18: Ground state of the NOT 00 model on a line at $s = 1$. Comparison of the exact Schmidt coefficients for a division across the middle of a finite chain and of the MPS values ($\chi = 32$) for an infinite line.

3.6.2 Infinite Tree

Here, I numerically investigate the NOT 00 model (3.43) on the Bethe lattice. As on the line, the numerics show continuous first and second derivatives of the energy with respect to s (see FIG.3-19) and a continuous decrease in the magnetization from 0 at $s = 0$ to -0.671 at $s = 1$ (see FIG.3-20). The correlation length behaves similarly as on the line, increasing with s , but it reaches a maximum at $s = 0.96$ for $\chi = 8$. The maximum value of ξ is apparently lower than $1/\ln 2$, (see FIG.3-20), meaning that on the tree, the correlation function $\langle \sigma_z^{(i)} \sigma_z^{(j)} \rangle - \langle \sigma_z^{(i)} \rangle \langle \sigma_z^{(j)} \rangle$ falls off with distance faster than $2^{-|i-j|}$ for all s .

3.7 Stability and Correlation Lengths on the Bethe Lattice

I have found that on the Bethe lattice, for both my models, the second eigenvalue μ_2 of the matrix B (3.33), which determines the correlation length, apparently is never greater than $\frac{1}{2}$. In this section I argue that this is a model-independent, and calculation method independent, consequence of assuming that a translation-invariant ground state is the stable limit of a sequence of ground states of finite Cayley trees as the size of the tree grows. For a related problem, the stability of recursions for Valence Bond States on Cayley trees has been investigated by Fannes et.al. in [30].

The Hamiltonians (3.1) and (3.2) each consist of sums of terms $H_k^{(x)}$, $H_m^{(z)}$, as in (3.21), where each term $H_k^{(x)}$ depends on a single σ_x and each $H_m^{(z)}$ on a neighboring pair of σ_z . I calculate the quantum partition function

$$Z(\beta) = \text{tre}^{-\beta H} \quad (3.46)$$

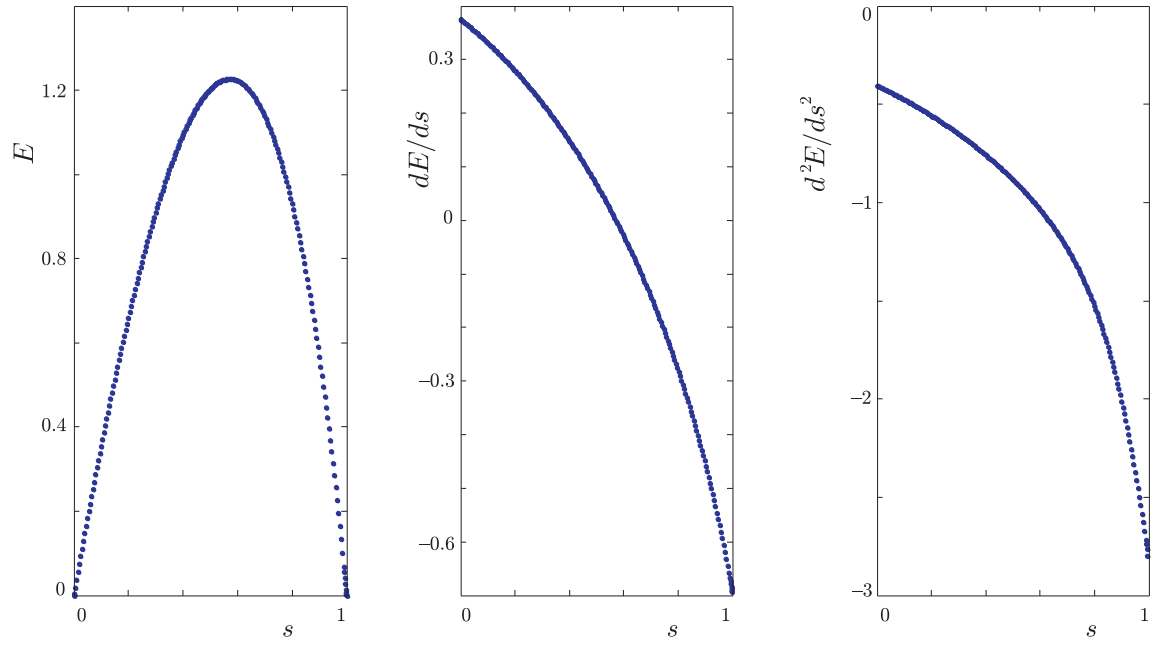


Figure 3-19: The NOT 00 model on an infinite tree. The ground state energy and its first two derivatives with respect to s .

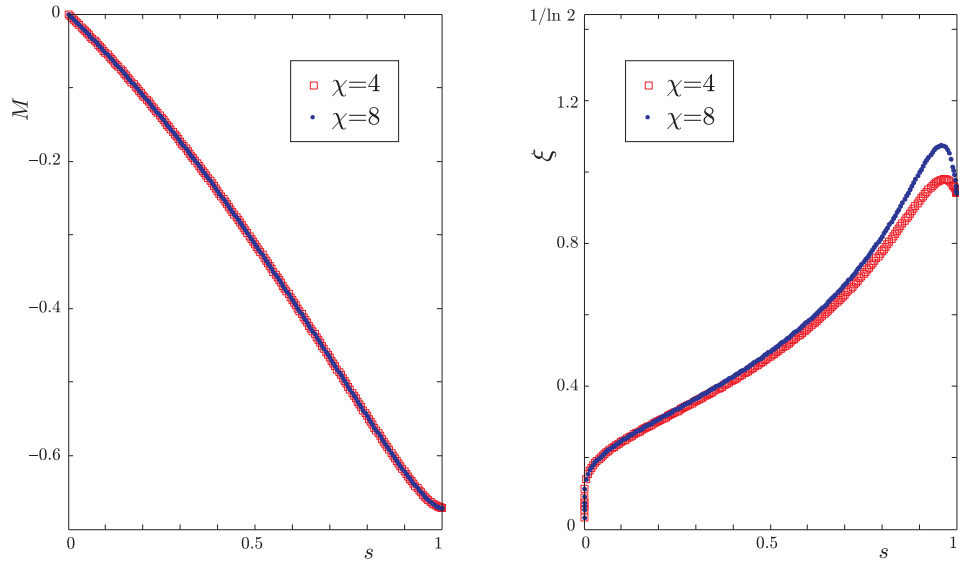


Figure 3-20: The NOT 00 model on an infinite tree. Magnetization as a function of s and correlation length as a function of s .

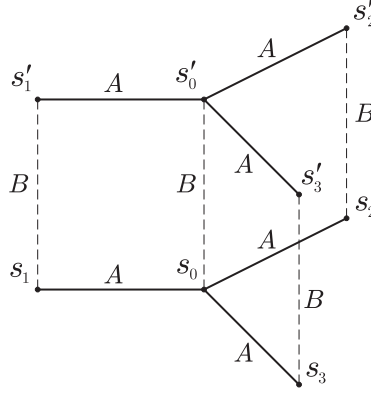


Figure 3-21: A system of classical spins on a lattice whose layers are Cayley trees. A and B denote the Boltzmann factors.

as the limit of

$$Z(N, \Delta t) = \text{tr} \left[\prod_k e^{-\Delta t H_k^{(x)}} \prod_m e^{-\Delta t H_m^{(z)}} \right]^N \quad (3.47)$$

as $\Delta t \rightarrow 0$, $N \rightarrow \infty$ with $N\Delta t = \beta$. To find the properties of the ground state, I take $\beta \rightarrow \infty$ so that I need $Z(N, \Delta t)$ as $\Delta t \rightarrow 0$, $N \rightarrow \infty$ with $N\Delta t \rightarrow \infty$ and $N(\Delta t)^3 \rightarrow 0$ (to make the error in using the Trotter-Suzuki formula go to zero). I interpret (3.47) as giving the *classical* partition function of a system of Ising spins ($s = \pm 1$) on a lattice consisting of N horizontal layers, each of which is a Cayley tree of radius M (i.e with a central node and concentric rings of $3, 3 \times 2, 3 \times 2^2, \dots, 3 \times 2^{M-1}$ nodes). I can write (3.47) as

$$Z(N, \Delta t) = \sum_{\{s\}} \prod A \prod B, \quad (3.48)$$

where the sum is over all configurations of $N \times (3 \times 2^M - 2)$ spins $s = \pm 1$ and the products are of a Boltzmann factor A for each horizontal link in the Cayley trees, and a Boltzmann factor B for each vertical link between corresponding nodes in neighboring layers (see FIG.3-21) (layer N is linked to layer 1 to give the trace). The factor A for the link between nodes i, j in the same horizontal layer is given by

$$A(s_i, s_j) = e^{-\Delta t H_{(ij)}^{(z)}(s_i, s_j)}. \quad (3.49)$$

The factor B for the link between nodes i, i' in the same vertical column is given by

$$B(s_i, s'_i) = \langle \sigma_z = s'_i | e^{-\Delta t H_{(i)}^{(x)}} | \sigma_z = s_i \rangle. \quad (3.50)$$

When all the terms $H_{(ij)}^{(z)}$ are of the same form, as are all the terms $H_{(i)}^{(x)}$, the form of the factors A and B does not depend on which particular links they belong to.

Each term in the sum, divided by Z , can be thought of as the probability of a configuration $\{s\}$. In what follows I will keep N and Δt fixed and consider the limit $M \rightarrow \infty$,

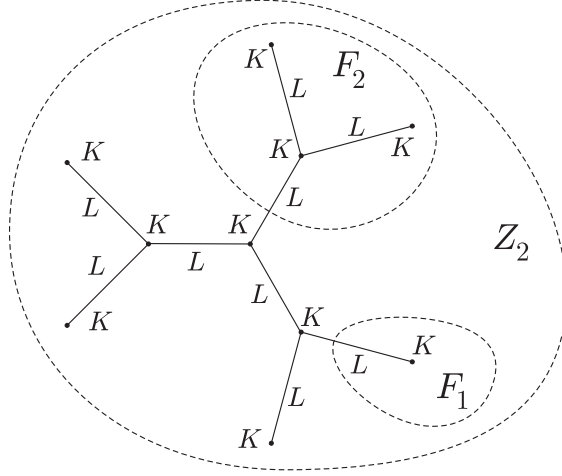


Figure 3-22: The $M = 2$ Cayley tree.

i.e. finite Cayley tree \rightarrow Bethe lattice. I will then suppose that my results, which are independent of the form of A and B (provided $A, B > 0$) will also hold after the $N \rightarrow \infty$ limit is taken, i.e. for the quantum ground state.

I will think of the lattice as a single tree, with each node being a vertical column of N spins. (A recent use of this technique to investigate the quantum spin glass on the Bethe lattice is in [59].) Denote by \vec{s} the vector of N values of s along a column. Let $K(\vec{s})$ be the product of the N factors $B(s, s')$ along a column and let $L(\vec{s}, \vec{s}')$ be the product of N factors $A(s, s')$ on the horizontal links between nearest neighbor columns. Let Z_M be the partition function for a tree of radius M . I can calculate Z_M by a recursion on M as follows:

$$Z_M = \sum_{\vec{s}} K(\vec{s}) [F_M(\vec{s})]^3, \quad (3.51)$$

$$F_M(\vec{s}) = \sum_{\vec{s}'} L(\vec{s}, \vec{s}') K(\vec{s}') [F_{M-1}(\vec{s}')]^2, \quad (3.52)$$

$$F_0(\vec{s}) = 1. \quad (3.53)$$

It is easy to see that this recursion gives the correct Z_M (the case $M = 2$ is shown in FIG.3-22). I also see that

$$P_M(\vec{s}) = \frac{1}{Z_M} K(\vec{s}) [F_M(\vec{s})]^3 \quad (3.54)$$

is the probability of the configuration \vec{s} along the central column. For the case $N = 1$, i.e. classical statistical mechanics on a tree, this is the well-known method to find an exact solution [14].

In order to have a well-defined translationally invariant limit as $M \rightarrow \infty$ I would like the recursion (3.52) for F_M to have an attractive fixed point F which F_M approaches as $M \rightarrow \infty$. ‘Attractive’ means that if I start the recursion with a different $F_0(\vec{s})$, sufficiently close to $F_0(\vec{s}) = 1$, the limiting value of $F_M(\vec{s})$ will be the same fixed point. This in turn implies that on a Cayley tree with large M , small changes in the Hamiltonian on the outer

edge will have small effects on the properties of the central region.

First however I need to fix the overall normalization of $F_M(\vec{s})$, since if $F_M(\vec{s})$ satisfies (3.52), so does $a^{2^M} F_M(\vec{s})$ which rules out an attractive fixed point.

Let

$$F_M(\vec{s}) = Z_M^{\frac{1}{3}} \hat{F}_M(\vec{s}), \quad (3.55)$$

so that

$$\sum_{\vec{s}} K(\vec{s}) [\hat{F}_M(\vec{s})]^3 = 1, \quad (3.56)$$

and

$$P_M(\vec{s}) = K(\vec{s}) [\hat{F}_M(\vec{s})]^3. \quad (3.57)$$

The recursion relation becomes

$$\hat{F}_M(\vec{s}) = \lambda_M \sum_{\vec{s}'} L(\vec{s}, \vec{s}') K(\vec{s}') [\hat{F}_{M-1}(\vec{s}')]^2, \quad (3.58)$$

with λ_M determined by the normalization condition (3.56). I can now suppose that

$$\hat{F}_M(\vec{s}) \rightarrow \hat{F}(\vec{s}) \quad \text{as} \quad M \rightarrow \infty, \quad (3.59)$$

with

$$\hat{F}(\vec{s}) = \lambda \sum_{\vec{s}'} L(\vec{s}, \vec{s}') K(\vec{s}') [\hat{F}(\vec{s}')]^2, \quad (3.60)$$

and

$$\sum_{\vec{s}} K(\vec{s}) [\hat{F}(\vec{s})]^3 = 1. \quad (3.61)$$

To determine whether \hat{F} is an *attractive* fixed point, let

$$\hat{F}_M(\vec{s}) = \hat{F}(\vec{s}) + f_M(\vec{s}), \quad (3.62)$$

$$\lambda_M = \lambda(1 + \epsilon_M), \quad (3.63)$$

with $f_M \rightarrow 0$ and $\epsilon_M \rightarrow 0$ as $M \rightarrow \infty$. To first order in f_M , ϵ_M , (3.58) and (3.56) become

$$f_M(\vec{s}) = \epsilon_M \hat{F}(\vec{s}) + 2 \sum_{\vec{s}'} T(\vec{s}, \vec{s}') f_{M-1}(\vec{s}'), \quad (3.64)$$

$$\sum_{\vec{s}} K(\vec{s}) [\hat{F}(\vec{s})]^2 f_M(\vec{s}) = 0, \quad (3.65)$$

where

$$T(\vec{s}, \vec{s}') = \lambda L(\vec{s}, \vec{s}') K(\vec{s}') \hat{F}(\vec{s}'). \quad (3.66)$$

From (3.60),

$$\sum_{\vec{s}'} T(\vec{s}, \vec{s}') \hat{F}(\vec{s}') = \hat{F}(\vec{s}), \quad (3.67)$$

and since $L(\vec{s}, \vec{s}') = L(\vec{s}', \vec{s})$,

$$\sum_{\vec{s}'} K(\vec{s}') [\hat{F}(\vec{s}')]^2 T(\vec{s}', \vec{s}) = K(\vec{s}) [\hat{F}(\vec{s})]^2, \quad (3.68)$$

i.e. the linear operator T has an eigenvalue one, with right eigenvector \hat{F} and left eigenvector $K\hat{F}^2$ (which from (3.61) have scalar product one). (3.64) now gives

$$\sum_{\vec{s}} K(\vec{s}) [\hat{F}(\vec{s})]^2 f_M(\vec{s}) = \epsilon_M + 2 \sum_{\vec{s}} K(\vec{s}) [\hat{F}(\vec{s})]^2 f_{M-1}(\vec{s}), \quad (3.69)$$

so from (3.65), $\epsilon_M = 0$. Let

$$T^\perp(\vec{s}, \vec{s}') = T(\vec{s}, \vec{s}') - \hat{F}(\vec{s}) K(\vec{s}') [\hat{F}(\vec{s}')]^2, \quad (3.70)$$

so that

$$\sum_{\vec{s}'} T^\perp(\vec{s}, \vec{s}') \hat{F}(\vec{s}') = 0, \quad (3.71)$$

and

$$\sum_{\vec{s}'} K(\vec{s}') [\hat{F}(\vec{s}')]^2 T^\perp(\vec{s}', \vec{s}) = 0. \quad (3.72)$$

(3.64) now becomes

$$f_M(\vec{s}) = 2 \sum_{\vec{s}'} T^\perp(\vec{s}, \vec{s}') f_{M-1}(\vec{s}'). \quad (3.73)$$

(3.73) shows that $\hat{F}(\vec{s})$ is an attractive fixed point if and only if

$$\|T^\perp\| < \frac{1}{2} \quad (3.74)$$

(for a tree with valence $p+1$ at each vertex, $\frac{1}{2}$ is replaced by $\frac{1}{p}$).

I can in fact prove that there does exist an $\hat{F}(\vec{s})$ satisfying (3.60) and (3.61), for which the corresponding T^\perp has a maximum eigenvalue less than $\frac{1}{2}$. Define a function Φ of $\hat{F}(\vec{s})$ by

$$\Phi[\hat{F}] = \sum_{\vec{s}, \vec{s}'} [\hat{F}(\vec{s})]^2 K(\vec{s}) L(\vec{s}, \vec{s}') K(\vec{s}') [\hat{F}(\vec{s}')]^2. \quad (3.75)$$

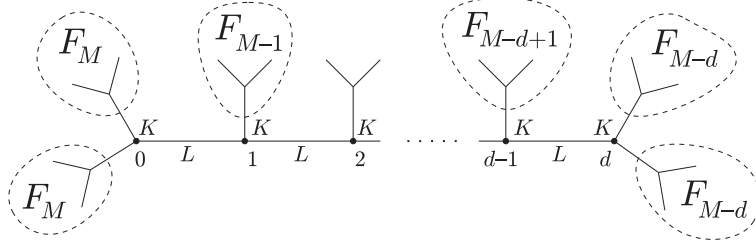


Figure 3-23: Computing the probability distribution of \vec{s}_0 and \vec{s}_d .

I look for a maximum of Φ with $\hat{F}(\vec{s})$ restricted to the region

$$\sum_{\vec{s}_i} K(\vec{s}) [\hat{F}(\vec{s})]^3 = 1, \quad (3.76)$$

$$\hat{F}(\vec{s}) \geq 0. \quad (3.77)$$

A maximum must exist, but it might be on the boundary of the region, i.e. it might have $\hat{F}(\vec{s}) = 0$ for some values of \vec{s} . Elementary calculations (omitted here) establish that stationary values of Φ on the boundary cannot be maxima. At stationary points in the interior of the region, i.e. with $\hat{F}(\vec{s}) > 0$ for all \vec{s} , (3.60) and (3.61) must be satisfied. If such a stationary point is a maximum, all the eigenvalues of $1 - 2T^\perp$ are ≥ 0 , i.e. all the eigenvalues of T^\perp are $\leq \frac{1}{2}$. This is weaker than the attractive fixed point condition, which also requires that no eigenvalue is less than $-\frac{1}{2}$, but does correspond to the observed property of μ_2 .

I now examine the joint probability distribution of \vec{s}_0 and \vec{s}_d , where 0 denotes the central column and d a column distance d from the center. One can see from FIG.3-23 that

$$P_M(\vec{s}_0, \vec{s}_d) = \frac{1}{Z_M} \sum_{\vec{s}_1, \dots, \vec{s}_{d-1}} [F_M(\vec{s}_0)]^2 K(\vec{s}_0) L(\vec{s}_0, \vec{s}_1) F_{M-1}(\vec{s}_1) \dots \dots L(\vec{s}_{d-1}, \vec{s}_d) K(\vec{s}_d) [F_{M-d}(\vec{s}_d)]^2. \quad (3.78)$$

As $M \rightarrow \infty$ (with d fixed) this becomes (using (3.66))

$$P(\vec{s}_0, \vec{s}_d) = C [\hat{F}(\vec{s}_0)]^2 K(\vec{s}_0) T^d(\vec{s}_0, \vec{s}_d) \hat{F}(\vec{s}_d), \quad (3.79)$$

where the normalization C is determined by

$$\sum_{\vec{s}_0, \vec{s}_d} P(\vec{s}_0, \vec{s}_d) = 1. \quad (3.80)$$

Using (3.67) I find

$$P(\vec{s}_0) = \sum_{\vec{s}_d} P(\vec{s}_0, \vec{s}_d) = C K(\vec{s}_0) [\hat{F}(\vec{s}_0)]^3, \quad (3.81)$$

so from (3.61), $C = 1$ (and $P(\vec{s}_0)$ agrees with the limit of (3.57)). Expressing (3.79) in

terms of T^\perp using (3.70), (3.71) and (3.72),

$$P(\vec{s}_0, \vec{s}_d) - P(\vec{s}_0)P(\vec{s}_d) = K(\vec{s}_0)[\hat{F}(\vec{s}_0)]^2(T^\perp)^d(\vec{s}_0, \vec{s}_d)\hat{F}(\vec{s}_d). \quad (3.82)$$

Thus the correlation between \vec{s}_0 and \vec{s}_d falls off as μ^d , where μ is the eigenvalue of T^\perp with maximum modulus, and so from (3.74), faster than $1/2^d$.

If this conclusion is correct (and clearly the argument is less than rigorous), it establishes more than my experimental observation that $\mu_2 < \frac{1}{2}$. The quantum limit of $P(\vec{s}_0, \vec{s}_d)$ encodes not only the static correlation $\langle \psi_0 | s_0 s_d | \psi_0 \rangle$ in the ground state, but also the imaginary time dependent correlation $\langle \psi_0 | e^{Ht} s_0 e^{-Ht} s_d | \psi_0 \rangle$ which in turn determines the linear response as measured by s_d to a time-dependent perturbation proportional to s_0 . If this indeed falls off faster than $1/2^d$, then there is some hope that the Bethe lattice can be used as a starting point for investigation of fixed valence random lattices.

Chapter 4

Quantum Satisfiability

The topic of this chapter is *Quantum Satisfiability*, the quantum analog of classical *Satisfiability*. I discussed the locally constrained classical problem in Section 1.1.3, where I also summarized the well known complexity of its many variants (see Table 1.1 in Section 1.1.2). On the other hand, there are still several open questions about Quantum Satisfiability, and I choose to investigate some of them here.

First, in Section 4.1.1 I introduce Quantum Satisfiability and summarize what is presently known about its complexity. Then I review the proof techniques I use throughout this chapter, focusing on several clock constructions in Section 4.1.3. Section 4.2 is based on the paper [70]

A New Construction for a QMA Complete 3-local Hamiltonian

Daniel Nagaj, Shay Mozes

We present a new way of encoding a quantum computation into a 3-local Hamiltonian. Our construction is novel in that it does not include any terms that induce legal-illegal clock transitions. Therefore, the weights of the terms in the Hamiltonian do not scale with the size of the problem as in previous constructions. This improves the construction by Kempe and Regev [54], who were the first to prove that 3-local Hamiltonian is complete for the complexity class QMA, the quantum analogue of NP.

Quantum k-SAT, a restricted version of the local Hamiltonian problem using only projector terms, was introduced by Bravyi [20] as an analogue of the classical k-SAT problem. Bravyi proved that quantum 4-SAT is complete for the class QMA with one-sided error (QMA_1) and that quantum 2-SAT is in P. We give an encoding of a quantum circuit into a quantum 4-SAT Hamiltonian using only 3-local terms. As an intermediate step to this 3-local construction, we show that quantum 3-SAT for particles with dimensions $3 \times 2 \times 2$ (a qutrit and two qubits) is QMA_1 complete. The complexity of quantum 3-SAT with qubits remains an open question.

One of the implications of this construction is that Adiabatic Quantum Computation can be used for simulating quantum circuits much more effectively than Aharonov et al. showed in [7] (see Section 2.2).

Next, in Section 4.3, I review the result of Aharonov et al. [5] who proved that Quantum 2-SAT on a line with 12-dimensional particles is QMA_1 complete. I decrease the required dimensionality of particles in their construction to $d = 11$ in Section 4.3.2.

In Section 4.4, I present a very recent result by Eldar and Regev, who proved that Quantum 2-SAT for a cinquit-qutrit pair (particles with $d = 5$ and $d = 3$) is QMA₁ complete, using a novel *triangle clock* construction. In Section 4.4.2, I use their result to obtain another QMA-complete 3-local Hamiltonian construction built from constant norm, 3-local terms of a restricted type.

Finally, I take yet another step towards figuring out the complexity of Quantum 3-SAT in Section 4.5. Using my *train switch* clock construction, I show that a simple Quantum 3-SAT Hamiltonian is universal for quantum computation in the Adiabatic Quantum Computing model of Chapter 2 and in the more general Hamiltonian Computer model of Section 2.2.1. Moreover, the required running time for this model scales only slightly worse than linearly with the length of the computation.

4.1 Introduction

4.1.1 Quantum k-SAT

The Quantum k -SAT (Q- k -SAT) promise problem was introduced by Bravyi [20] as an analogue of classical k -SAT. In classical *Satisfiability*, one needs to determine whether there exists a bit string satisfying all the boolean clauses of a given problem instance. In Q- k -SAT, the problem is to determine whether a Hamiltonian acting on n qubits has a zero eigenvalue, or whether all its eigenvalues are higher than $\epsilon \geq n^{-\alpha}$ for some constant α . Moreover, the Hamiltonian

$$H_{\text{Q-}k\text{-SAT}} = \sum P_i, \quad (4.1)$$

is composed of k -local projector terms

$$P_i = P_i^2 = \mathbb{I}^{\otimes(n-k)} \otimes |\psi_i\rangle \langle \psi_i|_{\{q_1^i \dots q_k^i\}}, \quad (4.2)$$

where each P_i acts nontrivially on k qubits $\{q_1^i \dots q_k^i\}$. For a ‘yes’ answer to a problem instance, the ground state of $H_{\text{Q-}k\text{-SAT}}$ must be *exactly zero*, i.e. a state which is annihilated by *all* the projectors P_i must exist. This shows the analogy to classical SAT, where all the boolean clauses have to be satisfied. If the states $|\psi_i\rangle$ are computational basis states and $\epsilon = 1$, the problem reduces to classical k -SAT. The special case of commuting projectors P_i has been analyzed in [21].

Local Hamiltonian, the problem I introduced in Section 1.3.2, is more general than Q- k -SAT. It is the quantum analogue of classical MAX- k -SAT, where one is interested in the properties of the ground state of the *total* Hamiltonian (the *sum* of the terms). Viewing Quantum k -SAT and k -local Hamiltonian as quantum analogues of classical k -SAT and MAX- k -SAT, it is interesting to investigate and compare their complexities. Known complexity results about the classical problems are summarized in Section 1.1.2, Table 1.1. In Table 4.1, I present the currently known results about Quantum Satisfiability and Local Hamiltonian problems. In this Chapter, I prove those marked with *.

In [20], Bravyi gave a classical polynomial algorithm for Quantum 2-SAT showing that Q-2-SAT belongs to P. He then proved that Quantum k -SAT is QMA₁ complete for $k \geq 4$. The class QMA₁ is a special case of QMA, where only one sided error is allowed (see Definition 8 in Section 1.2.2). When the answer to the problem instance is ‘yes’, the verifier circuit U can output ‘yes’ on some state with certainty. Thus, the ground state energy

Quantum	qubits	qudits
Q- k -SAT	$k = 2$: in P	Q-(5, 3)-SAT : QMA ₁ -complete Q-(11, 11)-SAT in 1D : QMA ₁ - complete*
	$k = 3$: contains NP, universal for BQP*	Q-(3, 2, 2)-SAT : QMA ₁ -complete*
	$k \geq 4$: QMA ₁ -complete	
k -local Hamiltonian	$k \geq 2$: QMA-complete even on a 2D grid $k = 3$ with constant norm terms : QMA-complete*	

Table 4.1: Known complexity of Quantum Satisfiability and Local Hamiltonian problems. My own results, presented later in this Chapter, are marked with *.

of a Quantum k -SAT Hamiltonian must be exactly zero in the ‘yes’ case. When I release the certainty requirement on the verifier circuit U , it becomes a verifier circuit for a QMA problem instead. Therefore, when I prove that a certain Q-SAT problem is QMA₁ complete, the QMA-completeness of the mother Local-Hamiltonian problem immediately follows as a side result. I present a result of this type in Section 4.2, where I give a new 3-local Hamiltonian construction with constant norm terms.

Bravyi’s original definition also required all of the terms in the Hamiltonian to be projectors. However, using k -local positive semidefinite operator terms H_i^+ with zero ground state and constant norm instead of projectors P_i in (4.1) is an equivalent problem. Quantum k -SAT with positive semidefinite operators contains Q- k -SAT with projectors. On the other hand, if one is able to solve Q- k -SAT with projectors, one can solve Q- k -SAT with positive semidefinite operators as well. For each positive semidefinite operator H_i^+ , define a projector P_i^+ with the same ground state subspace. If $H_P = \sum P_i^+$ has a zero ground state, so does $H_+ = \sum H_i^+$. If the ground state energy of $H_P = \sum P_i^+$ is greater than ϵ , the ground state energy of H_+ is greater than $c\epsilon$, where c is a constant. Therefore, instead of only projector terms, I can equivalently use positive semidefinite operators in Q- k -SAT Hamiltonians.

4.1.2 Constructing the Hamiltonians

The two types of results about Quantum k -SAT Hamiltonians that I present here are proofs of QMA (or QMA₁) completeness and proofs of universality of using a Hamiltonian for simulating quantum circuits. In both cases, the quest is to prove these results for ever simpler Hamiltonians. Better locality (lower k , the number of particles involved in each term), simpler geometry of interactions (2D grid, 1D), lower dimensionality of particles involved, and using only terms of a restricted type are what I am aiming at.

To prove that a Hamiltonian problem is QMA-complete, one first needs to show that it belongs to QMA, and then to show that it also contains QMA. I have presented Kitaev’s proof of QMA-completeness for the 5-Local Hamiltonian problem in Section 1.3.2, and because a lot of the results I present in the next two Chapters build on it, I now quickly summarize it. Before considering the special case of 5-local Hamiltonian, Kitaev proved that k -local Hamiltonian is in QMA for any constant k . Therefore, the question whether

a Hamiltonian problem belongs to QMA is usually answered quickly by a reduction to k -Local Hamiltonian. For the other direction, one needs to reduce every instance of a problem in QMA into an instance of the Hamiltonian problem in question. Every instance of a problem in QMA is a ‘yes/no’ question. Moreover, every problem instance has a verifier circuit U which verifies claims (supposed proofs) about the answer. This circuit must have the completeness and soundness properties. Completeness means that if the answer to the problem instance is ‘yes’, there exists a state that the circuit U will accept with high probability. The soundness property means that if the answer to the problem is ‘no’, there is no state that fools the verifier circuit U with high probability. I then need to construct a Hamiltonian whose ground state would have low energy (or exactly zero energy for Quantum k -SAT) if the answer is ‘yes’, and whose ground state would have considerably higher energy if the answer is ‘no’.

All the results in this Chapter use the same general form of this Hamiltonian, but differ in the implementation of some of the terms. Let me first look at what they have in common and focus on the differences in Section 4.1.3. The verifier quantum circuit U with L gates acts on n qubits. Consider now a quantum system whose Hilbert space \mathcal{H} is larger than $\mathcal{H}_{work} = \mathbb{C}^{\otimes n}$. A way to do this is to append a clock register to the system as

$$\mathcal{H} = \mathcal{H}_{work} \otimes \mathcal{H}_{clock}. \quad (4.3)$$

Then, encode the progression of the circuit U

$$|\psi_t\rangle = U_t U_{t-1} \dots U_2 U_1 |\psi_0\rangle, \quad t = 0, \dots, L \quad (4.4)$$

into states

$$|\Psi_t\rangle = |\psi_t\rangle \otimes |t\rangle_c \quad (4.5)$$

of the larger system. How exactly the states $|t\rangle_c$ and transitions between them are implemented constitutes a ‘clock construction’ and I explain the possibilities in detail in Section 4.1.3. Following Kitaev, the Hamiltonian I then construct for the larger system is

$$H = H_{prop} + H_{out} + H_{input} + H_{clock} + H_{clockinit}. \quad (4.6)$$

It is built in such a way that the history state for the quantum circuit U corresponding to some initial state $|\psi_0\rangle$

$$|\Psi_{history}\rangle = \frac{1}{\sqrt{L+1}} \sum_{t=0}^L |\Psi_t\rangle = \frac{1}{\sqrt{L+1}} \sum_{t=0}^L |\psi_t\rangle \otimes |t\rangle_c \quad (4.7)$$

is close to the ground state of H . In fact, for a ‘yes’ instance of Quantum k -SAT, it *is* the ground state of H .

The first term, H_{prop} is a Hamiltonian whose ground state has the form $|\Psi_{history}\rangle$:

$$H_{prop} = \sum_{t=1}^L H_{prop}^t \quad (4.8)$$

$$H_{prop}^t = \frac{1}{2} \left(\mathbb{I} \otimes (P_t + P_{t-1}) - U_t \otimes X_{t,t-1} - (U_t \otimes X_{t,t-1})^\dagger \right), \quad (4.9)$$

where $P_t = |t\rangle\langle t|_c$ is a projector onto the state $|t\rangle_c$ of the clock register, the operator $X_{t,t-1} = |t\rangle\langle t-1|_c$ increases the time register from $t-1$ to t and the unitary gate U_t acts on the corresponding work qubits. Observe that it is built from terms which check whether the amplitude of each $|\Psi_{t-1}\rangle$ and $|\Psi_t\rangle$ is the same. Therefore, the history state (4.7) is its eigenvector with eigenvalue 0.

