Spanning Tree Calculations on D-Wave 2 Machines

M.A. Novotny[†], L. Hobl[‡], J.S. Hall[†], and K. Michielsen[‡]

[†] Department of Physics and Astronomy, and Center for Computational Sciences, Mississippi State University, Mississippi State, MS 39762-5167, USA

[‡] Institute for Advanced Simulation, Jülich Supercomputing Centre, Forschungszentrum Jülich, D-52425 Jülich, Germany, and RWTH Aachen University, D-52056 Aachen, Germany

E-mail: `man40@msstate.edu`

Abstract. Calculations on D-Wave machines are presented, both for the 500-qubit and the 1000-qubit machines. Results are presented for spanning trees on the available $K_{4,4}$ Chimera graphs of both machines. Comparing trees of approximately the same size, the frequency of finding the ground state for the 1000-qubit machine is significantly improved over the 500-qubit older generation machine. Spanning trees are difficult problems for solution by adiabatic quantum computers, so the enhanced frequency of finding the ground state for newer machine generations and larger machines is encouraging for this immature technology.

1. Introduction

The availability of an Adiabatic Quantum Computer (AQC) with thousands of qubits would be a disruptive technology. In complexity theory the class P stands for polynomial time problems and NP for non-deterministic polynomial time problems [1]. We assume, although no proof is known, that P≠NP. The reason AQC would be a disruptive technology is that in theory an AQC would enable solutions to problems that belong to the complexity class NP-hard, whereas all classical computers allow solutions to problems of the complexity class P. In theory any NP-hard problem can be reduced in polynomial time to a chosen NP-hard problem. Therefore, a machine that can solve any particular NP-hard problem would be a disruptive technology.

The goal of an AQC is to solve one particular NP-hard problem, namely finding the ground state of an Ising spin glass. The AQC operates at a low temperature, but the functionality of an AQC is easiest understood by considering the zero-temperature dynamics of the qubit system that makes up the AQC. A bit can be represented as a vector of length two, with the binary variable 0 represented as the vector (0 1) and the binary variable 1 represented as the vector (1 0). A qubit can be represented as a vector of length two of the form $(a_1 \, a_2)$ with the normalization $a_1^* a_1 + a_2^* a_2 = 1$ where the a_i are in general complex numbers. In other words, a qubit is a superposition of the two classical binary bit values 0 and 1. An AQC works by initializing the vector of all N qubits as a direct product of the ground state of a Hamiltonian $\mathcal{H}_x$ of the Pauli matrix σ^x. The goal is to find the ground state of the Ising spin glass Hamiltonian $\mathcal{H}_z$, a function of the Pauli matrix σ^z. The Hamiltonian $\mathcal{H}_\lambda = \lambda \mathcal{H}_x + (1-\lambda)\mathcal{H}_z$ is adiabatically adjusted by changing the variable λ with time. At time t=0 one starts with λ=1, and at the end of the annealing cycle one has λ=0. After the annealing cycle the vector is measured to provide a ground state of $\mathcal{H}_z$.

Figure 1. An $L = 2$ $K_{4,4}$ Chimera graph $\mathcal{G}$ is shown. The graph $\mathcal{G}$ is represented as a bi-layer system, with blue nodes in the bottom layer and green nodes in the top layer. Each complete $K_{4,4}$ unit cell has four green nodes and four blue notes, with every blue node connected to every green node. Also shown are three stages of the construction of a spanning tree on $\mathcal{G}$. The initial graph with the chosen tree root (red), the tree with seven of the of the 27 sites occupied by the growing tree, and the spanning tree with all 27 sites occupied.

Obtaining the ground state of an Ising spin glass is an NP-hard problem. Consider a connected undirected graph $\mathcal{G}$ with N vertices (nodes) with an Ising variable $s_j = \pm 1$ placed on each vertex. The ground state of the Ising spin glass is the state $\{s\}$ of Ising spins on $\mathcal{G}$ that minimizes the energy function [2]

$$E = -\sum_{i=1}^{N-1}\sum_{j=i+1}^{N} J_{i,j} s_i s_j - \sum_{j=1}^{N} h_j s_j \tag{1}$$

for fixed bias fields h_j and fixed bond weights $J_{i,j}$. If there is no edge in $\mathcal{G}$ connecting vertices i and j then $J_{i,j}=0$. The $J_{i,j}$ and h_j are the parameters in the Hamiltonian $\mathcal{H}_z$. If all bonds are ferromagnetic, $J_{i,j} \geq 0$, and all bias fields are zero, the ground state is degenerate, as both the state with all $s_j = +1$ and the state with all $s_j = -1$ have the same energy. An Ising spin glass requires that there is frustration in the model, in other words that it is not possible to place all spins on $\mathcal{G}$ in such a fashion that every term in the sums in Eq. (1) is positive.

