

Random matrix model of adiabatic quantum computing

David R. Mitchell,^{1,2} Christoph Adami,^{1,4} Waynn Lue,³ and Colin P. Williams^{1,3}

¹*Jet Propulsion Laboratory MS 126-347, California Institute of Technology, Pasadena, California 91109, USA*

²*Northwestern State University, Natchitoches, Louisiana 71497*

³*Stanford University, Stanford, California 94305, USA*

⁴*Keck Graduate Institute, Claremont, California 91711, USA*

(Received 15 September 2004; published 20 May 2005)

We present an analysis of the quantum adiabatic algorithm for solving hard instances of 3-SAT (an NP-complete problem) in terms of random matrix theory (RMT). We determine the global regularity of the spectral fluctuations of the instantaneous Hamiltonians encountered during the interpolation between the starting Hamiltonians and the ones whose ground states encode the solutions to the computational problems of interest. At each interpolation point, we quantify the *degree of regularity* of the average spectral distribution via its *Brody parameter*, a measure that distinguishes regular (i.e., Poissonian) from chaotic (i.e., Wigner-type) distributions of normalized nearest-neighbor spacings. We find that for hard problem instances—i.e., those having a *critical* ratio of clauses to variables—the spectral fluctuations typically become irregular across a contiguous region of the interpolation parameter, while the spectrum is regular for easy instances. Within the hard region, RMT may be applied to obtain a mathematical model of the probability of avoided level crossings and concomitant failure rate of the adiabatic algorithm due to nonadiabatic Landau-Zener-type transitions. Our model predicts that if the interpolation is performed at a uniform rate, the average failure rate of the quantum adiabatic algorithm, when averaged over hard problem instances, scales exponentially with increasing problem size.

DOI: 10.1103/PhysRevA.71.052324

PACS number(s): 03.67.Lx, 0.5.40.–a

I. INTRODUCTION

Can quantum computers solve NP-complete problems using physical resources of time, space, and energy that are all bounded by *polynomials* in the size of the problem? Most computer scientists are skeptical of such a possibility [1]. However, recently a more physics-inspired perspective has arisen that is causing some to rethink this question. In 2002, Farhi *et al.* presented a quantum adiabatic algorithm for solving an NP-complete problem and showed via numerical simulations on a sequence of progressively larger problem instances that the running time of this algorithm appears to grow only as a polynomial in problem size [2]. By contrast, all known classical algorithms for solving NP-complete problems require a running time that scales exponentially with problem size in the worst case [3]. If the polynomial scaling of the adiabatic algorithm is correct, this would represent a monumental result for the field of quantum computing, as it would bring a host of useful but hard computations within the domain of computational tasks that can be performed exponentially faster on quantum computers than classical ones. With such extraordinary promise, it behooves us to understand the adiabatic algorithm in full detail. Unfortunately, it has proven to be exceedingly difficult to obtain analytic results on the scaling behavior of the quantum adiabatic algorithm. Instead, for the most part, researchers have relied upon numerical simulations of small problem instances, typically involving 23 variables or less [2]. By comparison, a modern-day classical algorithm for solving 3-satisfiability (3-SAT) can routinely solve instances containing several thousand variables [4,5]. It is questionable whether numerical results based on 23-variable simulations can be extrapolated reliably to 2000-variable instances. Thus a more analytic approach is needed.

In this paper we develop such an analysis based on random matrix theory (RMT) [6–8], which is a statistical description of complex quantum systems in which detailed knowledge about particle interactions is abandoned in favor of a description in terms of random interactions. Such a description is typically found to be applicable to complex Hamiltonians without fundamental symmetries and with very large phase space (high dimension). In this case the nearest-neighbor spacing (NNS) distribution assumes a universal law, and only the *statistical* properties of the levels are of interest. Such Hamiltonian systems are usually highly irregular, disordered, and chaotic.

RMT has proven to be a successful method for predicting properties of complex quantum systems that look superficially very different in terms of their energy eigenspectra. By working with the nearest-neighbor level spacing fluctuations rather than the raw eigenspectra, deep similarities between apparently different physical systems have been revealed and several hard to calculate properties, such as transition rates, have been determined. In essence, a tractable model Hamiltonian can be used to make predictions about an intractable one provided their energy spectra can be described by similar NNS distributions.

In standard RMT, Hamiltonians are drawn from an ensemble of orthogonal or unitary Gaussian matrices [the Gaussian orthogonal or unitary ensembles (GOE or GUE)], which corresponds to Hamiltonians with interactions of all possible particle ranks d (d -body interactions). In other words, every element of Hilbert space is assumed to be connected to every other by an interaction strength that is given by a random number. The level density corresponding to such a model is the celebrated “semicircle law” of Wigner [7], which predicts a power dependence between the density

of states and energy. Physical complex quantum systems (such as large atoms or nuclei), on the other hand, are better described by interactions of rank 2—that is, two-body interactions at the most. Such Hamiltonians can be obtained from the so-called “two-body random ensemble” [9,10] (or *embedded* GOE) that well describes complex nuclear and atomic level spectra, but gives rise to a Gaussian, rather than polynomial, density of states. For Hamiltonians drawn from either a GOE or an embedded GOE, the distribution of nearest-neighbor level separations is unimodal with a long tail, known as the *Wigner* distribution [6,8]. Such distributions are typical for complex systems without symmetries, which results in highly irregular energy spectra and chaotic dynamics. Conversely, Hamiltonians corresponding to physical systems subject to symmetries and conservation laws typically display *regular* energy spectra and *Poissonian* nearest-neighbor level separations. Such distributions fall off much faster than the Wigner distribution and decay monotonically. The *Brody* distribution [11] interpolates between those two distributions with a single Brody parameter q , where the limit $q=0$ corresponds to the Poissonian limit while $q=1$ gives the Wigner distribution.

