Robustness of adiabatic quantum computation

Andrew M. Childs* and Edward Farhi†
Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

John Preskill‡
Institute for Quantum Information, California Institute of Technology, Pasadena, California 91125
(Received 23 August 2001; published 14 December 2001)

We study the fault tolerance of quantum computation by adiabatic evolution, a quantum algorithm for solving various combinatorial search problems. We describe an inherent robustness of adiabatic computation against two kinds of errors, unitary control errors and decoherence, and we study this robustness using numerical simulations of the algorithm.

DOI: 10.1103/PhysRevA.65.012322 PACS number(s): 03.67.Lx, 03.65.Yz

I. INTRODUCTION

The method of quantum computation by adiabatic evolution has been proposed as a general way of solving combinatorial search problems on a quantum computer [1]. Whereas a conventional quantum algorithm is implemented as a sequence of discrete unitary transformations that form a quantum circuit involving many energy levels of the computer, the adiabatic algorithm works by keeping the state of the quantum computer close to the instantaneous ground state of a Hamiltonian that varies continuously in time. Therefore, an imperfect quantum computer implementing a conventional quantum algorithm might experience different sorts of errors than an imperfect adiabatic quantum computer. In fact, we claim that an adiabatic quantum computer has an inherent robustness against errors that might enhance the usefulness of the adiabatic approach.

The adiabatic algorithm works by applying a time-dependent Hamiltonian that interpolates smoothly from an initial Hamiltonian whose ground state is easily prepared to a final Hamiltonian whose ground state encodes the solution to the problem. If the Hamiltonian varies sufficiently slowly, then the quantum adiabatic theorem guarantees that the final state of the quantum computer will be close to the ground state of the final Hamiltonian, so a measurement of the final state will yield a solution of the problem with high probability. This method will surely succeed if the Hamiltonian changes slowly. But how slow is slow enough?

Unfortunately, this question has proved difficult to analyze in general. Some numerical evidence suggests the possibility that the adiabatic algorithm might efficiently solve computationally interesting instances of hard combinatorial search problems, outperforming classical methods [1–4]. Whether the adiabatic algorithm provides a definite speedup over classical methods remains an interesting open question. As we will discuss in Sec. II, the time required by the algorithm for a particular instance can be related to the minimum gap Δ between the instantaneous ground state and the rest of the spectrum. Roughly speaking, the required time goes like Δ^{-2}. Thus, if Δ^{-2} increases only polynomially with the size of the problem, then so does the time required to run the algorithm. However, determining Δ has not been possible in general.

Our objective in this paper is not to explore the computational power of the adiabatic model, but rather to investigate its intrinsic *fault tolerance*. Since quantum computers are far more susceptible to making errors than classical digital computers, fault tolerant protocols will be necessary for the operation of large-scale quantum computers. General procedures have been developed that allow any quantum algorithm to be implemented fault tolerantly on a universal quantum computer [5], but these involve a substantial computational overhead. Therefore, it would be highly advantageous to weave fault tolerance into the design of our quantum hardware.

We therefore will regard adiabatic quantum computation not as a convenient language for describing a class of quantum circuits, but as a proposed physical implementation of quantum information processing. We do not cast the algorithm into the conventional quantum computing paradigm by approximating it as a sequence of discrete unitary transformations acting on a few qubits at a time. Instead, suppose we can design a physical device that implements the required time-dependent Hamiltonian with reasonable accuracy. We then imagine implementing the algorithm by slowly changing the parameters that control the physical Hamiltonian. How well does such a quantum computer resist decoherence, and how well does it perform if the algorithm is imperfectly implemented?

Regarding resistance to decoherence, we can make a few simple observations. The phase of the ground state has no effect on the efficacy of the algorithm, and therefore dephasing in the energy eigenstate basis is presumably harmless. Only the interactions with the environment that induce transitions between eigenstates of the Hamiltonian might cause trouble. In principle, these may be well controlled by running the algorithm at a temperature that is small compared to the minimum gap Δ. (We use units in which Boltzmann's constant $k_B=1$, so that temperature has units of energy.) If Δ decreases slowly as the size of the problem increases, then the resources required to run at a sufficiently low tempera-

*Electronic address: amchilds@mit.edu
†Electronic address: farhi@mit.edu
‡Electronic address: preskill@theory.caltech.edu

ture may be reasonable. Since the adiabatic method is only efficient if Δ is not too small, we conclude that whenever the method works on a perfectly functioning quantum computer, it is robust against decoherence.

In addition to environmental decoherence, we must also consider the consequences of imperfect implementation. Our chosen algorithm may call for the time-dependent Hamiltonian $H(t)$, but when we run the algorithm, the actual Hamiltonian will be $H(t)+K(t)$, where $K(t)$ is an "error." An interesting feature of adiabatic quantum computation is that $K(t)$ need not remain small during the evolution in order for the algorithm to work effectively. A reasonably large excursion away from the intended Hamiltonian is acceptable, as long as $K(t)$ is slowly varying and has initial and final values that are not too large. A very rapidly fluctuating $K(t)$ may also be acceptable, if the characteristic frequency of the fluctuations is large compared to the energy scale of $H(t)$.

In this paper, we use numerical simulations to investigate the sensitivity of an adiabatic computer to decohering transitions and to a certain class of unitary perturbations induced by a Hamiltonian $K(t)$. The results are consistent with the idea that the algorithm remains robust as long as the temperature of the environment is not too high and $K(t)$ varies either sufficiently slowly or sufficiently rapidly. Thus, the adiabatic model illustrates the principle that when the characteristics of the noise are reasonably well understood, it may be possible to design suitable quantum hardware that effectively resists the noise. However, note that some of the effects of decoherence and unitary control error may not be significant for the small problems we are able to study—especially in the case of decoherence, where the time required by the simulation restricts us to systems with only four qubits—and hence, our data may not be indicative of the performance of the algorithm working on larger inputs.

A technique closely related to adiabatic computation was described by Kadowaki and Nishimori [6] and has been tested experimentally (in conjunction with a cooling procedure) by Brooke *et al.* [7]. In a different guise, the principles that make quantum adiabatic evolution robust also underlie the proposal by Kitaev [8] to employ nonabelian anyons for fault-tolerant quantum computation. The fact that adiabatic evolution incorporates a kind of intrinsic fault tolerance has also been noted in [9–14].

In Sec. II we review the adiabatic model of quantum computation, and in Sec. III we describe the specific combinatorial search problem (three-bit exact cover) that we use in our simulations. Sections IV and V report our numerical results on decoherence and unitary control error, and Sec. VI summarizes our conclusions.

