Discrepancies between asymptotic and exact spectral-gap analyses of quantum adiabatic barrier tunneling

Lucas T. Brady and Wim van Dam

Discrepancies between Asymptotic and Exact Spectral Gap Analyses of Quantum Adiabatic Barrier Tunneling

Lucas T. Brady

Department of Physics, University of California, Santa Barbara, CA 93106-5110, USA

Wim van Dam

Department of Computer Science, Department of Physics,
University of California, Santa Barbara, CA 93106-5110, USA

(Dated: April 28, 2017)

We study the asymptotic behavior of the spectral gap of simple barrier tunneling problems, which are related to using quantum annealing to find the global optimum of cost functions defined over n bits. Specifically we look at the problem of having an n qubit system tunnel through a barrier of width and height proportional to n^{α}. We show that with these quantum annealing problems, the asymptotic, $n \to \infty$, behavior of the spectral gap does not accurately describe the behavior of the gap at finite n until extremely large values of n ($n > 10^{12}$). We prove that this deficiency of the asymptotic expression is a feature of simple one-dimensional tunneling problems themselves, casting doubt on whether asymptotic analysis is an appropriate tool for studying tunneling problems in quantum annealing for reasonably sized systems.

I. INTRODUCTION

Quantum tunneling is one of the major aspects of quantum mechanics that sets it apart from classical mechanics. The ability of a quantum system to enter a classically disallowed region and, through it, reach an energetically favorable region has inspired many applications. In this article, we are concerned with tunneling as it pertains to quantum annealing, which is a quantum algorithm that uses quantum effects, such as tunneling, to relax a system into a desired ground state.

Quantum adiabatic optimization (QAO) [1] is a form of quantum annealing that solves optimization problems by adiabatically changing the Hamiltonian for a system from one with a known ground state to one with the optimization solution encoded in its ground state. Many studies [2–13] have sought to characterize what, if any, speed-up this algorithm has over classical alternatives such as simulated annealing or simulated quantum annealing.

A popular belief among the QAO community is that at least part of quantum annealing's power comes from its ability to tunnel through barriers more easily than classical algorithms can go over them. Several recent results [8, 9, 12–14] have shown that tunneling alone is not enough to ensure a quantum speed-up in one-dimensional stoquastic settings. It is currently unclear how much power quantum adiabatic computing with stoquastic Hamiltonians has, with several studies [5, 15] showing theoretical speed-ups, that rely on tunneling in higher dimensions, over some classical algorithms.

Nevertheless, adiabatic quantum computation is provably universal for quantum computation if non-stoquastic Hamiltonians are allowed [16], 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 article studies quantum tun-

neling as part of QAO in the simplest possible situation: tunneling in position-space through a one-dimensional barrier. In a previous article [17], we showed that this simple problem 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, we 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. We find that $n > 10^{12}$ is needed to ensure that asymptotic expressions accurately describe behavior, to the point where polynomial and exponential run time scaling can be confused for lower n. Given the current and near-term sizes of quantum computer implementations [18], 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 II we provide a brief review of QAO and introduce the details of the barrier tunneling problem to be studied. We mention both the original hypercube problem and the simpler continuous one-dimensional problem.

Section III gives the first evidence that the asymptotic limit may not be accurate even for moderately large n. We 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 [17].

We compare the exact spectral gap to its asymptotic approximation in Section IV 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 V uses properties of the derivatives of log-log data to analyze how well a curve approximates a power law versus an exponential fit. Furthermore, we show how easily finite data for tunneling could lead to a mislabel-

ing of the scaling behavior of the algorithm. Finally, Section VI concludes and provides an outlook for the future of such asymptotic analysis.

II. TUNNELING IN QUANTUM ADIABATIC OPTIMIZATION

Quantum annealing derives its name by comparison with classical simulated annealing, which starts a system at a high temperature and then slowly lowers the temperature, trying to relax the system into its ground state. Quantum annealing works on the same principle except it uses quantum effects instead of temperature.