A priori, the Hamiltonians arising in computational problems such as k -SAT appear to have such a special structure that they are unlikely to be described by random interaction matrices. We can assess this by characterizing the irregularity of the NNS distribution of the instantaneous Hamiltonian of the adiabatic algorithm. If we find a Brody parameter close to zero, RMT cannot be used, while a Brody parameter closer to 1 indicates that RMT can predict global properties of the Hamiltonian dynamics reliably. Note that the scaling of the level density itself (polynomial for GOE, exponential for embedded GOE) is irrelevant for this determination.

We determine the NNS distribution of the instantaneous Hamiltonians solving 3-SAT problems by generating random soluble problem instances (with exactly one solution) with a fixed ratio of clauses to variables, determining for each the eigenvalue distribution, from which the fluctuations can be obtained. In short, the results reveal a systematic change in the spectral regularity of the instantaneous Hamiltonians during the course of the adiabatic algorithm. In the initial phase of the interpolation for especially hard problem instances, the statistical NNS fluctuations conform to a regular, Poisson-type distribution. Later in the interpolation, the fluctuations conform to an irregular Wigner-type distribution instead. We also find that irregular spectra only occur for computationally *hard* problem instances.

In this paper, we predict the scaling of the failure rate of the adiabatic algorithm for a fixed ratio of clauses to variables at a given point in the interpolation process, as larger and larger problem instances are considered. The adiabatic algorithm fails when the system spontaneously transitions from its ground state into any excited state. If we make the conservative assumption that the only source of nonadiabatic transitions are of the Landau-Zener (LZ) type [12]—i.e., localized transitions between adjacent levels at avoided crossings where the energy levels locally assume the geometry of hyperbolae (see Fig. 1)—then we can obtain a lower bound on the transition probability using RMT. Additional failure modes can only make the failure rate of the algorithm worse.

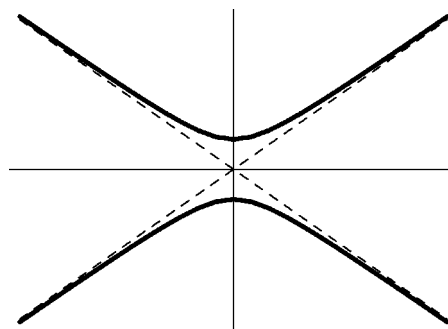


FIG. 1. Avoided level crossings can occur when adjacent levels assume the geometry of converging hyperbolae. Both the slope difference Δm and minimum gap ΔE affect the probability of a transition occurring.

With these assumptions and the use of RMT, we can then determine the transition rate from the ground state averaged over an ensemble of problem instances having the same ratio of clauses to variables. There are two model-dependent quantities in this result: the average ground-state level spacing and the typical size of LZ asymptotic slopes. Ultimately, these quantities are related to the parameters characterizing problem instances of the type being solved—i.e., the ratio of clauses to variables. Hence the transition rate, at a given point in the interpolation, is related to the difficulty of the problem instances at that point. The rest of the paper is organized as follows. Section II describes the quantum adiabatic algorithm and the Landau-Zener transition probability. Section III summarizes prior research on the scaling properties of the algorithm—i.e., whether the adiabatic criterion can be met if the interpolation is completed in polynomial time. Section IV introduces the concepts of random matrix theory and Landau-Zener transitions needed for our analysis. Section V reports on our numerical experiments and the gap fluctuation phenomena they reveal. Section VI uses the phenomena to justify a random matrix analysis of the adiabatic algorithm and describes the implications on the scaling of the quantum adiabatic algorithm. We discuss the distribution of gap energies in the Appendix.

II. ADIABATIC ALGORITHM

The idea behind the quantum adiabatic algorithm is as follows. If a quantum system is prepared in the ground state of a time-independent Hamiltonian H_0 and if we then cause the Hamiltonian to change from H_0 to a final form H_1 in T steps, e.g., by driving it linearly,

$$H\left(\frac{t}{T}\right) = \left(1 - \frac{t}{T}\right)H_0 + \frac{t}{T}H_1, \quad (1)$$

then the adiabatic theorem of quantum mechanics [13] guarantees that the system will remain in the ground state of the instantaneous Hamiltonians $H(t)$, provided the change is made sufficiently slowly—i.e., adiabatically. Thus, if the final Hamiltonian can be made to encode a computational problem such that the ground state of H_1 corresponds to the solution to this problem, then the natural quantum mechani-

cal evolution of the system under the slowly changing Hamiltonian $H(t)$ would carry our initial state into a final state corresponding to the solution. A final-state measurement would then reveal the solution. The key question is how quickly can one drive the interpolation between the initial and final Hamiltonians while keeping the system in the ground state of the instantaneous Hamiltonians passed through. If the shortest feasible interpolation time scales polynomially with increasing problem size, the quantum adiabatic algorithm would be deemed “efficient”; otherwise, it would be deemed “inefficient” [14]. An alternative way of stating this is to ask under what conditions the passage from H_0 and H_1 can be performed adiabatically [13]. If the minimum eigenvalue gap between the ground state E_0 and first excited state E_1 of the instantaneous Hamiltonians is given by g_{\min} , where

$$g_{\min} = \min_{0 \leq t \leq T} [E_1(t) - E_0(t)], \quad (2)$$

and the matrix element between the corresponding pair of eigenstates is

$$\left\langle \frac{dH}{dt} \right\rangle_{1,0} = \left\langle E_1; t \left| \frac{dH}{dt} \right| E_0; t \right\rangle, \quad (3)$$

then the adiabatic theorem asserts that the final state will be very close to the ground state of $H_1(T)$, i.e.,

$$|\langle E_0; T | \psi(T) \rangle|^2 \geq 1 - \epsilon^2, \quad (4)$$

provided that

$$\frac{\left| \left\langle \frac{dH}{dt} \right\rangle_{1,0} \right|}{g_{\min}^2} \leq \epsilon, \quad (5)$$

where $\epsilon \ll 1$. If this criterion is met, we can be sure the system will evolve into the desired state. But it is not immediately clear how quickly we can interpolate between H_0 and H_1 while ensuring this adiabaticity criterion is not violated.

