Necessary adiabatic run times in quantum optimization

Lucas T. Brady*
Department of Physics, University of California, Santa Barbara, California 93106-5110, USA

Wim van Dam
Department of Computer Science, Department of Physics, University of California, Santa Barbara, California 93106-5110, USA

(Received 6 February 2017; published 30 March 2017)

Quantum annealing is guaranteed to find the ground state of optimization problems provided it operates in the adiabatic limit. Recent work [S. Muthukrishnan *et al.*, Phys. Rev. X **6**, 031010 (2016)] has found that for some barrier tunneling problems, quantum annealing can be run much faster than is adiabatically required. Specifically, an n-qubit optimization problem was presented for which a nonadiabatic, or diabatic, annealing algorithm requires only a constant run time, while an adiabatic annealing algorithm requires a run-time polynomial in n. Here we show that this nonadiabatic speedup is the direct result of a specific symmetry in the studied problem. In the more general case, no such nonadiabatic speedup occurs and we show why the special case achieves this speedup compared to the general case. We also prove that the adiabatic annealing algorithm has a necessary and sufficient run time that is quadratically better than the standard quantum adiabatic condition suggests. We conclude with an observation about the required precision in timing for the diabatic algorithm.

DOI: 10.1103/PhysRevA.95.032335

I. INTRODUCTION

Recent work in quantum adiabatic optimization [1] has focused on a class of Hamming-symmetric problems that exhibits extremely strong nonadiabatic speedups over a slower adiabatic approach. Numerical evidence presented by Muthukrishnan, Albash, and Lidar [2] shows that for several barrier tunneling problems on n qubits, a well-calibrated constant time evolution of the quantum annealing Hamiltonian is sufficient. Thus, this algorithm significantly improves upon the slower adiabatic evolution of the Hamiltonian, which can take polynomial or even exponential time in n. 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.

Usually the sufficient run time of quantum adiabatic optimization is estimated using the standard adiabatic condition. This condition says that adiabaticity is ensured if the running time grows as

$$\tau \gg \max_{s\in[0,1]} \left\| \frac{\partial \hat{H}(s)}{\partial s} \right\| \Big/ g(s)^2 \tag{1}$$

for the spectral gap $g(s)$. More accurate versions of this condition have been proven [3], but all of them depend linearly on the matrix norm of $\hat{H}(s)$ or its derivatives with respect to s divided by a low-degree polynomial function of the gap $g(s)$.

The condition in Eq. (1) is merely a sufficient condition, and it is possible to have adiabatic evolutions with shorter running times than Eq. (1) describes. Furthermore, it is also possible to have a nonadiabatic evolution that succeeds in solving the optimization problem at hand. It is such a nonadiabatic speedup that is described by Muthukrishnan *et al.* [2].

A nonadiabatic speedup is obviously significant for near-term quantum computers where quantum annealing is a potential application. Kong and Crosson [4] have studied these diabatic transitions, and more recently the current authors presented complementary findings [5]. These recent results indicate that this nonadiabatic 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 [2,4,6–11] to study the more general properties of quantum annealing in the presence of a barrier.

Here we present results that indicate that even slightly more generalized versions of symmetric barrier tunneling problems do not exhibit this fast nonadiabatic speedup. The base Hamiltonian used to study this class of problems exists in a Hilbert space of n qubits and is given by

$$\hat{H}(s) = -(1-s)\sum_{i=1}^{n}\sigma_x^{(i)} + s\left[\sum_{i=1}^{n}\sigma_z^{(i)} + b\left(\sum_{i=1}^{n}\sigma_z^{(i)}\right)\right], \tag{2}$$

where $b(h)$ is some localized barrier or perturbation and $s = t/\tau$ is a normalized time variable representing the linear progression of time, t, from $t = 0$ to the algorithm stopping time τ. Current numerical evidence [2] suggests that the nonadiabatic speedup exists for many classes, shapes, and sizes of localized barriers $b(h)$. This article generalizes the problem slightly [ignoring $b(h)$ for the moment],

$$\hat{H}(s) = -(1-s)\sum_{i=1}^{n}\sigma_x^{(i)} + s\mu\sum_{i=1}^{n}\sigma_z^{(i)}, \tag{3}$$

by introducing a positive slope parameter μ, and we find that for the generic case $\mu \neq 1$, the nonadiabatic speedup no longer exists. We 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.