A. Quantum Adiabatic Optimization

As it is usually formulated [1], the quantum annealing Hamiltonian is a linear time interpolation between two static Hamiltonians:

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

where $s = t/\tau$ is the time parameter, t, divided by the total predetermined runtime of the algorithm, τ. $\hat{H}_0$ is the initial Hamiltonian and is usually chosen to have an easily prepared ground state. $\hat{H}_1$ encodes the solution to an optimization problem in its ground state configuration. The initial formulation of QAO [1] sets this in a Hilbert space of n qubits where $\hat{H}_0$ is a uniform field in the negative x direction:

$$\hat{H}_0 = -\sum_{i=1}^{n} \sigma_x^{(i)}, \tag{2}$$

and $\hat{H}_1$ is diagonal in the z basis with a cost function encoded along its diagonal:

$$\hat{H}_1 = \sum_{z \in \{0,1\}^n} f(z)|z\rangle\langle z|. \tag{3}$$

These canonical choices for $\hat{H}_0$ and $\hat{H}_1$ make the the algorithm easier to create and analyze, but they also ensure that $\hat{H}(s)$ is stoquastic or, alternatively, has no sign problem, meaning its off-diagonal terms are all non-positive. The universality proof [16] for QAO requires non-stoquastic Hamiltonians, so it is currently unclear how much power stoquastic Hamiltonians have in QAO.

Key to QAO is that the algorithm must run adiabatically, where adiabatically means that the total time, τ, scales according to a combination of matrix norms of time derivatives of $\hat{H}(s)$ and the spectral gap, $g(s)$, between the ground state and first excited state. The typically stated adiabatic condition says that adiabaticity is ensured if $\tau \gg \max_{s\in[0,1]} ||\frac{d\hat{H}(s)}{ds}||/(g(s))^2$. The derivation of this version of the adiabatic theorem is subtly incorrect, which has led to more robust statements and derivations of an adiabatic condition [19], but all of these more

accurate conditions claim that the worst-case bound on τ depends on matrix norms, which in our case are linear in n, and the inverse of $g_{\min} := \min_{s\in[0,1]} g(s)$. Therefore, the majority of this paper focuses on the calculation of $g_{\min}$ since this is the primary changing quantity that determines the worst-case bound on the run time of QAO.

B. Symmetric Barrier Tunneling Problem

Our 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. [2], and the general bit-symmetric barrier tunneling problem has been analyzed extensively by other groups [8, 9, 12–14, 17, 20, 22]. 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 a previous article [17], we showed that this Hamiltonian can be mapped onto the symmetric $(n + 1)$-dimensional subspace where it describes a single spin $1/\varepsilon := n/2$ particle. Furthermore, using a modified version of the Villain transformation [21], we 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 we only care about the spectral gap between the ground state and first excited state, this restriction is acceptable. In [17], we restrict ourselves to a square barrier so that our final, semiclassical Schrödinger equation becomes

$$\frac{d^2\psi}{dx^2} = \varepsilon^{-2}\left[V(x) - \varepsilon cE\right]\psi(x), \tag{4}$$

where

$$V(x) = \begin{cases} \omega\varepsilon^{1-\alpha} & \text{if } -a < x < a \\ \omega^2 x^2 & \text{otherwise} \end{cases} \tag{5}$$

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 $\varepsilon^{1-\alpha} =: 2a$ and height $\omega\varepsilon^{1-\alpha}$. The constants can either be taken as arbitrary or set to $c = 8/(3(\sqrt{3} - 1))$ and $\omega = 4/3$ if we want a direct correlation to the qubit problem.

While Eq. 4 was 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, we just have a particle moving in a one-dimensional quadratic well with a barrier in the middle. For most of this article, we treat the Schrödinger equation in Eq. 4 as its own independent problem.

The continuous well differential equation, Eq. 4, can be solved using exponentials and parabolic cylinder functions, $D_\nu(y)$, up to a transcendental equation for the eigenenergies:

$$k_\pm D_{\nu_\pm}\left(\sqrt{\frac{2\omega}{\varepsilon}}a\right)\left(e^{k_\pm a} \mp e^{-k_\pm a}\right) =$$

$$\sqrt{\frac{2\omega}{\varepsilon}}D'_{\nu_\pm}\left(\sqrt{\frac{2\omega}{\varepsilon}}a\right)\left(e^{k_\pm a} \pm e^{-k_\pm a}\right), \qquad (6)$$

where $\nu_\pm := \frac{cE_\pm}{2\omega} - \frac{1}{2}$, $k_\pm := \sqrt{\omega\varepsilon^{-1-\alpha} - \varepsilon^{-1}cE_\pm}$, and $E_\pm$ refers to even $(+)$ or odd $(-)$ energy solutions. In the limit of $n \to \infty$, Eq. 6 can be solved for the lowest two energy levels, resulting in a gap expression

$$g_{\min} = \frac{8\omega^{1/2}}{c\sqrt{\pi}}\varepsilon^{2\alpha-1/2} \qquad (7)$$

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

$$g_{\min} = \frac{16\omega}{c\sqrt{\pi}}\varepsilon^{\alpha/2}\exp\left(-\sqrt{\omega}\varepsilon^{\frac{1}{2}-\frac{3}{2}\alpha}\right) \qquad (8)$$

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 [2, 13, 14, 20].

