Journal of Physics: Conference Series **640** (2015) 012038 doi:10.1088/1742-6596/640/1/012038

# Numerical Studies of the Quantum Adiabatic Algorithm

**A. P. Young**[1] **and Itay Hen**[2]

[1]Department of Physics, University of California, Santa Cruz, California 95064
[2]Information Sciences Institute, University of Southern California, Marina del Rey, California 90292

E-mail: petery@ucsc.edu, itayhen@isi.edu

**Abstract.** Quantum annealers promise to solve practical optimization problems potentially faster than conventional classical computers. One of the major ongoing debates in this context pertains to their robustness against the decohering effects of finite temperature and interactions with the environment. We argue that even in an ideal setting of very low temperatures and in the absence of a decohering environment, quantum annealers do not necessarily perform better than classical heuristic solvers. Here, we numerically study the performance of the quantum adiabatic algorithm (QAA) on a variety of constraint satisfaction problems and a spin glass problem by studying the size dependence of the minimum energy gap during the evolution of the QAA. We do so by employing Quantum Monte Carlo schemes as these allow us to study these problems at much larger scales than exact methods would allow. We find that in all cases a quantum phase transition occurs and the minimum gap decreases exponentially with system size, leading to an exponentially large running time for the QAA. Based on these and other results, we briefly discuss potential modifications to the QAA that may improve the scaling of the minimum gap, leading to faster quantum adiabatic algorithms.

## 1. Introduction

In the standard paradigm for quantum computing, called the "gate" model, the quantum state of the qubits is acted on by a series of discrete unitary transformations [1, 2]. A general quantum state of $N$ qubits is a linear combination of $2^N$ basis states, and the unitary transformations act in parallel on all of them. Thus, *if phase coherence can be maintained*, these gate operations act in a *massively parallel* manner. Unfortunately, in order to get information out one must do a measurement, in which one observes just one state, with the appropriate probability, so a huge amount of information is lost at this point. It might therefore seem that the potential advantage of quantum parallelism is completely lost in practice. However, by doing some clever preprocessing *before* the measurement there are some problems for which quantum parallelism can produce spectacular gains. The best known, and most important, is Shor's algorithm [3] for factoring large integers. While the best classical algorithm for factoring an integer with $n$ bits takes a time proportional to the exponential of (a fractional power of) $n$ (exponential scaling) Shor's algorithm take a time proportional to a (fairly low) power of $n$ (polynomial scaling). Thus Shor's algorithm would be much more efficient than a classical one for large $n$. Unfortunately, there are huge experimental problems in maintaining phase coherence of more than a very small number of qubits for the time necessary to do a computation, and these are not yet solved.

Nonetheless, there is no doubt that, if they could be solved, a large *quantum speedup* would be obtained for certain problems by exploiting *quantum parallelism*.

Quantum parallelism arises because a quantum state can be a linear superposition of basis states (e.g. a single spin can be simultaneously "up" and "down"). However, this is only *one* of the ways quantum state differs from a classical state. Another difference is that a quantum system can *tunnel* through energy barriers to get to regions of configuration space that would be impossible classically. Consequently, an alternative paradigm for quantum computing has emerged which aims to take advantage of tunneling (rather than quantum parallelism). This approach is called *quantum annealing* [4, 5] (QA) and is very closely related to the *quantum adiabatic algorithm* (QAA) proposed by Farhi et al. [6]. These methods aim to solve optimization problems, in which one has to find the minimum of a function of many variables (which we will call the "energy") with constraints. Such problems occur widely in the sciences, engineering and industry. Practical problems are difficult because no configuration of the variables simultaneously minimizes all terms in the Hamiltonian, a concept known as "frustration" [7, 8] in the spin glass community, so one cannot find the ground state by repeatedly cycling through each of the variables and minimizing with respect to them one at a time. This "greedy" algorithm finds a local minimum but not, in general, the absolute minimum. For the hardest optimization problems, there are no known efficient algorithms to solve them in sub-exponential time. For these, the time $\tau$ to find the global minimum increases exponentially with the system size on a classical computer, i.e.

$$\tau \propto \exp(\mu N)\,. \tag{1}$$

The question of whether a quantum computer can do better than this using QA/QAA is one of the most intriguing open questions in the field to date, as there is no proof that a quantum computer exploiting quantum tunneling could solve optimization problems more efficiently than a classical computer. On the other hand, there is no proof that it could not. Hence, in practice one has to resort to numerics to try to answer this question. While it is probably too optimistic to expect that a quantum computer could convert the exponential dependence in $N$ in Eq. (1) to a power-law (polynomial) dependence, it would still be valuable if a quantum computer could reduce substantially the value of the coefficient $\mu$ since one would then be able to study larger sizes than those possible on a classical computer.

A general purpose *classical* method to avoid getting stuck in the nearest local minimum, using physics ideas, was proposed many years ago [9] and is called simulated annealing (SA). In this approach one simulates, using classical Monte Carlo methods [10], the Hamiltonian at a non-zero temperature $T$, which means that there is some probability to make a move which *increases* the energy, and so the system can escape from a local minimum. During the course of the simulation the temperature is gradually decreased to zero and one expects that the system will end up in the global minimum if the temperature reduction is done slowly enough.