II. ADIABATIC QUANTUM COMPUTATION

We briefly review the adiabatic model of quantum computation introduced in [1]. Let $h(z)$ be a function of n bits $z=(z_1,z_2,z_3,\ldots,z_n)$, and consider the computational problem of finding a value of z that minimizes $h(z)$. We will typically be interested in the case where this value of z is unique. We may associate with this function the Hermitian operator

$$H_P=\sum_{z=0}^{2^n-1} h(z)|z\rangle\langle z|, \tag{1}$$

so that the computational basis state $|z\rangle$ is an eigenstate of H_P with eigenvalue $h(z)$. Then the problem is to determine which state $|z\rangle$ is the ground state (eigenstate with lowest eigenvalue) of H_P. We refer to H_P as the *problem Hamiltonian.*

The strategy for finding the ground state of H_P is to prepare the ground state of some other *beginning Hamiltonian* H_B and slowly interpolate to H_P. A simple choice for the interpolation is given by the one-parameter family of Hamiltonians

$$\tilde{H}(s)=(1-s)H_B+sH_P \tag{2}$$

that interpolates between H_B and H_P as s varies from 0 to 1. We prepare the ground state of H_B at time $t=0$, and then the state evolves from $t=0$ to T according to the Schrödinger equation,

$$i\frac{\mathrm{d}}{\mathrm{d}t}|\psi(t)\rangle=H(t)|\psi(t)\rangle, \tag{3}$$

where the Hamiltonian is

$$H(t)=\tilde{H}(t/T). \tag{4}$$

At time T (the *run time* of the algorithm), we measure the state in the computational basis. If we let $|\varphi\rangle$ denote the (unique) ground state of H_P for a given instance of the problem, then the *success probability* of the algorithm for this instance is

$$\mathrm{Prob}(T)\equiv|\langle\varphi|\psi(T)\rangle|^2. \tag{5}$$

Does the algorithm work? According to the quantum adiabatic theorem [15,16], if there is a nonzero gap between the ground state and the first excited state of $\tilde{H}(s)$ for all $s\in[0,1]$, then $\mathrm{Prob}(T)$ approaches 1 in the limit $T\to\infty$. Furthermore, level crossings are nongeneric in the absence of symmetries, so a nonvanishing gap is expected if H_B does not commute with H_P. Thus, the success probability $\mathrm{Prob}(T)$ of the algorithm will be high if the evolution time T is large enough. The question is: how large a T is large enough so that $\mathrm{Prob}(T)$ is larger than some fixed constant?

We can reformulate this question in terms of

$$\Delta=\min_{s\in[0,1]}[E_1(s)-E_0(s)] \tag{6}$$

and

$$\mathcal{E}=\max_{s\in[0,1]}\left|\langle 1,s|\frac{\mathrm{d}\tilde{H}}{\mathrm{d}s}|0,s\rangle\right|, \tag{7}$$

where $E_0(s)$ is the lowest eigenvalue of $\tilde{H}(s)$, $E_1(s)$ is the second-lowest eigenvalue, and $|0,s\rangle$, $|1,s\rangle$ are the corresponding eigenstates. By calculating the transition probabil-

ity to lowest order in the adiabatic expansion [16], one finds that the probability of a transition from ground state to first excited state is small provided that the run time T satisfies

$$T \gg \frac{\mathcal{E}}{\Delta^2}. \tag{8}$$

If the spectrum consists of only two levels, then this condition is sufficient to ensure that the system remains in the ground state with high probability. In general, the required run time T will be bounded by a polynomial in n so long as Δ and $\mathcal{E}$ are polynomially bounded. For the problems we are interested in, $\mathcal{E}$ is polynomially bounded, so we only have to consider the behavior of Δ.

By rescaling the time, we can think of the evolution as taking place in the unit time interval between $s=0$ and 1, but in this case, the energy eigenvalues are rescaled by the factor T. Roughly speaking, we can think of $\mathrm{d}\tilde{H}(s)/\mathrm{d}s$ as a perturbation that couples the levels of the instantaneous Hamiltonian $\tilde{H}(s)$, and has the potential to drive a transition from $|0,s\rangle$ to $|1,s\rangle$. But if T is large, the effects of this perturbation are washed out by the rapid oscillations of the relative phase $\exp\{-iT\int_0^s \mathrm{d}s'[E_1(s')-E_0(s')]\}$.

Note that the Hamiltonian may be regarded as reasonable only if it is ''local,'' that is, if it can be expressed as a sum of terms, where each term acts on a bounded number of qubits (a number that does not grow with n). Indeed, in this case, the Hamiltonian evolution may be accurately and efficiently simulated by a universal quantum computer [17]. Many combinatorial search problems (e.g., 3SAT) can be formulated as a search for a minimum of a function that is local in this sense. Along with a local choice of H_B, this results in a full $H(t)$ that is also local.

A direct physical implementation of the continuously varying $H(t)$ would presumably be possible only under a somewhat stronger locality condition. We might require that each qubit is coupled to only a few other qubits, or perhaps that the qubits can be physically arranged in such a way that the interactions are spatially local. Fortunately, there are interesting computational problems that have such forms, such as 3SAT restricted to having each bit involved in only three clauses or the problem of finding the ground state of a spin glass on a cubic lattice [18]. However, for the purposes of our simulation, we will only consider small instances, and since we do not have a specific physical implementation in mind, we will not concern ourselves with the spatial arrangement of the qubits.

III. AN EXAMPLE: THE EXACT COVER PROBLEM

For definiteness, we study the robustness of the adiabatic algorithm via its performance on the problem known as ''three-bit exact cover'' (EC3). An n-bit instance of EC3 consists of a set of clauses, each of which specifies three of the n bits. A clause is said to be satisfied if and only if exactly one of its bits has the value 1. The problem is to determine if any of the 2^n assignments of the n bits satisfies all of the clauses.

For this problem, the function $h(z)$ is a sum

$$h(z)=\sum_C h_C(z_{i_C},z_{j_C},z_{k_C}) \tag{9}$$

of three-bit clauses, where

$$h_C(z_{i_C},z_{j_C},z_{k_C})=\begin{cases} 0, & (z_{i_C},z_{j_C},z_{k_C}) \quad \text{satisfies clause } C \\ 1, & (z_{i_C},z_{j_C},z_{k_C}) \quad \text{violates clause } C. \end{cases} \tag{10}$$

The value of the function $h(z)$ is the number of clauses that are violated; in particular, $h(z)=0$ if and only if z is an assignment that satisfies all the clauses.

To solve EC3 by the adiabatic algorithm, a sensible choice for the beginning Hamiltonian is

$$H_B=\sum_C H_{B,C}, \tag{11}$$

where

$$H_{B,C}=\frac{1}{2}(1-\sigma_x^{(i_C)})+\frac{1}{2}(1-\sigma_x^{(j_C)})+\frac{1}{2}(1-\sigma_x^{(k_C)}), \tag{12}$$

which has the ground-state

$$|\psi(0)\rangle=\frac{1}{2^{n/2}}\sum_{z=0}^{2^n-1}|z\rangle. \tag{13}$$

The resulting $H(t)$ is local in the sense that it is a sum of terms, each of which acts on only a few qubits. A stronger kind of locality may be imposed by restricting the instances so that each bit is involved in at most a fixed number of clauses. The computational complexity of the problem is unchanged by this restriction.