C. Definitions of Gap Analyses

In Subsection II B we outline three steps in the approximation procedure from which the gap can be extracted. In the rest of the paper we refer back to these three different models, so we take a moment here to label these three models and define exactly what we 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. 4, and the exact spectral gap can be extracted by numerically solving the transcendental equation for the energies, Eq. 6. For most of this article, the continuous problem is treated as an independent problem, and most of our efforts focus on comparing its exact spectral gap to the asymptotic limit.
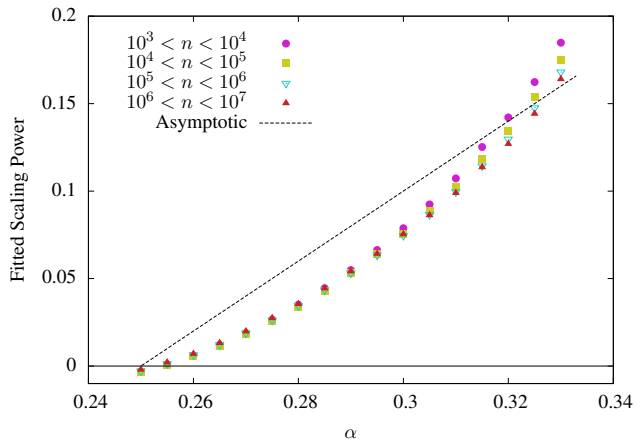

FIG. 1: Fitted scaling exponent of $g_{\min}$ for the symmetric qubit tunneling problem with a binomial shaped barrier. Each data point shows what scaling power, p, is extracted from a fit to $g_{\min} = An^{-p}$, assuming the fit is done on data from a finite range of n versus the scaling power of the 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.

- **Asymptotic analysis**: In the large n limit, the asymptotic expression for the spectral gap of the discrete and continuous problems can be derived. We did this derivation in a previous paper [17], and the gap expressions are reproduced in Eqs. 7 and 8. This version of the gap is the true asymptotic limit of the discrete and continuous gaps, but much of this article focuses on how large n must be for the asymptotic limit to be a good approximation of the true gap.

III. SCALING POWER COMPARISON

When analyzing an algorithm numerically or experimentally, a typical question to ask is how its run time, or in our case the gap, scales with the input size n. Therefore, in this section we analyze the gap numerically mostly ignoring our asymptotic expressions from Eqs. 7 and 8 and try to determine what asymptotic information we would extract from numerical data of n vs. $g_{\min}$. We then compare these finite n extrapolations to see how close these results are to our analytic expressions for the asymptotic behavior of $g_{\min}$.

We use symmetry to reduce the discrete qubit Hamiltonian to a tridiagonal $(n+1)$-dimensional matrix 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. 1 outlines the resulting p values for these fits done on simulations with $1/4 < \alpha < 1/3$, leading to polynomial scaling asymptotically.
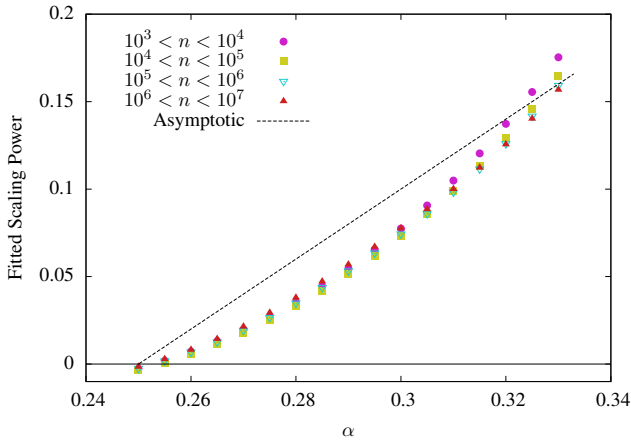
FIG. 2: Fitted scaling exponent of $g_{\min}$ for the continuous well DE, Eq. 4, with a square barrier. Each data point shows what scaling power, p, is extracted from a fit to $g_{\min} = An^{-p}$, assuming the fit is done on data from a finite range of n versus the scaling power of the barrier size, 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. 1, indicating that the problem in that plot was not in the approximations made for the Villain transformation.