Quantum annealing (QA) has a similar motivation to that of simulated (i.e. classical, or thermal) annealing (SA), but rather than including thermal fluctuations whose strength tends to zero at the end of the run, one adds a term to the Hamiltonian which does not commute with the "problem" Hamiltonian (the one that we are finding the ground state of), which thereby induces quantum fluctuations. The amount of this extra Hamiltonian, and hence the strength of the quantum fluctuations, is decreased to zero during the run. QA is often implemented using Quantum Monte Carlo simulations in which the time to perform the calculation is taken to be the annealing time (number of Monte Carlo sweeps) [5, 11] rather that being obtained from correlations in the "imaginary" time direction of the simulation. Following Ref. [11] we shall call this approach simulated quantum annealing (SQA).

The quantum adiabatic algorithm (QAA) shares many similarities with QA. However, it is designed to makes use of the adiabatic theorem of quantum mechanics [12] as follows. We denote

the problem Hamiltonian, whose ground state we want, by $\mathcal{H}_P$. Typically this will be a function of bits, $z_i = 0$ or 1, $(i = 1, 2, \cdots N)$, or equivalently Ising spins $\sigma_i^z = \pm 1$. One example would be a spin glass Hamiltonian where

$$\mathcal{H}_P = -\sum_{\langle i,j \rangle} J_{ij}\sigma_i^z\sigma_j^z - \sum_i h_i\sigma_i^z \,, \tag{2}$$

in which the two-spin interactions $J_{ij}$ can have either sign at random. In the spin glass context the fields $h_i$ are often set to zero. To induce quantum fluctuations, we add a non-commuting "driver" Hamiltonian $\mathcal{H}_D$, which, in the simplest case, is just a transverse field which takes the same value on each site, i.e.

$$\mathcal{H}_D = -\Gamma \sum_i \sigma_i^x \,, \tag{3}$$

where $\sigma_i^x$ and $\sigma_i^z$ are now Pauli spin operators. Typically we shall set the parameter $\Gamma$ to unity. The interactions between the qubits are adjusted so that they have the following Hamiltonian,

$$\mathcal{H}(s) = s(t)\mathcal{H}_P + [\,1 - s(t)\,]\mathcal{H}_D \,, \tag{4}$$

where $s(t)$ is a parameter varying smoothly with time, from 0 at $t = 0$ to 1 at the end of the algorithm, $t = \mathcal{T}$, so the initial Hamiltonian is $\mathcal{H}_D$ and the final Hamiltonian is $\mathcal{H}_P$. The qubits are initially prepared with $\sigma_i^x = 1$, so the system starts in its ground state. If the time evolution in Eq. (4) is done slowly enough, the adiabatic theorem of Quantum Mechanics [12] ensures that the system will stay close to the ground state of the instantaneous Hamiltonian throughout the evolution, so that one finally obtains a state close to the ground state of $\mathcal{H}_P$. At this point, measuring the state will give the solution of the original problem with high probability. The crucial question, which will be the focus of this article, is how does the time to maintain adiabicity, and hence to solve the problem with high probability, vary with system size $N$.

A bottleneck of the QAA occurs when the gap between the ground state and the first excited state becomes very small, because it is hard to maintain adiabatic evolution in this case. An upper bound for the run time needed to find the ground state of $\mathcal{H}_P$ is given in terms of the eigenstates $\{|n\rangle\}$ and eigenvalues $\{E_n\}$ of the Hamiltonian $\mathcal{H}(s)$, by [13, 14]

$$\mathcal{T} \gg \hbar \, \frac{|\max_s V_{10}(s)|}{(\Delta E_{\min})^2} \,, \tag{5}$$

where $\Delta E_{\min}$ is the minimum of the first excitation gap $\Delta E_{\min} = \min_s \Delta E$ with $\Delta E = E_1 - E_0$, and $V_{n0} = \langle 0|\mathrm{d}\mathcal{H}/\mathrm{d}s|n\rangle$. Typically, matrix elements of $\mathcal{H}$ scale as a low polynomial of the system size $N$, and the question of whether the running time depends polynomially or exponentially with $N$ therefore depends on how the minimum gap $\Delta E_{\min}$ scales with $N$. This means that if the gap becomes exponentially small at any point in the evolution, then the computation requires an exponential amount of time, rendering the QAA inefficient. The dependence of the minimum gap on the system size for a given problem is therefore a central issue in determining the running time of the QAA. The minimum gap will vary for different samples of a given size, so some averaging over instances of the same size is required. Here, we will determine performance of the QAA by looking at the "typical" minimum gap.

The study of QA/QAA is not just of theoretical interest, as there are now experiments on real hardware with several hundred qubits. A Vancouver-based company D-Wave has produced a machine, D-Wave One, with 128 superconducting qubits to implement the QAA, and recently has come out with bigger version, D-Wave Two, having 512 qubits [15, 16]. Thus one can compare the running time of real quantum hardware with that of numerical simulations, both classical and simulated quantum annealing, for non-trivial optimization problems [11, 17]. Since

the D-Wave machine is affected by the temperature being non-zero and by non-thermal noise, one can also investigate expderimentally whether quantum tunneling of a large number of qubits is destroyed by these effects, in which case the machine is effectively doing thermal annealing rather than quantum annealing. Since the system has an energy gap, the hope is that noise is not fatal to quantum annealing, as it is to the gate model of quantum computing [18, 19]. The qubits in the D-Wave machine are not claimed to maintain coherence during the run, so we shall describe this machine as a "quantum annealer" rather than a quantum computer, and use the latter term for devices which do maintain coherence.