The last of the n work qubits, q_n , is the output qubit for the circuit U . The term

$$H_{out} = |0\rangle\langle 0|_{q_n} \otimes P_L \quad (4.10)$$

takes care of whether the circuit U accepts some state or not. It gives an energy penalty to any history state built from a state $|\psi_{no}\rangle$ that is not accepted by U , i.e. whose output qubit after L gates is not in the state $|1\rangle$. The term

$$H_{input} = \sum_{k \in \text{ancilla}} |1\rangle\langle 1|_{q_k} \otimes P_0 \quad (4.11)$$

takes care of proper initialization of ancilla qubits, in the case they are required for U . The final two terms H_{clock} and $H_{clockinit}$ add an energy penalty to all states of the clock register \mathcal{H}_{clock} that are not in the subspace spanned by the states $|t\rangle_c$. The implementation of these two terms and how the states $|t\rangle_c$ are actually encoded is the topic of the next Section.

For the other type of result, universality of a Hamiltonian H for simulating a quantum circuit, one needs to show that time evolution (for not too long a time) from a simple initial state with the Hamiltonian H can produce the output of the circuit U with high probability. This type of proof again requires the encoding of the progression of a quantum circuit U into states $|\Psi_t\rangle$ of a larger system. However, proving this type of result is different in that I need to analyze the dynamics of the system, not just the ground state of H . However, what simplifies the problem is that I can choose the initial state myself. I present one such result in Section 4.5. Later, the whole Chapter 5 is dedicated to showing universality for Hamiltonians which are both time and translation-invariant.

4.1.3 Constructing Clocks

In this Section, I review several methods of implementing the clock register in (4.3). The goal of a *clock construction* is to encode the progression (4.4) of a quantum circuit U with L gates on n work qubits into a sequence of orthogonal states $|\Psi_t\rangle$ of a larger system \mathcal{H} . It is desirable for a clock construction to have:

- Easily implementable unique transitions

One needs to construct a Hamiltonian inducing transitions between the states $|\Psi_t\rangle$. It is important to ensure that each state $|\Psi_t\rangle$ can transition directly only to the states $|\Psi_{t+1}\rangle$ or $|\Psi_{t-1}\rangle$, i.e.

$$H |\Psi_t\rangle = \alpha_{t-1} |\Psi_{t-1}\rangle + \beta_t |\Psi_t\rangle + \gamma_{t+1} |\Psi_{t+1}\rangle, \quad (4.12)$$

and not into any other state. Undesirable transitions can be suppressed by adding large penalty terms as in the work of Kempe and Regev [54], but such constructions can not be used for proving results about Quantum k -SAT. In Sections 4.2-4.5 I give ways to construct the clock register with unique transition rules, allowing me to prove new results about Q- k -SAT and strengthen old results about the Local Hamiltonian problem.

- Locally checkable encoding, initialization and final detection

The full Hilbert space \mathcal{H} can be much larger than the subspace \mathcal{H}_{legal} spanned by the states $|\Psi_t\rangle$. I must be able to detect whether a state belongs to the subspace \mathcal{H}_{legal} using only local terms. In practice, this is done by adding an energy penalty to any state outside \mathcal{H}_{legal} . Also, I need to be able to make a local projection measurement onto $|\Psi_L\rangle$ and $|\Psi_L\rangle^\perp$.

- Local transition rules involving only a few (clock) particles,
- Low dimensionality of the (clock) particles,

Simple local interactions in the system \mathcal{H} and low dimensionality of the particles involved are essential for my goal of obtaining stronger results about ever simpler local Hamiltonians.

- Simple geometry of interactions,

Finally, for practical implementation, the k -particle interactions in H need to involve only spatially close particles. The 1D construction of [5] presented in Section 4.3 is one example of this.

Let me now present a few of clock register encodings and discuss their advantages and disadvantages.

Adding a clock register

The first option is to directly add a clock register to the system as in Feynman's construction in Section 1.3.1

$$\mathcal{H} = \mathcal{H}_{work} \otimes \mathcal{H}_{clock}. \quad (4.13)$$

A set of orthogonal states $|t\rangle_c$ of the clock register is then used to label the progress of the computation, while the contents of the work register holds the state of the work qubits after t gates have been applied to an initial state $|\psi_0\rangle$

$$|\Psi_t\rangle = \underbrace{(U_t U_{t-1} \dots U_1 |\psi_0\rangle)}_{|\psi_t\rangle} \otimes |t\rangle_c. \quad (4.14)$$

The space spanned by the clock states $|t\rangle_c$ then defines the *legal clock subspace* \mathcal{H}_{legal} . An essential ingredient of the clock constructions I present here is a local way of checking whether a state of the clock register belongs to the legal clock subspace. It is usually implemented by projector terms adding an energy penalty to states outside of \mathcal{H}_{legal} .

One could think of implementing the clock register by simply labeling the basis states of \mathcal{H}_{clock} by $|t\rangle_c$. The minimum required number of qubits in \mathcal{H}_{clock} would then be $\log L$. However, such a construction is not local, as the transition rules between different clock states necessarily involve $(\log L)$ qubits.

The simplest local implementation of the clock register is the *pulse clock* used by Feynman [38]. A state corresponding to time t is the state of $L + 1$ qubits c_0, \dots, c_L , with a single clock qubit c_t in the state $|1\rangle$

$$|t\rangle_c = |00 \dots 00100 \dots 00\rangle. \quad (4.15)$$

This single up-spin denotes the *active site* in the clock register. To transition between clock states $|t\rangle_c$ and $|t+1\rangle_c$ (and vice versa), I need to look only at 2 clock qubits and apply the operator

$$X_{t+1,t} + X_{t+1,t}^\dagger = |01\rangle\langle 10|_{c_t, c_{t+1}} + |10\rangle\langle 01|_{c_t, c_{t+1}}. \quad (4.16)$$

This low locality of transitions is a big advantage of the pulse clock. Proper encoding of the clock register is assured by operators $|11\rangle\langle 11|$ which give an energy penalty to incorrect clock states with two active spots (two spin up particles) in the system. These are also 2-local. On the other hand, the big disadvantage of the pulse clock is that it is impossible to check whether the clock is not in the state $|D\rangle = |0\dots 0\rangle_c$, without a single spin up, using only local, projector terms. This ‘dead’ state is annihilated by the transition operators (4.16), so it does not ‘move’ anywhere. For this reason, the pulse clock is suited only for constructing Hamiltonian computers (see Section 2.2.1), where one can pick the initial state of the clock and thus rule out the state $|D\rangle$. It is unusable for proving QMA-completeness results, unless one rules out the state $|D\rangle$ by adding terms that energetically favor a single spin up in the system over the state $|D\rangle$.

I already presented Kitaev’s unary *domain wall* clock in Section 1.3.2

$$|t\rangle_c = |11\dots 1100\dots 00\rangle. \quad (4.17)$$

The active spot in the clock register is now determined by the position of the domain wall 10. The unique transition rules for this clock are now 3-local, $|\dots 100\dots\rangle \leftrightarrow |\dots 110\dots\rangle$, but the advantage one gains are simple check operators for the clock. Illegal clock states are detected by the projector $|01\rangle\langle 01|$ on consecutive clock qubits, and the dead state $|D\rangle = |0\dots 0\rangle_c$ can be easily ruled out using the projector $|0\rangle\langle 0|_{c_0}$ on the first clock qubit. Also, one can use a 1-local transition rule $|0\rangle \leftrightarrow |1\rangle$ between clock states as in [54] or [7], with the caveat that the illegal transitions such as $|11100000\rangle \rightarrow |11100100\rangle$ can occur and need to be taken care of by using large penalty terms. I work around this problem while keeping the locality low in the following constructions.

It is desirable to decrease the locality of transition rules as much as possible. In Bravyi [20] introduced a clock made out of particles with $d = 4$. The upside of his construction is that some of the transition rules involve only one clock particle (albeit 4-dimensional). The interaction required to increment the clock and to apply a two-qubit gate at the same time thus involves particles with dimension $4 \times 2 \times 2$, and can be thought of as a 4-qubit interaction. I now present a combined domain wall and pulse clock made out of particles with $d = 2$ and $d = 3$, building on the idea of [20]. Consider first a pulse clock made out of particles with dimension $d = 3$, with the progression of states

$$|0\rangle_c = |\triangleright \square \square \square \square\rangle, \quad (4.18)$$

$$|1\rangle_c = |\triangleleft \square \square \square \square\rangle, \quad (4.19)$$

$$|2\rangle_c = |\square \triangleright \square \square \square\rangle, \quad (4.20)$$

$$|3\rangle_c = |\square \triangleleft \square \square \square\rangle, \quad (4.21)$$

$$|4\rangle_c = |\square \square \triangleright \square \square\rangle, \quad (4.22)$$

$$|5\rangle_c = |\square \square \triangleleft \square \square\rangle. \quad (4.23)$$

Some of the transitions are 1-local ($\triangleright \leftrightarrow \triangleleft$) while others are 2-local ($\triangleleft \square \leftrightarrow \square \triangleright$). Therefore, one can engineer a 3-local interaction for a qutrit and two qubits which increments

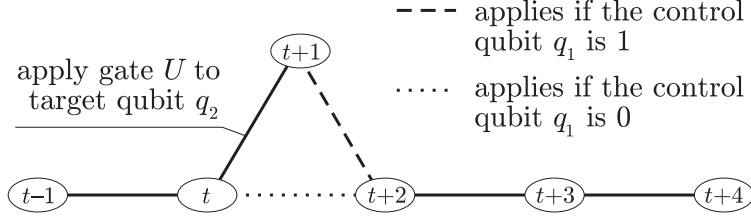


Figure 4-1: The ‘triangle’ clock transition rules allowing to apply a controlled gate to two work qubits using interactions that involve one work qubit at a time.

the clock and concurrently applies a unitary gate to two work qubits. The initialization problem of this pulse clock is then solved by combining it with a domain wall clock $|111000\rangle$ as follows:

$$|0\rangle_c = |10 \square 00 \square 00 \square 00\rangle, \quad (4.24)$$

$$|1\rangle_c = |11 \triangleright 00 \square 00 \square 00\rangle, \quad (4.25)$$

$$|2\rangle_c = |11 \triangleleft 00 \square 00 \square 00\rangle, \quad (4.26)$$

$$|3\rangle_c = |11 \square 10 \square 00 \square 00\rangle, \quad (4.27)$$

$$|4\rangle_c = |11 \square 11 \triangleleft 00 \square 00\rangle, \quad (4.28)$$

$$|5\rangle_c = |11 \square 11 \triangleright 00 \square 00\rangle, \quad (4.29)$$

$$|6\rangle_c = |11 \square 11 \square 10 \square 00\rangle. \quad (4.30)$$

The transition rules are 3-local for the transition of the active spot from the domain wall to the pulse clock $10 \square \leftrightarrow 11 \triangleleft$, while they stay 1-local for the $\triangleleft \leftrightarrow \triangleright$ transition. Also, note that 3-local operators suffice to detect whether the clock register state has a single active spot (either 10 in neighboring domain wall qubits or $1 \triangleleft 0$ or $1 \triangleright 0$ when there is a pulse clock particle at the domain wall). I implement this clock using qubits rather than qutrits in Section 4.2, which results in a new QMA-complete 3-local Hamiltonian construction without the requirement for large norm terms penalizing illegal clock states.

Very recently, Eldar and Regev [28] had a revolutionary idea and abandoned the linear progression of states $|\Psi_t\rangle$ to create a ‘triangle’ clock as in Figure 4-1. In this construction, the progression of states $|\Psi_t\rangle$ can take two different paths, depending on the state of one of the work qubits. Eldar and Regev use it to prove interesting results about Quantum 2-SAT with qudits, shattering the misconception that to apply a two-qubit gate, one has to interact with both qubits in the same step. I review their work in Section 4.4.1 and present my own results using this construction in Section 4.4.2. I then take this idea further, creating a more complex, but more symmetric, *train switch* clock in Section 4.5. There I use it to prove the final result of this Chapter: one can simulate BQP in the Hamiltonian Computer model of Section 2.2.1 using a Quantum 3-SAT Hamiltonian.

Geometric clocks

Besides directly adding a separate clock register to the system and leaving the work qubits at a particular location in the system, one has the option of having the n work qubits move around by using higher dimensional particles. Consider a simple 2D construction using

qutrits, where the state space of each particle is

$$\mathcal{H}_p = \mathcal{H}_{\odot} \oplus \mathcal{H}_{\boxtimes}, \quad (4.31)$$

with \mathcal{H}_{\odot} a 2-dimensional subspace (can hold the state of a qubit) and \mathcal{H}_{\boxtimes} a one-dimensional subspace. When a particle is in the state $|\boxtimes\rangle$, it indicates that no qubit is present at the site. The location of the qubits along the strip then determines the states $|\Psi_t\rangle$:

$$\begin{array}{cccccccccccc} \boxtimes & \boxtimes & \boxtimes & \odot & \boxtimes & \boxtimes & \boxtimes & \boxtimes & \boxtimes & \boxtimes & \boxtimes & \boxtimes \\ \boxtimes & \boxtimes & \boxtimes & \odot & \boxtimes & \boxtimes & \boxtimes & \boxtimes & \boxtimes & \boxtimes & \boxtimes & \boxtimes \\ \boxtimes & \boxtimes & \boxtimes & \odot & \boxtimes & \boxtimes & \boxtimes & \boxtimes & \boxtimes & \boxtimes & \boxtimes & \boxtimes \\ \boxtimes & \boxtimes & \boxtimes & \odot & \boxtimes & \boxtimes & \boxtimes & \boxtimes & \boxtimes & \boxtimes & \boxtimes & \boxtimes \end{array} \quad (4.32)$$

Observe that the state $|\Psi_t\rangle$ is necessarily orthogonal to any state $|\Psi_{t'}\rangle$ for which the location of the qubits is different. However, synchronization in this naïve geometric clock construction is a problem, resulting in transition rules which are not unique and need to be at least 4-local, if one wants to apply two qubit gates U_t while moving the qubits forward. The challenge is to make a geometric clock encoding with unique transition rules which would involve only a few particles with low dimension. In other words, the goal is to get the work qubits to move along in a coordinated fashion and to have the corresponding gates U_t applied to them along the way.

In the papers of Mizel [65], a geometric clock like this is planned to be implemented in a system of electrons hopping between quantum dots. The interactions in [65] involve 2 particles hopping through 4 sites.

Another geometric clock construction appeared in [7], where the authors used 6-dimensional particles on a 2D grid. However, it doesn't have unique transition rules and thus requires large penalty terms for illegal clock states. Their progression of states $|\Psi_t\rangle$ is

$$\begin{array}{c|c|c|c|c|c|c} \begin{array}{cc} \boxtimes & \boxtimes \\ \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \boxtimes \\ \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \boxtimes \\ \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \boxtimes \\ \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \boxtimes \\ \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \boxtimes \\ \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \boxtimes \\ \boxtimes & \odot \end{array} \\ \begin{array}{cc} \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \odot \end{array} \\ \begin{array}{cc} \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \odot \end{array} & \begin{array}{cc} \boxtimes & \odot \end{array} \end{array}$$

where the state space of each 6-dimensional particle is a direct sum $\mathcal{H}_{\boxtimes} \oplus \mathcal{H}_{\odot} \oplus \mathcal{H}_{\boxtimes} \oplus \mathcal{H}_{\square}$, with the 2-dimensional subspaces \mathcal{H}_{\boxtimes} and \mathcal{H}_{\odot} able to hold the state of a qubit. The transitions in this type of clock can be made unique and checkable without large penalty terms if one makes the particles 12-dimensional.

The recent construction of Aharonov, Gottesman, Irani and Kempe [5] shows that a geometric clock can be constructed even in 1D, on a chain of 12-dimensional particles. What is even better, the legal clock states are locally checkable and the transition rules are unique and 2-local. The underlying idea in getting the qubits to move is a 'pass the hat forward and then jump over it' technique illustrated in Figure 4-2. I explain this construction in detail in Section 4.3, where I also decrease the required dimensionality of particles to $d = 11$.

Janzing and Wocjan [48] also use a geometric clock in their complicated Hamiltonian Computer construction. Recently, myself and Wocjan [71] simplified their model greatly, and use a novel geometric clock construction with transition rules reminiscent of diffusion instead of the conventional quantum walk on a line. I present this result in Chapter 5 on Hamiltonian Quantum Cellular Automata.

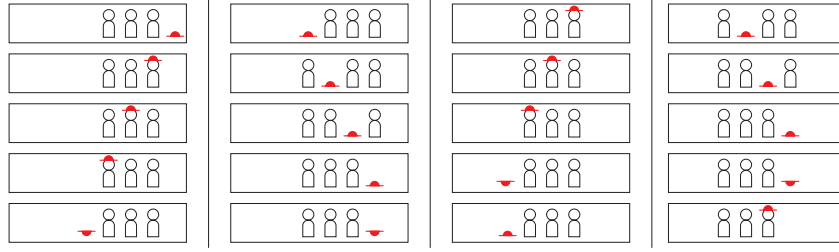


Figure 4-2: The ‘pass the hat forward and then jump over it’ technique for a geometric clock (qubit transport) on a line from [5]. The progression of 20 states is shown.

4.2 A new 3-local QMA complete Local Hamiltonian (Quantum 5-SAT from 3-local terms)

Quantum k -SAT, the special case where the Hamiltonian is a sum of local projectors was defined and studied by Bravyi [20] as a natural analogue of classical k -SAT. It is in P for $k = 2$, and it is QMA₁ complete for $k \geq 4$. However, the classification of quantum 3-SAT is still an open question. Its mother problem, 3-local Hamiltonian was shown to be QMA complete by Kempe and Regev in [54]. This result was further improved, showing that 2-local Hamiltonian is QMA complete in [53, 74]. These constructions use the Hamiltonian of the form (4.6), and encode the clock register using the unary domain wall clock (4.17). Since the terms in their Hamiltonian are no longer 5-local as in Kitaev’s proof in Section 1.3.2, the corresponding terms in H_{prop}^t do not only verify proper application of U_t , but also induce transitions from legal clock states into illegal ones. The subspace $\mathcal{H}_{work} \otimes \mathcal{H}_{legal}$ is thus no longer invariant under the action of H . To fix this, the penalty associated with illegal clock states is made high (scaling as a high polynomial in n), effectively forcing the ground state of the Hamiltonian to reside in the subspace of legal clock states. The Hamiltonian then contains $O(L)$ terms with weights that scale as $O(L^{12})$ and has thus norm $\|H\| = O(L^{13})$, where L is the number of gates in the computation [54]. This was later improved to weights of order $O(L^6)$, resulting in $\|H\| = O(L^7)$ in [53].

Note that there are two energy scales in this problem. The norm of the Hamiltonian $\|H\|$, and the energy difference $b - a \geq 1/\text{poly}(L)$ in the definition of local-Hamiltonian. It is important to keep track of both of these. One can always rescale the Hamiltonian such that $\|H\|$ is a constant, but this also shrinks the energy difference $b - a$. The quantity $\frac{b-a}{\|H\|}$ is an indicator of the strength of the result. Physically speaking, it relates the precision required in obtaining the ground state energy (finding whether it is below a or above b) to the strength of the interactions (as given by $\|H\|$). Proving QMA completeness for the local-Hamiltonian where $\frac{b-a}{\|H\|} = \text{const.}$ would give a quantum analogue of the famous PCP theorem [12, 11].

My goals are much simpler. I show in section 4.2.5 that the 3-local Hamiltonian problem is QMA complete using a Hamiltonian with $O(L)$ terms with only constant operator norms. This makes the norm of my Hamiltonian scale as $\|H\| = O(L)$, while I keep the same energy difference $b - a = 1/\text{poly}(L)$ as in the previous constructions. The $O(L^6)$ increase in the ratio of the two energy scales of the problem thus makes my construction physically much more interesting.

In this work, I show a new reduction from a verifier quantum circuit to a 3-local Hamiltonian. The novelty of my construction is that it leaves the space of legal clock-register states invariant. Therefore, the weights of the terms in my Hamiltonian do not scale with the size of the input problem. Such terms do appear in the constructions of [54] and [53]. As an intermediate step in the construction, I prove that quantum 3-SAT for qutrits is QMA₁-complete.

The rest of Section 4.2 is organized as follows. After reviewing Bravyi's proof that Quantum 4-SAT is QMA₁-complete in Section 4.2.1, I present a qutrit-clock construction in Section 4.2.2 and show that quantum 3-SAT for particles with dimensions $3 \times 2 \times 2$ (the interaction terms in the Hamiltonian couple one qutrit and two qubits) is QMA₁-complete. The existence of such a construction was previously mentioned but not specified by Bravyi and DiVincenzo in [20] as [27]. In Section 4.2.5 I show how to encode the qutrit clock particles from Section 4.2.2 into a pair of qubits in such a way that the Hamiltonian remains 3-local, obtaining a new construction of a QMA complete 3-local Hamiltonian. This Hamiltonian is composed of 4-local positive semidefinite operator terms. However, each of these 4-local operators is composed of only 3-local interaction terms. It is not a Quantum 3-SAT Hamiltonian, since the 3-local terms by themselves are not positive semidefinite operators. I discuss the complexity of Quantum 3-SAT and further directions in Section 4.2.8.

4.2.1 Bravyi's Quantum 4-SAT

In [20], Bravyi proved that quantum k -SAT belongs to QMA₁ for any constant k . Furthermore, he showed that quantum 4-SAT is QMA₁ complete using a new realization of the clock. In his construction Bravyi uses $L + 1$ clock particles with 4 states: unborn, active 1 (a_1 , input for a gate), active 2 (a_2 , output of a gate), and dead. These 4 states of a clock particle are easily realized by two qubits per clock particle. There are $2L$ legal clock states:

$$|C_{2k-1}\rangle = |\underbrace{d \dots d}_{k-1} a_1 \underbrace{u \dots u}_{L-k}\rangle, \quad \text{and} \quad |C_{2k}\rangle = |\underbrace{d \dots d}_{k-1} a_2 \underbrace{u \dots u}_{L-k}\rangle, \quad (4.33)$$

for $1 \leq k \leq L$. A clock Hamiltonian $H_{\text{clock}} = H_{\text{clockinit}} + \sum_{k=1}^{L-1} H_{\text{clock}}^{(k)}$ is required to check whether the states of the clock are legal.

$$\begin{aligned} H_{\text{clock}}^{(k)} &= |d\rangle \langle d|_k \otimes |u\rangle \langle u|_{k+1} \\ &+ |u\rangle \langle u|_k \otimes \left(|d\rangle \langle d| + |a_1\rangle \langle a_1| + |a_2\rangle \langle a_2| \right)_{k+1} \\ &+ \left(|a_1\rangle \langle a_1| + |a_2\rangle \langle a_2| \right)_k \otimes \left(|a_1\rangle \langle a_1| + |a_2\rangle \langle a_2| + |d\rangle \langle d| \right)_{k+1}, \end{aligned} \quad (4.34)$$

$$H_{\text{clockinit}} = |u\rangle \langle u|_1 + |d\rangle \langle d|_L. \quad (4.35)$$

The Hamiltonian checking the correct application of gates is $H_{\text{prop}} = \sum_{k=1}^L H_{\text{prop}}^{(k)}$, with

$$H_{\text{prop}}^{(k)} = \frac{1}{2} \left(\mathbb{I} \otimes |a_1\rangle \langle a_1|_k + \mathbb{I} \otimes |a_2\rangle \langle a_2|_k - U_k \otimes |a_2\rangle \langle a_1|_k - U_k^\dagger \otimes |a_1\rangle \langle a_2|_k \right). \quad (4.36)$$

Each such term verifies the correct application of the gate U_k between the states $|a_1\rangle$ and $|a_2\rangle$ of the k -th clock particle. This only requires interactions of the k -th clock particle (qubit pair) and the two work qubits the gate U_k is applied to. Each of the terms is thus a

4-local projector.

I need another Hamiltonian term to propagate the clock state $|C_{2k}\rangle$ into $|C_{2k+1}\rangle$ while leaving the work qubits untouched (that is, for the ground state $|\psi_{2k}\rangle_{work} = |\psi_{2k+1}\rangle_{work}$). This is done by the 4-local clock-propagation Hamiltonian $H_{clockprop} = \sum_{k=1}^{L-1} H_{clockprop}^{(k)}$, with

$$\begin{aligned} H_{clockprop}^{(k)} &= \frac{1}{2} \left(|a_2\rangle \langle a_2|_k \otimes |u\rangle \langle u|_{k+1} + |d\rangle \langle d|_k \otimes |a_1\rangle \langle a_1|_{k+1} \right) \\ &\quad - \frac{1}{2} \left(|d\rangle \langle a_2|_k \otimes |a_1\rangle \langle u|_{k+1} + |a_2\rangle \langle d|_k \otimes |u\rangle \langle a_1|_{k+1} \right). \end{aligned} \quad (4.37)$$

The final ingredients in this construction are

$$H_{init} = \sum_{n=1}^{N_a} |1\rangle \langle 1|_n \otimes |a_1\rangle \langle a_1|_1, \quad (4.38)$$

$$H_{out} = |0\rangle \langle 0|_{out} \otimes |a_2\rangle \langle a_2|_L. \quad (4.39)$$

Applying Kitaev's methods [56] to this construction, Bravyi shows that the quantum 4-SAT Hamiltonian (a sum of 4-local projectors)

$$H = H_{clock} + H_{clockprop} + H_{init} + H_{out} + H_{prop} \quad (4.40)$$

is QMA₁ complete.

Bravyi's original definition required all of the terms in the Hamiltonian to be projectors. However, as I have shown at the end of Section 4.1.1, using positive semidefinite operator terms in H instead of just restricting ourselves to projectors is an equivalent problem.

4.2.2 A qutrit clock implementation

In this section I present a new realization of the clock which builds on Bravyi's quantum 4-SAT realization described above. Using this clock construction, I prove that quantum 3-SAT for qutrits is QMA₁-complete.

First, I need to show that quantum 3-SAT with qutrits is in QMA. I can use Bravyi's proof that quantum k -SAT for qubits is in QMA₁ for any constant k . Given an instance of quantum 3-SAT for qutrits, I convert it into an instance of quantum 6-SAT for qubits by encoding each qutrit in two qubits and projecting out one of the four states. According to Bravyi, this problem is in QMA₁ and therefore so is the original quantum 3-SAT problem with qutrits.

For the other direction in the proof, I need to construct a quantum 3-SAT Hamiltonian for qutrits, corresponding to a given quantum verifier circuit U for a problem in QMA. The terms in the Hamiltonian I will construct act on the space of one qutrit and two qubits (particles with dimensions $3 \times 2 \times 2$).

4.2.3 Clock register construction

The clock-register construction in the previous section required 4 states for each clock particle: $|u\rangle, |a_1\rangle, |a_2\rangle$ and $|d\rangle$. Let me first explain why Bravyi's construction requires two "inactive" states: $|d\rangle$ and $|u\rangle$. If I only use $|d\rangle$ (i.e., have legal clock states of the form $|d \dots da_1 d \dots d\rangle$ and $|d \dots da_2 d \dots d\rangle$), I immediately get a 3-local Hamiltonian for qutrits.

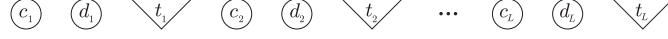


Figure 4-3: Clock register consisting of $2L$ qubits and L qutrits.

However, in Bravyi's construction, the first clock particle is never in the state $|u\rangle$, and the last one is never in the state $|d\rangle$ (see (4.35)). This ensures that at least one clock particle is in an active state. When not using the state $|u\rangle$, I can no longer exclude the state with no active particles $|dd\dots d\rangle$ in a simple local fashion.

I fix this by modifying the clock register as shown in in Fig.4-3. The clock register now consists of $2L$ qubits and L qutrits. The $2L$ qubits $c_1, d_1, \dots, c_k, d_k$ play the role of the usual unary $|1\dots 1100\dots 0\rangle$ clock representation, while the L qutrits t_1, t_2, \dots, t_k play the role of Bravyi's clock with just three states (d, a_1, a_2) .

I define the legal clock space \mathcal{H}_{legal} as the space spanned by the $3L$ states $|C_m\rangle$. These states are defined for $1 \leq k \leq L$ as follows:

$$\begin{aligned}
 |C_{3k-2}\rangle &= |\underbrace{(11d)(11d)\dots(11d)}_{k-1 \text{ times}} \underbrace{(10d)(00d)(00d)\dots(00d)}_{L-k \text{ times}}\rangle, \\
 |C_{3k-1}\rangle &= |\underbrace{(11d)(11d)\dots(11d)}_{k-1 \text{ times}} (11a_1) \underbrace{(00d)(00d)\dots(00d)}_{L-k \text{ times}}\rangle, \\
 |C_{3k}\rangle &= |\underbrace{(11d)(11d)\dots(11d)}_{k-1 \text{ times}} (11a_2) \underbrace{(00d)(00d)\dots(00d)}_{L-k \text{ times}}\rangle.
 \end{aligned} \tag{4.41}$$

The first state ($|C_{3k-2}\rangle$) corresponds to the time when the qubits are “in transport” from the previous gate to the current (k th) gate. The second one corresponds to the time right before application of gate U_k and the third corresponds to the time right after the gate U_k was applied. The structure of such clock register can be understood as two coupled “unary” clocks, the qubit one (c_k, d_k) of the $11\dots 11100\dots 00$ type and the qutrit one (the t_k 's) of the $00\dots 00100\dots 00$ type. Formally, the legal clock states satisfy the following constraints:

1. if d_k is 1, then c_k is 1.
2. if c_{k+1} is 1, then d_k is 1.
3. if t_k is active (a_1/a_2), then d_k is 1.
4. if t_k is active (a_1/a_2), then c_{k+1} is 0.
5. if d_k is 1 and c_{k+1} is 0, then t_k is not dead (d).
6. c_1 is 1.
7. if d_L is 1, then t_L is not dead.

The last two conditions are required to exclude the clock states $|(00d)(00d)\dots(00d)\rangle$ and $|(11d)(11d)\dots(11d)\rangle$ that have no active clock terms. The clock Hamiltonian $H_{clock} =$

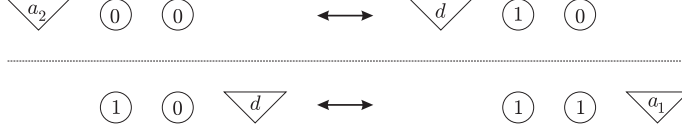


Figure 4-4: Illustration of the two-step clock pointer propagation.

$H_{clockinit} + \sum_{k=1}^L H_{clock1}^{(k)} + \sum_{k=1}^{L-1} H_{clock2}^{(k)}$ verifies the above constraints.

$$\begin{aligned}
H_{clock1}^{(k)} &= |01\rangle \langle 01|_{c_k, d_k} + |0\rangle \langle 0|_{d_k} \otimes \left(|a_1\rangle \langle a_1| + |a_2\rangle \langle a_2| \right)_{t_k}, \\
H_{clock2}^{(k)} &= |01\rangle \langle 01|_{d_k, c_{k+1}} + \left(|a_1\rangle \langle a_1| + |a_2\rangle \langle a_2| \right)_{t_k} \otimes |1\rangle \langle 1|_{c_{k+1}} \\
&\quad + |1d0\rangle \langle 1d0|_{d_k, t_k, c_{k+1}}, \\
H_{clockinit} &= |0\rangle \langle 0|_{c_1} + |1\rangle \langle 1|_{d_L} \otimes |d\rangle \langle d|_{t_L}.
\end{aligned} \tag{4.42}$$

Only the last term in $H_{clock2}^{(k)}$ is a 3-local projector, acting on the space of two qubits and one qutrit. The rest of the terms are 2-local projectors on two qubits, or a qubit and a qutrit. The space of legal clock states \mathcal{H}_{legal} is the kernel of the clock Hamiltonian H_{clock} .