Since this Hamiltonian describes a simple toy model, it is unlikely that a physical system will exhibit the exact $\mu = 1$ behavior, leading us to the conclusion that for realizable problems, this diabatic speedup will not exist. In this article, we focus on the $b(h) = 0$ case since it decouples all the

*Corresponding author: lbrady@physics.ucsb.edu

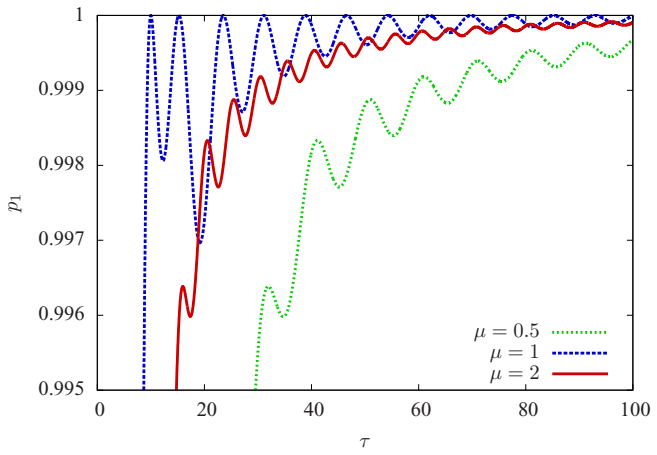

FIG. 1. The single qubit success probability, p_1, as a function of the total run time for several μ values. The dashed blue, $\mu = 1$ line corresponds to the model that has been studied in previous articles. Note that the $\mu = 1$ curve has several special properties, including that it goes to $p_1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

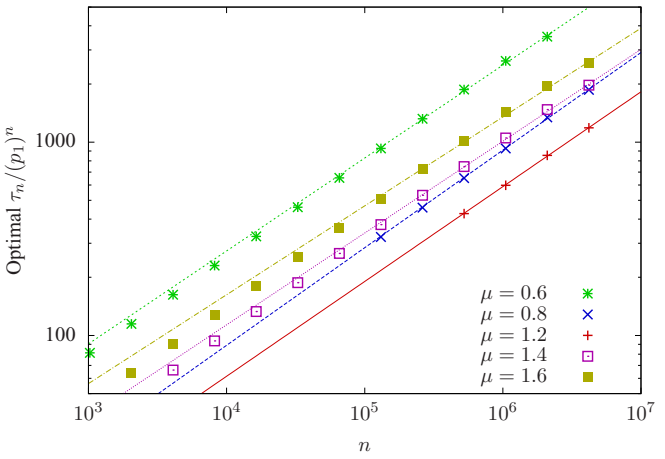

FIG. 2. Optimal expected running time of quantum annealing, $\tau_n/p_n(\tau_n)$, as a function of n for different μ values. Unlike the $\mu = 1$ case, τ_n increases with n for these μ values. Lines through the symbols are power-law fits of the form $\tau_n = An^r$, and the fitted r values in the order $\mu = (0.6,0.8,1.2,1.4,1.6)$ 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 [2] and our results below while being quadratically faster than the sufficient adiabatic condition.

qubits, allowing us to extract information about the system by studying the evolution of a single-qubit Hamiltonian. Since $\mu \neq 1$ disrupts the nonadiabatic speedup even in this $b(h) = 0$ case, we fully expect similar disruptions to occur for more complicated barriers and perturbations.

II. OPTIMAL RUN TIME

We first need to define our criteria for an optimal run time. 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 our case, we have 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 shown in Fig. 1, meaning that $p_n = 1$ at this value, leading to the nonadiabatic speedup noted in other studies. Figure 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 we note below, this failure to reach $p_1 = 1$ for finite τ ultimately leads to the breaking of the nonadiabatic speedup. Therefore, this speedup is restricted to the special case of $\mu = 1$.

Muthukrishnan *et al.* [2] 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. Our 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 nonadiabatic speedup in the $\mu \neq 1$ cases, consider Fig. 2, which shows the optimal expected run time, $\tau_n/p_n(\tau_n)$, 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 nonadiabatic algorithm that runs in constant time. The curves are fitted to power laws of the form $\tau_n = An^p$, and all of the fitted p values are close to $1/2$, indicating a running time of $O(\sqrt{n})$.