III. PRIOR ANALYTIC RESULTS

Prior analytic studies of the adiabatic quantum algorithm have yielded mixed results. Farhi *et al.* analyze several models in which the gap behavior can be computed analytically and be shown to decrease polynomially in the problem size [15]. However, they caution that the particular problems they studied have a high degree of structure that would also make them easy to solve classically. Nevertheless, the results show that the adiabatic algorithm scales favorably at least on easy problems. To show the gap is, at the very least, nonvanishing, Ruskai [16] provides a clever proof that the ground state of the instantaneous Hamiltonian must be unique. However, as she points out, this tells us nothing about the magnitude of the gap and how it scales with problem size. A less encouraging result was obtained by van Dam, Mosca, and Vazirani [17], who were able to construct a family of minimization problems for which they could prove an exponential lower bound on the running time of the (original) adiabatic algorithm on these problems. A subsequent paper by Farhi *et al.*

[18] challenged the inevitability of such results by arguing that they might be circumvented by choosing a different interpolation path between the initial Hamiltonian and the one encoding the problem to be solved. To date, the most sophisticated analysis of the running time of the adiabatic algorithm on NP-complete problems was provided by Roland and Cerf [19]. They found that by nesting one quantum adiabatic search algorithm within another, one can solve NP-complete problems more efficiently than naive use of an adiabatic version of Grover’s algorithm. Nevertheless, the run-time scaling is still exponential in problem size, albeit better than what is possible classically.

IV. APPLICABILITY OF RANDOM MATRIX THEORY

Random matrix theory is a statistical approach to Hamiltonian systems that are otherwise analytically intractable. For example, RMT focuses on universal model-independent properties of the system under study, such as the distribution of the spectral *fluctuations*. Many superficially different physical systems are found to have distributions of spectral fluctuations that fall into just a handful of categories. Such a characterization of spectra originated in the context of nuclear physics [20] and was applied later to complex many-body systems and quantum systems having a chaotic classical analog [21]. Once the distribution of spectral fluctuations of a physical system has been identified (and deemed to be irregular), RMT can be applied to make predictions about properties of interest, such as transition rates between different levels. For example, the problem of estimating transition rates has been examined in the context of nuclear dissipation, and the use of the LZ transition as a mechanism for nuclear dissipation was suggested originally by Hill and Wheeler [22]. The combination of the LZ transition probability with the RMT statistical approach was examined by Wilkinson [23,24], whose results we apply to the current problem. Although the LZ assumption is reasonable for adiabatic systems [25], the mechanism for dissipation in complex spectra continues to be investigated in the context of RMT, with more recent approaches using a non-LZ, propagator approach [26].

To determine the applicability of RMT, the regularity of the entire spectrum at each adiabatic interpolation point is measured by the Brody parameter (defined below). The nearest-neighbor spacing distribution is the most common measure of spectral regularity in quantum systems. This measure is quantified by the Brody parameter q that interpolates between a regular Poisson spectrum ($q=0$) and an irregular (quantum chaotic) Wigner distribution ($q=1$) [11]. A renormalized spectrum with a Brody parameter q is characterized by the following NNS probability distribution for spacing level spacing δ :

$$p_q(\delta) = (1+q)\beta\delta^q \exp(-\beta\delta^{1+q}), \quad \beta = \left[\Gamma\left(\frac{2+q}{1+q}\right) \right]^{1+q}. \quad (6)$$

The form of this distribution, for different values of the Brody parameter, is shown in Fig. 2.

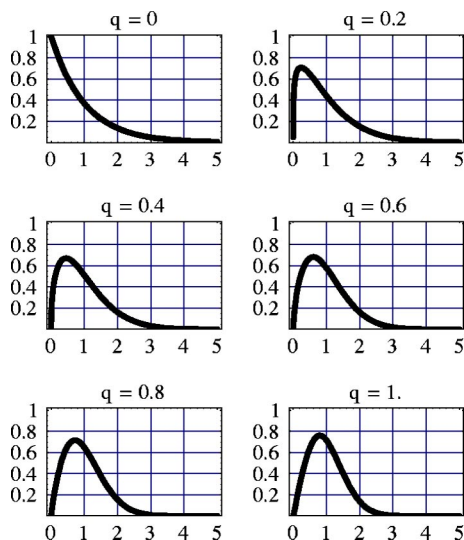


FIG. 2. (Color online) Nearest-neighbor probability distribution $p_q(\delta)$ for eigenvalue spacings δ , in unfolded spectra for values of the Brody parameter $q \in [0, 1]$. When $q=0$, the distribution resembles an exponential distribution, but changes to unimodal as q is increased.

Irregular RMT spectra are characterized by an abundance of avoided level crossings and a lack of level degeneracies. We now proceed to determine whether or not the distribution of spectral fluctuations is anywhere irregular during the interpolation process.