4.2.4 Checking correct application of gates and clock propagation

The gate-checking Hamiltonian $H_{prop} = \sum_{k=1}^L H_{prop}^{(k)}$ is an analogue of (4.36), with

$$H_{prop}^{(k)} = \frac{1}{2} \left(\mathbb{I}_{work} \otimes \left(|a_1\rangle \langle a_1|_{t_k} + |a_2\rangle \langle a_2|_{t_k} \right) - U_k \otimes |a_2\rangle \langle a_1|_{t_k} - U_k^\dagger \otimes |a_1\rangle \langle a_2|_{t_k} \right). \tag{4.43}$$

The clock propagation proceeds in two steps. First, the “active” spot in the clock register moves from the state $|a_2\rangle$ of the qutrit t_k to the $|10\rangle$ state of the next two qubits c_{k+1}, d_{k+1} . After this, it moves into the state $|a_1\rangle$ of the next qutrit t_{k+1} , as in Fig.4-4. The Hamiltonian checking whether this happened, while the work qubits were left untouched, is $H_{clockprop} = \sum_{k=1}^L H_{clockprop1}^{(k)} + \sum_{k=1}^{L-1} H_{clockprop2}^{(k)}$ with

$$\begin{aligned}
H_{clockprop1}^{(k)} &= \frac{1}{2} \left(|10\rangle \langle 10|_{c_k, d_k} \otimes |d\rangle \langle d|_{t_k} + |11\rangle \langle 11|_{c_k, d_k} \otimes |a_1\rangle \langle a_1|_{t_k} \right) \\
&\quad - \frac{1}{2} \left(|11\rangle \langle 10|_{c_k, d_k} \otimes |a_1\rangle \langle d|_{t_k} + |10\rangle \langle 11|_{c_k, d_k} \otimes |d\rangle \langle a_1|_{t_k} \right), \\
H_{clockprop2}^{(k)} &= \frac{1}{2} \left(|a_2\rangle \langle a_2|_{t_k} \otimes |00\rangle \langle 00|_{c_{k+1}, d_{k+1}} + |d\rangle \langle d|_{t_k} \otimes |10\rangle \langle 10|_{c_{k+1}, d_{k+1}} \right) \\
&\quad - \frac{1}{2} \left(|d\rangle \langle a_2|_{t_k} \otimes |10\rangle \langle 00|_{c_{k+1}, d_{k+1}} + |a_2\rangle \langle d|_{t_k} \otimes |00\rangle \langle 10|_{c_{k+1}, d_{k+1}} \right).
\end{aligned} \tag{4.44}$$

The input Hamiltonian checks whether the computation has properly initialized ancilla qubits.

$$H_{init} = \sum_{n=1}^{N_a} |1\rangle \langle 1|_n \otimes |a_1\rangle \langle a_1|_{t_1}. \tag{4.45}$$

Finally, the output Hamiltonian checks whether the result of the computation was 1.

$$H_{out} = |0\rangle\langle 0|_{out} \otimes |a_2\rangle\langle a_2|_{t_L}. \quad (4.46)$$

All of the terms coming from (4.42) – (4.46) in the Hamiltonian

$$H = H_{clock} + H_{clockprop} + H_{init} + H_{out} + H_{prop}. \quad (4.47)$$

are projectors. Therefore, the ground state has energy zero if and only if there exists a zero energy eigenstate of all of the terms. If there exists a witness $|\varphi\rangle$ on which the computation U gives the result 1 with probability 1, I can construct a computational history state (4.7) for a modified circuit $\tilde{U} = U_L \cdot \mathbb{I} \cdot \mathbb{I} \cdot U_{L-1} \cdot \mathbb{I} \cdot \mathbb{I} \cdots U_1 \cdot \mathbb{I}$, where the “identity” gates correspond to the clock propagation in my construction, with nothing happening to the work qubits. This state is a zero eigenvector of all of the terms in the Hamiltonian (4.47).

I now need to prove that if no witness exists (the answer to the problem is “no”), then the ground state energy of (4.47) is lower bounded by $1/\text{poly}(L)$. Let me decompose the Hilbert space into

$$\mathcal{H} = (\mathcal{H}_{work} \otimes \mathcal{H}_{legal}) \oplus (\mathcal{H}_{work} \otimes \mathcal{H}_{legal}^\perp). \quad (4.48)$$

where \mathcal{H}_{legal} is the space of legal clock states (on which $H_{clock}|\alpha\rangle = 0$). The Hamiltonian (4.47) leaves this decomposition invariant, because it does not induce transitions between legal and illegal clock states. Since any state in $\mathcal{H}_{legal}^\perp$ violates at least one term in H_{clock} , the lowest eigenvalue of the restriction of (4.47) to $\mathcal{H}_{work} \otimes \mathcal{H}_{legal}^\perp$ is at least 1. On the other hand, the restriction of H to the legal clock space is identical to the legal clock space restriction of Bravyi’s Hamiltonian (4.40) from the previous section. Therefore, his proof using the methods of Kitaev [56] applies to this case as well. He shows that if a no witness state for the quantum circuit U exists, then the ground state energy of the restriction of (4.40) to $\mathcal{H}_{work} \otimes \mathcal{H}_{legal}$ lower bounded by $1/\text{poly}(L)$. This means that if there is no witness state for the verifier circuit U , the ground state of (4.47) is lower bounded by $1/\text{poly}(L)$. This concludes the proof that quantum 3-SAT with qutrits (in fact, quantum 3-SAT on particles with dimensions $3 \times 2 \times 2$, a qutrit and two qubits) is QMA₁ complete.

The existence of another $3 \times 2 \times 2$ construction for quantum 3-SAT (i.e., a Hamiltonian with terms acting on one qutrit and two qubits) was already mentioned in [20] as [27], though that construction was not specified. I choose to write out my result explicitly as it serves as a natural intermediate step towards the new 3-local Hamiltonian construction described in the following section.

4.2.5 The new 3-local QMA complete construction (for qubits)

In Bravyi’s Quantum 4-SAT construction [20], the clock particles (qubit pairs) can be in 4 states. In the previous section, I required only 3 states of the clock particles and used qutrits as particles with these three states. I start with the clock-register construction (see Fig.4-3) from the previous section, with legal states as in (4.41). However, I now encode the three states of every clock qutrit t_k using a pair of qubits r_k, s_k . The new clock register

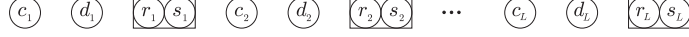


Figure 4-5: Clock register construction with $2L + 2L$ qubits.

is depicted in Fig.4-5.

$$\begin{aligned} |a_1\rangle_{t_k} &\rightarrow \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)_{r_k, s_k}, & |d\rangle_{t_k} &\rightarrow |00\rangle_{r_k, s_k}, \\ |a_2\rangle_{t_k} &\rightarrow \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)_{r_k, s_k}. \end{aligned} \quad (4.49)$$

This encoding allows me to obtain a new 3-local Hamiltonian construction (for the QMA-complete 3-local Hamiltonian problem). This Hamiltonian is a quantum 4-SAT Hamiltonian whose 4-local positive semidefinite operator terms consist of just 3-local interactions.

I am looking for 3-local terms that flip between the clock states $|a_1\rangle \leftrightarrow |a_2\rangle$, while simultaneously (un)applying a 2-qubit gate U_k on two work qubits. I encode the active states of a clock particle into entangled states, and thus I am able to flip between these clock states with a term like Z_1 involving only one of the clock particles. Thus my 2-qubit gate checking Hamiltonian involves only 3-local terms (acting on one clock qubit r_k or s_k and the two work qubits on which the gate U_k acts).

First, I define the legal clock space \mathcal{H}_{legal} as the space spanned by the $3L$ states $|C_m\rangle$. These states are defined for $1 \leq k \leq L$ as follows (compare to (4.41)):

$$\begin{aligned} |C_{3k-2}\rangle &= \underbrace{|(11)(00) \dots (11)(00)\rangle}_{k-1 \text{ times}} \otimes |10\rangle_{c_k, d_k} \otimes |00\rangle_{r_k, s_k} \otimes \underbrace{|(00)(00) \dots (00)\rangle}_{L-k \text{ times}}, \\ |C_{3k-1}\rangle &= \underbrace{|(11)(00) \dots (11)(00)\rangle}_{k-1 \text{ times}} \otimes |11\rangle_{c_k, d_k} \otimes \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)_{r_k, s_k} \otimes \underbrace{|(00)(00) \dots (00)\rangle}_{L-k \text{ times}}, \\ |C_{3k}\rangle &= \underbrace{|(11)(00) \dots (11)(00)\rangle}_{k-1 \text{ times}} \otimes |11\rangle_{c_k, d_k} \otimes \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)_{r_k, s_k} \otimes \underbrace{|(00)(00) \dots (00)\rangle}_{L-k \text{ times}}. \end{aligned} \quad (4.50)$$

Similarly to the construction of the previous section, the first state ($|C_{3k-2}\rangle$) corresponds to the time when the qubits are “in transport” from the previous gate to the current (k th) gate. The second one corresponds to the time right before application of gate U_k and the third corresponds to the time right after the gate U_k was applied.

Formally, the legal clock states for this construction satisfy the following constraints:

1. if d_k is 1, then c_k is 1.
2. if c_{k+1} is 1, then d_k is 1.
3. the pair r_k, s_k is not in the state $|11\rangle$.
4. if the pair r_k, s_k is active (in the state $(|01\rangle \pm |10\rangle)/\sqrt{2}$), then d_k is 1.
5. if the pair r_k, s_k is active (in the state $(|01\rangle \pm |10\rangle)/\sqrt{2}$), then c_{k+1} is 0.
6. if d_k is 1 and c_{k+1} is 0, then the pair r_k, s_k is not dead (in the state $|00\rangle$).

7. c_1 is 1.

8. if d_L is 1, then the pair r_L, s_L is not dead (in the state $|00\rangle$).

The last two conditions are required to make the clock states $|(00)(00)\dots(00)(00)\rangle$ and $|(11)(00)\dots(11)(00)\rangle$ with no active spots illegal. The clock Hamiltonian $H_{clock} = H_{clockinit} + \sum_{k=1}^L H_{clock1}^{(k)} + \sum_{k=1}^{L-1} H_{clock2}^{(k)}$ verifies the above constraints.

$$\begin{aligned} H_{clock1}^{(k)} &= |01\rangle\langle 01|_{c_k, d_k} + |0\rangle\langle 0|_{d_k} \otimes (|1\rangle\langle 1|_{r_k} + |1\rangle\langle 1|_{s_k}) + |11\rangle\langle 11|_{r_k, s_k}, \quad (4.51) \\ H_{clock2}^{(k)} &= |01\rangle\langle 01|_{d_k, c_{k+1}} + (|1\rangle\langle 1|_{r_k} + |1\rangle\langle 1|_{s_k}) \otimes |1\rangle\langle 1|_{c_{k+1}} + h_4^{(k)}, \\ h_4^{(k)} &= |1\rangle\langle 1|_{d_k} \otimes \frac{1}{2} (Z_{r_k} + Z_{s_k}) \otimes |0\rangle\langle 0|_{c_{k+1}} + |11\rangle\langle 11|_{r_k, s_k}, \\ H_{clockinit} &= |0\rangle\langle 0|_{c_1} + |1\rangle\langle 1|_{d_L} \otimes |00\rangle\langle 00|_{r_L, s_L}. \end{aligned}$$

All of the terms involve only 3-local interactions. All terms in $H_{clock1}^{(k)}$, $H_{clock2}^{(k)}$ and $H_{clockinit}$, are projectors. The term $h_4^{(k)}$ corresponds to the sixth legal state condition. It is a 4-local projector onto the space spanned by (illegal clock) states $|1_{d_k}(00)_{r_k, s_k}0_{c_{k+1}}\rangle$, $|0(11)0\rangle$, $|0(11)1\rangle$ and $|1(11)1\rangle$. Note that even though $h_4^{(k)}$ is a 4-local projector, it is only constructed of 3-local terms.

4.2.6 Checking gate application with 3-local terms

Let me start by writing out a Hamiltonian that checks the correct application of a single-qubit gate U_k .

$$H_{prop}^{(k), one-qubit} = \frac{1}{2} \begin{pmatrix} \mathbb{I} \otimes |01 - 10\rangle\langle 01 - 10|_{r_k, s_k} - U_k \otimes |01 + 10\rangle\langle 01 - 10|_{r_k, s_k} \\ \mathbb{I} \otimes |01 + 10\rangle\langle 01 + 10|_{r_k, s_k} - U_k^\dagger \otimes |01 - 10\rangle\langle 01 + 10|_{r_k, s_k} \end{pmatrix}, \quad (4.52)$$

where $|01 \pm 10\rangle$ is a shortcut notation for the normalized entangled states $(|01\rangle \pm |10\rangle)/\sqrt{2}$. This Hamiltonian is a 3-local projector. Note that in the case $U_k = \mathbb{I}$, this Hamiltonian becomes the projector $(\mathbb{I} - X)/2$ on the space of active clock states $\{|01 - 10\rangle, |01 + 10\rangle\}$.

For a two-qubit gate U_k , the above construction would be 4-local. However, I am able to construct this 4-local projector using only 3-local terms. To do this, I require the 2-qubit gate to be symmetric'' $U_k = U_k^\dagger$. This is a universal construction, since the symmetric gate CNOT (or C_ϕ) is universal. Now I can write

$$H_{prop}^{(k), two-qubit} = \frac{1}{2} \left(\mathbb{I} \otimes \frac{1}{2} (\mathbb{I} - Z_{r_k} Z_{s_k}) - U_k \otimes \frac{1}{2} (Z_{r_k} - Z_{s_k}) \right). \quad (4.53)$$

The first term in this Hamiltonian, $(\mathbb{I} - Z_{r_k} Z_{s_k})/2$, is a projector onto the space of active clock states, $|01 \pm 10\rangle_{r_k, s_k}$, as I needed. The second term contains $(Z_{r_k} - Z_{s_k})/2$, which has zero eigenvalues for the states $|00\rangle_{r_k, s_k}$ and $|11\rangle_{r_k, s_k}$, and flips between the states $|01 - 10\rangle_{r_k, s_k} \leftrightarrow |01 + 10\rangle_{r_k, s_k}$. Altogether, this is a 4-local projector made out of only 3-local terms.

4.2.7 Clock propagation

After a gate U_k is applied, I need to "propagate" the pointer (the active state of the qubit pair r_k, s_k) to the next pair of qubits r_{k+1}, s_{k+1} . This is done in two steps, as shown in

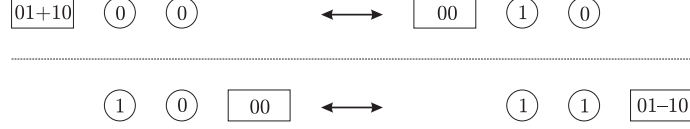


Figure 4-6: Illustration of the two-step clock pointer propagation.

Fig.4-6.

For each step, I want to write a 3-local positive semidefinite Hamiltonian with terms acting on 4 consecutive qubits $r_k, s_k, c_{k+1}, d_{k+1}$ (for the second step of the clock pointer propagation, the four qubits in play are c_k, d_k, r_k, s_k), with zero eigenvalue for the legal clock-propagation states, and perhaps also some illegal clock states, which will be disallowed by other terms in the Hamiltonian (H_{clock}). For the first step, these desired eigenvectors with zero eigenvalues are

$$\begin{aligned}
 |\alpha_1\rangle_{r_k, s_k, c_{k+1}, d_{k+1}} &= |00\rangle_{r_k, s_k} |00\rangle_{c_{k+1}, d_{k+1}}, \\
 |\alpha_2\rangle_{r_k, s_k, c_{k+1}, d_{k+1}} &= \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)_{r_k, s_k} |00\rangle_{c_{k+1}, d_{k+1}}, \\
 |\alpha_3\rangle_{r_k, s_k, c_{k+1}, d_{k+1}} &= \frac{1}{2}(|01\rangle + |10\rangle)_{r_k, s_k} |00\rangle_{c_{k+1}, d_{k+1}} + \frac{1}{\sqrt{2}}|00\rangle_{r_k, s_k} |10\rangle_{c_{k+1}, d_{k+1}}, \\
 |\alpha_4\rangle_{r_k, s_k, c_{k+1}, d_{k+1}} &= |00\rangle_{r_k, s_k} |11\rangle_{c_{k+1}, d_{k+1}}.
 \end{aligned} \tag{4.54}$$

The state that I want to exclude (make it a nonzero eigenvector) is the legal clock state with incorrect pointer propagation:

$$|\alpha^\perp\rangle_{r_k, \dots, d_{k+1}} = \frac{1}{2}(|01\rangle + |10\rangle)_{r_k, s_k} |00\rangle_{c_{k+1}, d_{k+1}} - \frac{1}{\sqrt{2}}|00\rangle_{r_k, s_k} |10\rangle_{c_{k+1}, d_{k+1}}. \tag{4.55}$$

Let me present the Hamiltonian penalizing this state.

$$\begin{aligned}
 H_{clockprop1}^{(k)} &= |10\rangle\langle 10|_{c_{k+1}, d_{k+1}} \\
 &+ \frac{1}{2}(|01\rangle + |10\rangle)(\langle 01| + \langle 10|)_{r_k, s_k} \\
 &- \frac{1}{\sqrt{2}}(|0\rangle\langle 1|_{r_k} + |0\rangle\langle 1|_{s_k}) \otimes |10\rangle\langle 00|_{c_{k+1}, d_{k+1}} \\
 &- \frac{1}{\sqrt{2}}(|1\rangle\langle 0|_{r_k} + |1\rangle\langle 0|_{s_k}) \otimes |00\rangle\langle 10|_{c_{k+1}, d_{k+1}} \\
 &+ 2|11\rangle\langle 11|_{r_k, s_k}.
 \end{aligned} \tag{4.56}$$

It is positive semidefinite, with eigenvalues 0 ($\times 7$), 1 ($\times 4$), 2 ($\times 3$) and 3 ($\times 2$). Its zero energy eigenvectors are $|\alpha_1\rangle, |\alpha_2\rangle, |\alpha_3\rangle, |\alpha_4\rangle$ expressed above, and three illegal clock states, $|00\rangle_{r_k, s_k} |01\rangle_{c_{k+1}, d_{k+1}}, (|01\rangle - |10\rangle)|01\rangle$ and $(|01\rangle - |10\rangle)|11\rangle$. The state $|\alpha^\perp\rangle$ is an eigenvector of $H_{clockprop1}^{(k)}$ with eigenvalue 2. This means that $H_{clockprop1}^{(k)}$ fulfills its job in punishing the legal states of the clock register (4.55), which do not correctly propagate the clock. This Hamiltonian term is a positive semidefinite operator, while Bravyi's original definition of quantum k -SAT requires the terms in the Hamiltonian to be projectors. However, as I have shown at the end of Section 4.1.1, quantum k -SAT with positive semidefinite operator terms

is equivalent to quantum k -SAT with only projector terms.

For the second step, the desired zero energy eigenvectors are

$$\begin{aligned}
|\beta_1\rangle_{c_k, d_k, r_k, s_k} &= |00\rangle_{c_k, d_k} |00\rangle_{r_k, s_k}, \\
|\beta_2\rangle_{c_k, d_k, r_k, s_k} &= \frac{1}{\sqrt{2}} |10\rangle_{c_k, d_k} |00\rangle_{r_k, s_k} + |11\rangle_{c_k, d_k} \frac{1}{2} (|01\rangle - |10\rangle)_{r_k, s_k}, \\
|\beta_3\rangle_{c_k, d_k, r_k, s_k} &= |11\rangle_{c_k, d_k} \frac{1}{\sqrt{2}} (|01\rangle + |10\rangle)_{r_k, s_k}, \\
|\beta_4\rangle_{c_k, d_k, r_k, s_k} &= |11\rangle_{c_k, d_k} |00\rangle_{r_k, s_k},
\end{aligned} \tag{4.57}$$

and the state I want to exclude is

$$|\beta^\perp\rangle_{c_k, d_k, r_k, s_k} = \frac{1}{\sqrt{2}} |10\rangle_{c_k, d_k} |00\rangle_{r_k, s_k} - |11\rangle_{c_k, d_k} \frac{1}{2} (|01\rangle - |10\rangle)_{r_k, s_k}. \tag{4.58}$$

The Hamiltonian with these properties is a simple analogue of (4.56):

$$\begin{aligned}
H_{clockprop2}^{(k)} &= |10\rangle \langle 10|_{c_k, d_k} \\
&+ \frac{1}{2} (|01\rangle - |10\rangle) (\langle 01| - \langle 10|)_{r_k, s_k} \\
&- |11\rangle \langle 10|_{c_k, d_k} \otimes \frac{1}{\sqrt{2}} \left(-|1\rangle \langle 0|_{r_k} + |1\rangle \langle 0|_{s_k} \right) \\
&- |10\rangle \langle 11|_{c_k, d_k} \otimes \frac{1}{\sqrt{2}} \left(-|0\rangle \langle 1|_{r_k} + |0\rangle \langle 1|_{s_k} \right) \\
&+ 2 |11\rangle \langle 11|_{r_k, s_k}.
\end{aligned} \tag{4.59}$$

This is again a positive semidefinite operator with eigenvalues $0 (\times 7)$, $1 (\times 4)$, $2 (\times 3)$ and $3 (\times 2)$. Its zero energy eigenvectors are $|\beta_1\rangle$, $|\beta_2\rangle$, $|\beta_3\rangle$, $|\beta_4\rangle$ expressed above, and three illegal clock states, $|01\rangle_{c_k, d_k} |00\rangle_{r_k, s_k}$, $|01\rangle (|01\rangle + |10\rangle)$ and $|00\rangle (|01\rangle + |10\rangle)$, which are penalized by H_{clock} . The state $|\beta^\perp\rangle$ is an eigenvector of $H_{clockprop2}^{(k)}$ with eigenvalue 2, which is what I intended $H_{clockprop2}^{(k)}$ to do.

Just as in the previous section, the total Hamiltonian leaves the decomposition into $\mathcal{H}_{legal} \oplus \mathcal{H}_{legal}^\perp$ invariant while all illegal clock states violate at least one term in H_{clock} . Again, up to a constant prefactor, the restriction of H to the legal clock space is the same as that of the Hamiltonian in (4.40). The proof of the necessary separation between positive and negative instances then follows the proof in the previous section. This concludes the proof that quantum 4-SAT with positive semidefinite operators made out of 3-local terms is QMA₁ complete.

4.2.8 Discussion and further directions

In the above, I proved that quantum 3-SAT for particles with dimensions $3 \times 2 \times 2$ is QMA₁ complete. I have shown in Section 4.1.1 that quantum k -SAT with positive semidefinite operator terms (not just projectors) is equivalent to quantum k -SAT with projector terms. I presented a new 3-local construction of a quantum 4-SAT Hamiltonian with positive semidefinite operator terms, proving that quantum 4-SAT with 3-local interactions is QMA₁ complete.

The currently known complexities of classical and quantum satisfiability problems are

shown in Table 4.1. Quantum 3-SAT contains classical 3-SAT and therefore is NP-hard. Unlike classical k -SAT, which is known to be NP-complete for $k \geq 3$, quantum k -SAT is only known to be QMA₁ complete for $k \geq 4$ [20]. In my opinion, it is unlikely that one can show that quantum 3-SAT ($k = 3$) is also complete for QMA₁. One indication for this arises in my numerical explorations, where random instances of quantum 3-SAT for a reasonable number of clauses generally have no solutions, unless the clauses exclude non-entangled states. This may suggest that the hardness of quantum 3-SAT actually lies only in the classical instances (3-SAT) and a classical verifier circuit for quantum 3-SAT might exist.

Another reason comes from dimension counting. This argument, however, is only valid for the specific encoding of a circuit into the Hamiltonian I used. I worked with a tensor product space $\mathcal{H}_{work} \otimes \mathcal{H}_{clock}$, encoding the computation in the history state (4.7). Encoding an interaction of two qubits requires at least an $4 + 4 = 8$ dimensional space (4 for the two qubits before the interaction and 4 for the qubits after the interaction). On a first glance, a three-local projector on the space of two work qubits and one clock qubit ($2 \times 2 \times 2 = 8$) seems to suffice. However, I must ensure that this interaction only occurs at a specific clock time. When the two work qubits interact with just a single clock qubit (flipping it between states before/after interaction) ambiguities and legal-illegal clock state transitions are unavoidable. This transforms the problem from the SAT-type (determining whether a *simultaneous* ground state of *all* terms in the Hamiltonian exists) to the MAX-SAT type problem (determining the properties of the ground state energy of the *sum of terms* in the Hamiltonian). A single clock qubit cannot both determine the exact time of an interaction and distinguish between the states before and after the interaction. I managed to overcome this obstacle by using three-dimensional clock particles with states d , a_1 and a_2 and a $2 \times 2 \times 3 = 12$ dimensional space for encoding the interactions. I believe that this can not be further improved with more clever clock-register realizations within the usual $\mathcal{H} = \mathcal{H}_{work} \otimes \mathcal{H}_{clock}$ framework. However, a recent novel idea by Eldar and Regev [28] presented in Section 4.4.1 could be a step in this direction. They introduce a novel ‘triangle’ clock construction, and show that Quantum 2-SAT for particles with dimension 5×3 is QMA₁-complete.

A different approach to encoding a quantum computation into the ground state of a Hamiltonian is the geometric clock (see Section 4.1.3 found in the work of Aharonov et.al [7]. Their idea is to lay out the qubits in space in such a way, that the shape of the state (the locations of the work qubits moving around in the system) uniquely corresponds to a clock time. Their motivation was to show that adiabatic quantum computation [7] is polynomially equivalent to the circuit model. As a side result, they showed that the nearest-neighbor 2-local Hamiltonian problem with 6-dimensional particles is QMA complete. Actually, as Kempe et al. proved, even 2-local Hamiltonian with 2-dimensional particles (qubits) is QMA complete [53]. I did not succeed to improve or reproduce the construction of quantum 3-SAT for particles with dimensions $3 \times 2 \times 2$ using the geometric clock framework. However, one can use the idea of a geometric clock to construct quantum 2-SAT for higher dimensional particles (qudits). I know that classical 2-SAT for particles with dimensions 3×3 contains graph coloring, and is thus NP-complete. Using a rather straightforward modification of the Aharonov et. al. construction, one can prove that quantum 2-SAT for 12-dimensional particles is QMA complete. Combining the work/clock and the geometric clock constructions, a much tighter result can be shown. Specifically, one can construct quantum 2-SAT for particles with dimensions 9×4 and prove that it is QMA₁ complete. However, this is now made obsolete by the result of Eldar and Regev presented in Section 4.4.1 which does not use a geometric clock. It remains to be seen what are the minimal

dimensions of particles for which quantum 2-SAT is QMA_1 complete.

4.3 Quantum 2-SAT on a line

In 2007, Aharonov, Gottesman, Irani and Kempe [5] proved a surprising result about the power of quantum systems on a line. One might have thought before that finding ground states or simulating time evolution of spin chains is not too hard a task with today's modern DMRG and MPS methods (see Chapter 3). However, [5] shows that finding ground states of 2-local Hamiltonians on a line of $d = 12$ dimensional particles is QMA complete. Utilizing this construction, recently Schuch, Cirac and Verstraete [84] proved that even finding ground states of 2-local Hamiltonians on a line whose ground states are known to be Matrix Product States with constant dimension is NP-hard.

In this section I revisit [5] and show that the problem remains QMA complete even for particles with dimension $d = 11$. Moreover, [5] shows that one can use a simplified nearest-neighbor Hamiltonian on a line of $d = 9$ -dimensional particles to perform universal Quantum Adiabatic Computation. This result was recently improved to a system of 8-dimensional particles by Chase and Landahl [22]. Concurrently with their work, myself with P. Wocjan in [71] have proved that a translationally invariant Hamiltonian on a line of 10-dimensional particles can be used for this purpose. When I release the translational invariance requirement, I get a Hamiltonian for a chain of 8-dimensional particles as well. I present this result later in Section 5.2.

4.3.1 A Line of $d = 12$ Dimensional Particles (Quantum (12,12)-SAT)

As discussed in Section 4.1.3, to show that a Local Hamiltonian (or Quantum-SAT) problem is QMA complete, one usually proceeds as follows. First, encode a quantum circuit into a progression of orthogonal states $|\Psi_t\rangle$ of a larger system, either by adding a clock register to the register holding the work qubits, or by using a geometric clock construction. The corresponding Hamiltonian is then constructed to check the proper transition rules for the progression of states $|\Psi_t\rangle$. What remains to be shown are the properties of the low-lying spectrum of H .

The $d = 12$ Geometric clock on a line

On a line, only nearest neighbor interactions are allowed. This rules out the use of a separate clock register, which means a geometric clock (see Section 4.1.3) is a necessity. This is also a reason for the dimensionality ($d = 12$) of the particles. The state space of each particle consists of four two-dimensional subspaces

$$\blacklozenge : \text{a qubit marked as 'active'}, \quad (4.60)$$

$$\oplus : \text{a qubit labeled 'messenger'}, \quad (4.61)$$

$$\otimes : \text{a qubit to the left of the active site}, \quad (4.62)$$

$$\odot : \text{a qubit to the right of the active site}, \quad (4.63)$$

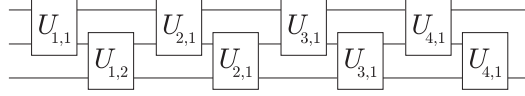


Figure 4-7: A quantum circuit for $n = 3$ qubits with $K = 4$ rounds of nearest neighbor gates $U_{k,s}$.

and four more states, each of them 1-dimensional

$$\boxtimes : \text{ to the left of the qubits ('done/dead'),} \quad (4.64)$$

$$\boxdot : \text{ to the right of the qubits ('ready/unborn'),} \quad (4.65)$$

$$\circlearrowleft : \text{ a 'turn' state at the qubit sequence boundary,} \quad (4.66)$$

$$\triangleleft : \text{ a 'push' state used to move the qubits to the right.} \quad (4.67)$$

Altogether, the state space of each particle is 12-dimensional, with the structure

$$\mathcal{H}_{12} = Q \oplus h = (q \otimes l) \oplus h = (\bigcirc \otimes \{\blacktriangleright, \triangleright, \times, \cdot\}) \oplus \{\boxtimes, \boxdot, \circlearrowleft, \triangleleft\}, \quad (4.68)$$

where Q is the subspace for the 4 types of qubits with q the internal state space of a qubit and l its four possible labels. The state space for the four other states is labeled h .

Without loss of generality, take a quantum circuit on n qubits which consists of K rounds of nearest neighbor gates as in Figure 4-7

$$U = (U_{K,n-1}U_{K,n-2}\dots U_{K,1})\dots (U_{2,n-1}U_{2,n-2}\dots U_{2,1})(U_{1,n-1}U_{1,n-2}\dots U_{1,1}), \quad (4.69)$$

where the gate $U_{k,s}$ acts on the pair of qubits $s, s+1$. I now encode the progression of U into a set of states of a line of qudits with length nK . The initial state $|\psi_1\rangle$ has a sequence of n qubits on the left, and the rest of the chain is in the 'ready' state \boxdot . The leftmost qubit is active \blacktriangleright . Here I write out the initial state corresponding to a quantum circuit on $n = 3$ qubits with $G = 4 \times 2$ nearest neighbor gates (here $K = 4$, as in Figure 4-7).

$$|\psi_1\rangle = \blacktriangleright \circlearrowleft \circlearrowleft \boxdot \boxdot \boxdot \boxdot \boxdot \boxdot \boxdot \boxdot \boxdot \boxdot \quad (4.70)$$

Using a few rules, a progression of states $|\psi_t\rangle$ is now constructed. The rules are prepared in such a way that there is always only one possible state $|\psi_{t+1}\rangle$ for a given state $|\psi_t\rangle$.

The first rule says that when the active spot \blacktriangleright is at position $kn + m$ where $1 \leq k \leq K$ and $1 < m \leq n$, it can pass to the right. When this happens, the gate $U_{k,m}$ is first applied to the state of the corresponding qubits $(kn + m, kn + m + 1)$, while their labels change from $\blacktriangleright \boxdot$ to $\boxtimes \blacktriangleright$.

$$\begin{aligned} 1a \quad : \quad \blacktriangleright \circlearrowleft &\longrightarrow U_{k,m}(\boxtimes \blacktriangleright) \\ &\text{on particles } (kn + m, kn + m + 1). \end{aligned} \quad (4.71)$$

On the other hand, the active spot \blacktriangleright just moves to the right, changing the qubit labels

without modifying their internal states.

$$1b : \blacktriangleright \circ \longrightarrow \circ \blacktriangleright \quad (4.72)$$

Here is the corresponding progression of states (only 1a applies here):

$$\blacktriangleright \circ \circ \square \square \square \square \square \square \square \square \square \quad (4.73)$$

$$\otimes \blacktriangleright \circ \square \square \square \square \square \square \square \square \square \quad (4.74)$$

$$\otimes \otimes \blacktriangleright \square \square \square \square \square \square \square \square \square \quad (4.75)$$

After these $n - 1$ applications of rule 1a, the internal state of the n qubits holds the state of the quantum circuit after the first round of gates.