It should be noted that the nonadiabatic 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, which would not be possessed in a realistic setting. The lack of robustness of our 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.* [2] show that a fine-tuned driver Hamiltonian can solve these problems nonadiabatically, but our work shows that this speedup does not allow for a general algorithm but only specific algorithms tailored to the problem at hand.

III. SINGLE-QUBIT SUCCESS PROBABILITY

We can extract the $\sqrt{n}$ running-time behavior from the curves in Fig. 1 as well because the qubits in our problem are completely decoupled. For sufficiently long running times τ, the curves of the single-qubit success probability p_1 as a function of τ shown in Fig. 1 are bounded above and below by envelopes of the form

$$1 - \frac{c_\ell(\mu)}{\tau^q} < p_1 < 1 - \frac{c_u(\mu)}{\tau^q}, \tag{4}$$

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 shown in Fig. 3, and for all our fits to different μ data, q is close to 2. Note that $c_u(1) = 0$, which, as we see,

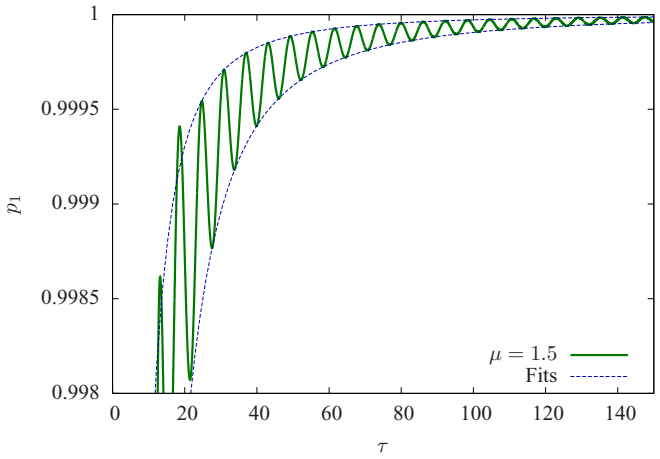

FIG. 3. A single-qubit success probability curve as a function of the total run time τ for $\mu = 1.5$, with upper and lower envelopes shown. These envelopes were obtained by first extracting all the local minima (maxima) and doing a power-law fit of the form $p_1 = 1 - c\tau^{-q}$. The first two minima 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 we use in the text.

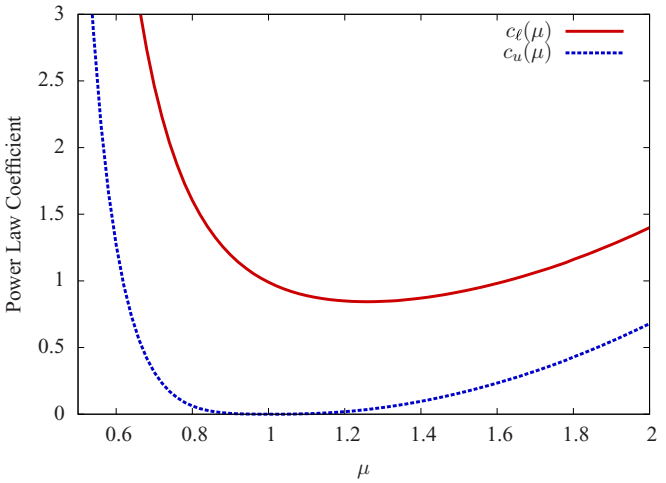

FIG. 4. Curves like the one in Fig. 3 are bounded above and below by curves of the form $1 - c/\tau^2$. We show the values of c for the upper, c_u, and lower, c_ℓ, bounding functions as obtained from numerical fits. These coefficients are a function of μ, and all of the fits used to obtain this data were good quality. In the text, we show that these bounding curves directly lead to a $O(\sqrt{n})$ running time for the algorithm in all cases except the $\mu = 1$ case, where $c_u(1) = 0$.

is one of the main reasons why the $\mu = 1$ diabatic speedup can occur.