If P$\neq$NP, then in the worst-case instance the solution of an NP-problem requires an exhaustive search of all possible states. For the Ising spin glass an exhaustive search requires calculating the energy in Eq. (1) 2^N times. This can be easily programmed in parallel for a Single-Instruction-Multiple-Data (SIMD) classical computer. The steps of the SIMD algorithm, assuming 2^N processors, are:

- Broadcast the same $J_{i,j}$ bond weights and h_j bias fields to each processor
- Scatter one configuration of the Ising variables s_j to each processor
- Each processor calculates the energy of Eq. (1) for its configuration of the Ising variables
- Gather the energies from each processor, and keep only the configuration(s) with the lowest energies.

The comparison (last) step can be performed with the order of the number of operations $\mathcal{O}(N)$. Therefore the number of computations required is given by the rate limiting step, the third step. The number of calculations which should be performed is $\mathcal{O}\left(N^2 2^N\right)$, since each energy calculation requires $\mathcal{O}\left(N^2\right)$ computations.

Figure 2. Configuration of the D-Wave 2X used to study spanning trees on 1000+ qubits. [Left] Each square is a $K_{4,4}$ unit cell of eight qubits. The chip is a 12×12 square array of unit cells (with free boundary conditions). The colors correspond to the number of functioning qubits in the unit cell: black=8, gray=7, pink=6, and yellow=5. [Right] The D-Wave 2X graph for the $K_{4,4}$ Chimera graph $\mathcal{G}$ used in this study, shown as a bilayer system. The blue qubits are in the bottom layer and the green qubits in the top layer. The chip has 1097 available qubits (1097 vertices in $\mathcal{G}$). The qubit labeled as zero on the D-Wave 2X is shown in red (top left).

The company D-Wave, Inc [3] recently announced a 1000-qubit machine [4]. The reason the SIMD algorithm above cannot be used in principle is that $2^{1000} \approx 10^{301}$, while the estimated number of particles in the visible universe is 'only' about 10^{80}. Clearly the SIMD algorithm is impracticle for 1000 vertices in $\mathcal{G}$. An ideal AQC can perform all 2^N calculations of the energy in parallel because of principles of quantum mechanics, in particular coherence, superposition, and tunneling [5]. Quantum mechanics is probabilistic, and hence any AQC is a probabilistic machine. An ideal AQC is possible only in the unphysical limits of the machine operating at zero degrees Kelvin (absolute zero) and taking an infinite time to perform the required adiabatic calculation. Therefore the question that needs to be addressed for any practical AQC is whether it provides groundstates with a high enough frequency of occurrence that it is useful to calculate NP-hard problems.

This conference proceedings article, based on a Sept. 9, 2015 plenary talk at the *International Conference on Computer Simulation in Physics and beyond* in Moscow, Russia, seeks to address the practical question of applicability of D-Wave machines. The results of that talk, as well as of this paper, include only preliminary results for the 1000 qubit D-Wave 2X machine due to the short amount of time the machine has been available. Nevertheless, even the preliminary results can shed light on the functionality of the D-Wave 2X.

The approach we have taken to test the D-Wave machines is to use spanning trees on the D-Wave graph $\mathcal{G}$. Section 2 details the graph $\mathcal{G}$ of the D-Wave machines. Section 3 presents results of placing ferromagnetic bonds on the 1000 qubit D-Wave 2X. Section 4 describes spanning trees on D-Wave machines. Results of solving spanning tree problems on D-Wave machines are presented in Section 5 for ferromagnetic bonds and in Section 6 for random ±1 bonds. Section 7 contains a discussion, outlook, and conclusions.

2. D-Wave Machines

The architecture of the D-Wave chip is an $L \times L$ $K_{4,4}$ Chimera lattice [3, 4]. The number of qubits on the chip is $N = 8L^2$. The Chimera lattice can be represented as a bilayer system, as shown in Fig. 1 for an $L=2$ lattice. Some of the qubits on a chip are not accessible to the programmer, so the $K_{4,4}$ Chimera lattice has disorder, just as depicted in Fig. 1.

In Fig. 1 the graph $\mathcal{G}$ is taken to be a $K_{4,4}$ Chimera lattice, composed of a 2×2 arrangement of unit cells. The graph is represented as a bilayer system. Each unit cell has four vertices in the top layer (green vertices) and four vertices in the bottom layer (blue vertices). In each unit cell the top four vertices are each connected (brown edges) to every one of the four vertices in the bottom layer. The vertices in the bottom layer are connected (blue edges) to corresponding bottom layer vertices in nearest neighbor unit cells. Similarly, the vertices in the top layer are connected (gray edges) to corresponding top layer vertices in nearest neighbor unit cells. The graph $\mathcal{G}$ in Fig. 1 has disorder, as only 27 of the possible $2^2 8 = 32$ vertices are present. If a vertex from the full $K_{4,4}$ Chimera lattice is missing, all of its associated edges are also missing in $\mathcal{G}$. The D-Wave machines have the same $K_{4,4}$ structure as in Fig. 1, but are on an 8×8 arrangement of unit cells for the 500 qubit machine (the total number of qubits is $8^3 = 512$) and on a 12×12 arrangement of unit cells for the 1000 qubit machine (the total number of qubits is $12^2 8 = 1152$).