The second set of rules involves the active spot reaching the front of the chain of qubits. Rule 2a applies to particle pairs $(kn - 1, kn)$ for integer k , while rule 2b applies everywhere else. The reason for this is to ensure that when working out the progression of states backwards, there is always only one of the rules that applies in a given situation.

$$2a : \blacktriangleright \square \longrightarrow \otimes \circ \text{ on particles } (kn - 1, kn), \quad (4.76)$$

$$2b : \blacktriangleright \square \longrightarrow \otimes \circ \text{ everywhere else,} \quad (4.77)$$

Thus, after applying rule 2a on particles (1, 2), I get

$$\otimes \otimes \blacktriangleright \square \square \square \square \square \square \square \square \square \quad (4.78)$$

$$\otimes \otimes \otimes \circ \square \square \square \square \square \square \square \square \square \quad (4.79)$$

The next three rules facilitate the sending of a message back to the left end of the chain.

$$3 : \circ \square \longrightarrow \triangleleft \square \quad (4.80)$$

$$4 : \otimes \triangleleft \longrightarrow \triangleleft \circ \quad (4.81)$$

$$5 : \boxtimes \triangleleft \longrightarrow \boxtimes \circ \quad (4.82)$$

Also note that rule 3 is simply $\circ \longrightarrow \triangleleft$ for the rightmost particle of the line and rule 5 is $\triangleleft \longrightarrow \circ$ for the leftmost particle of the line. This produces the state progression

$$\otimes \otimes \otimes \circ \square \square \square \square \square \square \square \square \square \quad (4.83)$$

$$\otimes \otimes \otimes \triangleleft \square \square \square \square \square \square \square \square \square \quad (4.84)$$

$$\otimes \otimes \triangleleft \circ \square \square \square \square \square \square \square \square \square \quad (4.85)$$

$$\otimes \triangleleft \circ \circ \square \square \square \square \square \square \square \square \square \quad (4.86)$$

$$\triangleleft \circ \circ \circ \square \square \square \square \square \square \square \square \square \quad (4.87)$$

$$\circ \circ \circ \circ \square \square \square \square \square \square \square \square \square \quad (4.88)$$

Observe now that the qubits have moved one step to the right from where they were in (4.73). Finally, the last set of rules activates the leftmost of the qubits. Depending on the position this qubit is at, it will become either \blacktriangleright or \blacktriangleright .

$$6a : \circ \circ \longrightarrow \boxtimes \blacktriangleright \text{ for particles } (kn, kn + 1), \quad (4.89)$$

$$6b : \circ \circ \longrightarrow \boxtimes \blacktriangleright \quad (4.90)$$

The progression of states thus continues as

$$\circlearrowleft \circ \circ \circ \square \square \square \square \square \square \square \square \quad (4.91)$$

$$\boxtimes \circlearrowright \circ \circ \square \square \square \square \square \square \square \square \quad (4.92)$$

$$\boxtimes \boxtimes \circlearrowright \circ \square \square \square \square \square \square \square \square \quad (4.93)$$

$$\boxtimes \boxtimes \boxtimes \circlearrowright \square \square \square \square \square \square \square \square \quad (4.94)$$

$$\boxtimes \boxtimes \boxtimes \boxtimes \circlearrowleft \square \square \square \square \square \square \square \square \quad (4.95)$$

$$\boxtimes \boxtimes \boxtimes \boxtimes \triangleleft \square \square \square \square \square \square \square \square \quad (4.96)$$

$$\boxtimes \boxtimes \boxtimes \triangleleft \circ \square \square \square \square \square \square \square \square \quad (4.97)$$

$$\boxtimes \boxtimes \triangleleft \circ \circ \square \square \square \square \square \square \square \square \quad (4.98)$$

$$\boxtimes \triangleleft \circ \circ \circ \square \square \square \square \square \square \square \square \quad (4.99)$$

$$\boxtimes \circlearrowleft \circ \circ \circ \square \square \square \square \square \square \square \square \quad (4.100)$$

$$\boxtimes \boxtimes \circlearrowright \circ \circ \square \square \square \square \square \square \square \square \quad (4.101)$$

and

$$\boxtimes \boxtimes \circlearrowright \circ \circ \square \square \square \square \square \square \square \square \quad (4.102)$$

$$\boxtimes \boxtimes \boxtimes \circlearrowright \circ \square \square \square \square \square \square \square \square \quad (4.103)$$

$$\boxtimes \boxtimes \boxtimes \boxtimes \circlearrowright \square \square \square \square \square \square \square \square \quad (4.104)$$

$$\boxtimes \boxtimes \boxtimes \boxtimes \boxtimes \circlearrowleft \square \square \square \square \square \square \square \square \quad (4.105)$$

$$\boxtimes \boxtimes \boxtimes \boxtimes \boxtimes \triangleleft \square \square \square \square \square \square \square \square \quad (4.106)$$

$$\boxtimes \boxtimes \boxtimes \boxtimes \triangleleft \circ \square \square \square \square \square \square \square \square \quad (4.107)$$

$$\boxtimes \boxtimes \boxtimes \triangleleft \circ \circ \square \square \square \square \square \square \square \square \quad (4.108)$$

$$\boxtimes \boxtimes \triangleleft \circ \circ \circ \square \square \square \square \square \square \square \square \quad (4.109)$$

$$\boxtimes \boxtimes \circlearrowleft \circ \circ \circ \square \square \square \square \square \square \square \square \quad (4.110)$$

$$\boxtimes \boxtimes \boxtimes \circlearrowright \circ \circ \square \square \square \square \square \square \square \square \quad (4.111)$$

Note that in steps (4.92) and (4.101) the last qubit is marked as a ‘messenger’ \circlearrowright by rule 6b. This rule is used at the left end of the qubit sequence until the qubits have moved n steps to the right. Then, in (4.111), the leftmost qubit is at position $kn + 1$, which makes it become \circlearrowright by rule 6a. The second round ($k = 2$) of gate applications now begins.

The Hamiltonian

As in Section 4.2, the Hamiltonian is an implementation of (4.6).

The part checking the propagation of the computation is constructed from the above rules as

$$H = R_{(1)}^{(5)} + R_{(Kn)}^{(3)} + \sum_{s=1}^{nK-1} \sum_{r=1a}^{6b} R_{(s,s+1)}^{(r)}, \quad (4.112)$$

where each term $R^{(r)}$ is a projector acting on one or two neighboring qubits corresponding to the rule r .

As an example, I write out the term $R^{(1a)}$ which checks the proper application of rule 1a.

This projector acts nontrivially on an 8-dimensional subspace of two neighboring particles spanned by $\otimes \odot$ and $\odot \otimes$. Using the notation of (4.68), this subspace can be thought of as

$$(q_1 \otimes l_1) \otimes (q_2 \otimes l_2), \quad (4.113)$$

where q_1 and q_2 are states of the two qubits, while their labels (l_1, l_2) are in the state $(\blacktriangleright, \cdot)$ or $(\times, \blacktriangleright)$. The corresponding projector (only for clock particles in position $kn+m, kn+m+1$) is

$$R_{(kn+m, kn+m+1)}^{(1a)} = \frac{1}{2} \mathbb{I}_{q_1, q_2} \otimes (|\blacktriangleright \cdot\rangle \langle \blacktriangleright \cdot| + |\times \blacktriangleright\rangle \langle \times \blacktriangleright|)_{l_1 l_2} \quad (4.114)$$

$$- \frac{1}{2} (U_{k,m})_{q_1, q_2} \otimes |\times \blacktriangleright\rangle \langle \blacktriangleright \cdot|_{l_1 l_2} \quad (4.115)$$

$$- \frac{1}{2} (U_{k,m}^\dagger)_{q_1, q_2} \otimes |\blacktriangleright \cdot\rangle \langle \times \blacktriangleright|_{l_1 l_2}, \quad (4.116)$$

where the gate $U_{k,m}$ is acting on the internal state of the two qubits.

The rest of the projectors are much simpler, therefore I write out only one of them. Rule 2a involves the 4-dimensional subspace of two particles spanned by $\blacktriangleright \boxplus$ and $\boxtimes \circ$.

$$P_{(kn-1, kn)}^{(2a)} = \frac{1}{2} (\mathbb{I} - \sigma_X) \quad (4.117)$$

The low-lying spectrum of H

What remains in the Hamiltonian construction are the legal clock state checking operators. First, analogously to the $|01\rangle \langle 01|$ check operators of the domain wall clock, I need

$$H_{check01} = \sum_{A \in S_a} \sum_{B \in S_d} |A\rangle \langle A| \otimes |B\rangle \langle B| \quad (4.118)$$

$$+ \sum_{C \in S_r} \sum_{A \in S_a} |C\rangle \langle C| \otimes |A\rangle \langle A|, \quad (4.119)$$

where

$$S_a = \{\otimes, \odot, \circ, \triangleleft\}, \quad (4.120)$$

$$S_d = \{\boxtimes, \otimes\}, \quad (4.121)$$

$$S_r = \{\boxplus, \odot\}. \quad (4.122)$$

Next, I need to check that there is an active site in the system at all by

$$H_{check1a0} = \sum_{A \in S_d} \sum_{B \in S_r} |A\rangle \langle A| \otimes |B\rangle \langle B|. \quad (4.123)$$

Next, I need to check that there is not more than one active site in the system by adding a term

$$H_{check11} = \sum_{A \in S_a} \sum_{B \in S_a} |A\rangle \langle A| \otimes |B\rangle \langle B|. \quad (4.124)$$

In fact, this is not sufficient to rule out that the two active sites wouldn't be at two places in the system that are spatially separated. However, Aharonov et al. [5] prove a *clairvoyance lemma*, which says that the expectation value of the complete Hamiltonian in a state with two active states is lower bounded by an inverse polynomial in n . This is because the two active sites have to move in the system, and thus will end up on neighboring sites. Also, one needs to check whether the number of qubits in the system is right. Whether there are not too few of them can be checked by

$$H_{few} = |\square\rangle\langle\square|_{c_n} \quad (4.125)$$

on the n -th particle from the left. On the other hand, whether there are not too many of them can be checked by

$$H_{many} = |\blacktriangleright\odot\rangle\langle\blacktriangleright\square|_{c_n, c_{n+1}}. \quad (4.126)$$

Finally, the initial state of the ancillae is checked by

$$H_{init} = \sum_{k=m+1}^n |1\rangle\langle 1|_{\square_{c_k}} \quad (4.127)$$

and the very first state in the sequence is initialized by

$$H_{clockinit} = |\square\rangle\langle\square|_{c_1} + |\blacktriangleright\rangle\langle\blacktriangleright|_{c_1}. \quad (4.128)$$

Altogether, for any quantum circuit U , Aharonov et. al. [5] constructed a Hamiltonian for $d = 12$ dimensional particles on a line with nearest neighbor interactions, whose ground state energy encodes whether the quantum circuit U can output 1 with high probability on some state. Thus, Quantum (12,12)-SAT on a line is QMA_1 complete and Local Hamiltonian for $d = 12$ particles on a line is QMA complete.

4.3.2 Q-(11,11)-SAT on a line is QMA_1 complete.

Here I present a modification of the above clock construction which requires only 11-dimensional qudits. The complexity of Q-(11,11)-SAT on the line is thus still QMA_1 complete (and Local Hamiltonian with $d = 11$ qudits on a line is QMA complete). This result is my own previously unpublished work.

There are 8 dimensions required for the qubits in the previous clock construction. I needed two types of active qubits, one to send a message to the right \blacktriangleright and the other one \blacktriangleleft to facilitate the application of gates. I also required two types of inactive qubits, one in the state 'done' \otimes (to the left of the active spot), and the other one 'ready' \odot (to the right of the active spot). This can be reduced by having only one type of inactive qubit \odot , whose property ('ready' or 'done') is determined by the parity of its position. To compensate for this, I introduce two more non-qubit states \boxplus and \boxminus . The initial state is

$$\boxtimes\boxtimes\boxtimes\boxtimes\boxtimes\boxtimes \underbrace{\blacktriangleright\boxminus\odot\boxminus\odot\boxminus}_{\text{3 qubits separated by } \boxminus} \odot\odot\odot\odot\odot\odot\odot\odot\odot\odot \quad (4.129)$$

where I now follow each qubit by the non-qubit state \boxminus . The progression of the states $|\Psi_t\rangle$ becomes somewhat more involved. To move the qubits two spaces to the right, one goes

changes according to

$$2a : \quad \blacktriangleright \square \longrightarrow \otimes \circlearrowleft \quad (4.153)$$

on particles $(kn, kn + 1)$,

while everywhere else it changes according to

$$2b : \quad \boxplus \square \longrightarrow \otimes \circlearrowleft. \quad (4.154)$$

The reason for this distinction between rules 2a and 2b is that the rules applied backwards also have to be unique. To send the active spot to the left now takes several steps:

$$3a : \quad \circlearrowleft \square \longrightarrow \boxplus \circlearrowleft \quad (4.155)$$

on particles $(2k + 1, 2k + 2)$,

on neighboring particles with (odd,even) locations, while on the other pairs of particles the transition rule is

$$3b : \quad \circlearrowleft \square \longrightarrow \triangleleft \square \quad (4.156)$$

on particles $(2k + 2, 2k + 3)$.

Next, the qubits are pushed to the right as the active symbol \triangleleft moves to the left. Then the active spot bounces back as it reaches the left end of the qubit sequence.

$$4a : \quad \boxplus \triangleleft \longrightarrow \triangleleft \boxplus \quad (4.157)$$

$$4b : \quad \odot \triangleleft \longrightarrow \triangleleft \odot \quad (4.158)$$

$$5a : \quad \boxtimes \triangleleft \longrightarrow \boxtimes \circlearrowleft \quad (4.159)$$

$$5b : \quad \circlearrowleft \boxplus \longrightarrow \boxtimes \circlearrowleft \quad (4.160)$$

Finally, the leftmost qubit becomes active. When the particle pair involved is at position $(kn, kn + 1)$, it means that the qubits have moved $2n$ spots to the right and another round of gate applications can start. This activation proceeds according to the rule

$$6a : \quad \circlearrowleft \odot \longrightarrow \boxtimes \blacktriangleright \quad (4.161)$$

on particles $(2kn - 1, 2kn)$,

while everywhere else the last of the qubits is activated into the ‘messenger’ state \boxplus as

$$6b : \quad \circlearrowleft \odot \longrightarrow \boxtimes \boxplus \quad (4.162)$$

4.4 Quantum 2-SAT in general geometry (for qudits)

4.4.1 The Triangle Clock and Q-(5,3)-SAT

In a recent paper, Eldar and Regev [28], present a novel idea. Their ‘triangle’ clock construction abandons a simple linear progression of states $|\Psi_t\rangle$ encoding the progression of the quantum circuit U . This allows them to prove that certain variant of Quantum 2-SAT for higher spins is QMA₁ complete.

The complexity of classical 2-Satisfiability with higher dimensional particles is known.

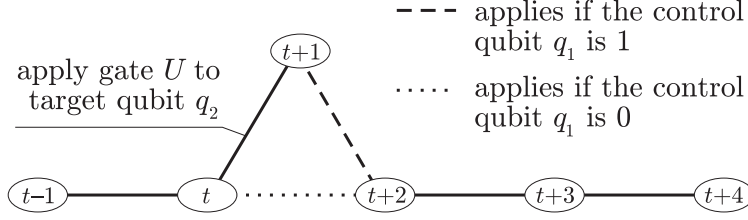


Figure 4-8: The ‘triangle’ clock transition rules allowing to apply a controlled gate to two work qubits using interactions that involve only one work qubit at a time.

Already the simplest variant, (3,2)-SAT, i.e. 2-SAT for a trit and a bit is NP-complete, as shown in Section 1.1.3. On the other hand, the complexity of Q-SAT for higher spins is an open question. In Section 4.2 I proved that Quantum (3,2,2)-SAT, i.e. Quantum 3-SAT for a qutrit and two qubits, is QMA₁ complete. Here I focus on the higher spin version of Quantum 2-SAT. Eldar and Regev’s result which I now review is that Quantum (5,3)-SAT is QMA₁-complete. Here (5,3) means that each projector in their Hamiltonian involves one cinquit (a particle with dimension $d = 5$) and one qutrit ($d = 3$).

The proof of QMA₁-hardness relies on the ‘triangle’ clock construction depicted in Figure 4-8. The state

$$|\Psi_t\rangle = |\psi_t\rangle \otimes |t\rangle_c \quad (4.163)$$

has an allowed transition to the state

$$|\Psi_{t+1}\rangle = \left(\sigma_x^{(q_2)} |\psi_t\rangle \right) \otimes |t+1\rangle_c, \quad (4.164)$$

which has the operator σ_x applied to a target work qubit q_2 . Next, the transition from the state $|\Psi_t\rangle$ to the state $|\Psi_{t+2}\rangle$ is conditioned on a control (work) qubit q_1 to be in the state $|0\rangle$. On the other hand, if the control qubit q_1 is in the state $|1\rangle$, the transition $|\Psi_{t+1}\rangle \rightarrow |\Psi_{t+2}\rangle$ applies. In this fashion, a two-qubit controlled gate, in this case a CNOT, is applied to two work qubits, while interacting with only one of them at a time. The corresponding ‘triangle’ Hamiltonian consists of three terms:

$$\begin{aligned} H_\Delta &= \frac{1}{2} \left[\mathbb{I} \otimes (P_t + P_{t+1}) - \sigma_x^{(q_2)} \otimes (X_{t+1,t} + X_{t+1,t}^\dagger) \right] \\ &+ |0\rangle\langle 0|_{q_1} \otimes \frac{1}{2} \left[P_t + P_{t+2} - (X_{t+2,t} + X_{t+2,t}^\dagger) \right] \\ &+ |1\rangle\langle 1|_{q_1} \otimes \frac{1}{2} \left[P_{t+1} + P_{t+2} - (X_{t+2,t+1} + X_{t+2,t+1}^\dagger) \right]. \end{aligned} \quad (4.165)$$

The first term checks that for the transition between states t and $t+1$, a σ_x gate is applied to the target qubit q_2 . The second and third terms check the transitions $|\Psi_t\rangle \rightarrow |\Psi_{t+2}\rangle$ and $|\Psi_{t+1}\rangle \rightarrow |\Psi_{t+2}\rangle$, conditioned on the state of the control work qubit q_1 . The reader is welcome to check that the ground state energy of H_Δ is exactly zero, and the states in the ground state subspace of H_Δ are of the form

$$|\Psi_\Delta\rangle = c_\Delta (\cdots + |\Psi_t\rangle + |\Psi_{t+1}\rangle + |\Psi_{t+2}\rangle + \cdots), \quad (4.166)$$

where

$$|\Psi_t\rangle = |\psi\rangle \otimes |t\rangle_c, \quad (4.167)$$

$$|\Psi_{t+1}\rangle = \left(\sigma_x^{(q_2)} |\psi\rangle\right) \otimes |t+1\rangle_c, \quad (4.168)$$

$$|\Psi_{t+2}\rangle = (\text{CNOT}_{q_1, q_2} |\psi\rangle) \otimes |t+2\rangle_c \quad (4.169)$$

and $|\psi\rangle$ is a state of the work register.

The goal of Eldar and Regev in [28] is to construct a Quantum 2-SAT Hamiltonian (for a qutrit and cinquit), and for that one needs to implement the clock in a 2-local fashion. The clock register is made of alternating particles with dimension $d = 3$ and $d = 5$. The 3-dimensional clock particle states are \square , \blacksquare and \boxtimes , while the 5-dimensional clock particle states are \circ , \blacktriangleleft , \blacktriangle , \blacktriangleright and \times . The progression of clock states is similar to the domain wall combined with the pulse clock, as discussed in Section 4.1.3:

$$|1\rangle_c = \blacksquare \circ \square \circ \square \circ \square \circ \quad (4.170)$$

$$|2\rangle_c = \boxtimes \blacktriangleleft \square \circ \square \circ \square \circ \quad (4.171)$$

$$|3\rangle_c = \boxtimes \blacktriangle \square \circ \square \circ \square \circ \quad (4.172)$$

$$|4\rangle_c = \boxtimes \blacktriangleright \square \circ \square \circ \square \circ \quad (4.173)$$

$$|5\rangle_c = \boxtimes \times \blacksquare \circ \square \circ \square \circ \quad (4.174)$$

$$|6\rangle_c = \boxtimes \times \boxtimes \blacktriangleleft \square \circ \square \circ \quad (4.175)$$

$$|7\rangle_c = \boxtimes \times \boxtimes \blacktriangle \square \circ \square \circ \quad (4.176)$$

$$|8\rangle_c = \boxtimes \times \boxtimes \blacktriangleright \square \circ \square \circ \quad (4.177)$$

$$|9\rangle_c = \boxtimes \times \boxtimes \times \blacksquare \circ \square \circ \quad (4.178)$$

$$|10\rangle_c = \boxtimes \times \boxtimes \times \boxtimes \blacktriangleleft \square \circ \quad (4.179)$$

$$|11\rangle_c = \boxtimes \times \boxtimes \times \boxtimes \blacktriangle \square \circ \quad (4.180)$$

$$|12\rangle_c = \boxtimes \times \boxtimes \times \boxtimes \blacktriangleright \square \circ \quad (4.181)$$

$$|13\rangle_c = \boxtimes \times \boxtimes \times \boxtimes \times \blacksquare \circ \quad (4.182)$$

$$|14\rangle_c = \boxtimes \times \boxtimes \times \boxtimes \times \boxtimes \blacktriangleleft \quad (4.183)$$

with a single active particle for a given state $|t\rangle_c$ (denoted by a black filled symbol). The three clock states $|t\rangle_c$, $|t+1\rangle_c$ and $|t+2\rangle_c$ involved in the triangle construction in Figure 4-8 correspond to three consecutive states of the clock register in whose the active site is on one 5-dimensional particle, such as $|2\rangle_c$, $|3\rangle_c$ and $|4\rangle_c$.

The propagation Hamiltonian H_{prop} in 4.6 now consists of four types of terms. For a neighboring clock qutrit and cinquit I have

$$H_{prop(3,5)}^t = \frac{1}{2} \left(|\blacksquare\rangle \langle \blacksquare|_{(3)} \otimes \mathbb{I}_{(5)} + \mathbb{I}_{(3)} \otimes |\blacktriangleleft\rangle \langle \blacktriangleleft|_{(5)} \right) \quad (4.184)$$

$$- \frac{1}{2} \left(|\boxtimes \blacktriangleleft\rangle \langle \blacksquare \circ|_{(3,5)} + |\blacksquare \circ\rangle \langle \boxtimes \blacktriangleleft|_{(3,5)} \right), \quad (4.185)$$

while for a neighboring clock cinquit and qutrit, the transition is checked by

$$H_{prop(5,3)}^t = \frac{1}{2} \left(|\blacktriangleright\rangle\langle\blacktriangleright|_{(5)} \otimes \mathbb{I}_{(3)} + \mathbb{I}_{(5)} \otimes |\blacksquare\rangle\langle\blacksquare|_{(3)} \right) \quad (4.186)$$

$$- \frac{1}{2} \left(|\times\blacksquare\rangle\langle\blacktriangleright\Box|_{(5,3)} - |\blacktriangleright\Box\rangle\langle\times\blacksquare|_{(5,3)} \right). \quad (4.187)$$

The third type of term in H_{prop} is the above described ‘triangle’ Hamiltonian (4.165), implemented using the following projectors and transition operators:

$$\begin{aligned} P_t &= |\blacktriangleleft\rangle\langle\blacktriangleleft|_{(5)}, \\ P_{t+1} &= |\blacktriangleup\rangle\langle\blacktriangleup|_{(5)}, \\ P_{t+2} &= |\blacktriangleright\rangle\langle\blacktriangleright|_{(5)}, \\ X_{t+1,t} &= |\blacktriangleup\rangle\langle\blacktriangleleft|_{(5)}, \\ X_{t+2,t} &= |\blacktriangleright\rangle\langle\blacktriangleleft|_{(5)}, \\ X_{t+2,t+1} &= |\blacktriangleright\rangle\langle\blacktriangleup|_{(5)}. \end{aligned} \quad (4.188)$$

This term checks the application of a controlled-not (CNOT) gate on two work qubits in a 2-local fashion, as each of the three terms in H_Δ is a projector acting nontrivially on the Hilbert space of one cinquit (5) and one qubit (2).

So far, I have shown how Eldar and Regev check the application of a 2-local controlled gate on two work qubits by using projectors on one cinquit and one qutrit at a time. However, for universality one also needs to apply single qubit unitary gates. The final, fourth type of term in H_{prop} is then

$$H_{prop(single)}^t = \frac{1}{2} \left(\mathbb{I} \otimes (P_t + P_{t+1}) - U \otimes X_{t+1,t} - (U \otimes X_{t+1,t})^\dagger \right) \quad (4.189)$$

$$+ \mathbb{I} \otimes \frac{1}{2} \left(P_{t+1} + P_{t+2} - (X_{t+2,t+1} + X_{t+2,t+1}^\dagger) \right), \quad (4.190)$$

and it is implemented using (4.188). The first line is a projector which checks the application of a unitary gate U on a single work qubit between states $|\Psi_t\rangle$ and $|\Psi_{t+1}\rangle$. The second line is a projector checking a transition from $|\Psi_{t+1}\rangle$ to $|\Psi_{t+2}\rangle$ without changing the work qubits.

Finally, the legal clock subspace checking operators are the analogue of the operators $|01\rangle\langle 01|$ for the domain wall clock. They read

$$H_{clock(3,5)} = |\boxtimes\rangle\langle\boxtimes|\otimes|\circ\rangle\langle\circ| \quad (4.191)$$

$$+ |\blacksquare\rangle\langle\blacksquare|\otimes(\mathbb{I}-|\circ\rangle\langle\circ|) \quad (4.192)$$

$$+ |\Box\rangle\langle\Box|\otimes(\mathbb{I}-|\circ\rangle\langle\circ|) \quad (4.193)$$

for a neighboring qutrit and cinquit, while for a neighboring cinquit and qutrit they are

$$H_{clock(5,3)} = |\times\rangle\langle\times|\otimes|\Box\rangle\langle\Box| \quad (4.194)$$

$$+ (\mathbb{I}-|\times\rangle\langle\times|)\otimes|\blacksquare\rangle\langle\blacksquare| \quad (4.195)$$

$$+ (\mathbb{I}-|\times\rangle\langle\times|)\otimes|\boxtimes\rangle\langle\boxtimes| \quad (4.196)$$

Whether there is an active site in the clock register at all is taken care of by

$$H_{clockinit} = |\Box\rangle\langle\Box|_{c_1} + |\boxtimes\rangle\langle\boxtimes|_{c_L}, \quad (4.197)$$

where c_1 and c_L are the leftmost and the rightmost qutrits in the clock register.

Eldar and Regev then proceed to show that if the circuit U accepts a state $|\psi_{yes}\rangle$, the history state corresponding to $|\psi_{yes}\rangle$ with extra terms for the states above the linear progression of states such as $|\Psi_{t+1}\rangle$ in Figure 4-8 is the zero-energy ground state of H . Also, they show that in the ‘no’ case, the ground state energy of H is bounded from below by an inverse polynomial in n . Thus, Quantum (5,3)-SAT is QMA₁-complete. In the following sections, I build on this result.

4.4.2 A QMA₁-complete 3-local Hamiltonian with restricted terms

Here I implement the 2-local *triangle clock* which originally uses cinquits and qutrits, using qubits. This way I obtain another QMA₁ complete 3-local Hamiltonian construction which does not require any large penalty terms. Moreover, all the terms involved are constructed from a restricted set of operators. This is an unpublished result obtained with Peter Love during the Computational Complexity of Quantum Hamiltonian Systems workshop in Leiden, Netherlands (July 2007).

As in Section 4.2, I use a combined domain wall and pulse clock. First, replace each $d = 3$ particle in the construction of the previous section with two qubits w_1, w_2 as

$$\square \longrightarrow 00, \quad (4.198)$$

$$\blacksquare \longrightarrow 10, \quad (4.199)$$

$$\boxtimes \longrightarrow 11, \quad (4.200)$$

as in the domain wall clock in Section 4.1.3. Next, replace each $d = 5$ particle in the above construction with three qubits p_1, p_2, p_3 of a pulse clock as

$$\circ \longrightarrow 000, \quad (4.201)$$

$$\blacktriangleleft \longrightarrow 100, \quad (4.202)$$

$$\blacktriangle \longrightarrow 010, \quad (4.203)$$

$$\blacktriangleright \longrightarrow 001, \quad (4.204)$$

$$\times \longrightarrow 000. \quad (4.205)$$

Note that the transitions between the active (pulse) states of the $d = 5$ clock involve only two of the three qubits p_i . Therefore, the clock register transitions between $|t\rangle_c$, $|t+1\rangle_c$ and $|t+2\rangle_c$ in the triangle construction of the previous section are now implemented 2-locally as

$$\blacktriangleleft \longrightarrow \blacktriangle : |10\rangle_{p_1, p_2} \longrightarrow |01\rangle_{p_1, w_2}, \quad (4.206)$$

$$\blacktriangleleft \longrightarrow \blacktriangleright : |10\rangle_{p_1, p_3} \longrightarrow |01\rangle_{p_1, p_3}, \quad (4.207)$$

$$\blacktriangle \longrightarrow \blacktriangleright : |10\rangle_{p_2, p_3} \longrightarrow |01\rangle_{p_2, p_3}. \quad (4.208)$$

Adding an interaction with a single work qubit at a time, this becomes 3-local. The propagation Hamiltonian is thus made from projectors on three qubits. This is true also for the transition from an active state of a qutrit clock particle to an active state of the neighboring cinquit clock particle. It is implemented 3-locally as

$$\blacksquare \circ \longrightarrow \boxtimes \blacktriangleleft : |10\rangle_{w_1, w_2} |0\rangle_{p_1} \longrightarrow |11\rangle_{w_1, w_2} |1\rangle_{p_1}, \quad (4.209)$$

while moving the active state from a cinquit clock particle to a qutrit is done as

$$\blacktriangleright \square \longrightarrow \times \blacksquare \quad : \quad |1\rangle_{p_3} |00\rangle_{w_1, w_2} \longrightarrow |0\rangle_{p_3} |10\rangle_{w_1, w_2} . \quad (4.210)$$

The legal clock state progression is now

$$|0\rangle_c = |10 \ 000 \ 00 \ 000 \ 00 \ 000 \ 00\rangle \quad (4.211)$$

$$|1\rangle_c = |11 \ 100 \ 00 \ 000 \ 00 \ 000 \ 00\rangle \quad (4.212)$$

$$|2\rangle_c = |11 \ 010 \ 00 \ 000 \ 00 \ 000 \ 00\rangle \quad (4.213)$$

$$|3\rangle_c = |11 \ 001 \ 00 \ 000 \ 00 \ 000 \ 00\rangle \quad (4.214)$$

$$|4\rangle_c = |11 \ 000 \ 10 \ 000 \ 00 \ 000 \ 00\rangle \quad (4.215)$$

$$|5\rangle_c = |11 \ 000 \ 11 \ 100 \ 00 \ 000 \ 00\rangle \quad (4.216)$$

$$|6\rangle_c = |11 \ 000 \ 11 \ 010 \ 00 \ 000 \ 00\rangle \quad (4.217)$$

$$|7\rangle_c = |11 \ 000 \ 11 \ 001 \ 00 \ 000 \ 00\rangle \quad (4.218)$$

$$|8\rangle_c = |11 \ 000 \ 11 \ 000 \ 10 \ 000 \ 00\rangle \quad (4.219)$$

$$|9\rangle_c = |11 \ 000 \ 11 \ 000 \ 11 \ 100 \ 00\rangle \quad (4.220)$$

$$|10\rangle_c = |11 \ 000 \ 11 \ 000 \ 11 \ 010 \ 00\rangle \quad (4.221)$$

$$|11\rangle_c = |11 \ 000 \ 11 \ 000 \ 11 \ 001 \ 00\rangle \quad (4.222)$$

$$|12\rangle_c = |11 \ 000 \ 11 \ 000 \ 11 \ 000 \ 10\rangle \quad (4.223)$$

However, this is not a proof that Quantum 3-SAT is QMA₁ complete. I have not found a way to make the terms checking the legal clock states 3-local projectors. If I want the term which ensures proper coupling of the domain wall and the pulse clock to be a non-negative operator, it necessarily becomes 5-local. On the other hand, I have found a 5-local positive operator made out of 2-local operators, which does the job. This is reminiscent of the idea presented in Section 4.2.

The operators checking whether a state is in a legal clock subspace are the following. First, for the domain-wall clock particles, the operator

$$H_{clock(wall)} = \sum_{\langle k, l \rangle} |01\rangle \langle 01|_{w_k, w_l} \quad (4.224)$$

where w_k, w_l are consecutive domain wall qubits, checks the domain wall clock is correctly encoded. Next, I check whether the pulse clocks have at most one pulse in each of them by

$$H_{clock(pulse)} = \sum_k \left(|11\rangle \langle 11|_{p_1^{(k)}, p_2^{(k)}} + |11\rangle \langle 11|_{p_1^{(k)}, p_3^{(k)}} + |11\rangle \langle 11|_{p_2^{(k)}, p_3^{(k)}} \right), \quad (4.225)$$

where $p_1^{(k)}, p_2^{(k)}, p_3^{(k)}$ is the k -th pulse clock triplet. Finally, I need to check that the domain wall and the pulse clocks are correctly coupled. In other words, I need to check that the state in which the pulse clock was *not* initialized at the domain wall

$$|\dots 11 \ 000 \ 00 \dots\rangle \quad (4.226)$$

is illegal, while the states

$$|\dots 00\ 000\ 00\dots\rangle \quad (4.227)$$

$$|\dots 10\ 000\ 00\dots\rangle \quad (4.228)$$

$$|\dots 11\ 100\ 00\dots\rangle \quad (4.229)$$

$$|\dots 11\ 010\ 00\dots\rangle \quad (4.230)$$

$$|\dots 11\ 001\ 00\dots\rangle \quad (4.231)$$

$$|\dots 11\ 000\ 10\dots\rangle \quad (4.232)$$

$$|\dots 11\ 000\ 11\dots\rangle \quad (4.233)$$

are perfectly fine. For this task, I now construct two versions of a 5-local positive semidefinite operator $H_{clockpulse}$, the first made from 3-local and the second made from 2-local terms. In both cases, $H_{clockpulse}$ acts nontrivially on the space of three pulse qubits p_i and two neighboring domain wall qubits w_i .

In the first case, $H_{clockpulse}$ uses at most 3-local interactions:

$$\begin{aligned} H_{clockpulse}^{(3-loc)} &= |1\rangle\langle 1|_{w_1} \otimes \left(\sigma_z^{(p_1)} + \sigma_z^{(p_2)} + \sigma_z^{(p_3)} - 1 \right) \otimes |0\rangle\langle 0|_{w_2} \\ &+ \sum_{i \neq j} |11\rangle\langle 11|_{p_i, p_j}. \end{aligned} \quad (4.234)$$

It checks whether there is a single 1 in the pulse clock, when the neighboring domain wall qubits are 1 and 0.

The second variant uses only 2-local interactions:

$$\begin{aligned} H_{clockpulse}^{(2-loc)} &= |1\rangle\langle 1|_{w_1} \otimes \left(\sigma_z^{(p_1)} + \sigma_z^{(p_2)} + \sigma_z^{(p_3)} + \sigma_z^{(w_2)} - 2 \right) \\ &+ \sum_{\langle i, j \rangle} |11\rangle\langle 11|_{p_i, p_j} + \sum_i |1\rangle\langle 1|_{p_i} \otimes |1\rangle\langle 1|_{w_2}, \end{aligned} \quad (4.235)$$

checking that when a domain wall qubit to the left of a pulse clock is 1, exactly one of the following 4 qubits (3 pulse and 1 domain wall) must be 1, as seen in (4.229)-(4.233).