Also, note that the data in Fig 1 comes from a problem with a binomial shaped barrier rather than the standard square barrier we consider. To collect these data points, we restrict ourselves 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 us to collect more data points than the square barrier, whose width grows with n^{α}. Therefore, we use the binomial barrier in order to collect more data points and create a cleaner data set.

In Fig. 1, we 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 [17].

However, the only possible source of error in the asymptotic analysis in [17] is in the Villain transformation and the other assumptions that lead to the derivation of the continuous well DE, Eq. 4. Therefore, if the spectral gap for the differential equation shows the same difficulty in reaching its asymptotic value, we know that extremely lare 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. We do in fact see that if we perform the same procedure that gave us Fig. 1 for nu-

merical $g_{\min}$ values found from solving Eq. 6, we obtain a very similar plot, Fig. 2.

The failure of $g_{\min}$ to quickly reach its asymptotic scaling even in Fig. 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, we focus exclusively on the continuous problem, ignoring the discrete problem that motivated it. In the next section, we examine what n are necessary to bring the scaling close to the asymptotic expressions.

IV. COMPARISON OF EXACT GAP AND ASYMPTOTIC EXPRESSIONS

Given the large discrepancy seen in Section III 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 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, we see similar results for different values of α.

A. Direct Gap Comparison

Since our asymptotic analysis of $g_{\min}$ in Eqs. 7 and 8 is exact enough to get the constant prefactor, we can compare the numerical gap directly to the asymptotic expression to visually and numerically see where agreement occurs.

In Fig. 3 we show the discrete, continuous, and asymptotic gaps for $\alpha = 3/10$ up to $n = 10^6$. 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. 3 shows how far away the actual gap is from the asymptotic expression, leading to the trough in Fig. 2.

With our current computational limits, we cannot push exact diagonalization of the symmeterized n qubit Hamiltonian to much higher n, but the condition in Eq. 6 can be solved for extremely large n with rounding error being the only block. Therefore, we can extend Fig. 3 to higher n to get Fig. 4 from which we conclude that the continuous gap does eventually approach its asymptotic value, but only for extremely large $n > 10^{12}$.
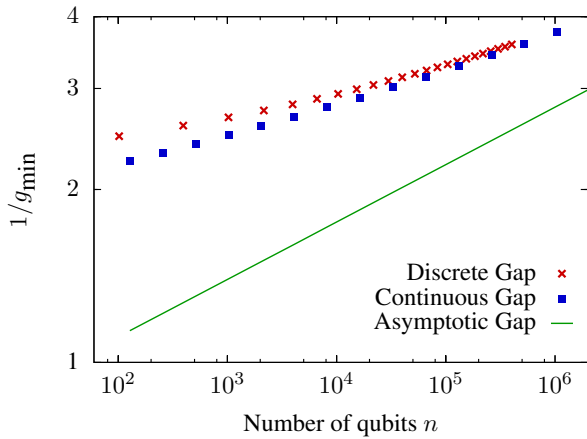
FIG. 3: Spectral gap g_{min} vs. n for $\alpha = 3/10$ for our three different gap extraction methods with a square barrier. The discrete gap and continuous gap well approximate each other at these relatively modes n, but neither of them are well approximated by the asymptotic expression for the gap. This difference seems to confirm the failure of the asymptotics observed in Figs. 1 and 2.
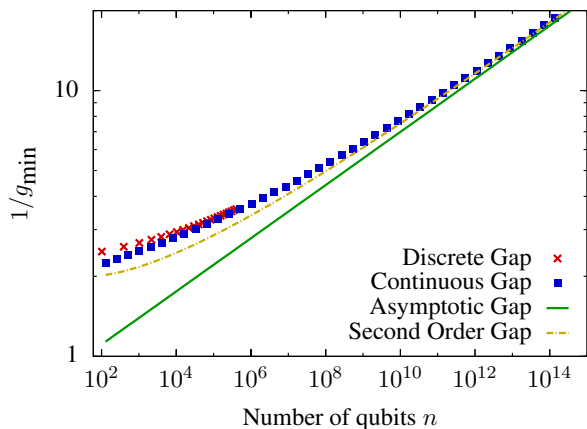

FIG. 4: Spectral gap g_{min} vs. n on a log-log scale for $\alpha = 3/10$ for our three different gap extraction methods with a square barrier. This plot covers a much larger range than Fig. 3, allowing us to see that the continuous gap does eventually approach the asymptotic gap. However, this approach occurs at extremely large n that are impractical for most current applications. In addition, we have included the second order expansion of the asymptotic large n expression from Eq. 9 that more closely approximates the continuous gap.