Results from three different D-Wave chips will be reported. Two of the chips are last generation 512 qubit chips. The number of functioning qubits is 496 of one chip [7] and of the other is 476 qubits. Both chips were operated at the company headquarters in Burnaby, British Columbia, Canada, and were so-called in-house 'test' chips. The 1152 qubit D-Wave 2X chip actually had 1097 accessible qubits. This computer is located at the Ames Research Center located at Moffett Field in California, USA. The D-Wave machine is managed by a collaboration of NASA (National Aeronautics and Space Administration) [8], USRA (Universities Space Research Association) [9], and Google [10].

The configuration of the 1000 qubit D-Wave 2X is shown in Fig. 2. Of the $L^2 = 144$ $K_{4,4}$ unit cells, 109 have all eight qubits available. The details of the number of available qubits in the unit cells are in Table 1. Any qubit j which is not available results in all the $J_{i,j}$ connections not available. In addition, there is one unavailable $J_{i,j}$ connection between two available qubits. This makes that there are 3060 bonds available from the 3360 bonds in an undisordered $12^2 K_{4,4}$ Chimera lattice. The bilayer representation of the graph of the 1097 qubit machine is shown in Fig. 2. In Fig. 2(Right) the qubit in red (top left) marks the qubit that is assigned to be qubit number zero. The D-Wave 2X is an analog machine, and has a 3% analog control error on the h_j and the $J_{i,j}$.

Total qubits in a unit cell	8	7	6	5
144	109	18	14	3

Table 1. The number of available qubits in each unit cell of the 1000 qubit D-Wave 2X.

3. Ferromagnetic Computations on a 1000 Qubit D-Wave Machine

The first test we performed on the 1097 qubit D-Wave 2X was to place $J_{i,j}=1$ on all available bonds, resulting in overall ferromagnetic bonds and to put the same bias fields h_j on each qubit. One thousand anneals were performed for each value of h_j (set to be uniform). All anneals found the expected ground state, each having the energy $E=-3060$. For zero field the two ground states with all $s_j=+1$ and $s_j=-1$ should be found equally often. However, as shown in Table 2 only one of the two ground states was found. By gradually changing the uniform value of the bias field h_j on each qubit, for the bias field $h_j=0.07$ both ground states were found. For

the D-Wave 2X there is a 3% analog control error, so quantitatively the results in Table 2 are reasonable.

h_j	0.00	0.06	0.07	0.08
$s_j{=}{-}1$	1000	1000	823	0
$s_j{=}{+}1$	0	0	177	1000

Table 2. The number of times each of the two ground states were found for different applied fields h_j for ferromagnetic $J_{i,j}{=}1$ bonds.

4. Spanning Trees

A tree is a graph with no loops. A spanning tree of $\mathcal{G}$ is a subgraph of $\mathcal{G}$ that has a tree structure and that includes all N vertices of $\mathcal{G}$. We have generated our spanning trees using an algorithm of A. Broder [6]. The algorithm has an expected run time of $\mathcal{O}(N\log N)$ for almost all graphs, and a run time per generated tree of $\mathcal{O}(N^3)$ for the worst graphs. Three steps during a run of the algorithm are sketched in Fig. 1.

The algorithm to find a spanning tree is

- Choose one vertex of $\mathcal{G}$ uniformly at random. In Fig. 1(left) the chosen vertex is red. Let this vertex be labeled by j.
- Repeat the steps below until all N vertices belong to the tree:
    - Pick an edge uniformly at random that leaves the vertex j
    - Walk along the chosen edge to a vertex labeled as k
    - If this is the first time vertex k has been visited, add vertex k and the edge j,k to the growing tree. In Fig. 1(center) this is shown by coloring the vertices and edges red.
    - Change the label j to the label k, whether or not the vertex k was visited for the first time.

In other words, a random walk starting at any vertex is generated, and the first time a vertex is visited it is added to the tree and the edge that led to this first visit of the vertex is added to the tree. The algorithm stops when all vertices belong to the tree [see Fig. 1(right)]. Any spanning tree has $N-1$ edges. The spanning trees generated by this algorithm, for example the spanning tree in Fig. 1, are those drawn uniformly from the distribution of all spanning trees of $\mathcal{G}$ [6].