Two types of questions need to be answered about QA/QAA:

(a) Suppose we consider an ideal situation in which the temperature is zero and also there isn't any non-thermal noise. In these ideal circumstances, does the time to solve a hard optimization problem with QA/QAA scale better with size than for a classical algorithm such as simulated annealing? Assuming that scaling is exponential, this is equivalent to asking if the coefficient $\mu$ in Eq. (1) is smaller for the quantum algorithm than for a classical algorithm.

(b) Now include the effects of non-zero temperature and non-thermal noise which are present in the D-Wave machine. Do these destroy quantum tunneling and render the machine effectively a thermal annealer? If not, does the run time on the D-Wave machine scale better with size than that of a classical simulation?

Concerning (b), some recent work [11, 20, 17] indicates that the D-Wave machine is doing quantum annealing, though it does not appear that the scaling with size is better than that of an optimized classical algorithm. Experiments with an even larger number of qubits, e.g. 2048, would be helpful in elucidating the asymptotic performance of the D-Wave hardware.

Here we will summarize our own work [21, 22, 23, 24, 25, 26, 27] which focuses on the ideal situation in (a) above. We shall consider several optimization problems, and compare the scaling with size of the run time of the QAA with that of a heuristic classical algorithm similar to simulated annealing. We will find that the basic implementation of the QAA which we use is not better than the classical algorithm for any of the problems studied. We will also consider how one might improve this implementation of the QAA to get better results.

## 2. Numerical results

### *2.1. Method*

To find the minimum gap for a specific instance of a specific problem, we perform quantum Monte Carlo simulations for a range of $s$ values ($s$ being the interpolating adiabatic parameter) that bracket the minimum gap. The specific method we use in this study is known as the stochastic series expansion (SSE) algorithm [28, 29] which involves a Taylor series expansion of the partition function $\mathrm{Tr}[\mathrm{e}^{-\beta\hat{H}}]$ and uses a discrete representation of continuous imaginary time. This discretization however does not introduce errors into the algorithm as is the case in the alternative path-integral formulation, where one usually performs a Trotter-type discretization of imaginary time, see e.g. Refs. [30, 21], though formulations in continuous imaginary time also exist [31, 32]. Here $\beta$ is the inverse temperature $1/T$ (in our units $k_B = 1$).

The SSE algorithm has several properties that are very useful in addressing the problems we focus on in this study. Firstly, it works in continuous imaginary time as discussed above. Secondly, it allows not only local updates of system configurations but also global cluster updates, which in most cases prove to be more efficient than single-spin-flip updates. These global updates are achieved by dividing the configurations of the system produced by the QMC into clusters and then flipping a fraction of them within each sweep of the simulation [33]. An important bonus of cluster updates is the existence of "improved estimators" for determining time-dependent correlation functions, for which the signal to noise is much better than with conventional measurements.

In addition, we speed up equilibration by implementing "parallel tempering" [34], where

simulations for different values of $s$ are run in parallel and spin configurations with adjacent values of $s$ are swapped with a probability satisfying the detailed balance condition. Traditionally, parallel tempering is performed for systems at different temperatures, but here the parameter $s$ plays the role of (inverse) temperature.

The gap of the system for a given instance and a given $s$ value is extracted by analyzing measurements of (imaginary) time-dependent correlation functions of the type

$$C_A(\tau) = \langle \mathcal{A}(\tau)\mathcal{A}(0)\rangle - \langle A\rangle^2 \,, \tag{6}$$

where the operator $\mathcal{A}$ is some measurable physical quantity. In the low temperature limit, $\Delta E \ll T$ where $\Delta E = E_1 - E_0$, the system is in its ground state so the imaginary-time correlation function is given by

$$C_A(\tau) = \sum_{m=1} |\langle 0|\mathcal{A}|m\rangle|^2 \left(e^{-\Delta E_m \tau} + e^{-\Delta E_m(\beta-\tau)}\right) \,, \tag{7}$$

where $\Delta E_m = E_m - E_0$. At long times, $\tau$, the correlation function is dominated by the smallest gap, $\Delta E \equiv \Delta E_1$, (as long as the matrix element $|\langle 0|\mathcal{A}|1\rangle|^2$ is nonzero). On a log-linear plot $C_A(\tau)$ then has a region where it is a straight line whose slope is the negative of the gap. This can therefore be extracted by linear fitting. The difficulty is that the straight-line region is at long times where the signal is very small and is easily dominated by the Monte Carlo noise. However, the signal to noise can be dramatically improved by optimizing the choice of the operator $\mathcal{A}$, as discussed in Ref. [25].

