Title
Multiqubit Tunneling and Evolution in Quantum Annealing

Author
Brady, Lucas

Publication Date
2018

University of California
Santa Barbara

Multiqubit Tunneling and Evolution in Quantum Annealing

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Physics

by

Lucas T. Brady

Committee in charge:

Professor Wim van Dam, Chair
Professor Mark Srednicki
Professor David Weld

June 2018

The Dissertation of Lucas T. Brady is approved.

Professor Mark Srednicki

Professor David Weld

Professor Wim van Dam, Committee Chair

June 2018

Multiqubit Tunneling and Evolution in Quantum Annealing

by

Lucas T. Brady

I would like to dedicate this dissertation to Beth and Carleton Mcmullin and Jackie and Dale Brady, my maternal and paternal grandparents, respectively. Throughout their lives, they all influenced me in their unique ways and helped shape me into the person capable of this work. Beth passed away when I was younger, and Carleton and Jackie died while I was completing my PhD. They will all be missed. As of this writing, Dale continues to enjoy life and hopefully will for years to come. Thank you.

Acknowledgements

I would like to thank the University of California, Santa Barbara's physics and computer science departments for their support of this work. Especially, I would like to thank my advisor, Professor Wim van Dam, for his support and guidance over the years. Without Wim, none of this research would have happened, and he has been an excellent advisor, especially for me. I also want to thank Tengiz Bibilashvili and, especially, Sathya Guruswamy who have both provided me with mentorship for teaching and how to pass on the knowledge I have learned.

I would also like to thank all my past advisors from other institutions, especially Vatche Sahakian, Wally Melnitchouk, and Alberto Accardi, who provided me with guidance and advise on my path even if I have left their areas of research.

Science is not done in a vacuum, so I would like to thank my friends and colleagues for their support and friendship these past few years. I must mention my roommates, Evan Bauer, Dan Ish, Michael Swift, and Emmy (cat). Thanks as well to my lab mates, both of the quantum variety and the crypto variety for putting up with my weird work habits. Finally, I want to thank all my friends, whether from board gaming, physics, hiking, or any other activity for making these years bearable.

Of course none of this research would have been possible without financial support. This dissertation is based upon work supported by the National Science Foundation under Grants No. 1314969 and No. 1620843. The numerical work in this dissertation was performed using computing resources provided by the Center for Scientific Computing from the CNSI, MRL: under NSF MRSEC (DMR-1720256) and NSF CNS-0960316.

Above all else, I want to thank my parents, John and Charis Brady, and my family as a whole. Without you I never would have made it to this point, thank you.

Curriculum Vitæ

Lucas T. Brady

Education

2018	Ph.D. in Physics (Expected), University of California, Santa Barbara.
2016	M.A. in Physics, University of California, Santa Barbara.
2013	B.S. in Physics, Harvey Mudd College.

Publications

- Lucas T. Brady, Wim van Dam, "Discrepancies between Asymptotic and Exact Spectral Gap Analyses of Quantum Adiabatic Barrier Tunneling." Phys. Rev. A **95**, 052350 (2017). arXiv:1609.09137
- Lucas T. Brady, Wim van Dam, "Necessary Adiabatic Run Times in Quantum Optimization." Phys. Rev. A **95**, 032335 (2017). arXiv:1611.02585
- Lucas T. Brady, Wim van Dam, "Spectral-gap analysis for efficient tunneling in quantum adiabatic optimization." Phys. Rev. A **94**, 032309 (2016). arXiv:1601.01720
- A. Accardi, L. T. Brady, W. Melnitchouk, J. F. Owens, N. Sato, "Constraints on large-x parton distributions from new weak boson production and deep-inelastic scattering data." Phys. Rev. D **93**, 114017 (2016). arXiv:1602.03154
- Lucas T. Brady, Wim van Dam, "Quantum Monte Carlo simulations of tunneling in quantum adiabatic optimization." Phys. Rev. A **93**, 032304 (2016). arXiv:1509.02562
- P. J. Ehlers, A. Accardi, L. T. Brady, W. Melnitchouk, "Nuclear effects in the proton-deuteron Drell-Yan process." Phys. Rev D **90**, 014010 (2014). arXiv:1405.2039
- Lucas Brady and Vatche Sahakian, "Scrambling with matrix black holes." Phys. Rev. D **88**, 046003 (2013). arXiv:1306.5200
- L. T. Brady, A. Accardi, W. Melnitchouk and J. F. Owens, "Impact of PDF uncertainties at large x on heavy boson production." JHEP **2012**, 019 (2012). arXiv:1110.5398
- L. T. Brady, A. Accardi, T. J. Hobbs and W. Melnitchouk, "Next-to leading order analysis of target mass corrections to structure functions and asymmetries." Phys. Rev. D **84**, 074008 (2011). arXiv:1108.4734

Pre-Print Publications

- Lucas T. Brady, Wim van Dam, "Evolution-Time Dependence in Near-Adiabatic Quantum Evolutions." (2018) arXiv:1801.04349

Abstract

Multiqubit Tunneling and Evolution in Quantum Annealing

by

Lucas T. Brady

Quantum computing seeks to use the powers of quantum mechanics to accomplish tasks that classical computers cannot easily accomplish. Adiabatic quantum computing is one flavor of quantum computing that slowly changes a system with time in such a way that the state remains close to the ground state for the entire evolution. An easy way to implement and study adiabatic computing is using stoquastic Hamiltonians with no-sign problem. These Hamiltonians can be readily simulated on classical computers using quantum Monte Carlo schemes; though, it is unknown whether such simulation captures the speed and power of the quantum computation.

In this work, I study stoquastic Hamiltonians, focusing on quantum mechanical tunneling through classically inaccessible regions and barriers. I study multi-qubit tunneling problems to determine what properties of the barrier make this problem quantum mechanically or classically hard or inefficient. I compare quantum adiabatic computing with a path-integral Quantum Monte Carlo algorithm, which is a classical algorithm designed to simulate the quantum mechanical dynamics of the adiabatic evolution. Limited numerical data shows strong correlation between adiabatic and Monte Carlo runtimes, but due to computational limits, only small system sizes could be sampled.

Additionally, I study the properties of the quantum adiabatic algorithm in both the asymptotic limit of a large number of qubits and long runtimes and more realistic finite system sizes and non-adiabatic conditions. I develop a modification of the existing Villain transformation that allows it to find the asymptotic running time of the quantum

evolution for different barrier sizes. Furthermore, I compare this to finite system sizes and discover the extremely large numbers of qubits are needed for these asymptotics to be useful.

When the quantum dynamics are run faster than adiabatically, new behavior arises that could potentially lead to enhancements in the quantum runtimes. I explore these enhancements, discovering when they do and do not occur. I also develop a broad framework for determining the nature of the near-adiabatic behavior of any quantum evolution, using only information about the energy gap between the ground state and the first excited state.

Contents

Chapter 1

Introduction

Quantum computing is a potentially very powerful application of physics to practical applications. Recent years have seen a blossoming of quantum computing research and development with new results and devices being produced at a rapid rate. As near-term quantum computers become more powerful than classical computers for certain tasks, it is important to understand the limitations and capabilities of quantum computing.

At its core, all computing is quantum mechanical since our world is inherently quantum mechanical. What distinguishes true quantum computing from classical computing is that quantum computing requires a quantum mechanical theory to model it whereas a classical computer can be modeled using classical physics. The key question is whether a quantum mechanical model gives more computing power and how much more power it gives.

Therefore, quantum computing is modeled in Hilbert spaces which come with their assortment of entanglement, multiple bases, superposition, measurement issues, tunneling, and all the other uniquely quantum mechanical features. Despite the plethora of seemingly quantum mechanical ideas, it is not known which of them give quantum computing its power, and in fact it is not proven (but is highly suspected) whether quantum mechanics provides a substantive advantage over classical mechanics in computational power.

The ultimate goal of this dissertation is to look at one specific paradigm of quantum computing, adiabatic quantum computing, and find evidence for or against its quantum advantage over classical computing and what aspects of quantum mechanics are useful in distinguishing quantum adiabatic computing from classical computing. To that end, I explore quantum adiabatic computing in a tunneling setting. The problem I look at is initially chosen because tunneling was thought to be a potential key to quantum advantage, but this problem eventually is just a testing bed for other ideas in quantum computing.

1.1 Current State of Quantum Computing

Quantum computing as a field started in 1982 when Feynman [1] proposed using a quantum system as a simulator for other, more complicated quantum systems, which was (and still is) thought to be a classically difficult problem for general systems. The field was further buoyed by the discovery of several early quantum algorithms, including the Deutsch-Jozsa algorithm [2], Grover's search [3], and Shor's algorithm [4], all of which demonstrated some improvement over all known or, in the first two cases, possible classical algorithms.

More recently, Hamiltonian based quantum computing as exemplified by quantum annealing [5], quantum adiabatic optimization [6], and quantum approximate optimization [7], have sought to use various methods to find the ground state of a Hamiltonian as a computational tool rather than modifying the properties of a quantum state. These algorithms are usable on much smaller problem instances than other quantum algorithms, and often they are touted (without direct evidence) as being more robust to computational errors, making them targets for near-term quantum devices.

Currently, there are about 60 known quantum algorithms that provide some form of quantum speedup over known classical algorithms [8].

With the promise that quantum algorithms can provide provable polynomial speedups, as highlighted by Grover's search [3], and potential exponential speedups, as exemplified by Shor's algorithm [4], quantum computing has garnered interest from industry and academy. At a recent American Physical Society meeting, Google [9] announced that they are testing a 72 qubit quantum device. If this device meets appropriate error thresholds, it will be capable of performing the first quantum advantage or supremacy test on a universal quantum computer [10].

Other experimental groups are developing quantum computers with similar numbers of qubits. IBM has a cloud-based quantum computer, accessible over the internet that currently has up to 20 qubits [11]. There are rumors that Microsoft's experimental topological computing teams are nearing a single working qubit. Meanwhile start-up companies are also producing viable quantum computers such as Rigetti's 19Q machine [12]. Ion trap quantum computers are lagging slightly; for instance, Chris Monroe's group currently has more than five qubits [13].

Of course, given that this dissertation discusses quantum annealing and quantum adiabatic computing, I would be remiss if I did not mention the D-Wave machine [14]. Currently, the most advanced, D-Wave computer, the 2000Q, has 2000 qubits, specifically designed to perform quantum annealing using stoquastic Hamiltonians. The fact that this computer only uses stoquastic Hamiltonians is a significant disadvantage theoretically, as discussed in Ch. 2, but beyond the types of Hamiltonians, the computer itself has significant thermal noise [15]. It is unclear how quantum mechanical the D-Wave computations are, and tests have shown no scaling advantage over classical algorithms [16, 17].

1.2 Organization

Chs. 2 & 3 both focus on background material for my subfield. The next chapter, Ch. 2, covers quantum adiabatic computing. This chapter goes over the basis of the algorithm, the terminology surrounding quantum annealing, quantum adiabatic computing, and quantum adiabatic optimization. I will also discuss stoquastic Hamiltonians and how stoquastic computing differs from general adiabatic quantum computing. Ch. 3 introduces a couple of simple models and problems that can be implemented on adiabatic computers. These models are relatively easy to analyze while still demonstrating essential features of adiabatic Hamiltonians, allowing analysis in a simple setting.

The remainder dissertation is focused on original work done in collaboration with my advisor, Wim van Dam. All of these chapters cover work that was previously published [18, 19, 20, 21, 22].

Ch. 4 (based on [18]) explores path-integral Quantum Monte Carlo, a classical algorithm inspired by quantum annealing. I numerically implement a path-integral QMC algorithm and compare its running time to the runtime for quantum adiabatic optimization for the same problem instance. I mostly find good correspondence between quantum and classical algorithms in this case which has subsequently been confirmed analytically [23, 24].

For a long time, the community had unproven results about the asymptotic runtime scaling of a symmetric multi-qubit tunneling problem with n qubits [25]. In Ch. 5, I go through, prove, and expand upon these scaling results. This chapter is closely related to App. A where I develop a version of the Villain transformation [26] that is useful in a quantum algorithm setting. Both Ch. 5 and App. A are adapted from [19].

In the previous chapter, I developed information about the asymptotic scaling of multi-qubit tunneling runtimes, but in Ch. 6 (based on [20]), I revisit the applicability of this asymptotic behavior to finite n behavior. I find that extremely large n are needed to see the true asymptotic behavior, and it is even possible mistake polynomial and exponential scaling.

The next two chapters go beyond the adiabatic limit and look at the behavior of this algorithm on instances run faster than adiabatically. In Ch. 7 (based on [21]), I examine a substantial non-adiabatic speedup that was discovered originally by Muthukrishnan *et al.* [27]. This non-adiabatic speedup allows for a problem that should require polynomial or exponential time to be solved in constant time, independent of n. In this chapter, I show that this speedup is atypical and due to the delicate balance between the initial and final Hamiltonians and spectra.

Ch. 8 continues looking at non-adiabatic behavior but this time with an eye towards characterizing the dependence of the success probability on the total runtime, τ. These results are more general than my specific toy model and show how avoided-level crossings effect the behavior of fine oscillations in the success probability. I also use these results to connect the adiabatic version of Grover's search more closely to the original digital version.

Finally in Ch. 9, I provide more context for my results and conclude this dissertation.

Chapter 2

Quantum Adiabatic Optimization

Quantum Adiabatic Optimization (QAO) is a quantum algorithm first formalized by Farhi *et al.* [6]. Rather than the standard quantum circuit model, QAO relies on the time evolution properties of quantum systems, as characterized by the quantum adiabatic theorem, originally proposed by Born and Fock [28].

QAO has been touted as more robust to noise and errors than quantum circuit models, leading to it often being proposed as a potential application for near-term quantum computers and devices. This robustness is largely folklore and is simply based on the facts that the adiabatic schedule is flexible and that in optimization problems it is often good enough to find solutions close to the optimum. Many near-term quantum devices have started adopting quantum adiabatic evolution as their early application, such as ion-trap computers [29] and the controversial D-Wave device [14].

This chapter focuses first on the adiabatic theorem and its various incarnations. After that, I discuss QAO and the related notion of quantum annealing. Finally, I end with a discussion of stoquastic Hamiltonians and why these are of key importance for QAO applications and theory.

2.1 Quantum Adiabatic Theorem

The quantum adiabatic theorem is old, dating back all the way to Born and Fock [28] in 1928; however, exact, rigorous proofs of the adiabatic theorem and its conditions are more recent [30, 31]. The basic statement of the adiabatic theorem is simple: a system initialized in the i^{th} energy state of a time dependent quantum Hamiltonian at time $t = 0$ will remain in the i^{th} energy state if the Hamiltonian changes adiabatically and there is a non-zero energy gap separating the i^{th} state from all other energy states.

QAO focuses exclusively on the ground state of the quantum system. Therefore, the quantum adiabatic theorem for our purposes can be summarized as: a quantum system will remain in its ground state if the system evolves adiabatically and the ground state remains non-degenerate. The two key ideas here that require further discussion are the conditions on the end. How quickly can a system evolve and still be considered adiabatic? The answer to this rests in the fact that the ground state is non-degenerate and has a spectral gap, $\Delta(t)$. The spectral gap, $\Delta(t)$, is just the instantaneous energy difference between the ground state and the first excited state, and it turns out to be one of the key quantities that defines the adiabatic speed-limit.

I will parameterize the time evolution using a normalized time quantity $s \in [0, 1]$. This quantity, s, will mark off the progress of the evolution from $t = 0$ to $t = \tau$, the total evolution-time or runtime of the algorithm, with $s = t/\tau$. The quantum adiabatic condition provides a sufficient (but not necessary) constraint on τ to ensure adiabatic evolution. A folklore version of the adiabatic condition that is often cited but not rigorous (or ultimately correct) is given by

$$\tau \gg \max_{s\in[0,1]} \frac{|\langle\psi_0| \frac{d\hat{H}}{ds} |\psi_1\rangle|}{\Delta(s)^2} \equiv \max_{s\in[0,1]} |\rho(s)|. \tag{2.1}$$

where $|\psi_i\rangle$ refers to the i^{th} eigenstate of the Hamiltonian $\hat{H}(s)$. I have defined the relevant quantity in the adiabatic condition to be $\rho(s)$ since this quantity will become important for my work in Ch. 8.

This folklore condition is often sited and used, and it is considered a good rule of thumb by the community. However, more rigorous conditions have been proven. One particularly relevant proof is provided by Lidar, Rezakhani, and Hamma [30]. Lidar, Rezakhani, and Hamma's proof relies on the separability of the Hilbert space into n distinct subspaces (e.g. n spin-$\frac{1}{2}$ qubits), and is specifically geared to application for boundary cancellation methods which I will briefly allude to in Ch. 8. Their proof has more power than just a simple restriction on τ, but for our purposes, their adiabatic condition can be summarized as

$$\tau \gg \max_{s\in[0,1]} A\frac{||\frac{d\hat{H}}{ds}||^2}{\Delta(s)^3}, \tag{2.2}$$

where $||\cdot||$ denotes the operator norm, and A is a constant which includes information about properties of the Hilbert space and the derivatives of the spectral gap at $s = 0, 1$.

Another relevant rigorous proof comes from Jansen, Ruskai, and Seiler [31]. Their results agree with Lidar's in the appropriate limits and work under slightly different conditions. As I will discuss in the next section, the exact forms of these adiabatic conditions does not matter much for my purposes in complexity theory. What does matter is that the adiabatic condition depends directly on a polynomial of the time derivative of the Hamiltonian and inversely on a polynomial of the spectral gap.

The adiabatic theorem has many uses throughout physics, but our main concern with it will be its ability to aid in computational tasks.

2.2 Quantum Adiabatic Algorithms

My primary focus is on the Quantum Adiabatic Optimization (QAO) algorithm first developed by Farhi, Goldstone, Gutmann, and Sipser [6]. This algorithm uses the power inherent in the quantum adiabatic theorem to find the solution to a classical optimization problem.

2.2.1 Statement of Algorithm

The starting point of QAO are two Hamiltonians, $\hat{H}_0$ and $\hat{H}_1$, both of which are time independent. The initial Hamiltonian, $\hat{H}_0$ should be an easy to analyze and study system with a known and easily preparable ground state. One extremely common choice of an initial Hamiltonian is a uniform transverse field acting on n spin-$\frac{1}{2}$ particles, or qubits:

$$\hat{H}_0 = -\sum_{i=1}^{n} \sigma_x^{(i)}. \tag{2.3}$$

This choice of initial Hamiltonian is simple and well studied, but as I discuss in Sec. 2.3, it can have severe limitations.

On the other hand, the final Hamiltonian, $\hat{H}_1$, should have a simple description but an unknown ground state. As its name suggests, QAO is designed for $\hat{H}_1$ to encode a classical optimization problem. A classical optimization problem starts with a function $f(x)$ that takes in length n bit strings and outputs some value. The goal is to find the bit string x, for which $f(x)$ is minimized (or maximized, but that problem is equivalent to minimizing $-f(x)$). Optimization can be reframed as a decision problem by instead asking whether there exists an x such that $f(x)$ is lower than a specified value.

The goal of encoding the optimization problem into $\hat{H}_1$ is to ensure that the ground state of $\hat{H}_1$ in the σ_z basis, or computational basis, of the n qubits is the bit string x that minimizes $f(x)$. The original method for performing this encoding, proposed by Farhi *et al.* [6] is to make $\hat{H}_1$ diagonal in the computational basis, given by $|x\rangle$ for $x \in 0, 1^n$:

$$\hat{H}_1 = \sum_{x \in 0,1^n} f(x) \left|x\right\rangle \left\langle x\right| . \tag{2.4}$$

QAO solves the optimization problem, $f(x)$, by adiabatically changing the Hamiltonian of the system between $\hat{H}_0$ and $\hat{H}_1$:

$$\hat{H}(s) = g(s)\hat{H}_0 + h(s)\hat{H}_1, \tag{2.5}$$

where $g(0) = h(1) = 1$ and $g(1) = h(0) = 0$. In this dissertation, I will restrict my algorithms to simple linear interpolation where $g(s) = (1 - s)$ and $h(s) = s$, but other choices are valid, especially when using boundary cancellation methods [30].

The adiabatic theorem guarantees that if the system starts in the ground state of $\hat{H}_0$ and evolves slowly enough in runtime τ, then the final state of the system will be the ground state of $\hat{H}_1$. Measurement of the final state should yield the desired solution to the optimization problem.

2.2.2 Adiabatic Runtimes

In quantum adiabatic computing, the inverse spectral gap, $\Delta(s)^{-1}$ is viewed as the key quantity determining the runtime, τ, of quantum algorithms based on the adiabatic conditions in Eqs. 2.1 & 2.2. The relevant matrix norms of the Hamiltonian that appear in the numerators of these conditions are almost always proportional to a low-degree polynomial in the log of the dimension of the Hilbert space, $\log_2 N$, which in a qubit system is equal to the number of qubits, $n = \log_2 N$.

On the other hand, $\Delta(s)$ can become exponentially small in the number of qubits,

n, leading to it being the key quantity of interest. In complexity theory, we care much more about the separation between an algorithm that has runtime that scales polynomially with n and one that scales exponentially in n. Therefore, The exact form of the adiabatic condition is not crucially important for my purposes, but the fact that the adiabatic condition depends inversely on a polynomial in the spectral gap, $\Delta(s)$ is important. Throughout this dissertation, I will analyze the runtime of the quantum adiabatic algorithm by reference to the size of the spectral gap.

Importantly, the adiabatic conditions listed in Eqs. 2.1 & 2.2 are merely sufficient conditions, not necessary conditions. It is possible that a quantum adiabatic algorithm can be run faster than the adiabatic condition suggests with no ill effects. Necessary conditions for adiabatic quantum computing are few and far between (for an example of one, see [32]). In Ch. 8, I will discuss some of the consequences for running QAO faster than the adiabatic condition suggests.

2.2.3 Quantum Annealing

Quantum Adiabatic Optimization is often associated with quantum annealing which is considered a broader class of quantum algorithms. The term quantum annealing dates back to the work of Finnila *et al.* [33] and Kadowaki and Nishimori [5] who proposed using a changing quantum system to prepare a solution to a desired problem.

Quantum annealing was originally proposed as a quantum version of classical simulated annealing, where the ground state of a system is found by slowly cooling the system. The thermal effects initially keep the system in an excited state and allow it to explore more of the state space, but as the system cools it gets frozen into the true ground state. Quantum annealing as proposed by [33, 5] sought to perform the same procedure but switching quantum affects, as characterized by a transverse field, for thermal affects.

In essence, these original quantum annealing proposals relied on the same principles as QAO as proposed by Farhi *et al.* [6]. However, Farhi *et al.* formalized the runtime dependence on the adiabatic condition and couched the technique as an algorithm within the larger framework of computer science.

In the literature, quantum annealing has taken on a broader definition than QAO, with QAO being viewed as a form of quantum annealing. Quantum annealing can also refer to algorithms that are run non-adiabatically, such as I cover in Ch. 7 and Ch. 8. Additionally, QAO refers specifically to ground state computation at zero temperature whereas quantum annealing can encompass computation with higher excited states and even thermal excitations.

2.2.4 Quantum Universality

In computational complexity theory, BQP (Bounded Quantum Polynomial Time) is the complexity class that contains all the computational tasks that are efficiently solvable by a quantum computer. BQP is the class of all problems which can be solved efficiently, meaning the runtimes scale polynomially with the number of qubits, n, on a quantum computer with success probability bounded by at least $2/3$. The quantum circuit model is known to be universal, meaning it can efficiently implement any algorithm in BQP. The question of whether QAO is universal for quantum computation is a less understood issue.

In their original paper proposing QAO, Farhi *et al.* [6] showed that any QAO algorithm can be efficiently implemented on a quantum circuit model. The key insight here is that the continuous time path of the quantum annealing algorithm can be broken into discreet time steps using a Suzuki-Trotter expansion [34, 35]. These discreet time steps can then be implemented on a circuit based quantum computer, and the details of doing so were worked out in [6]. This does not show that QAO is universal which requires the opposite direction, that an adiabatic quantum computer can implement any quantum circuit algorithm efficiently.

Aharonov *et al.* [36] proved that all quantum circuit algorithms can be implemented as QAO algorithms. However, their construction requires specific types of Hamiltonians, that among other issues are non-stoquastic. I will discuss stoquastic Hamiltonians in depth in the next section, but for now, it suffices to say that stoquastic Hamiltonians are preferred due to their ease of implementation and analysis. It is currently unknown by the quantum computing community whether QAO with stoquastic Hamiltonians is universal. Stoquastic QAO can implement an algorithm for any problem in BQP, but the question of universality relates to whether that implementation can be made efficiently, i.e. in polynomial time.

2.3 Stoquastic Hamiltonians

As mentioned in the last section, it is currently unknown how much power quantum adiabatic computing with stoquastic Hamiltonians has in relation to other quantum computing paradigms. A stoquastic Hamiltonian is one which does not suffers from the "sign problem" [37], and the concept of stoquasticity was first defined by Bravyi *et al.* [38].

A stoquastic Hamiltonian is one for which all the off-diagonal matrix entries are strictly real and non-positive, with no restrictions on the diagonal matrix entries. This definition shows that stoquasticity is incredibly basis dependent. Therefore, it is often common to describe Hamiltonians as being stoquastic with respect to a natural basis, in our case, the σ_z basis of the qubits. Furthermore, in time dependent Hamiltonians, such as quantum annealing uses, a Hamiltonian is stoquastic if it remains stoquastic for all time.

One of the early reasons why stoquacticity mattered is because it has an impact on the Gibbs distribution of the thermal system. The Gibbs density matrix, $\rho = \frac{e^{-\beta \hat{H}}}{\mathrm{Tr}(e^{-\beta \hat{H}})}$, for a stoquastic Hamiltonian has all positive or real matrix entries, which, as I discuss in Ch. 4, can be utilized to simulate stoquastic systems classically using Monte Carlo techniques, though the efficiency of this simulation is not known in general.

The standard form of QAO, proposed by Farhi *et al.* [6], is given by

$$\hat{H}(s) = -(1-s)\sum_{i=1}^{n} \sigma_x^{(i)} + s \sum_{x \in 0,1^n} f(x) \left|x\right\rangle \left\langle x\right| \tag{2.6}$$

and is completely stoquastic. In addition, one of the most well known and controversial implementations of quantum annealing, the D-Wave system [14], is only capable of implementing stoquastic Hamiltonians of the form Eq. 2.6.

An important question, and one that will be discussed in more depth in Ch. 4 is whether QAO with stoquastic Hamiltonians can achieve quantum universality. The proof that quantum adiabatic computing is universal [36] explicitly uses non-stoquastic Hamiltonians in its construction. Currently the answer to the question of stoquastic QAO's power is still unknown.

There is mounting evidence that stoquastic QAO does not achieve full quantum universality. Several studies [39, 40, 18, 23, 24] have provided either numerical or theoretical evidence that stoquastic Hamiltonians are no better than classical analoguous algorithms such as Quantum Monte Carlo. My own work [18] fits into this narrative and is outlined in depth in Ch. 4.

Very recently, Marvian, Lidar, and Hen [41] showed that the problem of finding a stoquastic basis for a non-stoquastic Hamiltonian is an NP-complete problem. This means that the problem of "curing the sign-problem" is as hard as any problem in the complexity class NP. In complexity theory, it is not proven but is highly suspected that NP-complete problems cannot be solved efficiently by either classical or quantum computers.

Chapter 3

Symmetric Barrier Tunneling Problems

In order to understand the power of QAO better, I will focus in on relatively simple toy problems to stand as a testbed for studying the power and deficiencies of quantum adiabatic computation. In choosing these problems, I focus on problems that exhibit barrier tunneling in a quantum setting.

Especially earlier in the developement of quantum annealing, it was expected that certain quantum affects such as quantum tunneling, entanglement, and superposition would be candidates for the power of quantum computing over classical computing. Thanks in part to the work presented in this dissertation, there is growing evidence that quantum tunneling alone is not sufficient to ensure quantum speedup over classical algorithms. Now, it is much more common to believe that a wholistic view of quantum mechanics is necessary to understand quantum advantage if it exists.

One key feature of barrier tunnelling problems that ends up being more general is that they lead to avoided level crossings in the energy spectrum. This means that during barrier tunneling, the spectral gap becomes small which, as I discussed in Ch. 2, means that QAO will have difficulty. In early chapters, I will characterize the spectral gap through this avoided level crossing, but in later chapters I will move on to methods for dealing with and attempting to avoid these avoided crossings.

I discuss two barrier tunneling models. Both are qubit symmetric, meaning the Hamiltonians are the same under any permutation of the qubits. I exploit this symmetry both for ease of analytic work and for computational techniques since qubit symmetry lends itself to effective reduction of the Hilbert space to the symmetric subspace. The first tunneling model is known as the multi-qubit symmetric barrier problem, and will be the focus of most of my study. The other model is the so called p-spin model, specifically the $p = 3$ case, which does not include an explicit classical barrier but does exhibit a barrier in the semi-classical potential picture.

3.1 Multi-Qubit Barrier Model

The first tunneling model I consider is the symmetric barrier tunneling problem. This problem has a long history with Farhi *et al.* analyzing a restricted version of it in their original paper [6] and many other groups looking at it since then [42, 40, 18, 43, 27, 23, 24, 19, 20, 21, 22]. This model relies on the standard stoquastic QAO construction with an initial Hamiltonian based on the transverse field acting on n qubits:

$$\hat{H}_0 = -\sum_{i=1}^{n} \sigma_x^{(i)}. \tag{3.1}$$

The final Hamiltonian contains the actual barrier information and is formulated in terms of the Hamming weight of the states of the system. The Hamming weight, $|x|$, of a string of bits (or qubits), x, is the number of ones in the string, usually measured in the σ_z basis, which is also referred to as the computational basis. Therefore, with n

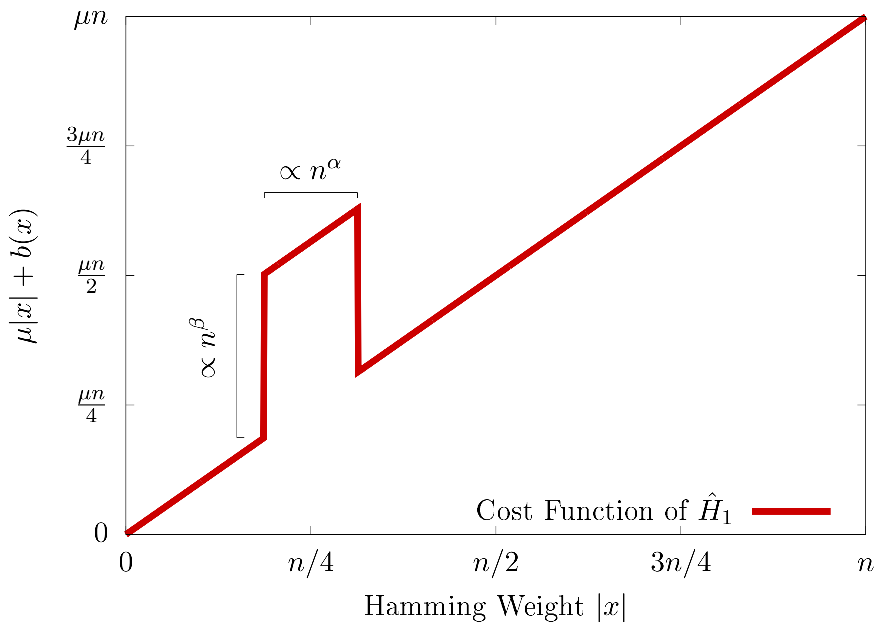

Figure 3.1: The cost function of the final Hamiltonian, $\hat{H}_1$. This cost function only depends on the Hamming weight of the bit strings. The barrier's height and width scale as polynomials with the number of qubits n.

qubits, the Hamming weight can run between 0 and n, and a given state of qubits can be a superposition of multiple Hamming weight states. The final Hamiltonian, shown in Fig. 3.1, is linear in the Hamming weight with a barrier function $b(|x|)$:

$$\hat{H}_1 = \sum_{x \in \{0,1\}^n} (\mu|x| + b(|x|)) \, |x\rangle \langle x| \, . \tag{3.2}$$

The constant μ determines the slope of the final cost function, which is relevant in Ch. 7, and $b(|x|)$ is a barrier function. The barrier function I use is centered around $|x| = n/4$, and has some width proportional to n^{α} and height proportional to n^{β} for $\alpha, \beta > 0$. In subsection 3.1.1, I will outline the different types of barriers I use.

One extremely important aspect of this Hamiltonian is that it is qubit symmetric. Both the initial and final Hamiltonians are symmetric under exchange of the qubits. This

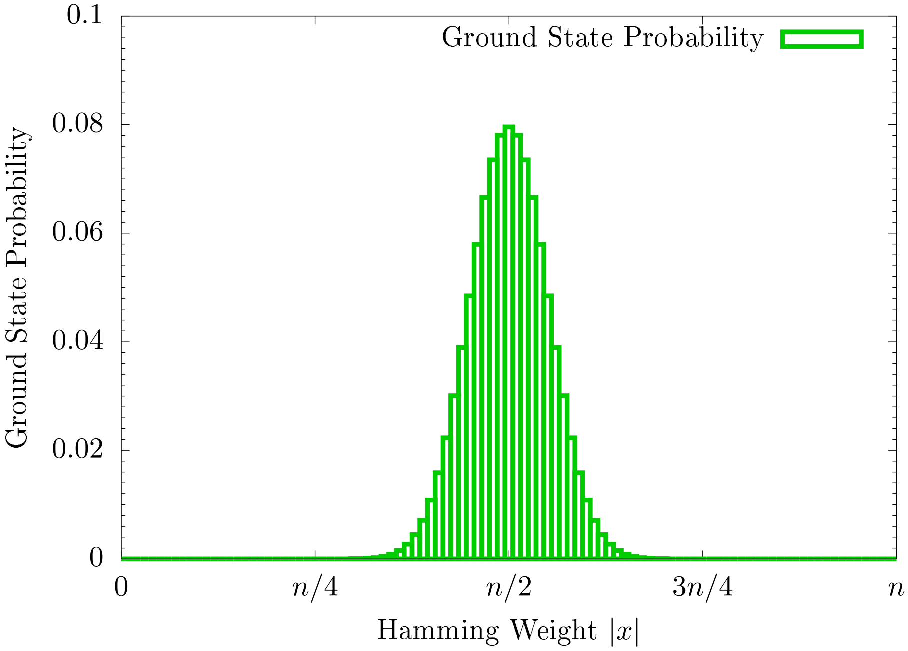

Figure 3.2: The ground state distribution the initial Hamiltonian, $\hat{H}_0$. The ground state is a binomial distribution in Hamming weight, centered on $n/2$. This distribution must tunnel through the barrier at $n/4$ to reach the true final ground state. This plot uses $n = 100$.

means that the Hamming weight of the current state fully determines the energy of the system. The initial Hamiltonian's ground state is a binomial distribution, centered on $|x| = n/2$, shown in Fig. 3.1. This ground state must tunnel through the barrier in the final Hamiltonian to reach the true final ground state, the all 0 state. In the Appendix A, I discuss the semi-classical version of this Hamiltonian picture.

The qubit symmetry of this problem allows it to be studied entirely in its symmetric subspace as discussed in subsection 3.1.2. The symmetric subspace is also $n + 1$ dimensional as opposed to the 2^n dimensionality of the full Hilbert space which allows numerical simulation up to a substantial number of qubits as discussed in subsection 3.1.3.