Altogether, I constructed a Quantum 5-SAT Hamiltonian from at most 3-local interaction terms. Each of the $O(n)$ terms in this Hamiltonian has constant norm. Moreover, I retained the properties of the previous construction, as I only implemented the Hamiltonian (4.6) in a different system. Therefore, Quantum 5-SAT made from 3-local interactions of a restricted type is QMA₁ complete. Moreover, 3-local Hamiltonian with constant norm terms of a restricted type is QMA-complete.

4.5 Train Switch Construction: Universal Quantum Computation using a Quantum 3-SAT Hamiltonian

In this Section, building on Regev's triangle construction idea, I prove that one can perform universal quantum computation using a Quantum 3-SAT Hamiltonian H_{3S} (for qubits). It is suitable for an Adiabatic Quantum Algorithm where I slowly change the Hamiltonian from a simple starting Hamiltonian H_B to H_{3S} , or for a Hamiltonian Computer model where I let an easily prepared starting state evolve with H_{3S} for not too long a time.

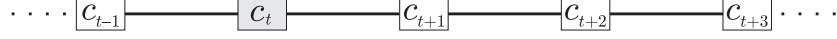


Figure 4-9: The ‘pulse’ clock construction consists of $T + 1$ clock qubits on a line, where only one of them is in the state $|1\rangle$ at a time.

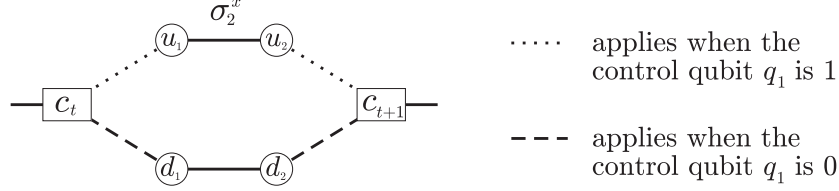


Figure 4-10: The ‘train switch’ clock construction for the application of a CNOT gate using only 3-local terms. There are four extra qubits inserted between c_t and c_{t+1} . The transitions from the original clock qubits to the gadget qubits are controlled by one of the work qubits. During the transition from u_1 to u_2 , the target work qubit is flipped. This gadget checks the application of a CNOT gate using 3-local projectors on qubits.

As I want to perform a quantum computation, as opposed to constructing a Hamiltonian with a unique ground state, I have the advantage that I can choose the initial state of the system. This allows me to use the pulse clock (see Section 4.1.3 and Figure 4-9)

$$|t\rangle_c = |0_{c_0} \dots 0_{c_{t-1}} 1_{c_t} 0_{c_{t+1}} \dots 0_{c_L}\rangle \quad (4.236)$$

as my starting point. The main innovation I add to it is the following *train switch* clock construction. My goal is to use a 3-local projector to check whether a 2-qubit controlled gate CNOT_{q_1, q_2} was properly applied to the work register when transitioning from the state $|\Psi_t\rangle$ to $|\Psi_{t+1}\rangle$. For this, I replace two clock qubits c_t and c_{t+1} with a six-qubit gadget as in Figure 4-10. If the control qubit q_1 is $|1\rangle$, the active site in the clock register follows the upper rail, and takes the lower rail otherwise. Moreover, when the active site in the clock register moves from u_1 and u_2 , the target qubit q_2 is flipped. This effectively applies a CNOT gate between q_1 and q_2 , between the states $|\Psi_t\rangle$ and $|\Psi_{t+1}\rangle$. In detail, the transition checking operators are

$$H_{t, u_1} = \frac{1}{2} \left(\mathbb{I} \otimes (P_t + P_{u_1}) - |1\rangle \langle 1|_{q_1} \otimes (X_{u_1, t} + X_{u_1, t}^\dagger) \right), \quad (4.237)$$

$$H_{u_1, u_2} = \frac{1}{2} \left(\mathbb{I} \otimes (P_{u_1} + P_{u_2}) - \sigma_x^{(q_2)} \otimes (X_{u_2, u_1} + X_{u_2, u_1}^\dagger) \right), \quad (4.238)$$

$$H_{u_2, t+1} = \frac{1}{2} \left(\mathbb{I} \otimes (P_{u_2} + P_{t+1}) - |1\rangle \langle 1|_{q_1} \otimes (X_{t+1, u_2} + X_{t+1, u_2}^\dagger) \right) \quad (4.239)$$

for the upper rail, with the projectors P and time-increase operators X implemented as before for the pulse clock (see Section 4.1.3):

$$P_t = |1\rangle \langle 1|_{c_t}, \quad (4.240)$$

$$X_{b, a} = |10\rangle \langle 01|_{b, a}. \quad (4.241)$$

Analogously, for the lower rail I write

$$H_{t,d_1} = \frac{1}{2} \left(\mathbb{I} \otimes (P_t + P_{u_1}) - |0\rangle \langle 0|_{q_1} \otimes (X_{d_1,t} + X_{d_1,t}^\dagger) \right), \quad (4.242)$$

$$H_{d_1,d_2} = \frac{1}{2} \left(\mathbb{I} \otimes (P_{d_1} + P_{d_2}) - \mathbb{I} \otimes (X_{d_2,d_1} + X_{d_2,d_1}^\dagger) \right), \quad (4.243)$$

$$H_{d_2,t+1} = \frac{1}{2} \left(\mathbb{I} \otimes (P_{d_2} + P_{t+1}) - |0\rangle \langle 0|_{q_1} \otimes (X_{t+1,d_2} + X_{t+1,d_2}^\dagger) \right). \quad (4.244)$$

Let me examine the action of this Hamiltonian on the state

$$|\Psi_t\rangle = |\psi_t\rangle \otimes |t\rangle_c = \left(a |0\rangle_{q_1} |\alpha\rangle_{q_2,\dots} + b |1\rangle_{q_1} |\beta\rangle_{q_2,\dots} \right) \otimes |t\rangle_c. \quad (4.245)$$

When I consider only the forward moving terms in

$$H_{train(t,t+1)} = H_{t,u_1} + H_{u_1,u_2} + H_{u_2,t+1} + H_{t,d_1} + H_{d_1,d_2} + H_{d_2,t+1} \quad (4.246)$$

(forgetting the identity and time-decreasing terms), I obtain the following three states:

$$|\Psi_t^{(u_1,d_1)}\rangle = a |0\rangle_{q_1} |\alpha\rangle_{q_2,\dots} \otimes |d_1\rangle_c + b |1\rangle_{q_1} |\beta\rangle_{q_2,\dots} \otimes |u_1\rangle_c, \quad (4.247)$$

$$|\Psi_t^{(u_2,d_2)}\rangle = a |0\rangle_{q_1} |\alpha\rangle_{q_2,\dots} \otimes |d_2\rangle_c + b |1\rangle_{q_1} |\beta'\rangle_{q_2,\dots} \otimes |u_2\rangle_c, \quad (4.248)$$

$$|\Psi_{t+1}\rangle = \left(a |0\rangle_{q_1} |\alpha\rangle_{q_2,\dots} + b |1\rangle_{q_1} |\beta'\rangle_{q_2,\dots} \right) \otimes |t+1\rangle_c, \quad (4.249)$$

where

$$|\beta'\rangle_{q_2,\dots} = \sigma_x^{(q_2)} |\beta\rangle_{q_2,\dots} \quad (4.250)$$

is the state obtained of the work qubits (besides q_1) after the work qubit q_2 was flipped. Observe that the work qubits of the state $|\Psi_{t+1}\rangle$ have the CNOT gate applied to q_1, q_2 , as I wanted, therefore

$$|\Psi_{t+1}\rangle = (\text{CNOT}_{q_1,q_2} |\psi_t\rangle) \otimes |t+1\rangle_c. \quad (4.251)$$

For a single qubit gate application, the transition checking is much simpler. At places where I do not insert the train switch gadget, I use the usual single qubit gate checking operator

$$H_{single(t)} = \frac{1}{2} \left(\mathbb{I} \otimes (P_t + P_{t+1}) - U \otimes X_{t+1,t} - (U \otimes X_{t+1,t}^\dagger) \right). \quad (4.252)$$

The Hamiltonian

$$H_{3S} = \sum_{t: U_t^{(q_1)}} H_{single(t)} + \sum_{t: U_t^{(q_1,q_2)}} H_{train(t,t+1)}, \quad (4.253)$$

corresponding to a quantum circuit U is a Quantum 3-SAT Hamiltonian, because it is composed of 3-local projectors on qubits.

The progression of states $|\Psi_t\rangle$ with the occasional intermediate states $|\Psi_t^{(u_1,d_1)}\rangle$ and $|\Psi_t^{(u_2,d_2)}\rangle$, when a CNOT gate is on order, has the geometry of a line. Each of the states $|\Psi_t\rangle$ is connected by a transition only to two neighboring states. Let me call \mathcal{H}_{legal} the

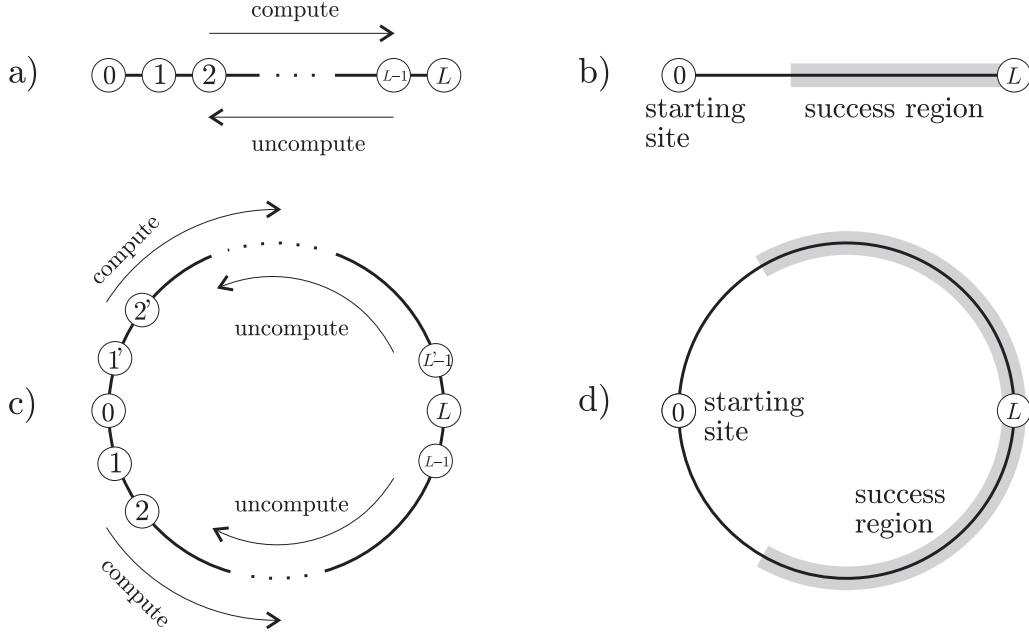


Figure 4-11: a) Computation on a line of states $|\Psi_t\rangle$. b) The line corresponding to a circuit U padded with extra identity gates. The region where the computation is done is marked. c) Computation on a cycle of states $|\Psi_t\rangle$ coming from two copies of the computation on a line. d) The cycle corresponding to a padded circuit U .

subspace of \mathcal{H} spanned by the states $|\Psi_t\rangle$. Because I constructed it so, the Hamiltonian H_{3S} (4.253) does not induce transitions between the subspace \mathcal{H}_{legal} and $\mathcal{H}_{legal}^\perp$. The time evolution of an initial state in \mathcal{H}_{legal} is thus governed solely by the restriction of H_{3S} to \mathcal{H}_{legal} . Moreover, in the basis $|\Psi_t\rangle$ this restriction has the form

$$H_{3S}|_{\mathcal{H}_{legal}} = \begin{bmatrix} \frac{1}{2} & -\frac{1}{2} & & & \\ -\frac{1}{2} & 1 & -\frac{1}{2} & & \\ & -\frac{1}{2} & 1 & -\frac{1}{2} & \\ & & -\frac{1}{2} & \ddots & \ddots \\ & & & \ddots & 1 & -\frac{1}{2} \\ & & & & -\frac{1}{2} & \frac{1}{2} \end{bmatrix}, \quad (4.254)$$

same as Kitaev's propagation Hamiltonian (see Section 1.3.2 and Appendix A). It is a Hamiltonian for a quantum walk on a line of length $L' = \text{poly}(n)$, where n is the number of qubits the quantum circuit U acts on. Here $L' = 1 + L_1 + 3L_2$ where L_1 is the number of single qubit gates in U and L_2 is the number of CNOT's in the circuit. To avoid complications in the analysis of the required running time coming from the endpoints, I can change the 'line' of states $|\Psi_t\rangle$ into a circle of length $2L'$ as in Figure 4-11. First, double the number of qubits in the clock register as

$$|c_0 c_1 \dots c_L\rangle \longrightarrow |c_0 c_1 \dots c_L\rangle \otimes |c'_0 c'_1 \dots c'_L\rangle, \quad (4.255)$$

with the second ‘line’ $|c'_0 c'_1 \dots c'_L\rangle$ having analogous transition rules. Second, identify the endpoint qubits, i.e. $c_0 \equiv c'_0$ and $c_L \equiv c'_L$. This gives the clock register the geometry of a cycle, with a unique state $|0\rangle_c = |1_{c_0} 0 \dots 0\rangle$ denoting time $t = 0$. The active site (spin up) in the clock register can proceed towards c_L both ways, as in Figure 4-11. This new set of states $|\Psi_t\rangle$ with the geometry of a circle then defines the subspace $\mathcal{H}_{legal}^\circ$. The Hamiltonian H_{3S} restricted to this subspace is

$$H_{3S}|_{\mathcal{H}_{legal}^\circ} = \frac{1}{2}(\mathbb{I} - B_\circ), \quad (4.256)$$

where B_\circ is the adjacency matrix for a cycle of length $2L$. The dynamics of this system is the quantum walk on a cycle, and I analyze it in detail in Appendix C. There I show that when starting from the state on site A of the line and letting the system evolve for a time chosen uniformly at random between zero and a number not larger than $O(L \log^2 L)$, the probability to find the state farther than $L/3$ from the starting point is close to $\frac{2}{3}$. This corresponds to finding a state with a spin up on clock qubit $c_{t \geq L/3}$ (lower part of the circle) or $c_{t' \geq L/3}$ (upper part of the circle). In both cases, because the circuit U was padded with identity gates, the state of the work qubits is then $|\psi_L\rangle$ and contains the output of the quantum circuit U . Therefore, one can simulate a quantum computation U with L_1 single qubit gates and L_2 CNOT’s with this system as follows:

1. Construct a clock register with $C = 1 + L_1 + 4L_2$ qubits and transition rules as described by the *train switch* construction.
2. Pad the clock register with extra $2C$ qubits (equivalent to padding the circuit U with identity gates).
3. Double the clock register and make it into a cycle with length $L = 6C$.
4. Initialize the system in the state

$$|\Psi_0\rangle = |00 \dots 0\rangle \otimes |1_{c_0} 0 \dots 0\rangle. \quad (4.257)$$

5. Let it evolve with H_{3S} (4.253) for a time chosen uniformly at random between zero and $\tau \leq O(L \log^2 L)$.
6. Measure the clock register qubits in the success area (farther than $L/3$ from the initial site, see also Figure 4-11). With probability close to $\frac{2}{3}$ you will find the active site (spin up) there. The work register now contains the output of the quantum circuit U . Restart otherwise.

With this I have shown that Quantum 3-SAT Hamiltonians are powerful enough to perform universal quantum computation in the Hamiltonian Computer model based on a quantum walk. When the number of qubits and gates involved in the circuit is L , what I need for my Quantum 3-SAT Hamiltonian computer are thus $O(L)$ qubits, a Hamiltonian with $O(L)$ projector terms with norm $O(1)$, resulting in $\|H\| = O(L)$, and a running time of order $O(L \log^2 L)$. The rescaled required resources (time \times energy) scale as $\tau \cdot \|H\| = O(L^2 \log^2 L)$. The initial part of this Hamiltonian, turning off the initial Hamiltonian pinning the initial state in $|\Psi_0\rangle$ can be turned off and H_{3S}° turned on adiabatically or quickly, as a quick change will not destroy the computation as discussed in Section 2.2.1. Moreover, Seth Lloyd [63] showed that this Hamiltonian Computer model is protected by an energy gap between the

legal subspace $\mathcal{H}_{legal}^\circ$ and the subspace orthogonal to it. Noise in the system will induce transitions away from $\mathcal{H}_{legal}^\circ$, but their energy cost is going to be scaling like $O(L^{-1})$. The computation is thus protected by an energy barrier of the order $O(L^{-1})$.

Chapter 5

Hamiltonian Quantum Cellular Automata

Can universal quantum computation be performed by time evolving a simple system by a translationally invariant, time-independent Hamiltonian? I present ways how to do it in this chapter, largely based on the paper [71], so far posted on the arXiv.

Hamiltonian Quantum Cellular Automata in 1D

Daniel Nagaj, Pawel Wocjan

We construct a simple translationally invariant, nearest-neighbor Hamiltonian on a chain of 10-dimensional qudits that makes it possible to realize universal quantum computing without any external control during the computational process. We only require the ability to prepare an initial computational basis state which encodes both the quantum circuit and its input. The computational process is then carried out by the autonomous Hamiltonian time evolution. After a time polynomially long in the size of the quantum circuit has passed, the result of the computation is obtained with high probability by measuring a few qudits in the computational basis. This result also implies that there cannot exist efficient classical simulation methods for generic translationally invariant nearest-neighbor Hamiltonians on qudit chains, unless quantum computers can be efficiently simulated by classical computers (or, put in complexity theoretic terms, unless $BPP=BQP$).

This chapter is organized as follows. First, in Section 5.1 I introduce classical cellular automata, quantum cellular automata and the Hamiltonian Quantum Cellular Automaton model. Then, in Section 5.2.1 I present a HQCA in 1D with cell size $d = 10$. I give another HQCA construction in 1D with cell size $d = 20$ in Section 5.3.1. Throughout the required runtime analysis, Appendix C on the quantum walk on a line and Appendix D about diffusion of free fermions on a line is referenced.

5.1 Introduction

One of the most important challenges in quantum information science is to identify quantum systems that can be controlled in such a way that they can be used to realize universal quantum computing. The quantum circuit model abstracts from the details of concrete

physical systems and states that the required elementary control operations are: (i) initialization in basis states, (ii) implementation of one and two-qubit gates, and (iii) measurement of single qubits in basis states. Meanwhile, many other models have been proposed such as measurement-based quantum computing [81, 72, 60, 24], adiabatic quantum computing [33, 7], or topological quantum computing [55, 79, 78, 66, 67, 41, 40] that reduce or modify the set of elementary control operations. However, the common principle underlying all these models is that the computation process is always driven by applying a sequence of control operations.

Instead, I consider a model that does not require any control during the computational process. This model consists of a quantum system with a Hamiltonian that makes it possible to realize universal quantum computing by the following protocol: (1) prepare an initial state in the computational basis that encodes both the program and input, (2) let the Hamiltonian time evolution act undisturbed for a sufficiently long time, and (3) measure a small subsystem in the computational basis to obtain the result of the computation with high probability. I refer to this model as a *Hamiltonian quantum computer* and more specifically as a *Hamiltonian quantum cellular automaton* (HQCA) provided that the Hamiltonian acts on qudits that are arranged on some lattice, is invariant with respect to translations along the symmetry axis of the lattice, and contains only finite range interactions. Most natural Hamiltonians have these properties, so it is important to construct HQCA that are as close as possible to natural interactions.

Hamiltonian QCA are related to the more usual discrete-time QCA (for further review of the different types of quantum cellular automata I refer the reader to [76]). However, while the evolution of discrete-time QCA proceeds in discrete update steps (corresponding to tensor products of local unitary operations, see e.g. [80, 87]), the states of Hamiltonian QCA change in a continuous way according to the Schrödinger equation (with a time-independent Hamiltonian). For this reason, Hamiltonian QCA are also called *continuous-time* QCA [76]. Also, in the HQCA model, all the couplings (interactions) are present all the time, while for the discrete-time QCA, the execution of updates on overlapping cells is synchronized by external control. Therefore, the nearest-neighbor interactions of a HQCA have to include a mechanism that ensures that the logical transformations are carried out in the correct order.

The motivation to consider Hamiltonian computers is threefold. First, it is a fundamental question in the thermodynamics of computation how to realize computational processes within a closed physical system. Such Hamiltonian computers were presented and discussed by Benioff [16], Feynman [39], and Margolus [64]. Second, Hamiltonian quantum cellular automata could lead to new ideas for reducing the set of necessary control operations in current proposals for quantum computing by using the inherent computational power of the interactions. HQCA are at one end of the spectrum of possible implementations; more realistic perspectives for quantum computing could arise by combining this model with more conventional models involving external control operations throughout the computation. Third, this model can show the limitations of current and future methods in condensed matter physics for simulating the time evolution of translationally invariant systems. If evolving with a certain Hamiltonian can realize universal quantum computing, then there cannot exist any classical method for efficiently simulating the corresponding time evolution unless classical computers are as powerful as quantum computers (BPP=BQP).

The first theoretical computational models based on a single time-independent Hamiltonian go back to [16, 39, 64]. However, these Hamiltonian computers were not explicitly designed for realizing universal *quantum* computing. Margolus' model [64] has the attrac-

tive feature that it is laid out on a 2-dimensional lattice with translationally invariant, finite-range interactions. (In [18] it was argued that the part of the Hamiltonian responsible for the synchronization in a 1-dimensional variant is close to real interaction in solid states.) However, this scheme does not satisfy the requirement (1) since its initial state has to be prepared in a superposition. Building upon Margolus' idea, a translationally invariant Hamiltonian universal for quantum computing even if the initial state is restricted to be a canonical basis state was given in [48]. This model requires 10-local, finite-range interactions among qubits on a 2-dimensional rectangular lattice wrapped around a cylinder. Subsequently, it was established in [47] that nearest-neighbor interactions among qutrits on a 2-dimensional lattice suffice. However, the Hamiltonian of [47] is translationally invariant only when translated over several lattice sites. A different approach was taken by Vollbrecht and Cirac in [107], showing that one can implement universal quantum computation with a translationally invariant, nearest-neighbor Hamiltonian on a chain of 30-dimensional qudits.

I present two different simplified HQCA constructions on one-dimensional qudit chains. In both models, I think of the qudit chain as composed of two registers, data and program. The work qubits I compute on are located at a static location in the data register. Driven by the autonomous Hamiltonian time evolution, the program sequence contained in the program register moves past the work qubits and the gates are applied to them. After I let the system evolve for a time not larger than a polynomial in the length of the program, I measure one or two qudits in the computational basis to read out the output of the computation with high probability.

My first construction is for a chain of 10-dimensional qudits and is related to the ideas of [107]. The mechanism behind the progress of the program sequence in this particular model can be thought of as the diffusion of a system of free fermions on a line. My second construction uses qudits with dimension $d = 20$ and is inspired by [49], utilizing a technique of [5] to transport the program. Here, the mechanism for the progress of the computation can be thought of as a quantum walk on a line.

5.2 The HQCA in 1D with cell size $d = 10$

I present a simple universal HQCA on a chain of qudits with dimension $d = 10$. First, I encode the progression of a quantum circuit U on N qubits into a set of states $|\varphi_\sigma\rangle$ of a chain of qudits with length $L = \text{poly}(N)$. Second, I give a translationally invariant nearest-neighbor Hamiltonian on this chain of qudits, which induces a quantum walk on the set of states $|\varphi_\sigma\rangle$. Finally, using a mapping to a system of free fermions in 1D, I prove that when I initialize the qudit chain in an easily determined computational basis state and let the system evolve for a time $\tau \leq \tau_{10} = O(L \log L)$ chosen uniformly at random, I can read out the result of the quantum circuit U with probability $p_{10} \geq \frac{5}{6} - O\left(\frac{1}{\log L}\right)$ by measuring one of the qudits in the computational basis. I then show that this is enough to ensure universality of my HQCA for the class BQP.

5.2.1 The Construction

The gate set {Toffoli, Hadamard} is universal for quantum computation [88]. With only polynomial overhead, one can simulate a circuit consisting of these gates using only the gate

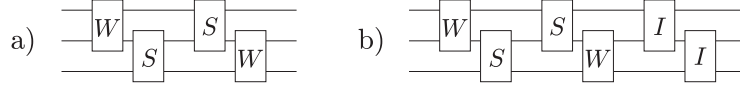


Figure 5-1: a) A quantum circuit consisting of two sequences of gates acting on nearest neighbors. b) The previous circuit with a third sequence of identity gates added.

W (controlled $\frac{\pi}{2}$ rotation about the y -axis)

$$W = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ 0 & 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} \quad (5.1)$$

if it can be applied to any pair of qubits. Let me consider implementing universal quantum computation on a qubit chain using only nearest neighbor gates. Let me also restrict the use of the W gate so that the control qubit has to be to the left of the target qubit. Using only polynomially many additional swap gates S , one can still do universal quantum computation on a qubit chain. Thus given a quantum circuit U' on N' qubits with $\text{poly}(N')$ generic two-qubit gates, I can transform it into a circuit U on a chain of $N = \text{poly}(N')$ qubits with nearest neighbor gates W (with control on the left) and S without loss of universality. I then add identity to my gate set and further transform the circuit U to have the following form (see Figure 5-1). Rewrite the circuit as K sequences of nearest neighbor gates $U_{k,g} \in \{W, S, I\}$, where gate $U_{k,g}$ belongs to the k -th sequence and acts on the pair of qubits $(g, g + 1)$:

$$U = (U_{K,N-1} \dots U_{K,1}) \dots (U_{1,N-1} \dots U_{1,1}). \quad (5.2)$$

I wish to encode the progression of the circuit U into the states of a chain of qudits with dimension 10, with length $L = \text{poly}(N)$. The basis states of each qudit $|q\rangle = |p\rangle \otimes |d\rangle$ are constructed as a tensor product of a 5-dimensional program register $p \in \{ \cdot, \blacktriangleright, W, S, I \}$ and a 2-dimensional data register $d \in \{0, 1\}$. I start by writing the initial product state

$$|\varphi\rangle = \bigotimes_{j=1}^L (|p_j\rangle \otimes |d_j\rangle)_j, \quad (5.3)$$

with p_j and d_j as follows (here I give an example for the circuit in Figure 5-1a):

j	1	\dots	M	\dots	$2M$
p_j	\cdot	\cdot	\blacktriangleright	\cdot	\cdot
d_j	0	0	0	0	0

(5.4)

The qudit chain has length

$$L = 2M = 2KN, \quad (5.5)$$

where K is the number of gate sequences in (5.2). The left half of the top (program) register contains K pointer symbols \blacktriangleright at positions kN for $k = 1 \dots K$ and empty symbols \cdot everywhere else. The right half holds the program in the form

$$I \underbrace{U_{1,1} \dots U_{1,N-1}}_{\text{gate sequence 1}} I \underbrace{U_{2,1} \dots U_{2,N-1}}_{\text{gate sequence 2}} I \dots I \underbrace{U_{K,1} \dots U_{K,N-1}}_{\text{last sequence}}, \quad (5.6)$$

with $U_{k,g} \in \{W, S, I\}$ and each sequence preceded by an identity gate. The bottom (data) register contains N work qubits (labeled w_n in the table) at positions $M+n$ for $n = 1 \dots N$ and qubits in the state $|0\rangle$ everywhere else. I designate w_N as the readout qubit.

I now give two simple rules, many applications of which generate the set of states $\{|\varphi_\sigma\rangle\}$ from the initial state $|\varphi\rangle$. I label the states by σ , the description of the sequence of rules I choose to apply to the initial state. I invite the reader to work out what happens to $|\varphi\rangle$ (5.4) as the rules are applied, noting that there are usually several possible rules that can be applied to a given state. The first rule says that the symbols $A \in \{W, S, I\}$ (from now on I call them gates and think of them as particles) in the program register can move one step to the left, if there is an empty spot there:

$$1 : \begin{array}{|c|c|} \hline \cdot & A \\ \hline \end{array} \longrightarrow \begin{array}{|c|c|} \hline A & \cdot \\ \hline \end{array} \quad (5.7)$$

The second rule concerns what happens when a gate meets a pointer symbol:

$$2 : \begin{array}{|c|c|} \hline \blacktriangleright & A \\ \hline x & y \\ \hline \end{array} \longrightarrow \begin{array}{|c|c|} \hline A & \blacktriangleright \\ \hline A(x, y) & \\ \hline \end{array} \quad (5.8)$$

In this case, the gate moves to the left, the pointer \blacktriangleright gets pushed to the right, and the gate $A \in \{W, S, I\}$ is applied to the qubits in the data register below.

The initial state has K pointers, one for each sequence of gates in the circuit. I constructed the initial state (5.4) in such a way that as a gate $U_{k,g}$ from the k -th gate sequence moves to the left, it meets the k -th (counting from the right) pointer \blacktriangleright exactly above the work qubits to which the gate was intended to be applied ($g, g+1$). However, one also needs to consider what happens when a gate meets a pointer for a different sequence. If this happens over a pair of extra qubits in the state $|0\rangle|0\rangle$, the qubits stay unchanged because the gate is either the controlled gate W , a swap gate or the identity. The second possibility is that a gate meets a pointer above the boundary of the work qubits (i.e. $|0\rangle_M |w_1\rangle_{M+1}$ or $|w_N\rangle_{M+N} |0\rangle_{M+N+1}$). I ensured that it is going to be the identity gate by inserting one in front of each gate sequence in (5.6). This implies that the leftmost (or the rightmost) work qubit again stays unchanged. Therefore, I am sure that the state $|\varphi_\sigma\rangle$ in which the gate particles have all moved to the left half of the chain contains the result of the quantum computation U in the state of the work qubits, while the additional qubits in the data register remain in the state $|0\rangle$.

Let me now allow all the rules to be applied backwards as well, opening the possibility of returning back to the initial state $|\varphi\rangle$ (undoing the computation). Although the number of possible sequences of rule (and backward rule) applications then becomes infinite, the space of states $\{|\varphi_\sigma\rangle\}$ is nevertheless finite-dimensional. In every state $|\varphi_\sigma\rangle$, the M gates $\{W, S, I\}$ in the program register that started at positions $\{M+1, \dots, 2M\}$ occupy some combination $C = \{a_1^{(C)}, \dots, a_M^{(C)}\}$ of the L sites of the chain. Given the rules (5.7)-(5.8), the order of the gates in the program register cannot change, i.e. $a_k^{(C)} < a_m^{(C)}$ for $k < m$. Because the

positions of the pointers \blacktriangleright and the state of the data register are also uniquely determined by $\{a_1^{(C)}, \dots, a_M^{(C)}\}$, I can label the states I constructed by $|\varphi_C\rangle$ with $C = 1 \dots \binom{L}{M}$.

I now give the universal translationally invariant Hamiltonian as a sum of translationally invariant terms

$$H_{10} = - \sum_{j=1}^{L-1} \left(R + R^\dagger \right)_{(j,j+1)}, \quad (5.9)$$

where R corresponds to the rules (5.7)-(5.8) and acts on two neighboring qudits as

$$R = \sum_{A \in \{W, S, I\}} \left[|A \cdot \rangle \langle \cdot A|_{p_1, p_2} \otimes \mathbb{I}_{d_1, d_2} + |A \blacktriangleright \rangle \langle \blacktriangleright A|_{p_1, p_2} \otimes A_{d_1, d_2} \right], \quad (5.10)$$

where p stands for the program register and d for the data register of the respective qudit.

Recall that my model of computation consists of initializing the qudit chain in the state $|\varphi\rangle$ and evolving the system for a time $\tau \leq \tau_{10}$ which I will determine in Section 5.2.2. Finally, I want to show that when I measure the output qubit w_N in the data register, I will read out the result of the quantum computation U with high probability.

5.2.2 Required Time Analysis

The time evolved state $|\varphi(\tau)\rangle$ (obtained from $|\varphi\rangle$ by evolving with H_{10} for time τ) is a superposition of the states $|\varphi_C\rangle$. I can write it as

$$|\varphi(\tau)\rangle = \sum_{C=1}^{\binom{L}{M}} \varphi_C(\tau) |\pi_C\rangle_{\text{program}} \otimes |\theta_C\rangle_{\text{work}} \otimes |\alpha\rangle_{\text{extra}}, \quad (5.11)$$

where $|\pi_C\rangle$ is the state of the program register of the chain in the state $|\varphi_C\rangle$, the corresponding state of the work qubits is $|\theta_C\rangle$ and $|\alpha\rangle$ is the state (all zero) of the extra data qubits. The state $|\theta_C\rangle$ of the work qubits holds the output state of the computation U , if all of the gate “particles” have moved to the left of the work qubits. Let me now choose some number f and pad the qubit chain with $(f-1)M$ empty sites on the left and M sites containing I in the program register on the right:

$$|\varphi^{pad}\rangle = \begin{bmatrix} \cdot \\ 0 \end{bmatrix}^{\otimes (f-1)M} \otimes |\varphi\rangle \otimes \begin{bmatrix} I \\ 0 \end{bmatrix}^{\otimes M}. \quad (5.12)$$

The original chain had length $2M$, so the length of this padded chain is $L = (f+2)M$. The program register of the initial state $|\varphi^{pad}\rangle = |\varphi_{C=1}\rangle$ now has $2M$ gate “particles” $\{W, S, I\}$ at positions $a_m^{(1)} = fM + m$ with $m = 1 \dots 2M$. On this modified chain, every state $|\varphi_C\rangle$ in which the first M gate “particles” are located in the first fM sites of the chain ($a_m^{(C)} \leq fM$ for $m \leq M$) contains the finished computation in the state of its work qubits $|\theta_C\rangle$. Note that for all these states, the state of the work qubits is the same and equal to $|\theta_U\rangle = U|w_1 \dots w_N\rangle$, as it does not change under the extra identity gates I added. I can

now rewrite the time evolved state (5.11) as

$$|\varphi(\tau)\rangle = \left(\sum_{a_M^{(C)} > fM} \varphi_C(\tau) |\varphi_C\rangle \right) + |\pi\rangle_{prog} \otimes |\theta_U\rangle_{work} \otimes |\alpha\rangle_{extra}, \quad (5.13)$$

where the sum in the first term is over the set of positions of the $2M$ gate particles in which the M -th particle is still near the right end of the chain and the computation is thus not finished yet. Meanwhile,

$$|\pi\rangle_{prog} = \sum_{a_M^{(C)} \leq fM} \varphi_C(\tau) |\pi_C\rangle_{prog}. \quad (5.14)$$

is a superposition of the program register states which correspond to an executed computation.