Muthukrishnan *et al.* [2] showed that the lower envelope with $c_\ell(\mu)$ guarantees that the worst-case running time for the $\mu = 1$ case scales as $O(\sqrt{n})$. We employ their method to show that a relationship such as Eq. (4) provides both the necessary and the sufficient condition for the running time. Muthukrishnan *et al.* also applied methods created by Boixo and Somma [12] 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. (4) tells us that

$$\left(1 - \frac{c_\ell(\mu)}{\tau^q}\right)^n \leqslant p \leqslant \left(1 - \frac{c_u(\mu)}{\tau^q}\right)^n. \tag{5}$$

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

$$\left(\frac{c_u(\mu)}{\ln 1/p}\, n\right)^{1/q} \leqslant \tau \leqslant \left(\frac{c_\ell(\mu)}{\ln 1/p}\, n\right)^{1/q}. \tag{6}$$

Therefore, since $q = 2$ in our cases, having a running time that scales as $\sqrt{n}$ is both a necessary and a 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 nonadiabatic 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. (1), would predict $O(n)$ scaling; whereas, our results indicate that a faster $O(\sqrt{n})$ running time is sufficient. This result was shown in [2] for $\mu = 1$, and our 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 our problem more specifically and that provide a stricter bound that matches our results. Jansen, Ruskai, and Seiler [3] 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 $0 \leqslant t \leqslant \tau$ is bounded by

$$p = 1 - O(\tau^{-2}). \tag{7}$$

If we take this to be the probability of success for a single-qubit case, our results in Eqs. (5) and (6) imply that $\tau \in O(\sqrt{n})$ is sufficient for an adiabatic evolution. This shows that the result of Jansen *et al.* provides a stricter sufficient condition than the standard adiabatic condition for our optimization problem with decoupled qubits.

In Fig. 4 we 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. 3, making us confident in the $1/\tau^2$ scaling of the error. Note that as we approach the special case $\mu = 1$ we see that $c_u(\mu) \to 0$ and observe that around $\mu = 1$ the coefficient $c_u(\mu)$ stays close to 0. Hence for μ approximately (but not exactly) 1, the nonadiabatic speedup will persist for a large range of n until the adiabatic running time of $O(\sqrt{n})$ is required again at very large n.

IV. SCALING OF THE TRUE ADIABATIC RUN TIME

All of our work so far has shown that the optimal running time of this algorithm is $O(\sqrt{n})$, but this does not imply that the optimal running time results from adiabatic evolution. If we look at the occupancy of the energy states for these optimal runs, we in fact see the ground state being depopulated during the s evolution. Therefore, a remaining question to ask is whether this behavior also holds if we require the system to

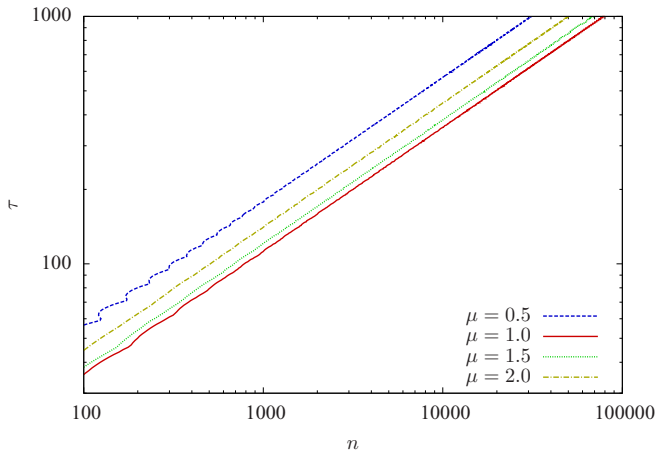

FIG. 5. This plot shows the run time, τ, needed to ensure that the state of the system is at least 75% in the ground state over the entire s evolution. This growth of τ with n comes closest to a true adiabatic evolution, and we can see that the $\tau \in O(\sqrt{n})$ behavior holds even in this case. 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 we see over the sufficient adiabatic condition is a property of adiabatic evolution in this system, not the specific τ/p_n criteria we used.

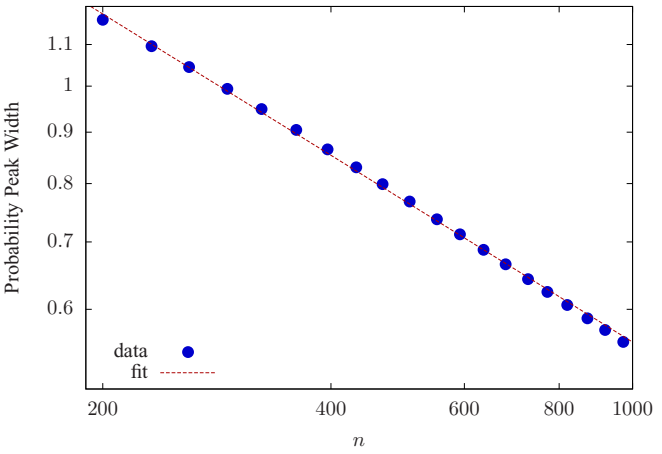

FIG. 6. The width of the nonadiabatic success probability peak in the curve of p_n vs τ is shown as a function of n. This problem tunnels through a binomial barrier of width and height $\propto n^{0.3}$. The width of the peaks in success probability are extracted by performing a Gaussian fit to the top 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.

stay within a certain range of its ground state for the entire $s \in [0,1]$ evolution.