V. SPECTRAL FLUCTUATION EXPERIMENTS

Our first task is to determine an appropriate ensemble of random soluble 3-SAT problem instances to use. We would like to use computationally hard problem instances, because we are most interested in assessing the scaling of the failure rate of the adiabatic algorithm on hard problems. Hard, in this sense, is a relative term. When solving random instances of soluble 3-SAT problems having n variables and m clauses, typically the hardest instances are encountered at a critical value of the clause to variable ratio m/n . For the 3-SAT problem as $n \rightarrow \infty$, the hardest instances are clustered around the ratio $m/n \approx 4.2$. However, small problem instances (having, say, $n < 40$), typically, have a somewhat displaced transition point. Figure 3 shows the mean computational cost of solving 3-SAT problems containing from $n=8$ to $n=50$ variables using either the Davis-Putnam (DP) algorithm [27] or the GSAT algorithm [4]. Regardless of the algorithm used, an easy-hard-easy pattern is apparent when the number of clauses is increased at fixed number of variables. In the limit of infinite problem size, the easy and hard instances are separated by a phase transition (see, e.g., [28]). But as can be seen from Fig. 3, the location of the phase transition point is extremely ill defined for problem instances having $n < 20$. Thus inferring any reliable cost scaling by extrapolating costs from such small instances would be exceedingly unreliable. As the simulation of the adiabatic algorithm solving an n -variable 3-SAT problem involves $(2^n \times 2^n)$ -dimensional matrices, we cannot simulate very large cases. Hence, we

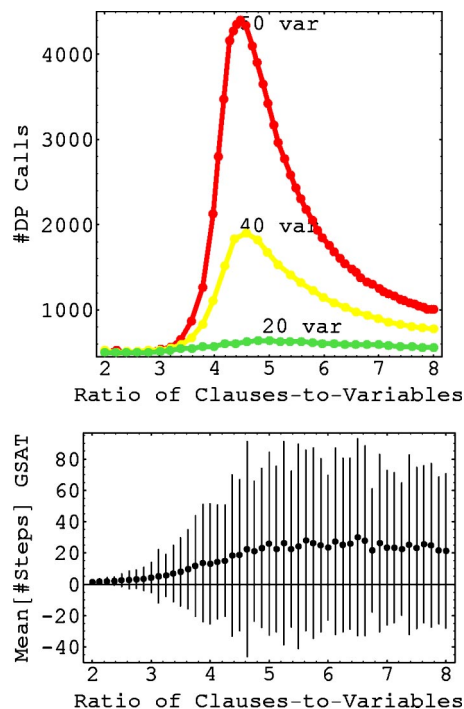


FIG. 3. (Color online) Upper plot shows the mean computational cost of solving $n=20, 40$, and 50 variable 3-SAT problems using the Davis-Putnam algorithm (the data on the scaling of this algorithm are due to Bart Selman; see [28]). The data show that for small problems (i.e., $n < 50$) the phase transition region is smeared out. The lower plot shows the mean computational cost of solving 1000 random instances of guaranteed soluble 3-SAT problems having $n=8$ variables and $m=8-80$ clauses, using the GSAT algorithm. Again, for such small-sized problems the region of hard problem instances (relative to other instances) is quite spread out, but roughly centered on $m/n=6$, rather than the asymptotic value of 4.2.

begin by first determining the *actual* location of the hardest problems for 3-SAT problems involving a more tractable $n=8$ variables, rather than relying on the asymptotically known result whose applicability is suspect at small n . Specifically, we generated 72 000 random 3-SAT problem instances all having $n=8$ variables, but with the number of clauses ranging from $m=8$ to 80, corresponding to $1 \leq m/n \leq 10$, and solved them using the GSAT algorithm. Each data point was computed from an average of 1000 problem instances. The results are shown in the lower portion of Fig. 3. Although the difference between easy cases and hard cases is not as pronounced as it is for much larger n , nevertheless the data suggest that problem instances centered around $m/n=6$ will yield relatively hard cases [29]. We will use such problem instances to create the ensemble we need in our numerical studies of the distribution of spectral fluctuations of the instantaneous Hamiltonians encountered during the interpolation phase of the adiabatic algorithm.

Next we turn our attention to the global spectral properties of the instantaneous Hamiltonians encountered in the quantum adiabatic algorithm for easy and hard problem instances. Specifically, we obtain the NNS distribution of the adiabatic Hamiltonian $H(s)$ after renormalization of the spec-

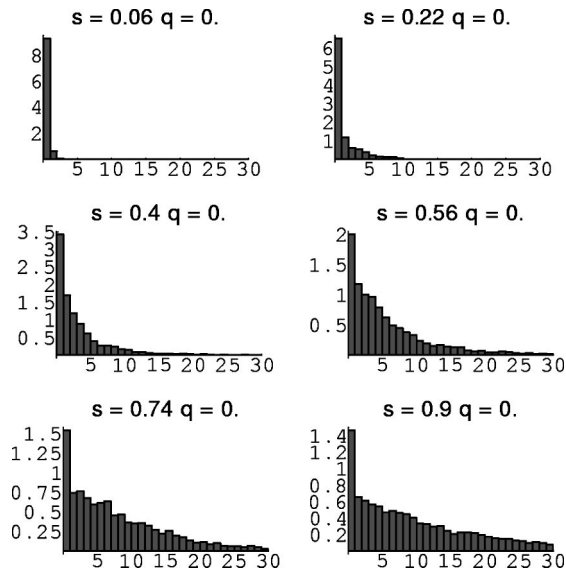


FIG. 4. Distribution of eigenvalue gaps in the unfolded spectra of the instantaneous Hamiltonians during the course of the adiabatic algorithm solving easy instances of 3-SAT having $n=8$ variables and $m=4$ clauses, as the interpolation parameter varies from 0 to 1. As s increases, the Brody parameter remains zero, implying that RMT is not applicable to easy problem instances. The horizontal axes are nearest-neighbor eigenvalue spacings (NNS), expressed in tenths of the mean local spacing.