One way to test an AQC machine is to give problems to the machine for which one knows the answer. This strategy has been followed by others in terms of chains of qubits [11] or frustrated optimization problems with tunable hardness [12]. The spanning tree planted graphs have the advantage that:

- All vertices are included in every spanning tree.
- The ensemble of spanning trees covers all edges.
- The ensemble of spanning trees is well defined.
- Spanning trees exist for any connected graph $\mathcal{G}$.
- The ground state for a spanning tree is known. For example, if all $h_j{=}0$ and all $J_{i,j}{=}{\pm}1$ in the tree, then the ground state energy is the negative of one less than the number of edges in the spanning tree, $E = -N + 1$.
- The spin arrangement of the ground state of the spanning tree is easy to calculate.
- Obtaining a ground state for a spanning tree problem is a difficult problem for an AQC.
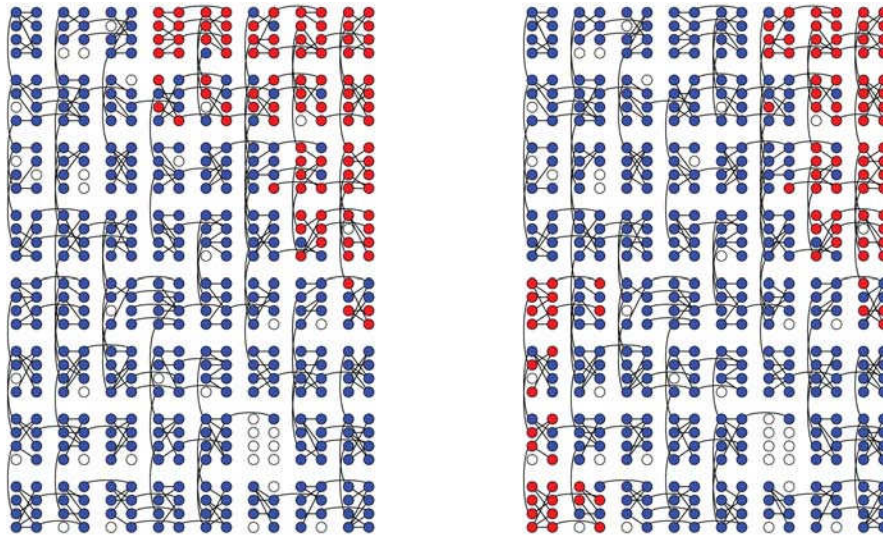
Figure 3. Two solutions of a ferromagnetic spanning tree on the 476 qubit D-Wave 2. The 8×8 arrangment of unit cells are represented with all qubits in a single plane, with blue qubits $+1$ and red qubits -1. All available bonds have non-zero weight $J_{i,j}=1$, and only these bonds are drawn (black lines and curves). The ground state should have all qubits in the same state (all qubits should be either blue or all should be red). In order to bias which ground state has the minimum energy, a non-zero bias field $h_{452}=-1$ was placed on qubit number 452 (at the bottom-left), while all other bias fields were zero. This bias should have forced all qubits to be $s_j=-1$ (red). This is not true for either solution of this spanning tree.

Some results for spanning trees on a 500 qubit D-Wave machine have recently been reported [7]. There are two disadvantages of using spanning trees to test AQC machines. The trees do not contain any loops, and hence there is no frustration in the model, in other words the problem is in P not NP. The NP-hard problems that an AQC should solve contain P problems like the spanning tree, but they are designed to provide answers to the NP-hard problem of finding the ground state of the Ising spin glass. Nevertheless, the spanning trees give a feeling for how much to trust the answers obtained from an AQC.

5. Ferromagnetic Spanning Trees on D-Wave Machines

On the 476 chip we placed one spanning tree. Figure 3 shows the spanning tree on the 476 qubit chip, here represented with all qubits in a single plane. The 8×8 arrangement of unit cells is shown, with only bonds in the spanning tree shown. Note that the spanning tree graph as represented here is not drawn in a planar fashion, even though any spanning tree can be drawn as a planar graph with non-crossing bonds. The $K_{4,4}$ Chimera graph is shown in a way the D-Wave 2 presents the qubits on the web interface. All bonds were ferromagnetic, $J_{i,j}=+1$.

We examined the number of times each energy E was found following the anneal cycle. All bias fields were set to zero. A total of 10^3 submissions, each with 10^3 anneals, was performed. The ground state, with energy $E_{\mathrm{gs}}=-475$, was only found 237 times out of the 10^6 anneals. Fig. 4 shows the occurrence of each energy for this particular spanning tree.

Since all bias fields were set to zero, the ferromagnetic tree for the 476 qubit chip should find the two ground states the same number of times. In that case the magnetization averaged over all anneals should be $\langle m_j \rangle=0$. Alternatively if only one state should be found, as in the ferromagnetic case of Sec. 4, then for the ground state all qubits should have the same $\langle m_j \rangle$. For

Figure 4. Occurrence in % of the energy found for a 476 qubit D-Wave 2 for a ferromagnetic spanning tree. The tree was submitted 10^3 times each with 10^3 anneals. The ground state has an energy $E_{\mathrm{gs}} = -475$ and has a relative occurrence frequency of 0.02%.

the 476 qubit machine, as seen in Fig. 5, neither of these assumptions are true. The behavior of $\langle m_j \rangle$ may be understood as follows. The analog nature of the D-Wave machine may leave a non-zero field h_j on each qubit even when the h_j value has been set to zero. This slightly different from zero field for each qubit is sufficient to decrease the number of times the ground state is found.