In Fig. 5, we 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 $s \in [0,1]$ evolution. 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 cutoff stricter than 75% is used.

Figure 5 shows that the run-time relationships we observe are in fact indicative of how adiabatic evolution behaves as well. Therefore, we are led to the conclusion that for general $\mu \neq 1$, the run time $\tau \in \Theta(\sqrt{n})$ 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.

V. WIDTH OF THE NONADIABATIC SUCCESS PEAK

Our 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 $p_1 = 1$. We show that this narrowness implies that to be successful for large n, one has to be very precise in using the right running time τ.

We know that there is a critical run time 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

$$p_1 = 1 - \delta = 1 - k(\tau - \tau_c)^2, \quad \delta \ll 1, \tag{8}$$

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:

$$p_n = [1 - k(\tau - \tau_c)^2]^n \approx 1 - nk(\tau - \tau_c)^2. \tag{9}$$

If we want the probability of failure to be less than ε, we must have that

$$1 - \varepsilon < 1 - kn(\tau - \tau_c)^2 \Rightarrow |\tau - \tau_c| < (\varepsilon/kn)^{1/2}. \tag{10}$$

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.

We have run simulations using binomial-shaped barriers to get a sense of this width scaling when a barrier is present. For $\mu = 1$, we find that the narrowing of the spiked success probability p_n around the critical τ_c running time is close to $n^{-1/2}$. In Fig. 6, we show the width of the probability peak as a function of n for a binomial barrier with a height and width proportional to $n^{0.3}$. This width is extracted by looking at the nonadiabatic 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 us to extract an approximation of the width of the peak.

The widths for the barrier case in Fig. 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 nonadiabatic success probability peak.

VI. CONCLUSION

While the $\mu = 1$ case does exhibit a surprising nonadiabatic 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 [2]. We performed simple simulations of spin vector dynamics for no barrier and $\mu \neq 1$ and observed the same breaking of the nonadiabatic speedup seen in quantum annealing. Therefore, the nonadiabatic speedup does not survive in either the classical or the quantum setting.

In addition, we show that even in the $\mu = 1$ case, achieving the nonadiabatic speedup requires inverse polynomial precision in the run time. Thus, it is difficult to hit the speedup if it exists. However, even in the adiabatic setting, this problem shows a running time of $\sqrt{n}$, 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 permit analytic work. The fact that simple changes can be made to this system to eliminate an atypical nonadiabatic speedup shows the robustness of this toy model. A small change to the model is enough to bring it into 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.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under Grants No.1314969 and No.1620843.

[1] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, arXiv:quant-ph/0001106.
[2] S. Muthukrishnan, T. Albash, and D. A. Lidar, Phys. Rev. X **6**, 031010 (2016).
[3] S. Jansen, M. Ruskai, and R. Seiler, J. Math. Phys. **48**, 102111 (2007).
[4] L. Kong and E. Crosson, arXiv:1511.06991 [quant-ph].
[5] W. van Dam and L. T. Brady, presentation at Adiabatic Quantum Computing Conference, Venice, CA (2016).
[6] E. Farhi, J. Goldstone, and S. Gutmann, arXiv:quant-ph/0201031.
[7] B. W. Reichardt, in *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC '04)* (ACM Press, New York, 2004).
[8] E. Crosson and M. Deng, arXiv:1410.8484 [quant-ph].
[9] E. Crosson and A. Harrow, in *Proceedings of the FOCS' 16* (IEEE, New York, 2016), pp. 714–723.
[10] Z. Jiang, V. N. Smelyanskiy, S. V. Isakov, S. Boixo, G. Mazzola, M. Troyer, and H. Neven, Phys. Rev. A **95**, 012322 (2017).
[11] L. T. Brady and W. van Dam, Phys. Rev. A **94**, 032309 (2016).
[12] S. Boixo and R. D. Somma, Phys. Rev. A **81**, 032308 (2010).