Recall the definition of the quantum complexity class BQP (Definition 6 in Section 1.2.2). There is a verifier circuit U associated with each problem. Denote x a problem instance. As the problems in BQP are decision problems, $x \in (L_{yes} \cup L_{no})$, where L_{yes} is the set of instances with the answer ‘yes’. Now, when $x \in L_{yes}$, the probability of the circuit U outputting ‘yes’ is not smaller than $\frac{2}{3}$. On the other hand, when $x \in L_{no}$, the probability of the circuit outputting ‘yes’ is not greater than $\frac{1}{3}$. Let me assume the worst case for the circuit U , i.e. that the circuit outputs ‘yes’ on a good proof of a ‘yes’ instance with probability $p_U = \frac{2}{3}$. In the language of spins, the circuit U outputs ‘yes’ when I measure spin up on the output qubit. Therefore, the expected value of measuring σ_z on output qubit of the circuit U is bounded from below by

$$\langle \sigma_{w_N}^{(z)} \rangle_{yes}^{circuit} \geq 1 \times p_U + (-1) \times (1 - p_U) = 2p_U - 1 = \frac{1}{3} \quad (5.15)$$

when $x \in L_{yes}$. Analogously, when $x \in L_{no}$, it is bounded from above by

$$\langle \sigma_{w_N}^{(z)} \rangle_{no}^{circuit} \leq -2p_U + 1 = -\frac{1}{3}. \quad (5.16)$$

To solve BQP problems with my automaton, I need to distinguish the ‘yes’ from the ‘no’ cases, i.e. I need to show that the expectation value of measuring $\sigma^{(z)}$ on the output qubit of my automaton at a random time $\tau \leq \tau_{10}$ is greater than zero in the ‘yes’ case, and smaller than zero in the ‘no’ case. First, call p_{10} the probability to find a state where the computation is finished. Let me consider a *yes* instance ($x \in L_{yes}$). Recalling (5.13) and observing that $|\pi\rangle_{prog}$ is orthogonal to the states of the program register in which the computation is not finished, I have

$$\langle \sigma_{w_N}^{(z)} \rangle_{yes} = \langle \varphi(\tau) | \sigma_{w_N}^{(z)} | \varphi(\tau) \rangle = p_{10} \underbrace{\langle \theta_U | \sigma_{w_N}^{(z)} | \theta_U \rangle_{work}}_{\text{output of } U} + (1 - p_{10}) \langle \varphi' | \sigma_{w_N}^{(z)} | \varphi' \rangle \quad (5.17)$$

where $|\varphi'\rangle$ is the normalized first term in (5.13). The first term is the circuit output (5.15), therefore

$$\langle \theta_U | \sigma_{w_N}^{(z)} | \theta_U \rangle_{work} = \langle \sigma_{w_N}^{(z)} \rangle_{yes}^{circuit} \geq 2p_U - 1. \quad (5.18)$$

The second term can be bounded from below (adversarially, i.e. for every time the computation is not finished, the output qubit gives the opposite of the correct answer) by

$$\langle \varphi' | \sigma_{w_N}^{(z)} | \varphi' \rangle \geq -1. \quad (5.19)$$

Putting it together, the expectation value is bounded by

$$\langle \sigma_{w_N}^{(z)} \rangle_{yes} \geq p_{10}(2p_U - 1) - (1 - p_{10}) = 2p_{10}p_U - 1. \quad (5.20)$$

Analogously, for the $x \notin L$ case, I obtain

$$\langle \sigma_{w_N}^{(z)} \rangle_{no} \geq -2p_{10}p_U + 1. \quad (5.21)$$

I will now prove that when I choose the time τ uniformly at random in $(0, \tau_{10})$, with $\tau_{10} = \text{poly}(M)$, the probability of finding a state with the computation executed (with $a_M \leq fM$) is $p_{10} \geq \frac{5}{6} - O\left(\frac{L}{\tau_{10}}\right)$ with $L = (f + 2)M$.

Let me analyze the time evolution of $|\varphi^{pad}\rangle$ (5.12) under H_{10} (5.9). I can restrict the analysis to the program register of the chain, as the content of the data register in the time-evolved state $|\varphi(\tau)\rangle$ is completely determined by the content of the program register. The data register does not hinder the time evolution of the program register in any way. In fact, there exist bases, in which H_{10} is identical (as a matrix) to H_{10} restricted to the program register. Moreover, let me consider a further mapping of the system restricted to the program register to a line of qubits with length $L = (f + 2)M$ as follows. Map the states $\{\blacktriangleright, \bullet\}$ to the state $|0\rangle$, and the states $\{W, S, I\}$ to the state $|1\rangle$. The mapping of H_{10} to this system is a sum of hopping terms

$$H_q = - \sum_{j=1}^{L-1} (|10\rangle \langle 01| + |01\rangle \langle 10|)_{j,j+1}. \quad (5.22)$$

Using the Wigner-Jordan transformation I can define the operators

$$b_j^\dagger = \sigma_1^z \dots \sigma_{j-1}^z \otimes |1\rangle \langle 0|_j \otimes \mathbb{I}_{j+1, \dots, L}, \quad (5.23)$$

$$b_j = \sigma_1^z \dots \sigma_{j-1}^z \otimes |0\rangle \langle 1|_j \otimes \mathbb{I}_{j+1, \dots, L}. \quad (5.24)$$

As b_j^\dagger and b_j have the required properties $\{b_i, b_j^\dagger\} = \delta_{ij}\mathbb{I}$ and $b_j^2 = b_j^{\dagger 2} = 0$, they can be viewed as the creation and annihilation operators for a fermion at site j . Rewriting (5.22) in terms of (5.23)-(5.24), I obtain

$$H_f = - \sum_{j=1}^{L-1} b_j^\dagger b_{j+1} + h.c., \quad (5.25)$$

a Hamiltonian for a system of free fermions in second quantization. Following the mapping, the initial state $|\varphi^{pad}\rangle$ of the qudit chain thus corresponds to the state of the fermionic system $|\Psi\rangle = b_{fM+1}^\dagger \dots b_{(f+2)M}^\dagger |0\rangle$ with $2M$ fermions on the right end of the line (here $|0\rangle$ is the state with no fermions). I now use the following Lemma proved in Appendix D:

Lemma 3. *Consider the state $|\Psi\rangle$ of $2M$ fermions on the right end of a line with $L = (f + 2)M$ sites. Let the system evolve for a time chosen uniformly at random between 0 and*

τ_{10} with the Hamiltonian given by (5.25) and measure the number of fermions in the region $1 \leq x \leq fM$. The probability to measure a number greater than M is $p_{10} \geq \frac{f-2}{f+2} - O\left(\frac{L}{\tau_{10}}\right)$.

Let me choose $f = 22$ and $\tau_{10} = O(L \log L) = O(M \log M)$. Following the mapping I did from my qudit chain backwards, this implies that when I initialize the qudit chain of length $L = 24M$ in $|\varphi^{pad}\rangle$ as in (5.12) and let it evolve with H_{10} (5.9) for a random time $\tau \leq \tau_{10}$, the probability for the chain to be in a state where the gate particles have moved sufficiently to the left for the computation to be done ($a_M \leq fM$) is

$$p_{10} > \frac{5}{6} - O\left(\frac{1}{\log M}\right). \quad (5.26)$$

Therefore, equations (5.20) and (5.21) now read

$$\begin{aligned} \langle \sigma_{w_N}^{(z)} \rangle_{yes} &\geq \frac{1}{9} - O\left(\frac{1}{\log M}\right), \\ \langle \sigma_{w_N}^{(z)} \rangle_{no} &\leq -\frac{1}{9} + O\left(\frac{1}{\log M}\right). \end{aligned} \quad (5.27)$$

Therefore, one can recognize any language in BQP using the HQCA I described above.

As an aside, note that there is a way to determine that I obtained a state in which the computation has been done with certainty (and thus getting rid of the second term in (5.17)). I could have chosen to measure all the program qudits to the right of the first work qubit and check whether all the S and W are gone. This happens with the above probability p_{10} , and the postselected state of the work qubits now surely contains the output of the circuit U . Note also that I can think of the state of all the work qubits as the circuit output, as compared to only the last work qubit. Nevertheless, thinking only about the last work qubit is enough to ensure universality of my HQCA for the class BQP.

5.3 The HQCA in 1D with cell size $d = 20$

I now present the second construction, a HQCA for a chain of 20-dimensional qudits. As in Section 5.2.1, I describe an encoding of the progression of a quantum circuit U into a set of states of a qudit chain. However, the geometry of this set of states $|\psi_t\rangle$ will be now much simpler, as I can label them by a “time” label $t = 1 \dots L = \text{poly}(N)$, thinking of the set of states as a “line”. The Hamiltonian H_{20} I construct induces a quantum walk on this “line” of states. I conclude by proving that when I let the initial state $|\psi_0\rangle$ evolve with H_{20} for a time τ chosen uniformly at random between 0 and $\tau_{20} = O(L \log L)$, I can read out the result of the quantum computation U with probability $p_{20} \geq \frac{5}{6} - O\left(\frac{L}{\tau_{20}}\right)$ by measuring two of the qudits in the computational basis.

5.3.1 The Construction

I encode the progression of a quantum circuit U in the form (5.2) (see also Figure 5-1) into a set of states $|\psi_t\rangle$ of a qudit chain with length $L = (2K - 1)(N + 1) + 2$. As in Section 5.2.1, each qudit consists of a program register and a data register. The data register is

again two-dimensional, but the program register can now be in the following 10 states:

- W, S, I : the program sequence,
- $\textcircled{W}, \textcircled{S}, \textcircled{I}$: marked characters in the program sequence, used to propagate the active spot to the front (left) of the program sequence,
- \blacktriangleright : apply gate symbol,
- \triangleright : shift program forward,
- \circlearrowleft : a turn-around symbol,
- \cdot : empty spot (before/after the program).

Similarly to (5.3) and (5.4), the initial product state $|\psi_0\rangle = \bigotimes_{j=1}^L (|p_j\rangle \otimes |d_j\rangle)_j$ is given by (the following is an example for the circuit in Figure 5-1a)

$$\begin{array}{c|cccccccccccccc}
 j & 1 & \cdots & & & & & & & & & \cdots & L \\
 \hline
 p_j & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & I & W & S & I & I & S & W & \circlearrowleft \\
 d_j & 0 & 1 & 0 & 0 & 0 & 1 & w_1 & w_2 & w_3 & 1 & 0 & 0 & 0 & 1
 \end{array} \quad (5.28)$$

In general, the data register contains N work qubits (labeled w_n in my example) at positions $(K-1)(N+1)+2+n$ for $n = 1 : N$ (counting from the left). Qubit w_N is the designed output qubit for the computation, i.e. once the computation is done, w_N contains the output of U . Next, the data register contains qubits in the state $|1\rangle$ at positions $(k-1)(N+1)+2$ for $k = 1 \dots 2K$ and qubits in the state $|0\rangle$ everywhere else. The 1's serve as sequence boundary markers. The program register has empty symbols \cdot on the left, and then it contains the program in the form

$$I \underbrace{U_{1,1} \dots U_{1,N-1}}_{\text{1st gate sequence}} I I \underbrace{U_{2,1} \dots U_{2,N-1}}_{\text{2nd gate sequence}} I I \dots I I \underbrace{U_{K,1} \dots U_{K,N-1}}_{\text{last gate sequence}}, \quad (5.29)$$

with the program written from left to right. In the example given in (5.28), the first gate sequence (see Figure 5-1a) is WS and the second gate sequence is SW . Finally, the last qudit in the program register is in the state \circlearrowleft , marking an active spot in the computation.

I now give the rules to obtain the sequence of states $|\psi_t\rangle$ from $|\psi_0\rangle$. These rules are constructed so that there is always only one of them that can be applied to a given state $|\psi_t\rangle$, thus giving me a unique state $|\psi_{t+1}\rangle$. (Also, using the rules backwards, one obtains a unique $|\psi_{t-1}\rangle$ from $|\psi_t\rangle$). The first three are

$$\begin{array}{ll}
 1 : & \boxed{A} \boxed{\circlearrowleft} \longrightarrow \boxed{\textcircled{A}} \boxed{\cdot} \\
 2 : & \boxed{A} \boxed{\textcircled{B}} \longrightarrow \boxed{\textcircled{A}} \boxed{B} \\
 3 : & \boxed{\cdot} \boxed{\textcircled{A}} \longrightarrow \boxed{\circlearrowleft} \boxed{A}
 \end{array} \quad (5.30)$$

where A, B stands for either W, S or I . These rules ensure the passing of the active spot from the back end (right side) of the program to the front (left side), without modifying the data register or the order of the gates in the program sequence. Next,

$$\begin{array}{ll}
 4a : & \boxed{\cdot} \boxed{\circlearrowleft} \longrightarrow \boxed{\cdot} \boxed{\blacktriangleright} \\
 & \boxed{1} \longrightarrow \boxed{1} \\
 4b : & \boxed{\cdot} \boxed{\circlearrowleft} \longrightarrow \boxed{\cdot} \boxed{\triangleright} \\
 & \boxed{0} \longrightarrow \boxed{0}
 \end{array} \quad (5.31)$$

After the active spot has moved to the front of the program, there are two possibilities. The turn symbol \circlearrowleft can change to the apply gate symbol \blacktriangleright (rule 4a), or to the shift program symbol \triangleright (rule 4b), depending on whether the data qubit below contains the sequence boundary marker state 1. Afterwards, for the states containing the apply gate symbol \blacktriangleright , I have:

$$5a : \begin{array}{|c|c|} \hline \blacktriangleright & A \\ \hline x & y \\ \hline \end{array} \longrightarrow \begin{array}{|c|c|} \hline A & \blacktriangleright \\ \hline A(x,y) & \\ \hline \end{array} \quad 6a : \begin{array}{|c|c|} \hline \blacktriangleright & \cdot \\ \hline 1 & \\ \hline \end{array} \longrightarrow \begin{array}{|c|c|} \hline \circlearrowleft & \cdot \\ \hline 1 & \\ \hline \end{array} \quad (5.32)$$

When applying rule 5a, the apply gate symbol \blacktriangleright moves to the right, while a gate from the program sequence is applied to the qubits in the data register below. Applying the rule repeatedly, the \blacktriangleright symbol moves to the right end of the program sequence. As an example, I now write out the state $|\psi_{12}\rangle$ that I obtained from the state $|\psi_0\rangle$ applying rules 1, 2 (6 times), 3, 4a and 5a (3 times) from the state $|\psi_0\rangle$.

$$|\psi_{12}\rangle = \left[\begin{array}{cccccc} \cdot & \cdot & \cdot & \cdot & \cdot & I & W & S & \blacktriangleright & I & I & S & W & \cdot \\ 0 & 1 & 0 & 0 & 0 & 1 & | \dots \theta \dots \rangle & 1 & 0 & 0 & 0 & 0 & 1 \end{array} \right], \quad (5.33)$$

where $|\dots \theta \dots\rangle$ stands for the state of the three work qubits after the gates W_{12} and then S_{23} were applied to them. Let me now take a closer look at the marker qubits (all qubits in the data register except for the work qubits w_n) and the application of rule 5a. The marker qubits stay unchanged for all $|\psi_t\rangle$. The gate applied to pairs $|0\rangle|1\rangle$ and $|1\rangle|0\rangle$ of marker qubits or the pairs of qubits $|1\rangle|q_1\rangle$ and $|q_N\rangle|1\rangle$ (the left and right ends of the work qubit sequence) is always I , because of the identity gates I inserted between sequences of gates in the program (5.29). Finally, the qubit pairs $|0\rangle|0\rangle$ between the 1 markers do not change under the swap operation or the W gate (a controlled gate).

After the apply gate \blacktriangleright symbol gets to the end of the sequence, it changes into the turn symbol \circlearrowleft via rule 6a. Note that the boundary markers in the data register are spaced in such a way, that the \blacktriangleright symbol will arrive at the right end of the sequence when the qubit below is in the state 1. Note that at the very end of the line, rule 6a needs to be modified to involve only the two particles directly above each other. Using rule 6a, \blacktriangleright will then change into the turn symbol \circlearrowleft . After applying rules 1, 2 (6 times) and 3, the active spot again moves to the left of the program. Because the \circlearrowleft symbol is now above a 0 marker qubit, rule 4b can be used, and I get a state with the shift program symbol \triangleright . Finally, here are the last two rules:

$$5b : \begin{array}{|c|c|} \hline \triangleright & A \\ \hline & \\ \hline \end{array} \longrightarrow \begin{array}{|c|c|} \hline A & \triangleright \\ \hline & \\ \hline \end{array} \quad 6b : \begin{array}{|c|c|} \hline \triangleright & \cdot \\ \hline 0 & \\ \hline \end{array} \longrightarrow \begin{array}{|c|c|} \hline \circlearrowleft & \cdot \\ \hline 0 & \\ \hline \end{array} \quad (5.34)$$

where again A stands for either W , S or I . Rule 5b makes the program shift to the left while the \triangleright symbol moves to the right. Finally, rule 6b deals with what happens when the \triangleright symbol arrives at the end of the program sequence. Because of the way I constructed the data register in $|\psi_0\rangle$, the data qubit below the \triangleright symbol will then be in the state 0, so that the \triangleright symbol changes to the turn symbol \circlearrowleft . The reason why I need to look at the qubit in the data register below the \circlearrowleft symbol in rules 6a and 6b is that when I apply the rules backwards (making $|\psi_{t-1}\rangle$ from $|\psi_t\rangle$), again only one of them applies for each $|\psi_t\rangle$.

After applying rule 1, 2a (6 times), 3 and 4b, the \triangleright symbol appears again and starts shifting the program further to the left. After several rounds of this, when the program shifts to the left by $N + 1$, rule 4a can be used again (as the \circlearrowleft symbol will be above a

1 marker qubit), and subsequently, the \blacktriangleright symbol facilitates the application of the second sequence of gates to the work qubits.

After many applications of the above rules, I arrive at the state $|\psi_L\rangle$, for which none of the (forward) rules apply.

$$|\psi_L\rangle = \begin{bmatrix} \circlearrowleft & I & W & S & I & I & S & W & \cdot & \cdot & \cdot & \cdot & \cdot \\ 0 & 1 & 0 & 0 & 0 & 1 & |\dots\theta'\dots\rangle & 1 & 0 & 0 & 0 & 1 \end{bmatrix}. \quad (5.35)$$

This is the state in which the program has moved to the left of the qudit chain, and all sequences of gates have been applied to the qubits in the data register. The state $|\dots\theta'\dots\rangle$ is thus the output state of the circuit U and the last of the work qubits (w_N) holds the output of the quantum computation.

Starting from (5.28), I have constructed the set of states $|\psi_t\rangle$ for $t = 0 \dots L$ with $L = O(K^2 N^2) = \text{poly}(N)$. As t grows, these states encode the progress of a quantum circuit U . What is the geometry of this set of states? They are labeled by a discrete label t , with the state $|\psi_t\rangle$ obtainable only from the states $|\psi_{t-1}\rangle$ and $|\psi_{t+1}\rangle$ using the above rules and their backward applications. Therefore, the states $|\psi_t\rangle$ can be thought of as position basis states on a line of length $L + 1$

$$|\psi_t\rangle \leftrightarrow |t\rangle_{\text{line}}, \quad (5.36)$$

where $t = 0 \dots L$.

Let me choose a Hamiltonian H_{20} for this system as a sum of translationally invariant terms:

$$H_{20} = - \sum_{i=1}^{L-1} \sum_{k=1}^{6b} \left(P_k + P_k^\dagger \right)_{(i,i+1)} \quad (5.37)$$

where the terms P_k correspond to the rules 1-6b (5.30),(5.31),(5.32) and (5.34) and act on two neighboring qudits as

$$P_1 = \sum_{A \in \{W, S, I\}} |\textcircled{A} \cdot\rangle \langle A \circlearrowleft|_{p_1, p_2} \otimes \mathbb{I}_{d_1, d_2}, \quad (5.38)$$

$$P_2 = \sum_{A, B \in \{W, S, I\}} |\textcircled{A} B\rangle \langle A \textcircled{B}|_{p_1, p_2} \otimes \mathbb{I}_{d_1, d_2}, \quad (5.39)$$

$$P_3 = \sum_{A \in \{W, S, I\}} |\circlearrowleft A\rangle \langle \cdot \textcircled{A}|_{p_1, p_2} \otimes \mathbb{I}_{d_1, d_2}, \quad (5.40)$$

and

$$P_{4a} = |\bullet \blacktriangleright\rangle \langle \bullet \circlearrowleft|_{p_1, p_2} \otimes \mathbb{I}_{d_1} \otimes |1\rangle \langle 1|_{d_2}, \quad (5.41)$$

$$P_{4b} = |\bullet \triangleright\rangle \langle \bullet \circlearrowleft|_{p_1, p_2} \otimes \mathbb{I}_{d_1} \otimes |0\rangle \langle 0|_{d_2}, \quad (5.42)$$

$$P_{5a} = \sum_{A \in \{W, S, I\}} |A \blacktriangleright\rangle \langle \blacktriangleright A|_{p_1, p_2} \otimes A_{d_1, d_2}, \quad (5.43)$$

$$P_{5b} = \sum_{A \in \{W, S, I\}} |A \triangleright\rangle \langle \triangleright A|_{p_1, p_2} \otimes \mathbb{I}_{d_1, d_2}, \quad (5.44)$$

$$P_{6a} = |\circlearrowleft \bullet\rangle \langle \blacktriangleright \bullet|_{p_1, p_2} \otimes |1\rangle \langle 1|_{d_1} \otimes \mathbb{I}_{d_2}, \quad (5.45)$$

$$P_{6b} = |\circlearrowleft \bullet\rangle \langle \triangleright \bullet|_{p_1, p_2} \otimes |0\rangle \langle 0|_{d_1} \otimes \mathbb{I}_{d_2}. \quad (5.46)$$

When thinking of the set of states $|\psi_t\rangle$ as the set of positions of a particle on a line (5.36), H_{20} becomes

$$H_{line} = - \sum_{t=0}^{L-1} (|t\rangle \langle t+1| + |t+1\rangle \langle t|). \quad (5.47)$$

This is the Hamiltonian of a (continuous-time) quantum walk on a line of length $L + 1$. Therefore, H_{20} induces a quantum walk on the “line” of states $|\psi_t\rangle$.

5.3.2 Required Evolution Time Analysis

After initializing the qudit chain in the state $|\psi_0\rangle$ and evolving with H_{20} for time τ , the final step is to read out the output of the computation. As in Section 5.2.2, I need to ensure that the probability of finding the chain of qudits in a state where the computation was performed completely is high. To raise this probability, I choose to pad the program (K sequences of gates) with another $5K$ sequences of identity gates and redo the construction in the previous section. The length of the qudit chain thus becomes $L = (2(6K) - 1)(N + 1) + 2$. The states $|\psi_{t > L/6}\rangle$ (with L modified) now all contain the result of the quantum circuit U in the readout qubit w_N , as the relevant gates have been applied to the work qubits in those states. Note that as the extra identity gates pass by, the state of the work qubits does not change.

The readout procedure consists of two steps. First, measure the qudit $p_{L-K(N+1)}$ in the program register (the qudit with distance from the right end of the chain equal to the length of the original program). Let me call p_{20} the probability to measure \bullet (which would mean the program has moved to the left of the qudit I just measured). When this happens, I am assured that the work register is in a state in which the computation is done. Second, I measure w_N , the last of the work qubits, and read out the result of the computation U . I now prove that when I choose to measure $p_{L-K(N+1)}$ at a random time $0 \leq \tau \leq \tau_{20}$ with $\tau_{20} = \text{poly}(N)$, the probability p_{20} of obtaining the state \bullet is close to $\frac{5}{6}$.

To simplify the notation, let me label the states $|\psi_t\rangle$ as $|t\rangle$. In this basis, the Hamiltonian (5.37) is the negative of the adjacency matrix of a line graph with $L + 1$ nodes. I now use the following lemma about the quantum walk on a line proved in Appendix C:

Lemma 4. *Consider a continuous time quantum walk on a line of length $L + 1$, where the Hamiltonian is the negative of the adjacency matrix for the line. Let the system evolve for a time τ chosen uniformly at random between 0 and τ_{20} , starting in a position basis state*

$|c\rangle$. The probability to measure a state $|t\rangle$ with $t > L/6$ is then $p_{20} \geq \frac{5}{6} - O\left(\frac{L+1}{\tau_{20}}\right)$.

This implies that when I initialize the qudit chain in the state $|\psi_0\rangle$ (corresponding to the leftmost state on the line $|c\rangle = |1\rangle$) and let it evolve with H for a random time $\tau \leq \tau_{20}$ with $\tau_{20} = O(L \log L)$, the probability to find a state with $t > L/6$ is close to $\frac{5}{6}$. Therefore, when I measure the program qudit $p_{L-K(N+1)}$, I will obtain \cdot with probability close to $\frac{5}{6}$. Finally, when I subsequently measure the work qubit w_N , I will obtain the result of the quantum circuit U .

Note that I can also avoid this postselection procedure and simply measure the output qubit. The analysis of the outcome would then follow what I did above in Section 5.2.2, resulting in (5.27) again, with M replaced by L .

Chapter 6

Conclusions

In this thesis I presented three ways to look at local Hamiltonians in quantum computation. Some of these Hamiltonians are useful for building a quantum computer, some of them can be simulated classically, while the properties of others are hard to determine even on a quantum computer.

First, in Chapter 2 I showed that local structure is a necessity for Adiabatic Quantum Computing to work, and also that a simple 3-local Hamiltonian can be used in the AQC model to simulate any quantum circuit much more effectively than had been known. Second, in Chapter 3 I presented a numerical method to find the ground state of a translationally invariant Hamiltonian with an infinite tree geometry of interactions, investigated two particular AQC Hamiltonian models, and found encouraging results about the phase transitions in these systems. Third, in Chapter 4 I obtained several new results about the complexity of Quantum k -SAT and Local Hamiltonian, implying that finding the ground state properties of rather simple local Hamiltonians is unlikely to be possible even on a quantum computer. Finally, returning to the idea of using local Hamiltonians for building a quantum computer, in Chapter 5 I presented two Hamiltonian Quantum Cellular Automata, showing a way to perform quantum computation in continuous time with a translationally invariant, time-independent Hamiltonian on a line, reminiscent of universal Turing Machines.

Adiabatic Quantum Computing

– how to use it and how to make it fail

My results about how not to design quantum adiabatic algorithms are a response to researchers who made a bad choice of the initial Hamiltonian and then showed that the algorithm doesn't work. I showed that throwing away local structure in the Hamiltonian maps the AQC algorithm into running the unstructured search problem backwards. The running time of any such attempt must then take time of order at least $\sqrt{2^n}$, dashing the hopes of obtaining an exponential speedup. On the other hand, it has not been ruled out that a properly chosen initial Hamiltonian could lead to algorithmic success. A recent numerical study of the NP-complete *Exact Cover* problem up to 128 bits shows promising, inverse-polynomial scaling of the eigenvalue gap governing the required runtime of the AQC algorithm [113]. Besides using AQC for solving locally constrained optimization problems, it remains an interesting alternative to the conventional quantum circuit model, because it is possible to simulate quantum circuits using AQC Hamiltonians. Moreover, I showed that the previously estimated required running time $\tau \propto L^{14}$ (where L is the number of gates in

the circuit and the norm of the Hamiltonian is also proportional to L) for such simulations was grossly overestimated. When using an AQC Hamiltonian made from 3-local projectors I constructed in Section 4.5, the running time needs to scale only as $\tau \cdot \|H\| \propto L^2 \log^2 L$. In fact, I believe the log factors in the required running time can be omitted, as I required them to prove strong convergence of the underlying quantum walk (see Appendix C), while showing weak convergence would be enough¹. Moreover, such computation is inherently protected against decoherence by an energy gap of the order L^{-1} . On the other hand, we do not yet know how to make the AQC model fault tolerant [51]. It will be interesting to see whether the first useful quantum computer will be based on the quantum circuit model, use measurement-based computation [24], utilize topological quantum computing [55], or whether it will be based on “analog” adiabatic evolution [36].

Matrix Product State ansatz based numerics

– investigating spins on an infinite tree

It is hard classically to find the spectral properties of (and to simulate the time evolution with) AQC Hamiltonians for many qubits. This is why I chose to investigate two simple translationally invariant systems. The method I used is based on a local approximate description of a quantum state, the Matrix Product State ansatz. I improved the imaginary time evolution method for finding the ground state energy of a Hamiltonian within the MPS ansatz, and showed how to apply it for a system with an infinite tree geometry of interactions. I then investigated the phase transitions in the transverse field Ising model and the Not 00 model on the Bethe lattice. The phase transitions I found are not first-order. This may have positive implications for the minimum gap scaling of the corresponding AQC Hamiltonians on finite trees. While the two models I considered do not encode computationally interesting optimization problems, they are an important stepping stone for further research. Showing that the systems with infinite tree geometry have similar properties as large finite trees, adding randomness and loops to the underlying graph structure, and finding the properties of more interesting Hamiltonians are now open for research.

Matrix Product State methods were thought to work especially well in 1D. However, after Aharonov et. al. [5] showed that Quantum 2-SAT on a line for particles with dimension $d = 12$ is QMA₁ complete, it became clear that MPS methods are not a panacea for 1D systems, especially for spins with higher dimensions. Moreover, recently Schuch et. al. showed that the ground state of a particular frustration free Hamiltonian in 1D can encode solutions to NP-complete problems, even though the Hamiltonian has a polynomial eigenvalue gap and the ground state is exactly described by a low-dimension MPS [84]. Classical MPS-based methods for finding this ground state are thus bound to fail. How is this possible? The imaginary time evolution method I proposed is not complete. Its weakness must lie in the method for bringing the state back to the MPS ansatz, implying that the procedure can get stuck in local minima, although the required dimensions of the MPS description of the state are small. However, this didn’t happen for the two Hamiltonians I investigated. DMRG (and MPS) based methods remain a very useful tool for spin- $\frac{1}{2}$ systems in 1D and on trees. Nevertheless, new methods such as PEPS are necessary in higher dimensional geometry. At this front, exciting new results have been recently found

¹Here weak convergence means that the time averaged probability distribution $p_\tau(x|0)$, averaged between $\tau = 0$ and $\tau = \tau_0$, converges to the limiting distribution $\pi(x|0)$ as $|\sum_{x \in SR} (\bar{p}_{\tau_0}(x|0) - \pi(x|0))| \leq \epsilon$, where SR is a region covering two thirds of the line. On the other hand, strong convergence means that $\bar{p}_\tau(x|0)$ converges to $\pi(x|0)$ at every point x , i.e. $\sum_{x \in A} |\bar{p}_{\tau_0}(x|0) - \pi(x|0)| \leq \epsilon$.

by Verstraete et al. [102, 68].

Quantum Satisfiability

– an interesting analogue of a classical problem

The complexity of the Local Hamiltonian problem was sorted out by the paper by Oliveira and Terhal [74], who showed that 2-local Hamiltonian on a 2D grid is QMA complete. As k -local Hamiltonian is the quantum analogue of MAX- k -SAT, it is interesting to look at Quantum k -SAT problem, constructed from a sum of k -local projectors, which is the analogue of classical Satisfiability. The solutions to Q- k -SAT are states which are annihilated by all of the terms in the Hamiltonian. This problem seems easier than k -local Hamiltonian, but already for $k = 3$ it contains the NP-complete problem 3-SAT. Is Quantum 3-SAT then QMA₁ complete? This big question remains unanswered here. As showed by Cook and Levin, the answer to a classical 3-SAT instance can encode the evaluation of a classical circuit, because classical information can be freely copied. However, quantum information cannot be cloned, forcing us to store the information about a computation in orthogonal states such as $|\psi_t\rangle \otimes |t\rangle_c$ by adding a clock register to the workspace. So far, no one has been able to show how to encode the clock register in such a way that its transitions would be checkable by 3-local projectors while being coupled to the application of a quantum circuit. So far, my attempt using projectors on one qutrit and two qubits, the QMA₁-complete problem (3,2,2)-SAT comes closest to this goal. Wholly new ideas, such as Eldar and Regev's triangle construction, which I successfully built on in Section 4.5, are required to answer the question of complexity of Q-3-SAT. There I have shown that a Q-3-SAT Hamiltonian can be used as a Hamiltonian computer to perform any quantum computation. However, the question of the complexity of the original Q-3-SAT problem (as a question about a Hamiltonian: is there a zero energy ground state of H) remains open. It might even be possible that a classical verifier circuit exists for it.