Much of QAO and quantum annealing relies on knowledge about the spectral gap, especially the minimum spectral gap. At the beginning of my research, the minimum spectral gap for this problem had been postulated, without formal proof, by Jeffery Goldstone. Subsequent work [42, 19, 23, 24] has confirmed the result of Goldstone, and Ch. 5 outlines my contribution to this proof, along with the mathemtical tools I developed to do so. The minimum spectral gap is given by

$$g_{\min} \propto \begin{cases} 1 \text{ if } \alpha+\beta \leq \frac{1}{2}, \\ n^{1/2-\alpha-\beta} \text{ if } \alpha+\beta > \frac{1}{2} \text{ and } 2\alpha+\beta \leq 1, \\ \operatorname{poly}(n) \exp\left(-C\, n^{(2\alpha+\beta-1)/2}\right) \text{ if } 2\alpha+\beta > 1, \end{cases} \tag{3.3}$$

Ch. 5 goes into detail about the history of this spectral gap problem as well as several of the techniques used in solving it.

3.1.1 Barrier Types

Most of my work uses the square barrier model. This model is nothing more than a simple rectangular bump located at $|x| = n/4$, like the one shown in Fig. 3.1. Mathematically, this barrier is given by

$$b(z) = \begin{cases} n^{\beta} & \text{when } \left(\frac{n}{4} - \frac{1}{2}c\, n^{\alpha}\right) < z < \left(\frac{n}{4} + \frac{1}{2}c\, n^{\alpha}\right) \\ 0 & \text{otherwise} \end{cases}, \tag{3.4}$$

where c is an n independent constant. Depending on the model being studied, I may set $\alpha = \beta$ or fix $c = 1$. Due to this barrier the full cost function has a global minimum at $|x| = 0$ and a local minimum at $|x| = \left\lfloor \frac{n}{4} + \frac{1}{2}cn^{\alpha} \right\rfloor + 1$.

One primary issue with this barrier, especially for low-n simulation is that the width of the barrier grows in discrete jumps due to the discrete nature of the Hamming weight. Thus, to avoid this issue in numerical simulations, I often restrict to values of n that are divisible by 4 and such that $\lfloor 1 + cn^\alpha \rfloor$ is greater than $\lfloor 1 + c(n-4)^\alpha \rfloor - \lceil 1 - \frac{1}{2}c(n-4)^\alpha \rceil$. This restriction ensures that the size of the barrier has just increased so that comparisons over different n are capturing similar behavior. This restriction becomes less relevant for larger n.

Most of my early work and analytic work uses the square barrier because it is easier to utilize; however, the square barrier exhibits more small n effects due to its discontinuous nature. The solution to this is to also use a barrier with a much smoother form. In a discrete setting, the smoothest curve is a binomial distribution, so many of my results also use binomial barriers.

The binomial distribution in a finite range can be centered at any point in the range, but this change of centering does not change the width of the distribution which is proportional to the square root of its range. Therefore, instead of defining the binomial barrier over the entire Hamming space, I define it only in a small region centered about $n/4$. To get the proper width scaling, the width of the range of the binomial distribution depends on $n^{2\alpha}$, so that the binomial barrier is defined between $\lceil \frac{n}{4} - \frac{1}{2}n^{2\alpha} \rceil$ and $\lfloor \frac{n}{4} + \frac{1}{2}n^{2\alpha} \rfloor$.

To get the height scaling correct, the binomial distribution is normalized so that $b(n/4) = n^\beta$. This distribution still suffers from the discrete width scaling from the square barrier. However, since the actual width of the distribution scales with $n^{2\alpha}$, the number of useful n increases dramatically, allowing denser computational studies with the binomial barrier. I switch to using this barrier exclusively in the later chapters due to its better small-n behavior and its smoother computational use.

3.1.2 Symmetric Subspace

The Hamiltonian, $\hat{H}(s)$, on n qubits can be simplified by considering just the symmetric subspace. For each Hamming weight $0 \leq h \leq n$, the Hamiltonian is degenerate on the subspace of $\{|x\rangle : |x| = h\}$, so there will only be one degenerate energy level for each Hamming Weight h. This symmetry can be utilized to rewrite the full $2^n \times 2^n$ Hamiltonian as an $(n+1) \times (n+1)$ symmetric Hamiltonian:

$$\begin{aligned}\hat{H}_{\mathrm{sym}}(s) = &\sum_{h=0}^{n} \left[s(h + b(h))\right] |h\rangle \langle h| \\ &- \frac{(1-s)}{2} \sum_{h=0}^{n-1} \sqrt{(h+1)(n-h)}\, |h\rangle \langle h+1| \\ &- \frac{(1-s)}{2} \sum_{h=0}^{n-1} \sqrt{(h+1)(n-h)}\, |h+1\rangle \langle h| .\end{aligned} \tag{3.5}$$

The form of this symmetric Hamiltonian can be derived by looking at the action of the full Hamiltonian on symmetric states. A careful consideration of normalization factors leads to most of the form of this Hamiltonian.

When $b(z) = 0$, the ground state of the symmetric Hamiltonian is explicitly

$$|GS_{b(z)=0}\rangle = \frac{1}{\left(2\delta(\delta+s)\right)^{n/2}} \sum_{h=0}^{n} \sqrt{\binom{n}{h}} (s+\delta)^{n-h} (1-s)^h \,|h\rangle , \tag{3.6}$$

where $\delta \equiv \sqrt{1 - 2s + 2s^2}$ is the unperturbed spectral gap. This distribution is a binomial for all s and a symmetric binomial for $s = 0$, with the width remaining proportional to $\sqrt{n}$ for $0 \leq s < 1$. The maximum amplitude $|h\rangle$ state here corresponds with the zero amplitude state in the first excited state and can be thought of as the center of the distribution. The center coincides with $h = \frac{n}{4}$ when $s = s^* \equiv \frac{1}{2}(\sqrt{3} - 1)$.

In the large n limit with a non-zero barrier, $b(z)$ becomes extremely narrow relative to the dimension of the Hilbert space, so for most s values, the energy states are unaffected by the barrier. It is only when the energy states get close to the barrier that the perturbation becomes important. Therefore, in the large n limit, the location of the minimum spectral gap becomes this critical s^*.

3.1.3 Numerical Simulation

As stated, the symmetric version of the Hamiltonian, Eq. 3.5, aids greatly in numerical simulation of this problem. This symmetric Hamiltonian captures the entire spectrum of the problem but eliminates the unneeded degeneracy. The three main numerical techniques I use are direct Schrödinger evolution, numerical diagonalization, and path-integral quantum Monte Carlo (QMC).

For direct Schrödinger evolution, I just take the Schrödinger equation

$$i\frac{d}{dt}\left|\psi\right\rangle = \hat{H}(t)\left|\psi\right\rangle \tag{3.7}$$

and interpret it as a differential equation problem for a $2(n+1)$ dimensional system, treating the real and imaginary parts as separate. This simulation is fairly straightforward, and I utilized GSL, the GNU Scientific Library, to perform most of the actual numerical integration.

For numerical diagonalization, I used LAPACK, a numerical linear algebra package. In my case, I programmed in Fortran, but LAPACK libraries are available for several languages. In addition, my purposes did not require me to know the form of the energy states, just the eigenvalues. Therefore, I saved memory usage by only calculating eigenvalues.

Path-integral QMC is an integral part of the work in this dissertation and will be covered on its own in Ch. 4.

3.2 p-Spin Model

The second model I consider works with a final potential that is cubic but without an explicit barrier. This model is the $p = 3$ case of the p-spin model that has been used by numerous groups [44, 45, 46, 47, 48, 49], and in the language of Hamming weight it is given by the final Hamiltonian

$$\hat{H}_1 = n \sum_{x \in \{0,1\}^n} \left(2\frac{|x|}{n} - 1 \right)^p |x\rangle \langle x| \,. \tag{3.8}$$

The initial Hamiltonian is still the same transverse field used for the barrier tunneling problem.

This cost function still has the all-zero bit string as its ground state, so the annealing evolution should still take the ground state from being localized around $|z| = n/2$ at $s = 0$ to $|z| = 0$ at $s = 1$. This problem does not include a barrier in the final potential but can still be visualized as a barrier tunneling problem in a semi-classical large-n limit, using such methods as the Villain transformation described in Appendix A. Notably, Jörg *et al.* [44] showed that if the exponent $p \geq 3$, then the spectral gap becomes exponentially small in this problem. Therefore, finding the ground state of this cost function through quantum adiabatic optimization is a difficult task.

Many of the useful properties from the n-qubit barrier model also carry over to this system. The final cost function remains symmetric between qubits, so symmetry simplifications can be employed to make this problem numerically tractable to solve for large n.

The cubic model is only used in Ch. 8 as another example to numerically test my analytic results.

Chapter 4

Quantum Monte Carlo Comparison

In this chapter, I outline path-integral Quantum Monte Carlo (QMC) and use it as a comparison to QAO to determine whether there is a quantum advantage for the barrier tunneling model. The results in this chapter are numerical in nature. Therefore, I run Monte Carlo simulations long enough for them to equilibrate and compare that equilibration time to the numerical spectral gap to compare the runtimes of QAO and QMC.

In Ch. 5, I derive the asymptotic scaling behavior of the spectral gap for the barrier tunneling problem, but as discussed in Ch. 6, this asymptotic behavior is often not indicative of the behavior for low numbers of qubits. Therefore, in the first section I numerically find the spectral gap for a range of n that I can simulate QMC on. In sections 4.2 & 4.3, I discuss and develop the tools of QMC. Finally, in section 4.4 I look at the actual comparison of QMC and QAO to determine whether QMC is an efficient simulator of QAO. If there is a polynomial relationship between QMC runtime and QAO runtime ($\propto g_{\min}^{-2}$), then QMC is an efficient quantum simulator of this problem, and there is no quantum advantage.

All the work in this chapter is done in the regime where the height and width of the barrier scale together so that $\alpha = \beta$. However, I do explore what happens when the constant multiplying the width of the barrier is allowed to change so that $c \neq 1$.

The work presented in this chapter has appeared previously in an article by myself and Wim van Dam [18]. This work is numerical in nature, but after its publication, it was confirmed analytically by Crosson and Harrow [23] and at non-zero temperature by Jiang *et al.* [24].

4.1 Exact Spectral Gap

Using the symmetry of the Hamiltonian for the barrier tunneling model with a square barrier, I can numerically diagonalize the Hamiltonian accurately in the same range of finite n that the Quantum Monte Carlo simulations access. As a result, I can compare the QMC run-times directly to the inverse spectral gap, rather than having to rely on extrapolations to large n behavior.

4.1.1 Numerical Results

In Fig. 4.1 I show the minimum gap for a barrier of size $\propto n^{0.5}$ as a function of n. The line drawn through the points is a linear best fit to the log-log data, and the plot below shows the residuals for this fit. Since the residuals curve downward, the gap is decreasing faster than a power law can account for; therefore, the running time for QAO, which depends on the gap $g_{\min}^{-2}$, is superpolynomial in n for $\alpha = 0.5$. In Figs. 4.2 and 4.3, I show similar plots for $\alpha = 0.4$ and 0.3 respectively.

In Fig. 4.1 I show the minimum gap for a barrier of size $n^{0.5}$ as a function of n. The line drawn through the points is a linear best fit to the log-log data, and the plot below shows the residuals for this fit. The downward curve of the residuals indicates that a linear fit is insufficient here and that some higher order, nonlinear term is necessary

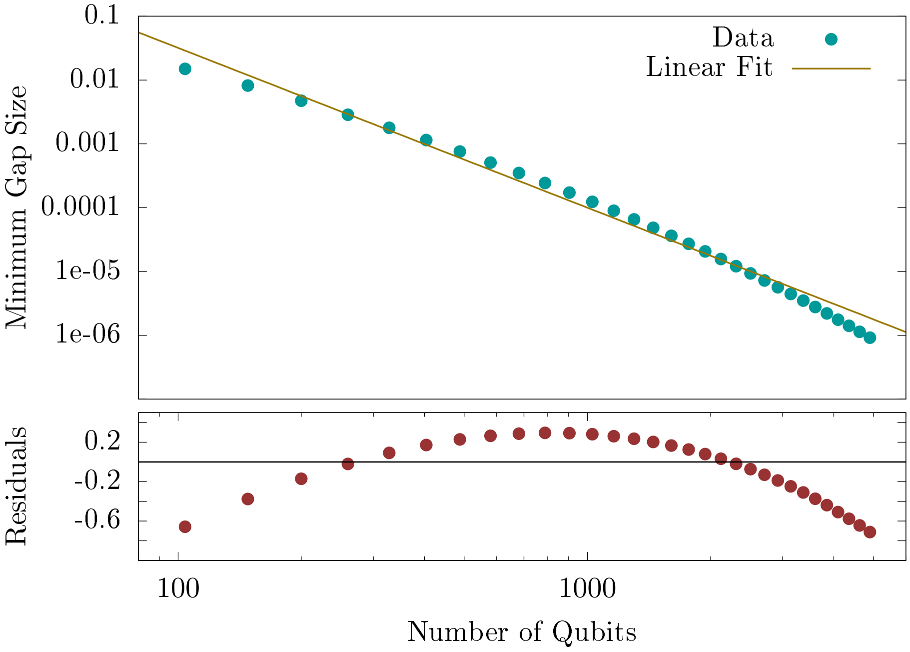

Figure 4.1: $g_{\min}$ vs. n for barrier size $n^{0.5}$: I show a best fit linear regression through the log-log data and plot the residuals of that linear fit versus the log-log data. The fact that the residuals curve down means that $g_{\min}$ is decreasing faster than a power law with n, indicating superpolynomial growth in the QAO run-time

to describe this behavior. As an example, let us approximate the residual shape to be quadratic, in which case, I would need to add a term

$$\ln g_{\min} = -m \ln n + b + a(\ln n)^2 \tag{4.1}$$

$$g_{\min} = e^b n^{-m} e^{a(\ln n)^2}, \tag{4.2}$$

for some $p > 1$. If a is negative, as it would be in my case, then, this extra term would cause the gap to decrease faster than could be accounted for by a polynomial, meaning that the running time for QAO, which depends on the gap $g_{\min}^{-2}$, is superpolynomial in n. This quadratic approximation is just an estimate; for instance, an exponentially decreasing gap would create a similar pattern. I do not claim to determine whether the

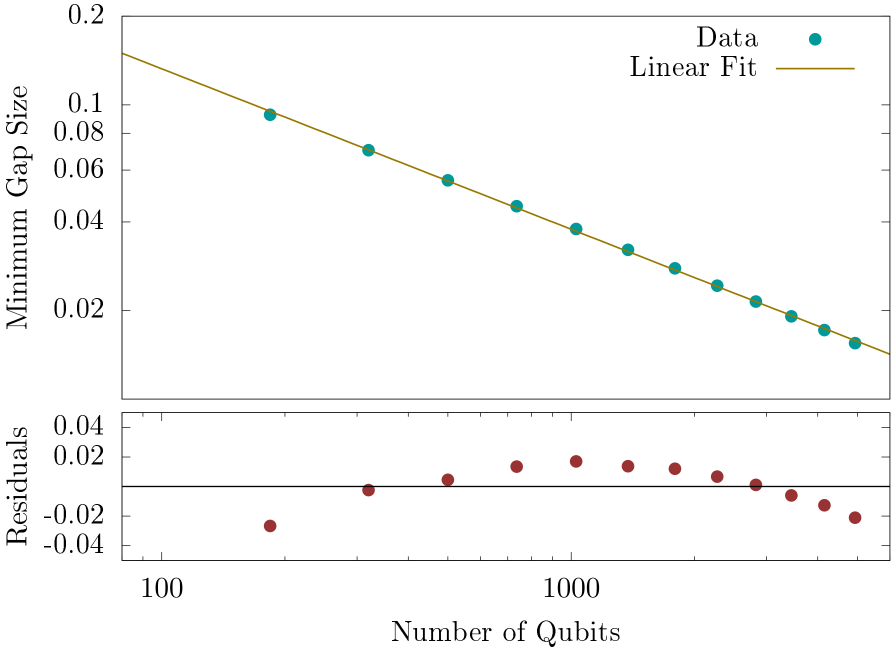

Figure 4.2: $g_{\min}$ vs. n for barrier size $n^{0.4}$: The best fit linear regression to the log-log data has residuals that curve downwards indicating superpolynomial growth in the QAO run-time. Also, notice that y-axis scale on the residual plot is much smaller than in Fig. 4.1, indicating that this scaling is not as strong as n the higher α case.

gap is truly exponentially decreasing here, merely that its decrease is superpolynomial. Furthermore, this is not an asymptotic statement but merely a statement about the local behavior of the gap in the region of n I am examining.

Therefore, I can conclude that QAO is superpolynomial for $\alpha = 0.5$ for low n such as those shown here. In Figs. 4.2 and 4.3, I show similar plots for $\alpha = 0.4$ and 0.3 respectively.

Varying α, I do the same procedure, sweeping through a range of n from 100 to 5000 with $c = 1$. The second derivative of these log-log residuals can be used to estimate the curvature of those residuals (i.e. whether they are concave up or down), and these second derivatives are plotted in Fig. 4.4. This tells us just what the second order correction to the fit should be and therefore gives me information on whether there is

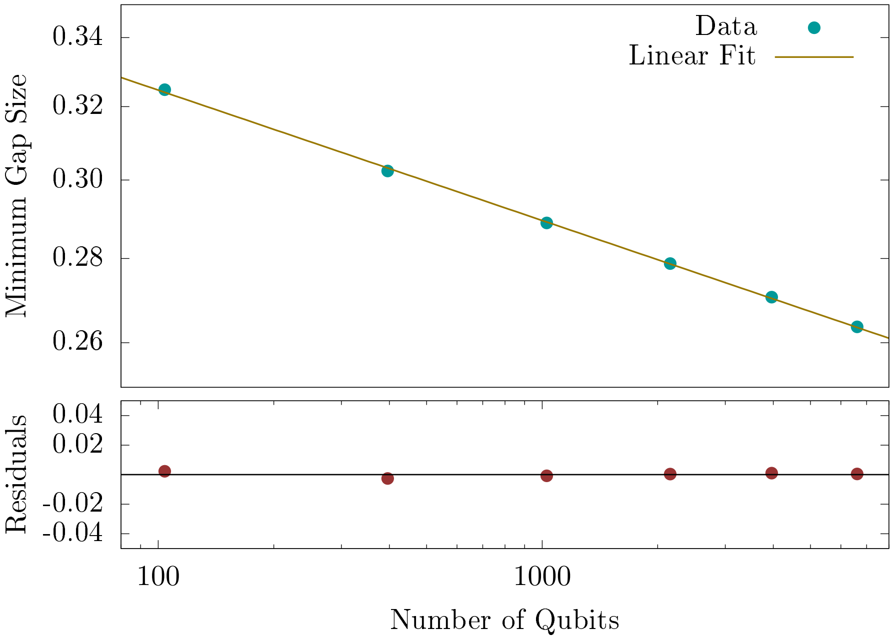

Figure 4.3: $g_{\min}$ vs. n for barrier size $n^{0.3}$: The best fit linear regression to the log-log data has residuals that are essentially zero, indicating polynomial scaling with n. I have used the same residual scale as in Fig. 4.2 to indicate just how small these residuals are.

some superpolynomial effect, but there could be higher order terms, leading possibly even to exponentially decreasing gaps. The residual plots in Figs. 4.1, 4.2, and 4.3 are all used in the construction of Fig. 4.4. Since the second derivative varies over the range of n, the second derivative is averaged and the standard error is used as the error bars. A negative number indicates superpolynomial running time, whereas zero represents polynomial scaling.

The curvature in Fig. 4.4 becomes negative by more than one error bar starting at $\alpha = 0.34$, which indicates that the quantum adiabatic algorithm undergoes a transition from polynomial to exponential scaling somewhere between $\alpha = 0.33$ and 0.34.

It is a folklore result [25] by Goldstone that as n grows for a barrier with height and

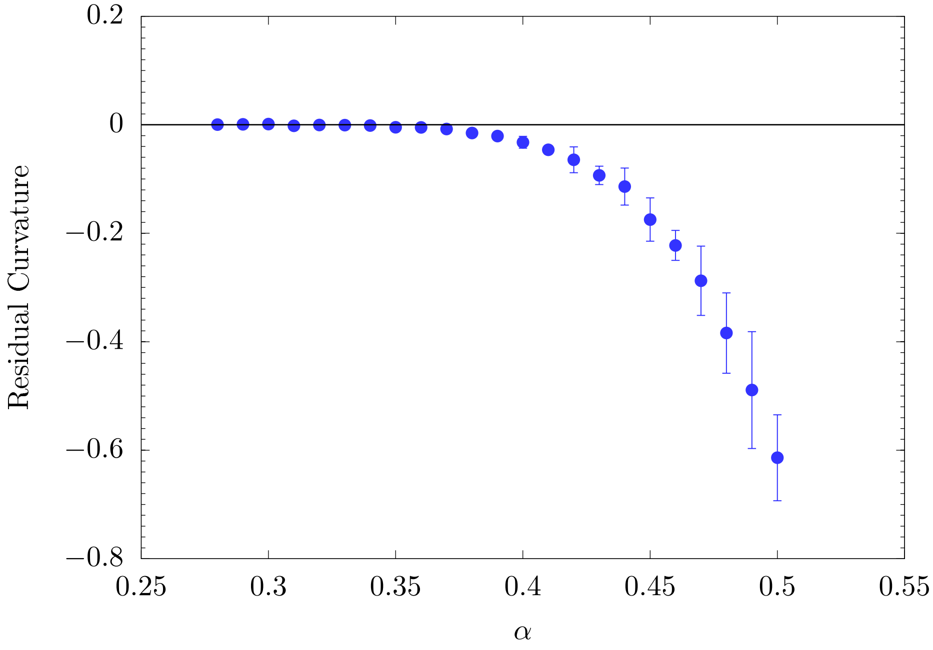

Figure 4.4: Deviation of g_{min} from Power Law in n: For each barrier scaling power α at $c = 1$ I find the spectral gap for n between 100 and 5000 and tried to fit a linear curve to the log-log plot of spectral gap versus n. What is displayed here is the curvature of the residuals from those fits. If the residuals are concave down (meaning negative curvature on this figure), the spectral gap is decreasing faster than a power law in n. Therefore, QAO becomes superpolynomial in n somewhere between $\alpha = 0.33$ and 0.34.

width proportional to n^α, the spectral gap decreases asymptotically as:

$$g_{\text{min}} = \begin{cases} \text{constant} & \text{if } \alpha < \frac{1}{4} \\ 1/\text{polynomial}(n) & \text{if } \frac{1}{4} < \alpha < \frac{1}{3} \\ 1/\text{exponential}(n) & \text{if } \frac{1}{3} < \alpha \end{cases} \tag{4.3}$$

Hence I expect a transition from polynomial to exponentially small gaps to occur when $\alpha = 1/3$. The existence of constant scaling for $\alpha < 1/4$ has been proven rigorously by Reichardt [42]. More recently, the exponential scaling region for $\alpha > 1/3$ was proven by Kong and Crosson [43], and Wim van Dam and I [19] (Ch. 5) proved the n dependence

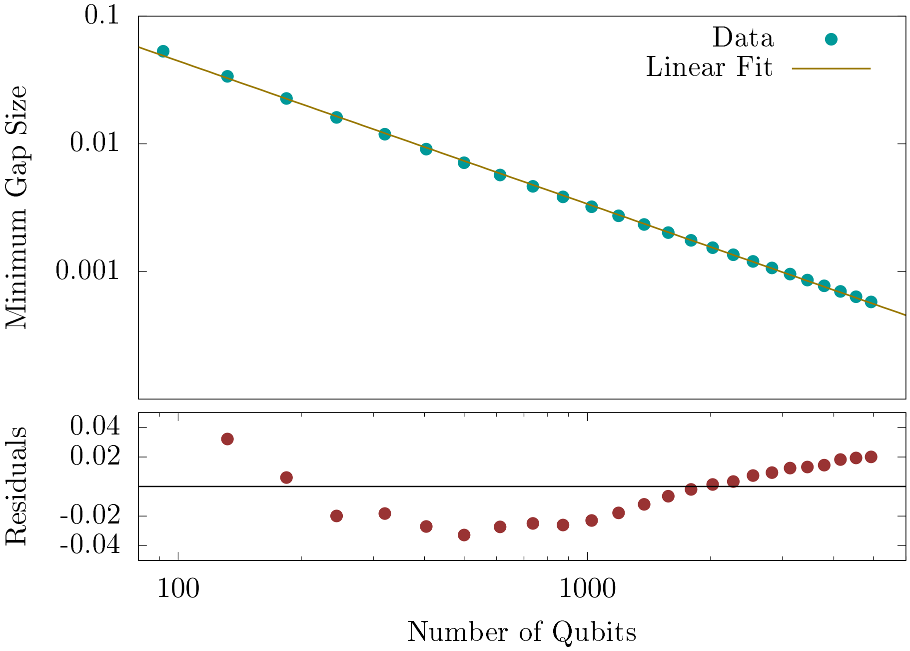

Figure 4.5: $g_{\min}$ vs. n for barrier size $2n^{0.4}$: The best fit linear regression to the log-log data has residuals that curve upwards indicating polynomial or subpolynomial decrease with n. At the end of the n range, the residuals begin to curve down again, indicating the beginning of the superpolynomial region indicated by Figs. 4.4 and 4.2.

of the gap in the polynomial scaling region between $\alpha > 1/4$ and $\alpha < 1/3$. The cutoffs in Eq. 4.3 mesh exactly with the numerical results for when the gap is polynomial versus when it transitions to superpolynomial behavior.

The numerical results are still useful in their own rights since the QMC calculations access finite n values and it is important to compare the QMC results with equivalent gap results, and I need these results to be aware of any possible small n phenomena.

Additional numerical results indicate that the large n scaling behavior in Fig. 4.4 does not hold for smaller n when c is large. For instance Fig. 4.2 does display the large n superpolynomial behavior, but Fig. 4.5 does not. In Fig. 4.5, if I consider just the largest n, there are indications that the residuals are becoming concave down at the end, indicating that the superpolynomial scaling is starting at the end of the n range considered.

The computational limits of the QMC algorithm and computing facilities mean that some of the QMC simulations in this chapter are at lower n where the large n scaling behavior is not yet dominant. In cases where I can access the large n scaling behavior, such as $\alpha = 0.4$ and $c = 1$ in Fig. 4.2, I mention so in subsequent analysis. Largely, I compare QMC running times with $g_{\min}^{-2}$ directly so that I can see if QMC running time scales polynomially with QAO running time.

4.2 Path-integral Quantum Monte Carlo

The path-integral QMC algorithm [35] is a method of simulating a quantum mechanical system at finite inverse temperature β. The procedure uses the Trotter expansion to take an n qubit quantum system to a classical system of n bits evolving in a discretized "imaginary time" dimension. These time evolving states can then be treated as states in a Monte Carlo simulation that samples possible paths of the system.

The Monte Carlo algorithm then picks paths with probability proportional to their Boltzmann weights, so from these states, an expectation value for the ground state energy can be obtained. I run the Monte Carlo algorithm for fixed s until I reach the ground state at that s value and then transition to a new s. This so called annealing schedule captures the same adiabaticity that makes QAO so powerful.

4.2.1 Trotter Expansion

To start, I take the partition function for the Hamiltonian at finite inverse temperature β and Trotter expand it into T "time"-slices

$$Z = \mathrm{Tr}\left\{e^{-\beta\hat{H}}\right\} = \lim_{T\to\infty} \sum_{x^{(0)},\ldots,x^{(T-1)}} \left[\prod_{\xi=0}^{T-1} \langle x^{(\xi)}| \, e^{-\frac{\beta}{T}\hat{H}} \, |x^{(\xi+1)}\rangle\right], \tag{4.4}$$

where the sums go over each $x^{(\xi)} \in \{0,1\}^n$. In order to be in the ground state, the temperature needs to be low, which means high β, but T also needs to be much greater than β in order for the Trotter approximation to work well. In practice, I take $\beta = 32$ and $T \propto n$ for reasons that are discussed in subsection 4.2.2. I also have periodic boundary conditions $x^{(0)} = x^{(T)}$ as dictated by the Trotter decomposition. The goal is to have the operators act on these $|x\rangle$ basis states so that I can get a partition function in terms of c-numbers rather than quantum operators. Each of the T bases corresponds to a different imaginary "time" slice of the system, so I am transforming my n qubit system into an $n \times T$ lattice of classical bits with interactions between adjacent time slices.

4.2.2 Exponential Approximation

The Hamiltonian includes terms diagonal in the computational basis, which I call $\hat{H}_d$, and off-diagonal terms, which I call $\hat{H}_o$. The goal is to separate out these terms so that each operator can act on its own eigenbasis. I consider two approximations that can be used here: either a linear approximation or an exponential approximation for $\beta/T \to 0$:

$$e^{-\frac{\beta}{T}(\hat{H}_d+\hat{H}_o)} = 1 - \frac{\beta}{T}(\hat{H}_d + \hat{H}_o) + O((\beta/T)^2) \tag{4.5a}$$

$$e^{-\frac{\beta}{T}(\hat{H}_d+\hat{H}_o)} = e^{-\frac{\beta}{T}\hat{H}_d} e^{-\frac{\beta}{T}\hat{H}_o} + O((\beta/T)^2). \tag{4.5b}$$

To first order these are both the same, but the additional terms in the exponential change the algorithm significantly. The linear approximation only includes one copy of the off-diagonal Hamiltonian, so adjacent Trotter time slices would differ by at most a single bit. Single bit flips between adjacent sites lend a nice sense of continuity to the time dimension, but they necessitate larger T. The off-diagonal part of the Hamiltonian manifests itself in the simulation as bit flips between adjacent time-slices, so in order to get enough bit flips in the linear approximation, T must be larger, whereas the exponential approximation, with multiple adjacent bit flips, can be more compact.

In Fye [50], there is a discussion of these two approximation methods where they find that for local Hamiltonians the exponential approximation is more robust and can be used with an n-independent T. The linear approximation requires T to increase with increasing n, making it less desirable. The barrier-tunneling Hamiltonian relies on the Hamming Weight, which is a non-local quantity, so these results do not hold perfectly. I found that the exponential approximation did require T to have some dependence on n; however, numerically, I found that dependence to be much smaller than the dependence of the linear approximation. Therefore, I use the exponential approximation in this article.

Eventually, I want to interpret the product of these exponentials as a Boltzmann factor or probability for the given $n \times T$ configuration of the system. In order to do this, the Boltzmann factors must be positive. In order to ensure that the approximated exponentials remain positive, the Hamiltonian must be one with "no sign problem." This means that all the off-diagonal terms in the Hamiltonian must be non-positive. To see why, consider Eq. 4.5a; if the off-diagonal Hamiltonian contained negative terms, then this operator would lead to negative terms if it were between non-identical states. This same logic is true in Eq. 4.5b. The Hamiltonian has no sign problem, so I am free to use these methods.

As discussed in Ch. 2, it is unclear how much quantum computational power Hamiltonians without a sign problem have. These so called stoquastic Hamiltonians can be simulated with QMC by utilizing the positive Boltzmann factors as probabilities. Although stoquastic Hamiltonians can be simulated by path-integral QMC, there is no theoretical guarantee that this simulation is efficient.

4.2.3 Final Partition Function

For an in depth derivation of the partition function see section 4.3. Here, I just cite the resulting partition function

$$
\begin{aligned}
Z = \lim_{T\to\infty} \sum_{x^{(0)},\ldots,x^{(T-1)}} \Bigg[\prod_{\xi=0}^{T-1} & e^{-\frac{\beta}{T}\left((1-s)\frac{n}{2}+sf(|x^{(\xi)}|)\right)} \\
\times \prod_{d=1}^{n} & \left(e^{\frac{\beta}{T}\frac{(1-s)}{2}} + (-1)^{x_d^{(\xi)}-x_d^{(\xi+1)}} e^{-\frac{\beta}{T}\frac{(1-s)}{2}} \right) \Bigg] .
\end{aligned}
\tag{4.6}
$$

The first summation can be thought of as a sum over possible states, where a state is a full configuration of the $n \times T$ bit lattice. The expression in the square brackets is the Boltzmann factor for that configuration. The Boltzmann factors are the unnormalized probabilities for the states, so they can be used in a Metropolis algorithm to create a Monte Carlo simulation. The Quantum Monte Carlo method consists of performing standard Monte Carlo methods on this classical partition function which can then be used to gain information about the original quantum system (e.g. see subsection 4.3.1 for how to extract the energy from this Monte Carlo simulation).

4.2.4 Update Rules

I follow the same update rule as Crosson and Deng [40], sweeping through the $n \times T$ classical bits. One sweep consists of $n \times T$ updates, where I go through each bit in the lattice separately. For that bit I try flipping its value, and then compare the Boltzmann weight of the lattice before and after the bit-flip. The acceptance rate of this bit flip is then equal to the ratios of the Boltzmann factors before and after the flip. Once the sweep has gone through every bit in the lattice, the sweep ends, and the algorithm calculates the current ground state energy of the entire lattice based on the results of subsection 4.3.1.

For the annealing schedule, I have a fixed $\Delta s = \frac{1}{100}$ and change how much time I spend on each s value. The algorithm calculates the quantum mechanical energy (see subsection 4.3.1) of the system after each sweep and moves onto the next s value when the energy gets close enough to the true ground state energy. The algorithm then starts over again at $s+\Delta s$, using the ending configuration of the system at s as the new starting configuration at $s+\Delta s$.

This annealing schedule does use information that the QMC algorithm would not have in a normal simulation (namely the ground state energy and spectral gap), but since my goal is to judge how long it takes to reach the ground state rather than how long it takes the algorithm to realize it has reached the ground state, this is appropriate.

The algorithm judges it is close enough to the true ground state when the average energy over the last 100 sweeps, $\langle E(s)\rangle_{100}$, is within 0.4 spectral gaps, $g(s)$, of the true ground state energy, $E_{GS}(s)$:

$$\frac{|\langle E(s)\rangle_{100} - E_{GS}(s)|}{g(s)} < 0.4 \tag{4.7}$$

In subsequent graphs, I report the number of sweeps for each s value. If the algorithm has already satisfied this condition after the first 100 steps, it extrapolates back to when it first met the update condition and reports that as the number of sweeps.

In Fig. 4.6, the results are shown for simulations using a barrier of size $3n^{0.5}$, and

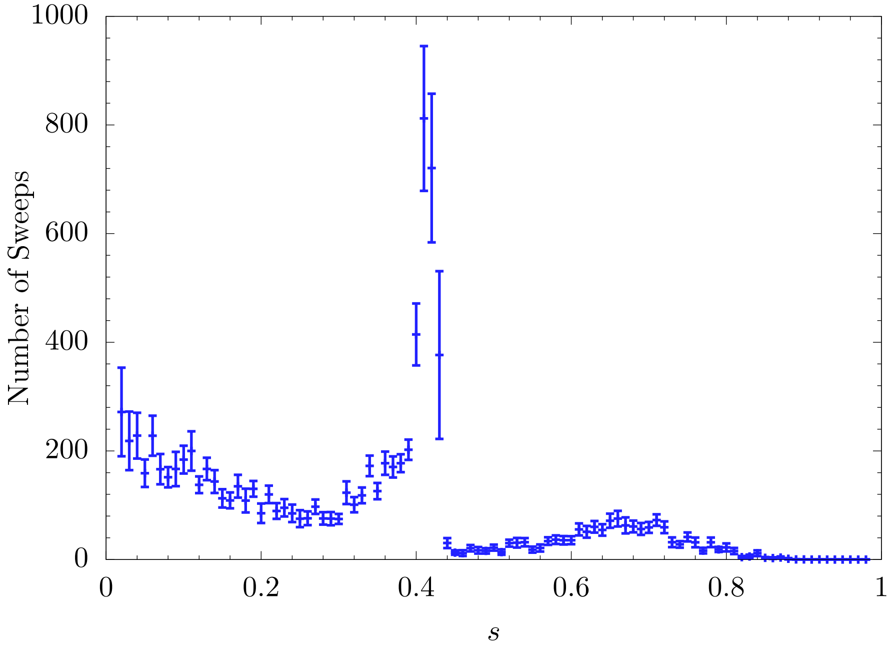

Figure 4.6: QMC Sweeps vs. s for barrier size $3n^{0.5}$ at $n = 116$: This is averaged over 30 simulations. The spike is roughly located at the s value where QAO would undergo tunneling in this problem. This correspondence between the QMC run-time spike and the QAO minimum gap, indicates that there may be a similar process occurring in both algorithms.