Numerical studies of the adiabatic algorithm applied to this problem were reported in [2,4]. Instances of EC3 with n bits were generated by adding random clauses until there was a unique satisfying assignment, giving a distribution of instances that one might expect to be computationally difficult to solve. The results for a small number of bits ($n \lesssim 20$) were consistent with the possibility that the adiabatic algorithm requires a time that grows only as a polynomial in n for typical instances drawn from this distribution. If this is the case, then the gap Δ does not shrink exponentially. Although the typical spacing between levels must be exponentially small, since there are an exponential number of levels in a polynomial range of energies, it is possible that the gap at the bottom is larger. For example, Fig. 1 shows the spectrum of a randomly generated seven-bit instance of EC3. The gap at the bottom of the spectrum is reasonably large compared to the typical spacing. This feature is not specific to this one instance, but is characteristic of randomly generated instances, at least for $n \lesssim 10$, beyond which the repeated matrix diagonalization required to create a picture of the spectrum becomes computationally costly. A large gap makes an in-

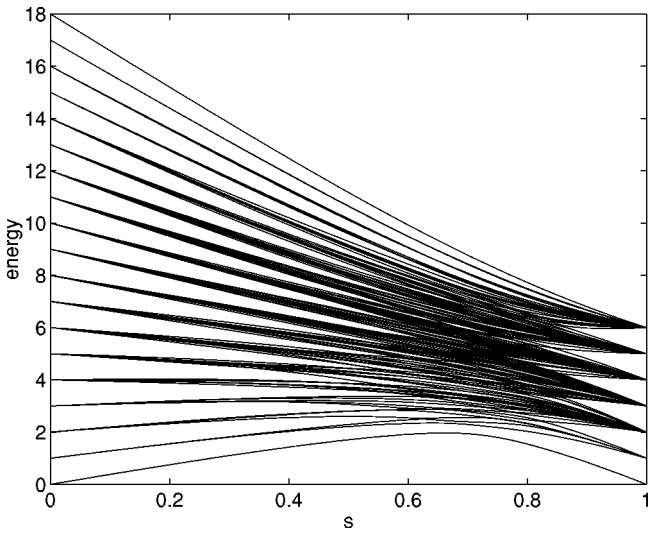

FIG. 1. Spectrum of a randomly generated $n=7$ bit instance of EC3 with a unique satisfying assignment. Note that the energy gap between the ground state and the first excited state is significantly larger than all other gaps. An expanded view would show that there are no level crossings anywhere in the spectrum (except for the degeneracies at $s=0$ and 1).

stance readily solvable by the adiabatic algorithm, and also provides robustness against thermal transitions out of the ground state.

IV. DECOHERENCE

Perhaps the most significant impediment to building a large-scale quantum computer is the problem of decoherence. No quantum device may be perfectly isolated from its environment, and interactions between a device and its environment will inevitably introduce noise. Fortunately, such effects can be countered using fault-tolerant protocols, but as we have already mentioned, these protocols may be costly. Therefore, we would like to consider quantum systems with inherent resistance to decohering effects. If the ground state of our adiabatic quantum computer is separated from the excited states by a sizable energy gap, then we expect it to exhibit such robustness. Here, we consider how the adiabatic algorithm for EC3 is affected by decoherence.

First, we briefly review the master equation formalism for describing the decohering effects of an environment on a quantum system. Suppose that our quantum computer is a collection of spin-1/2 particles interacting with each other according to the Hamiltonian H_S and weakly coupled to a large bath of photons. The total Hamiltonian of the quantum computer and its environment is

$$H=H_S+H_E+\lambda V, \tag{14}$$

where H_E is the Hamiltonian of its environment, V is an interaction that couples the quantum computer and the photon bath, and λ is a coupling constant. We may describe the state of the quantum computer alone by the density matrix ρ found by tracing over the environmental degrees of freedom. In general, the time evolution of ρ is complicated, but under reasonable assumptions, we may approximate its evolution using a Markovian master equation.

One way of deriving such a master equation is to consider the weak coupling limit, in which $\lambda \ll 1$ [19]. If the environment is very large and only weakly coupled to the quantum computer, it will be essentially unchanged by the interaction. Furthermore, in this limit, we expect the evolution of the quantum computer to be Markovian, or local in time, if we filter out high-frequency fluctuations by some coarse-graining procedure. Assuming that the combined state of the quantum computer and its environment begins in a product state $\rho(0)\otimes\rho_E$, Davies derives the master equation

$$\frac{d\rho}{dt}=-i[H_S,\rho]+\lambda^2 K^\natural \rho, \tag{15}$$

where

$$K\rho=-\int_0^\infty dx\, \mathrm{Tr}_E[U(-x)VU(x),[V,\rho\otimes\rho_E]], \tag{16}$$

$$K^\natural\rho=\lim_{x\to\infty}\frac{1}{x}\int_0^x dy\, U(-y)\{K[U(y)\rho U(-y)]\}U(y), \tag{17}$$

with

$$U(x)=e^{-ix(H_S+H_E)}, \tag{18}$$

where we have (temporarily) assumed that H_S is time independent. Although the $\natural$ operation defined by Eq. (17) does not appear in some formulations of the Markovian master equation, it appears to be essential for the equation to properly describe the weak-coupling limit [20], and in particular, for it to capture the physics of relaxation to thermal equilibrium. The master equation (15) has the property that if the environment is in thermal equilibrium at a given temperature, then the decohering transitions drive the quantum computer towards the Gibbs state of H_S at that temperature. While not an exact description of the dynamics, Eq. (15) should provide a reasonable caricature of a quantum computer in a thermal environment.

Note that Eq. (15) is derived assuming a time-independent Hamiltonian H_S; with a time-varying $H_S(t)$, we should expect the generator of time evolution at any particular time to depend on the Hamiltonian at all previous times [21]. However, if $H_S(t)$ is slowly varying, then it is a good approximation to imagine that the generator at any particular time depends only on H_S at that time [22]. In particular, since we are interested in nearly adiabatic evolution, $H_S(t)$ varies slowly, so Eq. (15) remains a good approximation, where at any given time t we compute $K^\natural$ using only $H_S(t)$. Note that with $H_S(t)$ time dependent, $U(x)$ defined by Eq. (18) is not the time evolution operator; it depends on the time t only implicitly through $H_S(t)$.