trum to unit average local level density [6,8] for an ensemble of easy cases and for an ensemble of hard cases. The regularity of the spectrum is determined at each point in the adiabatic evolution by fitting it to Eq. (6) and obtaining the Brody parameter $q(s)$ of the NNS distribution at that point. We begin by determining the spectral distribution of an ensemble of easy problems as the interpolation parameter in the quantum adiabatic algorithm ranges from $s=0$ to $s=1$, for instances of soluble 3-SAT having $n=8$ variables and $m=4$ clauses—i.e., problems for which $m/n=0.5$. Each histogram in Fig. 4 is based on the spectral behavior of an ensemble of 20 problem instances having fixed values of n and m . For the easy problems, the spectrum of each instantaneous Hamiltonian conforms to a Poisson (regular) spectral fluctuation distribution and small Brody parameter $q=0$. Hence, we conclude that RMT is not applicable in the easy region—i.e., for $m/n < 5$ ($n=8$).

In contrast, Fig. 5 shows the eigenvalue gap fluctuations of instantaneous Hamiltonians induced from “hard” instances of soluble 3-SAT with $n=8$ variables and $m=48$ clauses. Again, our ensemble averages over 20 instances of a fixed n and m . For $m/n=6$ (hard problems) we observe Poisson behavior for $s < 0.5$, but for $s > 0.5$ the spectra become increasingly irregular and the Brody parameter becomes significant. In other words, the instantaneous Hamiltonians induced by random, hard 3-SAT instances appear to have a qualitatively different spectrum from those of easy problems of the same size. In particular, at a certain point in the interpolation process between the initial and final Hamiltonians, the spectrum becomes irregular and the NNS distribution resembles a Wigner distribution with a relatively large Brody

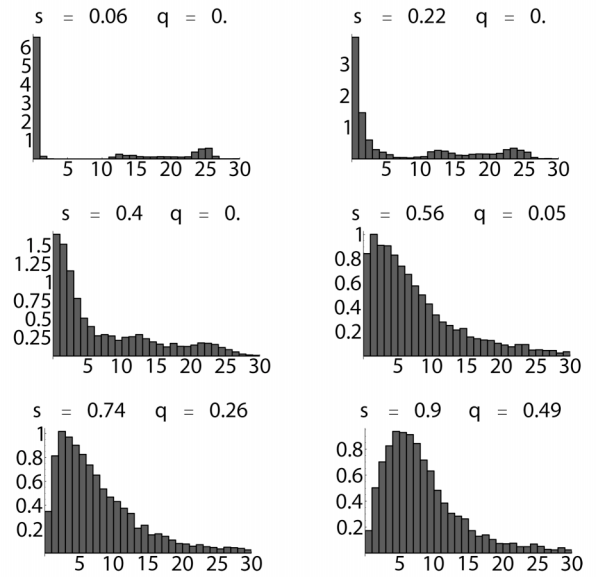


FIG. 5. Distribution of eigenvalue gaps in the unfolded spectra of the instantaneous Hamiltonians during the course of the adiabatic algorithm solving hard instances of 3-SAT having $n=8$ variables and $m=48$ clauses, as the interpolation parameter varies from 0 to 1. As s increases, the Brody parameter varies from 0 to 0.49 and back to 0. The horizontal axes are nearest-neighbor eigenvalue spacings (NNS), expressed in tenths of the mean local spacing.

parameter. Here, RMT can be applied to estimate transition rates between levels.

Finally, one can repeat these experiments for problems having $n=8$ variables at a clause to variable ratio of $m/n=9$ (easy region again—data not shown). Here, we observe results similar to $m/n=6$, but with a slightly smaller maximum Brody parameter. We attribute this to the fact that $m/n=9$ is easier to solve than $m/n=6$, although for $n=8$ variables, this difference is slight.

Figure 6 summarizes the Brody parameter as a function of the interpolation parameter s for easy and hard problems with $n=8$ variables and from 4 to 72 clauses. The critical question in deciding if the quantum adiabatic algorithm can be completed in polynomial time is whether such a (fast) interpolation would induce level transitions. If they do occur, the system will not reside in the ground state of H_1 upon completion of the adiabatic path, and the algorithm will have failed to find the solution. In the next section, we compute the probability that level transitions occur during the course of the interpolation from H_0 to H_1 using RMT, for problems of a given degree of difficulty.

VI. RANDOM MATRIX ANALYSIS OF QUANTUM ADIABATIC ALGORITHM

Previous analyses of the adiabatic algorithm (summarized in Sec. III) placed the greatest significance on the scaling of the $E_1 - E_0$ gap with increasing problem size. However, as Eq. (5) shows, it is the ratio of the matrix element $\langle dH/dt \rangle_{1,0}$ to the square of the minimum gap that determines whether