$n = 116$. The spike in run-time here corresponds roughly to s^*, the location of the minimum spectral gap, and this correspondence becomes closer at larger n. This spike in difficulty for the QMC algorithm right at the time when QAO undergoes tunneling indicates that there may be a process going on in the QMC algorithm that is analogous to QAO tunneling.

In the subsequent sections when I report the run-time of the QMC simulations, I report the total number of sweeps taken between $s = 0.3$ and $s = 0.5$. For all of my simulations, this s range captures the run-time spike and some of the surrounding area while ignoring any low s initialization artifacts or high s tailing-off.

4.3 QMC Partition Function

I start with Eq. 4.4, and my goal is to derive Eq. 4.6 as well as an estimator for the quantum mechanical ground state energy. My first step involves inserting the exponential approximation scheme so that

$$Z = \lim_{T\to\infty} \sum_{x^{(0)},\ldots,x^{(T-1)}} \left[\prod_{\xi=0}^{T-1} \langle x^{(\xi)} | \, e^{-\frac{\beta}{T}\hat{H}_d} e^{-\frac{\beta}{T}\hat{H}_o} \, | x^{(\xi+1)} \rangle \right], \tag{4.8}$$

where the sums go over each $x^{(\xi)} \in \{0,1\}^n$.

Here $\hat{H}_o$ and $\hat{H}_d$ are the off-diagonal and diagonal parts of the Hamiltonian, given by

$$\begin{aligned} \hat{H}_d &\equiv \sum_{x\in\{0,1\}^n} \left[(1-s)\frac{n}{2} + s(|x| + b(|x|) \right] |x\rangle \langle x| \\ \hat{H}_o &\equiv \sum_{\langle x,y\rangle} \left[-\frac{(1-s)}{2} \right] |x\rangle \langle y| \, . \end{aligned}$$

The sum in $\hat{H}_o$ is over nearest neighbor sites (i.e. bit strings x and y that differ by one bit flip). Since $\hat{H}_d$ is diagonal in the computational basis, I can just have it act on the basis states pulling out the eigenvalues $H_d(x) = (1-s)\frac{n}{2} + s(|x| + b(|x|))$.

$$\begin{aligned} Z = \lim_{T\to\infty} \sum_{x^{(0)},\ldots,x^{(T-1)}} & \left[\prod_{\xi=0}^{T-1} e^{-\frac{\beta}{T}H_d(x^{(\xi)})} \right] \\ \times & \left[\prod_{\xi=0}^{T-1} \langle x^{(\xi)} | \, e^{-\frac{\beta}{T}\hat{H}_o} \, | x^{(\xi+1)} \rangle \right]. \end{aligned} \tag{4.9}$$

Next, I claim that there is an orthonormal basis $|k^{(\xi)}\rangle$ that is the eigenbasis for $\hat{H}_o$, whose eigenvalues are $H_o(k^{(\xi)})$. I can insert a complete set of these states at ever time slice to get

$$Z = \lim_{T\to\infty} \left[\sum_{x^{(0)},\ldots,x^{(T-1)}} \sum_{k^{(0)},\ldots,k^{(T-1)}} \right] \left[\prod_{\xi=0}^{T-1} e^{-\frac{\beta}{T} H_d(x^{(\xi)})} e^{-\frac{\beta}{T} H_o(k^{(\xi)})} \langle x^{(\xi)} | k^{(\xi)} \rangle \langle k^{(\xi)} | x^{(\xi+1)} \rangle \right]. \tag{4.10}$$

To find these $|k\rangle$ states, I just need to diagonalize $\hat{H}_o$. This operator can be represented by a translationally invariant matrix on an n dimensional hypercubic lattice (where each dimension is two sites long) with periodic boundary conditions and nearest neighbor interactions. These properties mean that the eigenstates of $\hat{H}_o$ are simply the Brillouin Zone lattice sites. If I represent each Brillouin Zone lattice site using $k \in \{0,1\}^n$, then these lattice sites can be represented in the $|x\rangle$ basis by

$$|k\rangle = \sum_{x\in\{0,1\}^n} e^{i\pi\vec{k}\cdot\vec{x}} |x\rangle . \tag{4.11}$$

Using standard Brillouin Zone methods for translationally invariant matrices, you can work out that the eigenvalues of the off-diagonal Hamiltonian are

$$H_o(k) = -\frac{(1-s)}{2} \sum_{d=1}^{n} (1 - 2k_d). \tag{4.12}$$

Furthermore, the overlap between $|x\rangle$ and $|k\rangle$ states is given by

$$\langle x|k\rangle = (-1)^{\vec{k}\cdot\vec{x}}. \tag{4.13}$$

Inserting Eqs. 4.13 and 4.12 back into the partition function gives

$$\begin{aligned} Z = \lim_{T\to\infty} \sum_{x^{(0)},\ldots,x^{(T-1)}} & \left[\prod_{\xi=0}^{T-1} e^{-\frac{\beta}{T}H_d(x^{(\xi)})}\right] \\ \times & \left[\prod_{\xi=0}^{T-1}\sum_{k^{(\xi)}}\prod_{d=1}^{n} e^{\frac{\beta}{T}\frac{(1-s)}{2}(1-2k_d^{(\xi)})}(-1)^{k_d^{(\xi)}(x_d^{(\xi)}-x_d^{(\xi+1)})}\right]. \end{aligned} \tag{4.14}$$

I can rewrite $\sum_{k^{(\xi)}}\prod_{d=1}^{n} \to \prod_{d=1}^{n}\sum_{k_d^{(\xi)}=0,1}$. Focusing on just the important part and dropping the ξ labels in favor of labeling the two bit strings by x and y, I get

$$\begin{aligned} & \prod_{d=1}^{n}\sum_{k_d=0,1} e^{\frac{\beta}{T}\frac{(1-s)}{2}(1-2k_d)}(-1)^{k_d(x_d-y_d)} \\ & = \prod_{d=1}^{n}\left[e^{\frac{\beta}{T}\frac{(1-s)}{2}} + (-1)^{x_d-y_d}e^{-\frac{\beta}{T}\frac{(1-s)}{2}}\right]. \end{aligned} \tag{4.15}$$

Note that I have now eliminated the k variables entirely. Inserting this simplification recovers Eq. 4.6:

$$\begin{aligned} Z = \lim_{T\to\infty} \sum_{x^{(0)},\ldots,x^{(T-1)}} & \left[\prod_{\xi=0}^{T-1} e^{-\frac{\beta}{T}\left((1-s)\frac{n}{2}+sf(|x^{(\xi)}|)\right)}\right. \\ \times & \left.\prod_{d=1}^{n}\left(e^{\frac{\beta}{T}\frac{(1-s)}{2}} + (-1)^{x_d^{(\xi)}-x_d^{(\xi+1)}}e^{-\frac{\beta}{T}\frac{(1-s)}{2}}\right)\right]. \end{aligned} \tag{4.16}$$

4.3.1 Derivation of Energy Estimators

Next, I need to look at what the expectation value of a quantum operator is in the Trotter expanded formalism. By definition, I have

$$\left\langle \hat{O} \right\rangle = \frac{1}{Z}\mathrm{Tr}\left\{\hat{O}e^{-\beta\hat{H}}\right\}. \tag{4.17}$$

When I do the Trotter expansion I do not and should not expand $\hat{O}$ as we do the exponential. In fact after the Trotter expansion, I still only have one copy of $\hat{O}$, so the original copy of $\hat{O}$ is just with one of the time slices. For convenience, I put it with the very first time slice, so that after Trotterization:

$$
\begin{aligned}
\left\langle \hat{O} \right\rangle &= \lim_{T\to\infty} \frac{1}{Z} \sum_{x^{(0)},\ldots,x^{(T-1)}} \langle x^{(T-1)}| \, e^{-\frac{\beta}{T}\hat{H}_d} e^{-\frac{\beta}{T}\hat{H}_o} \hat{O} \, |x^{(0)}\rangle \\
&\times \left[\prod_{\xi=0}^{T-2} \langle x^{(\xi)}| \, e^{-\frac{\beta}{T}\hat{H}_d} e^{-\frac{\beta}{T}\hat{H}_o} \, |x^{(\xi+1)}\rangle \right], \\
\left\langle \hat{O} \right\rangle &= \lim_{T\to\infty} \frac{1}{Z} \sum_{x^{(0)},\ldots,x^{(T-1)}} \frac{\langle x^{(T-1)}| \, e^{-\frac{\beta}{T}\hat{H}_d} e^{-\frac{\beta}{T}\hat{H}_o} \hat{O} \, |x^{(0)}\rangle}{\langle x^{(T-1)}| \, e^{-\frac{\beta}{T}\hat{H}_d} e^{-\frac{\beta}{T}\hat{H}_o} \, |x^{(0)}\rangle} \\
&\times \left[\prod_{\xi=0}^{T-1} \langle x^{(\xi)}| \, e^{-\frac{\beta}{T}\hat{H}_d} e^{-\frac{\beta}{T}\hat{H}_o} \, |x^{(\xi+1)}\rangle \right],
\end{aligned} \tag{4.18}
$$

Next consider the probability of obtaining a specific configuration, $\{x^{(\xi)}\}$, of the $n \times T$ lattice of bits:

$$
p\left(\{x^{(\xi)}\}\right) \equiv \frac{1}{Z} \left[\prod_{\xi=0}^{T-1} \langle x^{(\xi)}| \, e^{-\frac{\beta}{T}\hat{H}_d} e^{-\frac{\beta}{T}\hat{H}_o} \, |x^{(\xi+1)}\rangle \right] \tag{4.19}
$$

Using Eq. 4.19, the average becomes

$$
\begin{aligned}
\left\langle \hat{O} \right\rangle = \lim_{T\to\infty} & \sum_{x^{(0)},\ldots,x^{(T-1)}} p\left(\{x^{(\xi)}\}\right) \\
& \times \frac{\langle x^{(T-1)}| \, e^{-\frac{\beta}{T}\hat{H}_d} e^{-\frac{\beta}{T}\hat{H}_o} \hat{O} \, |x^{(0)}\rangle}{\langle x^{(T-1)}| \, e^{-\frac{\beta}{T}\hat{H}_d} e^{-\frac{\beta}{T}\hat{H}_o} \, |x^{(0)}\rangle}.
\end{aligned} \tag{4.20}
$$

The QMC method specifically uses the average energy: $\left\langle \hat{H} \right\rangle = \left\langle \hat{H}_d \right\rangle + \left\langle \hat{H}_o \right\rangle$. Starting with $\left\langle \hat{H}_d \right\rangle$, the operator is already acting on its eigenstates, so the average becomes

$$
\left\langle \hat{H}_d \right\rangle = \lim_{T\to\infty} \sum_{x^{(0)},\ldots,x^{(T-1)}} \left[H_d(x^{(0)}) p\left(\{x^{(\xi)}\}\right) \right]. \tag{4.21}
$$

In actual simulations, the estimator $H_d(x^{(0)}) \rightarrow \frac{1}{T}\sum_{\xi=0}^{T-1} H_d(x^{(\xi)})$ is used so that information from the entire time dimension can enter the statistics.

Moving onto $\left\langle \hat{H}_o \right\rangle$ and focusing on just the relevant piece I have (replacing $x^{(T-1)} \rightarrow x$ and $x^{(0)} \rightarrow y$ for notational convenience):

$$\frac{\langle x| e^{-\frac{\beta}{T}\hat{H}_d} e^{-\frac{\beta}{T}\hat{H}_o} \hat{H}_o |y\rangle}{\langle x| e^{-\frac{\beta}{T}\hat{H}_d} e^{-\frac{\beta}{T}\hat{H}_o} |y\rangle},$$

I can insert k resolutions of the the identity in the top and bottom to get

$$\begin{aligned} &\frac{\sum_{k\in\{0,1\}^n} e^{-\frac{\beta}{T}H_o(k)} H_o(k)\langle x|k\rangle\langle k|y\rangle}{\sum_{k'\in\{0,1\}^n} e^{-\frac{\beta}{T}H_o(k')}\langle x|k'\rangle\langle k'|y\rangle} \\ =&\frac{-\frac{(1-s)}{2}\sum_{k\in\{0,1\}^n} e^{\frac{\beta}{T}\frac{(1-s)}{2}\sum_{d=1}^n(1-2k_d)}\sum_{p=1}^n(1-2k_p)(-1)^{k\cdot(x-y)}}{\sum_{k'\in\{0,1\}^n} e^{\frac{\beta}{T}\frac{(1-s)}{2}\sum_{d=1}^n(1-2k'_d)}(-1)^{k'\cdot(x-y)}} \end{aligned} \tag{4.22}$$

Next, I pull out what I can and switch $\sum_{k\in\{0,1\}^n}\prod_{d=1}^n \rightarrow \prod_{d=1}^n \sum_{k_d=0,1}$:

In a given p element, the term in the product is the same in the numerator and denominator if $d \neq p$, so the terms in the product cancel except in the case where $d = p$:

$$-\frac{(1-s)}{2}\sum_{p=1}^n \frac{e^{\frac{\beta}{T}\frac{(1-s)}{2}} - (-1)^{(x_p-y_p)} e^{-\frac{\beta}{T}\frac{(1-s)}{2}}}{e^{\frac{\beta}{T}\frac{(1-s)}{2}} + (-1)^{(x_p-y_p)} e^{-\frac{\beta}{T}\frac{(1-s)}{2}}} \tag{4.23}$$

Inserting Eq. 4.23 into the off-diagonal energy estimator gives

$$\begin{aligned} \left\langle \hat{H}_o \right\rangle = &\lim_{T\rightarrow\infty} \sum_{x^{(0)},\dots,x^{(T-1)}} p\left(\{x^{(\xi)}\}\right) \\ &\times\left[-\frac{(1-s)}{2}\sum_{p=1}^n \frac{e^{\frac{\beta}{T}\frac{(1-s)}{2}} - (-1)^{(x_p^{(0)}-x_p^{(T-1)})} e^{-\frac{\beta}{T}\frac{(1-s)}{2}}}{e^{\frac{\beta}{T}\frac{(1-s)}{2}} + (-1)^{(x_p^{(0)}-x_p^{(T-1)})} e^{-\frac{\beta}{T}\frac{(1-s)}{2}}}\right] \end{aligned} \tag{4.24}$$

Again, I typically average over the result for the different time slices in the actual simulation.

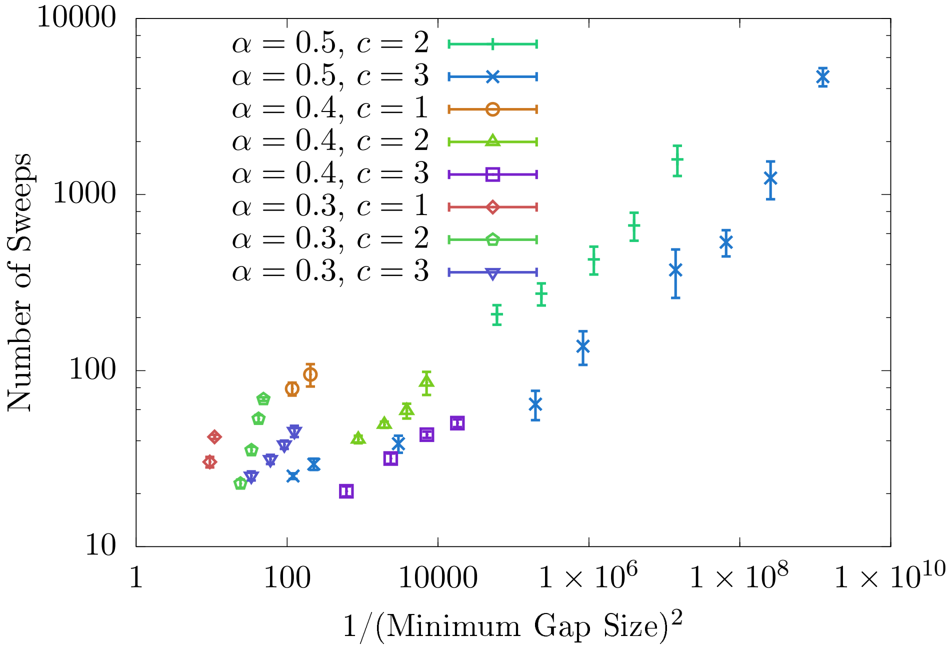

Figure 4.7: QMC Sweeps vs. $g_{\min}$: This lists all my data together; a further breakdown of this data is available in Figs. 4.8, 4.9, and 4.10. Notice that there is an obvious strong correlation between required and sufficient QMC sweeps and the gap. More analysis, specific to the different α values can be found in Figs. 4.8, 4.9.

4.4 Numerical Monte Carlo Results

In this section, I explore a few different values of the barrier scaling power, α, and the width scaling coefficient, c, using the QMC methods developed in the previous sections. For most of the simulations considered here the number of Trotter slices is related to the number of qubits through $T = 4n$. In reporting QMC times, I report the number of sweeps each simulation took while going through the critical s region. There are $n \cdot T$ Metropolis steps per sweep, so the actual run-time of the algorithm depends polynomially on the number of sweeps.

In Fig. 4.7, I show the full results of the QMC simulations, comparing the run-times of these algorithms to the corresponding $g_{\min}^{-2}$. There is a strong correlation between these two quantities, which at least indicates some relation. The following sections breakdown this data by α value and analyze it independently.

4.4.1 Barriers Proportional to $n^{0.5}$

To start, I focus on $\alpha = 0.5$. Based on the spectral gap scaling discussed in Chs. 2 & 5, this size of barrier has QAO run-times that scale superpolynomially with n. Practically, I am able to run QMC simulations with n ranging up to ~ 220 qubits. For this regime of n, small n effects mask the superpolynomial scaling of the gap for $c = 3$ but not for $c = 2$.

Note that $c = 2$ leads to smaller spectral gaps than $c = 3$ at fixed n. From trial and error, I found that the smaller gap sizes mean that the Trotter approximation needs to be better in order to get sensible results. Thus, for $c = 2$, $T = 16n$ rather than the usual $T = 4n$. This necessity to improve the QMC for simulations with smaller gap sizes lends significant credence to the idea that the QMC algorithm depends heavily on the spectral gap itself.

The QMC run-time (averaged over multiple simulations) as a function of g_{min}^{-2} is shown in Fig. 4.8 Notice that the data in this figure does not lie along a straight line, so the QMC run-times seem to be increasing at a rate faster than polynomially in the inverse gap. This lack of a power law could be caused by three possible effects.

It is possible that this means the QMC algorithm does indeed scale superpolynomially with g_{min}^{-2}. An alternative is that this curvature is due to small n effects that are still prevalent even for n in the several hundreds. Especially for $c = 3$ and lower n, there is overlap between the initial $s = 0$ ground state distribution and the barrier, which could account for the apparent deviation from a power law here. Additionally, this curvature could be an indication of deficiencies in my QMC implementation specifically. As discussed in the next section, this algorithm has some notable approximations and simplifications

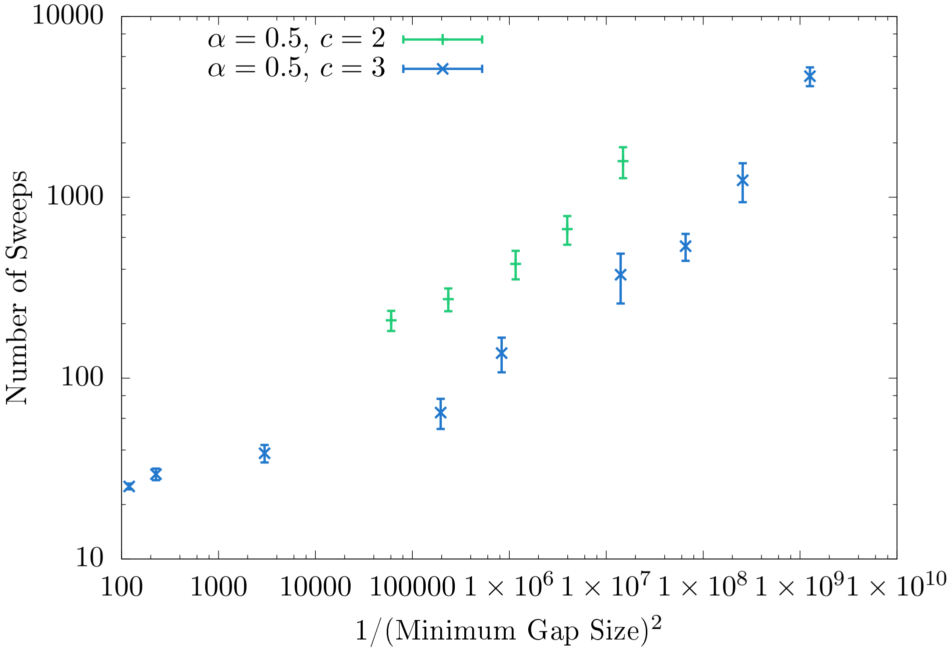

Figure 4.8: QMC Sweeps vs. $g_{\min}$ for barrier size $cn^{0.5}$: The number of sweeps is increasing faster than a power law with the inverse gap size, indicating the specific QMC algorithm is worse than QAO in this case. For $c = 2$ (green plus), n ranges from 84 to 172, and for $c = 3$ (blue cross), n ranges from 88 to 216.

that could be leading to this discrepancy.

4.4.2 Barriers Proportional to $n^{0.4}$

For $\alpha = 0.4$, the QMC simulations are able to go up to ~ 320 qubits. In this regime of n, small n effects mean that the gap is not superpolynomial for $c = 3, 2$ (see Fig. 4.5) but it is for $c = 1$ (see Fig. 4.2). In Fig. 4.9 I have compared the QMC run-times directly to the spectral gap. Notice that in this case, there does seem to be a linear relationship between the log-log data. Many of the deficiencies in my specific implementation are less pronounced in this case than in the $\alpha = 0.5$ case since the barrier is smaller. There is less overlap between the initial ground state and the barrier, which could also mean these simulations suffer less from small n effects than the $\alpha = 0.5$ simulations.

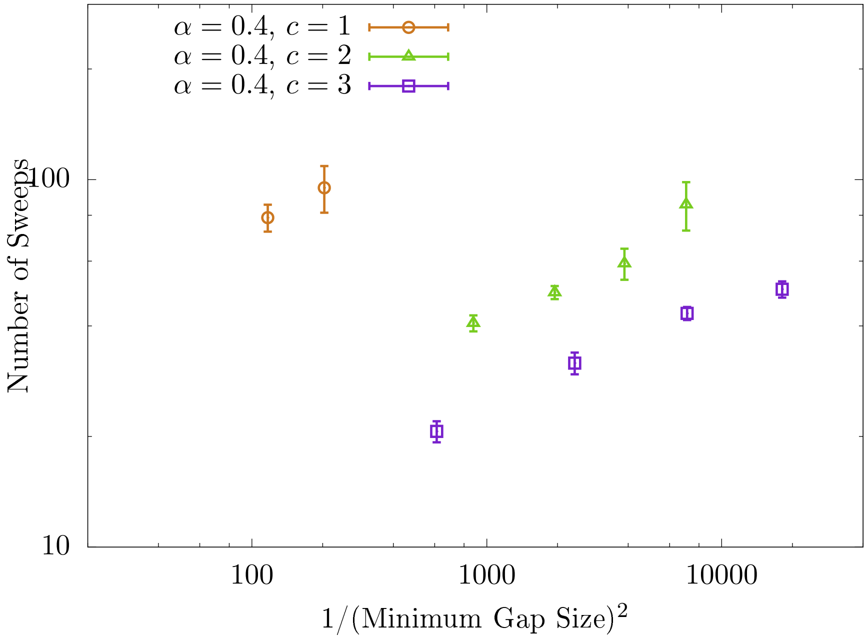

Figure 4.9: QMC Sweeps vs. g_{min} for barrier size $cn^{0.4}$: There appears to be a linear relationship here, indicating that QMC performance and QAO performance are polynomially related in this region. For $c = 1$ (red circle), n ranges from 184 to 320, for $c = 2$ (green triangle), n ranges from 132 to 320, and for $c = 3$ (blue square), n ranges from 116 to 224.

4.4.3 Barriers Proportional to $n^{0.3}$

Finally for $\alpha = 0.3$, numerical diagonalization indicates the gap decreases polynomially in n for low and high n, no matter what c is chosen. Since the width of the barrier does not increase often for such a low scaling power α, the number of n accessible to the QMC simulations is low here. The data is displayed in Fig. 4.10. Notice that there does seem to be a linear relationship on the log-log scale between inverse gap size and run-time here, though it is partially masked by the dearth of data points. However, this does seem to indicate a polynomial relationship between QMC run-time and g_{min}^{-2}.

Additionally, a plot of run-time versus s for higher powers, such as in Fig. 4.6, shows a noticeable spike near where QAO has its minimum gap. For lower powers, such as $\alpha = 0.3$

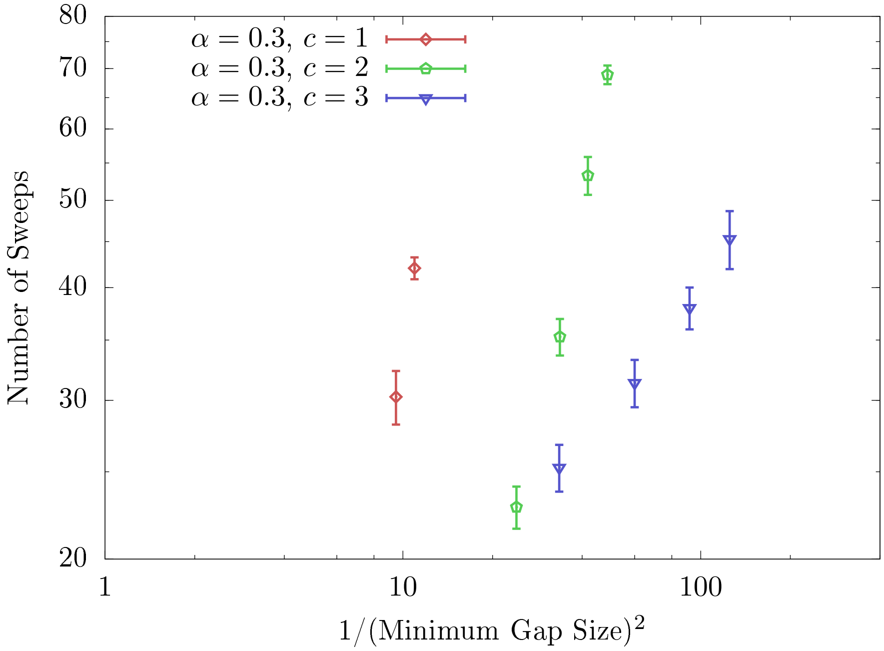

Figure 4.10: QMC Sweeps vs. $g_{\min}$ for barrier size $cn^{0.3}$: There appears to be a linear relationship here, indicating that QMC performance and QAO performance are polynomially related in this region. For $c = 1$ (red diamond), n ranges from 104 to 396, for $c = 2$ (green pentagon), n ranges from 104 to 660, and for $c = 3$ (blue triangle), n ranges from 104 to 396.

as shown in Fig. 4.11, there is no noticeable spike in the run-time. From my simulation results, it seems that the distinction between spikes and no spikes corresponds with the superpolynomial scaling cutoff I saw in the spectral gap in Section 4.1.

4.5 Discussion

First, in Section 4.1, I numerically verified a folklore result [25] about the relationship between n and the minimum gap $g_{\min}$. I showed that $g_{\min}$ scales polynomially with n for barriers whose height and width grow like $\alpha < \frac{1}{3}$ but that for $\alpha > \frac{1}{3}$, the minimum gap decreases faster than a power law. This indicates that QAO can succeed in finding the true ground state in polynomial time only for $\alpha < \frac{1}{3}$.

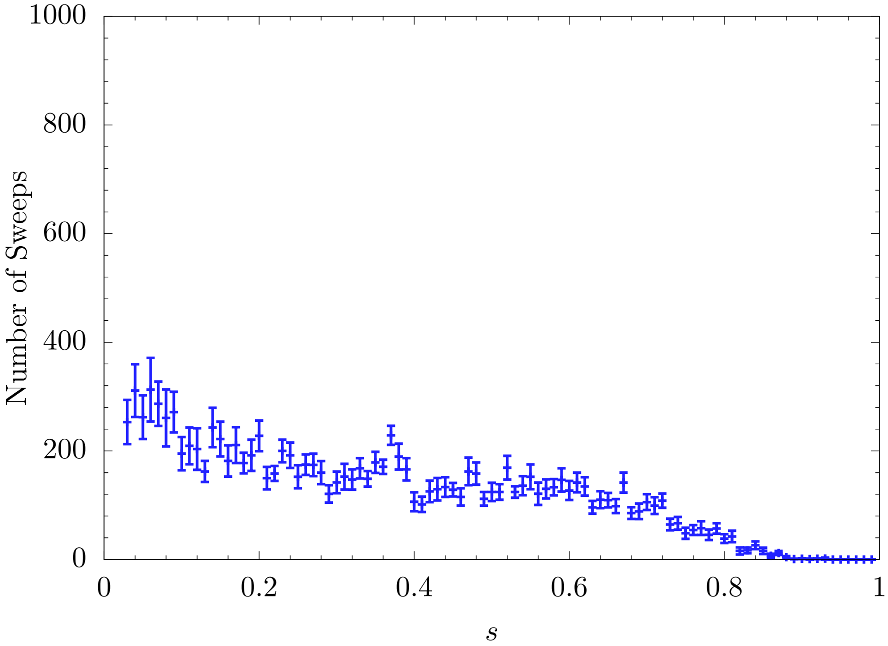

Figure 4.11: QMC Sweeps vs. s for barrier size $3n^{0.3}$ at $n = 116$: This is averaged over 30 simulations. Notice that unlike Fig. 4.6, there is no noticeable spike here corresponding to the location of the minimum gap in QAO

The numerical results with Quantum Monte Carlo simulations show that above $\alpha = \frac{1}{3}$, there is a clear slowdown in the QMC algorithm (see Fig. 4.6) whose location in s corresponds well with the location of the minimum gap in QAO. This slowdown all but disappears for lower α (see Fig. 4.11) where the QMC algorithm has little trouble overcoming the potential barrier. This is strong evidence that there is a correlation between spectral gap and QMC performance.

Furthermore, in Section 4.4, I showed that there is indeed a correlation between gap size and QMC run-time. For αs less than $\frac{1}{3}$, I see data consistent with a polynomial relationship between QMC run-time and g_{min}^{-2}. This relationship is more difficult to discern for $\alpha > \frac{1}{3}$ with there seeming to be either a polynomial or superpolynomial relationship. The lack of a solid polynomial relationship could be due to small n effects which are more prevalent in my simulations for higher α, or it could also be due to inadequacies in my

QMC implementation rather than QMC algorithms in general.

Most notably the algorithm keeps a fixed Δs throughout its annealing schedule and relies on spending more time on each s value rather than decreasing the size of the s step. A more advanced algorithm could also dynamically update s to move more slowly through problem regions.

For the most part, the simulations also keep the number of Trotter time steps $T = 4n$. While $T = 4n$ is sufficient for the region of parameter space discussed in this chapter, it is possible that other Trotterization divisions would be more efficient.

Furthermore, the numerical results presented here comparing QAO and QMC have since been superseded by analytic results from Crosson and Harrow [23]. Using the Markov chain properties of QMC, they analytically showed that there is a polynomial (albeit large polynomial) relationship between QMC runtime and the inverse spectral gap for the barrier tunneling problem. Therefore, the results in this chapter that indicate a non-polynomial relationship point to a discrepancy between the asymptotic behavior proven in [23] and the actual small n scaling behavior observed here. In Ch. 6, I will discuss some of the idiosyncrasies and problems with small-n behavior in this problem.

Chapter 5

Barrier Spectral Gap Characterization

It has been conjectured that a large part of QAO's power comes from the ability of quantum systems to tunnel through potential barriers. In this chapter, I focus on the barrier tunneling model and prove information about its asymptotic spectral gap. Previously, Goldstone [25] had presented the form of the spectral gap as a function of barrier size, but he never published a proof. The results of this chapter mostly formalize Goldstone's conjecture and make his statement rigorous. Additionally, this chapter is heavily connected with Appendix A which takes the preexisting Villain transformation and makes it rigorous and applicable to QAO problems. The results of the Villain transformation are cited in this chapter, but its rigorous mathematical development and usage are reserved for the Appendix.

Reichardt [42] showed that QAO can tunnel in constant time ($g_{\min} \in \mathcal{O}(1)$) provided that the area (width × height) of the barrier is bounded by $\mathcal{O}(\sqrt{n})$. More recently, Crosson and Deng [40] examined thin barriers of varying height, and Kong and Crosson [43] found that sufficiently large barriers lead to exponential run-times. Jiang et al. [24]

showed that Quantum Monte Carlo (QMC) can reproduce the exponential run-time behavior of thermally assisted quantum tunneling through such large barriers. Crosson and Harrow [23] have shown that QMC can efficiently solve barrier tunneling problems for spike barriers and barriers of the size considered by Reichardt's proof. Independently the Wim van Dam and I have numerically found [18] (Ch. 4) that the transitions between constant, polynomial, and exponential run-time scaling for QMC simulations coincide with the same transitions for QAO.

In this chapter, I consider barriers with width proportional to n^{α} and height proportional to n^{β} and mainly focus on barriers with $1/2 \leq \alpha + \beta$, which is above Reichardt's [42] constant scaling region, and $2\alpha + \beta < 1$ which is below Kong and Crosson's [43] exponential scaling region. I show that barriers in this intermediate size regime lead to polynomial scaling of the minimum spectral gap with $g_{\min} \in \mathcal{O}\left(n^{1/2-\alpha-\beta}\right)$. Through the quantum adiabatic theorem this scaling implies that a polynomial running time is sufficient for the QAO algorithm to tunnel through such barriers. Additionally, the Villain method also confirms Kong and Crosson's exponential scaling and provides an exact form for the polynomial prefactor on the exponential.

Section 5.1 presents details of previous work on this problem and highlights both the polynomial scaling region between $1/2 \leq \alpha + \beta$ and $2\alpha + \beta < 1$ where few solid results have been published and the unexplored region for $\alpha > \beta$.

The barrier tunneling problem lends itself to a large spin analysis using either spin coherent states [51] or the Villain transformation [26]. Section 5.2 briefly touches on spin coherent states, which have been used to analyze this problem before [43], and presents an in depth analysis using the Villain transformation, resulting in a semi-classical Hamiltonian that describes the dynamics for large n.

Focusing on just the critical region of the problem where the spectral gap is smallest, Section 5.3 derives a model that approximates the semi-classical Hamiltonian in the large n limit. I provide several arguments for why this model accurately represents the asymptotic behavior of the original problem. Finally, in Section 5.4, I use this model to derive an exact asymptotic expression for the scaling behavior of the spectral gap.

The work presented in this chapter appeared previously in an article by Wim van Dam and myself [19]. This article also included details of the Villain transformation which have been moved to Appendix A.

5.1 Previous Asymptotic Results

A 2002 folklore result by Goldstone [25] says that for $\alpha < \beta$ and $\alpha < 1/2$ the minimum gap for tunneling through an $n^\alpha \times n^\beta$ barrier scales as a function of n like

$$g_{\min} \propto \begin{cases} 1 \text{ if } \alpha + \beta \leq \frac{1}{2}, \\ n^{1/2-\alpha-\beta} \text{ if } \alpha + \beta > \frac{1}{2} \text{ and } 2\alpha + \beta \leq 1, \\ p(n) \exp\left(-C\, n^{(2\alpha+\beta-1)/2}\right) \text{ if } 2\alpha + \beta > 1, \end{cases} \tag{5.1}$$

where $p(n)$ is some polynomial in n.

While this result has never been published, its derivation is known to use "large spin and instanton methods" [25]. Fig. 5.1 shows the scaling behavior according to [25]. Parts of Goldstone's result have been verified by several other sources.