In order to have a non-degenerate ground state, we set all bias fields to $h_j = 0$, except for qubit number 452, which was set to $h_{452}=-1$. The bias field should have enabled the D-Wave 2 to find the ground state with all $s_j=-1$ (red in Fig. 3). For this chip and spanning tree, not even the qubit number 452 was forced to be in the $s_j=-1$ (red) state. Neither the ground state energy ($E_{\mathrm{gs}}=-476$) nor the first excited state was returned by the D-Wave machine. Fig. 3 shows two representative returned low-energy configurations, both of a second excited state (energy $E_2=-472$). Fig. 3 shows that this chip seems to be inhomogeneous, since $h_{452}=-1$ did not bias it to find the non-degenerate ground state. This chip may not have been fine tuned, because it was an in-house 'test' chip.

The analog nature of the D-Wave 2 presents certain issues in terms of programming the machine. One question one could ask is whether setting bonds to zero disconnects the two qubits. Because of the 3% analog setting, the degree of isolation should be tested. We have considered the 476 qubit chip to be four $L=4$ $K_{4,4}$ Chimera lattices. A missing qubit on any $L=4$ Chimera quadrant was replicated by setting the $J_{i,j}=0$, leading to four identical graphs $\mathcal{G}_{L=4}$ on each quadrant. The same ferromagnetic spanning tree was placed on each quadrant. As seen in Fig. 6 each quadrant was found to obtain the ground state a different fraction of time. The success frequency of each quadrant was labeled as P_1, P_2, P_3, and P_4. Assuming the success frequency of finding the ground state in each quadrant was independent of the ground state found in the other quadrants, the frequency of simultaneously finding the ground state in quadrants 1 and 4 would be

$$P_{1\,4} = P_1 \, P_4 \,. \tag{2}$$

Similarly for any two quadrants. Fig. 6 shows results for $P_{1\,2}$, $P_{1\,3}$, $P_{1\,4}$, $P_{2\,3}$, $P_{2\,4}$, and $P_{3\,4}$. The

Figure 5. The average magnetization $\langle m_j \rangle$ for a 476 qubit D-Wave 2 for a ferromagnetic spanning tree. The tree was submitted 10^3 times each with 10^3 anneals. For a system with no bias, the average magnetization should be zero.

measured frequency of success (red bars in Fig. 6) correspond within the measurement errors to the predicted results. Similarly for simultaneously finding the ground state in three quadrants, labeled $P_{1\,2\,3}$, $P_{1\,2\,4}$, $P_{1\,3\,4}$, and $P_{2\,3\,4}$. The frequency of simultaneously finding the ground state in all four quadrants is $P_{1\,2\,3\,4}$. If the probability of finding a ground state in each quadrant is independent, $P_{1\,2\,3\,4} = P_1 P_2 P_3 P_4$. Fig. 6 compares the measured frequencies (red bars) to the expected probabilities obtained by multiplying the frequency of finding the ground state in each quadrant (right yellow bars). The conclusion is that within the measurement errors the ground state in each quadrant is found to be independent of what is occuring in the other quadrants.

6. Spanning Tree Computations with ± 1 Bonds on D-Wave Machines

We performed spanning tree calculations on $L \times L$ $K_{4,4}$ Chimera subgraphs of two different chips, the one with 496 qubits and the one with 1097 qubits. The bond weights were uniformly randomly chosen to be $J_{i,j} = \pm 1$. All magnetic fields h_j were set to zero. The $L \times L$ subgraphs were chosen to all include the $K_{4,4}$ unit cell in the upper-left corner of the chip. This choice is similar to the gauge transformations used by others [13] to remove local biases on the chip. Results for the 496 qubit machine were published previously [7]. We submitted a number of spanning trees for each L, and each submission requested 10^3 ground states (10^3 anneal cycles for the given tree and $J_{i,j}$ values). Each tree was submitted with different $J_{i,j}$ values up to 10^2 times. The probability the D-Wave 2 returned the ground state, p_{gs}, was estimated by the frequency of occurrence of success from all anneals. Results are shown in Table 3 and Fig. 7.

The D-Wave machine can be modeled to a first approximation as an uncorrelated probabalistic machine. Let the random variaible be $X = 1$ if the ground state solution is returned in a particular anneal and $X = 0$ otherwise. The expectation value for X is