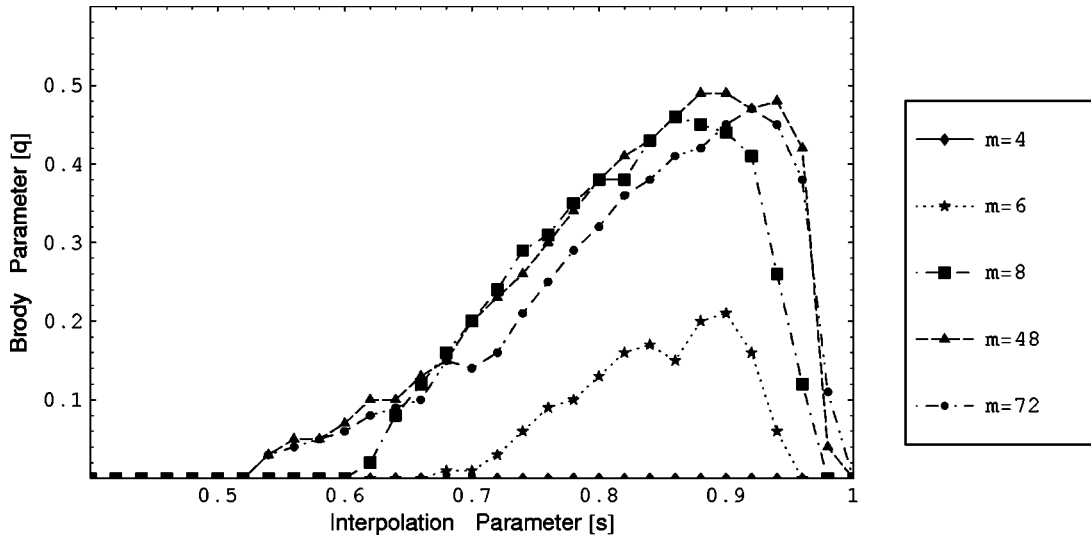


FIG. 6. Brody parameter as a function of the interpolation parameter s , for easy and hard instances of 3-SAT with $n=8$ variables from $m=4$ to 72 clauses. Note that the largest values of the Brody parameter coincides with the hardest problems. For soluble 3-SAT problems of size $n=8$, these are found around $m/n=6$ rather than at $m/n=4.2$, which is the asymptotically valid transition point for 3-SAT as $n \rightarrow \infty$.

the adiabatic theorem applies. Here, we calculate instead the *probability of a transition* from the ground state, which is a reliable proxy for the failure rate of the adiabatic algorithm. In regions where the Brody parameter is significant, we assume that any nonadiabatic transitions are of Landau-Zener type—i.e., confined to adjacent levels at avoided crossings where the energy levels assume the geometry of convergent hyperbolas [12] (see Fig. 1). When approached in this manner, the probability for a single transition, anywhere in the spectrum, can be parametrized by the minimum gap size ΔE , the difference in the asymptotic slopes Δm , and the rate of change of the adiabatic evolution parameter $ds/dt = \dot{s}$. Specifically, the transition probability is [12]

$$P = e^{-2\pi\gamma}, \quad \gamma = \frac{1}{4\hbar} \frac{\Delta E^2}{|\Delta m| \dot{s}}. \quad (7)$$

Typical values for the parameters ΔE and Δm will vary with the difficulty of the problem instance being solved (reflected by the clause to variable ratio m/n) as well as with the interpolation parameter.

In order to exploit this transition probability to predict the rate of transition from the ground state, we need to verify that the $E_0 - E_1$ gap fluctuations follow the same distribution as those in the body of the spectrum. In the Appendix, we show that the distribution of gap fluctuations is characterized by a Brody parameter comparable to the typical fluctuations in the body of the spectrum, establishing this point. To incorporate the model dependence (nonuniversality) of the problem, the average local level density of the spectral region of interest (here the levels) is included. Thus, the RMT-LZ calculation results in an ensemble-averaged transition rate with the LZ process as the principle transition mechanism.

The transition rate from the i th eigenstate of the instantaneous Hamiltonian with energy E_i , dP_i/dt , can be written in terms of the second moment of the occupation probability distribution $P_i(t)$ [23] because it essentially is the rate of

diffusion of the occupation probability. As discussed in the Introduction, the appropriate RMT ensemble to use for physical systems with exponential densities of states is the embedded GOE, consisting of real, orthogonal matrices having Gaussian-distributed random matrix elements with at most two-body interactions. Since in the calculation of transition rates in RMT only the fluctuation properties of the spectrum (rather than the density of states proper) enter [23], we can safely substitute a GOE to obtain the probability to transition from the ground state

$$\frac{dP_0}{dt} \propto \sigma^{3/2} \rho^2 |\dot{s}|^{3/2}, \quad (8)$$

where ρ is the level density of $E_0 - E_1$ levels averaged over an ensemble of problem instances having the same clause to variable ratio and σ is the typical size of the asymptotic slopes of the LZ avoided level crossings in the region of interest. Indeed, as shown in Ref. [24], embedded GOE's give rise to NNS distributions very similar to those arising in a GOE, except that the maximal Brody parameter is limited to $q \approx 0.8$. Note that this is also implied by a *numerical* calculation [25] of the dissipation rate of probability distributions of physical instantaneous Hamiltonians (i.e., those drawn from and embedded GOE), which agree with the scaling in Eq. (8).

We emphasize that Eq. (8) expresses the transition rate from the ground state in terms of the average $E_0 - E_1$ level density, rather than the minimum gap (i.e., the maximum level density). The average level density is the more relevant parameter for assessing the typical behavior of the adiabatic algorithm. As the adiabatic algorithm fails if the system transitions from the ground state, Eq. (8) can be interpreted as the average failure rate of the adiabatic algorithm when driven with uniform interpolation velocity ds/dt .