Reichardt [42] rigorously proved the existence of the constant region, and his results apply to the entire region where $\alpha + \beta < 1/2$ not just for $\alpha < \beta$. Recently, Kong and Crosson [43] verified the behavior of the gap in the exponential region for $2\alpha+\beta > 1$ using the instanton method [52], and Jiang et al. [24] have also found the same exponential scaling behavior for the runtime of thermally assisted quantum annealing on this barrier problem using a WKB approach. In Ch. 4 and [18] I numerically analyze the transition

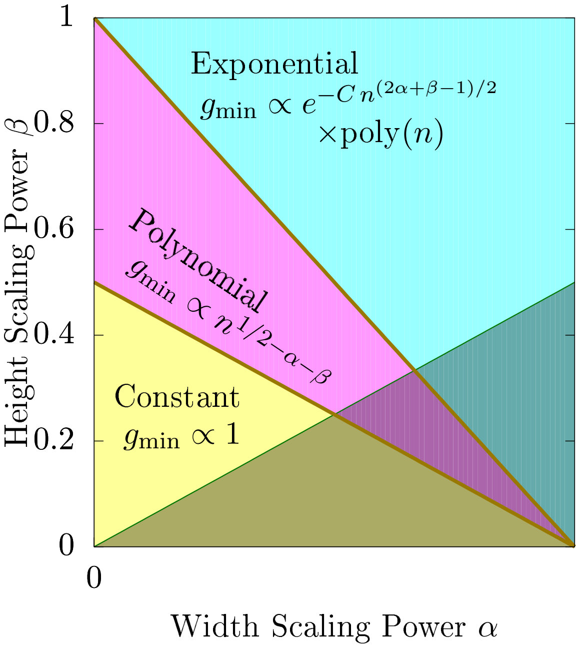

Figure 5.1: The spectral gap scaling of QAO according to the original folklore result by Goldstone [25]. This large n behavior describes tunneling through a barrier of size $n^{\alpha} \times n^{\beta}$ in the setting of n symmetric qubits. The folklore result is restricted as it only works for $\alpha < \beta$ and $\alpha < 1/2$, and it predicts constant, polynomial, or exponential scaling of the minimum gap $g_{\min}$ depending on the barrier size. The proof of this result has not been formally published.

between the polynomial and exponential regions. Notably, no previously published work before my article with Wim van Dam [19] has been able to verify the polynomial region, and while Kong and Crosson [43] proved the exponential region scaling, they restricted their proof to $\alpha < \beta$ and did not derive the polynomial prefactor. The different scaling regions in α and β are shown in Fig. 5.2, with references in the figure caption to which sources proved that region's scaling behavior, including what is proven in this paper.

The goal of the current chapter is to explore both the polynomial region between $\alpha + \beta > 1/2$ and $2\alpha + \beta < 1$ and in general the region where $\alpha > \beta$. I develop elementary techniques to analyze the spectral gap and verify the polynomial, $n^{1/2-\alpha-\beta}$, scaling behavior, and I show that the results of Eq. 5.1 are valid even when $\alpha > \beta$.

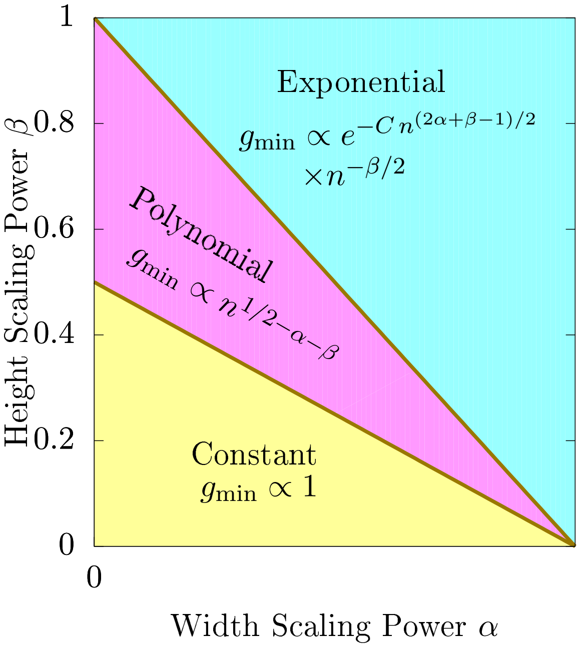

Figure 5.2: The spectral gap scaling during QAO tunneling through a barrier of size $n^{\alpha} \times n^{\beta}$. Unlike Fig. 5.1, this figure displays all current knowledge of each region, which includes the case $\alpha > \beta$. The yellow "Constant" region was proven by Reichardt [42], and the blue "Exponential" region was shown in [43, 24] up to the polynomial prefactor. The current chapter proves the polynomial scaling $g_{\min} \propto n^{1/2-\alpha-\beta}$ for the red region with $1/2 < \alpha + \beta$ and $2\alpha + \beta < 1$. This chapter also determines the polynomial prefactor for the blue exponential region described by $1 < 2\alpha + \beta$ with $g_{\min} \propto n^{-\beta/2} \exp(-Cn^{(2\alpha+\beta-1)/2})$.

5.2 Large Spin Approximation

The Hamiltonian

$$
\begin{aligned}
H_{\text{sym}}(s) = & \sum_{h=0}^{n} \left[\frac{(1-s)}{2} n + s(h + b(h)) \right] |h\rangle \langle h| \\
& - \frac{(1-s)}{2} \sum_{h=0}^{n-1} \sqrt{(h+1)(n-h)}\, |h\rangle \langle h+1| \\
& - \frac{(1-s)}{2} \sum_{h=0}^{n-1} \sqrt{(h+1)(n-h)}\, |h+1\rangle \langle h|
\end{aligned}
\tag{5.2}
$$

readily lends itself to reinterpretation as the Hamiltonian for a single particle with spin $J = n/2$. A common analytic technique for dealing with a spin Hamiltonian is to use spin coherent states [51] to create a semi-classical continuous version of the Hamiltonian. Several groups [53, 43] have used spin coherent states to analyze qubit systems, and Kong and Crosson [43] have employed spin coherent states to analyze the symmetric barrier problem for exponentially small gaps. I use a similar technique employing a modified and formalized version of the Villain transformation [26]. The Villain transformation has been used for similar problems [54, 55, 56]; I present a more formal approach to this transformation, making clear all its assumptions and needs in Appendix A.

If you re-imagine the Hilbert space as representing a spin $J = n/2$ particle and associate $\hat{J}_z$ eigenstates $|m\rangle$ with $|h\rangle$ states through $|m\rangle = |h - J\rangle$, then the symmetric Hamiltonian can be rewritten in terms of spin operators $\hat{J}_i$ as

$$\hat{H}(s) = -(1-s)\hat{J}_x + s\hat{J}_z + s\,b(\hat{J}_z + J) + J\Theta, \tag{5.3}$$

where Θ represents some constant. Since I only care about energy differences, this constant Θ can be arbitrary, and later, I use it to ensure that the bottom of the potential energy well sits at zero energy.

Large spin techniques then pull a factor of $J = n/2 \equiv 1/\epsilon$ out of the Hamiltonian so that the operators can be represented using $\hat{\jmath}_i = \epsilon\hat{J}_i$ that have eigenvalues that run from -1 to $+1$. Specifically, I call the $\hat{\jmath}_z$ eigenvalue $-1 \leq q \leq 1$, and in the large J (i.e. small ϵ) limit, q can be treated as a continuous variable. I also introduce $r(q) \equiv \epsilon b(Jq + J)$ that is zero everywhere except in the vicinity of $q = -1/2$ where there is a bump of width $\epsilon^{1-\alpha}$ and height $\epsilon^{1-\beta}$. The Hamiltonian can be rewritten as

$$\epsilon\hat{H}(s) = -(1-s)\hat{\jmath}_x + s\hat{\jmath}_z + sr(\hat{\jmath}_z) + \Theta. \tag{5.4}$$

At this point, I can write an approximate Schrödinger equation for this Hamiltonian using the Villain transformation. In Appendix A, I have taken the standard Villain transformation and made its logic more formal, applying it specifically to Eq. 5.4. In making the logic more formal, I have held off taking the continuum limit of q as long as possible. The end result of the Villain transformation itself, before making any assumptions about the properties of the eigenstates, gives a continuum Schrödinger equation:

$$\begin{aligned}\epsilon E\psi(q) = \Big(&sq + sr(q) + \Theta - (1-s)\sqrt{1-q^2} \\ &-(1-s)\frac{\epsilon^2}{2}\sqrt{1-q^2}\frac{\partial^2}{\partial q^2} + \mathcal{O}(\epsilon)\Big)\,\psi(q).\end{aligned} \tag{5.5}$$

The first line includes a potential energy, and the next one contains the kinetic term for the problem. Note that the norm of the second derivative operator, $\frac{\partial^2}{\partial q^2}$, is proportional to ϵ^{-2} which is why this term survives. At this point, the problem cannot be simplified without making reference to the desired eigenstates. Notably, if I assume that the system is near $s^* = \frac{1}{2}(\sqrt{3}-1)$ where the minimum of the potential energy is at $q = -\frac{1}{2}$ in the $\epsilon \to 0$ limit and make reasonable assumptions about the nature of the ground state and first excited state, then Eq. 5.5 can be simplified even more. In Appendix A, I formalize these approximations, and in Sec. 5.3 I analyze the resulting approximate differential equation.

5.3 Quadratic Potential Approximation

In Appendix A, I continue the approximation of Eq. 5.5 by noting that the low-lying energy states for $s^* = \frac{1}{2}(\sqrt{3}-1)$ are centered in the extremely close vicinity of $q = -1/2$. This allows me to focus on the variable $x \equiv q + \frac{1}{2}$ and the region near $x = 0$. For the low-lying energy states, such as the ground state and first-excited state that determine the spectral gap, the approximate differential equation representing the barrier tunneling problem in the small ϵ limit is

$$\frac{\partial^2\psi}{\partial x^2} = \frac{1}{\epsilon^2}\left[\omega^2 x^2 + \tfrac{4}{3}r(x-\tfrac{1}{2}) - c\epsilon E\right]\psi(x), \tag{5.6}$$

where $c \equiv 8/(3(\sqrt{3}-1))$ and $\omega \equiv 4/3$.

The potential has become an ordinary quadratic well, so I use standard techniques from the quantum harmonic oscillator to solve the Schrödinger equation. Furthermore, since the width of the barrier $r(x-1/2)$ is proportional to $\epsilon^{1-\alpha}$ and the height is proportional to $\epsilon^{1-\beta}$, in the region of the barrier, it will overshadow the quadratic potential in the small ϵ limit if $(\epsilon^{1-\alpha})^2 < \epsilon^{1-\beta}$ which translates to $1 > 2\alpha - \beta$. If I restrict to $\alpha < 1/2$ and $\beta > 0$, this is always true, so that I can treat the barrier as the dominant factor in the region where $|x| = \mathcal{O}(\epsilon^{1-\alpha})$. Therefore, I can say that the following is a good approximation for the problem in the large n limit:

$$\frac{\partial^2\psi}{\partial x^2} = \epsilon^{-2}\left[V(x) - \epsilon c E\right]\psi(x) \tag{5.7}$$

where

$$V(x) = \begin{cases} \epsilon^{1-\beta} & \text{if } -a < x < a \\ \omega^2 x^2 & \text{otherwise} \end{cases}, \tag{5.8}$$

where $a \equiv \frac{1}{2}\epsilon^{1-\alpha}$. In Eq. 5.8 I have settled on a form of $r(q)$ that is just a square barrier. I have focused on the square barrier (as opposed to the binomial barrier) since it makes the differential equation in Eq. 5.7 easy to solve, but I have numerically looked at the binomial (Gaussian) barrier which shows signs that it has similar scaling behavior of $g_{\min}$

with n.

5.4 Asymptotic Expansion

In this section, I focus on the differential equation in Eq. 5.7 and find the spectral gap. Since Eq. 5.7 describes the original n dimensional hypercube problem in the large n limit, an asymptotic analysis of Eq. 5.7 in the small ϵ limit will give the correct asymptotics for the original problem.

Outside of the barrier, the Schrödinger equation looks like that of an ordinary quantum harmonic oscillator problem, but I cannot use the standard harmonic oscillator solutions since these have already had boundary conditions imposed, ensuring that the wave-functions go to zero as $x \to \pm\infty$. To get the solutions for arbitrary boundary conditions, I compare the harmonic oscillator equation to the Weber equation

$$\frac{d^2 D_\nu(z)}{dz^2} + \left(\nu + \frac{1}{2} - \frac{1}{4}z^2\right) D_\nu(z) = 0, \tag{5.9}$$

where ν is an arbitrary eigenvalue, and $D_\nu(z)$ is known as a parabolic cylinder function. Note that when ν is a positive integer and z is real, these functions become the standard Gaussians times Hermite polynomials expected from the harmonic oscillator.

When ν is not a positive integer and z is real, these functions blow up as $z \to -\infty$ but go to zero as $z \to \infty$, so I can use these as the solution to the DE for $x > a$. Furthermore, to get the solution in the $x < -a$ region, I employ the symmetry of the problem about $x = 0$ to say that the eigenfunctions are either symmetric or anti-symmetric. Therefore, the eigen-solutions to the differential equation will have the form

$$\psi(x) = \begin{cases} \pm A_1 D_{\nu_\pm}\left(-\sqrt{\frac{2\omega}{\epsilon}}x\right) & \text{if } x < -a \\ A_2 \mathrm{e}^{k_\pm x} \pm A_2 \mathrm{e}^{-k_\pm x} & \text{if } -a < x < a \\ A_1 D_{\nu_\pm}\left(\sqrt{\frac{2\omega}{\epsilon}}x\right) & \text{if } x > a \end{cases}, \tag{5.10}$$

where $\nu_\pm \equiv \frac{cE_\pm}{2\omega} - \frac{1}{2}$ and $k_\pm \equiv \sqrt{\epsilon^{-1-\beta} - \epsilon^{-1} cE_\pm}$.

By applying continuity in the wave-function and its derivative across the boundary at $x = \pm a$, I can find a transcendental equation for the energies, denoted by $E_\pm$ representing the two lowest level energy states:

$$\begin{aligned} & k_\pm D_{\nu_\pm}\left(\sqrt{\frac{2\omega}{\epsilon}}a\right)\left(\mathrm{e}^{k_\pm a} \mp \mathrm{e}^{-k_\pm a}\right) = \\ & \sqrt{\frac{2\omega}{\epsilon}} D'_{\nu_\pm}\left(\sqrt{\frac{2\omega}{\epsilon}}a\right)\left(\mathrm{e}^{k_\pm a} \pm \mathrm{e}^{-k_\pm a}\right). \end{aligned} \tag{5.11}$$

This transcendental equation can be solved numerically for the lowest energy levels, and a comparison of this numerical solution to the full spectral gap of the Hamiltonian in Eq. 3.5 is shown in Fig. 5.3. In the rest of this section, I show that I can do better than numerical solutions to Eq. 6.3 by finding an asymptotic solution in the limit of large n.

I expect the energies to be close to the unperturbed first excited state energy of $E_1 = 3\omega/c$, so I set $E_\pm = (3\omega + \delta_\pm)/c$ and find $\delta_\pm$ in the limit of small ϵ. In this limit $\nu_\pm = \frac{\delta_\pm}{2\omega} + 1$ and $k_\pm \approx \epsilon^{-\frac{1}{2}-\frac{\beta}{2}}$.

At this point, I want to calculate $g_{\min} = |\delta_+ - \delta_-|/c$ up to leading order in ϵ using these approximations. If I assume that $2\alpha + \beta < 1$, $\alpha < 1/2$, and $\alpha + \beta > 1/2$ (this corresponds to the polynomial region in Fig. 5.2), then the gap becomes

$$g_{\min} = \frac{8(\omega)^{3/2}}{c\sqrt{\pi}} \epsilon^{\alpha+\beta-\frac{1}{2}} \propto n^{\frac{1}{2}-\alpha-\beta}. \tag{5.12}$$

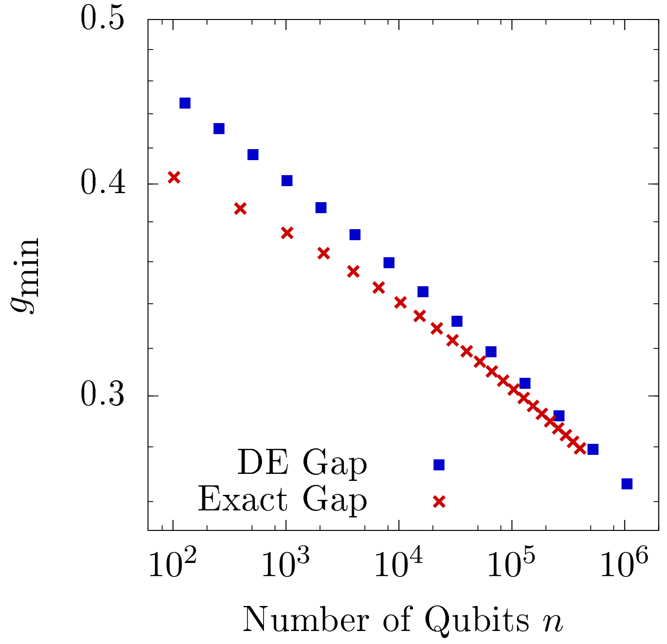

Figure 5.3: Comparison between the scaling of the true spectral gap obtained by diagonalization of Eq. 3.5 and the gap obtained by solving for the eigenenergies of the differential equation in Eq. 5.7. The latter is calculated by numerically solving the transcendental equation in Eq. 6.3. Note that these scalings converge for $n > 10^5$, confirming that the derivation of Eq. 5.7 is indeed valid in the limit of large n. This data was obtained for a rectangular barrier with $\alpha = \beta = 0.3$.

Similarly, if I assume $2\alpha + \beta > 1$ and $\alpha < 1/2$, then the gap becomes

$$\begin{aligned} g_{\min} &= \frac{16(\omega)^{3/2}}{c\sqrt{\pi}} \epsilon^{\beta/2} \exp\left(-\epsilon^{\frac{1}{2}-\alpha-\frac{\beta}{2}}\right) \\ &\propto n^{-\frac{\beta}{2}} \exp\left(-(n/2)^{\alpha+\frac{\beta}{2}-\frac{1}{2}}\right). \end{aligned} \tag{5.13}$$

This result matches the exponentially small gap found by Kong and Crosson [43] and Jiang et al. [24] and determines the prefactor that those groups did not find.

The dependences on ϵ in Eqs. 5.12 and 5.13 are exactly what I would expect given Eq. 5.1. Notice that I do not need to assume $\alpha < \beta$ as in Eq. 5.1, so the result extends farther than Goldstone's result and covers the entire area bounded by $0 < \alpha < 1/2$ and

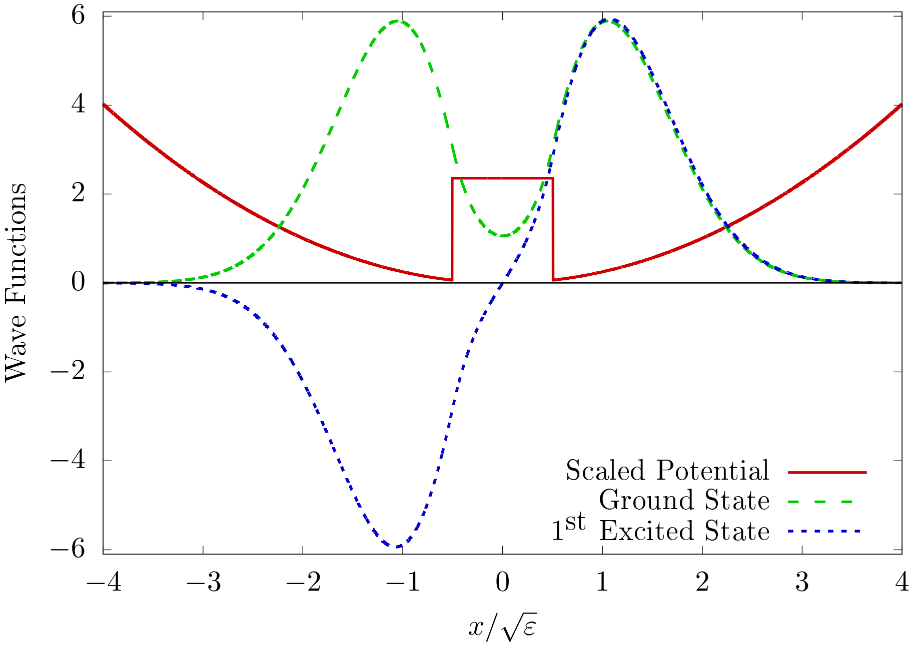

Figure 5.4: The ground state and first excited state wave functions for the quadratic well approximation of Eq. 5.7. I also display the potential energy from Eq. 5.8 multiplied by a factor of $10/\sqrt{\epsilon}$ so that it is fully visible. Notice that the ground state looks like a Gaussian with the middle dragged downward and the first excited state looks unchanged from the unperturbed quantum harmonic oscillator since the barrier sits in a region where this function was already small.

$\alpha + \beta > 1/2$.

In Fig. 5.4, I plot the exact ground state and first excited state for the quadratic approximation for $\epsilon = 1/5000$, a bump width of $1/70$, and a bump height of $1/300$. These values were chosen to provide visibility for the bump and its effect on the eigenfunctions. The potential is also plotted, multiplied by $10/\sqrt{\epsilon}$ so that it is visible. Here I are using the exact energies, obtained by solving the transcendental equation, Eq. 6.3, numerically. Notice that the ground state looks like a Gaussian with its center pulled down whereas the first excited state looks almost unchanged from its unperturbed state. The first excited state is unchanged because it was already small in the vicinity of $x = 0$, so the barrier does not alter this state much by making that region more unfavorable. This is also reflected in the approximation in Eq. 5.12 where the leading order term shown here is due to the ground state rather than the first excited state.

I have taken the original n qubit barrier tunneling problem and, through a series of large-n approximations, have arrived at an elementary barrier tunneling problem in one continuous dimension. The resulting approximate Schrödinger equation gives a transcendental equation for the energies that in the large-n limit gives a spectral gap that is proportional to $1/n^{\alpha+\beta-\frac{1}{2}}$ for $1/2 < \alpha+\beta$ and $2\alpha+\beta < 1$, and $n^{-\beta/2}\exp(-Cn^{(2\alpha+\beta-1)/2})$ when $1 < 2\alpha+\beta$. The gap scaling result verifies and provides a solid proof the 2002 folklore result by Goldstone [25].

Chapter 6

Large n Behavior

Adiabatic quantum computation is provably universal for quantum computation if non-stoquastic Hamiltonians are allowed [36], and an understanding of tunneling in a quantum annealing setting can shed light on the power and limitations of QAO, even if stoquastic tunneling alone is not enough for a quantum speed-up. To that end, this chapter studies quantum tunneling as part of QAO in a very simple setting: tunneling in position-space through a one-dimensional barrier. In Ch. 5, I showed that the simple problem of one-dimensional barrier tunneling is equivalent to barrier tunneling of n qubits in a bit-symmetric setting in the large n limit.

Despite the simplicity of one-dimensional position-space tunneling, in this chapter I find that an asymptotic analysis of QAO spectral gaps for $n \to \infty$, which give bounds on the adiabatic run time, fail to accurately describe the behavior of the tunneling for moderately large n. I find that very large n (around $n > 10^{12}$) are needed to ensure that asymptotic expressions accurately describe behavior, to the point where polynomial and exponential run time scaling can be confused at lower n. Given the current and near-term sizes of quantum computer implementations [10], this failure of asymptotics casts doubt on the efficacy of using asymptotics to analyze the output of these systems or to predict their capabilities.

In section 6.1, I review the barrier tunneling problem and outline the terminology that I use throughout the rest of the chapter. Section 6.2 gives the first evidence that the asymptotic limit may not be accurate even for moderately large n. I show that the failure of the asymptotics is the same in both the hypercube and one-dimensional cases, proving that the failure is in the tunneling problem itself and not in the approximation scheme from Ch. 5 and App. A.

I compare the exact spectral gap to its asymptotic approximation in Section 6.3 to see where the issues discussed in the previous section arise. This section also provides estimates for the n necessary to see asymptotic behavior and explores what effect next-order corrections to the asymptotics have on the analysis.

Section 6.4 uses properties of the derivatives of log-log data to analyze how well a curve approximates a power law versus an exponential fit. Furthermore, I show how easily finite data for tunneling could lead to a mislabeling of the scaling behavior of the algorithm.

The work presented in this chapter is joint work performed with Wim van Dam. The information in this chapter previously appeared in an article by Wim van Dam and me [20].

6.1 Symmetric Barrier Tunneling Problem

The one-dimensional barrier tunneling problem is inspired by tunneling using n qubits in a bit-symmetric setting. A thin spike version of this bit-symmetric problem was originally proposed by Farhi et al. [53], and the general bit-symmetric barrier tunneling problem has been analyzed extensively by other groups [42, 40, 18, 23, 43, 19, 24, 25]. The final cost function $f(z)$ is a function only of the Hamming weight $|z|$ as $f(z) = |z| + b(|z|)$

where $b(w)$ is a function centered around $w = n/4$ with width and height proportional to n^{α}.

In Ch. 5, I showed that the multi-qubit symmetric barrier tunneling Hamiltonian can be mapped onto the symmetric $(n+1)$-dimensional subspace where it describes a single spin $1/\epsilon \equiv n/2$ particle. Furthermore, using a modified version of the Villain transformation [26] (developed in App. A), I showed that this system can be mapped, in the $n \to \infty$ limit, onto a semiclassical wave equation in one continuous variable x. This mapping only works for low-lying energy states, but because I only care about the spectral gap between the ground state and first excited state, this restriction is acceptable. In Ch. 5, I restrict to a square barrier with $\alpha = \beta$ so that the final, semiclassical Schrödinger equation becomes

$$\frac{d^2\psi}{dx^2} = \epsilon^{-2}\left[V(x) - \epsilon c E\right]\psi(x), \tag{6.1}$$

where

$$V(x) = \begin{cases} \omega\epsilon^{1-\alpha} & \text{if } -a < x < a \\ \omega^2 x^2 & \text{otherwise} \end{cases} \tag{6.2}$$

The wave function $\psi(x)$ is the continuous version of the eigenvector components in the spin-$n/2$ problem, and E is the eigenenergy. The variable x is a rescaled and shifted version of the Hamming weight, taken in the continuum limit, and the barrier has also been rescaled to have width $2a \equiv \epsilon^{1-\alpha}$ and height $\omega\epsilon^{1-\alpha}$. The constants can either be taken as arbitrary or set to $c = 8/(3(\sqrt{3}-1))$ and $\omega = 4/3$ to get a direct correlation to the qubit problem.

While Eq. 6.1 is derived from the n qubit problem, it can also stand on its own, describing the critical tunneling moment in a one-dimensional continuous system. In this context, I just have a particle moving in a one-dimensional quadratic well with a barrier in the middle. For most of this chapter, I treat the Schrödinger equation in Eq. 6.1 as its own independent problem.

The continuous well differential equation, Eq. 6.1, can be solved using exponentials and parabolic cylinder functions, $D_\nu(y)$, up to a transcendental equation for the eigenenergies:

$$k_\pm D_{\nu_\pm}\left(\sqrt{\frac{2\omega}{\epsilon}}a\right)\left(\mathrm{e}^{k_\pm a} \mp \mathrm{e}^{-k_\pm a}\right) = \sqrt{\frac{2\omega}{\epsilon}}D'_{\nu_\pm}\left(\sqrt{\frac{2\omega}{\epsilon}}a\right)\left(\mathrm{e}^{k_\pm a} \pm \mathrm{e}^{-k_\pm a}\right), \tag{6.3}$$

where $\nu_\pm \equiv \frac{cE_\pm}{2\omega} - \frac{1}{2}$, $k_\pm \equiv \sqrt{\omega\epsilon^{-1-\alpha} - \epsilon^{-1}cE_\pm}$, and $E_\pm$ refers to even (+) or odd (−) energy solutions. In the limit of $n \to \infty$, Eq. 6.3 can be solved for the lowest two energy levels, resulting in a gap expression

$$g_{\min} = \frac{8\omega^{1/2}}{c\sqrt{\pi}}\epsilon^{2\alpha-1/2} \tag{6.4}$$

for $1/4 < \alpha < 1/3$, and

$$g_{\min} = \frac{16\omega}{c\sqrt{\pi}}\epsilon^{\alpha/2}\exp\left(-\sqrt{\omega}\epsilon^{\frac{1}{2}-\frac{3}{2}\alpha}\right) \tag{6.5}$$

for $1/3 < \alpha < 1/2$. These asymptotic gap expressions are consistent with results for the original symmetric hypercube tunneling problem from other groups using various methods [53, 25, 43, 24].

6.1.1 Definitions of Gap Analyses

In the earlier parts of this section I outline three steps in the approximation procedure from which the gap can be extracted. In the rest of the chapter I refer back to these three different models, so I take a moment here to label these three models and define exactly what I mean by each term.

- **Discrete analysis**: The original n qubit hypercube problem consists of a discrete Hilbert space that can be interpreted in terms of Hamming weight or spin states. The gap can be extracted from this problem through exact diagonalization of the $(n + 1)$-dimensional symmeterized problem. While the discrete problem was the original inspiration for this work, much of the subsequent text focuses exclusively on the continuous and asymptotic gaps.

- **Continuous analysis**: After the Villain transformation, the discrete problem transforms into a parabolic well potential for a continuous, position-like variable. The Schrödinger equation governing this continuous problem is shown in Eq. 6.1, and the exact spectral gap can be extracted by numerically solving the transcendental equation for the energies, Eq. 6.3. For most of this chapter, the continuous problem is treated as an independent problem, and most of my efforts focus on comparing its exact spectral gap to the asymptotic limit.

- **Asymptotic analysis**: In the large n limit, the asymptotic expression for the spectral gap of the discrete and continuous problems can be derived. I did this derivation in a Ch.5, and the gap expressions are reproduced in Eqs. 6.4 and 6.5. This version of the gap is the true asymptotic limit of the discrete and continuous gaps, but much of this chapter focuses on how large n must be for the asymptotic limit to be a good approximation of the true gap.

6.2 Scaling Power Comparison

When analyzing an algorithm numerically or experimentally, a typical question to ask is how its run time, or in the QAO case the gap, scales with the input size n. Therefore, in this section I analyze the gap numerically mostly ignoring the asymptotic expressions

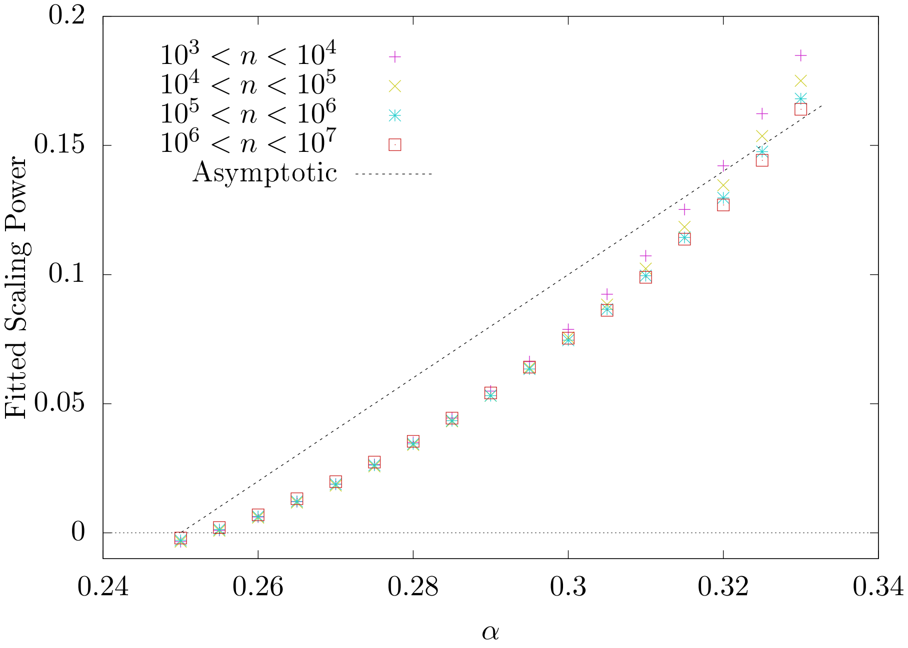

Figure 6.1: Fitted scaling exponent of $g_{\min}$ for the symmetric qubit tunneling problem with a binomial shaped barrier. Each data point shows what scaling power, p, is extracted from a fit to $g_{\min} = An^{-p}$, assuming the fit is done on data from a finite range of n versus the scaling power of the barrier size, α. Each type of point represents data binned from a different range of n, and the dashed line represents the power p predicted by the asymptotic expression. Based on this data alone, the scaling of $g_{\min}$ does not seem to approach the asymptotic expression.

from Eqs. 6.4 and 6.5 and try to determine what asymptotic information can be extracted from numerical data of n vs. $g_{\min}$. I then compare these finite n extrapolations to see how close these results are to the analytic expressions for the asymptotic behavior of $g_{\min}$.

I use symmetry to reduce the discrete qubit Hamiltonian to a tridiagonal $(n+1)$-dimensional matrix (as described in Ch. 3) and then numerically diagonalize it for a finite range of n before performing power law fits of the form $g_{\min} = An^{p}$ in order to extract the exponent p. Fig. 6.1 outlines the resulting p values for these fits done on simulations with $1/4 < \alpha < 1/3$, leading to polynomial scaling asymptotically.

Also, note that the data in Fig 6.1 comes from a problem with a binomial shaped barrier rather than the standard square barrier. To collect these data points, I restrict to n such that the discrete width of the barrier has just increased, so the binomial barrier's full width, which grows with $n^{2\alpha}$, allows collection of more data points than the square barrier, whose width grows with n^{α}. Therefore, I use the binomial barrier in order to collect more data points and create a cleaner data set.

In Fig. 6.1, I show p values for fits done on different bins of n data, so the red filled triangle should most closely approximate the asymptotic value of p. Unfortunately, the scaling powers do not seem to approach their asymptotic values, shown as a dashed black line in the figure, even for the relatively large problem size of one million qubits. The lack of agreement with the asymptotics, even for relatively large numbers of qubits casts doubt on the validity of the asymptotic analysis techniques used in Ch. 5 and App. A.

However, the only possible source of error in the asymptotic analysis in Ch. 5 is in the Villain transformation and the other assumptions that lead to the derivation of the continuous well DE, Eq. 6.1. Therefore, if the spectral gap for the differential equation shows the same difficulty in reaching its asymptotic value, then extremely large n are necessary to observe asymptotic behavior because this is a simple tunneling problem, not because the Villain transformation was used in the derivation of these asymptotics. I do in fact see that if I perform the same procedure that gave Fig. 6.1 for numerical $g_{\min}$ values found from solving Eq. 6.3, I obtain a very similar plot, Fig. 6.2.