For a system of spins coupled to photons, we choose the interaction

$$V=\sum_i \int_0^\infty d\omega[g(\omega)a_\omega\sigma_+^{(i)}+g^*(\omega)a_\omega^\dagger\sigma_-^{(i)}], \qquad (19)$$

where Σ_i is a sum over the spins, $\sigma_\pm^{(i)}$ are raising and lowering operators for the ith spin, a_ω is the annihilation operator for the photon mode with frequency ω, and $\lambda g(\omega)$ is the product of the coupling strength and spectral density for that mode. Note that if the coupling strength is frequency dependent, we may absorb that dependence into $g(\omega)$, leaving λ as a frequency-independent parameter. With this specific choice for V, we can perform the integrals and trace in Eqs. (15–18). If we assume that all spacings between eigenvalues of H_S are distinct, the resulting expression simplifies considerably, and we find

$$\frac{d\rho}{dt}=-i[H_S,\rho]-\sum_{i,a,b}[N_{ba}|g_{ba}|^2\langle a|\sigma_-^{(i)}|b\rangle\langle b|\sigma_+^{(i)}|a\rangle + (N_{ab}+1)|g_{ab}|^2\langle b|\sigma_-^{(i)}|a\rangle\langle a|\sigma_+^{(i)}|b\rangle]\{(|a\rangle\langle a|\rho) + (\rho|a\rangle\langle a|)-2|b\rangle\langle a|\rho|a\rangle\langle b|\}, \qquad (20)$$

where the states $|a\rangle$ are the time-dependent instantaneous eigenstates of H_S with energy eigenvalues ω_a,

$$N_{ba}=\frac{1}{\exp[\beta(\omega_b-\omega_a)]-1} \qquad (21)$$

is the Bose-Einstein distribution at temperature $1/\beta$, and

$$g_{ba}=\begin{cases}\lambda g(\omega_b-\omega_a), & \omega_b>\omega_a,\\ 0, & \omega_b\leqslant\omega_a.\end{cases} \qquad (22)$$

We simulated the effect of thermal noise by numerically integrating the master Eq. (20) with a Hamiltonian H_S given by Eq. (4) and with the initial pure state density matrix $\rho(0)=|\psi(0)\rangle\langle\psi(0)|$ given by Eq. (13). For simplicity, we chose $g(\omega)=1$ for $\omega\geqslant 0$ and zero otherwise. Although we would expect that $g(\omega)\to 0$ as $\omega\to\infty$, for the small systems we are able to simulate, it should be a reasonable approximation to treat $g(\omega)$ as constant and tune the overall coupling strength using λ^2.

How should we expect the success probability $\langle\varphi|\rho(T)|\varphi\rangle$, where $|\varphi\rangle$ is the ground state of H_P, to depend on the run time T and the temperature? If the run time T is sufficiently long, then regardless of its initial state, the quantum computer will come to thermal equilibrium. At the time of the final readout, it will be close to the Gibbs state

$$\lim_{T\to\infty}\rho(T)=\frac{e^{-\beta H_P}}{\mathrm{Tr}\,e^{-\beta H_P}}\equiv\rho_P \qquad (23)$$

of the problem Hamiltonian H_P, and the success probability will be approximately $\langle\varphi|\rho_P|\varphi\rangle$. This probability may be appreciable if the temperature is small compared to the gap between the ground state and first excited state of H_P. Thus, one way to find the ground state of H_P is to prepare the computer in any initial state, put it in a cold environment, wait a long time, and measure. However, this thermal relaxation method is not an efficient way to solve hard optimization problems. Although it may work well on some instances of a given problem, this method will not work in cases where the computer gets stuck in local minima from which downward transitions are unlikely. In such cases, the time for equilibration is expected to be exponentially large in n.

Consider an instance with a long equilibration time so that cooling alone is not an efficient way to find the ground state of H_P. It is possible that the minimum gap Δ associated with the quantum algorithm is not small, and the idealized quantum computer, running without decohering effects, would find the ground state of H_P in a short time. In this situation, if we include the coupling of the system to the environment and we run at a temperature much below Δ, then thermal transitions are never likely, and the adiabatic algorithm should perform nearly as well as in the absence of decoherence. But if the temperature is comparable to Δ, then the performance may be significantly degraded.

On the other hand, consider an instance for which the equilibration time is short, so that cooling alone is a good algorithm. Furthermore, suppose that the adiabatic algorithm would find the ground state of H_P in a short time in the absence of decohering effects. In this case, the combined effects of cooling and adiabatic evolution will surely find the ground state of H_P in a short time. But note that Δ alone does not control the success of the algorithm. Even if $H(t)$ changes too quickly for the evolution to be truly adiabatic so that a transition occurs where the gap is smallest, the system may be cooled back into its ground state at a later time.

Typical results of the simulation are shown in Fig. 2 for two $n=4$ bit instances of EC3 with unique satisfying assignments. These two instances have minimum gaps of $\Delta\approx 0.301$ and $\Delta\approx 0.425$. For each instance, we plot the success probability as a function of the run time T. With $\lambda^2=0.1$, we consider five temperatures: 1/10, 1/2, 1, 2, and 10. We also present the data with no decoherence ($\lambda^2=0$) for comparison.

Unfortunately, the time required to integrate Eq. (20) grows very rapidly with n. Whereas a state vector contains 2^n entries, the density matrix contains 4^n entries; and in addition, calculating $d\rho/dt$ at each timestep requires evaluating a double sum over 2^n energy eigenstates. For this reason, we were only able to consider instances with $n\leqslant 4$.

The results are consistent with our general expectations. In the absence of decoherence, the success probability becomes appreciable for sufficiently long run times. This probability rises faster for the problem with a larger gap. When we add decoherence at high temperature, the success probability never becomes very large (note the lowest curves in Fig. 2). As the temperature is decreased to a value of order one, the presence of decoherence has a less significant effect on the success probability. In fact, for sufficiently low temperatures, the success probability may actually be higher in the presence of decoherence than when there is no decoherence. This is because the primary effect of decoherence at low temperature is to drive transitions towards the ground state, improving performance.

However, these results do not illustrate a definitive connection between the minimum gap Δ and the temperature