In the course of investigating Q-3-SAT, I obtained another result greatly strengthening previous results about the Local Hamiltonian problem. The 3-local Hamiltonian was proved to be QMA complete in [54]. When the Hamiltonian in [54] is rescaled to have norm $O(1)$, the separation in the question about its ground state energy (whether E_0 is lower than a or greater than b) scales like $b - a \propto L^{-10}$. On the other hand, when using my new 3-local Hamiltonian construction of Section 4.2, I showed that even for $b - a \propto L^{-4}$ (and $\|H\| = O(1)$), the 3-local Hamiltonian problem remains QMA-complete. If one could show this for $b - a \propto 1$, it would result in the quantum analogue of the PCP theorem², the basis of many results about hardness of approximation and a great achievement of theoretical computer science [12, 11]. It is unlikely that a method based on perturbation theory gadgets, such as one employed in when transforming 3-local Hamiltonian to 2-local Hamiltonian [53], can lead to a quantum PCP theorem. After rescaling the Hamiltonian back to norm $O(1)$, each round of perturbation theory necessarily shrinks the gap in the spectrum. Although the separation $b - a$ is not the gap $\Delta = E_2 - E_1$ of the Hamiltonian, it is related to it as follows. The low energy states in the 'yes' instance of a Local Hamiltonian problem encode a correct computation with a positive result. The excited states encode either a wrong computation or a negative final result. The eigenvalue separation between these is Δ . On the other hand, in the 'no' instance of Local Hamiltonian, the ground state must again have

² PCP (Probabilistically Checkable Proofs): Any 3-SAT instance can be reformulated so that any purported proof of the new instance is probabilistically checkable by looking at only a constant number of randomly chosen clauses.

the computation history wrong or the final result negative. It is very likely that the ground state energy is going to be smaller than the gap Δ in the ‘yes’ instance. It surely was the case for the Hamiltonians we have encountered so far, where $\Delta \propto L^{-3}$, while the largest separation was $b - a \propto L^{-4}$ (after rescaling the Hamiltonian to $\|H\| = O(1)$). Is it then possible to ever obtain a better scaling of the necessary separation $b - a$? If it is tied to the Δ as I described, the outlook is bleak. If the ground state of a Hamiltonian with gap scaling even like $\Delta \propto L^{-1}$ could encode any quantum computation, it would mean we could use the Hamiltonian as a computer and obtain the result in time shorter than what is physically possible. Thus, when searching for a PCP-like quantum theorem, it seems necessary to find new Hamiltonian encodings which do not tie Δ and $b - a$, or whose underlying idea is not that of Feynman’s Hamiltonian computer.

I also showed that Quantum 2-SAT in 1D for particles with dimension $d = 11$ is QMA₁ complete. It would be interesting to look at the problem from the other end and classify the hardness of Quantum 2-SAT on a line for lower-dimensional particles, perhaps relating it to the *matrix consistency* problem in 1D [62]. My preliminary results show that the ground states of random instances of Quantum 2-SAT in 1D for particles with dimension $d = 4$, with 4 clauses per spin could be hard to approximate classically.

Hamiltonian Quantum Cellular Automata

– universal computing with a simple Hamiltonian

I proposed a model of quantum computation based on a time-independent, translationally invariant Hamiltonian in Chapter 5. The best realization of a HQCA that I found requires 10 dimensional particles placed on a line with nearest neighbor interactions. I initialize the program and data by preparing a computational basis state of the chain, and use a diffusion-like mechanism to run the program through the data. What errors could plague this model? First, it is prone to an Anderson localization problem in the motion of the program, requiring sensitive tuning of the transitions. Other error preventing measures (such as encoding an already fault tolerant circuit into the HQCA) can be taken, but the model cannot be made fault tolerant, unless the number of ‘gate particles’ in the system is kept fixed. This could be done by using actual particles whose number is conserved. However, three kinds of ‘gate particles’ are needed. Unless this is remedied, this model will remain purely theoretical. On the other hand, it shows interesting results about the hardness of simulation of simple translationally invariant Hamiltonians. It would be interesting to find out how much can the dimension of the particles be lowered, keeping the quantum computing universality of the model. One possible new research direction I have embarked upon so far is a programmable scheme in 2D, based on 2-local nearest neighbor interactions of particles with dimensions 3 and 5. For this, I must increase the size of the unit cell (for translational invariance) to about 20 sites.

There is a lot to be learned when looking at the underlying machinery of quantum mechanical systems – local Hamiltonians. They can lead to discoveries of new quantum algorithms such as the NAND tree algorithm [32], classical simulation methods such as iTEBD [106], proofs of complexity results such as the QMA completeness of Quantum (12,12)-SAT on a line [5], or a scalable Hamiltonian computer design based on a feasible physical Hamiltonian such as the superconducting AQC of [52]. I hope this thesis will serve as a pointer to these directions not just for myself, but for the reader as well.

Appendix A

Kitaev's propagation Hamiltonian

In this Appendix, I look at Kitaev's propagation Hamiltonian (1.53) in more detail and prove a lower bound on its second lowest eigenvalue.

Consider the system consisting of two registers, the first of which holds n work qubits and the second of which holds $L + 1$ clock qubits:

$$\mathcal{H} = \mathcal{H}_{work} \otimes \mathcal{H}_{clock}. \quad (\text{A.1})$$

Kitaev's propagation Hamiltonian (1.53) is

$$H_{prop} = \frac{1}{2} \sum_{t=0}^L H_{prop}^t, \quad (\text{A.2})$$

$$H_{prop}^t = \mathbb{I} \otimes (P_t + P_{t+1}) - (X_{t+1,t} \otimes U_t + X_{t+1,t}^\dagger \otimes U_t^\dagger), \quad (\text{A.3})$$

where

$$P_t = |t\rangle \langle t| \quad (\text{A.4})$$

is the projector onto the state $|t\rangle$ of the clock register,

$$X_{t+1,t} = |t+1\rangle \langle t|, \quad (\text{A.5})$$

$$X_{t+1,t}^\dagger = |t\rangle \langle t+1|, \quad (\text{A.6})$$

increase and decrease the state of the clock register by one, and U_t acts on two of the work qubits. To diagonalize this Hamiltonian, it is convenient to make a basis transform

$$|s\rangle \otimes |t\rangle \rightarrow |\Psi_{s,t}\rangle = (U_t U_{t-1} \dots U_2 U_1) |s\rangle \otimes |t\rangle, \quad (\text{A.7})$$

where $|s\rangle$ are computational basis states of the work qubit register. In this new basis, H_{prop} (A.2) has a block diagonal form,

$$H_{prop} = \begin{bmatrix} A & & & \\ & A & & \\ & & \ddots & \\ & & & A \end{bmatrix}, \quad (\text{A.8})$$

with each $(L + 1) \times (L + 1)$ block

$$A = \begin{bmatrix} \frac{1}{2} & -\frac{1}{2} & & & & \\ -\frac{1}{2} & 1 & -\frac{1}{2} & & & \\ & -\frac{1}{2} & 1 & -\frac{1}{2} & & \\ & & -\frac{1}{2} & \ddots & \ddots & \\ & & & \ddots & 1 & -\frac{1}{2} \\ & & & & -\frac{1}{2} & \frac{1}{2} \end{bmatrix} \quad (\text{A.9})$$

corresponding to a basis state $|s\rangle$. Note that the Hamiltonian does not mix sectors with different $|s\rangle$. This matrix can be decomposed as

$$A = \mathbb{I} - \frac{1}{2} (B + P_0 + P_L), \quad (\text{A.10})$$

where B is the adjacency matrix for a line of length $L + 1$ and P_0 and P_L are projectors on its ends. The eigenvectors of A are plain waves [56],

$$|\phi_0\rangle = \frac{1}{\sqrt{L+1}} \sum_{t=0}^L |t\rangle, \quad (\text{A.11})$$

$$|\phi_j\rangle = \sqrt{\frac{2}{L+1}} \sum_{t=0}^L \cos\left(p_j \left(t + \frac{1}{2}\right)\right) |t\rangle, \quad (\text{A.12})$$

with momenta

$$p_j = \frac{\pi j}{L+1}, \quad j = 0, \dots, L, \quad (\text{A.13})$$

corresponding to the eigenvalues

$$\lambda_j = 1 - \cos p_j. \quad (\text{A.14})$$

The smallest eigenvalue is $\lambda_0 = 0$, and the first nonzero eigenvalue is

$$\lambda_1 = 1 - \cos\left(\frac{\pi}{L+1}\right), \quad (\text{A.15})$$

For large L , this eigenvalue is bounded from below by

$$\lambda_1 \geq \frac{c_1}{L^2}, \quad (\text{A.16})$$

for some constant c_1 . The eigenvalue gap between the ground state and the first excited state of H_{prop} is thus bounded from below by

$$\Delta_{H_{prop}} \geq \frac{c_1}{L^2}. \quad (\text{A.17})$$

The dynamics of this system is a quantum walk on a line with the two additional boundary terms.

Appendix B

Dyson Series for an AQC Hamiltonian

Consider a time-dependent AQC Hamiltonian (see Chapter 2)

$$H(s) = (1-s)H_B + sH_P = \underbrace{H_B}_A + \frac{t}{T} \underbrace{(H_P - H_B)}_B, \quad (\text{B.1})$$

where $s = t/T$ and H_B does not commute with H_P . Assume also that all the terms inside H_B commute with each other, and so do the terms within H_P . My goal in this section is to find a formula of the form

$$W_{t_0+\Delta t, t_0} = e^{-ia_2 H_B} e^{-ib_2 H_P} e^{-ic_2 H_B}, \quad (\text{B.2})$$

approximating the time evolution operator $U_{t_0+\Delta t, t_0}$ corresponding to (B.1) up to order $(\Delta t)^2$ and possibly higher orders as well.

The time evolution operator for a time-dependent Hamiltonian $H(t)$ that doesn't commute at different times can be expressed using the Dyson series

$$U_{t_0+\Delta t, t_0} = \mathbb{I} + \sum_{n=1}^{\infty} \left(\frac{-i}{\hbar} \right)^n \int_{t_0}^{t_0+\Delta t} dt_1 \int_{t_0}^{t_1} dt_2 \cdots \int_{t_0}^{t_{n-1}} dt_n H(t_1) H(t_2) \cdots H(t_n). \quad (\text{B.3})$$

In the following, I take $\hbar = 1$.

The zeroth order term in (B.3) for H (B.1) is just $U_0 = \mathbb{I}$. The first order term is

$$\begin{aligned} U_1 &= (-i) \int_{t_0}^{t_0+\Delta t} H(t_1) dt_1 \\ &= \Delta t (-i) \left(A + B \left(\frac{2t_0 + \Delta t}{2T} \right) \right), \end{aligned} \quad (\text{B.4})$$

and the second order term is

$$U_2 = - \int_{t_0}^{t_0+\Delta t} H(t_1) \left(\int_{t_0}^{t_1} H(t_2) dt_2 \right) dt_1 \quad (\text{B.5})$$

$$\begin{aligned} &= -\frac{(\Delta t)^2}{2} A^2 - \frac{(\Delta t)^2}{2T} \left(t_0 + \frac{\Delta t}{3} \right) AB \\ &\quad - \frac{(\Delta t)^2}{T} \left(\frac{t_0}{2} + \frac{(\Delta t)}{3} \right) BA \\ &\quad - \frac{(\Delta t)^2}{2T^2} \left(t_0^2 + \frac{(\Delta t)t_0}{2} + \frac{(\Delta t)^2}{4} \right) BB. \end{aligned} \quad (\text{B.6})$$

Neglecting all terms of order higher than $(\Delta t)^2 \frac{t_0^2}{T^2}$ in $U_0 + U_1 + U_2$, I obtain

$$\begin{aligned} V_2 &= 1 - i(\Delta t)A - i(\Delta t) \left(\frac{t_0}{T} + \frac{\Delta t}{2T} \right) B \\ &\quad - \frac{(\Delta t)^2}{2} AA - \frac{(\Delta t)^2 t_0}{2T} (AB + BA) - \frac{(\Delta t)^2 t_0^2}{2T^2} BB. \end{aligned} \quad (\text{B.7})$$

Rewriting A and B back in terms of H_B and H_P , I obtain

$$\begin{aligned} V_2 &= \mathbb{I} \\ &\quad - i \left(1 - \frac{t_0}{T} - \frac{\Delta t}{2T} \right) (\Delta t) H_B \\ &\quad - i \left(\frac{t_0}{T} + \frac{\Delta t}{2T} \right) (\Delta t) H_P \\ &\quad - \frac{1}{2} \left(1 - \frac{t_0}{T} \right)^2 (\Delta t)^2 H_B H_B \\ &\quad - \frac{t_0}{2T} \left(1 - \frac{t_0}{T} \right) (\Delta t)^2 (H_B H_P + H_P H_B) \\ &\quad - \frac{1}{2} \left(\frac{t_0}{T} \right)^2 (\Delta t)^2 H_P H_P. \end{aligned} \quad (\text{B.8})$$

I now want to find a formula of the form (B.2), approximating (B.8). The simplest version is

$$\begin{aligned} W_1 &= e^{-ia_1 H_B} e^{-ib_1 H_P} \\ &\approx (1 - ia_1 H_B)(1 - ib_1 H_P), \end{aligned} \quad (\text{B.9})$$

which approximates (B.8) to order (Δt) when

$$a_1 = \left(\frac{t_0}{T} \right) \Delta t, \quad (\text{B.10})$$

$$b_1 = \left(1 - \frac{t_0}{T} \right) \Delta t. \quad (\text{B.11})$$

This agrees with the usual first-order Trotter formula.

Next, let me look at

$$\begin{aligned} W_2 &= e^{-ia_2 H_B} e^{-ib_2 H_P} e^{-ic_2 H_B} \\ &\approx \left(1 - ia_2 H_B - \frac{a_2^2}{2} H_B^2\right) \left(1 - ib_2 H_P - \frac{b_2^2}{2} H_P^2\right) \left(1 - ic_2 H_B - \frac{c_2^2}{2} H_B^2\right). \end{aligned} \quad (\text{B.12})$$

When choosing

$$a_2 = c_2 = \left(1 - \frac{t_0}{T} - \frac{\Delta t}{2T}\right) \frac{\Delta t}{2}, \quad (\text{B.13})$$

$$b_2 = \left(\frac{t_0}{T} + \frac{\Delta t}{2T}\right) \Delta t, \quad (\text{B.14})$$

this formula agrees with (B.8) to order $(\Delta t)^2$. It has extra corrections terms, not found in the usual second order Trotter formula derived for a time-independent Hamiltonian H

$$W_2^T = e^{-ia_2^T H_B} e^{-ib_2^T H_P} e^{-ic_2^T H_B}, \quad (\text{B.15})$$

$$a_2^T = c_2^T = \left(1 - \frac{t_0}{T}\right) \frac{\Delta t}{2}, \quad (\text{B.16})$$

$$b_2^T = \left(\frac{t_0}{T}\right) \Delta t. \quad (\text{B.17})$$

Appendix C

Continuous-time Quantum Walks in 1D

First, in Section C.1 I analyze the continuous-time quantum walk on a line and prove two useful lemmas about the mixing of this walk used in Section 5.3.2 and Appendix D. Second, in Section C.2 I analyze the quantum walk on a cycle and prove another mixing lemma used in the proof of universality of the *train switch* Hamiltonian Computer in Section 4.5.

C.1 Quantum Walk on a Line

Consider a continuous time quantum walk on a line of length L , where the Hamiltonian is the negative of the adjacency matrix for the line

$$H_1 = - \sum_{j=1}^{L-1} (|j\rangle \langle j+1| + |j+1\rangle \langle j|). \quad (\text{C.1})$$

The eigenvalues of this Hamiltonian are

$$\lambda_j = -2 \cos \left(\frac{j\pi}{L+1} \right), \quad (\text{C.2})$$

for $j = 1 \dots L$, while the corresponding eigenvectors $|\phi^{(j)}\rangle = \sum_{k=1}^L \phi_k^{(j)} |k\rangle$ have components

$$\phi_k^{(j)} = \sqrt{\frac{2}{L+1}} \sin \left(\frac{jk\pi}{L+1} \right). \quad (\text{C.3})$$

Consider the time evolution of a particular basis state $|c\rangle$. The probability of finding the system in the basis state $|m\rangle$ at some time τ can be found by expanding $|c\rangle$ and $|m\rangle$ in the basis of the eigenvectors (C.3):

$$p_\tau(m|c) = \left| \langle m | e^{-iH\tau} | c \rangle \right|^2 = \sum_{j,k=1}^L e^{-i(\lambda_j - \lambda_k)\tau} \phi_m^{(j)} \phi_c^{(j)*} \phi_m^{(k)*} \phi_c^{(k)}. \quad (\text{C.4})$$

Because the time evolution (according to the Schrödinger equation) is unitary, this probability $p_\tau(m|c)$ does not converge. On the other hand, let me define the time average of

$p_\tau(m|c)$ for time $0 \leq \tau \leq \tau_{20}$ as

$$\bar{p}_{\tau_{20}}(m|c) = \frac{1}{\tau_{20}} \int_0^{\tau_{20}} p_\tau(m|c) d\tau. \quad (\text{C.5})$$

As I will show below in Lemma 4, this average probability distribution does converge to a limiting distribution $\pi(m|c)$, defined as the $\tau_{20} \rightarrow \infty$ limit of the average probability distribution (C.5). All the eigenvalues (C.2) are different, so I can express the limiting distribution as

$$\pi(m|c) = \lim_{\tau_{20} \rightarrow \infty} \bar{p}_{\tau_{20}}(m|c) = \sum_{j=1}^L |\phi_m^{(j)}|^2 |\phi_c^{(j)}|^2, \quad (\text{C.6})$$

which in this case is

$$\pi(m|c) = \frac{2 + \delta_{m,c} + \delta_{m,L+1-c}}{2(L+1)}. \quad (\text{C.7})$$

According to the following lemma, the average probability distribution (C.5) converges to the limiting distribution $\pi(m|c)$.

Lemma 5. *Consider a continuous time quantum walk on a line of length L , where the Hamiltonian is the negative of the adjacency matrix for the line. Let the system evolve for time $\tau \leq \tau_{20}$ chosen uniformly at random, starting in a position basis state $|c\rangle$. The average probability distribution $\bar{p}_{\tau_{20}}(\cdot|c)$ converges to the limiting probability distribution $\pi(\cdot|c)$ as*

$$\sum_{m=1}^L |\bar{p}_{\tau_{20}}(m|c) - \pi(m|c)| \leq O\left(\frac{L}{\tau_{20}}\right). \quad (\text{C.8})$$

Proof. First, recall Lemma 4.3 of [4] for the total variation distance of the probability distribution $\bar{p}_{\tau_{20}}$ from the limiting distribution, saying

$$\sum_m |\bar{p}_{\tau_{20}}(m|c) - \pi(m|c)| \leq \frac{2}{\tau_{20}} \sum_{\lambda_j \neq \lambda_k} \frac{|\phi_c^{(j)}|^2}{|\lambda_j - \lambda_k|}. \quad (\text{C.9})$$

Using (C.2) and (C.3), I can bound the expression on the right of (C.9). When j is close to k , i.e. $|j - k| \leq C_1$, I can write

$$\frac{|\phi_c^{(j)}|^2}{|\lambda_j - \lambda_k|} < 2. \quad (\text{C.10})$$

On the other hand, for $|j - k| > C_1$ I can bound

$$\frac{|\phi_c^{(j)}|^2}{|\lambda_j - \lambda_k|} < \frac{C_2}{L+1}, \quad (\text{C.11})$$

with C_1 and C_2 constants independent of L . Inserting into (C.9), I have

$$\sum_{m=1}^L |\bar{p}_{\tau_{20}}(m|c) - \pi(m|c)| \leq \frac{8C_1L}{\tau_{20}} + \frac{C_2L}{\tau_{20}} = O\left(\frac{L}{\tau_{20}}\right). \quad (\text{C.12})$$

□

Using Lemma 5, I now prove a useful result utilized in the time analysis of the $d = 20$ HQCA in Section (5.3).

Lemma 4. *Consider a continuous time quantum walk on a line of length L , where the Hamiltonian is the negative of the adjacency matrix for the line. Let the system evolve for a time $\tau \leq \tau_{20}$ chosen uniformly at random, starting in a position basis state $|c\rangle$. The probability to measure a state $|t\rangle$ with $t > L/6$ is then bounded from below as $p_{20} \geq \frac{5}{6} - O\left(\frac{L}{\tau_{20}}\right)$.*

Proof. The probability to measure a state $|t\rangle$ with $t > L/6$ at time $\tau \leq \tau_{20}$ chosen uniformly at random is

$$p_{20} = \sum_{m > \frac{L}{6}} \bar{p}_{\tau_{20}}(m|c). \quad (\text{C.13})$$

Starting with (C.8), I have

$$O\left(\frac{L}{\tau_{20}}\right) \geq \sum_{m=1}^L |\bar{p}_{\tau_{20}}(m|c) - \pi(m|c)| \quad (\text{C.14})$$

$$\geq \sum_{m > \frac{L}{6}} |\bar{p}_{\tau_{20}}(m|c) - \pi(m|c)| \quad (\text{C.15})$$

$$\geq \left| \sum_{m > \frac{L}{6}} \bar{p}_{\tau_{20}}(m|c) - \sum_{m > \frac{L}{6}} \pi(m|c) \right| \quad (\text{C.16})$$

$$= \left| p_{20} - \frac{5}{6} + O\left(\frac{1}{L}\right) \right|. \quad (\text{C.17})$$

Therefore, the probability of finding the chain in state $|\psi_{t > L/6}\rangle$ at a random time $\tau \leq \tau_{20}$ is thus bounded from below by

$$p_{20} \geq \frac{5}{6} - O\left(\frac{L}{\tau_{20}}\right). \quad (\text{C.18})$$

□

Also, Lemma 4 can be easily generalized for any desired probability q . Measuring a state which started in a position basis state $|c\rangle$ at a random time $\tau \leq \tau_q$ chosen uniformly at random, the probability to measure a state $|t\rangle$ with $t > (1 - q)L$ is then bounded from below by $p \geq q - O\left(\frac{L}{\tau_q}\right)$.

C.2 Quantum Walk on a Circle

If the geometry of the system is a closed loop of length L instead of a line, the Hamiltonian (C.1) gets an additional wrap-around term.

$$H_{loop} = -(|L\rangle\langle 1| + |1\rangle\langle L|) - \sum_{j=1}^{L-1} (|j\rangle\langle j+1| + |j+1\rangle\langle j|). \quad (\text{C.19})$$

The eigenvalues of this Hamiltonian are

$$\lambda_j = -2 \cos(p_j), \quad (\text{C.20})$$

corresponding to plain waves with momenta

$$p_j = \frac{2\pi j}{L}, \quad (\text{C.21})$$

for $j = 0 \dots L-1$. The corresponding eigenvectors $|\phi^{(j)}\rangle = \sum_{k=1}^L \phi_k^{(j)} |k\rangle$ have components

$$\phi_k^{(0)} = \frac{1}{\sqrt{L}}, \quad (\text{C.22})$$

$$\phi_k^{(j)} = \sqrt{\frac{2}{L}} e^{ip_j k}, \quad j = 1, \dots, L-1, \quad (\text{C.23})$$

These can be combined to make real eigenvectors. For my analysis, it will be enough to consider a line with even length, and only the cosine plain waves:

$$\phi_k^{(0)} = \frac{1}{\sqrt{L}}, \quad (\text{C.24})$$

$$\phi_k^{(j)} = \sqrt{\frac{2}{L}} \cos(ip_j k), \quad j = 1, \dots, \frac{L}{2}. \quad (\text{C.25})$$

The limiting distribution on the cycle when starting from site c is

$$\pi(m|c) = \frac{1}{L} - \frac{2}{L^2}. \quad (\text{C.26})$$

for all points m except for $m = c$ (return back) and $m = L+1-c$ (the point across the cycle), where I have

$$\pi(m|c) = \frac{2}{L} - \frac{2}{L^2}. \quad (\text{C.27})$$

As in Section C.1, I will again utilize Lemma 4.3 of [4] to prove the convergence of the time-averaged probability distribution to this limiting distribution. On the right side of (C.9), I now have

$$|\phi_0^{(j)}|^2 \leq \frac{2}{L}. \quad (\text{C.28})$$

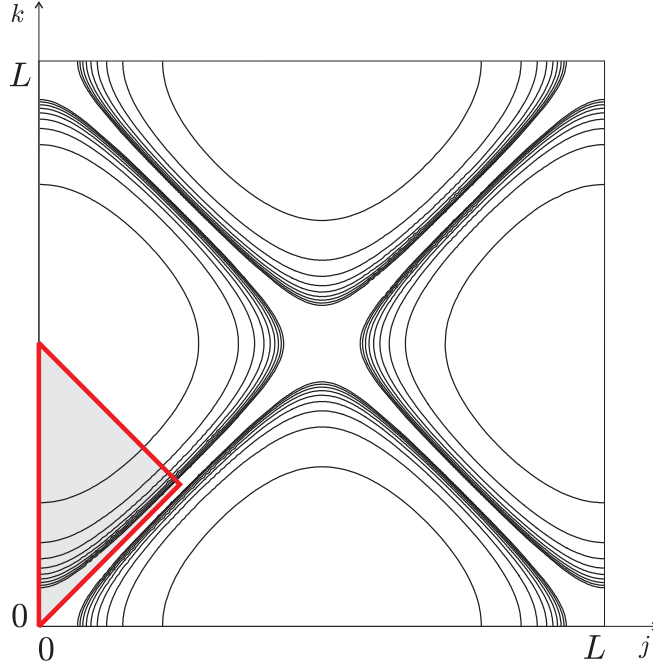


Figure C-1: A contour plot of $|\lambda_j - \lambda_k|^{-1}$ with λ_j from (C.20). The sum in (C.31) is over the marked region.

The sum over the non-equal eigenvalues

$$S_{\circ} = \sum_{\lambda_j \neq \lambda_k} \frac{1}{|\lambda_j - \lambda_k|} \quad (\text{C.29})$$

is now more complicated, because of the degeneracy of the spectrum. I plot $|\lambda_j - \lambda_k|^{-1}$ in Figure C-1. Because of the symmetrical way λ_j (C.20) arise, assuming L is divisible by 4, I can express (C.29) as

$$S_{\circ} = \sum_{\lambda_j \neq \lambda_k} \frac{1}{|\lambda_j - \lambda_k|} \quad (\text{C.30})$$

$$\leq 16 \sum_{j=0}^{L/4} \sum_{k=j+1}^{L/2-j} \frac{1}{|\lambda_j - \lambda_k|} \quad (\text{C.31})$$

$$= \underbrace{16 \sum_{k=1}^{L/2} \frac{1}{|\lambda_0 - \lambda_k|}}_{A_0} + \underbrace{16 \sum_{j=1}^{L/4} \sum_{k=j+1}^{L/2-j} \frac{1}{|\lambda_j - \lambda_k|}}_{A_1}. \quad (\text{C.32})$$

The term A_0 comes from $j = 0$. I can bound the sum by an integral, taking $x = \frac{\delta k}{L}$,

obtaining

$$A_0 = \frac{1}{2} \sum_{k=1}^{L/2} \frac{1}{1 - \cos(2\pi \frac{k}{L})} \leq \frac{L}{2} \int_{\frac{1}{L}}^{\frac{1}{2}} \frac{dx}{1 - \cos(2\pi x)}. \quad (\text{C.33})$$

For $x \in [0, \frac{1}{2}]$, I can bound

$$1 - \cos 2\pi x \geq 2\pi x^2, \quad (\text{C.34})$$

resulting in

$$A_0 \leq \frac{L}{4\pi} \int_{\frac{1}{L}}^{\frac{1}{2}} \frac{dx}{x^2} = \frac{1}{4\pi} (L^2 - 2L) = O(L^2). \quad (\text{C.35})$$

I bound the other term, A_1 , in (C.32) by an integral as well:

$$A_1 = \frac{1}{2} \sum_{j=1}^{L/4} \sum_{k=j+1}^{L/2-j} \frac{1}{\cos(2\pi \frac{k}{L}) - \cos(2\pi \frac{j}{L})} \quad (\text{C.36})$$

$$\leq \frac{L^2}{2} \int_{\frac{1}{L}}^{\frac{1}{4}-\frac{1}{L}} dy \int_{\frac{1}{L}+y}^{\frac{1}{2}-y} \frac{dx}{\cos(2\pi x) - \cos(2\pi y)}. \quad (\text{C.37})$$

Again, I can lower bound the difference in eigenvalues for $y \in [0, \frac{1}{4}]$ and $x \in [y, \frac{1}{2} - y]$ by

$$\cos(2\pi x) - \cos(2\pi y) \geq 2\pi(x^2 - y^2), \quad (\text{C.38})$$

allowing me to write

$$A_1 \leq \frac{L^2}{4\pi} \int_{\frac{1}{L}}^{\frac{1}{4}-\frac{1}{L}} dy \int_{\frac{1}{L}+y}^{\frac{1}{2}-y} \frac{dx}{x^2 - y^2} \quad (\text{C.39})$$

$$= \frac{L^2}{4\pi} \int_{\frac{1}{L}}^{\frac{1}{4}-\frac{1}{L}} dy \left[\frac{\log\left(\frac{x-y}{x+y}\right)}{2y} \right]_{\frac{1}{L}+y}^{\frac{1}{2}-y} \quad (\text{C.40})$$

$$= \frac{L^2}{4\pi} \int_{\frac{1}{L}}^{\frac{1}{4}-\frac{1}{L}} \frac{dy}{2y} \underbrace{[\log(1-4y) + \log(1+2yL)]}_R. \quad (\text{C.41})$$

As y in (C.41) is at most $\frac{1}{4} - \frac{1}{L}$, I can bound R by

$$|R| \leq \log L. \quad (\text{C.42})$$

Finally, this results in

$$A_1 \leq \frac{L^2 \log L}{2\pi} \int_{\frac{1}{L}}^{\frac{1}{4}-\frac{1}{L}} \frac{dy}{2y} = \frac{L^2 \log L}{2\pi} \underbrace{\log\left(\frac{\frac{1}{4}-\frac{1}{L}}{\frac{1}{L}}\right)}_{\leq \log L} \leq O(L^2 \log^2 L). \quad (\text{C.43})$$

Putting (C.35) and (C.43) into (C.32), I obtain

$$S_{\circ} \leq O(L^2 \log^2 L). \quad (\text{C.44})$$

Lemma 4.3 of [4] (see (C.9)) then reads

$$\sum_m |\bar{p}_{\tau_{20}}(m|c) - \pi(m|c)| \leq \frac{2}{\tau_{\circ}} \frac{1}{L} O(L^2 \log^2 L) = O\left(\frac{L \log^2 L}{\tau_{\circ}}\right). \quad (\text{C.45})$$

Thus, for $\tau_{\text{circ}} = \epsilon O(L \log^2 L)$, the time-averaged distribution becomes ϵ -close to the limiting distribution. Using the bound on the total variation distance I just proved, it is straightforward to obtain the following lemma which I use in Section 4.5:

Lemma 6. *Consider a continuous time quantum walk on a cycle of length L (divisible by 4), where the Hamiltonian is the negative of the adjacency matrix for the cycle. Let the system evolve for a time $\tau \leq \tau_{\circ}$ chosen uniformly at random, starting in a position basis state $|0\rangle$. The probability to measure a position state $|t\rangle$ farther than $L/6$ from the starting point (the farther two thirds of the cycle) is then bounded from below as $p_{\circ} \geq \frac{2}{3} - \frac{1}{3L} - O\left(\frac{L \log^2 L}{\tau_{\circ}}\right)$.*

Proof. The proof is analogous to the proof of Lemma 5 in Section C.1. Let me call the farther two thirds of the cycle (see Figure 4-11 in Section 4.5) the success region (SR). When I choose the time $\tau \leq \tau_{\circ}$ uniformly at random, the probability to measure a state $|t\rangle$ with $t \in SR$ is

$$p_{\circ} = \sum_{m \in SR} \bar{p}_{\tau_{\circ}}(m|c), \quad (\text{C.46})$$

Using the bound on the total variation distance (C.45) I just proved and the formulae for the limiting distribution (C.26),(C.27), I have

$$O\left(\frac{L \log^2 L}{\tau_{\circ}}\right) \geq \sum_{m=1}^L |\bar{p}_{\tau_{\circ}}(m|c) - \pi(m|c)| \quad (\text{C.47})$$

$$\geq \sum_{m \in SR} |\bar{p}_{\tau_{\circ}}(m|c) - \pi(m|c)| \quad (\text{C.48})$$

$$\geq \left| \sum_{m \in SR} \bar{p}_{\tau_{\circ}}(m|c) - \sum_{m \in SR} \pi(m|c) \right| \quad (\text{C.49})$$

$$= \left| p_{\circ} - \frac{2}{3} + \frac{1}{3L} + O\left(\frac{1}{L}\right) \right|. \quad (\text{C.50})$$

Therefore, the probability of finding the chain in state $|\psi_{t \in SR}\rangle$ at a random time $\tau \leq \tau_{\circ}$ is thus bounded from below by

$$p_{\circ} \geq \frac{2}{3} - \frac{1}{3L} - O\left(\frac{L \log^2 L}{\tau_{\circ}}\right). \quad (\text{C.51})$$

□

It is thus enough to wait a random time not larger than $O(L \log^2 L)$ to find the state of the system in the success region with probability close to two thirds.

Appendix D

Diffusion of Fermions on a Line

Here I prove Lemma 3, a result about the mixing of a discrete free fermion gas, used in Section 5.2.

Lemma 3. *Consider the state*

$$|\Psi_0\rangle = b_{fM+1}^\dagger b_{fM+2}^\dagger \cdots b_{fM+2M}^\dagger |0\rangle. \quad (\text{D.1})$$

of $2M$ fermions on the right end of a line with $L = (f+2)M$ sites. Let the system evolve for a time chosen uniformly at random between 0 and τ_{10} with the Hamiltonian

$$H_f = - \sum_{j=1}^{L-1} b_j^\dagger b_{j+1} + h.c. \quad (\text{D.2})$$

and measure the number of fermions in the region $1 \leq x \leq fM$. The probability to measure a number greater than M is $p_{10} \geq \frac{f-2}{f+2} - O\left(\frac{L}{\tau_{10}}\right)$.