The failure of $g_{\min}$ to quickly reach its asymptotic scaling even in Fig. 6.2 indicates that there is something special about simple tunneling problems that defies the conventional wisdom that asymptotic behavior sets in quickly. Based on these figures alone, an experimental or numerical approach to this problem would mislabel the polynomial scaling of this algorithm. From this point, I focus exclusively on the continuous problem, ignoring the discrete problem that motivated it. In the next section, I examine what n

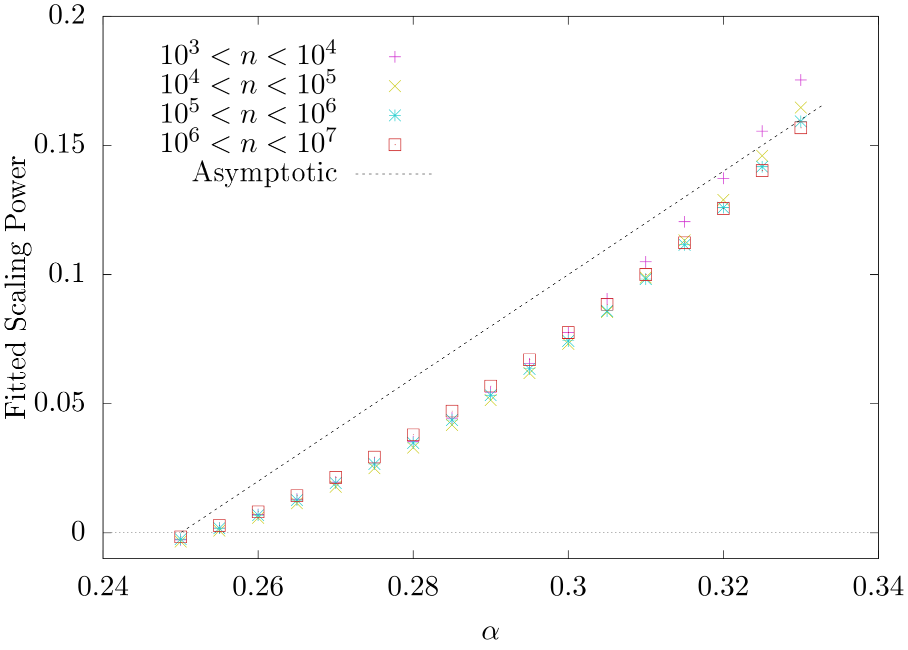

Figure 6.2: Fitted scaling exponent of $g_{\min}$ for the continuous well DE, Eq. 6.1, with a square barrier. Each data point shows what scaling power, p, is extracted from a fit to $g_{\min} = An^{-p}$, assuming the fit is done on data from a finite range of n versus the scaling power of the barrier size, α. Each type of point represents data binned from a different range of n, and the dashed line represents the power p predicted by the asymptotic expression. The trough shows the same behavior as seen in Fig. 6.1, indicating that the problem in that plot was not in the approximations made for the Villain transformation.

are necessary to bring the scaling close to the asymptotic expressions.

6.3 Comparison of Exact Gap and Asymptotic Expressions

Given the large discrepancy seen in Section 6.2 between finite n runtime scaling and asymptotic scaling, it makes sense to explore what values of n are necessary to get close to the true asymptotic scaling behavior. The discrepancy seen in the last section is largest in the polynomial scaling region near the middle of the region far from the boundaries

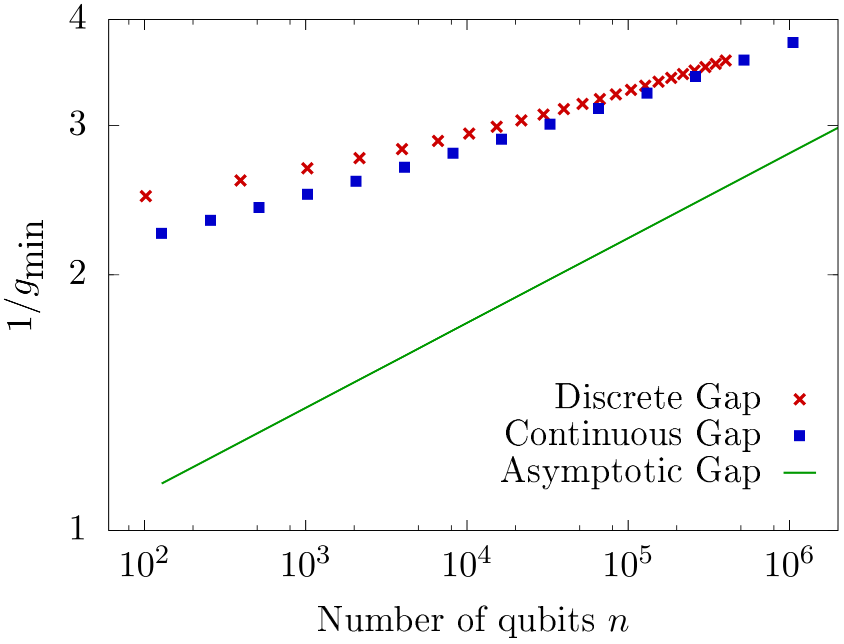

Figure 6.3: Spectral gap $g_{\min}$ vs. n for $\alpha = 3/10$ for the three different gap extraction methods with a square barrier. The discrete gap and continuous gap well approximate each other at these relatively modes n, but neither of them are well approximated by the asymptotic expression for the gap. This difference seems to confirm the failure of the asymptotics observed in Figs. 6.1 and 6.2.

with the constant and exponential regions. Much of the analysis in this section is done for $\alpha = 3/10$ which is a simple value in the middle of the polynomial region where this asymptotic discrepancy effect is large. However, the results are similar for different values of α.

6.3.1 Direct Gap Comparison

Since the asymptotic analysis of $g_{\min}$ in Eqs. 6.4 and 6.5 is exact enough to get the constant prefactor, I can compare the numerical gap directly to the asymptotic expression to visually and numerically see where agreement occurs.

In Fig. 6.3 I show the discrete, continuous, and asymptotic gaps for $\alpha = 3/10$ up to

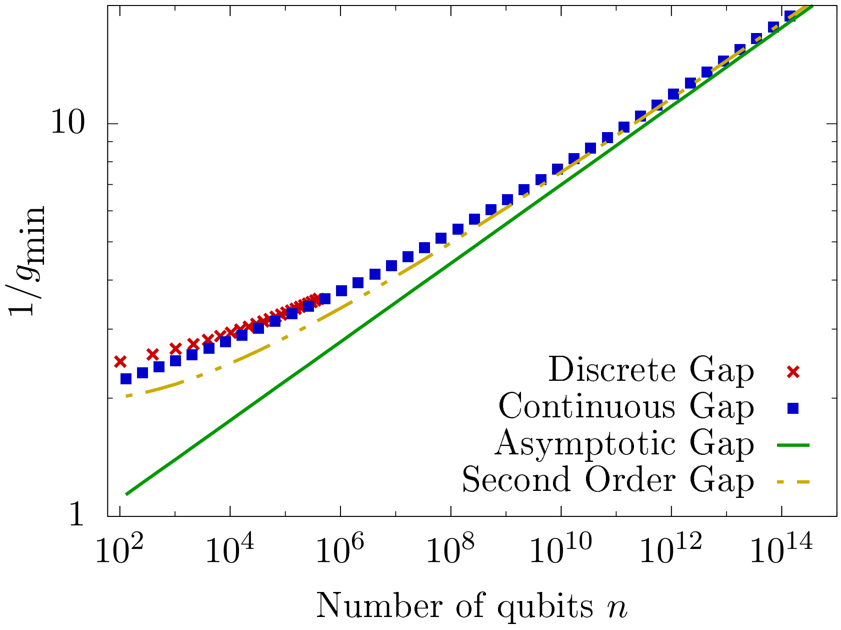

Figure 6.4: Spectral gap $g_{\min}$ vs. n on a log-log scale for $\alpha = 3/10$ for the three different gap extraction methods with a square barrier. This plot covers a much larger range than Fig. 6.3, allowing us to see that the continuous gap does eventually approach the asymptotic gap. However, this approach occurs at extremely large n that are impractical for most current applications. In addition, I have included the second order expansion of the asymptotic large n expression from Eq. 6.6 that more closely approximates the continuous gap.

$n = 10^6$. Since this is a log-log plot, a straight line corresponds to a power law with the slope equal to the scaling power. The continuous and discrete gaps approach each other for the reasonable values of n shown in this plot and have very similar slopes. However, neither the continuous nor the discrete gaps are close to the asymptotic expression, nor do their slopes appear similar to the asymptotic slope. Fig. 6.3 shows how far away the actual gap is from the asymptotic expression, leading to the trough in Fig. 6.2.

With my current computational limits, I cannot push exact diagonalization of the symmeterized n qubit Hamiltonian to much higher n, but the condition in Eq. 6.3 can be solved for extremely large n with rounding error being the only block. Therefore, I can extend Fig. 6.3 to higher n to get Fig. 6.4 from which I conclude that the continuous gap does eventually approach its asymptotic value, but only for extremely large $n > 10^{12}$.

Fig. 6.4 also shows another curve, labeled “Second Order Gap.” This curve shows the asymptotic expression for the gap in the large-n limit, but it keeps both the leading order term and the second highest order term so that this second order gap expression is given by

$$g_{\mathrm{min}}^{(2^{\mathrm{nd}}\mathrm{O})} = \frac{8\omega^{1/2}}{c\sqrt{\pi}}\epsilon^{2\alpha-1/2} - \frac{16}{c\pi}(\ln 2 + \psi_0(-\frac{1}{2}))\epsilon^{4\alpha-1} \tag{6.6}$$

for $1/4 < \alpha < 1/3$ where $\psi_m(z)$ is the polygamma function.

The second order term does improve agreement with the continuous gap, but this fact also shows that the second order (and higher) terms are still relevant and important for determining run time behavior even at the very large end of what quantum computers are expected to be able to accomplish in the next few decades. Since even a simple tunneling problem needs more than just its first-order asymptotic limit to accurately describe its behavior at any reasonable n, this data casts doubt on the efficacy of asymptotic analysis.

I have exclusively used parabolic wells up to this point, so the argument could be made that this large-n behavior is an artifact of the choice of wells. To counter this argument, I have done this same analysis using an appropriately sized square well instead of a parabolic well. The asymptotic scaling behavior of g_{min} has the same n and α dependence in the square well case, and extremely large n are still needed for the asymptotics to accurately describe the exact gap for the square well.

As a general rule, the smoother the potential is the faster g_{min} begins to approximate its asymptotic value. I have chosen the quadratic potential since it makes the differential equation solvable, but I have run a few other simulations with different shaped wells and barriers, specifically a square well and binomial barriers. While a smoother barrier (such as a binomial or Gaussian) does have a gap that reaches its asymptotic behavior faster

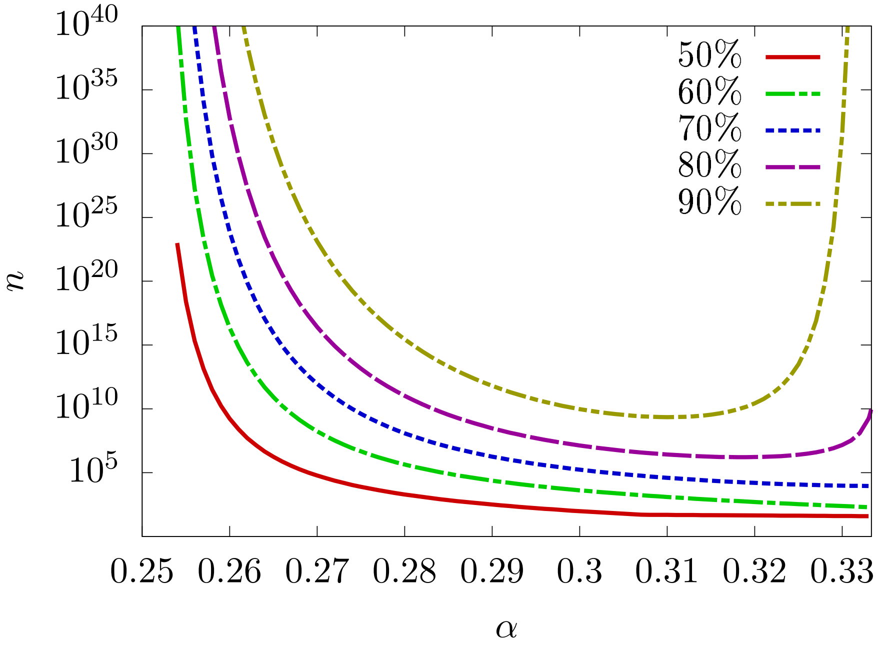

Figure 6.5: I show the n necessary for the ratio of the continuous gap and the asymptotic gap to reach the indicated percentages as a function of α. Especially close to the boundaries of the polynomial region, the system requires extremely large n to reach its asymptotic value. When coupled with Fig. 6.2, this plot indicates that for the middle values of α the continuous gap reaches its asymptotic value faster by deviating more from its asymptotic scaling.

than the square barrier's gap, the gap still requires very large n to have approximately the correct slope. Exact analysis is more difficult for different barrier and potential shapes since I cannot get exact asymptotic scaling expressions for them up to the constant prefactors.

6.3.2 Ratio of Exact Gap to Asymptotic

It is also of interest to know what size of n is necessary before the continuous gap is a certain fraction of the asymptotic gap. In Fig. 6.5, I plot the n necessary for the ratio of the continuous gap to the asymptotic gap to reach the specified value. When I want 50%, the asymptotic gap is a good approximation relatively quickly for higher α, but even for lower α, extremely large n are required. For ratios around 80% or 90% the necessary n becomes prohibitively large, even for the best α values. Additionally, for these higher ratios, high α also start requiring astronomical n in addition to the low α values.

The behavior noted in Fig. 6.5 at first seems to contradict the information in Fig. 6.2 since these two plots claim that opposite regions are worse. This is not a contradiction because the bad cases in Fig. 6.2 have scaling behavior very different from the asymptotic behavior that paradoxically allows them to approach the asymptotic value of the gap faster. Thus, the middle range of α approachs the true asymptotic values of the gap faster at the cost of having vastly different scaling behavior until they reach the appropriate value. However, this faster approach to the asymptotic value is still extremely slow in comparison to the concerns of practical quantum computers, meaning that the benefit of a faster approach cannot be realized.

To explore this phenomena more, analytically consider the ratio between the leading order asymptotic gap and the second order asymptotic gap

$$\begin{aligned}\frac{g_{\min}^{(2^{\mathrm{nd}}\mathrm{O})}}{g_{\min}^{(1^{\mathrm{st}}\mathrm{O})}} &= 1 - \frac{2}{\sqrt{\pi\omega}}(\ln 2 + \psi_0(-\frac{1}{2}))\epsilon^{2\alpha-1/2} \\ &= 1 - \kappa\epsilon^{2\alpha-1/2}, \end{aligned} \tag{6.7}$$

where $\kappa \equiv \frac{2}{\sqrt{\pi\omega}}(\ln 2 + \psi_0(-\frac{1}{2})) \approx 0.8233/\sqrt{\omega}$.

In the case $\omega = 4/3$, so the prefactor, κ, is close to 0.7, meaning any large-n issues with this expression are a direct result of the exponents themselves rather than the coefficients on higher order terms. If I want this ratio to reach some threshold value $0 < v < 1$, then the necessary $n(v)$ becomes

$$n(v) = 2\left(\frac{\kappa}{1-v}\right)^{1/(2\alpha-1/2)}. \tag{6.8}$$

Since $\kappa \approx 0.713$, the expression $\frac{\kappa}{1-v}$ is greater than one for most possible values of $0 < v < 1$, so the size of $n(v)$ depends mostly on the exponent. When $2\alpha - 1/2$ is close to zero, meaning the problem is close to the boundary between the constant and polynomial scaling regions, the exponent is extremely large, leading to large n necessary to reach a specific v.

However, the expression in Eq. 6.8 would seem to indicate that the necessary n should be relatively small for $2\alpha - 1/2$ farther from zero as it is close to the boundary between the polynomial and exponential scaling regions. Therefore the second order expansion cannot explain the large spike in $n(v)$ seen for high $2\alpha - 1/2$ in Fig. 6.5 for large v.

This analysis can be elucidated by looking at the general case of a gap described by $g = A\epsilon^{\gamma_1} + B\epsilon^{\gamma_2}$ for $\gamma_1 < \gamma_2$. This analysis says that the first order expansion with γ_1 becomes poor if $\gamma_2 - \gamma_1$ is small. Thus, much of the analysis is valid for other systems where the second order scaling is close to the first order scaling. Nevertheless, the behavior for $\alpha \approx 1/3$, close to the boundary with the exponential region is not well described by this phenomena, leading to the need for more caution.

6.4 Polynomial vs. Exponential Scaling Issues

Given the discrepancies between the actual gap and the asymptotic gap, an important question becomes how easily scaling behavior can be mislabeled. Given a data set, coming from an experiment or numerical simulation, how difficult is it to tell apart a polynomial scaling algorithm from an exponential or super-polynomial algorithm. Fig. 6.2 already explored how easily the asymptotic scaling power could be misidentified, but that analysis assumed a power law fit from the beginning. In this section, I explore ideas for analyzing how much a data set follows a power law model or an exponential model to judge if finite data indicates polynomial or super-polynomial scaling with n.

Whereas Fig. 6.2 looked at a broad sweep of data in an attempt to extract scaling information, this section is concerned with small scale features of curves, asking how well a curve can locally be approximated by a certain functional form. The ultimate goal is to be able to take data like that presented in Fig. 6.3 and judge how likely it is that the data is obeying a power law or exponential relationship.

I look at both a power law and exponential relationship, and for ease of writing I define $f \equiv \ln g_{\text{min}}$ and $x \equiv \ln n$:

$$g_{\text{min}} = An^{-p} \Rightarrow f_{\text{poly}} = \ln A - px, \tag{6.9}$$

$$g_{\text{min}} = B\exp(-Cn^{q}) \Rightarrow f_{\text{exp}} = \ln B - Ce^{qx}. \tag{6.10}$$

My analysis technique relies on properties of f and its derivatives with respect to x. In Fig. 6.2, when I do linear fits to log-log data, I am already looking at the properties of $f'(x)$.

Two natural other quantities to look at would be $f''(x)$ and the curvature of the function $f(x)$. The curvature is not a good metric since the curvatures of both f_{poly} and f_{exp} go to zero for $x \to \infty$. The second derivative $f''(x)$ is a better metric since $f''_{\text{poly}}(x)$ goes is zero; whereas, $f''_{\text{exp}}(x)$ goes to $-\infty$ as $x \to \infty$. The problem with $f''(x)$ is that you cannot immediately extract any additional information from it about the exponential.

6.4.1 Quality of Fits

This proposed analysis technique relies on another important feature of the derivatives of these log-log functions. Specifically, I can look at the ratio, R, of the second derivative to the derivative:

$$
\begin{aligned}
R_{\text{poly}} &= \frac{f''_{\text{poly}}(x)}{f'_{\text{poly}}(x)} = 0 \\
R_{\text{exp}} &= \frac{f''_{\text{exp}}(x)}{f'_{\text{exp}}(x)} = q.
\end{aligned} \tag{6.11}
$$

Note that only exponential scaling gives a non-zero, constant derivative ratio which gives me another metric to judge if a function has exponential scaling and allows me to discover q. In general, sub-exponential scaling takes this ratio to zero for large n. True polynomial scaling has a ratio of exactly zero, but given that the ratio does not immediately reach asymptotic scaling, asymptotic polynomial and sub-exponential scaling are not immediately distinguishable through this method.

In Fig. 6.6, I plot the ratio $R = f''(x)/f'(x)$ for the continuous gap for α values that should have asymptotically exponential scaling. The thin dashed lines represent what the curve of the same color should approach in its asymptotic limit. For lower α, close to the boundary with the polynomial region, the R vs. n curves do not come close to their asymptotic values, even for the extremely large n shown, up to $n = 10^{25}$. Some of the more strongly exponential systems do approach their asymptotic values for the n displayed, but these systems still require extremely large n.

Based on the nature of the ratio R as a function of $x = \ln n$, an asymptotically increasing or non-zero constant ratio would imply a super-exponential or exponential scaling of g_{min} with n, so it could be argued that these plots imply non-polynomial behavior much earlier. Even with this argument, many of the borderline α cases continue decreasing for a large range of n and do not have an upturn until much larger n. Additionally, if the goal is to extract information on the type of exponential scaling, Fig. 6.6 implies that very large n are needed.

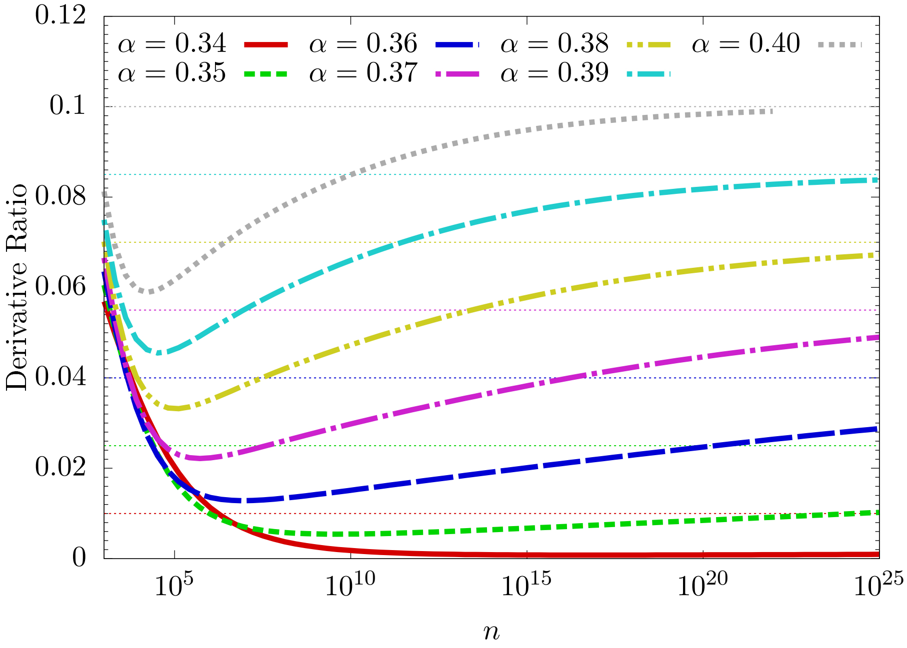

Figure 6.6: The derivative ratio $R = f''(x)/f'(x)$ where $f \equiv \ln g_{\min}$ and $x \equiv \ln n$ for the continuous gap. For true exponential scaling, this ratio should approach a constant equal to the power of n in the exponential. The dashed lines are what this ratio should be asymptotically. None of these curves do a good job approximating their asymptotic constant value, and even the very large barrier cases require large n to even approach the correct constant. This plot indicates that even the exponential scaling region does not approach its asymptotic expression at reasonable n.

Asymptotically polynomial scaling should give an R ratio that is zero, so in Fig. 6.7, I have plotted R as a function of n for several α values that asymptotically have polynomial scaling. All of these curves, except the solid line for $\alpha = 0.34$, should approach zero, but many of them take extremely large n to actually approach this asymptotic value. Notice especially that α closer to 1/4, the boundary between constant and polynomial scaling, require even larger n to reach zero. This behavior of small α corresponds well with what is seen in Fig. 6.5 in that lower α take a long time to reach the correct asymptotic value.

Also notice that the solid, red line for $\alpha = 0.34$ in Fig. 6.7 appears to approach zero faster than many of the ratios for α actually in the polynomial scaling region. Therefore,

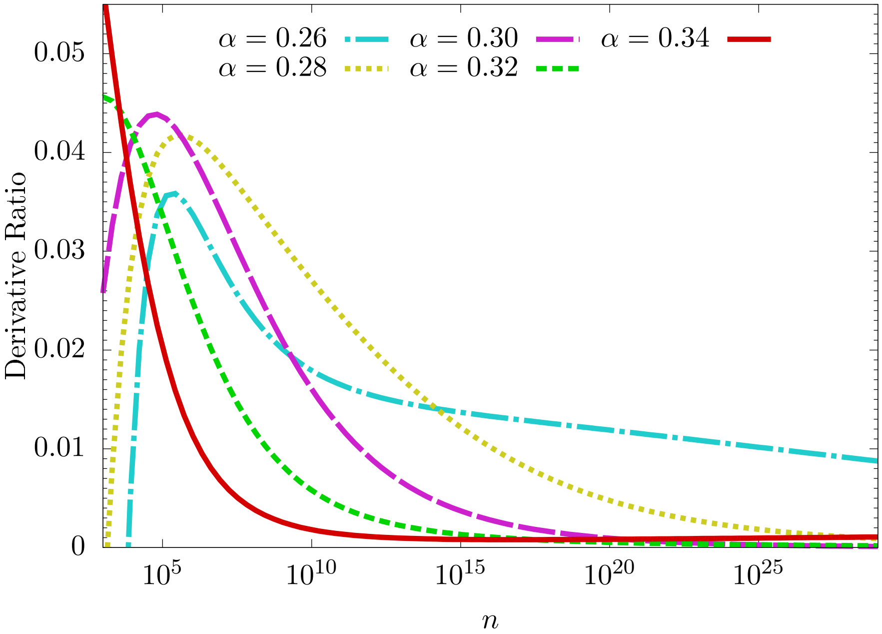

Figure 6.7: The derivative ratio $R = f''(x)/f'(x)$ where $f \equiv \ln g_{\min}$ and $x \equiv \ln n$ for the continuous gap. These α values were chosen in the polynomial scaling region, and for true power law scaling, I would expect the ratio $R = 0$. I do see the curves approaching zero, but all of them require extremely large n to do so, with several of the smaller barrier cases being far away from zero even at the extreme end of the simulation range. The solid, red line is for $\alpha = 0.34$ in the exponential region; notice that the ratio appears to approach zero faster than many of the ratios for αs in the polynomial scaling region.

for a large range of n, it would be easier to interpret $\alpha = 0.34$ as having a power law scaling behavior than many of the α close to $1/4$ which take very large n to display the appropriate ratio behavior.

The slowness of lower α systems in reaching their asymptotic value is at least partly due to the size of their scaling powers. Fig. 6.2 indicates that these α values close to the constant boundary reach their asymptotic scaling behavior more quickly. However, their asymptotic scaling behavior is a very small power law, meaning that any change in these functions will be small. By reaching their asymptotic scaling earlier, it will take these low α systems even longer to reach the correct asymptotic value for their scaling, as seen in Figs. 6.5 and 6.7.

Chapter 7

Non-Adiabatic Runtimes

Numerical evidence presented by Muthukrishnan, Albash, and Lidar [27] shows that for the symetric multiqubit barrier tunneling problem on n qubits, a well-calibrated constant time (i.e. τ is independent of n) evolution of quantum annealing Hamiltonian is sufficient. Thus, this so-called diabatic evolution significantly improves upon the slower adiabatic evolution of the Hamiltonian, which could take polynomial or even exponential time in n, as discussed in Ch. 5. Muthukrishnan *et al.* attribute this speedup to a diabatic cascade in which the ground state is quickly depopulated in favor of higher excited states and then repopulated right at the end of the diabatic evolution.

The adiabatic condition, discussed in Ch. 2, given by

$$\tau \gg \max_{s\in[0,1]} \frac{|\left\langle\psi_0\right| \frac{d\hat{H}}{ds} \left|\psi_1\right\rangle|}{\Delta(s)^2}, \tag{7.1}$$

is merely a sufficient condition, and it is possible to have adiabatic evolutions with shorter running times. Furthermore it is also possible to have a non-adiabatic evolution that succeeds in solving the optimization problem at hand. It is such a non-adiabatic speedup that is described by Muthukrishnan *et al.* [27].

A non-adiabatic speedup is obviously significant for near-term quantum computers where quantum annealing is a potential application. Kong and Crosson [43] have studied these diabatic transitions, and more recently my advisor presented joint work with me that provided complementary findings [57]. These recent results indicate that this non-adiabatic speedup can provide an alternate and efficient way of solving an important class of Hamming-symmetric barrier tunneling problems that are being used as toy models [27, 43, 53, 42, 40, 23, 24, 19] to study the more general properties of quantum annealing in the presence of a barrier.

In this chapter I present results that indicate that even slightly more generalized versions of symmetric barrier tunneling problems do not exhibit this fast non-adiabatic speedup. The base Hamiltonian used to study this class of problems is more general than the barrier problem I cited in Ch. 3,

$$\hat{H}(s) = -(1-s)\sum_{i=1}^{n}\sigma_x^{(i)} + s\left[\sum_{i=1}^{n}\sigma_z^{(i)} + b\left(\sum_{i=1}^{n}\sigma_z^{(i)}\right)\right], \tag{7.2}$$

since $b(h)$ is some localized barrier, that is not necessarily localized around $n/4$, but could be more generalized. Current numerical evidence [27] suggests that the non-adiabatic speedup exists for many classes, shapes, and sizes of localized barriers $b(h)$. This chapter generalizes the problem slightly (ignoring $b(h)$ for the moment):

$$\hat{H}(s) = -(1-s)\sum_{i=1}^{n}\sigma_x^{(i)} + s\mu\sum_{i=1}^{n}\sigma_z^{(i)}, \tag{7.3}$$

by introducing a positive slope parameter μ, and I find that for the generic case $\mu \neq 1$, the non-adiabatic speedup no longer exists. I call μ a slope as it relates linearly the energy of the system with the Hamming weight $\sum_i \sigma_z^{(i)}$ of the n qubits. In Ch. 3, I went ahead and included μ, but the work presented in this chapter is the first chronologically to consider this generalization of the barrier tunneling problem.

Since this Hamiltonian describes a simple toy model, it is unlikely that a physical system will exhibit the exact $\mu = 1$ behavior, leading to the conclusion that for realizable problems, this diabatic speedup will not exist. In this chapter, I focus on the $b(h) = 0$ case since it decouples all the qubits, allowing the extraction of information about the system by studying the evolution of a single qubit Hamiltonian. Since $\mu \neq 1$ disrupts the non-adiabatic speedup even in this $b(h) = 0$ case, I fully expect similar disruption to occur for more complicated barriers and perturbations.

The work presented in this chapter is joint work with Wim van Dam and appears in the article [21].

7.1 Optimal Runtime

I first need to define my criteria for an optimal runtime. If an algorithm on n qubits runs for time τ and has a probability of success of $p_n(\tau)$ at the end of that time, its expected running time is $\tau/p_n(\tau)$, and the optimal running time is the τ_n that minimizes $\tau/p_n(\tau)$ for n qubits. In my case, there are n independent qubits, each of which has a probability of success of p_1, hence $p_n = p_1^n$, which is where the n dependence comes into the minimization.

In the $\mu = 1$ case, p_1 goes to 1 for finite τ, as seen in Fig. 7.1, meaning that $p_n = 1$ at this value, leading to the non-adiabatic speedup noted in other studies. Fig. 7.1 also shows $\mu = 0.5$ and $\mu = 2$ curves. Note that for these curves the success probability does not achieve $p_1 = 1$ at finite τ. Similar plots can be obtained for other $\mu \neq 1$ and, as noteed below, this failure to reach $p_1 = 1$ for finite τ, ultimately leads to the breaking of the non-adiabatic speedup. Therefore, this speedup is restricted to the special case of $\mu = 1$.

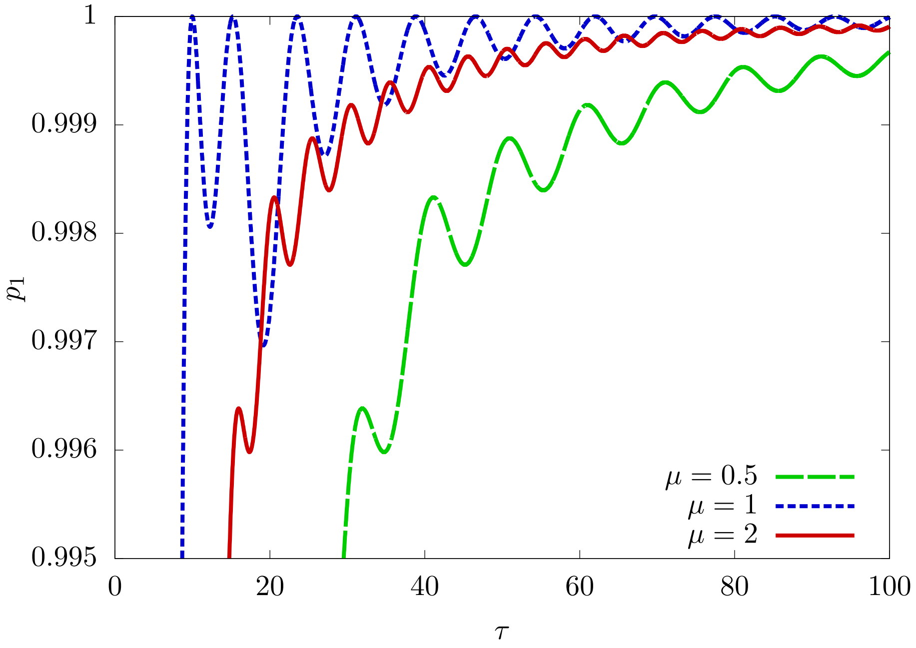

Figure 7.1: The single qubit success probability, p_1, as a function of the total runtime for several μ values. The blue, dashed, $\mu = 1$ line corresponds to the model that has been studied in previous articles. Notice that the $\mu = 1$ curve has several special properties, including that it goes to $p_1 = 1$ at finite τ, resulting in the non-adiabatic speedup noted in other papers. The $\mu \neq 1$ curves do not exhibit this $p_1 = 1$ behavior.

Muthukrishnan *et al.* [27] note that for $\mu = 1$ the optimal running time decreases asymptotically to a constant in the case with a barrier because p_n increases for fixed τ at the optimal running time. The decoupled model does not exhibit this behavior because the success probability is $p_n = 1$ independent of n. Therefore, in the barrier cases, the success probability for $\mu = 1$ seems to be approaching its value in the no barrier case.

To demonstrate the lack of a non-adiabatic speedup in the $\mu \neq 1$ cases, consider Fig. 7.2, which shows the optimal expected runtime, $\tau_n/p_n(\tau_n)$, as a function of n. All of the μ curves shown are increasing, meaning that the running time increases with n, and there is no non-adiabatic algorithm that runs in constant time. The fitted curves are to power laws of the form $\tau_n = An^p$, and all of the fitted p values are close to 1/2,

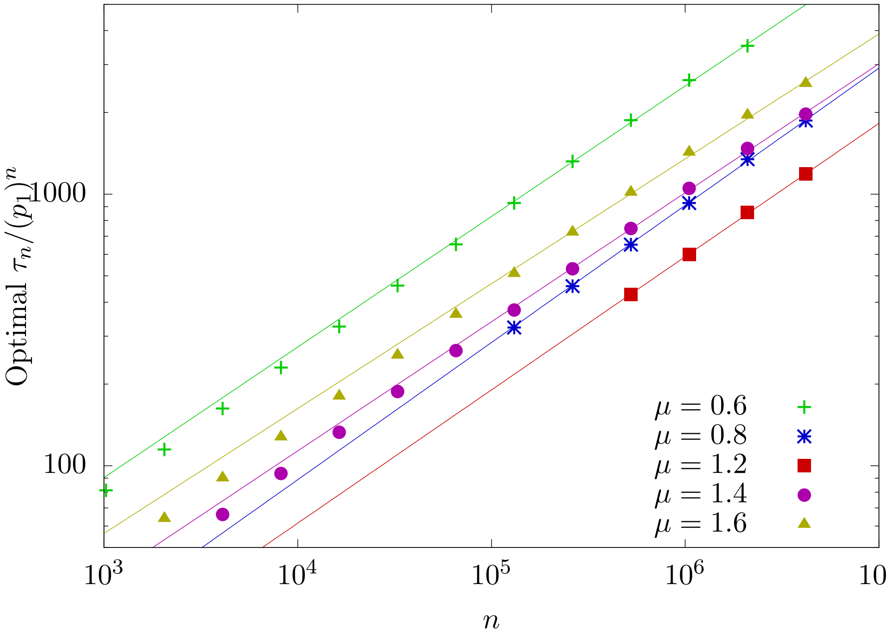

Figure 7.2: Optimal expected running time of quantum annealing, $\tau_n/p_n(\tau_n)$, as a function of n for different μ values. Unlike the $\mu = 1$ case, τ_n increases with n for these μ values. lines through the data are power law fits of the form $\tau_n = An^r$, and the fitted r values in the order $\mu = (0.6, 0.8, 1.2, 1.4, 1.6)$ are $(0.48, 0.51, 0.49, 0.48, 0.46)$, all close to $1/2$. A scaling power of $1/2$ is consistent with the adiabatic scaling of the $\mu = 1$ case as found in [27] and my results below while being quadratically faster than the sufficient adiabatic condition.

indicating a running time of $\mathcal{O}(\sqrt{n})$.