Fig. 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

$$g_{min}^{(2^{nd}O)} = \frac{8\omega^{1/2}}{c\sqrt{\pi}}\varepsilon^{2\alpha-1/2} - \frac{16}{c\pi}(\ln 2 + \psi_0(-\frac{1}{2}))\varepsilon^{4\alpha-1} \quad (9)$$

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 asymptotic limit to accurately describe its behavior at any reasonable n, this data casts doubt on the efficacy of asymptotic analysis.

We have exclusively used parabolic wells up to this point, so the argument could be made that this large-n behavior is an artifact of our choice of wells. To counter this argument, we 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, we have noted that the smoother the potential is the faster g_{min} begins to approximate its asymptotic value. We have chosen the quadratic potential since it makes the differential equation solvable, but we have run a few other simulations with different shaped wells and barriers. While a smoother barrier (such as a Gaussian) does have a gap that reaches its asymptotic behavior faster 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 we cannot get exact asymptotic scaling expressions for them up to the constant prefactors.

B. Ratio of Exact Gap to Asymptotic

We may also want to know what size of n is necessary before the continuous gap is a certain fraction of the asymptotic gap. In Fig. 5, we plot the n necessary for the ratio of the continuous gap to the asymptotic gap to reach the specified value. When we 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. 5 at first seems to contradict the information in Fig. 2 since these two plots claim that opposite regions are worse. This is not a contradiction because the bad cases in Fig. 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
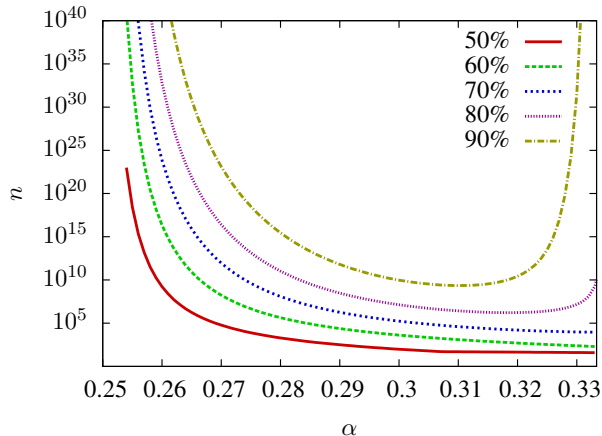
FIG. 5: We show the n necessary for the ratio of the continuous gap and the asymptotic gap to reach the indicated percentages as a function of α. Especially close to the boundaries of the polynomial region, the system requires extremely large n to reach its asymptotic value. When coupled with Fig. 2, this plot indicates that for the middle values of α the continuous gap reachs its asymptotic value faster by deviating more from its asymptotic scaling.

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, let us analytically consider the ratio between the leading order asymptotic gap and the second order asymptotic gap

$$\frac{g_{\min}^{(2^{\text{nd}}\text{O})}}{g_{\min}^{(1^{\text{st}}\text{O})}} = 1 - \frac{2}{\sqrt{\pi\omega}}(\ln 2 + \psi_0(-\frac{1}{2}))\varepsilon^{2\alpha-1/2}$$

$$= 1 - \kappa\varepsilon^{2\alpha-1/2}, \qquad (10)$$

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

In our 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 we want this ratio to reach some threshhold value $0 < v < 1$, then the necessary $n(v)$ becomes

$$n(v) = 2\left(\frac{\kappa}{1-v}\right)^{1/(2\alpha-1/2)}. \qquad (11)$$

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. 11 would seem to indicate that the necessary n should be relatively small for

$2\alpha-1/2$ farther from zero as we are 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. 5 for large v.

This analysis can be elucidated by looking at the general case of a gap described by $g = A\varepsilon^{\gamma_1} + B\varepsilon^{\gamma_2}$ for $\gamma_1 < \gamma_2$. Our analysis says that the first order expansion with γ_1 becomes poor if $\gamma_2 - \gamma_1$ is small. Thus, much of our 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.