Finally, we need to determine the scaling of the transition rate from the ground state (i.e., the failure rate) with increas-

ing problem size. We do this in two stages: first we show how the transition rate must scale with ρ in order to keep the transition rate bounded then, we show how ρ scales with problem size. Although the values of ρ and σ are model-dependent quantities, we note that σ is the characteristic slope in the energy-parameter space (i.e., $\sigma \approx dE/ds$). To make explicit the dependence on level density, we work with unfolded energies ($E_i = \Delta \times \epsilon_i$) where Δ is the mean-level spacing ($\Delta = 1/\rho$) and ϵ_i are the unfolded energy levels having mean-level spacing 1. Under this transformation, $\sigma^{3/2} \rightarrow (\Delta \times \tilde{\sigma})^{3/2}$, and we therefore write

$$\frac{dP}{dt} = \propto \tilde{\sigma}^{3/2} \rho^{1/2} |s|^{3/2}. \quad (9)$$

In general, while the unfolded level slope $\tilde{\sigma}$ and the average level density ρ are model-dependent quantities [30], Eq. (9) nevertheless exhibits the explicit dependence of the transition rate on ρ . We can establish a lower bound on the interpolation time required to evolve the system through the irregular region by noting first that

$$\frac{dP}{dt}(s, \rho) \geq \frac{dP}{dt}(s, \rho_{\min}), \quad (10)$$

where we defined $\rho_{\min} = \min_s \rho(s)$ and where the minimization is carried out over only those values of s that give rise to an irregular spectrum. It is important to note that ρ_{\min} is not the minimum of a particular problem instance; it is an average level density at a particular value of s , whose value is constant for a given problem parameter set. The failure probability during an evolution path of length T is then bounded by

$$P \geq \text{const} \times \tilde{\sigma}^{3/2} \Delta s^{3/2} \left[\frac{\rho_{\min}}{T} \right]^{1/2}. \quad (11)$$

To ensure a given transition probability over a given range Δs in an irregular region, the interpolation time T must scale as ρ_{\min} . If we can now estimate how ρ_{\min} (an average quantity) scales with increasing problem size, we can estimate how the time needed to complete a particular part of the adiabatic algorithm must scale in order to keep the transition rate from the ground state small. This part is precisely the region where one would need to go most slowly to avoid an unwanted transition from the ground state. What can we say about the scaling of ρ_{\min} with problem size? As pointed out by Ruskai [16], in regions characterized by a lack of level degeneracies (such as irregular spectral regions) the interpolating Hamiltonian $H(s)$ must fit 2^n eigenvalues into a range that is polynomial in n , and consequently the level density must scale *exponentially* with problem size. Our numerical simulations provide strong evidence for the existence of such irregular spectral regions by the fact that we obtain significant values of the Brody parameter for distributions averaged over instances that have a sizable degree of difficulty (as measured by their clause to variable ratio). In other words, difficult problem instances display spectra characterized by a large Brody parameter and few level degeneracies, and therefore show a lack of symmetry. Thus, with the final assumption that nondegenerate spectra (irregular spectra having a large Brody parameter) have a minimum average level den-

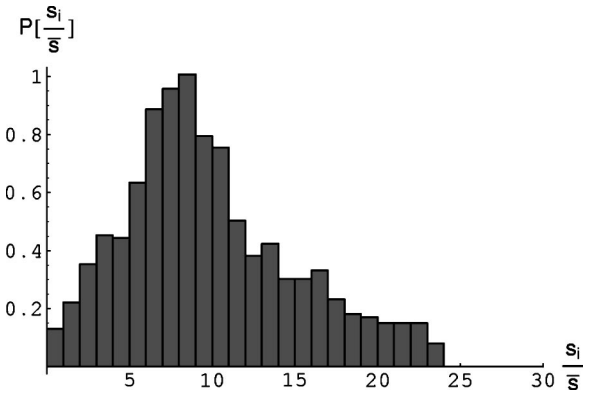


FIG. 7. NNS distribution of $E_0 - E_1$ gaps for 1000 random problem instances in the hard region for problems of size $n=8$. The horizontal axis is the eigenvalue spacing in tenths of the mean-level spacing.

sity that scales exponentially with problem size, we find the transition rate (11) out of the ground state (the algorithmic failure rate) and thus that the adiabatic evolution time cannot scale polynomially with problem size.

The information regarding the parameter regions having significant transition probability may be used in a variable-speed approach to the adiabatic quantum algorithm in a manner similar to [19]. In this way, a speed-up over classical algorithms for solving NP-complete problems still appears possible. We wish to emphasize that Eq. (11) and the entire RMT analysis hold only in those regions of the interpolation process where we find irregular spectra—i.e., nontrivial Brody parameters. In particular, the formula does not apply to those regions of the interpolation with a small Brody parameter, which typically occur in the very early and very late stages of a hard problem or at all stages of an easy problem. Therefore, our formula is not inconsistent with previous results (e.g., those mentioned in Sec. III, which find that the adiabatic algorithm scales favorably on easy problems), because in such cases Eq. (11) is invalid.

VII. CONCLUSION

Our approach to the analysis of adiabatic quantum computation is statistical in nature and relies on the applicability of random matrix theory to large Hamiltonians. With such an approach, we examine the behavior of an averaged failure rate and interpolation time when averaged over an ensemble of problem instances of a given difficulty. We find a spectral transition from orderly to disorderly as problem difficulty is increased. Our analysis of these results suggests using the Landau-Zener transition probability in a statistical RMT approach. For those regions where RMT applies, degeneracies are lacking and we show that the average failure rate and average interpolation time do not scale polynomially with problem size. Nevertheless, the quantum adiabatic algorithm appears to be more efficient than any known classical algorithm for solving NP-complete problems, with a speed-up commensurate with Grover's search algorithm. Moreover, the insights we gain from the spectral fluctuation analysis

may be used to design a *variable* interpolation rate, which has the potential to improve the adiabatic algorithm over its performance using a uniform interpolation rate.