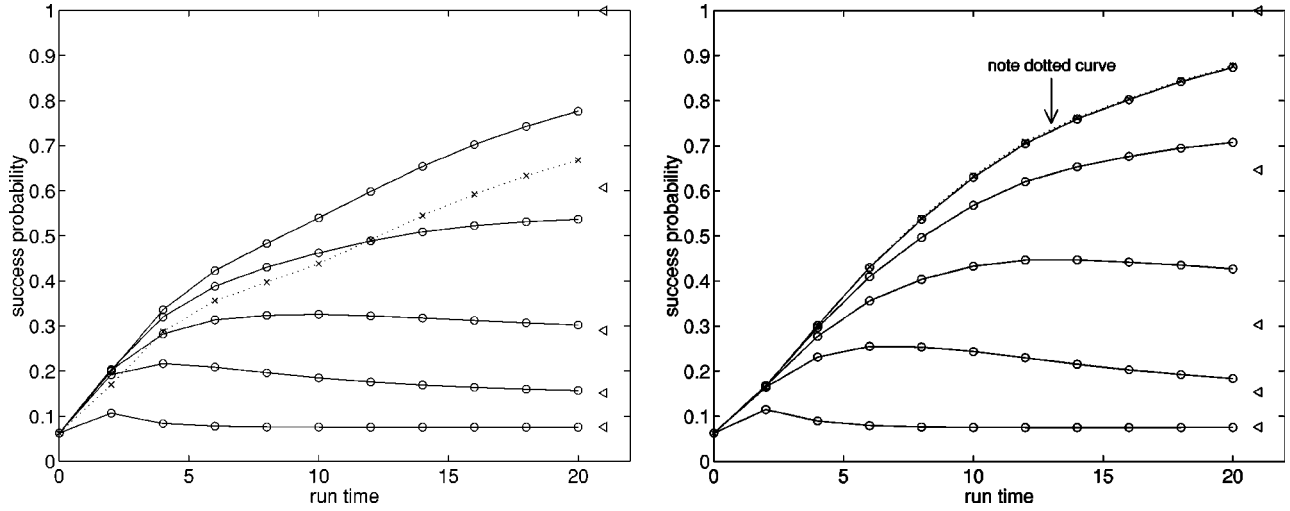

FIG. 2. The success probability as a function of run time T for two instances of EC3 with $n=4$ bits. The instance on the left has a gap of $\Delta_1 \approx 0.301$ and the instance on the right has a gap of $\Delta_2 \approx 0.425$. The dotted line shows the behavior of the algorithm with no decoherence, i.e., $\lambda^2=0$. Note that in the figure on the right, the dotted curve is partially obscured but can be seen slightly above the topmost solid curve. The solid lines show the behavior of the algorithm in the presence of decoherence with $\lambda^2=0.1$ for five different temperatures. The triangles at the far right show the thermal success probabilities $\langle\varphi|\rho_P|\varphi\rangle$ at each of these temperatures. From top to bottom, the temperatures are 1/10, 1/2, 1, 2, and 10.

above which the algorithm no longer works. These simple $n=4$ bit instances fall into the second category discussed above: the equilibration time is short, so cooling alone is a good algorithm. In other words, no sharp distinction can be drawn between the run time required for the adiabatic algorithm to perform well in the absence of decoherence and the run time required for equilibration. Accordingly, the dependence of the success probability on temperature and run time is similar for the two instances shown in Fig. 2, even though the minimum gaps for these instances are somewhat different.

V. UNITARY CONTROL ERROR

We now consider how the performance of the adiabatic algorithm for EC3 is affected by adding three different kinds of perturbations to the Hamiltonian. Each perturbation we consider is a sum of single-qubit terms, where each term can be interpreted as a magnetic field pointing in a random direction. To simplify our analysis, we assume that the magnitude of the magnetic field is the same for all qubits, but its direction varies randomly from qubit to qubit. The perturbations we consider are

$$\tilde{K}_1(s)=C_1 s\sum_{i=1}^{n}\hat{m}_i\cdot\vec{\sigma}^{(i)}, \tag{24}$$

$$\tilde{K}_2(s)=C_2\sin(\pi s)\sum_{i=1}^{n}\hat{m}_i\cdot\vec{\sigma}^{(i)}, \tag{25}$$

$$\tilde{K}_3(s)=\frac{1}{2}\sin(C_3\pi s)\sum_{i=1}^{n}\hat{m}_i\cdot\vec{\sigma}^{(i)}, \tag{26}$$

which are added to Eq. (2) and give a time-dependent Hamiltonian according to Eq. (4). Each $\hat{m}_i$ is a randomly generated real three-component vector with unit length, C_1 and C_2 are real numbers, and C_3 is a nonnegative integer.

The adiabatic algorithm was simulated by numerically solving the time-dependent Schrödinger equation with initial state $|\psi(0)\rangle$ given by Eq. (13) and Hamiltonian $\tilde{H}(t/T)+\tilde{K}_j(t/T)$ for a given $j\in\{1,2,3\}$. As in [2–4], we used a fifth-order Runge-Kutta method with variable step size, and checked the accuracy by verifying that the norm of the state was maintained to one part in a thousand. For a specified value of n, we randomly generated an instance of EC3 with a unique satisfying assignment. Then we randomly generated several different values of the magnetic field directions $\{\hat{m}_i\}$. For each instance of the problem and the magnetic field, the run time was chosen so that the success probability without the perturbation was reasonably high. With this run time fixed, we then determined the success probability for varying values of the relevant C_j.

First, we consider the perturbation K_1. Since it turns on at a constant rate, this perturbation can be thought of as an error in H_P. Note that with $C_1\neq 0$, the final Hamiltonian is not simply H_P, so the algorithm will not work exactly even in the adiabatic limit $T\to\infty$. This perturbation is potentially dangerous because of the way its effect scales with the number of bits n. Indeed, consider the case where H_P can be separated into a sum of Hamiltonians acting separately on each qubit. If adding K_1 reduces the overlap of the ground-state $|\varphi\rangle$ of H_P with the perturbed ground-state $|\varphi'\rangle$ by some fixed value ϵ for each of the n qubits, then the total overlap is $(1-\epsilon)^n$, which is exponentially small in the number of bits. Thus, the algorithm clearly fails in this factorized case. In general, if the magnitude of K_1 is independent of n, then we expect the algorithm to fail. However, if the magni-