V. 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. 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, we 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. 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. Our ultimate goal is to be able to take data like that presented in Fig. 3 and judge how likely it is that the data is obeying a power law or exponential relationship.

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

$$g_{\min} = An^{-p} \Rightarrow f_{\text{poly}} = \ln A - px, \qquad (12)$$
$$g_{\min} = B\exp(-Cn^q) \Rightarrow f_{\text{exp}} = \ln B - Ce^{qx}. \qquad (13)$$

Our analysis technique relies on properties of f and its derivatives with repspect to x. In Fig. 2, when we did linear fits to log-log data, we were 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 we cannot immediately extract any additional information from it about the exponential.
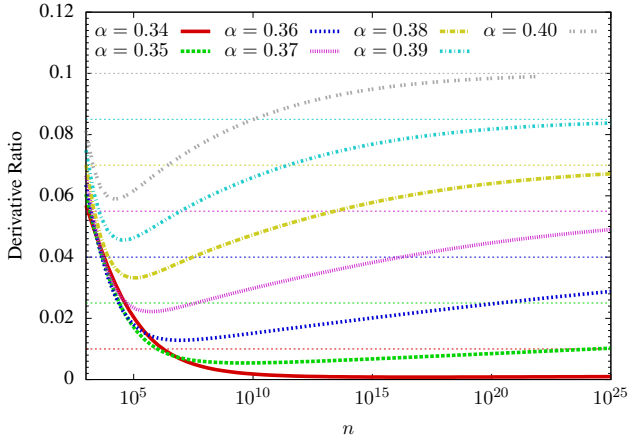
FIG. 6: The derivative ratio $R = f''(x)/f'(x)$ where $f :=$ $\ln g_{\min}$ and $x := \ln n$ for the continuous gap. For true exponential scaling, this ratio should approach a constant equal to the power of n in the exponential. The dashed lines are what this ratio should be asymptotically. None of these 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.

FIG. 7: The derivative ratio $R = f''(x)/f'(x)$ where $f :=$ $\ln g_{\min}$ and $x := \ln n$ for the continuous gap. These α values were chosen in the polynomial scaling region, and for true power law scaling, we would expect the ratio $R = 0$. We 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 our 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.

A. Quality of Fits

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

$$R_{\text{poly}} = \frac{f''_{\text{poly}}(x)}{f'_{\text{poly}}(x)} = 0$$

$$R_{\text{exp}} = \frac{f''_{\text{exp}}(x)}{f'_{\text{exp}}(x)} = q. \tag{14}$$

Note that only exponential scaling gives a non-zero, constant derivative ratio which gives us another metric to judge if a function has exponential scaling and allows us 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 we do not immediately reach asymptotic scaling, asymptotic polynomial and sub-exponential scaling are not immediately distinguishable through this method.

In Fig. 6, we 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 asymp-
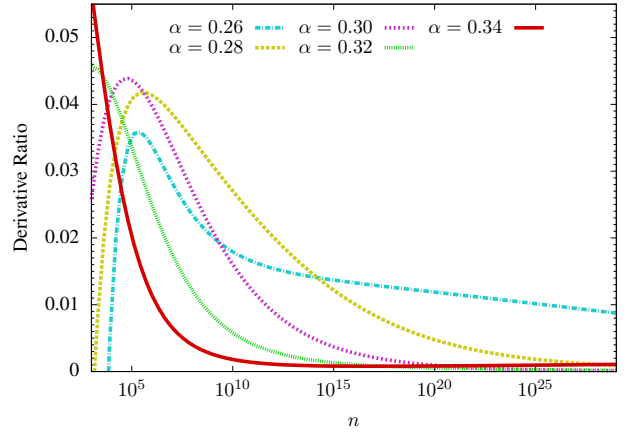
totic values for the n displayed, but these systems still require extremely large n.

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 our goal is to extract information on the type of exponential scaling, Fig. 6 implies that we need very large n.

Asymptotically polynomial scaling should give us an R ratio that is zero, so in Fig. 7, we 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 as we see, 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. 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. 7 appears to approach zero faster than many of the ratios for α actually in the polynomial scaling region. Therefore, 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. 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. 5 and 7.