$$\mathcal{E}\left(X\right) \; = \; 1 \cdot p_{gs} + 0 \cdot (1 - p_{gs}) \; = \; p_{gs}\,. \tag{3}$$
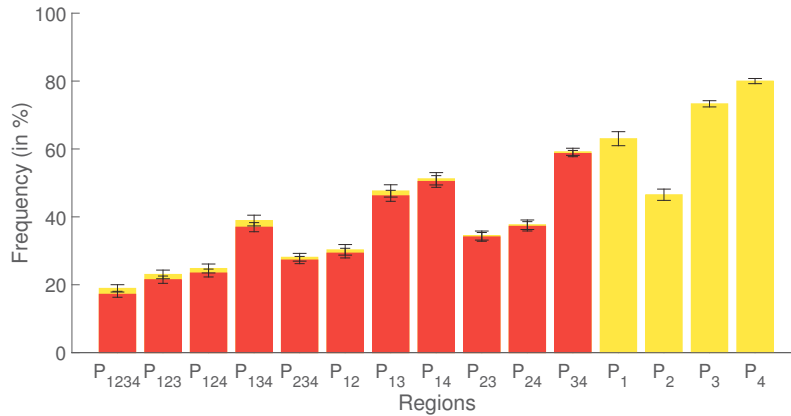
Figure 6. Results on the 476 qubit D-Wave 2. The success frequencies of finding the ground state for each $\mathcal{G}_{L=4}$ quadrant is labeled as P_1 through P_4 (yellow bars on right). The success frequency of simultaneously finding the ground state in two quadrants m and m' is labeled as $P_{m\,m'}$ (red middle bars) and should be equal to $P_{m\,m'} = P_m P_{m'}$ (the prediction of the superimposed yellow bars) if the solution frequencies are independent of each other. Similarly for simultaneously finding the ground state for three quadrants, and for all four quadrants (red left bar labeled P_{1234}).

Similarly, $\mathcal{E}\left(X^2\right) = p_{\mathrm{gs}}$. Therefore the standard deviation σ is

$$\sigma^2 \;=\; \mathcal{E}\left(X^2\right) - \left[\mathcal{E}\left(X\right)\right]^2 \;=\; p_{\mathrm{gs}}\left(1 - p_{\mathrm{gs}}\right) . \tag{4}$$

The standard error from the mean is thus estimated to be $\sigma/\sqrt{M-1}$ when there are M samples. The number of anneals are not uniform for the 496 qubit chip due to program or chip or network issues. There are two ways of deciding on the value of M. One way would be to use $M = M_{\mathrm{anneal}}$ to be all anneal requests for a given L. The other way would be to use $M = M_{\mathrm{tree}}$, the number of different spanning trees submitted for each L. The error estimate with $M = M_{\mathrm{tree}}$, with each tree submitted with up to 10^2 different $J_{i,j} = \pm 1$ and each submitted for 10^3 anneal cycles, is shown in Fig. 7 and listed in Table 3.

Two striking results are present in Fig. 7 and Table 3. The first is that the $8{\times}8$ $K_{4,4}$ Chimera lattice results are very different for the 496 qubit D-Wave 2 and the 1097 qubit D-Wave 2X. A marked improvement in the success frequency of finding the groundstate, p_{gs}, is obtained for the newer generation machine. Some of these differences may be due to the smaller errors in setting the values of the h_j and $J_{i,j}$ on the D-Wave 2X compared with the D-Wave 2.

The second striking result is the small frequency of success for the full graphs, particularly for the 1097 qubit D-Wave 2X. As described above, a spanning tree is a difficult problem for any AQC to solve. That the ground state is found at all for the full graph, even with this low average success frequency is encouraging.

Assume that the probabilities of success on any anneal cycle are independent random variables, and are the same for every tree. Consequently the probability of not finding the ground state is $1 - p_{\mathrm{gs}}$. For M measurements the probability of not finding the ground state is $(1 - p_{\mathrm{gs}})^M$. The probability of finding the ground state after M measurements is

$$P_{\mathrm{gs}} = 1 - (1 - p_{\mathrm{gs}})^M . \tag{5}$$

# qubits	L	# trees	Million anneals	p_{gs}	standard error
496	2	100	10.0	1.0000	0.0
496	3	85	8.4	0.9999	0.0011
496	4	38	3.8	0.9779	0.0242
496	5	100	10.0	0.8953	0.0308
496	6	100	7.9	0.7335	0.0444
496	7	80	10.0	0.5889	0.0554
496	8	100	10.0	0.3261	0.0471
1097	7	100	10.0	0.9944	0.0075
1097	8	100	10.0	0.8443	0.0364
1097	9	100	10.0	0.5800	0.0496
1097	10	100	10.0	0.3444	0.0478
1097	11	100	10.0	0.2017	0.0403
1097	12	100	10.0	0.0443	0.0207

Table 3. Details of the results plotted in Fig. 7. The number of spanning trees as well as the total number of annealing cycles is shown for each $K_{4,4}$ Chimera size, together with the frequency a ground state was found and the calculated standard deviation of the results. The spanning trees all had $h_j=0$ and randomly chosen bond weights $J_{i,j}=\pm 1$.

Thus the number M of measurements required to obtain a given frequency P_{gs} that the ground state was found is

$$M = \frac{\ln(1-P_{gs})}{\ln(1-p_{gs})} . \tag{6}$$

Every submission to the D-Wave by default requests 10^3 anneal cycles. Even for an extremely small average frequency of success $p_{gs}=0.0005$, the number of measurements required to find the ground state 99.99% of the time would require about 18 submissions. However, we find that the 1097 qubit D-Wave 2X probability of successfully finding the ground state is most probably not independent of the particular spanning tree. We have performed 10^6 anneal cycles (10^3 submissions with different $J_{i,j}=\pm 1$) on at least one spanning tree without the D-Wave returning the ground state.