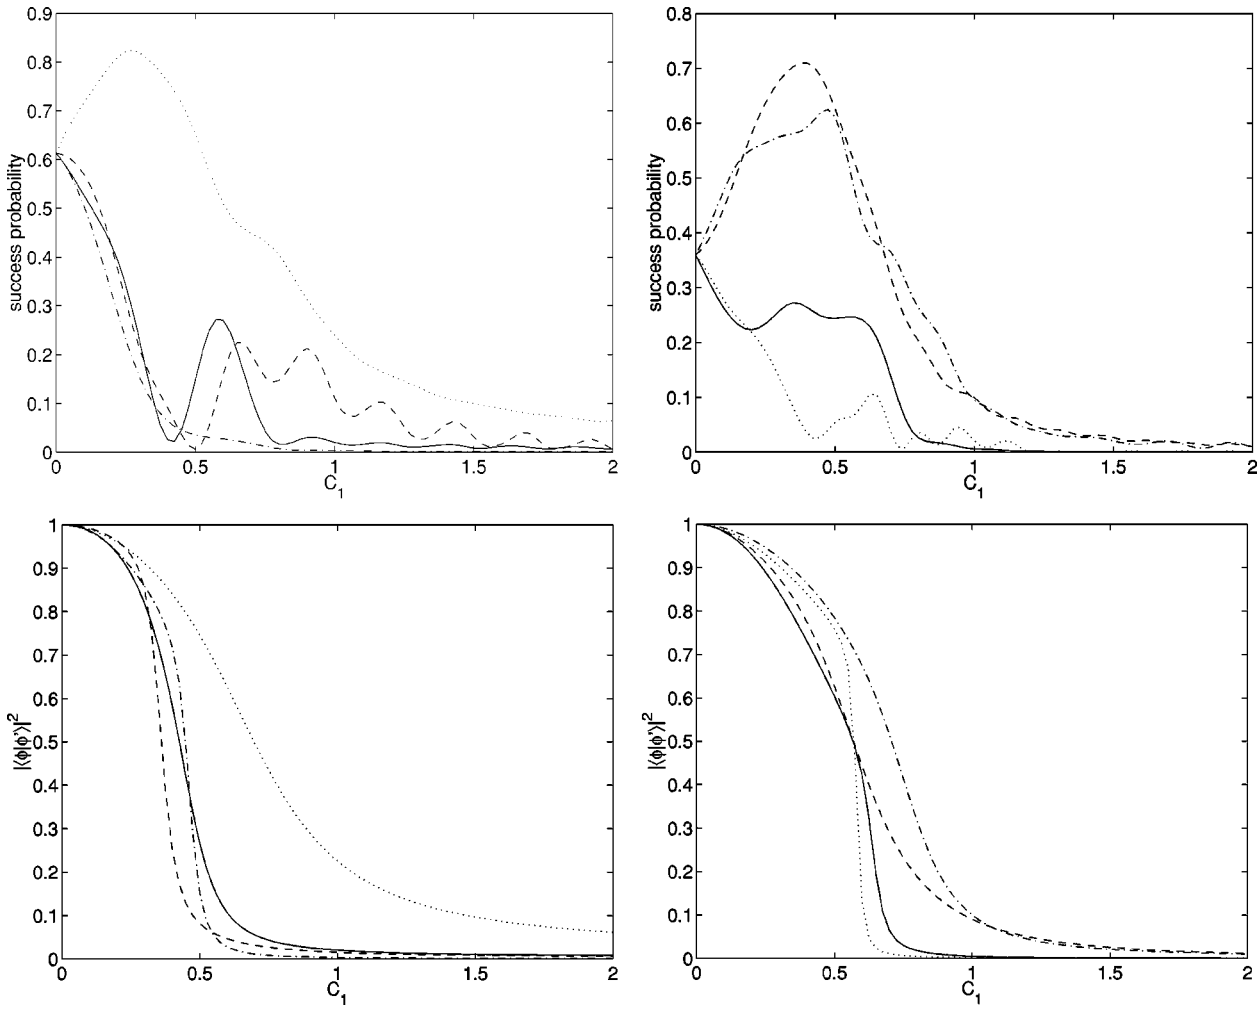

FIG. 3. (Top) The success probability of the adiabatic algorithm for two randomly generated instances of EC3 with $n=7$ bits (left) and $n=10$ bits (right) under the perturbation K_1 defined by Eq. (24) for four different sets of magnetic-field directions. For each n, the run time is the same for each random perturbation. (Bottom) The corresponding overlaps $|\langle\varphi|\varphi'\rangle|^2$ of the ground-state $|\varphi\rangle$ of H_P with the perturbed ground-state $|\varphi'\rangle$ at $s=1$.

tude of K_1 falls as $1/n$ or faster, then the shift of the ground state may be small enough (as it would be in the factorized case) that the algorithm is not significantly affected. Note that for any n there is some value of C_1 that is small enough that the disadvantage of reduced overlap with the ground state of H_P may be overcome if the perturbation happens to increase the minimum gap Δ. For this reason, we expect to sometimes see an increase in success probability for small C_1 that goes away as C_1 is increased.

The effect of the perturbation K_1 is shown in Fig. 3 for $n=7$ and 10 bit instances of EC3, with four different randomly generated sets of magnetic-field directions for each instance. The run time is chosen such that for $C_1=0$, the success probability is around 1/2. The top plots show that for small C_1, the success probability is not strongly suppressed; in fact, in some cases it is significantly enhanced. For large enough C_1, the success probability is heavily suppressed. The bottom plots show the overlap $|\langle\varphi|\varphi'\rangle|^2$ between the ground state of H_P and the actual ground state in the presence of the perturbation. As we expect, the suppression of the success probability is correlated with the amount of overlap. We also studied a similar perturbation in which s is replaced by $1-s$, which may be thought of as an error in H_B. Unsurprisingly, the results were qualitatively similar.

Next, we consider the low-frequency perturbation K_2. The period of oscillation is chosen such that the perturbation vanishes at $t=0$ and T, so the perturbation does not affect the algorithm in the adiabatic limit. Since the success probability is quite sensitive to the value of the minimum gap Δ, and it is not *a priori* obvious whether a perturbation will increase or decrease Δ, we can guess that turning on a nonzero value of C_2 may either increase the success probability or decrease it. In fact, it would be surprising if Δ decreased for all perturbations K_2. The Hamiltonian $\tilde{H}(s)+\tilde{K}_2(s)$ is another way to interpolate from H_B to H_P, and we know of no reason why the choice $\tilde{K}_2=0$ should always be optimal, even when the number of bits is large and C_2 is not decreasing with n.

Figure 4 shows the effect of the perturbation K_2, using the same instances, magnetic field directions, and run times as in Fig. 3. The top plots show the success probability as a function of C_2. As in the case of K_1, some perturbations may