It should be noted that the non-adiabatic speedup could be reinstated by modifying the driver Hamiltonian, $\hat{H}_0$, by multiplying it by μ as well. However, this kind of fine-tuning of the driver Hamiltonian to match up with the properties of the final Hamiltonian implies a large amount of knowledge about the final problem that would not be known in a realistic setting. The lack of robustness in the system to changes in μ could be viewed alternatively as a lack of robustness in the fine-tuning of the driver Hamiltonian. The examples of Muthukrishnan *et al.* [27] show that a fine-tuned driver Hamiltonian can solve these problems non-adiabatically, but this chapter shows that this speedup does

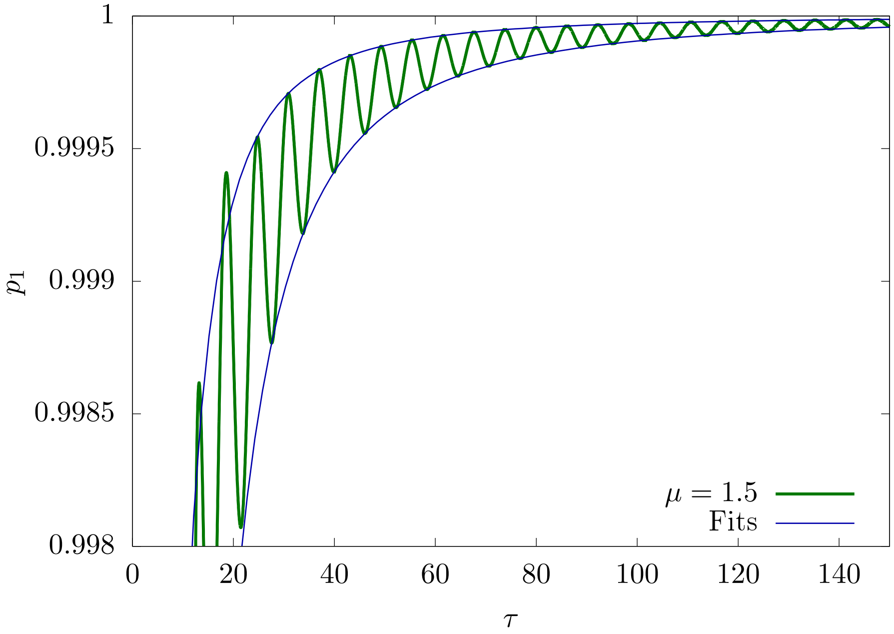

Figure 7.3: A single qubit success probability curve as a function of total runtime τ for $\mu = 1.5$ with upper and lower envelopes shown. These envelopes are obtained by first extracting all the local minima (maxima) and doing a power law fit of the form $p_1 = 1 - c\tau^{-q}$. The first two minima and maxima were excluded from this fit and others since they tend to be more abnormal. In this case, the upper envelope has a fitted $q = 1.998$ and the lower envelope has a fitted $q = 1.996$, both of which are extremely close to the 2 noted in the text.

not allow for a general algorithm but only specific algorithms tailored to the problem at hand.

7.2 Single Qubit Success Probability

The $\sqrt{n}$ running time behavior can be extracted from the curves in Fig. 7.1 as well because the qubits in this problem are completely decoupled. For sufficiently large running times τ, the curves of the single qubit success probability p_1 as a function of τ shown in Fig. 7.1 are bounded above and below by envelopes of the form

$$1 - \frac{c_\ell(\mu)}{\tau^q} < p_1 < 1 - \frac{c_u(\mu)}{\tau^q}, \tag{7.4}$$

with constants $c_\ell(\mu)$ and $c_u(\mu)$. This relationship is extracted by performing numerical fits to the minima and maxima in curves like those seen in Fig. 7.3, and for all the fits to different μ data, q is close to 2. The form of these bounds, including $q = 2$, arise naturally from large τ asymptotic expansion of the success probability, as seen in Ch. 8. Note that $c_u(1) = 0$, which is one of the main reasons why the $\mu = 1$ diabatic speedup can occur.

Muthukrishnan *et al.* [27] showed that the lower envelope with $c_\ell(\mu)$ guarantees that the worst case running-time for the $\mu = 1$ case scales as $\mathcal{O}(\sqrt{n})$. I employ their method to show that a relationship such as Eq. 7.4 provides both the necessary and sufficient condition for the running time. Muthukrishnan *et al.* also apply methods created by Boixo and Somma [58] to show that at least $\Omega(n^{1/2})$ is necessary for adiabatic evolution.

If for n qubits a total success probability of p is desired from the algorithm, then Eq. 7.4 gives

$$\left(1 - \frac{c_\ell(\mu)}{\tau^q}\right)^n \leq p \leq \left(1 - \frac{c_u(\mu)}{\tau^q}\right)^n. \tag{7.5}$$

I can manipulate this inequality, performing an expansion for small $c_*(\mu)/\tau^q$ since τ will be large. The result of these manipulations gives a relationship between the running time and n:

$$\left(\frac{c_u(\mu)}{\ln 1/p}\, n\right)^{1/q} \leq \tau \leq \left(\frac{c_\ell(\mu)}{\ln 1/p}\, n\right)^{1/q}. \tag{7.6}$$

Therefore, since $q = 2$ in this case, having a running time that scales as $\sqrt{n}$ is both a necessary and sufficient condition to reaching a desired probability. Note that when $\mu = 1$, $c_u(1) = 0$, so one side is no longer bounded, leading to the possibility of a non-adiabatic speedup.

In the Hamming weight problem, the gap is constant with n, and all matrix norms of the Hamiltonian and its derivatives will depend linearly on n. Therefore, the adiabatic condition, Eq. 2.1, would predict $\mathcal{O}(n)$ scaling; whereas, my results indicate that a faster $\mathcal{O}(\sqrt{n})$ running time is sufficient. This result was shown in [27] for $\mu = 1$, and my results indicate that this quadratic speedup holds for general slopes μ.

While the standard adiabatic condition overestimates the running time, there are other derivations that apply to this problem more specifically and that provide a stricter bound that matches these results. Jansen, Ruskai, and Seiler [31] showed that for fixed Hamiltonians $\hat{H}_0$ and $\hat{H}_1$ with time evolution $\hat{H}(t) = (1 - t/\tau)\hat{H}_0 + t/\tau\hat{H}_1$, the success probability p of remaining in the ground state throughout $0 \leq t \leq \tau$ is bounded by

$$p = 1 - \mathcal{O}(\tau^{-2}). \tag{7.7}$$

If this is the probability of success for a single qubit case, the results in Eqs. 7.5 and 7.6 imply that $\tau \in \mathcal{O}(\sqrt{n})$ is sufficient for an adiabatic evolution. This shows that the result from Jansen *et al.* provides a stricter sufficient condition than the standard adiabatic condition for the optimization problem with decoupled qubits.

In Fig. 7.4 I plot the coefficients $c_u(\mu)$ and $c_\ell(\mu)$ obtained from numerical fits. The fits used to obtain these values are akin to those shown in Fig. 7.3, giving confidence in the $1/\tau^2$ scaling of the error. Notice that as you approach the special case $\mu = 1$, $c_u(\mu) \to 0$ and observe that around $\mu = 1$ the coefficient $c_u(\mu)$ stays close to zero. Hence for μ approximately (but not exactly) 1, the non-adiabatic speedup persists for a large range of n until the adiabatic running time of $\mathcal{O}(\sqrt{n})$ is required again at very large n.

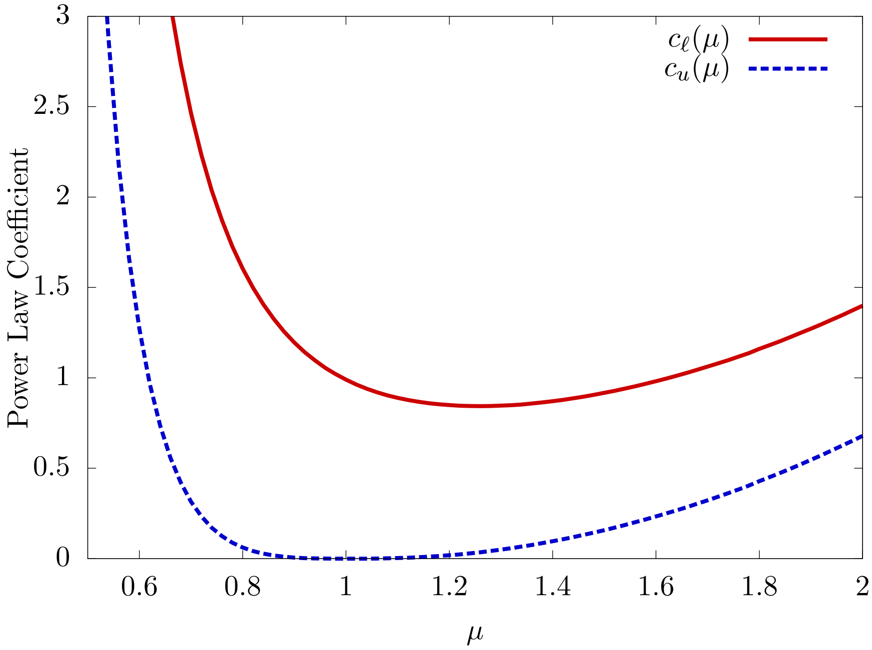

Figure 7.4: The curves like the one in Fig. 7.3 are bounded above and below by curves of the form $1 - c/\tau^2$. I show the values of c for the upper, c_u, and lower, c_ℓ, bounding functions as obtained from numerical fits. These coefficients are a function of μ, and all of the fits used to obtain this data were good quality. In the main text, I show that these bounding curves directly lead to a $\mathcal{O}(\sqrt{n})$ running time for the algorithm in all cases except the $\mu = 1$ case where $c_u(1) = 0$.

7.3 Scaling of True Adiabatic Run Time

All of the work so far has shown that the optimal running time of this algorithm is $\mathcal{O}(\sqrt{n})$, but this does not imply that the optimal running time results from adiabatic evolution. Looking at the occupancy of the energy states for these optimal runs, the ground state is depopulated during the s evolution. Therefore, a remaining question to ask is whether this behavior also holds if the system is required to stay within a certain range of its ground state for the entire $s \in [0, 1]$ evolution.

In Fig. 7.5, I show the time, τ, needed to ensure that the system has at least a 75% chance of being measured in its ground state for the entire $s \in [0, 1]$ evolu-

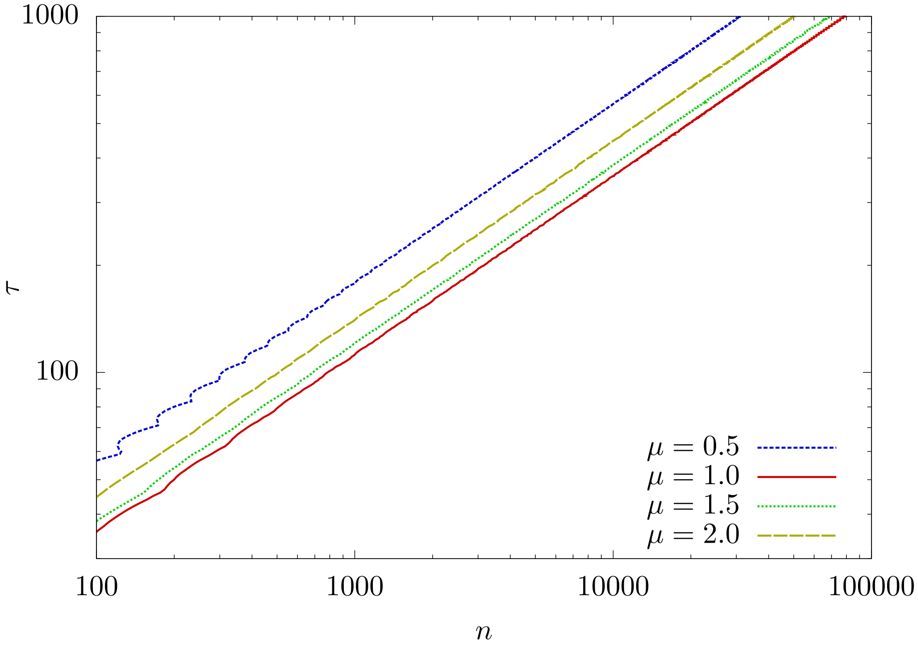

Figure 7.5: These plots shows the runtime, τ, needed to ensure that state of the system is at least 75% in the ground state over the entire s evolution. This growth of τ with n comes closest to a true adiabatic evolution, and the $\tau \in \mathcal{O}(\sqrt{n})$ behavior holds even in this case. Power law fits to these data sets show that the exponent for these curves, in the order $\mu = (0.5, 1.0, 1.5, 2.0)$, are $(0.497, 0.502, 0.501, 0.500)$. Therefore, the quadratic speedup over the sufficient adiabatic condition is a property of adiabatic evolution in this system, not the specific τ/p_n criteria.

tion. All of these curves exhibit power law relationships, $\tau = Bn^r$, with fitted $r = (0.497, 0.502, 0.501, 0.500)$ for $\mu = (0.5, 1.0, 1.5, 2.0)$ respectively. A similar plot can be obtained if a stricter cutoff than 75% is used.

Fig. 7.5 shows that the runtime relationships I observe are in fact indicative of how adiabatic evolution behaves as well. Therefore, I am led to the conclusion that for general $\mu \neq 1$, the runtime $\tau \in \Theta(\sqrt{n})$ is both necessary and sufficient to ensure finding the true ground state. The $\mu = 1$ case remains a special case that goes against this rule, allowing for an extreme speedup to a constant running time.

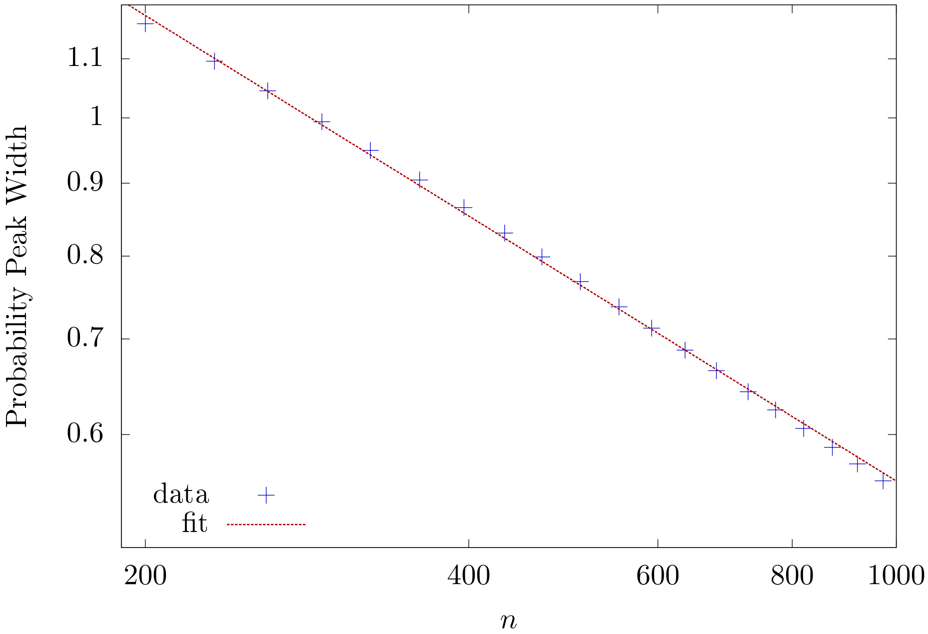

Figure 7.6: The width of the non-adiabatic success probability peak in the curve of p_n vs. τ is shown as a function of n. This problem tunnels through a binomial barrier of width and height $\propto n^0.3$. The width of the peaks in success probability are extracted by performing a Gaussian fit to the top of the peak and extracting the width of the fitted Gaussian. The width is modeled well by the curve $w = An^{-p}$ where $A = 13.97$ and $p = 0.467$. This width scaling with n is close to the width scaling of $n^{-1/2}$ extracted analytically for the no barrier case.

7.4 Width of Non-adiabatic Success Peak

My last goal is to understand the width of the success probability spike of p_1 in the unperturbed, $\mu = 1$ case when it reaches the optimal $p_1 = 1$. This narrowness implies that to be successful for large n, one has to be very precise in using the right running time τ.

There is a critical runtime τ_c such that $p_1 = 1$ for a single qubit. For run times close to this τ_c, the probability of success can be modeled by

$$p_1 = 1 - \delta = 1 - k(\tau - \tau_c)^2, \qquad \delta \ll 1, \tag{7.8}$$

where $|\tau - \tau_c|$ is the required stopping precision of the algorithm.

Scaling the system to n qubits, the probability of success is $p_n = p_1^n$ since the qubits are uncoupled in the unperturbed case:

$$p_n = \left(1 - k(\tau - \tau_c)^2\right)^n \approx 1 - nk(\tau - \tau_c)^2. \tag{7.9}$$

If you want the probability of failure to be less than ϵ, then

$$1 - \epsilon < 1 - kn(\tau - \tau_c)^2 \Rightarrow |\tau - \tau_c| < (\epsilon/kn)^{1/2}. \tag{7.10}$$

Thus, maintaining the same success probability as n increases requires the acceptable imprecision $|\tau - \tau_c|$ to shrink according to $n^{-1/2}$. Note that this $\sqrt{n}$ width scaling behavior is independent and unrelated to the adiabatic scaling of run time.

I have run simulations using binomial shaped barriers to get a sense of this width scaling when a barrier is present. For $\mu = 1$, the narrowing of the spiked success probability p_n around the critical τ_c running time is close to $n^{-1/2}$. In Fig. 7.6, I show the width of the probability peak as a function of n for a binomial barrier with height and width proportional to $n^{0.3}$. This width is extracted by looking at the non-adiabatic optimal success probability peak and doing a Gaussian fit to the data close to the peak. This fit can ignore the fact that the peak is not directly at $p_n = 1$ and allows me to extract an approximation of the width of the peak.

The widths for the barrier case in Fig. 7.6 are well modeled by the curve $w = An^{-p}$, with a fitted value of $p = 0.467$. This fitted exponent is close to what our analytics for the no barrier case predict. Other barrier shapes and sizes exhibit similar scaling in the width of their non-adiabatic success probability peak.

7.5 Discussion

While the $\mu = 1$ case does exhibit a surprising non-adiabatic speedup that could potentially be exploited, this diabatic speedup is not a general feature of this class of quantum annealing problems. Even small changes to μ are enough to alter the evolution and eliminate the speedup.

This non-adiabatic speedup had also been noted in the semi-classical limit of the Hamming-symmetric tunneling problem in a classical method called Spin Vector Dynamics [27]. I performed simple simulations of Spin Vector Dynamics for no barrier and $\mu \neq 1$, and observed the same breaking of the non-adiabatic speedup seen in quantum annealing. Therefore, the non-adiabatic speedup does not survive in either the classical or quantum settings.

In addition, I show that even in the $\mu = 1$ case, the non-adiabatic speedup requires inverse polynomial precision in the runtime to achieve the speedup. Thus, it is difficult to hit the speedup if it exists. However, even in the adiabatic setting, this problem shows $\sqrt{n}$ running time which is better than the linear running time predicted by the adiabatic condition.

The Hamming-symmetric qubit problem has been well studied explicitly because it is simple enough to admit analytic work. The fact that simple changes can be made to this system to eliminate an atypical non-adiabatic speedup shows the robustness of this toy model. A small change to the model is enough to bring the model in line with what should be expected of most physical systems. Running these algorithms adiabatically remains the best and only option to achieve success in general.

Chapter 8

Near Adiabatic Evolution

The Quantum Adiabatic Theorem [28] is a powerful tool for analyzing dynamical quantum systems. For slowly evolving Hamiltonians it ensures the system will closely track its originally initialized energy state throughout the entire evolution. The key point of the adiabatic theorem is that it also gives a condition for how slowly the system needs to evolve. This chapter will explore what happens when you run quantum annealing faster than the adiabatic limit suggests.

In this chapter, I will focus on the near-adiabatic regime of quantum evolution and quantum annealing, where the system evolves slowly enough to mostly stay in the desired eigenstate but with noticeable leakage. Non-adiabatic evolution has garnered much interest because it can potentially lead to speed-ups that the adiabatic theorem does not account for. In quantum annealing, recent work has focused on a rapid diabatic speed-up in certain barrier tunneling models [43, 27, 21] (also Ch. 7).

While I couch my findings in the language of quantum computing, my results are more general. Many of the assumptions and approximations are general and can apply to a large number of problems and settings. QAO relies on ground state evolution, and even this condition can be relaxed in my results. I will use both the Hamming weight barrier model and the cubic p-spin model (see Ch. 3), but these are used solely as examples to provide support for my analytic results.

In the near-adiabatic limit, my findings show that in the absence of a small spectral gap, the probability of transitioning out of the initialized state oscillates as a function of the total evolution time, τ. Furthermore, I confirm previous results [59, 30, 60] that show that the frequency of this oscillation depends on the integral of the spectral gap over the evolution.

My new results add a layer of depth by considering the case where the spectral gap becomes small during one portion of the evolution, as would occur in a Landau-Zener avoided crossing [61, 62]. This crossing is localized, so that it only effects the system during a short period of time. I find that the avoided crossing effectively splits the evolution in two, resulting in a superposition of oscillatory behavior in τ.

In Section 8.1, I develop much of the mathematical machinery that will be used in later sections. I study the large gap regime in Section 8.2, examining how the oscillatory transition probability behavior arises, and I back up these results with numerical evidence from quantum algorithm settings. In Section 8.3, I add in an avoided level crossing and explore analytically and numerically how this leads to a more complicated superposition of oscillatory behavior in total evolution time, τ. Section 8.4 shows an application of the large gap oscillations to the adiabatic version of Grover's search.

The work in this chapter originally appeared in an article co-written with Wim van Dam [22].

8.1 Setup

Suppose some quantum system obeying the Schrödinger equation

$$i\frac{d}{dt}\left|\psi\right\rangle = \hat{H}(t)\left|\psi\right\rangle, \tag{8.1}$$

where time runs between $t = 0$ and τ, the evolution time of the system or the run time in a quantum computing setting. The Hamiltonian should follow the same trajectory even for different values of τ, so I can use the "normalized time" $s \in [0, 1]$ to determine where in the Hamiltonian's evolution the system is.

The normalized time s relates to the actual time through $t = s\tau$. Since the form of the Hamiltonian depends only on s and not τ, I can rewrite the Schrödinger equation as

$$i\frac{d}{ds}\left|\psi\right\rangle = \tau\hat{H}(s)\left|\psi\right\rangle. \tag{8.2}$$

Now all the evolution time information has been pulled out into one parameter, τ, which can vary to run the evolution more slowly or quickly.

The next step involves rewriting the equation in the eigenbasis of the Hamiltonian. I take the instantaneous eigenbasis to be given by $\left|\varphi_j(s)\right\rangle$ with associated eigenenergies $\lambda_j(s)$. Then a general state of the system is written as

$$\left|\psi(s,\tau)\right\rangle = \sum_j C_j(s,\tau)\left|\varphi_j(s)\right\rangle. \tag{8.3}$$

In terms of the eigenbasis the Schrödinger equation gives

$$i\left[\frac{dC_k}{ds} + \sum_j C_j \left\langle\varphi_k\right| \frac{d}{ds}\left|\varphi_j\right\rangle\right] = \tau\lambda_k C_k. \tag{8.4}$$

I want to know how the state evolves if the system starts in the ground state at $s = 0$. These same arguments work for higher excited states as well, but for my purposes, the ground state is sufficient and simpler. The adiabatic theorem tells me that if τ is large enough, the system remains in the ground state. I relax the adiabatic condition slightly and allow near-adiabatic evolution where the majority of the state remains in $|\varphi_0(s)\rangle$, but some small amount leaks into the first excited states, with all other states being essentially unvisited.

The ground state needs to be nondegenerate, but the first excited eigenstate can be degenerate. The possible m-fold degeneracy is denoted with a superscript $C_1^{(a)}$.

I want to restrict down Eq. 8.4 to just those probability amplitudes that are assumed to be relevant, namely, those close to the ground state. In doing this, remember that $\langle\varphi_j|\frac{d}{ds}|\varphi_j\rangle = 0$. I also shift the Hamiltonian by an overall (s-dependent) constant so that $\lambda_0 = 0$, which means that $\lambda_1 = \Delta$ is just given by the spectral gap. Using all this information, I obtain the following coupled equations for the relevant amplitudes

$$i\frac{dC_0}{ds} + i\sum_{a=1}^{m} C_1^{(a)} \langle\varphi_0|\frac{d}{ds}|\varphi_1^{(a)}\rangle = 0 \tag{8.5}$$

$$\begin{aligned} i\frac{dC_1^{(a)}}{ds} &+ iC_0 \langle\varphi_1^{(a)}|\frac{d}{ds}|\varphi_0\rangle \\ &+ i\sum_{b\neq a} C_1^{(b)} \langle\varphi_1^{(a)}|\frac{d}{ds}|\varphi_1^{(b)}\rangle = \tau\Delta C_1^{(a)} \end{aligned} \tag{8.6}$$

For $\langle\varphi_1^{(a)}|\frac{d}{ds}|\varphi_1^{(b)}\rangle$, I can freely choose the basis within the degenerate eigenspace, and it is possible and desirable to choose the degenerate eigenbasis such that $\langle\varphi_1^{(a)}|\frac{d}{ds}|\varphi_1^{(b)}\rangle = 0$ for all a and b.

Just from the definition of the eigenvalues of $\hat{H}(s)$ and the orthonormality of its eigenvectors, I can look at

$$
\begin{aligned}
\frac{\partial}{\partial s}\left(\hat{H}\left|\varphi_i\right\rangle\right) &= \frac{\partial}{\partial s}\left(\lambda_i\left|\varphi_i\right\rangle\right) \\
\dot{\hat{H}}\left|\varphi_i\right\rangle + \hat{H}\frac{\partial}{\partial s}\left|\varphi_i\right\rangle &= \frac{\partial \lambda_i}{\partial s}\left|\varphi_i\right\rangle + \lambda_i\frac{\partial}{\partial s}\left|\varphi_i\right\rangle .
\end{aligned} \tag{8.7}
$$

I can look at the inner product of this time derivative with an eigenstate $j \neq i$ and utilize the orthonormatility of the eigenstates

$$
\begin{aligned}
\left\langle\varphi_j\right|\dot{\hat{H}}\left|\varphi_i\right\rangle + \left\langle\varphi_j\right|\hat{H}\frac{\partial}{\partial s}\left|\varphi_i\right\rangle &= \frac{\partial \lambda_i}{\partial s}\left\langle\varphi_j|\varphi_i\right\rangle + \lambda_i\left\langle\varphi_j\right|\frac{\partial}{\partial s}\left|\varphi_i\right\rangle \\
\left\langle\varphi_j\right|\dot{\hat{H}}\left|\varphi_i\right\rangle &= \left(\lambda_i - \lambda_j\right)\left\langle\varphi_j\right|\frac{\partial}{\partial s}\left|\varphi_i\right\rangle .
\end{aligned} \tag{8.8}
$$

which means that

$$
\left\langle\varphi_0(s)\right|\frac{d}{ds}\left|\varphi_1^{(a)}(s)\right\rangle = -\left\langle\varphi_1^{(a)}(s)\right|\frac{d}{ds}\left|\varphi_0(s)\right\rangle \tag{8.9}
$$

$$
=\frac{\left\langle\varphi_0(s)\right|\frac{d\hat{H}}{ds}\left|\varphi_1^{(a)}(s)\right\rangle}{\Delta(s)} \equiv \frac{\gamma_a(s)}{\Delta(s)}. \tag{8.10}
$$

Thus, the differential equations can be reduced to

$$
\frac{dC_0}{ds} + \sum_{a=1}^{m} C_1^{(a)}\frac{\gamma_a}{\Delta} = 0 \tag{8.11}
$$

$$
i\frac{dC_1^{(a)}}{ds} - iC_0\frac{\gamma_a}{\Delta} = \tau\Delta C_1^{(a)}. \tag{8.12}
$$

At this point it is obvious that all the m-fold degenerate first excited states will behave in the exact same way. Therefore, for notational convenience, I will drop the a index and make the m-fold degeneracy explicit in these differential equations:

$$\frac{dC_0}{ds} + mC_1\frac{\gamma}{\Delta} = 0 \tag{8.13}$$

$$i\frac{dC_1}{ds} - iC_0\frac{\gamma}{\Delta} = \tau\Delta C_1. \tag{8.14}$$

8.2 Large Gap

8.2.1 Analytic Approximation

Now, I take the near-adiabatic limit by assuming that the vast majority of of the state remains in the ground state. In this limit, $C_0 \gg C_1$ so that Eq. 8.13 reduces to $\frac{dC_0}{ds} = 0$. This reduction assumes that the gap Δ does not become too small. If the gap becomes small, then this problem can be approximated using different techniques, such as the Landau-Zener transition in the next section, but I focus on the large gap case in this section.

Given this approximation that $\frac{dC_0}{ds} = 0$, I assume that it is a good approximation that $C_0(s) = 1$ for the entire evolution. Essentially, the majority of the amplitude remains in the ground state with minimal changes to its value. With this assumption, Eq. 8.14 becomes

$$i\frac{dC_1}{ds} - i\frac{\gamma}{\Delta} = \tau\Delta C_1. \tag{8.15}$$

Notice that C_1 is kept in this equation despite being disregarded in the C_0 equation. It is kept both because I am now looking at the change in C_1 itself and because this otherwise small term is multiplied by τ which is taken to be large in the near-adiabatic limit. This differential equation has an integral solution when $C_1(s=0) = 0$

$$C_1(s,\tau) = \int_0^s dx \frac{e^{-i\tau\int_x^s dz\,\Delta(z)}}{\Delta(x)/\gamma(x)}. \tag{8.16}$$

What I am actually going to care about is the final amplitude after the total evolution time τ, so the quantity the matters is

$$C_1(1,\tau) = \int_0^1 ds \frac{e^{-i\tau \int_s^1 dz\, \Delta(z)}}{\Delta(s)/\gamma(s)} \tag{8.17}$$

The goal is to approximate this integral in the limit of large τ to find the probability amplitude for leaving the ground state and entering one of these excited states.

Throughout this approximation, I need to assume that $\Delta(s)$ does not become exceedingly small. The integral becomes

$$C_1(1,\tau) = \int_0^1 ds \frac{\frac{d}{ds} e^{-i\tau \int_s^1 dz\, \Delta(z)}}{i\tau \Delta(s)^2/\gamma(s)} \tag{8.18}$$

Integration by parts yields

$$\begin{aligned} C_1(1,\tau) = &\left[\frac{e^{-i\tau \int_s^1 dz\, \Delta(z)}}{i\tau \Delta(s)^2/\gamma(s)} \right]_{s=0}^{1} \\ &- \int_0^1 ds\, e^{-i\tau \int_s^1 dz\, \Delta(z)} \frac{d}{ds} \frac{1}{i\tau \Delta(s)^2/\gamma(s)}. \end{aligned} \tag{8.19}$$

By the properties of oscillatory integrals, the last integral here is $\mathcal{O}(\tau^{-2})$, so that

$$C_1(1,\tau) = \frac{1}{i\tau \Delta(1)^2/\gamma(1)} - \frac{e^{-i\tau \int_0^1 ds\, \Delta(s)}}{i\tau \Delta(0)^2/\gamma(0)} + \mathcal{O}(\tau^{-2}). \tag{8.20}$$

For convenience, I define $\rho(s) = \gamma(s)/\Delta(s)^2$, and this value $\rho(s)$ is related to the naive adiabatic condition that $\tau \gg |\rho(s)|$. Therefore, this $\rho(s)$ can be thought of as the gauge to determine whether the adiabatic limit is satisfied.

Then the probability of transitioning into one of the m-fold degenerate first excited states is given by

$$\begin{aligned} \frac{P(\tau)}{m} = & \frac{\rho(1)^2 + \rho(0)^2}{\tau^2} \\ & - \frac{2\rho(0)\rho(1)}{\tau^2} \cos(\omega\tau) + \mathcal{O}(\tau^{-3}), \end{aligned} \tag{8.21}$$

where

$$\omega \equiv \int_0^1 ds\, \Delta(s). \tag{8.22}$$

Therefore, the final probability of failure (and success) are oscillating functions with a frequency dependent on the integral of the spectral gap. This result is a previously known [30, 59, 60]. Additionally, Wiebe and Babcock [60] have proposed using this oscillating behavior to enhance quantum adiabatic computing.

Notice that the rule of thumb for the adiabatic condition states that the evolution time, τ, needs to grow with $\rho(s)$, and indeed this formula captures this since the probability of failure depends on $\frac{\rho(s)}{\tau}$. The new and interesting behavior here is not the overall τ dependence but the oscillating dependence.

8.2.2 Numerical Confirmation

To numerically test the predictions of the previous subsection, I return to the quantum computational problem of Hamming weight barrier tunneling introduced in Ch. 3.

One important and instructive way to get a large gap out of the barrier tunneling problem is to set the barrier to zero, $b(|z|) = 0$. This decouples all the qubits from each other, making the problem effectively a two-level system, independent of the number of qubits. As well, the spectral gap has a simple closed form expression for this case

$$\Delta_{NB}(s) = \sqrt{1 - 2s + (1 + \mu)s^2} \tag{8.23}$$

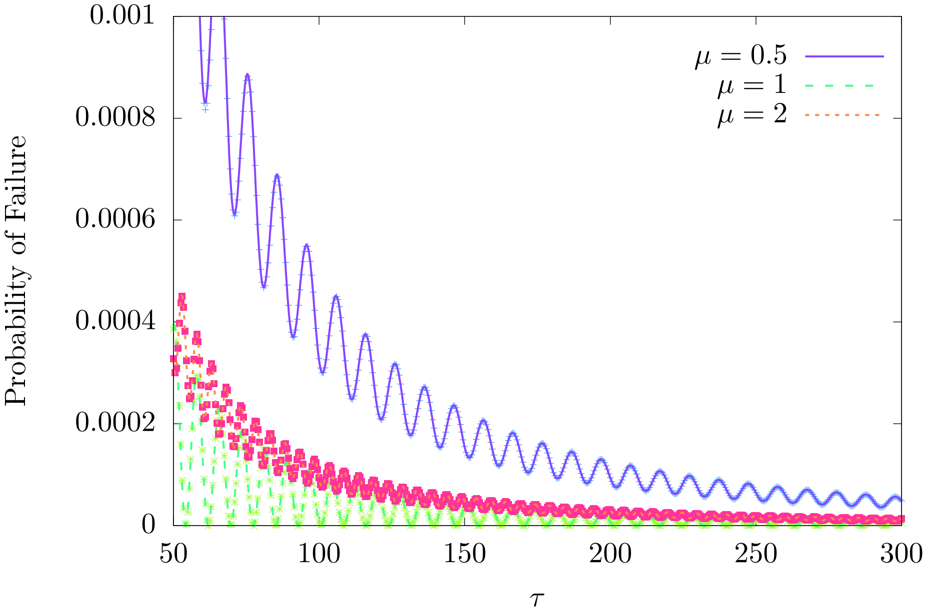

Figure 8.1: The probability of transition to the first excited state versus the evolution time τ for no barrier. The solid lines represent the theoretical predictions coming from Eq. 8.21, and the circles represent data obtained from direct integration of the Schrödinger equation. Data is shown for various μ and $n = 1$; though this problem has decoupled qubits, so this is representative of arbitrary n. Notice especially the oscillatory behavior that depends on the integral of the spectral gap over the entire evolution.

This gap is also useful even in the cases with a barrier because it can well approximate the gap far away from where the tunneling event occurs. Also note that in the no barrier case $\gamma(s) = \frac{\mu}{2\Delta(s)}$. In later analyses of cases with barriers, I use this expression for $\gamma(s)$ again when evaluating $\gamma(0)$ and $\gamma(1)$ which are far from the tunneling event.