7. Discussion, Outlook, and Conclusions

We have performed tests of D-Wave adiabatic quantum computers (AQCs) by detailing the results of placing spanning trees on the underlying $L\times L$ $K_{4,4}$ disordered Chimera graph of the D-Wave machines. Other researchers have put planted solutions on D-Wave machines [12]. A spanning tree test has a number of advantages over other planted solutions on AQCs. First, spanning trees are a difficult problem for any AQC machine. The other advantages of spanning trees are that they work on any graph, each tree includes all qubits, the ensemble of spanning trees include all couplings $J_{i,j}$, and even for large graphs it is easy to obtain a spanning tree drawn uniformly from the ensemble of all spanning trees. The disadvantage is that spanning trees have no loops, and hence no frustration is present.

The comparison of the two generations of D-Wave machines is encouraging. Spanning trees for 8×8 $K_{4,4}$ Chimera graphs were obtained both on the 512 qubit D-Wave 2 (with 496 functioning qubits) and the 1152 qubit D-Wave 2X (with 1097 functioning qubits). The frequency of finding the ground state for a tree with one anneal increased from $p_{gs}=0.33$ for the D-Wave 2 to $p_{gs}=0.84$ for the D-Wave 2X, both for the $L=8$ Chimera lattice. Some of this increase may be due to the

Figure 7. Success frequency of finding the ground state as a function of the size of the Chimera $K_{4,4}$ subgraph. The data are for the D-Wave 2 with 496 qubits (black squares) and the D-Wave 2X with 1097 qubits (red circles). The error bars for the success frequency are calculated as described in the text. The dashed and solid lines are guides for the eye. The results are tabulated in Table 3.

increased accuracy of setting the h_j and the $J_{i,j}$ parameters on the D-Wave 2X.

AQCs hold the promise that they can find the ground state of an Ising spin glass, and in the unphysical limits of operating at zero degrees Kelvin and taking an infinitely long time for each anneal you are gauranteed to obtain the ground state. Classical computers and advanced algorithms can obtain solutions that are close to the Ising spin glass ground state, and this property has been used to benchmark a D-Wave 2X machine [14]. The property of obtaining *the* global minimum is what, in the authors opinion, makes AQCs a disruptive computer technology. The ability of AQCs to find the solution of one particular NP-hard optimization problem, *the* ground state of an Ising spin glass, means that in principle other NP-hard problems can be solved by mapping them onto an Ising spin glass NP-hard problem. Any (non-ideal) AQC is a probabilistic machine, since it operates using quantum mechanical principles and quantum mechanics is a theory that only allows an observer to make predictions about probabilities of measureable quantities.

In this paper we only deal with spanning trees, so the number of edges between qubits is one, and there are no loops. This could also be called a tree with a tree width of one. Numerical transfer methods have been used to study Ising systems in two and three dimensions [15, 16] as well as in non-integer dimensions [17, 18], including studies of metastable phenomena [19, 20]. It would be possible to utilize numerical transfer matrix methods to study spanning trees with different widths. For example, one could place a tree of width four on the $12^2=144$ square array

of $K_{4,4}$ unit cells shown in Fig. 2. The size of the requried transfer matrix, since every unit cell would connect to its (at most two) neighboring unit cells with four $J_{i,j}$ connections can be obtained by multiplying matrices of size $2^4 \times 2^4$ (after utilizing generalized matrix products to exactly renormalize the tree branches [21, 22]). These types of numerical transfer matrix calculations would allow loops within the subgraphs, and hence can be used to study the model with frustration. Numerical transfer matrix methods have been used previously to study Ising glasses on regular lattices [23, 24, 25, 26] and on hierarchical lattices with non-integer fractal dimensions [26].

The question to be asked is whether the glass is half-full or half-empty in terms of usability of the current generation of D-Wave machines.

The glass-half-full outlook for AQCs is the promise given by Eq. (6). Even for a very low average success frequency, say p_{gs}=0.0005, if you want to find the ground state 99.99% of the time, you only need to submit the tree to the D-Wave 2X eighteen times with the default 10^3 anneals per submission. Hence, AQCs are a disruptive technology as no known classical computer with advanced algorithms could make a similar claim for the $2^{1097} \approx 1.7 \times 10^{330}$ state space of the D-Wave 2X machine. The fact that this paper utilizes a D-Wave machine with more than 1000 qubits while published work on D-Wave machines with about 100 qubits appeared in print in 2014 [11, 13] shows the speed at which AQC is becoming a reality.