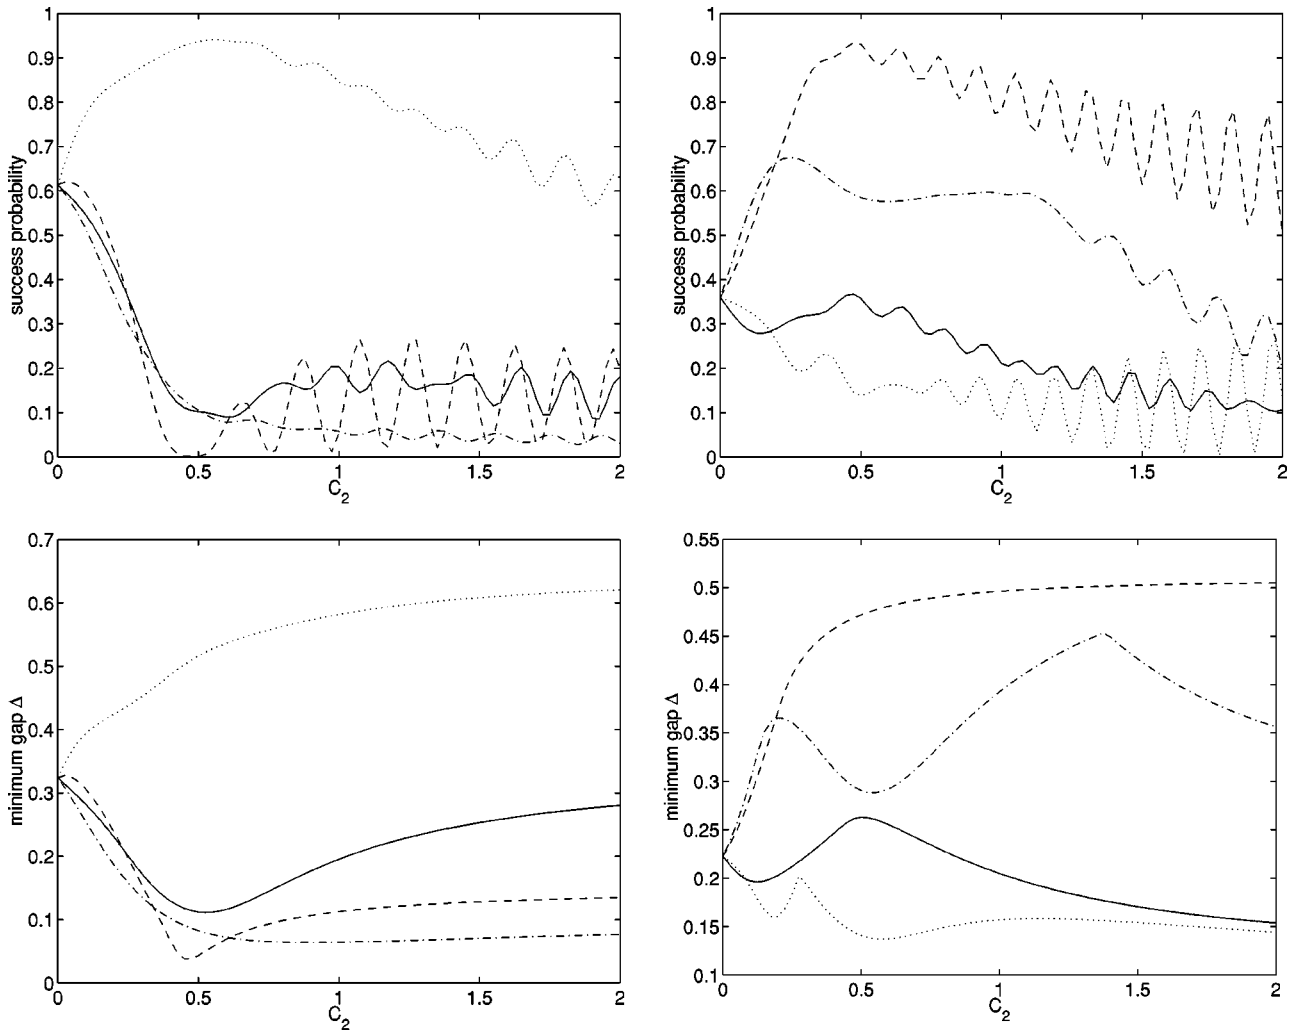

FIG. 4. (Top) The success probability of the adiabatic algorithm for the same instances used in Fig. 3 under the perturbation K_2 defined by Eq. (25). The four different magnetic field directions for each instance are also the same as in Fig. 3. (Bottom) The minimum gap Δ in the perturbed problem.

raise the success probability and some suppress it. Perhaps unsurprisingly, a particular set of magnetic field directions that raises the success probability under K_1 is also likely to help when K_2 is applied. But unlike K_1, K_2 may improve the success probability even with $C_2 \simeq 2$, where the size of the perturbation is comparable to the size of the unperturbed Hamiltonian. The bottom plots show the minimum gap Δ when the perturbation is added. Note that there is a strong correlation between the success probability and Δ.

For both perturbations K_1 and K_2, similar results have been observed (with fewer data points) for instances with as many as $n=14$ bits. Figures 3 and 4 present typical data. For example, for a given instance, typically one or two out of four sets of randomly chosen magnetic-field directions led to an improvement in the success probability for some values of C_1 and C_2, compared to the unperturbed case.

Finally, we consider the perturbation K_3, in which the magnitude of the oscillating component is fixed, but we may vary its frequency by varying C_3. As for K_2, the frequency is chosen so that the perturbation vanishes at $t=0$ and T. We expect that for C_3 of order one, the perturbation will be likely to excite a transition, and that the success probability will be small. But since both H_B and H_P have a maximum eigenvalue of order n, we may anticipate that for

$$C_3 \gg \frac{nT}{\pi}, \tag{27}$$

the perturbation will be far from any resonance. Then the probability that the perturbation drives a transition will be low, and the success probability should be comparable to the case where the perturbation vanishes.

Some representative plots of the dependence of the success probability on C_3 are shown in Fig. 5. Each plot corresponds to a particular randomly generated instance of EC3 (with either $n=8$ bits or $n=10$ bits) and a randomly generated set of magnetic field directions. In the top row of plots, the run time is chosen so that the success probability is around 1/8 with the perturbation absent (i.e., $C_3=0$). In the bottom row, the run time is doubled. All of the data exhibit