As we shall see, a couple of the problems we analyze have global bit-flip symmetry in which case, a subtle difficulty arises in the computation of the gap. Eigenstates of the Hamiltonian are either even or odd under bit-flip symmetry (in particular, the ground state is even). In the $s \to 1$ limit, states occur in even-odd pairs with an exponentially small gap (see Fig. 1 for an illustration). Therefore, the quantity of interest is the gap to the first even state. To do that, we consider correlation functions of even quantities, so there are only matrix elements between states of the same parity. However, the lowest odd level becomes very close to the ground state near where the gap to the first even excited state has a minimum, hence this lowest odd state becomes thermally populated, with the result that odd-odd gaps are present in the data as well. We have eliminated these undesired contributions by projecting out the symmetric subspace of the Hamiltonian: In standard quantum Monte Carlo simulations one imposes periodic boundary conditions in imaginary time. To project out the symmetric subspace one imposes, instead, free boundary conditions. The properties of the symmetric subspace can then be obtained, for $\beta \to \infty$, by measurements far from the boundaries. We have incorporated this idea into our QMC scheme, and used this modified algorithm in the simulations of problems with global bit-flip symmetry (the reader is referred to Ref. [23] for a more detailed discussion of this issue).

For each of the studied $s$ values we extract the gap and interpolate the minimum value using a simple quadratic fit. We determine the typical minimum gap for different sizes by considering about 50 instances for each size and extracting the minimum gap for each of them. We take the *median* value of the minimum gap among the different instances for a given size as a measure of the "typical" minimum gap for that size.

### *2.2. SAT problems*

In Ref. [22], we studied three optimization problems which had previously been suggested [35, 36, 37] as good potential candidates for detailed investigation. These are of the "constraint satisfaction" (SAT) type in which there are $N$ bits and $M$ "clauses", each clause being a logical condition on a small number of the bits. The problem is to decide whether there is an assignment of the $N$ bits which satisfies all of $M$ clauses. Each bit is represented in the Hamiltonian by

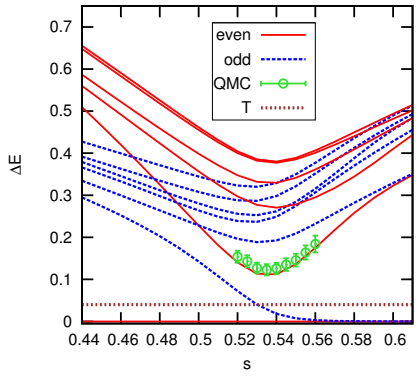


**Figure 1.** (Color online) Energy gaps to even (solid, red) and odd (dashed, blue) excited states for an $N = 16$ instance of the locked 2-in-4 problem, which has bit-flip symmetry as discussed in the text. The dotted line shows a characteristic value of another important energy scale in the problem, temperature. In the region where the gap to the first even state has a minimum, the gap to the first odd state becomes very small and is inevitably thermally populated. Hence, odd-odd gaps appear in this region as well as even-even gaps. This is the reason why we use a non-standard Monte Carlo algorithm for this problem which projects out the symmetric subspace, so only even-even gaps are present in the data. The figure also shows the gap obtained from the even-subspace projected QMC in vicinity of the minimum. It agrees with exact diagonalization within the error bars.

the $z$-component of a Pauli matrix, $\sigma_i^z$, where $i$ labels the spin. Each clause is converted to an energy function which depends on the spins associated with the clause, such that the energy is zero if the clause is satisfied and is a positive integer if it is not. The general structure of the problem Hamiltonian $\mathcal{H}_p$ is therefore

$$\mathcal{H}_P = \sum_{a=1}^{M} \mathcal{H}_a \,, \tag{8}$$

where $a$ is the clause index and $\mathcal{H}_a$ is the energy associated with the clause and is a function of the spins $\sigma_i^z$ belonging to it.

Clearly, it is easy to satisfy all clauses if the ratio $\alpha \equiv M/N$ is small enough. In fact, one expects an exponentially large number of satisfying assignments in this region. Conversely, if $M/N$ is very large, with high probability there will be a conflict between different clauses. Hence there is a "satisfiability transition" at some value $\alpha_s$ where the number of satisfying assignments goes to zero. It is particularly hard to solve satisfiability problems close to the transition [38], so we will work in this region. Furthermore, when studying the efficiency of the QAA numerically [6, 30, 21], it is convenient to consider instances with a unique satisfying assignment (USA), i.e. the ground state of the problem Hamiltonian is non-degenerate. This also means that we must take a value of $\alpha$ very close to $\alpha_s$.

We studied three different SAT problems which go under the arcane (for the physics community) names of "locked 1-in-3 SAT", "locked 2-in-4 SAT", and "3-reg-3-XORSAT". The first two problems are "locked" problems – a term first introduced by Zdeborová and Mézard [35, 36] for problems with instances having the following two properties: (i) every variable is in at least two clauses, and (ii) one can not get from one satisfying assignment to another by flipping a single bit. These locked problems have several properties that make them eminently suitable as benchmarks: They are analytically "simple", but are computationally hard. Also, fluctuations between instances are smaller than with "unlocked" problems. In the