The glass half-empty outlook for AQCs is that for the 1097 qubit D-Wave 2X the ground state average success frequency for the spanning trees studied was only found on the 1097 qubit D-Wave 2X with an average frequency of p_{gs}=0.0443. Furthermore, there was at least one spanning tree on the 1097 D-Wave 2X that did not find the ground state after 10^3 submissions (10^6 anneal cycles). One might expect that general spin-glass problems on the D-Wave 2X may have a similar difficulty, namely some ground states are much more difficult for the D-Wave 2X to find than are other ground states. There is still work to do to quantify whether the conclusions based on defining and attempting to detect quantum speedup in a D-Wave machine with 503 qubits [27] hold when applied to a 1000 qubit D-Wave 2X.

Acknowledgements

References
[1] C.H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity* (Prentis-Hall, New Jersey, 1982), Dover edition published 1998.
[2] D.L. Stein and C.M. Newman *Spin Glasses and Complexity* (Princeton University Press, Princeton, NJ, 2013).
[3] D-Wave Systems, Inc, http://www.dwavesys.com
[4] "Quantum chip doubles the qubits", Science, volume 349, page 906 August 28, 2015.
[5] E.G. Rieffel and W.H. Polak, *Quantum Computing: A Gentle Introduction* (MIT Press, Cambridge, MA, 2011).
[6] A. Broder "Generating random spanning trees", in FOCS, pages 442-447 (1989).
[7] J.S. Hall, M.A. Novotny, T. Neuhaus, and K. Michielsen, "A study of spanning trees on a D-Wave quantum computer", *Physics Procedia*, volume 68, pages 56-60 (2015).
[8] NASA Quantum Artificial Intelligence Laboratory (QuAIL), http://www.nas.nasa.gov/quantum/quantumcomp.html
[9] Quantum computing at USRA, http://www.usra.edu/quantum/
[10] Google quantum A.I. lab, https://plus.google.com/+QuantumAILab/posts
[11] K.L. Pudenz, T. Albash, and D.A. Lidar, "Error-corrected quantum annealing with hundreds of qubits", Nature Communications **5**, article 3243 [10 pages] (2014).
[12] I. Hen, J. Job, J. Job, T.F. Røoonw, M. Troyer, and D. Lidar, "Probing for quantum speedup in spin glas problems with planted solutions", arXiv:1502.01663v2.

[13] S. Boizo, T.F. Rönnow, S.V. Isakov, Z. Wang, D. Wecker, D.A. Lidar, J.M. Marinis, and M. Troyer, "Evidence for quantum annealing with more than one hundred qubits", Nature Physics **10**, pages 218-224 (2014).

[14] J. King, S. Yarkoni, M.M. Nevisi, J.P. Hilton, and C.C. McGeoch, "Benchmarking a quantum annealing processor with the time-to-target metric", arXiv:1508.05087.

[15] M.P. Nightingale, "Transfer matrices, phase transitions, and critical phenomena: numerical methods and applications", in *Finite Size Scaling and Numerical Simulations of Statistical Systems*, edited by V. Privman, pages 287-351 (World Scientific, Singapore, 1990).

[16] M.P. Nightingale and H.W.J. Blöte, "Transfer-matrix Monte Carlo estimates of critical points in the simple-cubic Ising, planar, and Heisenberg models", Physical Review B **54**, pages 1001-1008 (1996).

[17] M.A. Novotny, "Numerical transfer matrix study of the Ising model between one and two dimensions", Europhysics Letters **17**, pages 297-302 (1992); Erratum **17**, page 297 (1992).

[18] M.A. Novotny, "Critical exponents for the Ising model between one and two dimensions", Physical Review B **46**, pages 2939-2950 (1992).

[19] C.C.A. Günther, P.A. Rikvold, and M.A. Novotny, "Numerical transfer-matrix study of metastability in the $d=2$ Ising model", Physical Review Letters **71**, pages 3898-3901 (1993).

[20] T. Fiig, B.M. Gorman, P.A. Rikvold, and M.A. Novotny, "Numerical transfer-matrix study of a model with competing metastable states", Physical Review E **50**, pages 1930-1947 (1994)

[21] M.A. Novotny, "Matrix products with applications to classical statistical mechanics", Journal of Mathematical Physics **20**, pages 1146-1150 (1979).

[22] M.A. Novotny, unpublished.

[23] D.A. Huse and I. Morgenstern, "Finite-size scaling study of the two-dimensional Ising spin glass", Physical Review B **32**, pages 3032-3034 (1985).

[24] L. Saul and M. Kardar, "Exact integer algorithm for the two-dimensional $\pm J$ Ising spin glass", Physical Review E **48**, pages R3221-R3224 (1993).

[25] H. Rieger, L. Santen, U. Blasum, M. Diehl, M. Jünger, and G. Rinaldi, "Critical exponents of the two-dimensional Ising spin glass revisited: exact ground-state calculations and Monte Carlo simulations", Journal of Physics A: Mathematical and General **29**, pages 3939-3950 (1996).

[26] R.F.S. Andrade, E. Nogueira, Jr., and S. Coutinho, "Ising sping glass by the transfer matrix approach", Physical Review B **68**, article 104523 [9 pages] (2003).

[27] 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 **345**, pages 420-424 (2014).