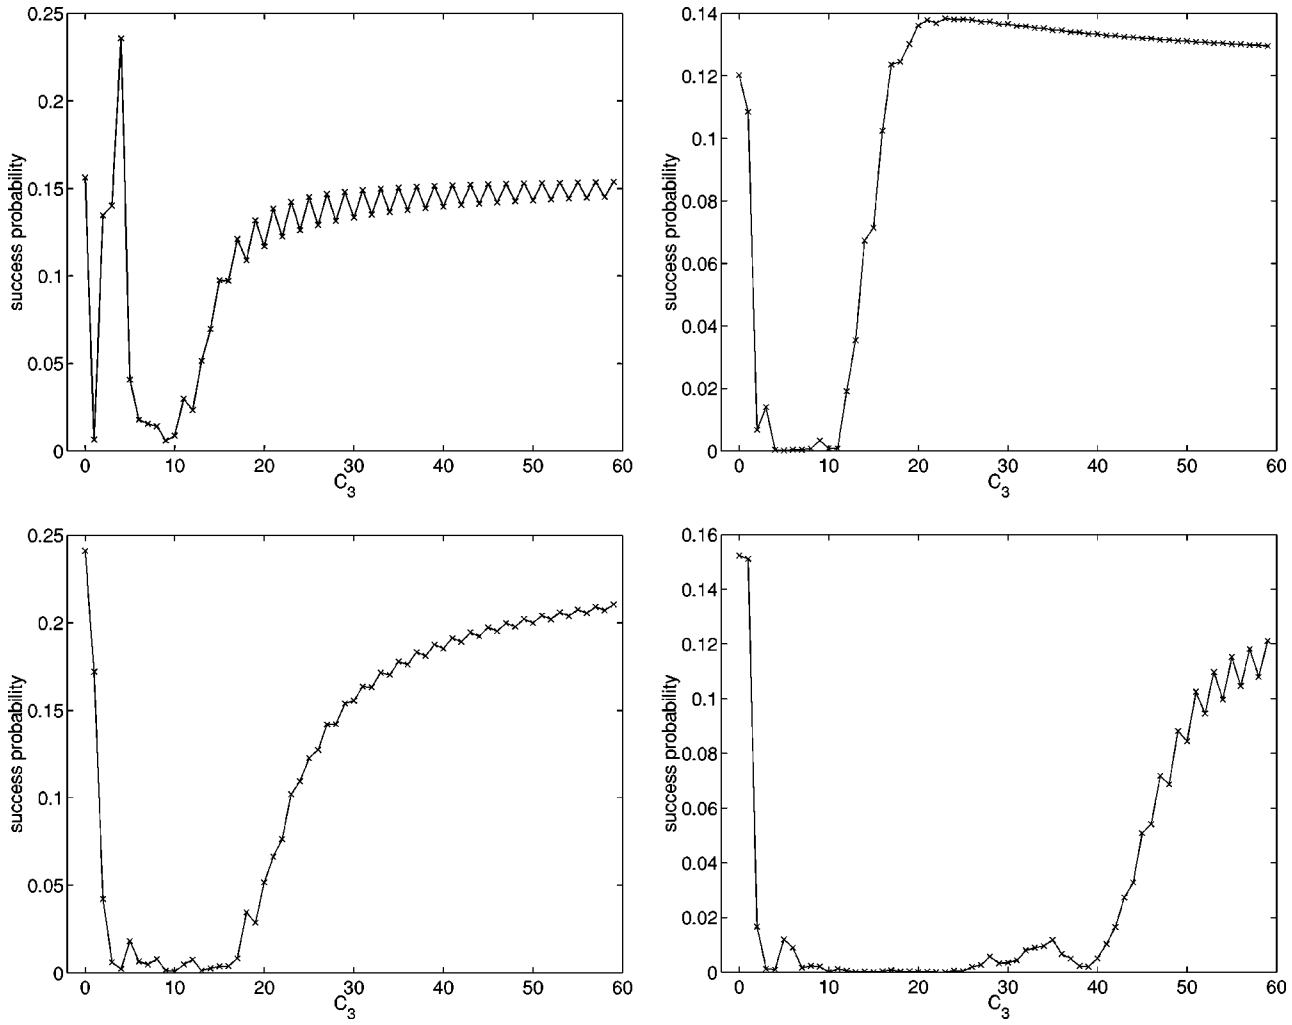

FIG. 5. The success probability as a function of the frequency C_3 of the perturbation K_3 defined in Eq. (26). The data in each plot were obtained for a randomly generated instance of EC3 with randomly generated magnetic-field directions. The data in the left column are for two instances with $n=8$ bits, and the data in the right column are for two instances with $n=10$ bits. For the top row, the run time is chosen so that the success probability is around 1/8 for $C_3=0$, and for the bottom row, the run time is twice as long. The leftmost points in each plot correspond to $C_3=0$, so the perturbation is absent for all t. C_3 takes integer values, so the lines are included only to guide the eye.

the expected qualitative trend. The leftmost point corresponds to $C_3=0$. For the smallest values of $C_3>0$, the success probability may not be too badly damaged; for somewhat larger values of C_3, it is heavily suppressed; and for sufficiently large C_3, it recovers to a value near the success probability in the absence of the perturbation. The value of nT/π is around 19 and 39 for the upper and lower $n=8$ plots and is around 38 and 76 for the upper and lower $n=10$ plots, so the estimate (27) turns out to be reasonable.

Another conspicuous feature of the plots in Fig. 5 is that the success probability tends to oscillate between even and odd values of C_3, though whether even or odd values are favored varies from case to case. This occurs because the perturbation's time average vanishes for C_3 even, so that its integrated effect is weaker than for C_3 odd. Since a small perturbation might either help or hurt, the success probability is slightly enhanced for odd C_3 in some cases, and is slightly suppressed in other cases.

VI. CONCLUSIONS

We have conducted numerical simulations to investigate the fault tolerance of adiabatic quantum computation, and our results are consistent with the claim that this algorithm is robust against decoherence and certain kinds of random unitary perturbations. Thus, if a physical system could be engineered with interactions reasonably well described by a Hamiltonian that smoothly interpolates from an initial H_B to a final H_P corresponding to an interesting combinatorial search problem, and if the gap remains large throughout the interpolation, that system might be a powerful computational device.

Although we have viewed unitary perturbations as noise, the fact that they sometimes raise the success probability suggests a possible way to speed up the adiabatic algorithm. The algorithm finds the ground state of H_P by starting the system in the ground state of H_B. The quantum state evolves

as the system Hamiltonian smoothly interpolates from H_B to H_P. However, there are many possible choices for H_B and many smooth paths from a given H_B to H_P. The choices (11) and (2) are convenient but arbitrary, so choosing an alternate route to H_P might speed up the algorithm. An example of this is seen in [23], where it is shown that optimizing the time-dependent coefficients of H_B and H_P allows the adiabatic algorithm to achieve a square-root speedup for an unordered search problem. More generally, the interpolating Hamiltonian might involve terms that have nothing to do with H_B or H_P, but that increase Δ and therefore improve performance. For example, the perturbation K_2 sometimes increases the success probability, as seen in Fig. 4. Rather than being thought of as a source of error, such a perturbation could be applied intentionally and might sometimes enhance the effectiveness of the adiabatic algorithm.

ACKNOWLEDGMENTS

We thank Todd Brun, Evan Fortunato, Jeffrey Goldstone, Sam Gutmann, Jeff Kimble, Alesha Kitaev, and Seth Lloyd for helpful discussions. A.M.C. gratefully acknowledges the support of the Fannie and John Hertz Foundation. This work was supported in part by the Department of Energy under Grant No. DE-FG03-92-ER40701 and Grant No. DE-FC02-94-ER40818, by the National Science Foundation under Grant No. EIA-0086038, by the Caltech MURI Center for Quantum Networks under ARO Grant No. DAAD19-00-1-0374, by the National Security Agency (NSA) and Advanced Research and Development Activity (ARDA) under Army Research Office (ARO) Contract No. DAAD19-01-1-0656, and by an IBM Faculty Partnership Award.

[1] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, e-print quant-ph/0001106.
[2] E. Farhi, J. Goldstone, and S. Gutmann, e-print quant-ph/0007071.
[3] A.M. Childs, E. Farhi, J. Goldstone, and S. Gutmann, e-print quant-ph/0012104.
[4] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, Science **292**, 472 (2001).
[5] P. W. Shor, in *Proceedings of the 37th Annual Symposium on Foundations of Computer Science* (IEEE Press, Los Alamitos, CA, 1996), pp. 56–65.
[6] T. Kadowaki and H. Nishimori, Phys. Rev. E **58**, 5355 (1998).
[7] J. Brooke, D. Bitko, T.F. Rosenbaum, and G. Aeppli, Science **284**, 779 (1999).
[8] A.Yu. Kitaev, e-print quant-ph/9707021.
[9] J. Preskill, in *Introduction to Quantum Computation and Information*, edited by H.-K. Lo, S. Popescu, and T. Spiller (World Scien tific, Singapore, 1998).
[10] W. Ogburn and J. Preskill, Lect. Notes Comput. Sci. **1509**, 341 (1999).
[11] P. Zanardi and M. Rasetti, Phys. Lett. A **264**, 94 (1999).
[12] A. Ekert, M. Ericsson, P. Hayden, H. Inamori, J.A. Jones, D.K.L. Oi, and V. Vedral, e-print quant-ph/0004015.
[13] S. Lloyd, e-print quant-ph/0004010.
[14] M.H. Freedman, A. Kitaev, M.J. Larsen, and Z. Wang, e-print quant-ph/0101025.
[15] T. Kato, Phys. Soc. Jap. **5**, 435 (1950).
[16] A. Messiah, *Quantum Mechanics* (North-Holland, Amsterdam, 1961), Vol. II.
[17] S. Lloyd, Science **273**, 1073 (1996).
[18] F. Barahona, J. Phys. A **15**, 3241 (1982).
[19] E.B. Davies, Commun. Math. Phys. **39**, 91 (1974).
[20] R. Dümcke and H. Spohn, Z. Phys. B: Condens. Matter **34**, 419 (1979).
[21] E.B. Davies and H. Spohn, J. Stat. Phys. **19**, 511 (1978).
[22] G. Lindblad, *Non-Equilibrium Entropy and Irreversibility* (Reidel, Dordrecht, 1983), pp. 65–71.
[23] J. Roland and N.J. Cerf, e-print quant-ph/0107015.