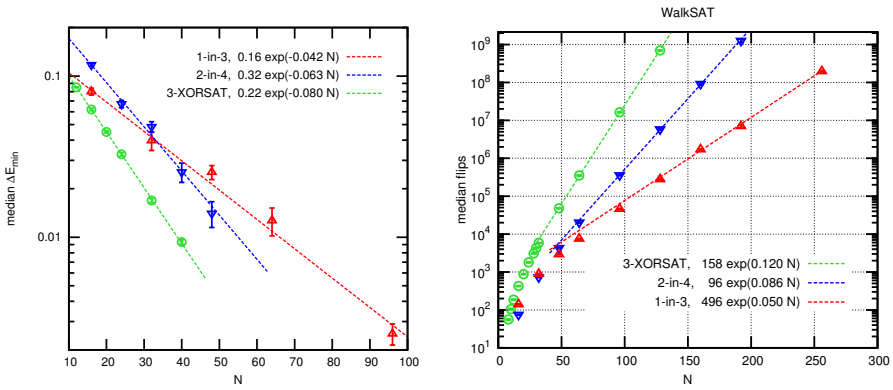


**Figure 2.** (Color online) Left: Median gap on a log-linear scale for the locked 1-in-3 (red), locked 2-in-4 (blue) and 3-reg-3-XORSAT (green) problems. The straight-line fits are fairly reasonable, indicating that the complexity for these problems is indeed exponential. Right: Median number of flips to solution of the WalkSAT algorithm applied to the same three problems also on a log-linear scale.

| model | $\mu$(QAA) | $\mu$(WalkSAT) | Ratio |
|---|---|---|---|
| locked 1-in-3 SAT | 0.084(3) [22] | 0.0505(5) [40] | 1.66 |
| locked 2-in-4 SAT | 0.126(5) [22] | 0.0858(8) [40] | 1.47 |
| 3-reg-3-XORSAT | 0.159(2) [23] | 0.1198(20) [40] | 1.32 |

**Table 1.** Values of $\mu$, the coefficient of $N$ in the exponential running time expression in Eq. (1), for the Quantum Adiabatic Algorithm (QAA) and for the analogous classical algorithm WalkSAT, as well as the ratios between them, for three SAT problems.

1-in-3 SAT problem each clause consists of three bits chosen randomly, and the clause is satisfied if one of the bits is one and the others are zero. Similarly, in locked 2-in-4 instances, a clause has four bits, and is satisfied if two are zero and two are one. The reader is referred to Ref. [22] for a more detailed description of the problems. For our purposes here we only need to know that these are three models with which we tested the performance of the QAA. The "3-reg-3-XORSAT" is an example of a $p$-spin model with $p = 3$ of the type that has been studied in the spin glass community (see, e.g., Ref. [39]) The latter model was studied for small sizes, up to $N = 24$, using diagonalization by Jörg *et al.* [37] while Refs. [22] and [23] extended the range of sizes up to $N = 40$ by Quantum Monte Carlo simulations, obtaining consistent results. For all three models we find [22, 23] that the minimum gap decreases exponentially with $N$ (see left pane of Fig. 2).

As discussed in Sec. 1, the computation time is proportional to $1/\Delta E_{\min}^2$ (neglecting $N$ dependence of matrix elements) and since we find [22, 23] that $\Delta E_{\min} \sim \exp(-c\,N)$, the running time can be written in the form of Eq. (1) with $\mu = 2\,c$.

It is interesting to compare the efficiency of the QAA with that of a classical algorithm. In Ref. [40] it was argued that a reasonable classical algorithm to compare with QAA is the heuristic local search algorithm known as WalkSAT [41], which is similar in spirit to simulated annealing in that both make moves which reduce the "energy", but also sometimes make moves which increase it to avoid being trapped in the nearest local minimum. In the WalkSAT algorithm, the running time is proportional to the number of "bit flips" the algorithm makes (for more details, the reader is referred to Ref [40]). These results are summarized in Fig. 2 (right pane). Writing the median number of flips as $N_{\text{flips}} \propto e^{\mu N}$, we compare in Table 1 [22] the values of $\mu$ defined by Eq. (1) for the QAA and WalkSAT.

As Table 1 indicates, the exponent coefficients obtained with WalkSAT are somewhat smaller than those of the QAA, suggesting that the latter algorithm, in the specific way it was implemented in Ref. [22], is slightly less efficient than its corresponding classical one for these three problems. It is also evident from the table that the harder the problem is for WalkSAT, the harder it also is for QAA.

Previous work [42, 37, 43] had argued that a first order quantum phase transition occurs for a broad class of random SAT problems, and indeed our results are consistent with the gap being a minimum *at* a first order quantum phase transition.

*2.3. A spin glass problem*

SAT problems are standard fare for computer scientists working on optimization, but are rather unfamiliar to physicists. We have therefore also studied [23] a problem in the physics domain, a spin glass. A spin glass model should have both disorder and "frustration". Most spin glass models [44] achieve this by having the spins on a regular lattice and choosing the interactions between neighbors to be random with a significant probability of having either sign. Here, by contrast, we put the spins on the vertices of a regular random graph of coordination number (degree) equal to three, rather than a regular lattice, and make all the interactions equal and antiferromagnetic. Locally the graph looks like a tree but, since there are no free ends, there are large loops with a typical size of order $\ln N$. Disorder and frustration come entirely from these large loops. Disorder arises because the connections between the sites are chosen at random so the structure of the loops is random. Frustration arises because half the loops, on average, have an odd number of links and these are frustrated. Because of these odd-size loops, it is not possible to form a regular antiferromagnetic state and the system is actually a spin glass [45]. The problem Hamiltonian is very simple:

$$\mathcal{H}_P = \sum_c \sigma^z_{i_1,c}\sigma^z_{i_2,c}\,, \tag{9}$$

where $c$ runs over the $3N/2$ links, and $(i_1, c)$ and $(i_2, c)$ refer to the two sites on a link. Again we choose samples (instances) for which $\mathcal{H}_P$ has a unique ground state, apart from the global symmetry of flipping all the $\sigma^z_i$.

For small $s$ the system is dominated by the transverse field and equilibration is fast; the system is in a quantum paramagnetic phase. As $s$ increases, the system goes through a quantum spin glass phase transition at a critical value $s_c \simeq 0.36$, while for still larger values of $s$ the system is in a quantum spin glass phase. We were not able to equilibrate the system all the way in the spin glass phase up to the end point $s = 1$. However, even considering just the part of the spin glass phase which could be equilibrated ($s_c = 0.36 < s < 0.5$) we found that the minimum gap decreases exponentially with size, as shown in Fig. 3 (left pane).

Interestingly, if we just consider data right in the vicinity of the critical point, the minimum gap in this region only seems to decrease with a power of the system size [23], at least for the range of sizes that could be studied.

We see that the bottleneck for the QAA occurs at a different place for the SAT problems and the spin glass. For SAT, the gap becomes very small *at* the (first order) quantum phase transition, while for the spin glass it becomes very small *beyond* the (second order) transition in the spin glass phase.

## 3. Modifications to the simplest approach

The numerical studies described above have made it clear that a simple application of the QAA, i.e., an equal-weight transverse-field driver Hamiltonian and a linear interpolating schedule, does not lead to quantum speedups, even in the ideal setting of zero-temperature and in the absence

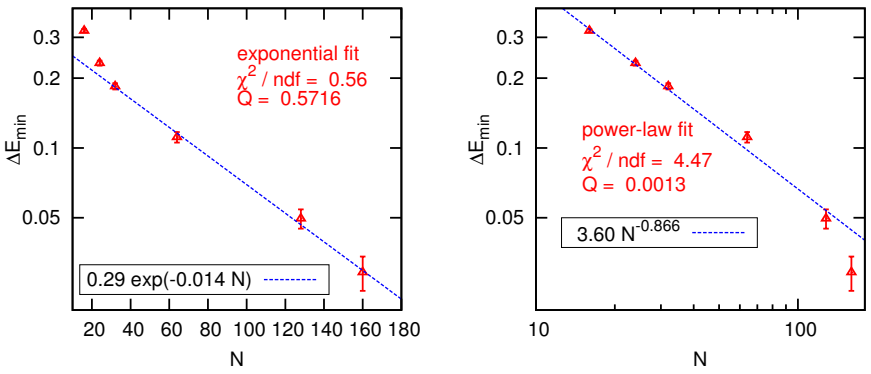


**Figure 3.** (Color online) Left panel is the median minimum gap on a log-lin scale for the spin glass problem. The quality of the straight-line fit is good, ($Q = 0.57$), indicating that the minimum gap decreases exponentially with $N$. (The two smallest sizes are omitted from the fit since they seem to be affected by corrections to scaling.) A power-law fit works less well as shown in the right hand panel.

of coupling to a decohering environment. One could then ask whether other, perhaps more appropriate, implementations of the QAA could lead to a better complexity scaling and what is the effort that would be involved in finding such variants of the QAA.

There are many ways in which one could generalize and expand on the QAA approach taken here to possibly greatly increase the success probability of the algorithm, or equivalently, shorten its required runtime. To mention a few examples, it is now known that a carefully chosen annealing schedule could improve the error-scaling of the algorithm [46] and/or yield a better scaling of the runtime with respect to the minimum gap [47]. Other studies show that taking different paths in 'Hamiltonian space' between the initial and final Hamiltonians may also prove beneficial [31, 48]. Finally, it has been shown that in some cases, thermal effects (i.e., nonzero temperature) that are normally thought of having destructive effects, populating unwanted excited states, may in fact assist in the annealing process by providing an additional mechanism for traversing energy barriers [49].

The quantum adiabatic algorithm studied here did not utilize any a priori knowledge about the SAT or spin-glass instances being solved in the construction of the algorithm (as both the driver Hamiltonian and the annealing schedule had been determined independently of those beforehand). It is well known that for some problems, the utilization of such knowledge in the implementation of the QAA may prove advantageous. One notable example is the oracular problem of searching for a marked item in an unstructured database addressed by Roland and Cerf [47] later extended to other problems [26, 27], which showed that tailoring the driver Hamiltonian as well as the adiabatic schedule to the problem at hand (which for these problems translated to an explicit calculation of the gap as a function of the adiabatic parameter) may lead to quantum speedups.

In satisfiability problems, one can not compute the value or location of the minimum gap beforehand as these vary between instances and doing so for each instance is likely to be as difficult as solving the problem itself. However, it is plausible to assume that one could utilize the fact that these Hamiltonians are "frustration-free" [50], i.e., that the solutions of the total Hamiltonian are also solutions of the individual clauses, to gain a favorable scaling of the gap. One approach would be to consider for each instance the class of problem Hamiltonians

$$\mathcal{H}_p = \sum_{a=1}^{M} c_a \mathcal{H}_a \,, \tag{10}$$

where $c_a$ are positive real-valued weights. That the solutions of the original satisfiability problem are also solutions of the individual clauses ensures that for any given set of instances $\{\mathcal{H}_a\}$, the ground states of the entire class of problem Hamiltonians are solutions of the original problem as well, and one therefore has the freedom to choose these positive weights as deemed fit. The above is also true for the family of driver Hamiltonians

$$\mathcal{H}_D = -\sum_i \Gamma_i \sigma_i^x \,, \tag{11}$$

where different values for the nonzero weights $\Gamma_i$ may also be considered.

An appropriate choice of values for the weights $c_a$ and/or $\Gamma_i$ could be sought by 'blind' trial and error, in which case a polynomial number of QAAs with different weight sets is attempted and the configuration of weights producing the shortest runtime (or largest minimum gap) is then picked out. A presumably more efficient method would be to adjust the weight values adaptively after running the QAA and examining the obtained bit configuration and the clauses it violated. Assigning larger weights to the violated clauses will render them more unlikely to be violated in the next round, thereby presumably also making the minimum gap larger. It should be noted that this type of techniques can be straightforwardly tested using the QMC techniques used here, however it should be emphasized that these are worthy of implementing only if the computational effort involved in adjusting the weights turns out to be smaller than the computational gain associated with the growth of the minimum gap.

## 4. Conclusions

We have studied the efficiency of the QAA in its most basic implementation, when applied to several optimization problems, ignoring the effects of finite temperature and non-thermal noise which occur in experiments on the D-Wave hardware. We used Quantum Monte Carlo simulations to compute the minimum gap and then took the run time to be proportional to one over the square of the minimum gap. For each problem we found that the QAA was not more efficient than a classical heuristic solver, WalkSAT, in the sense that the coefficient $\mu$ in Eq. (1) was a little larger for the QAA than for WalkSAT. We also discussed possible modifications to the QAA that might improve the situation.

## Acknowledgments

We thank Eddie Farhi and Vadim Smelyanskiy for many stimulating discussions about quantum annealing. This work is supported in part by the National Science Foundation under Grant No. DMR-DMR-1207036.

## References

[1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge, England: Cambridge University Press, 2000.
[2] N. D. Mermin, *Quantum Computer Science*. Cambridge: Cambridge University Press, 2007.
[3] P. W. Shor, "Algorithms for quantum computing: discrete logarithms and factoring," in *Proc. 35th Symp. on Foundations of Computer Science* (S. Goldwasser, ed.), (Los Alamitos CA), p. 124, IEEE Computer Society Press, 1994.
[4] T. Kadowaki and H. Nishimori, "Quantum annealing in the transverse Ising model," *Phys. Rev. E*, vol. 58, p. 5355, 1998.
[5] G. Santoro, R. Martoňák, E. Tosatti, and R. Car, "Theory of quantum annealing of an Ising spin glass," *Science*, vol. 295, p. 2427, 2002.
[6] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, "A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem," *Science*, vol. 292, p. 472, 2001. A longer version of the paper appeared in arXiv:quant-ph/0104129.
[7] G. Toulouse, "Theory of the frustration effect in spin glasses," *Commun. Phys.*, vol. 2, p. 115, 1977.

[8] K. Binder and A. P. Young, "Spin glasses: Experimental facts, theoretical concepts and open questions," *Rev. Mod. Phys.*, vol. 58, p. 801, 1986.
[9] S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, p. 671, 1983.
[10] M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics*. New York, USA: Oxford University Press Inc., 1999.
[11] S. Boixo et. al., "Quantum annealing with more than one hundred qubits," *Nature Phys.*, vol. 10, p. 2018, 2014.
[12] A. Messiah, *Quantum Mechanics*. Amsterdam: North-Holland, 1962.
[13] G. H. Wannier, "Probability of violation of the Ehrenfest principle," *Physics*, vol. 1, p. 251, 1965.
[14] E. Farhi, J. Goldstone, and S. Gutmann, "Quantum adiabatic algorithms versus simulated annealing." (arXiv:quant-ph/0201031), 2002.
[15] M. W. Johnson *et al.*, "Quantum annealing with manufactured spins," *Nature*, vol. 473, pp. 194–198, 2011.
[16] A. J. Berkley *et al.*, "Tunneling spectroscopy using a probe qubit," *Phys. Rev. B*, vol. 87, p. 020502(R), 2013.
[17] T. F. Rønnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer, "Defining and detecting quantum speedup," *Science*, vol. 345, no. 6195, pp. 420–424, 2014.
[18] A. M. Childs, E. Farhi, and J. Preskill, "Robustness of adiabatic quantum computation," *Phys. Rev. A*, vol. 65, p. 012322, 2001.
[19] M. Amin, D. V. Averin, and J. A. Nesteroff, "Decoherence in adiabatic quantum computation," *Phys. Rev. A*, vol. 79, p. 022107, 2009. (arXiv:0708.0384).
[20] L. Wang et. al., "Comment on "Classical signature of quantum annealing"." (unpublished), 2013.
[21] A. P. Young, S. Knysh, and V. Smelyanskiy, "First order phase transition in the Quantum Adiabatic Algorithm," *Phys. Rev. Lett.*, vol. 104, p. 020502, 2010.
[22] I. Hen and A. P. Young, "Exponential complexity of the quantum adiabatic algorithm for certain satisfiability problems," *Phys. Rev. E.*, vol. 84, p. 061152, 2011.
[23] E. Farhi, D. Gosset, I. Hen, A. Sandvik, P. Shor, A. P. Young, and F. Zamponi, "The performance of the quantum adiabatic algorithm on random instances of two optimization problems on regular hypergraphs," *Phys. Rev. A*, vol. 86, p. 052334, 2012.
[24] I. Hen and A. P. Young, "Solving the graph-isomorphism problem with a quantum annealer," *Phys. Rev. A.*, vol. 86, p. 042310, 2012.
[25] I. Hen, "Excitation gap from optimized correlation functions in quantum monte carlo simulations," *Phys. Rev. E.*, vol. 85, p. 036705, 2012.
[26] I. Hen, "Continuous-time quantum algorithms for unstructured problems," *J. Phys. A*, vol. 47, p. 045305, 2014.
[27] I. Hen, "How fast can quantum annealers count?," *J. Phys. A*, vol. 47, p. 235304, 2014.
[28] A. Sandvik, "Stochastic series expansion method with operator-loop update," *Phys. Rev. B*, vol. 59, p. R14157, 1999.
[29] A. Sandvik, "A generalization of Handscomb's quantum Monte Carlo scheme — Application to the 1-D Hubbard model," *J. Phys, A*, vol. 25, p. 3667, 1992.
[30] A. P. Young, S. Knysh, and V. Smelyanskiy, "Size dependence of the minimum excitation gap in the Quantum Adiabatic Algorithm," *Phys. Rev. Lett.*, vol. 101, p. 170503, 2008.
[31] E. Farhi, J. Goldstone, D. Gosset, S. Gutmann, H. B. Meyer, and P. Shor, "Quantum adiabatic algorithms, small gaps, and different paths," *Quant. Inf. Comp.*, vol. 11, p. 181, 2009. (arXiv:0909.4766).
[32] F. Krzakala, A. Rosso, G. Semerjian, and F. Zamponi, "On the path integral epresentation for quantum spin models and its application to the quantum cavity method and to Monte Carlo simulations," *Phys. Rev. B*, vol. 78, p. 134428, 2008.
[33] A. Sandvik, "Stochastic series expansion method for quantum Ising models with arbitrary interactions," *Phys. Rev. E*, vol. 68, p. 056701, 2003.
[34] K. Hukushima and K. Nemoto, "Exchange Monte Carlo method and application to spin glass simulations," *J. Phys. Soc. Japan*, vol. 65, p. 1604, 1996.
[35] L. Zdeborová and M. Mézard, "Locked constraint satisfaction problems," *Phys. Rev. Lett.*, vol. 101, p. 078702, 2008.
[36] L. Zdeborová and M. Mézard, "Constraint satisfaction problems with isolated solutions are hard," *J. Stat. Mech.*, vol. 12, p. P12004, 2008.
[37] T. Jörg, F. Krzakala, G. Semerjian, and F. Zamponi, "First-order transitions and the performance of quantum algorithms in random optimization problems," *Phys. Rev. Lett.*, vol. 104, p. 207206, 2010. (arXiv:0911.3438).
[38] S. Kirkpatrick and B. Selman, "Critical behavior in the satisfiability of random Boolean expressions," *Science*, vol. 264, p. 1297, 1994.

[39] C. K. Thomas and H. G. Katzgraber, "Optimizing glassy p-spin models," *Phys. Rev. E*, vol. 83, p. 046709, 2011.
[40] M. Guidetti and A. P. Young, "Complexity of several constraint satisfaction problems using the heuristic, classical, algorithm, WalkSAT," *Phys. Rev. E*, vol. 84, p. 011102, 2011. (arXiv:1102.5152).
[41] For information about WALKSAT see `http://www.cs.rochester.edu/~kautz/walksat/`.
[42] T. Jörg, F. Krzakala, J. Kurchan, and A. C. Maggs, "Simple glass models and their quantum annealing," *Phys. Rev. Lett.*, vol. 101, p. 147204, 2008.
[43] T. Jörg, F. Krzakala, J. Kurchan, A. C. Maggs, and J. Pujos, "Energy gaps in quantum first-order mean-field-like transitions: The problems that quantum annealing cannot solve," *Europhys. Lett.*, vol. 89, p. 40004, 2010. (arXiv:0912.4865).
[44] S. F. Edwards and P. W. Anderson, "Theory of spin glasses," *J. Phys. F*, vol. 5, p. 965, 1975.
[45] Y. Matsuda, H. Nishimori, and L. Zdeborova, "Random-field p-spin-glass model on regular random graphs," *J. Phys. A: Math. and Theor.*, vol. 44, p. 185002, 2007.
[46] N. Wiebe and N. S. Babcock, "Improved error-scaling for adiabatic quantum evolutions," *New Journal of Physics*, vol. 14, no. 1, p. 013024, 2012.
[47] J. Roland and N. J. Cerf, "Quantum search by local adiabatic evolution," *Phys. Rev. A*, vol. 65, p. 042308, 2002.
[48] A. Perdomo-Ortiz, S. E. Venegas-Andraca, and A. Aspuru-Guzik, "A study of heuristic guesses for adiabatic quantum computation," *Quantum Information Processing*, vol. 10, p. 33–52, mar 2010. n/a.
[49] N. Dickson *et al.*, "Thermally assisted quantum annealing of a 16-qubit problem," *Nature Comm.*, vol. 4, p. 1903, 2013.
[50] S. Bravyi and B. Terhal, "Complexity of stoquastic frustration-free hamiltonians," *SIAM J. Comput.*, vol. 39, pp. 1462–1485, Nov. 2009.