Finally, we believe the connection shown here between difficult computational problems and physical spectral irregularity is a powerful path towards understanding and classifying physical approaches to computation.

ACKNOWLEDGMENT

The research described in this paper was performed at the Jet Propulsion Laboratory (JPL), California Institute of Technology, under contract with the National Aeronautics and Space Administration (NASA). We thank the JPL Supercomputing Project for the use of the Cray supercomputer used in computations. D.M. received support through from NASA. C.P.W. thanks the Advanced Research and Development Activity and the National Security Agency for support. C.A. is

supported by the Army Research Office's Grant No. DAAD19-03-1-0207.

APPENDIX

Figure 7 shows the ground-state level distribution of $E_0 - E_1$ gaps for 1000 random problem instances in the hard region for problems of size $n=8$. The horizontal axis is the eigenvalue spacing $s_i = E_0 - E_1$ normalized by the mean-level spacing. The distribution shows that spectral irregularity definitely extends to the lowest-lying levels for hard instances of 3-SAT problems, implying that RMT is applicable to hard instances of 3-SAT across the entire spectrum of levels, including the $E_0 - E_1$ gap, which is the relevant gap for assessing the scaling of the quantum adiabatic algorithm. Therefore, our use of RMT to characterize the transition rate out of the ground state and, hence, to estimate the cost scaling of the quantum adiabatic algorithm, is valid.

-
- [1] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani, *SIAM J. Comput.* **26**, 1510 (1997).
 - [2] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, *Science* **292**, 472 (2002).
 - [3] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness* (Freeman, San Francisco, 1979).
 - [4] B. Selman, H. J. Levesque, and D. G. Mitchell, in *Proceedings of the Tenth National Conference on Artificial Intelligence* (AAAI Press/MIT Press, Cambridge, MA, 1992), pp. 440–444.
 - [5] I. Lynce and J. Marques-Silva, in *Proceedings of the 15th European Conference on Artificial Intelligence*, edited by F.van Harmelen (IOS Press, Amsterdam, 2002).
 - [6] M. L. Mehta, *Random Matrices* (Academic Press, San Diego, 1991).
 - [7] E. P. Wigner, *SIAM Rev.* **9**, 1 (1967).
 - [8] T. A. Brody, J. Flores, J. B. French, P. A. Mello, A. Pandey, and S. S. M. Wong, *Rev. Mod. Phys.* **53**, 385 (1981).
 - [9] K. K. Mon and J. B. French, *Ann. Phys. (N.Y.)* **95**, 90 (1975).
 - [10] L. Benet and H. A. Weidenmüller, *J. Phys. A* **36**, 3569 (2003).
 - [11] T. A. Brody, *Lett. Nuovo Cimento Soc. Ital. Fis.* **7**, 482 (1973).
 - [12] G. Zener, *Proc. R. Soc. London, Ser. A* **137**, 696 (1932).
 - [13] L. Schiff, *Quantum Mechanics* (McGraw-Hill, New York, 1955).
 - [14] Note that, even if the quantum adiabatic algorithm proves to be inefficient by this measure, it may, nevertheless, be superior to any known classical algorithm.
 - [15] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, e-print quant-ph/0001106.
 - [16] M. B. Ruskai, *Contemp. Math.* **307**, 265 (2002).
 - [17] W. van Dam, M. Mosca, and U. Vazirani, in *Proceedings of the 42nd Annual Symposium on the Foundations of Computer Science* (IEEE Computer Science Press, 2001), pp. 279–287.
 - [18] E. Farhi, J. Goldstone, and S. Gutmann, e-print quant-ph/0208135.
 - [19] J. Roland and N. J. Cerf, *Phys. Rev. A* **68**, 062312 (2003).
 - [20] F. J. Dyson, in *Statistical Theories of Spectra: Fluctuations*, edited by C. E. Porter (Academic Press, New York, 1965).
 - [21] O. Bohigas, in *Chaos and Quantum Physics*, edited by M. Giannoni, A. Voros, and J. Zinn-Justin (North-Holland, New York, 1991), pp. 87–199.
 - [22] D. L. Hill and J. A. Wheeler, *Phys. Rev.* **89**, 1102 (1952).
 - [23] M. Wilkinson, *J. Phys. A* **21**, 4021 (1988).
 - [24] M. Wilkinson, *Phys. Rev. A* **41**, 4645 (1990).
 - [25] M. J. Sanchez, E. Vergini, and D. A. Wisniacki, *Phys. Rev. E* **54**, 4812 (1996).
 - [26] A. Bulgac, G. DoDang, and D. Kusnezov, *Phys. Rev. E* **54**, 3468 (1996).
 - [27] M. Davis, G. Logemann, and D. Loveland, *Commun. ACM* **5**, 394 (1962).
 - [28] S. Kirkpatrick and B. Selman, *Science* **264**, 1297 (1994).
 - [29] It is not our intention to use simulation results to infer the cost scaling of the adiabatic algorithm. Rather, we use simulations merely to infer the applicability of a particular mathematical technique—random matrix theory. Hence, while using problem instances with variables is not ideal, it is unlikely that we gain much additional insight by using (say) instances having $n=20$ rather than $n=8$, because (as Fig. 3 shows) even at $n=20$ we would still be grappling with an anomalously broad phase transition phenomenon.
 - [30] D. Mitchell, Y. Alhassid, and D. Kusnezov, *Phys. Lett. A* **215**, 21 (1996).