When there is no barrier, I can analytically integrate Eq. 8.23 and obtain a form for the probability of failure, Eq. 8.21. In Fig. 8.1, I compare this analytic expression for the near-adiabatic probability of transitioning to an excited state with the exact result obtained by numerical integration of the Schrödinger equation. The analytic approximations predicts the actual data extremely well. It should be noted that without a barrier, the qubits are decoupled, so this is a two-level system, meaning many of my approximations are exact.

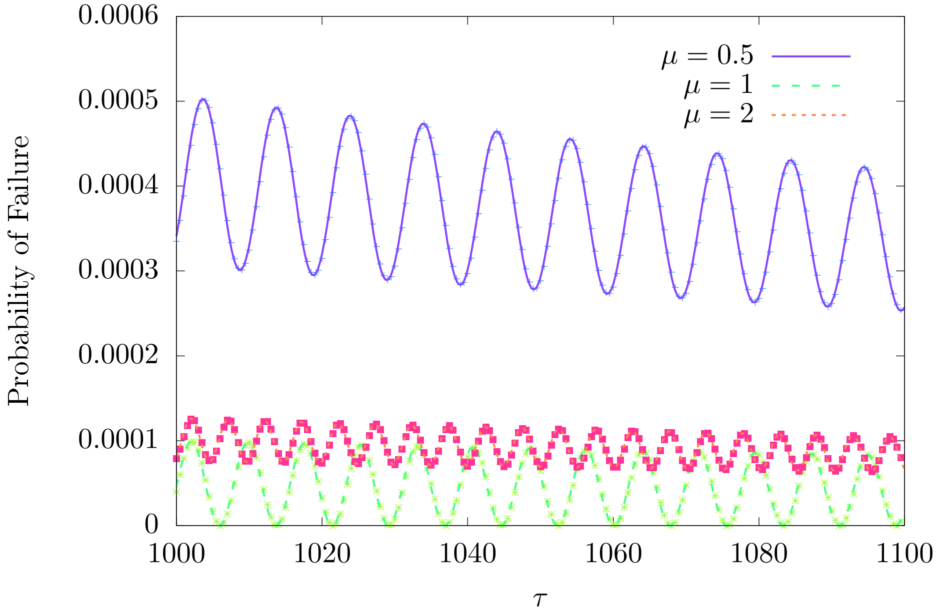

Figure 8.2: The probability of transition to the first excited state versus the evolution time τ for a barrier that grows with $\alpha = \beta = 1/10$. The solid lines represent the theoretical predictions coming from Eq. 8.21, and the circles represent data obtained from direct integration of the Schrödinger equation. Data is shown for various μ and $n = 100$. Even in this case with more approximations than in Fig. 8.1, the oscillatory analytic prediction is still quite accurate.

In Fig. 8.2, I look at a case, where the barrier is present and the qubits are not decoupled. In this case, I take a barrier with scaling exponents $\alpha = \beta = 1/10$. Based on previous work [42], this barrier should be easy to tunnel through, with only a constant gap as n increases. In this instance, I have taken $n = 100$ qubits and see that good agreement between the direct data and the analytic approximation from Eq. 8.21 that was calculated using numerical integration of the spectral gap and approximating $\rho(1)$ and $\rho(0)$ by the unperturbed value since they are evaluated far from where the barrier is relevant.

8.3 Small Gap

For small gaps, I examine the case where the gap remains large everywhere except in a region right around a critical s^*. At this critical s^*, the system has an avoided level crossing, which traditionally is handled by a formalism such as the Landau-Zener problem. In this section, I focus on a system with a single avoided level-crossing, but these methods can easily be generalized to systems with multiple level crossings.

Note that this section discusses the near-adiabatic limit, so big-O notation is not appropriate here. In the $\tau \to \infty$ limit, the results of the large gap section are accurate. This section focuses on the behavior of the success probability in regions where the inverse gap is small relative to the evolution-time, τ. Thus, I examine an intermediate region, and all of the results for modifications to $C_1(1, \tau)$ tend to zero faster than Eq. 8.20 in the asymptotic limit of τ.

8.3.1 Frequency Splitting

An integral of the form Eq. 8.17 can often be treated with the stationary phase approximation. If the phase function, $\int_s^1 dz\, \Delta(z)$, is ever stationary as would occur when $\Delta(s) = 0$, then the stationary phase approximation says that asymptotically, the integral is dominated by the value close to that stationary point.

Unfortunately, since the gap never goes to zero, $\Delta(s) \neq 0$, the system never has a point of true stationary phase. However, there is an avoided level crossing where the gap becomes very small in the vicinity of s^*. Additionally, near s^*, the denominator of the integral is also small since it is proportional to $\Delta(s)$ ($\gamma(s)$ also often depends inversely on the gap). Therefore, the region around s^* should still contribute more to the integral than other regions, but since this is not a true stationary phase point, the contribution near s^* is drowned out in the asymptotic limit of τ.

Therefore, I make an ansatz that the region around s^* also contributes significantly to the final probability amplitude in the near-adiabatic limit. The contribution to the probability amplitude in the vicinity of s^* is roughly of the form

$$\lambda \equiv e^{-i\tau \int_{s^*}^{1} dz\, \Delta(z)} \int_{s^*-\epsilon}^{s^*+\epsilon} ds \frac{e^{-i\tau \int_{s}^{s^*} dz\, \Delta(z)}}{\Delta(s)/\gamma(s)}, \tag{8.24}$$

where I have pulled out the contribution to the phase due to getting to the critical point. I can generalize this further (potentially allowing me to relax some of the assumptions that led to Eq. 8.17) to

$$\lambda = e^{-i\tau \int_{s^*}^{1} dz\, \Delta(z)} \Lambda(\tau). \tag{8.25}$$

Later in this section I present numeric evidence supporting this ansatz using example problems from Ch. 3. Furthermore, my numerics indicate that $\Lambda(\tau)$ is real, allowing me to ignore any potential extra phases.

Then, the conjectured probability amplitude in the near-adiabatic limit is

$$C_1(1,\tau) \approx \frac{\rho(1)}{i\tau} - \frac{\rho(0)e^{-i\tau(\omega_+ + \omega_-)}}{i\tau} + \Lambda(\tau)e^{-i\tau\omega_+}, \tag{8.26}$$

$$\omega_+ = \int_{s^*}^{1} dz\, \Delta(z), \qquad \omega_- = \int_{0}^{s^*} dz\, \Delta(z).$$

This probability amplitude leads to a probability of transition of

$$\begin{aligned}\frac{P(\tau)}{m} \approx{}& \Lambda(\tau)^2 + \frac{\rho(0)^2 + \rho(1)^2}{\tau^2} \\ &+ \frac{2\Lambda(\tau)}{\tau}\left(\rho(0)\sin(\omega_-\tau) + \rho(1)\sin(\omega_+\tau)\right) \\ &- \frac{2\rho(0)\rho(1)}{\tau^2}\cos((\omega_+ + \omega_-)\tau)\end{aligned} \tag{8.27}$$

where m is the degeneracy of the first excited state.

Most importantly the final probability amplitude now has sinusoidal motion dependent on two frequencies, $\omega_{\pm}$. Thus, the final probability no longer has a simple sinusoidal behavior but depends on the superposition of multiple sinusoids.

This splitting of the frequency, creating a superposition of sinusoids when the gap is small, is a well realized feature in actual problems, as seen numerically in Section 8.3.3. In fact, this frequency splitting seems to persist even when most of the simplifying assumptions that went into Eqs. 8.21 & 8.27 fail. In the next section, I take a Landau-Zener approach to the small gap, but in other small gap models I examined, this frequency splitting persisted.

8.3.2 Analytic Approximation

In this section, I approximate the avoided level crossing as a Landau-Zener transition. In the language of the previous subsection, I use the ansatz that $\Lambda(\tau)$ is related to the Landau-Zener transition probability. The Landau-Zener problem works with Hamiltonians of the form

$$\hat{H}_{LZ}(s) = \frac{v}{2}(s - s^*)\hat{\sigma}_z + \frac{g}{2}\hat{\sigma}_x, \tag{8.28}$$

where $g \equiv \Delta(s^*)$ is the minimum spectral gap, and v is slope of the spectral gap far from s^*. The Landau-Zener formula says that the probability of transitioning from the ground state to the excited state going from $t = -\infty$ to $t = \infty$ is

$$P_{LZ} = e^{-2\pi \frac{g^2}{4v}\tau}, \tag{8.29}$$

where the τ is coming in because the Landau-Zener transition is formulated in actual time, t, whereas Eq. 8.28 is formulated in s.

I take the probability amplitude of transition through the avoided level crossing to be proportional to $\sqrt{P_{LZ}}$. I also include a real parameter, A, to account for non-idealnesses in the Landau-Zener transition such as the finite nature of the transition and the fact that I do not start the avoided level-crossing in exactly the ground state. Therefore, I take

$$\Lambda(\tau) = Ae^{-\pi\frac{g^2}{4v}\tau}. \tag{8.30}$$

Then, my ansatz for the final failure probability when there is an avoided-level crossing is

$$\begin{aligned}\frac{P(\tau)}{m} \approx\; & A^2 e^{-2\pi\frac{g^2}{4v}\tau} + \frac{\rho(0)^2+\rho(1)^2}{\tau^2} \\ & + \frac{2Ae^{-\pi\frac{g^2}{4v}\tau}}{\tau}\left(\rho(0)\sin(\omega_-\tau)+\rho(1)\sin(\omega_+\tau)\right) \\ & - \frac{2\rho(0)\rho(1)}{\tau^2}\cos((\omega_+ + \omega_-)\tau).\end{aligned} \tag{8.31}$$

In an actual setting, g, v, and $\omega_\pm$ can be determined from the shape of the spectral gap of the problem in question. I leave A as a fitted parameter that accounts for non-idealness in the system. In the next subsection, I determine all these parameters in specific computational settings.

8.3.3 Numerical Confirmation

In this subsection, I present numeric data confirming the usefulness of the analytics in the rest of the section. The approximations rely on two key assumptions, namely that $|C_0(s)|$ is close to one, meaning the system primarily stays in the grounds state and that τ is large. I will test these approximations as well as the Landau-Zener anstaz modification by comparing Eq. 8.31 to direct data.

Each numeric simulation is based on direct Schrödinger evolution of the wavefunction

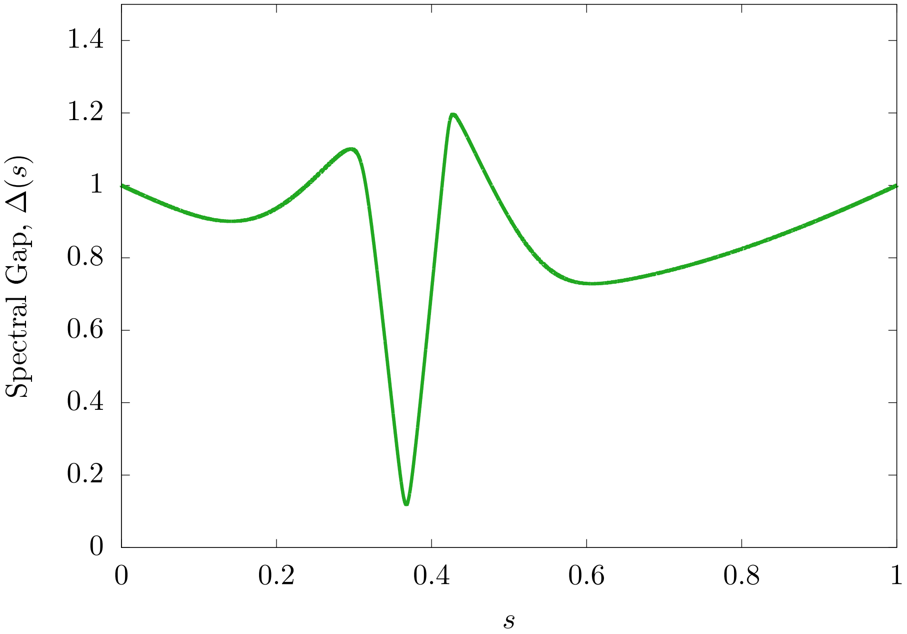

Figure 8.3: The spectral gap, $\Delta(s)$, versus s, for a binomial barrier with $\alpha = 0.3$, $\beta = 0.5$, $\mu = 1$, and $n = 84$. The minimum spectral gap, g is obtained from the minimum here, and the Landau-Zener slope, v, is approximated by the almost linear sections near the minimum gap. As well the frequencies, $\omega_\pm$, come from numerical integration of this curve.

from $s = 0$ to $s = 1$. This evolution then gives data of $P(\tau)$ (actually I calculate $1-P(\tau)$, the probability of staying in the ground state) versus τ. This data is then fit using a function of the form of Eq. 8.31 to obtain a fitted value of A. Across numerous trials with different barrier shapes and sizes, I determine that A can be taken as real.

For each of the simulations, I numerically calculate the gap as a function of s and use it to extract g, v, and $\omega_\pm$. For v, I base its value on the slope of $\Delta(s)$ close to the avoided level crossing where the gap is increasing or decreasing effectively linearly. The frequencies $\omega_\pm$ are obtained through numerical integration of the spectral gap before and after the critical s^*.

I have done several simulations for the Hamming weight barrier problem with a binomial barrier. For a representative plot, see Fig. 8.4 which shows the close correspondence

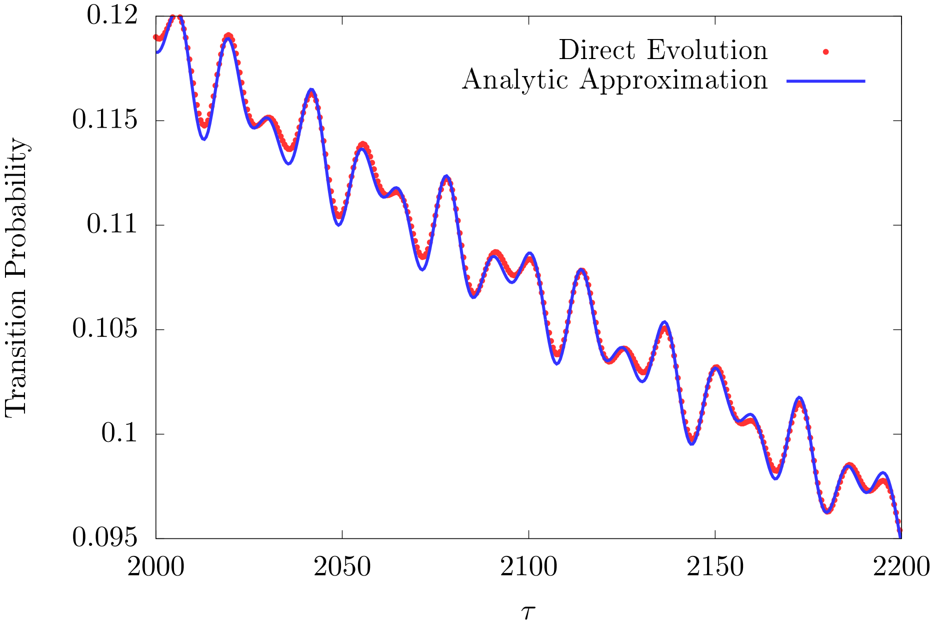

Figure 8.4: The probability of transitioning out of the ground state as a function of the total evolution or run time, τ, for a binomial barrier with $\alpha = 0.3$, $\beta = 0.5$, $\mu = 1$, and $n = 84$. The red dots represent data obtained through direct evolution of the Schrödinger equation, and the blue curve is the result of applying Eq. 8.31 to the problem. To obtain the blue curve, the parameter A was fitted to $A = 0.107$, but all other parameters were calculated from the spectral gap directly. Notice that the analytic expression matches up quite well with the data, mimicing the frequency behavior. Also notice that this analytic expression seems to hold even when the probability of failure is relatively high, around 10%.

between the analytic expression in Eq. 8.31 and the direct data, for a binomial barrier with height and width scaling like $n^{0.5}$ and $n^{0.3}$ respectively and with $\mu = 1$ and $n = 84$. Notably, the frequency splitting behavior is evident here in the superposition of two sinusoids, and the overall scaling matches well with the Landau-Zener probability of transition exponential.

The fitted parameter A in Fig. 8.4 is $A = 0.107$. In a pure Landau-Zener transition, the value of A is 1; the fact that the value is less than 1, is likely an indicator of the finite scale of the Landau-Zener region. A true Landau-Zener transition occurs from

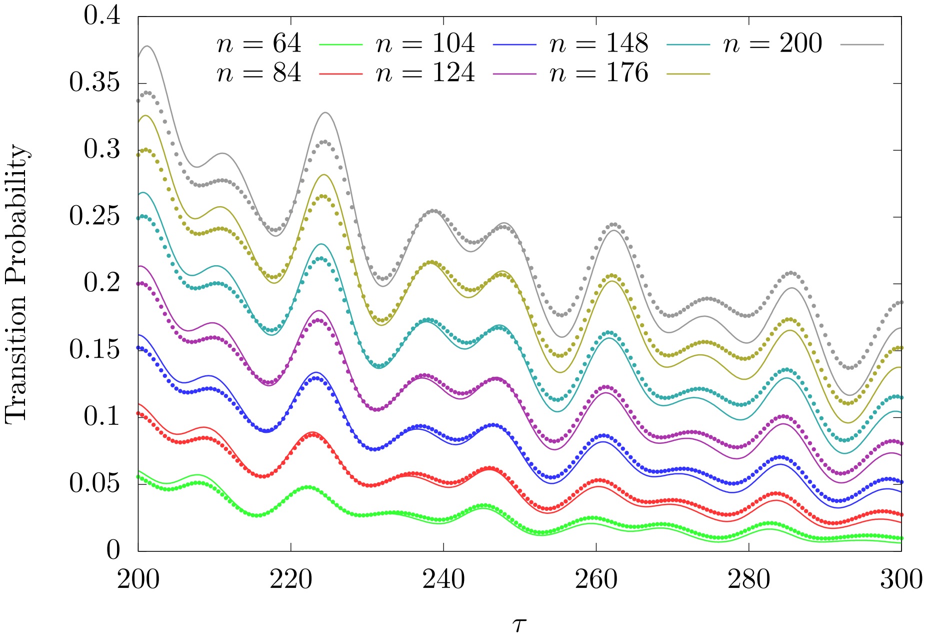

Figure 8.5: The probability of transitioning out of the ground state versus the evolution time, τ, compared to the theoretical predictions from Eq. 8.31. The dots are the direct Schrödinger evolution data, and the lines are the thoeretical predictions. This data all comes from the Hamming weight barrier problem with a binomial barrier with ehight and width both growing with $n^{0.3}$ and $\mu = 1$. The frequency behavior seems to be captured by the theoretic predictions well, but the overall amplitude is less well-predicted especially for higher n where the probability of trnasitioning is higher.

$t = -\infty$ to $t = \infty$, so I have a finite range, which might influence the value of A. Also the A parameter is probably absorbing discrepancies caused by the other assumptions in deriving Eq. 8.31, including those independent of the Landau-Zener-like transition.

I performed similar trials for other barrier sizes and values of n and μ, and for each trial, I find a very good correspondence between my anstaz and the direct Schrödinger data. One of the largest assumptions I use in deriving Eq. 8.31 is that the majority of the probability remains in the ground state so that I can approximate $C_0(s) \approx 1$ in the differential equations. My numerics show that this is somewhat of a loose condition with, for instance, Fig. 8.4 showing good correspondance even when $|C_0(s)|^2 \approx 0.9$.

Additionally, my approximations also rely on the fact that τ is large. In Fig. 8.5 I display data compared to predictions for a variety of n values at much lower τ. The agreement is not as clear as with larger τ, and especially at larger n (thus larger transition probability), the agreement is noticeably degraded. However, there is still correspondence, and especially the frequencies, if not the amplitudes line up well.

In Fig. 8.5, all this data is taken for a barrier with height and width scaling with $n^{0.3}$. The spectral gap $\Delta(s)$ is very similar for all these n values. As n increases, the gap around the avoided level-crossing changes, lowering g and raising v, but the majority of the spectral gap remains the same, meaning that $\omega_\pm$ are virtually the same for different n. The persistence of the same $\omega_\pm$ leads to similar frequency behavior across n, even though the enveloping probability scaling with τ changes with g, in accordance with the standard adiabatic theorem.

For the cubic potential, the spectral gap still goes through an avoided level crossing as shown in Fig. 8.6, but in this case, the avoided level crossing is much more assymetric for $n = 30$. The slopes on either side of the avoided level crossing are different, making it difficult to determine v. Therefore, in this case, I leave v as a fitted parameter. This essentially eliminates the benefit of the Landau-Zener ansatz, but the frequency splitting of Eq. 8.27 is still valid. I also calculate $\rho(0)$ and $\rho(1)$ numerically using equivalent diagonalization methods to how I calculate the spectral gap itself.

Using Eq. 8.31 with fitted A and v, I predict how the transition probability changes as a function of τ in Fig. 8.7. The correspondance between my prediction and the direct evolution data is quite good, but much of this is due to two fitted parameters. The important part of this figure is that the oscillatory behavior matches quite well between the predicition and direct data.

Also, Fig. 8.7 once again matches well even at the relatively high transition probability of around 0.24, indicating that the frequency splitting behavior is fairly robust to my

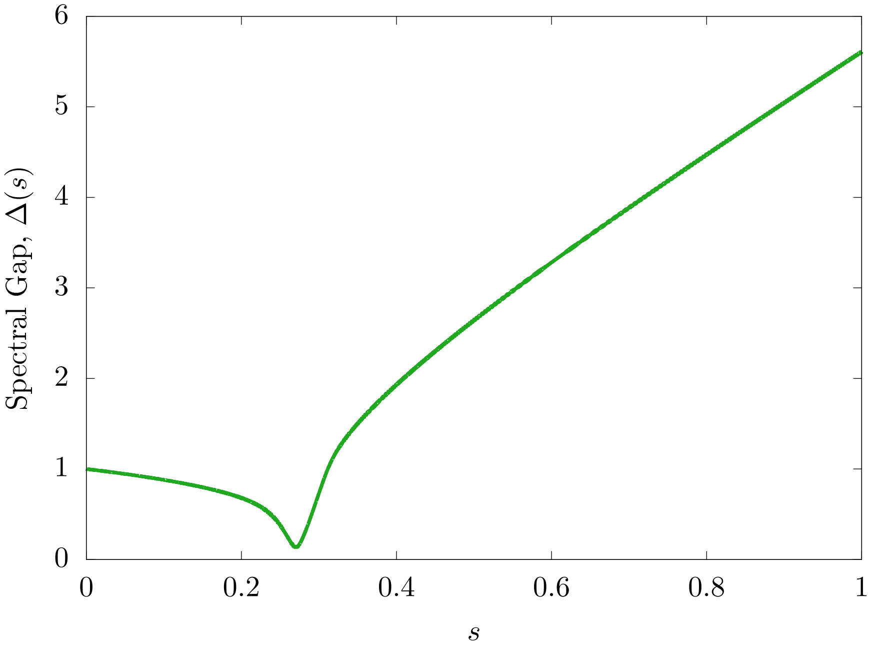

Figure 8.6: The spectral gap, $\Delta(s)$, versus s, for the cubic potential and $n = 30$. The frequencies, $\omega_{\pm}$, come from numerical integration of this curve. Notice that the avoided level crossing is assymmetric for $n = 30$, leaving me unable to properly use the Landau-Zener transition probability to approximate this crossing.

exact approximations. The correspondance in this cubic potential problem as well as the barrier problem indicate the correctness of the frequency splitting ansatz.

8.4 Adiabatic Grover Search

The Grover search algorithm [3, 63] is a digital quantum algorithm for searching an unstructured set of N elements for one of M target elements, where $M \ll N$. Classically, $\mathcal{O}(N/M)$ queries must be made to find one of the target states, but Grover's algorithm requires only $\mathcal{O}(\sqrt{N/M})$ oracular calls to find a target with high probability [63]. Notably, the likelihood of success of this algorithm is periodic in the number of queries with period $\propto \sqrt{N/M}$. Therefore, timing needs to be exact to achieve success; though, there

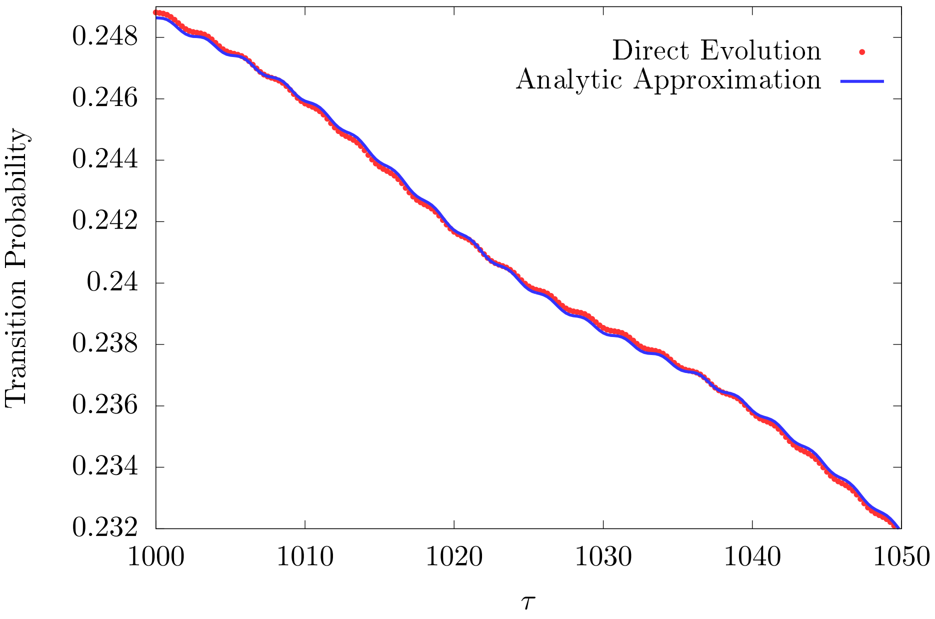

Figure 8.7: The probability of transitioning out of the ground state as a function of the total evolution or run time, τ, for the cubic potential and $n = 30$. The red dots represent data obtained through direct evolution of the Schrödinger equation, and the blue curve is the result of applying Eq. 8.31 to the problem. Due to the assymetric gap during the avoided level crossing as seen in Fig. 8.6, the fit needed more than just A as a fitted parameter. I also included v as a fitted parameter to obtain this agreement. Notice that the oscillatory pattern is still correct, confirming the robustness of the frequency splitting behavior from section 8.3.1.

are methods of alleviating this periodicity in favor of a larger constant scaling factor.

An adiabatic version of Grover's search algorithm has been developed [64, 65] that mimics the $\mathcal{O}(\sqrt{N/M})$ scaling. However, previous studies of this algorithm have relied on studying the adiabatic theorem's asymptotic scaling, and to my knowledge, no groups have looked into whether the periodic nature of digital quantum search carries over to adiabatic Grover in some way. In this section, I demonstrate that the Grover oscillations exist in adiabatic search and are a direct result of the large gap oscillations described in section 8.2.

8.4.1 Adiabatic Grover Background

The adiabatic grover algorithm is setup on a Hilbert space with dimension N (not to be confused with n qubits discussed in previous sections). The initial Hamiltonian is just a simple connection Hamiltonian between all basis states:

$$\hat{H}_0 = -\frac{1}{N}\sum_{i,j=1}^{N} |i\rangle \langle j| \,. \tag{8.32}$$

The ground state of this Hamiltonian is just the uniform superposition of all states. For ease of analysis, typically the first M basis states are chosen to be the target states, and the final Hamiltonian gives them a preferential energy term such that

$$\hat{H}_1 = \hat{\mathbb{I}} - \sum_{m=1}^{M} |m\rangle \langle m| \,. \tag{8.33}$$

Note that the ground state of the final Hamiltonian is M-fold degenerate; whereas, the ground state of the initial Hamiltonian is non-degenerate. Seemingly, this means that the spectral gap would go to zero at some point during the evolution; however, this is not an issue due to the nature of the degeneracy. The true ground state throughout the evolution is symmetric between target states so that in the end it is a uniform superposition of all targets. The degeneracy in the final ground state is due to states that are non-symmetric in target states transitioning between the first excited state and the ground state. Since the Grover Hamiltonian and initial ground state are symmetric between target states, and since I am concerned with coherent evolution, these non-symmetric states are inaccessible, meaning that I can largely ignore them.

A linear interpolation between these two Hamiltonians does not lead to the square root speedup of Grover's algorithm [64]. Instead, a more generalized interpolation needs to be considered

$$\hat{H}(s) = (1 - g(s))\hat{H}_0 + g(s)\hat{H}_1. \tag{8.34}$$

The spectral gap for this Hamiltonian is given by

$$\Delta(s) = \sqrt{1 - 4(1 - g(s))g(s)\frac{(N - M)}{N}}. \tag{8.35}$$

The optimal annealing schedule [64], utilizes the adiabatic condition, slowing down when the spectral gap is small and speeding up when it is large. By optimizing for the amount of time spent in the small region, the ideal annealing schedule uses

$$g(s) = \frac{1}{2}\left(1 - \frac{\tan\left((1 - 2s)\tan^{-1}\left(\sqrt{\frac{N-M}{M}}\right)\right)}{\sqrt{\frac{N-M}{M}}}\right) \tag{8.36}$$

This optimal annealing schedule leads to an adiabatic runtime that is $\mathcal{O}\left(\sqrt{N/M}\right)$.

8.4.2 Large Gap Oscillations

My contribution is to show that the large gap oscillations described in section 8.2 lead to the same periodic behavior in the adiabatic algorithm as in the standard digital Grover algorithm. Note that the gap in the Grover problem does become small at $s = 1/2$, where $\Delta(1/2) = \sqrt{M/N}$, but in this section I consider only the large gap oscillations. Also, Wiebe and Babcock [60], have examined the large gap oscillations of adiabatic Grover in the case of $M = 1$. There, they show that timing the adiabatic algorithm based on the large gap oscillations improves the performance of the algorithm; however, they never explicitly state the oscillation frequency or connect it back to digital Grover. Furthermore, my results make the simple extension to $M > 1$.

Looking at Eq. 8.31, it is easy to see that the large gap oscillations will dominate if

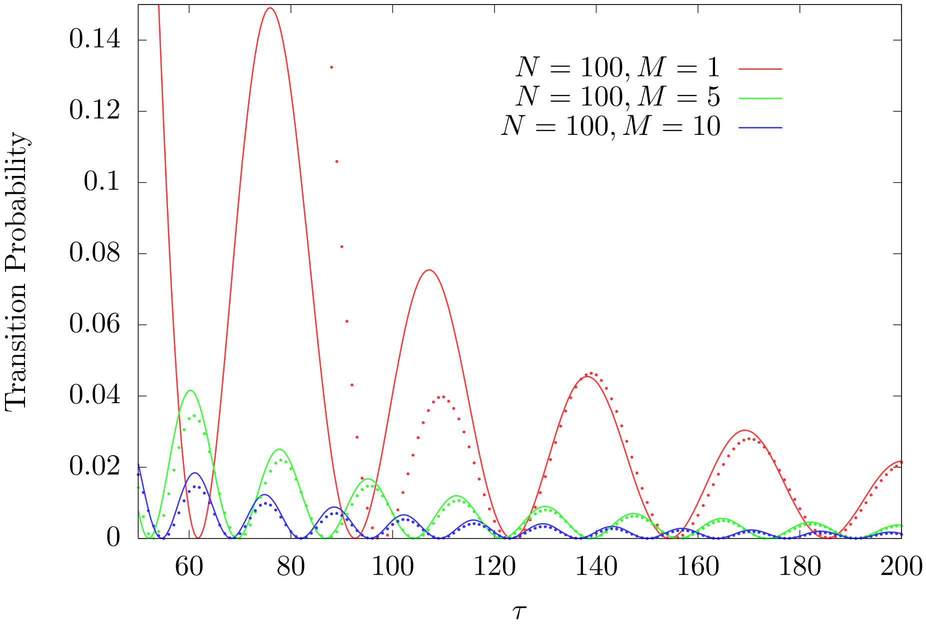

Figure 8.8: The probability of transitioning out of the ground state as a function of the total evolution or run time, τ, for the adiabatic Grover search problem. Notice that there is good agreement between the analytics and the data when the probability of transitioning is low but that the agreement is worse for shorter evolution times. The analytics here are exact and included no fitting parameters.

the standard adiabatic condition is met. Therefore, in this section the system is evolving adiabatically so that $\tau \in \mathcal{O}\left(\sqrt{N/M}\right)$, and the main goal is to determine what the fine structure of the oscillations in this limit are.

Using the spectral gap and the annealing schedule, Eqs. 8.35 & 8.36, I can calculate the frequency of the large gap oscillations as described by Eq. 8.22:

$$\omega = \sqrt{\frac{M}{N}} \frac{\tanh^{-1}\sqrt{\frac{N-M}{N}}}{\tan^{-1}\sqrt{\frac{N-M}{M}}} \tag{8.37}$$

The period of oscillations $T \equiv 1/\omega$ in the limit of large N/M is given by

$$T \to \frac{\pi}{2\ln 2 + \ln\frac{N}{M}}\sqrt{\frac{N}{M}}. \tag{8.38}$$

Therefore, up to logarithmic factors, the period of oscillations for adiabatic Grover's search is $\widetilde{\mathcal{O}}(\sqrt{\frac{N}{M}})$ which matches the digital equivalent.

Additionally, we can look at the amplitudes of the oscillations in the probability of transitioning, Eq. 8.21. In order to get $\rho(s)$, we need $\Delta(s)$ which was listed in the previous subsection and $\gamma(s)$. Assuming the annealing schedule in Eq. 8.35

$$\begin{aligned}\gamma(s) = {} & 4M\tan^{-1}\left(\sqrt{\frac{N-M}{M}}\right) \\ & \times \frac{\left(\sqrt{\frac{N-M}{M}} - \tan\left((1-2s)\tan^{-1}\left(\sqrt{\frac{N-M}{M}}\right)\right)\right)}{N-2M+N\cos\left(2(1-2s)\tan^{-1}\left(\sqrt{\frac{N-M}{M}}\right)\right)}.\end{aligned} \tag{8.39}$$

The truly important part of this equation is that this $\gamma(s)$ is symmetric about $s = 1/2$ so that notably $\gamma(0) = \gamma(1)$. Since the spectral gap is $\Delta(0) = \Delta(1) = 1$ at the end points as well, this leaves $\rho \equiv \rho(0) = \rho(1)$. Therefore, the probability of transitioning out of the ground state reduces to

$$P(\tau) = \frac{4\rho^2}{\tau^2}\sin^2\left(\frac{\omega\tau}{2}\right) + \mathcal{O}(\tau^{-3}). \tag{8.40}$$

Notably, this probability goes to $\mathcal{O}(\tau^{-3})$ periodically according to ω. Therefore, adiabatic Grover's search can be timed to get perfect success probability in the adiabatic limit. In Fig. 8.8, I show the agreement between the analytic predictions of Eq. 8.40 and direct evolution data. There is good agreement between the analytics and the data, especially for lower transition probabilities.

Chapter 9

Conclusion

Quantum adiabatic computing provides a novel framework for quantum computing that requires very different structure and programming expertise to the digital quantum gate model. By relying solely on Hamiltonians and non-commuting time-evolution, quantum annealing can utilize all the prior theory of quantum dynamics to create and analyze quantum algorithms with an eye toward previous intuition for inspirations and development.

Quantum annealing with stoquastic Hamiltonians is especially exciting because stoquastic Hamiltonians provide a relatively straightforward analytic and experimental testing bed for quantum ideas. Between the D-Wave system [14] and more recent quantum annealing initiatives, such as IARPA's Quantum Enhanced Optimization project, and implementation on near-term digital quantum circuits, it is highly likely that quantum annealing and adiabatic optimization will be (or already is) one of the earliest applications of near-term devices.

Therefore, the question of what quantum advantage can be achieved with just stoquastic Hamiltonians is important for determining how complicated quantum annealing algorithms need to be in order to achieve quantum advantage. The first part of this dissertation focused on developing the power of quantum adiabatic computing to tunnel through difficult barriers. I examined what sizes of barriers are easy or hard for quantum adiabatic computing to tunnel through, and I performed a numerical comparison of quantum adiabatic runtimes and Quantum Monte Carlo runtimes.

Evidence is mounting that stoquastic quantum adiabatic computing is not exponentially more powerful than classical algorithms, at least in Hamming-symmetric models. There is some evidence that topological obstructions in models with higher effective dimensions [39], but there has been little evidence of stoquastic quantum advantage beyond these contrived models which at the moment only have proven advantage over path-integral quantum Monte Carlo.

Therefore, a large portion of research has moved into how to enhance the power of quantum annealing, especially with stoquastic Hamiltonians. The avenue that I considered in this dissertation was using non-adiabatic processes to speed up the algorithms. Specifically, I studied how and why a fine balance between the initial and final Hamiltonians leads to a speedup that probably is not general or useful in real algorithmic settings.

I ended with a description of how to extend the adiabatic theorem to describing near-adiabatic behavior. This work has vast avenues by which it could be expanded. Firstly, my results are still reliant on ansatzes which could potentially have theoretic explanations; moreover, the dependence of the transition probability on the exact form of the avoided-level crossing can be explored further. This work could be useful for precise timing of the adiabatic theorem in a similar way to what Wiebe and Babcock describe [60].

As quantum devices become more and more powerful with more qubits and longer coherence times, it is important to study what algorithms are possible and useful on these devices. Stoquastic Hamiltonians and quantum annealing lend themselves well to solving classical optimization problems. Classical optimization problems are at the heart of many outstanding scientific and industrial problems, and being able to improve our methods of optimization can greatly effect society.

Approaches to improving quantum adiabatic optimization are being developed, from my own non-adiabatic approaches to changes in the annealing schedule and processes. The main concern is discerning how much our found speedups rely on specifics about the problem rather than general properties about quantum annealing itself. A speedup to one specific problem instance is not useful, but a speedup, even a polynomial speedup to a class of optimization problems could have profound consequences.

Appendix A

Villain Transformation

The goal of this appendix is to fill in the gaps in the derivation of Eq. 5.6, starting from Eq. 5.4 in Ch. 5. The Villain representation [26] is a standard semi-classical approximation which first takes the continuum limit of the eigenvalues of $\hat{\jmath}_z$ and then defines a conjugate momentum to this continuous "position" variable. This technique, as it is commonly implemented, has many subtleties that are ignored, so this appendix formalizes the assumptions implicit in the Villain representation. Furthermore, I extend these results and show that certain assumptions about the ground and first excited states allows a refining of the approximations and improving of the analysis for the barrier tunneling problem. I also derive all of the results in the discrete case, only resorting to the continuum limit at the end, elucidating the nature of the assumptions.

The work presented in this appendix is closely connected with the material in Ch. 5. Both this appendix and Ch. 5 originally appeared in joint work I did with Wim van Dam in [19].

A.1 Original Villain Transformation

The original Villain transformation [26] says that $\hat{\jmath}_z$ and $\hat{\jmath}_\pm \equiv \hat{\jmath}_x \pm \mathrm{i}\hat{\jmath}_y$ act on the $\hat{\jmath}_z$ eigenkets with the following relations

$$\begin{aligned}
\hat{\jmath}_z \left|q\right\rangle &= q \left|q\right\rangle \\
\hat{\jmath}_+ \left|q\right\rangle &= \mathrm{e}^{-\mathrm{i}p}\sqrt{1+\epsilon-q(q+\epsilon)} \left|q\right\rangle \\
\hat{\jmath}_- \left|q\right\rangle &= \sqrt{1+\epsilon-q(q+\epsilon)}\mathrm{e}^{\mathrm{i}p} \left|q\right\rangle
\end{aligned} \tag{A.1}$$

Here p is the conjugate momentum for q, and in q-space it can be represented as $p = -\mathrm{i}\epsilon\frac{\partial}{\partial q}$. Many users [54, 55, 56] of the Villain representation employ the small ϵ limit to say that the square root factors in the $\hat{\jmath}_\pm$ expressions are approximately equal and that the commutators between q and p are negligible. Using these approximations, they find that

$$\hat{\jmath}_x = \frac{1}{2}\left(\hat{\jmath}_+ + \hat{\jmath}_-\right) = \sqrt{1+q^2}\cos p. \tag{A.2}$$

These small ϵ approximations ignore many subtleties, most centering around how big p is. Eq. A.2 is true only to zeroeth order in ϵ, but because the derivative operator has an operator norm that is proportional to ϵ^{-1}, there is a relevant term to second order in p. However, this expression is incorrect at all higher orders of ϵ and in fact includes terms linear in p that are not included here. One of the most misleading errors is that the true expression for $\hat{\jmath}_x$ contains no terms with higher order of p than p^2.

Below, I go through a more formal derivation of $\hat{\jmath}_x$'s expansion using the underlying matrices and proceed with a description of how the discrete Villain transformation can be used in the setting of the barrier tunneling problem.

A.2 Discrete j-operators

I remind the reader that $\epsilon = 2/n = 1/J$, and I start by examining $\hat{\jmath}_x$ in the eigenbasis of $\hat{\jmath}_z$ given by $|q\rangle$ where $q = \epsilon m \in [-1, +1]$ for $m \in \{-J, -J+1, \ldots, J\}$. The goal will be to determine how $\hat{\jmath}_x$ acts on a general state $|\psi\rangle$:

$$\hat{\jmath}_x |\psi\rangle = \sum_q \hat{\jmath}_x \psi_q |q\rangle \tag{A.3}$$

I introduce three new operators which will simplify the representation of the raising and lowering operators $\hat{\jmath}_\pm$:

$$\hat{P} = \sum_{q\in[-1,1-\epsilon]} |q+\epsilon\rangle \langle q| \quad \text{and} \quad \hat{M} = \sum_{q\in[-1+\epsilon,1]} |q-\epsilon\rangle \langle q| \tag{A.4}$$

$$\hat{q} = \sum_q q |q\rangle \langle q| .$$

Since $\hat{q}$ is diagonal, functions of the operator are diagonal themselves, and the first two operators extract just raising and lowering of indices without any prefactors. Therefore, $\hat{\jmath}_x$ is represented as

$$\hat{\jmath}_x = \frac{1}{2}(\hat{\jmath}_+ + \hat{\jmath}_-) = \frac{1}{2}\left(\sqrt{(1-\hat{q})(1+\hat{q}+\epsilon)}\hat{P} + \sqrt{(1+\hat{q})(1-\hat{q}+\epsilon)}\hat{M}\right). \tag{A.5}$$

The eventual goal is to take a continuum limit of q and then represent $\hat{\jmath}_x$ in terms of the continuous q and its derivatives. I can create a dictionary for the matrices that lead to derivatives in the limit that $\langle q|\psi\rangle = \psi_q \to \psi(q)$:

$$\frac{\partial\psi}{\partial q} = \lim_{\epsilon\to 0}\frac{\psi(q+\epsilon)-\psi(q-\epsilon)}{2\epsilon} = \lim_{\epsilon\to 0}\frac{\psi_{q+\epsilon}-\psi_{q-\epsilon}}{2\epsilon} = \lim_{\epsilon\to 0}(\hat{A}\vec{\psi})_q, \tag{A.6}$$

$$\frac{\partial^2\psi}{\partial q^2} = \lim_{\epsilon\to 0}\frac{\psi(q+\epsilon)-2\psi(q)+\psi(q-\epsilon)}{\epsilon^2} = \lim_{\epsilon\to 0}\frac{\psi_{q+\epsilon}-2\psi_q+\psi_{q-\epsilon}}{\epsilon^2} = \lim_{\epsilon\to 0}(\hat{B}\vec{\psi})_q. \tag{A.7}$$

Here I have defined two new operators that correspond to the discrete versions of the first and second derivatives:

$$\epsilon\hat{A} = \frac{\hat{P}-\hat{M}}{2} \quad \text{and} \quad \epsilon^2\hat{B} = \hat{P} - 2\hat{I} + \hat{M}. \tag{A.8}$$

Throughout this appendix, I refer to the relative sizes of operators using their matrix norm. The definition of the matrix norm we are using is the maximum absolute value of any eigenvalue of the operator. Therefore, the norm of $\hat{A}$ is proportional to ϵ^{-1}, and the norm of $\hat{B}$ is proportional to ϵ^{-2}. Section A.3 revisits this concept later in the context of the specific eigenstates and how a problem instance can alter this assessment.

Based on the definitions in Eq. A.8, $\hat{P}$ and $\hat{M}$ are written in terms of $\hat{I}$, $\hat{A}$, and $\hat{B}$ as:

$$\hat{P} = \hat{I} + \epsilon\hat{A} + \frac{\epsilon^2}{2}\hat{B} \quad \text{and} \quad \hat{M} = \hat{I} - \epsilon\hat{A} + \frac{\epsilon^2}{2}\hat{B}. \tag{A.9}$$

Notice that in the continuous form of the Villain representation, Eq. A.1, the operators, $\hat{P}$ and $\hat{M}$, correspond to $e^{\mp ip}$, but in Eq. A.9 the operators only correspond to the first two terms in the Taylor expansion of the exponentials. In addition, the matrix norms of $\hat{A}$ and $\hat{B}$ further complicate the issue, making it deceptively appear that the later terms in the series are smaller when in fact every term in this series is roughly equivalent in size, relative to ϵ.

I use the expansions in Eq. A.9 with Eq. A.5 to obtain a discrete form of the Villain representation. It should be noted that the following expression is exact and includes no approximations yet

$$\hat{j}_x = \frac{1}{2}\left(\sqrt{(1-\hat{q})(1+\hat{q}+\epsilon)}\left[\hat{I} + \epsilon\hat{A} + \frac{\epsilon^2}{2}\hat{B}\right] + \sqrt{(1+\hat{q})(1-\hat{q}+\epsilon)}\left[\hat{I} - \epsilon\hat{A} + \frac{\epsilon^2}{2}\hat{B}\right]\right) \tag{A.10}$$

Next, I begin taking the large spin limit by expanding the square root prefactors in orders of ϵ

$$\sqrt{(1\mp\hat{q})(1\pm\hat{q}+\epsilon)} = \sqrt{1-\hat{q}^2} + \epsilon\frac{1\mp\hat{q}}{2\sqrt{1-\hat{q}^2}} + \mathcal{O}(\epsilon^2) \tag{A.11}$$

Combining Eq. A.10 and Eq. A.11 results in a discrete form of the Villain representation for $\hat{j}_x$:

$$\begin{aligned}\hat{j}_x = \frac{1}{2}\Bigg(&\left[2\sqrt{1-\hat{q}^2} + \frac{\epsilon}{\sqrt{1-\hat{q}^2}}\right] - \epsilon^2\frac{\hat{q}}{\sqrt{1-\hat{q}^2}}\hat{A} \\ &+ \epsilon^2\left(\sqrt{1-\hat{q}^2} + \frac{\epsilon}{2\sqrt{1-\hat{q}^2}}\right)\hat{B} + \mathcal{O}(\epsilon^2)\Bigg).\end{aligned} \tag{A.12}$$

Here I have kept terms that are at most proportional to ϵ, remembering that the norms of $\hat{A}$ and $\hat{B}$ are proportional to ϵ^{-1} and ϵ^{-2} respectively. The expression in Eq. A.12 can be thought of as a more accurate version of Eq. A.2, and up to zeroth order in ϵ, Eq. A.12 and Eq. A.2 agree.

Some knowledge of the specific eigenstates can restrict the form of Eq. A.12 even more, allowing us to consider the maximum eigenvalue of $\hat{A}$ and $\hat{B}$ relevant to the low energy eigenvectors of the problem, rather than the maximum eigenvalues obtainable for a general problem.

There are a few key things to note about Eq. A.12 in relation to Eq. A.2. First, there are only terms up to the second derivative, even including all orders in ϵ. Second, this form of the operator makes no assumptions about the specific form of the Hamiltonian or its energy states. If, as I do in the next section, you make assumptions about the energy states of the barrier tunneling problem, you can further approximate the expression for $\hat{\jmath}_x$ into a simpler form.

A.3 Hamiltonian and Eigenstate-based Approximations

Eq. A.12 can be used with the barrier tunneling Hamiltonian:

$$\epsilon \hat{H} = -(1-s)\hat{\jmath}_x + s\hat{\jmath}_z + sr(\hat{\jmath}_z) + \Theta \tag{A.13}$$

The operator $\hat{\jmath}_z$ is equivalent to $\hat{q}$, and $\hat{\jmath}_x$ can be replaced using Eq. A.12, interpreting $\vec{\psi}$ as an eigenstate with eigenenergy E, so that the Schrödinger equation becomes

$$\begin{aligned} \epsilon E\left|\psi\right\rangle = \Bigg(& s\hat{q} + sr(\hat{q}) + \Theta - (1-s)\left[\sqrt{1-\hat{q}^2} + \frac{\epsilon}{2\sqrt{1-\hat{q}^2}}\right] \\ & +(1-s)\frac{\epsilon^2}{2}\frac{\hat{q}}{\sqrt{1-\hat{q}^2}}\hat{A} \\ & -(1-s)\frac{\epsilon^2}{2}\left(\sqrt{1-\hat{q}^2} + \frac{\epsilon}{2\sqrt{1-\hat{q}^2}}\right)\hat{B} + \mathcal{O}(\epsilon^2)\Bigg)\left|\psi\right\rangle . \end{aligned} \tag{A.14}$$

Eq. 5.5 is a continuum limit version of this equation, where q is treated as a continuous variable so that $\hat{q} \to q$, $\vec{\psi} \to \psi(q)$, $\hat{A} \to \frac{\partial \psi}{\partial q}$, and $\hat{B} \to \frac{\partial^2 \psi}{\partial q^2}$.

Next, I define the operator $\hat{x} = \hat{q} + \frac{1}{2}$ and call the diagonal entries of this operator $x = q + \frac{1}{2}$. The advantage of this variable, x, is that it is small in the vicinity of $q = -\frac{1}{2}$ where the tunneling event occurs. Since the critical tunneling moment and minimum spectral gap occur at $s^* = \frac{1}{2}(\sqrt{3}-1)$, the rest of this appendix will set $s = s^*$.

I look more closely at $\hat{x}$ and its relationship to $|\psi\rangle$. The low energy eigenvectors, $|\psi\rangle$, are essentially zero for most of their entries except right around the location of the primary bump in the distribution. The reasoning behind this comes from the fact that for the low-lying energy states, their energy is lower than the potential energy function for the entire range of x, except in an extremely narrow range around the barrier, leading to exponential suppression of the wave functions outside this region. For the no barrier case, the ground state and first excited state both have width $\mathcal{O}(\sqrt{\epsilon})$ and are centered around $x = 0$ with exponential suppression farther away from $x = 0$.

Since the widths of the ground state and first excited state ($\mathcal{O}(\sqrt{\epsilon})$) are larger than the width of the barrier ($\mathcal{O}(\epsilon^{1-\alpha})$), the range of x over which the components of $|\psi\rangle$ are non-zero is $\mathcal{O}(\sqrt{\epsilon})$. Therefore, focusing on the diagonal terms in the Schrödinger equation that do not include $\hat{A}$ or $\hat{B}$, I expand these to order ϵ by treating $\|\hat{x}\,|\psi\rangle\,\| \in \mathcal{O}(\sqrt{\epsilon})$ since for the non-zero components of $|\psi\rangle$, the typical x values will be of order $\sqrt{\epsilon}$. I also use an arbitrary constant Θ to cancel out the constant terms in the expansion, physically ensuring that the bottom of the potential well is at zero energy:

$$\begin{aligned}&\left(s^*\hat{q} + \Theta - (1-s^*)\left[\sqrt{1-\hat{q}^2} + \frac{\epsilon}{2\sqrt{1-\hat{q}^2}}\right]\right)|\psi\rangle = \\ &\left(\frac{2}{3}(\sqrt{3}-1)\hat{x}^2\right)|\psi\rangle + \mathcal{O}(\epsilon^{3/2}),\end{aligned} \tag{A.15}$$

where $\Theta = -\frac{(\sqrt{3}-1)}{24}(24 + 12\epsilon)$.

Next, I focus on the derivative terms of Eq. A.14 that include $\hat{A}$ and $\hat{B}$. The size of the derivative is governed by the inverse of the length scale over which the eigenvector components change. In the unperturbed case, the eigenvector (which is a binomial distribution) components change on a length scale of $\sqrt{\epsilon}$ which means that $\hat{A}$ scales like $1/\sqrt{\epsilon}$ and $\hat{B}$ scales like $1/\epsilon$. Note that these would then correspond to the norms $\|\hat{A}\,|\psi\rangle\,\|$

and $\|\hat{B}\,|\psi\rangle\,\|$ not $\|\hat{A}\|$ and $\|\hat{B}\|$ which as discussed in the last section can be much larger. This scaling behavior requires the restriction to the low-lying energy states.

In the perturbed case, I expect the shortest length scale in the problem to be governed by the exponential decay inside the barrier. In the prototypical barrier tunneling problem of plane waves tunneling through a square barrier, the Schrödinger equation inside the barrier, which will have extent $-\xi < y < \xi$, will be of the form

$$\frac{d^2\varphi}{dy^2} = \frac{2}{\hbar^2}(V_0 - \mathcal{E})\varphi(y) = k^2\varphi(y). \tag{A.16}$$

The simple Schrödinger equation in Eq. A.16 can be transformed into my problem by taking $\hbar \to \epsilon$, setting the width of the barrier ξ to be proportional to $\epsilon^{1-\alpha}$, setting the height of the barrier, V_0, to be proportional to $\epsilon^{1-\beta}$, and making the energy, $\mathcal{E}$, much smaller than V_0 in the small ϵ limit. If we compare this to the expression in Eq. A.14, at the very least there is still a factor of ϵ^2 in the ratio between the potential barrier $r(q)$ and the second derivative term $\hat{B}$. This is a rough comparison, but it can inform what the exponential decay inside the barrier looks like.

Specifically, Eq. A.16 is solved by $e^{\pm ky}$, where in this case k is proportional to $\epsilon^{-\frac{1}{2}-\frac{\beta}{2}}$, assuming that $V_0 \gg \mathcal{E}$, which is a good assumption in my problem. Therefore, the length scale over which the wavefunction changes inside the barrier is proportional to $\epsilon^{\frac{1}{2}+\frac{\beta}{2}}$. This means that the derivative, and therefore $\hat{A}$, scales like $\epsilon^{-\frac{1}{2}-\frac{\beta}{2}}$, and similarly $\hat{B}$ scales like $\epsilon^{-1-\beta}$. Alternatively, the derivatives are proportional to k and k^2 for $\hat{A}$ and $\hat{B}$ respectively.

One other thing to note is that if you look at a derivative that scales like $\epsilon^{\frac{1}{2}+\frac{\beta}{2}}$, it will only have this extreme scaling in the region close to the edge of the barrier, which means that $x \in \mathcal{O}(\epsilon^{1-\alpha})$ when you care about derivatives that are this large. Thus in keeping track of the order, the higher order terms in x are even more exacerbated in this region of the barrier.

The next two equations focus on just the lowest order terms that include $\hat{A}$ and $\hat{B}$. Using the approximations stated in the last few paragraphs, the lowest order term containing $\hat{A}$ becomes

$$\begin{aligned}&(1-s^*)\epsilon^2\frac{\hat{q}}{2\sqrt{1-\hat{q}^2}}\hat{A}=\\&-\frac{1}{2}\left(\sqrt{3}-1\right)\epsilon^2\hat{A}+\mathcal{O}(\epsilon^{\frac{5}{2}-\alpha-\frac{\beta}{2}})=\mathcal{O}(\epsilon^{\frac{3}{2}-\frac{\beta}{2}}),\end{aligned}\tag{A.17}$$

while for $\hat{B}$ the term becomes

$$\begin{aligned}&-(1-s^*)\frac{\epsilon^2}{2}\left(\sqrt{1-\hat{q}^2}+\frac{\epsilon}{2\sqrt{1-\hat{q}^2}}\right)\hat{B}=\\&-\frac{3}{8}\left(\sqrt{3}-1\right)\epsilon^2\hat{B}+\mathcal{O}(\epsilon^{2-\alpha-\beta})=\mathcal{O}(\epsilon^{1-\beta})\end{aligned}\tag{A.18}$$

The condition for whether a term is discounted as too small or not depends on whether it is larger than the energy term. The unperturbed energies are constant with ϵ, but notice that E is multiplied by ϵ in the Schrödinger equation. Thus, this energy term is proportional to ϵ with some polynomially or exponentially small corrections. Therefore, if a term is higher order than linear in ϵ, discard it since it is smaller than the energy term.

I assume $\beta < 1$, in which case, the $\hat{A}$ terms are all small. For the $\hat{B}$ terms, I need $1 < 2-\alpha-\beta$ in order for the next highest term to contribute, so I need to restrict to $\alpha+\beta < 1$. The only remaining thing to consider is how much the $r(\hat{q})$ term will contribute. The height of the barrier scales like $\epsilon^{1-\beta}$ in this setup, so as long as $\beta > 0$, the barrier term remains relevant as well. With all of these approximations, the Schrödinger equation becomes

$$\begin{aligned} \epsilon E\vec{\psi} = \Bigg(& s^* r\left(\hat{x} - \frac{1}{2}\right) + \frac{2}{3}(\sqrt{3}-1)\hat{x}^2 \\ & -\frac{3}{8}\left(\sqrt{3}-1\right)\epsilon^2\hat{B} + \mathcal{O}\left(\max\{\epsilon^{2-\alpha-\beta}, \epsilon^{\frac{3}{2}-\frac{\beta}{2}}\}\right)\Bigg)\vec{\psi} \end{aligned} \tag{A.19}$$

Eq. A.19 is the final form of the approximated Schrödinger equation in the discrete setting. The continuum limit allows us to treat this as a differential equation which can be solved exactly, so the final step is to take the continuum limit. This limit takes x to a continuous variable in the small ϵ limit, taking $\hat{B} \to \frac{\partial^2}{\partial x^2}$, $\hat{x} \to x$, and $|\psi\rangle \to \psi(x)$. Doing this gives the differential equation

$$\begin{aligned} \epsilon E\psi(x) = \Bigg(& s^* r\left(x - \frac{1}{2}\right) + \frac{2}{3}(\sqrt{3}-1)x^2 \\ & -\frac{3}{8}\left(\sqrt{3}-1\right)\epsilon^2\frac{\partial^2}{\partial x^2} \\ & +\mathcal{O}\left(\max\{\epsilon^{2-\alpha-\beta}, \epsilon^{\frac{3}{2}-\frac{\beta}{2}}\}\right)\Bigg)\psi(x) \end{aligned} \tag{A.20}$$

This can be solved using the parabolic cylinder equations as shown in Ch. 5.

Bibliography

[1] R. Feynman, *Simulating Physics with Computers*, *International Journal of Theoretical Physics* **21** (1982).

[2] D. Deutsch and R. Jozsa, *Rapid solutions of problems by quantum computation*, *Proc. of the Royal Society of London A* **439** (1992).

[3] L. K. Grover, *A fast quantum mechanical algorithm for database search*, *Proc. 28th Annual ACM Symposium on the Theory of Computing* (1996) 212.

[4] P. Shor, *Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer*, *SIAM J. Comput.* **26** (1997) 1484–1509.

[5] T. Kadowaki and H. Nishimori, *Quantum Annealing in the Transverse Ising Model*, *Phys. Rev. E* **58** (1998) 5355, [cond-mat/9804280].

[6] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, *Quantum Computation by Adiabatic Evolution*, quant-ph/0001106.

[7] E. Farhi, J. Goldstone, and S. Gutmann, *A Quantum Approximate Optimization Algorithm*, quant-ph/1411.4028.

[8] S. Jordan, *https://math.nist.gov/quantum/zoo/.* Quantum Algorithm Zoo, accessed 4/5/18.

[9] J. Kelly, *Engineering superconducting qubit arrays for quantum supremacy*, presentation at APS March Meeting, 2018.

[10] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, J. M. Martinis, and H. Neven, *Characterizing Quantum Supremacy in Near-Term Devices*, quant-ph/1608.00263.

[11] IBM Q Experience, *https://quantumexperience.ng.bluemix.net/qx/devices*, accessed 4/5/18.

[12] J. S. Otterbach, R. Manenti, N. Alidoust, A. Bestwick, M. Block, B. Bloom, S. Caldwell, N. Didier, E. S. Fried, S. Hong, P. Karalekas, C. B. Osborn, A. Papageorge, E. C. Peterson, G. Prawiroatmodjo, N. Rubin, C. A. Ryan,

D. Scarabelli, M. Scheer, E. A. Sete, P. Sivarajah, R. S. Smith, A. Staley, N. Tezak, W. J. Zeng, A. Hudson, B. R. Johnson, M. Reagor, M. P. da Silva, and C. Rigetti, *Unsupervised Machine Learning on a Hybrid Quantum Computer*, quant-ph/1712.05771.

[13] N. M. Linke, D. Maslov, M. Roetteler, S. Debnath, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, *Experimental Comparison of Two Quantum Computing Architectures*, *PNAS* **114** (2017) 3305–3310, [quant-ph/1702.01852].

[14] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, C. J. S. Truncik, S. Uchaikin, J. Wang, B. Wilson, and G. Rose, *Quantum annealing with manufactured spins*, *Nature* **473** (2011) 194–198.

[15] T. Albash, V. Martin-Mayor, and I. Hen, *Temperature scaling law for quantum annealing optimizers*, *Phys. Rev. Lett.* **119** (2017) 110502, [quant-ph/1703.03871].

[16] V. S. Denchev, S. Boixo, S. V. Isakov, N. Ding, R. Babbush, V. Smelyanskiy, J. Martinis, and H. Neven, *What is the Computational Value of Finite Range Tunneling?*, *Phys. Rev. X* **6** (2016) 031015, [quant-ph/1512.02206].

[17] S. Mandrá, H. G. Katzgraber, and C. Thomas, *The pitfalls of planar spin-glass benchmarks: Raising the bar for quantum annealers (again)*, *Quantum Sci. Technol.* **2** (2017) 038501, [quant-ph/1703.00622].

[18] L. Brady and W. van Dam, *Quantum Monte Carlo Simulations of Tunneling in Quantum Adiabatic Optimization*, *Phys. Rev. A* **93** (2016) 032304, [quant-ph/1509.02562].

[19] L. Brady and W. van Dam, *Spectral Gap Analysis for Efficient Tunneling in Quantum Adiabatic Optimization*, *Phys. Rev. A* **94** (2016) 032309, [quant-ph/1601.01720].

[20] L. Brady and W. van Dam, *Discrepancies between Asymptotic and Exact Spectral Gap Analyses of Quantum Adiabatic Barrier Tunneling*, *Phys. Rev. A* **95** (2017) 052350, [quant-ph/1609.09137].

[21] L. Brady and W. van Dam, *Necessary Adiabatic Run Times in Quantum Optimization*, *Phys. Rev. A* **95** (2017) 032335, [quant-ph/1611.02585].

[22] L. Brady and W. van Dam, *Evolution-Time Dependence in Near-Adiabatic Quantum Evolutions*, quant-ph/1801.04349.

[23] E. Crosson and A. Harrow, *Simulated Quantum Annealing Can Be Exponentially Faster than Classical Simulated Annealing*, in *Proc. of FOCS 2016*, pp. 714–723, 2016. quant-ph/1601.03030.

[24] Z. Jiang, V. N. Smelyanskiy, S. V. Isakov, S. Boixo, G. Mazzola, M. Troyer, and H. Neven, *Scaling analysis and instantons for thermally-assisted tunneling and Quantum Monte Carlo simulations*, *Phys. Rev. A* **95** (2017) 012322, [quant-ph/1603.01293].

[25] J. Goldstone *Private Communication.*

[26] J. Villain, *Quantum theory of one- and two-dimensional ferro- and antiferromagnets with an easy magnetization plane. I. ideal 1-d or 2-d lattices without in-plane anisotropy*, *J . Physique* **35** (1974) 27–47.

[27] S. Muthukrishnan, T. Albash, and D. A. Lidar, *Tunneling and speedup in quantum optimization for permutation-symmetric problems*, *Phys. Rev. X* **6** (2016) 031010, [quant-ph/1511.03910].

[28] M. Born and V. Fock, *Beweis des adiabatensatzes*, *Zeitschrift für Physik* **51** (1928) 165–180.

[29] A. Safavi-Naini, R. J. Lewis-Swan, J. G. Bohnet, M. Garttner, K. A. Gilmore, E. Jordan, J. Cohn, J. K. Freericks, A. M. Rey, and J. J. Bollinger, *Exploring adiabatic quantum dynamics of the Dicke model in a trapped ion quantum simulator*, quant-ph/1711.07392.

[30] D. A. Lidar, A. T. Rezakhani, and A. Hamma, *Adiabatic approximation with exponential accuracy for many-body systems and quantum computation*, *J. Math. Phys.* **50** (2009) 102106, [quant-ph/0808.2697].

[31] S. Jansen, M. Ruskai, and R. Seiler, *Bounds for the adiabatic approximation with applications to quantum computation*, *J. Math. Phys.* **48** (2007) 102111, [quant-ph/0603175].

[32] T. D. Kieu, *A New Class of Time-Energy Uncertainty Relations for Time-dependent Hamiltonians*, quant-ph/1702.00603.

[33] A. B. Finnila, M. A. Gomez, C. Sebenik, C. Stenson, and J. Doll, *Quantum Annealing: A New Method for Minimizing Multidimensional Functions*, *Chem. Phys. Lett.* **219** (1994) 343–348, [chem-ph/9404003].

[34] H. F. Trotter, *On the product of semi-groups of operators*, *Proc. Am. Math. Soc.* **10** (1959) 545551.

[35] M. Suzuki, *Relationship between d-Dimensional Quantal Spin Systems and $(d+1)$-Dimensional Ising Systems: Equivalence, Critical Exponents and Systematic Approximants of the Partition Function and Spin Correlations*, *Prog. Theor. Phys.* **56** (1976) 1454–1469.

[36] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev, *Adiabatic Quantum Computation is Equivalent to Standard Quantum Computation*, *SIAM J. Comp.* **37** (2007) 166–194, [quant-ph/0405098].

[37] E. Y. Loh, Jr., J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar, *Sign problem in the numerical simulation of many-electron systems*, *Phys. Rev. B* **41** (1990) 9301.

[38] S. Bravyi, D. P. DiVincenzo, R. I. Oliveira, and B. M. Terhal, *The Complexity of Stoquastic Local Hamiltonian Problems*, *Quant. Inf. Comp.* **8** (2008) 0361–0385, [quant-ph/0606140].

[39] M. B. Hastings and M. H. Freedman, *Obstructions To Classically Simulating The Quantum Adiabatic Algorithm*, *Quant. Inf. & Comp. 13, 11-12 (2013)* **13** (2013) 1038, [quant-ph/1302.5733].

[40] E. Crosson and M. Deng, *Tunneling through high energy barriers in simulated quantum annealing*, quant-ph/1410.8484.

[41] M. Marvian, D. A. Lidar, and I. Hen, *On the Computational Complexity of Curing the Sign Problem*, quant-ph/1802.03408.

[42] B. W. Reichardt, *The Quantum Adiabatic Optimization Algorithm and Local Minima*, in *Proc. of the 36th annual ACM Symposium on Theory of Computing (STOC)*, pp. 502–510, 2004.

[43] L. Kong and E. Crosson, *The performance of the quantum adiabatic algorithm on spike Hamiltonians*, quant-ph/1511.06991.

[44] T. Jörg, F. Krzakala, J. Kurchan, A. C. Maggs, and J. Pujos, *Energy gaps in quantum first-order mean-fieldlike transitions: The problems that quantum annealing cannot solve*, *Europhys. Lett.* **89** (2010) 40004, [quant-ph/0912.4865].

[45] Y. Seki and H. Nishimori, *Quantum annealing with antiferromagnetic fluctuations*, *Phys. Rev. E* **85** (2012) 051112, [quant-ph/1203.2418].

[46] V. Bapst and G. Semerjian, *On quantum mean-field models and their quantum annealing*, *J. Stat. Mech.* **2012** (2012) [quant-ph/1203.6003].

[47] B. Seoane and H. Nishimori, *Many-body transverse interactions in the quantum annealing of the p-spin ferromagnet*, *Phys. Rev. A* **45** (2012) 435301, [quant-ph/1207.2909].

[48] Y. Susa, J. F. Jadebeck, and H. Nishimori, *Relation between quantum fluctuations and the performance enhancement of quantum annealing in a nonstoquastic Hamiltonian*, *Phys. Rev. A* **95** (2017) 042321, [quant-ph/1612.08265].

[49] Y. Susa, Y. Yamashiro, M. Yamamoto, and H. Nishimori, *Exponential Speedup of Quantum Annealing by Inhomogeneous Driving of the Transverse Field*, *J. Phys. Soc. Jpn.* **87** (2018) 023002, [quant-ph/1801.02005].

[50] R. M. Fye, *New results on Trotter-like approximations*, *Phys. Rev. B* **33** (1986) 6271.

[51] A. Auerbach, *Interacting Electrons and Quantum Magnetism.* Springer-Verlag, 1994.

[52] S. Coleman, *Aspects of Symmetry.* Cambridge University Press, 1985.

[53] E. Farhi, J. Goldstone, and S. Gutmann, *Quantum Adiabatic Evolution Algorithms versus Simulated Annealing*, quant-ph/0201031.

[54] M. Enz and R. Schilling, *Magnetic field dependence of the tunnelling splitting of quantum spins*, *J. Phys. C: Solid State Phys.* **19** (1986).

[55] A. Boulatov and V. N. Smelyanskiy, *Quantum adiabatic algorithms and large spin tunnelling*, *Phys. Rev. A* **68** (2003) 062321, [quant-ph/0309150].

[56] S. Boixo, V. N. Smelyanskiy, A. Shabani, S. V. Isakov, M. Dykman, V. S. Denchev, M. Amin, A. Smirnov, M. Mohseni, and H. Neven, *Computational Role of Collective Tunneling in a Quantum Annealer*, quant-ph/1411.4036.

[57] W. van Dam and L. T. Brady, *Merits and limitations of asymptotic analyses for quantum adiabatic optimization*, presentation at Adiabatic Quantum Computing Conference, 2016.

[58] S. Boixo and R. D. Somma, *Necessary condition for the quantum adiabatic approximation*, *Phys. Rev. A* **81** (2010) 032308, [quant-ph/0911.1362].

[59] C. De Grandi and A. Polkovnikov, *Adiabatic perturbation theory: from Landau-Zener problem to quenching through a quantum critical point*, in *Quantum Quenching, Annealing and Computation*, vol. 802, pp. 75–114. Springer.

[60] N. Wiebe and N. S. Babcock, *Improved error-scaling for adiabatic quantum evolutions*, *New J. Phys.* **18** (2012) 013024, [quant-ph/1105.6268].

[61] L. Landau, *Zur theorie der energieubertragung. ii*, *Physikalische Zeitschrift der Sowjetunion* **2** (1932) 46–51.

[62] C. Zener, *Non-Adiabatic Crossing of Energy Levels*, *Proc. of the Royal Society of London A.* **137** (1932) 696–702.

[63] M. Boyer, G. Brassard, P. Hoeyer, and A. Tapp, *Tight bounds on quantum searching*, *Fortsch. Phys.* **46** (1998) 493–506, [quant-ph/9605034].

[64] J. Roland and N. J. Cerf, *Quantum Search by Local Adiabatic Evolution*, *Phys. Rev. A* **65** (2002) 042308, [quant-ph/0107015].

[65] W. van Dam, M. Mosca, and U. Vazirani, *How Powerful is Adiabatic Quantum Computation?*, *Proc. of 42nd FOCS* (2001) 279–287, [quant-ph/0206003].