VI. CONCLUSION

The problem of tunneling through a barrier in a quadratic well is a simple quantum mechanical system with elements that are not unfamiliar to an introductory quantum mechanics class, so it is surprising that asymptotics have difficulty describing such an elementary system at reasonable system sizes. Asymptotic analysis is deemed useful because it is assumed that the run time of an algorithm is well approximated by its asymptotic behavior starting at relatively small system sizes, n.

Our study shows that barrier tunneling is a case where asymptotic behavior is not a good descriptor of spectral gap and therefore run time scaling until extremely large n. This failure of asymptotic analysis is worsened by the direction of the failure: asymptotically polynomial algorithms appear to locally scale in a super-polynomial manner as seen in Fig. 7, making it more difficult to justify the use of asymptotic analysis for small system sizes.

In terms of the actual effect on the adiabatic algorithm, which depends on the inverse square of the gap, the gaps we see that are two or three times lower than the asymptotic prediction would lead to algorithms running for or nine times slower for small system sizes. More importantly, the scaling behavior of these algorithms for finite n is not close to the asymptotic scaling. This discrepancy would lead a naive analysis of a numerically or experimentally implemented algorithm to misinterpret the asymptotic scaling.

Our results seem to hold numerically for other one-dimensional tunneling problems, with smoother barriers giving only a few orders of magnitude improvement in the n needed for agreement with asymptotics. Our analysis of a second order gap correction in section IV B is general, and gap sizes will be far from the asymptotics whenever the scaling exponent on the second order correction is close to that on the first order correction. We do see in section V that scaling discrepancies are larger close to scaling boundaries, such as the boundary between the polynomial and exponential scaling regions. It is possible that asymptotic scaling is more accurate farther from such boundaries, such as in the spike tunneling analysis of [2], but more study is needed to confirm this conjecture in other systems.

Quantum tunneling is a key part of many near-term quantum algorithms, including quantum adiabatic optimization. These algorithms are gaining popularity because of their applicability on smaller, possibly non-error-corrected, quantum computers. Our results on the difficulty of describing barrier tunneling in small systems with asymptotic analysis suggest that care should be taken when analyzing the potential power of these algorithms for near-term experiments and systems.

[1] E. Farhi, J. Goldstone, S. Gutmann, M. Sipser, quant-ph/0001106 (2000).
[2] E. Farhi, J. Goldstone, S. Gutmann, quant-ph/0201031 (2002).
[3] E. Farhi, J. Goldstone, S. Gutmann, D. Nagaj. Int. J. Quantum Inf. **6**, 3 (2008).
[4] R. Martoňák, G. E. Santoro, E. Tosatti. Phys. Rev. B **66**, 094203 (2002).
[5] M. B. Hastings, M. H. Freedman, Quant. Inf. & Comp. **13**, 11-12 (2013).
[6] S. Boixo, T. F. Rønnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, M. Troyer, Nature Phys. **10**, 218 (2014).
[7] D. Battaglia, G. Santoro, E. Tosatti, Phys. Rev. E **71**, 066707 (2005).
[8] E. Crosson, M. Deng, quant-ph/1410.8484 (2014).
[9] E. Crosson, A. Harrow, Proc. of FOCS 2016, pp. 714-723 (2016).
[10] B. Heim, T. F. Rønnow, S. V. Isakov, M. Troyer, Science **348**, 6231 (2015).
[11] S. Muthukrishnan, T. Albash, D. A. Lidar, Phys. Rev. X **6**, 031010 (2016).
[12] L. Brady, W. van Dam, Phys. Rev. A **93**, 032304 (2016).
[13] L. Kong, E. Crosson, quant-ph/1511.06991 (2015).
[14] Z. Jiang, V. N. Smelyanskiy, S. V. Isakov, S. Boixo, G. Mazzola, M. Troyer, and H. Neven, Phys. Rev. A **95**, 012322 (2017). (2016).
[15] E. Andriyash, M. H. Amin, quant-ph/1703.09277 (2017).
[16] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Llyod, O. Regev, SIAM J. Comp. **37**, 166 (2007).
[17] L. Brady, W. van Dam, Phys. Rev. A **94**, 032309 (2016).
[18] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, J. M. Martinis, H. Neven, quant-ph/1608.00263 (2016).
[19] S. Jansen, M. Ruskai, R. Seiler, J. Math. Phys. **48**, 102111 (2007).

[20] J. Goldstone (unpublished).
[21] J. Villain, J . Phys. France, **35**, 27 (1974).
[22] B. W. Reichardt, in Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC'04), ACM Press(2004).