Proof. Let us start with the outline of the proof. We look at the fermionic system in both first and second quantization to obtain an expression for the time evolution of the creation and annihilation operators in the Heisenberg picture, mapping it to a quantum walk on a line. We then consider the observable X , the number of particles sufficiently far from the right end of the line. We will show that when we choose the time to measure X uniformly at random between 0 and τ_{10} , the expected value we will obtain is approaching a number close to $2M$. To show this, we will express the expected value of X in the time-averaged state of the system using the results from a quantum walk on a line. Finally, because the number of particles in the system is $2M$, we will deduce that the probability to measure a number less than M is then small.

Observe that H_f is the Hamiltonian of a free fermion gas on a line in second quantization (a special case of the XY model). The time evolution of the state $|\Psi_0\rangle$ can be obtained by looking at the problem back in the first quantization, where we write $|\Psi_0\rangle$ as

$$|\Psi_0\rangle = \left[|\phi_{fM+1}\rangle \otimes |\phi_{fM+2}\rangle \otimes \cdots \otimes |\phi_{fM+2M}\rangle \right]^{-}, \quad (\text{D.3})$$

with $|\phi_j\rangle = |j\rangle$ in the position basis and $[\cdot]^{-}$ the standard antisymmetrization operator. We first solve for the time evolution of the corresponding one-particle wavefunction $|\phi_j(\tau)\rangle$

with the Hamiltonian

$$H_1 = - \sum_{j=1}^{L-1} (|j\rangle \langle j+1| + |j+1\rangle \langle j|), \quad (\text{D.4})$$

and then obtain the solutions for the many-particle problem by antisymmetrization as

$$|\Psi(\tau)\rangle = \left[|\phi_{fM+1}(\tau)\rangle \otimes |\phi_{fM+2}(\tau)\rangle \otimes \cdots \otimes |\phi_{fM+2M}(\tau)\rangle \right]^{-}. \quad (\text{D.5})$$

The eigenfunctions of H_1 (quantum walk on a line) are plain waves (as in (C.2) and (C.3)), and the time evolved states $|\phi_j(\tau)\rangle$ thus readily available. Let us define the unitary matrix $u(\tau)$ by

$$|j(\tau)\rangle = \sum_{k=1}^L u_{jk}(\tau) |k\rangle. \quad (\text{D.6})$$

Returning to the second quantized system, the time evolution of the creation and annihilation operators in the Heisenberg picture is then

$$b_j^\dagger(\tau) = \sum_{k=1}^L u_{jk}(\tau) b_k^\dagger, \quad b_j(\tau) = \sum_{k=1}^L u_{jk}^*(\tau) b_k. \quad (\text{D.7})$$

Consider now the observable X , the number of particles in the first fM sites of the line with length $L = (f+2)M$

$$X = \sum_{m=1}^{fM} \hat{n}_m. \quad (\text{D.8})$$

Its expectation value at time τ is

$$E_\tau(X) = \sum_{m=1}^{fM} \langle \Psi(\tau) | \hat{n}_m | \Psi(\tau) \rangle. \quad (\text{D.9})$$

The number operator for site m is $\hat{n}_m = b_m^\dagger b_m$. We can go to the Heisenberg picture and use (D.7) to write

$$\langle \Psi(\tau) | \hat{n}_m | \Psi(\tau) \rangle = \langle \Psi_0 | b_m^\dagger(\tau) b_m(\tau) | \Psi_0 \rangle \quad (\text{D.10})$$

$$= \sum_{c=1}^L \sum_{d=1}^L u_{mc}(\tau) u_{md}^*(\tau) \langle \Psi_0 | b_c^\dagger b_d | \Psi_0 \rangle \quad (\text{D.11})$$

$$= \sum_{c=1}^L |u_{mc}(\tau)|^2 \langle \Psi_0 | b_c^\dagger b_c | \Psi_0 \rangle \quad (\text{D.12})$$

$$= \sum_{c=fM+1}^L \underbrace{|u_{mc}(\tau)|^2}_{p_\tau(m|c)}, \quad (\text{D.13})$$

where each term $|u_{mc}(\tau)|^2 = p_\tau(m|c)$ can be thought of as the probability of finding a particle at site m at time τ when it started from the site c and performed a quantum walk on a line, according to (D.4). Inserting this into (D.9), the expected number of particles not in the rightmost part of the chain at time τ is

$$E_\tau(X) = \sum_{c=fM+1}^L \left(\sum_{m=1}^{fM} p_\tau(m|c) \right). \quad (\text{D.14})$$

Let us now choose the time τ uniformly at random between 0 and τ_{10} . The average value of X (the expectation value in the time-average state) is

$$\bar{E}_{\tau_{10}}(X) = \frac{1}{\tau_{10}} \int_0^{\tau_{10}} E_\tau(X) d\tau. \quad (\text{D.15})$$

For a quantum walk on a line, the time-averaged probability (C.5) of finding a particle that started at position c at final position m converges to the limiting distribution (C.7) according to Lemma 1 (C.8) proven in Appendix C. Using this fact, we can show that the expectation value $\bar{E}_{\tau_{10}}(X)$ in the time-averaged state converges to the limiting expectation value

$$\bar{E}(X) = \sum_{m \leq fM} \sum_{c > fM} \pi(m|c) \quad (\text{D.16})$$

as

$$|\bar{E}_{\tau_{10}}(X) - \bar{E}(X)| \leq O\left(\frac{LM}{\tau_{10}}\right). \quad (\text{D.17})$$

Recalling the limiting probability distribution for a quantum walk on a line of length L (C.7), we have

$$\bar{E}(X) = \sum_{m \leq fM} \sum_{c > fM} \pi(m|c) \quad (\text{D.18})$$

$$= fM \times 2M \times \frac{2}{2(L+1)} + 2M \times \frac{1}{2(L+1)} \quad (\text{D.19})$$

$$= 2M \left(\frac{f}{f+2} \right) + O(1). \quad (\text{D.20})$$

Putting this into (D.17), the average value of X when the time $\tau \leq \tau_{10}$ is chosen uniformly at random is bounded from below as

$$\bar{E}_{\tau_{10}}(X) \geq 2M \left(\frac{f}{f+2} \right) - O\left(\frac{LM}{\tau_{10}}\right). \quad (\text{D.21})$$

We want to find the probability of measuring $X > M$. First, the maximum possible value we could measure at any time is $2M$, the number of particles in the system. Second, the average value $\bar{E}_{\tau_{10}}(X)$ at time τ chosen randomly is close to $2M$. Therefore, the fraction Δ of times at which we measure a number significantly lower than $2M$ must be small. Let us bound Δ in the worst case scenario. This is when each unsuccessful measurement yields

$X = M$, and each successful measurement gives us $2M$. We then have

$$\Delta M + (1 - \Delta)2M \geq \bar{E}_{\tau_{10}}(X), \quad (\text{D.22})$$

$$\Delta \leq \frac{2\bar{E}_{\tau_{10}}(X) - M}{M}. \quad (\text{D.23})$$

Hence we arrive at the desired bound on the probability to measure $X > M$:

$$p_{10} = 1 - \Delta \geq \frac{2M \left(\frac{f}{f+2} \right) - O\left(\frac{LM}{\tau_{10}} \right) - M}{M} = \frac{f-2}{f+2} - O\left(\frac{L}{\tau_{10}} \right). \quad (\text{D.24})$$

□

Appendix E

Transitions in a two level system

In this Appendix I give the analysis of transitions in a simple two-level system referenced in Section 2.3.1. Consider a two level system with Hamiltonian

$$H(s) = E_0(s) |\phi_0(s)\rangle \langle \phi_0(s)| + E_1(s) |\phi_1(s)\rangle \langle \phi_1(s)|, \quad (\text{E.1})$$

which varies smoothly with $s = t/T$. Here $|\phi_0(s)\rangle$ and $|\phi_1(s)\rangle$ are orthonormal for all s . The Schrödinger equation reads

$$i \frac{d}{ds} |\psi\rangle = TH(s) |\psi\rangle. \quad (\text{E.2})$$

The two energy levels in the system are separated by a gap

$$g(s) = E_1(s) - E_0(s), \quad (\text{E.3})$$

which I assume is always larger than 0. Let me introduce θ (with the dimension of energy) as

$$\theta(s) = \int_0^s g(s') ds', \quad (\text{E.4})$$

and let

$$|\psi(s)\rangle = c_0(s) e^{-iT \int_0^s E_0(s') ds'} |\phi_0(s)\rangle + c_1(s) e^{-iT \int_0^s E_1(s') ds'} |\phi_1(s)\rangle. \quad (\text{E.5})$$

I pick the phases of $|\phi_1(s)\rangle$ and $|\phi_0(s)\rangle$ such that $\langle \phi_1(s) | \frac{d}{ds} | \phi_1(s) \rangle = \langle \phi_0(s) | \frac{d}{ds} | \phi_0(s) \rangle = 0$. Plugging E.5 into the Schrödinger equation gives

$$\frac{dc_0}{ds} = c_1 e^{-iT\theta} \left\langle \phi_1 \left| \frac{d}{ds} \right| \phi_0 \right\rangle^*, \quad (\text{E.6})$$

$$\frac{dc_1}{ds} = -c_0 e^{iT\theta} \left\langle \phi_1 \left| \frac{d}{ds} \right| \phi_0 \right\rangle, \quad (\text{E.7})$$

or equivalently,

$$\frac{dc_0}{d\theta} = c_1 e^{-iT\theta} f^*, \quad (\text{E.8})$$

$$\frac{dc_1}{d\theta} = -c_0 e^{iT\theta} f, \quad (\text{E.9})$$

where

$$f(\theta) \equiv \left\langle \phi_1 \left| \frac{d}{d\theta} \right| \phi_0 \right\rangle = -\frac{1}{g} \left\langle \phi_1 \left| \frac{dH}{d\theta} \right| \phi_0 \right\rangle = -\frac{1}{g^2} \left\langle \phi_1 \left| \frac{dH}{ds} \right| \phi_0 \right\rangle. \quad (\text{E.10})$$

Now let $\theta(1) = \bar{\theta}$. I started with $c_1(0) = 0$ and I want to find the transition amplitude at $s = 1$ which is

$$c_1(\bar{\theta}) = - \int_0^{\bar{\theta}} c_0 e^{iT\theta} f d\theta \quad (\text{E.11})$$

$$= \left[-c_0 f \frac{e^{iT\theta}}{iT} \right]_0^{\bar{\theta}} + \frac{1}{iT} \int_0^{\bar{\theta}} e^{iT\theta} \left(\frac{dc_0}{d\theta} f + c_0 \frac{df}{d\theta} \right) d\theta \quad (\text{E.12})$$

$$= \frac{1}{T} \left(\left[ic_0 f e^{iT\theta} \right]_0^{\bar{\theta}} - i \int_0^{\bar{\theta}} \left(c_1 f f^* + e^{iT\theta} c_0 \frac{df}{d\theta} \right) d\theta \right). \quad (\text{E.13})$$

Now $|c_0| \leq 1$ and $|c_1| \leq 1$. As long as the gap does not vanish $|f(\theta)|$ and $\left| \frac{df}{d\theta} \right|$ are bounded, therefore $|c_1(\bar{\theta})| = O\left(\frac{1}{T}\right)$. The probability of transition to the excited state for a two-level system with a nonzero gap is thus

$$|c_1(\bar{\theta})|^2 = O(T^{-2}). \quad (\text{E.14})$$

Bibliography

- [1] Scott Aaronson and Greg Kuperberg. The Complexity ZOO. http://qwiki.stanford.edu/wiki/Complexity_Zoo, 2008.
- [2] Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. *Journal of the ACM*, 51 (4), p. 595–605, 2004.
- [3] Manindra Agrawal, Neeraj Kayal, and Nitin Saxena. Primes is in p. *Ann. of Math. Volume 160, Number 2*, 781–793, 2004.
- [4] Dorit Aharonov, Andris Ambainis, Julia Kempe, and Umesh Vazirani. Quantum walks on graphs. In *Proc. of ACM Symposium on Theory of Computation (STOC'01)*, p. 50–59, 2001.
- [5] Dorit Aharonov, Daniel Gottesman, Sandy Irani, and Julia Kempe. The power of quantum systems on a line. *quant-ph/0705.4077*, 2007.
- [6] Dorit Aharonov and Amnon Ta-Shma. Adiabatic quantum state generation and statistical zero knowledge. *Proc. 35th Annual Symposium on Theory of Computing*, pp. 20–29, 2003.
- [7] Dorit Aharonov, Wim van Dam, Julia Kempe, Zeph Landau, Seth Lloyd, and Oded Regev. Adiabatic quantum computation is equivalent to standard quantum computation. *arXiv:quant-ph/0405098*, 2005.
- [8] Andris Ambainis. Quantum lower bounds by quantum arguments. *Proc. 32th ACM Symposium on the Theory of Computing (STOC)*, p. 636–643, 2000.
- [9] Andris Ambainis and Oded Regev. An elementary proof of the quantum adiabatic theorem. *arXiv:quant-ph/0411152*, 2004.
- [10] Andris Ambainis, Robert Špalek, and Ronald de Wolf. A new quantum lower bound method with applications to direct product theorems and time-space tradeoffs. *Proc. 38th ACM Symposium on the Theory of Computing (STOC)*, p. 618–633, 2006.
- [11] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.
- [12] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of np. *J. ACM*, 45(1):70–122, 1998.
- [13] Francisco Barahona. On the computational complexity of ising spin glass models. *J. Phys. A: Math. Gen.* 15, 3241, 1982.

- [14] Rodney J. Baxter. *Exactly Solved Models in Statistical Mechanics*, p.47-59. Academic Press, 1982.
- [15] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *Proc. 39th Annual Symposium on Foundations of Computer Science, Palo Alto, CA*, 1998.
- [16] Paul Benioff. The computer as a physical system: a microscopic quantum mechanical model of computers represented by turing machines. *J. Stat. Phys.*, vol. 5, no. 5, pp. 562-591, 1980.
- [17] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh V. Vazirani. Strengths and weaknesses of quantum computing. *SIAM Journal on Computing* 26:1510-1523, 1997.
- [18] Michael Biafore. Can quantum computers have simple hamiltonians? *In Proc. Workshop on Physics of Computation*, pp. 63-86, Los Alamitos, CA, IEEE Computer Society Press, 1994.
- [19] Alfredo Braunstein, Marc Mezard, and Riccardo Zecchina. Survey propagation: an algorithm for satisfiability. *Random Structures and Algorithms* 27, 201-226, 2005.
- [20] Sergey Bravyi. Efficient algorithm for a quantum analogue of 2-sat. *quant-ph/0602108*, 2006.
- [21] Sergey Bravyi and Mikhail Vyalyi. Commutative version of the local hamiltonian problem and common eigenspace problem. *Quantum Inf. and Comp.*, Vol. 5, No. 3, pp.187-215, 2005.
- [22] Bradley A. Chase and Andrew J. Landahl. Universal quantum walks and adiabatic algorithms by 1d hamiltonians. *arXiv:0802.1207*, 2008.
- [23] Andrew Childs, Edward Farhi, and John Preskill. Robustness of adiabatic quantum computation. *Phys.Rev. A* 65, 012322, 2002.
- [24] Andrew M. Childs, Debbie W. Leung, and Michael A. Nielsen. Unified derivations of measurement-based schemes for quantum computation. *Phys. Rev. A* 71, 032318, 2005.
- [25] Andrew J. Daley, Corinna Kollath, Ulrich Schollwöck, and Guifre Vidal. Time dependent dmrg using adaptive effective hilbert spaces. *J.Stat.Mech.: Theor. Exp.* P04005, cond-mat/0403313, 2004.
- [26] Percy Deift, Mary Beth Ruskai, and Wolfgang Spitzer. Improved gap estimates for simulating quantum circuits by adiabatic evolution. *Quantum Information Processing*, Vol. 6, No. 2, 2007.
- [27] David DiVincenzo. private communication. *unpublished*, 2006.
- [28] Lior Eldar and Oded Regev. Quantum sat for a qutrit-cinquit pair is qma₁-complete. *unpublished, personal communication*, 2008.
- [29] Mark Fannes, Bruno Nachtergaele, and Reinhard F. Werner. Finitely correlated states on quantum spin chains. *Comm. Math. Phys.* 144, 3, pp. 443-490, 1992.

- [30] Mark Fannes, Bruno Nachtergaele, and Reinhard F. Werner. Ground states of vbs models on cayley trees. *J. Stat. Phys.* 66, 939-973, 1992.
- [31] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. Quantum adiabatic evolution algorithms with different paths. *arXiv:quant-ph/0208135*, 2002.
- [32] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum algorithm for the hamiltonian nand tree. *arXiv:quant-ph/0702144*, 2007.
- [33] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, Joshua Lapan, Andrew Lundgren, and Daniel Preda. A quantum adiabatic evolution algorithm applied to random instances of an np-complete problem. *Science* 292, 472, 2001.
- [34] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Daniel Nagaj. How to make the quantum adiabatic algorithm fail. *arXiv:quant-ph/0512159*, 2005.
- [35] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. A limit on the speed of quantum computation in determining parity. *Phys.Rev.Lett.* 81, 5442-5444, 1998.
- [36] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. Quantum computation by adiabatic evolution. *quant-ph/0001106*, 2000.
- [37] Edward Farhi and Sam Gutmann. Analog analogue of a digital quantum computation. *Phys. Rev. A* 57 2403, 1998.
- [38] Richard P. Feynman. Simulating physics with computers. *Internat. J. Theoret. Phys.* 21 (6/7), 467-488, 1982.
- [39] Richard P. Feynman. Quantum mechanical computers. *Opt. News*, vol. 11, pp. 11-46, 1985.
- [40] Michael Freedman, Alexei Kitaev, and Zhenghan Wang. Simulation of topological field theories by quantum computers. *Commun. Math. Phys.* 227, no.3, 587-603, 2002.
- [41] Michael Freedman, Michael Larsen, and Zhenghan Wang. A modular functor which is universal for quantum computation. *Commun. Math. Phys.* 227, no.3, 605-622, 2002.
- [42] Michael R. Garey and David S. Johnson. *Computers and Intractability. A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, San Francisco, CA, 1979.
- [43] Jeffrey Goldstone. Notes on the adiabatic theorem. *unpublished, personal communication*, 2008.
- [44] Tad Hogg. Adiabatic quantum computing for random satisfiability problems. *Phys. Rev. A* 67, 022314, 2003.
- [45] Lawrence Ioannou and Michele Mosca. Limitations of some simple adiabatic quantum algorithms. *arXiv:quant-ph/0702241*, 2007.

- [46] Sabine Jansen, Ruedi Seiler, and Mary-Beth Ruskai. Bounds for the adiabatic approximation with applications to quantum computation. *arXiv:quant-ph/0603175*, 2006.
- [47] Dominik Janzing. Spin 1/2 particles moving on a two dimensional lattice with nearest-neighbor interactions can realize an autonomous quantum computer. *Phys. Rev. A* **75**, 012307, 2007.
- [48] Dominik Janzing and Pawel Wocjan. Ergodic quantum computing. *Quantum Inf. Process.* **4**, 129, 2005.
- [49] Dominik Janzing, Pawel Wocjan, and Shenyu Zhang. A single-shot measurement of the energy of product states in a translation invariant spin chain can replace any quantum computation. *quant-ph/0710.1615*, 2007.
- [50] Jacob Jordan, Roman Orús, Guifre Vidal, Frank Verstraete, and J. Ignacio Cirac. Classical simulation of infinite size quantum lattice systems in two spatial dimensions. *cond-mat/070378*, 2007.
- [51] Stephen P. Jordan. Quantum computation beyond the circuit model. *Ph.D. thesis, MIT*, 2008.
- [52] William M. Kaminsky and Seth Lloyd. Scalable architecture for adiabatic quantum computing of np-hard problems. *Quantum Computing and Quantum Bits in Mesoscopic Systems 229-236*, A.J. Leggett, B. Ruggiero, and P. Silvestrini, Eds. New York: Kluwer Academic, *quant-ph/0211152*, 2004.
- [53] Julia Kempe, Alexei Kitaev, and Oded Regev. The complexity of the local hamiltonian problem. *Proc. 24th FSTTCS*, pp.372-383, 2004.
- [54] Julia Kempe and Oded Regev. 3-local hamiltonian is qma-complete. *Quantum Inf. Comput.*, **3**(3):258-264, 2003.
- [55] Alexei Yu. Kitaev. Fault-tolerant quantum computation by anyons. *Annals Phys.* **303**, pp.2-30, 2003.
- [56] Alexei Yu. Kitaev, Alexander H. Shen, and Michael N. Vyalyi. *Classical and Quantum Computation, Graduate Studies in Mathematics Vol.47*. American Mathematical Society, Providence, RI, 2002.
- [57] Emanuel Knill. Quantum randomness and nondeterminism. *Technical Report LAUR-96-2186, Los Alamos National Laboratory*, 1996.
- [58] Manfred Kudlek. Small deterministic turing machines. *Theoretical Computer Science*, **168**(2):241-255, 1996.
- [59] Christopher Laumann, Antonello Scardicchio, and Shivaji L. Sondhi. Cavity method for quantum spin glasses on the bethe lattice. *arXiv/0706.4391*, 2007.
- [60] Debbie W. Leung. Quantum computation by measurements. *Int. J. Quant. Inf.* **2**, 33, 2004.
- [61] Daniel Lidar. Towards fault tolerant adiabatic quantum computation. *arXiv:0707.0021*, 2007.

- [62] Yi-Kai Liu. The local consistency problem for stoquastic and 1-d quantum systems. *submitted to IEEE Conference on Computational Complexity (CCC), arXiv:0712.1388*, 2007.
- [63] Seth Lloyd. Robustness of adiabatic quantum computing. *unpublished, private communication*, 2003.
- [64] Norman Margolus. Parallel quantum computation. *In W. Zurek, editor, Complexity, Entropy and the Physics of Information, Addison Wesley*, 1990.
- [65] Ari Mizel, Daniel A. Lidar, and Morgan Mitchell. Simple proof of equivalence between adiabatic quantum computation and the circuit model. *Phys. Rev. Lett.* **99**, 070502, 2007.
- [66] Carlos Mochon. Anyons from non-solvable groups are sufficient for universal quantum computation. *Phys. Rev. A* **67**, 022315, 2003.
- [67] Carlos Mochon. Anyon computers with smaller groups. *Phys. Rev. A* **69**, 032306, 2004.
- [68] Valentin Murg, Frank Verstraete, and J. Ignacio Cirac. Variational study of hard-core bosons on a 2d optical lattice using projected entangled pair states. *Phys. Rev. A* **75**, 033605, 2007.
- [69] Daniel Nagaj, Edward Farhi, Jeffrey Goldstone, Peter Shor, and Igor Sylvester. The quantum transverse field ising model on an infinite tree from matrix product states. *arXiv:0712.1806*, 2007.
- [70] Daniel Nagaj and Shay Mozes. New construction for a qma complete three-local hamiltonian. *J. Math. Phys.* **48**, 072104, 2007.
- [71] Daniel Nagaj and Pawel Wocjan. Hamiltonian quantum cellular automata in 1dt. *arXiv:0802:0886*, 2008.
- [72] Michael A. Nielsen. Quantum computation by measurement and quantum memory. *Phys. Lett. A* **308**, 96, 2003.
- [73] Michael A. Nielsen and Isaac L. Chuang. *Quantum Information and Computation*. Cambridge University Press, Cambridge, UK, 2000.
- [74] Roberto Oliveira and Barbara Terhal. The complexity of quantum spin systems on a two-dimensional square lattice. *SIAM Journal of Computing*, Vol. **35**(5), p. 1070-1097, 2006.
- [75] Stellan Östlund and Stefan Rommer. Thermodynamic limit of dmr. *Phys. Rev. Lett.* **75**, pp. 3537, 1995.
- [76] Carlos Perez-Delgado and Donny Cheung. Models of quantum cellular automata. *quant-ph/0508164*, 2005.
- [77] David Perez-Garcia, Frank Verstraete, Michael M. Wolf, and J. Ignacio Cirac. Matrix product state representations. *quant-ph/0608197*, 2006.

- [78] John Preskill. Fault-tolerant quantum computation. *Introduction to Quantum Computation and Information*, edited by H.-K. Lo, S. Popescu, and T. P. Spiller (World Scientific, Singapore), 1998.
- [79] John Preskill and R. Walter Ogburn. Topological quantum computation. *Lecture Notes In Computer Science; Vol. 1509, Selected papers from the First NASA International Conference on Quantum Computing and Quantum Communications*, pp. 341–356, 1998.
- [80] Robert Raussendorf. Quantum cellular automaton for universal quantum computation. *Phys. Rev. A* **72**, 022301, 2005.
- [81] Robert Raussendorf and Hans Briegel. A one-way quantum computer. *PRL* **86**, 5188, 2000.
- [82] Subir Sachdev. *Quantum Phase Transitions*. Cambridge University Press, Cambridge, UK, 1999.
- [83] Ullrich Schollwöck. The density matrix renormalization group. *Rev.Mod.Phys* **77**, 259, 2005.
- [84] Norbert Schuch, Ignacio Cirac, and Frank Verstraete. The computational difficulty of finding mps ground states. *arXiv:0802.3351*, 2008.
- [85] Bart Selman, Henry Kautz, and Bram Cohen. Local search strategies for satisfiability testing. *Cliques, Coloring, and Satisfiability: Second DIMACS Implementation Challenge*, 1993.
- [86] Bart Selman, Hector Levesque, and David Mitchell. A new method for solving hard satisfiability problems. In *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI'92)*, pp. 440–446, 1992.
- [87] Daniel J. Shepherd, Torsten Franz, and Reinhard F. Werner. Universally programmable quantum cellular automaton. *PRL* **97**, 020502, 2006.
- [88] Yaoyun Shi. Both toffoli and controlled-not need little help to do universal quantum computation. *quant-ph/020115*, 2002.
- [89] Yaoyun Shi, Luming Duan, and Guifre Vidal. Classical simulation of quantum many-body systems with a tree tensor network. *cond-mat/0605597*, 2005.
- [90] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *Proc. 35th Annual Symposium on Foundations of Computer Science, Santa Fe, NM*, 1994.
- [91] Michael Sipser. *Introduction to the Theory of Computation*. Course Technology, Boston, MA, 2006.
- [92] Peter W. Shor Stephen P. Jordan, Edward Farhi. Error correcting codes for adiabatic quantum computation. *Phys. Rev. A* **74**, 052322, 2006.
- [93] Masuo Suzuki. Fractal decomposition of exponential operators with applications to many-body theories and monte carlo simulations. *Phys. Lett. A* **146**, 319, 1990.

- [94] Masuo Suzuki. General theory of fractal path integrals with applications to many-body theories and statistical physics. *J. Math. Phys.* **32**, 400, 1991.
- [95] Masuo Suzuki. Mathematical basis of computational statistical physics and quantum analysis. *Commun. Math. Phys.* **163**, 491, 1994.
- [96] Masuo Suzuki and Takashi Yamauchi. Convergence of unitary and complex decompositions of exponential operators. *J. Math. Phys.* **34**, 4892, 1993.
- [97] Tommaso Toffoli. Reversible computing. *MIT Technical Report MIT/LCS/TM-151*, 1980.
- [98] Hale F. Trotter. On the product of semi-groups of operators. *Proc. Am. Math. Soc.* **10**, 545, 1959.
- [99] Alan M. Turing. On computable numbers, with an application to the entscheidungsproblem. *Proc. London Math. Soc. Ser. 2* **42**, 230-265, 1937.
- [100] Alan M. Turing. Correction to: On computable numbers, with an application to the entscheidungsproblem. *Proc. London Math. Soc. Ser. 2* **43**, 544-546, 1938.
- [101] Wim van Dam, Michele Mosca, and Umesh Vazirani. How powerful is adiabatic quantum computation? *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science*, pp. 279-287, 2001.
- [102] Frank Verstraete and J. Ignacio Cirac. Renormalization algorithms for quantum many body systems in two and higher dimensions. *cond-mat/0407066*, 2004.
- [103] Frank Verstraete, Diego Porras, and J. Ignacio Cirac. Dmrg and pbc: A quantum information perspective. *Phys. Rev. Lett.* **93**, 227205, 2004.
- [104] Guifre Vidal. Efficient classical simulation of slightly entangled quantum computations. *Phys. Rev. Lett.* **91**, 147902, 2003.
- [105] Guifre Vidal. Efficient simulation of one-dimensional quantum many-body systems. *Phys. Rev. Lett.* **93**, 040502, 2004.
- [106] Guifre Vidal. Classical simulation of infinite-size quantum lattice systems in one spatial dimension. *cond-mat/0705597*, 2006.
- [107] Karl G. H. Vollbrecht and J. Ignacio Cirac. Quantum simulators, continuous-time automata, and translationally invariant systems. *quant-ph/0704.3432*, 2007.
- [108] Marko Žnidarič and Martin Horvat. Exponential complexity of an adiabatic algorithm for an np-complete problem. *Phys. Rev. A* **73** 022329, 2006.
- [109] Zhaohui Wei and Mingsheng Ying. Quantum adiabatic evolutions that can't be used to design efficient algorithms. *arXiv:quant-ph/0604077*, 2006.
- [110] Steven R. White. Density matrix formulation for quantum renormalization groups. *PRL* **69**, 2863, 1992.
- [111] Steven R. White and Adrian E. Feiguin. Real time evolution using the dmrg. *PRL* **93**, 076401, 2004.

- [112] Stephen Wolfram. *A New Kind of Science*. Wolfram Media, Champaign, IL, 2002.
- [113] A. Peter Young, Sergey Knysh, and Vadim N. Smelyanskiy. Size dependence of the minimum excitation gap in the quantum adiabatic algorithm. *arXiv:0803.3971*, 2008.

Massachusetts Institute of Technology
Department of Electrical Engineering
and Computer Science

Proposal for Thesis Research in Partial
Fulfillment
of the Requirements for the Degree of
Doctor of Philosophy

TITLE: Parallel Processor Architecture
SUBMITTED BY: Peter Nuth
305 Memorial Drive, 606C
Cambridge, MA 02139

(SIGNATURE OF AUTHOR)

DATE OF SUBMISSION: August 25, 2008
EXPECTED DATE OF COMPLETION: September 1990
LABORATORY: Artificial Intelligence Laboratory

BRIEF STATEMENT OF THE PROBLEM:

The proposed research is a study of processor architecture for large scale parallel computer systems. The thesis introduces mechanisms for fast context switching, synchronization between tasks, and run-time binding of variable names to processor memory. Various design tradeoffs are evaluated through simulation of a processor running a typical load. This work contains estimates of the speed and complexity of the different alternatives as implemented in VLSI.

Massachusetts Institute of Technology
Department of Electrical Engineering
and Computer Science
Cambridge, Massachusetts 02139

Doctoral Thesis Supervision Agreement

TO: Department Graduate Committee
FROM: Professor William J. Dally

The program outlined in the proposal:

TITLE: Parallel Processor Architecture
AUTHOR: Peter Nuth
DATE: August 25, 2008

is adequate for a Doctoral thesis. I believe that appropriate readers for this thesis would be:

READER 1: Professor Arvind
READER 2: Professor Thomas Knight

Facilities and support for the research outlined in the proposal are available.
I am willing to supervise the thesis and evaluate the thesis report.

SIGNED: _____
ASSOCIATE PROFESSOR OF ELECTRICAL ENGINEERING
AND COMPUTER SCIENCE

DATE: _____

Comments:

Massachusetts Institute of Technology
Department of Electrical Engineering
and Computer Science
Cambridge, Massachusetts 02139

Doctoral Thesis Reader Agreement

TO: Department Graduate Committee
FROM: Professor Arvind

The program outlined in the proposal:

TITLE: Parallel Processor Architecture
AUTHOR: Peter Nuth
DATE: August 25, 2008
SUPERVISOR: Professor William J. Dally
OTHER READER: Professor Thomas Knight

is adequate for a Doctoral thesis. I am willing to aid in guiding the research
and in evaluating the thesis report as a reader.

SIGNED: _____
PROFESSOR OF ELECTRICAL ENGINEERING
AND COMPUTER SCIENCE

DATE: _____

Comments:

Massachusetts Institute of Technology
Department of Electrical Engineering
and Computer Science
Cambridge, Massachusetts 02139

Doctoral Thesis Reader Agreement

TO: Department Graduate Committee
FROM: Professor Thomas Knight

The program outlined in the proposal:

TITLE: Parallel Processor Architecture
AUTHOR: Peter Nuth
DATE: August 25, 2008
SUPERVISOR: Professor William J. Dally
OTHER READER: Professor Arvind

is adequate for a Doctoral thesis. I am willing to aid in guiding the research
and in evaluating the thesis report as a reader.

SIGNED: _____

ASSISTANT PROFESSOR OF ELECTRICAL ENGINEERING
AND COMPUTER SCIENCE

DATE: _____

Comments:

Massachusetts Institute of Technology
Department of Electrical Engineering
and Computer Science
Cambridge, Massachusetts 02139

Doctoral Thesis Reader Agreement

TO: Department Graduate Committee
FROM: Professor William J. Dally

The program outlined in the proposal:

TITLE: Parallel Processor Architecture
AUTHOR: Peter Nuth
DATE: August 25, 2008
SUPERVISOR: Professor William J. Dally
OTHER READER: Professor Arvind
OTHER READER: Professor Thomas Knight

is adequate for a Doctoral thesis. I am willing to aid in guiding the research
and in evaluating the thesis report as a reader.

SIGNED: _____
ASSOCIATE PROFESSOR OF ELECTRICAL ENGINEERING
AND COMPUTER SCIENCE

DATE: _____

Comments:
