

28th Annual CSP Workshop on “Recent Developments in Computer Simulation Studies in Condensed Matter Physics”, CSP 2015

A Study of Spanning Trees on a D-Wave Quantum Computer

J.S. Hall^a, M.A. Novotny^{a,*}, T. Neuhaus^b, Kristel Michielsen^b

^aMississippi State University, 75 B.S. Hood Dr., Mississippi State, MS 39762, USA

^bInstitute for Advanced Simulation, Jülich Supercomputing Centre, Research Centre Jülich, D-52425 Jülich, Germany

Abstract

The performance of a 496 qubit D-Wave Two quantum computer was investigated for spanning tree problems. The chip has a Chimera interaction graph G , an 8×8 lattice of clusters of eight qubits. Problem input consists of values for the fields h_j and for the two-qubit interactions $J_{i,j}$ of an Ising spin-glass problem formulated on G . Output is returned in terms of a spin configuration $\{s_j\}$, with $s_j = \pm 1$. A tree is a connected, undirected subgraph of G that contains no cycles, and a spanning tree is a tree which includes all of the vertices of G . We generated random spanning trees (RSTs), uniformly distributed over all spanning trees of G . One hundred RSTs with random $J_{i,j} = \{-1, 1\}$ and $h_j = 0$ were generated on the full 8×8 graph G of the chip. Each RST problem was solved up to one hundred times and the number of times the ground state energy was found was recorded. This procedure was repeated for square subgraphs G' , thereby providing results for portions of the chip with dimensions ranging from 2×2 to 8×8 .

© 2015 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

Peer-review under responsibility of The Organizing Committee of CSP 2015 Conference

Keywords: D-Wave; Spanning Tree; Chimera; Quantum Computer; Quantum Computing

The goal of this research was to explore ways in which to test the performance of the D-Wave Two [1]. In particular, our goal was to find ways to test and to validate the performance as the number of qubits are scaled in future generations of adiabatic quantum computers. It is envisioned that, in the future, the number of qubits will exceed the limits of simulation and comparison on a classical supercomputer. Random spanning tree problems [2][3][4] are a good candidate for such validation and testing. The ground state energy and spin configuration are known, and thus the solutions returned by the D-Wave Two can be checked without the use of classical supercomputers. Furthermore, every spanning tree incorporates every vertex of the graph, thereby testing every qubit, and the ensemble of all random spanning trees utilizes every edge of the graph, thereby testing all the couplers

* Corresponding author. Tel.: +1-662-325-2688; fax: +1-662-325-8898.

E-mail address: man40@ra.msstate.edu

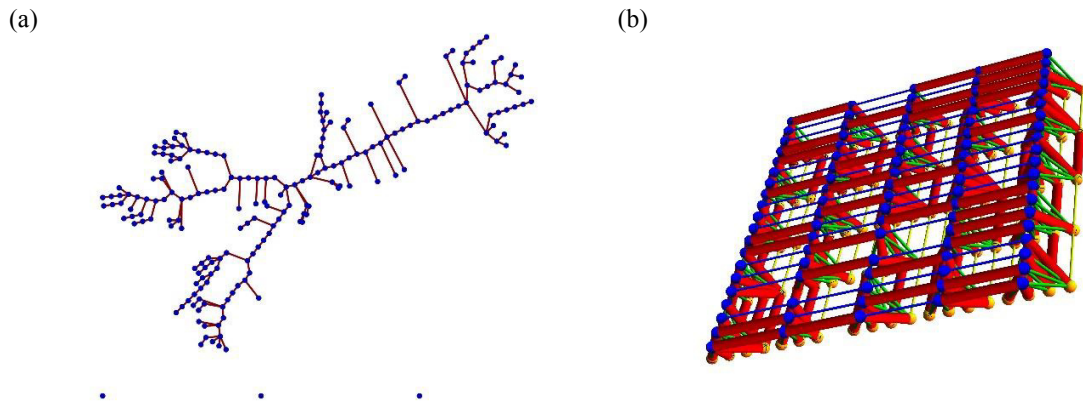


Fig. 1. (a) A spanning tree on a 5x5 Chimera lattice; (b) The same spanning tree (thick connections) embedded into the 5x5 subgraph of the Chimera lattice of the D-Wave Two.

between the qubits. Because of this, these types of problems can be used as a common test for different sizes of lattices; additionally, the spanning tree approach works for any graph G . For smaller chips, the results of experiments on the D-Wave Two can be compared with simulations on classical computers, namely simulated annealing and simulated quantum annealing [5].

The D-Wave Two utilizes what has been named the Chimera architecture [6][7][8]. Clusters of eight qubits are arranged into rows and columns, and the interactions among the qubits within each cluster are described by a complete bipartite graph $K_{4,4}$. That is, each cluster can be further divided into two groups of four; we can name these two groups the red group and the blue group. Every qubit in the red group interacts with every qubit in the blue group, but there are no interactions among members of the same group. Furthermore, each red qubit in a cluster is connected with corresponding red qubits in the neighboring rows, and each blue qubit in a cluster is connected with corresponding blue qubits in the neighboring columns. Therefore, each qubit interacts via a coupler with at most six other qubits. Each coupler can be assigned a weight J_{ij} , and each qubit can be assigned a magnetic field bias h_j . Sixteen of the qubits and one coupler were non-functional on the chip used in our study. A complete bipartite graph $K_{m,m}$ has m^{2m-2} spanning trees, which for $K_{4,4}$ is 4096. This number can be found on the OEIS website [9].

In the initial experiment, we studied the ferromagnet and the antiferromagnet on the full 8x8 Chimera graph (the entire chip) by applying equal interaction weights to each qubit pair; no magnetic fields were applied. For the ferromagnet, the interaction weights were varied from -1.5 to 0 in increments of 0.1; for the antiferromagnet, the values were varied from 1 to 0 in increments of 0.5. The results were similar in both cases: for interaction weights greater than 0.5, the D-Wave Two displayed an average success probability of 90% or greater for finding the ground state energy; the success probability decreased quickly for smaller weights. For the ferromagnet, every ground state configuration found consisted of all spin-ups; no spin-down ground state was observed, suggesting that some bias exists in the finding of the ground state spins. The D-Wave machine used was not a delivered product, and was intended for use internal to or approved by the company; so, the tuning procedure had not been performed to remove this up-down bias. Next, different random spanning trees were applied to the Chimera graph. Ferromagnetic interaction weights were applied and were again varied from -1.5 to 0 in increments of 0.1. We observed the average success probability drop significantly as the interaction weight decreased. No clear pattern was observed among the data points, with the average success probability fluctuating. These fluctuations were most likely due to the up-down bias.

In the second experiment, random spanning trees were solved with equal, constant interaction weights applied to each qubit pair in the tree. This was done for the 3x3, 4x4, 5x5, 6x6, and 7x7 subgraphs. Up to one hundred different spanning trees were solved on each subgraph, and each tree was submitted up to one hundred times. Each submission, or trial, resulted in the execution of one thousand annealing trajectories. We observed the success

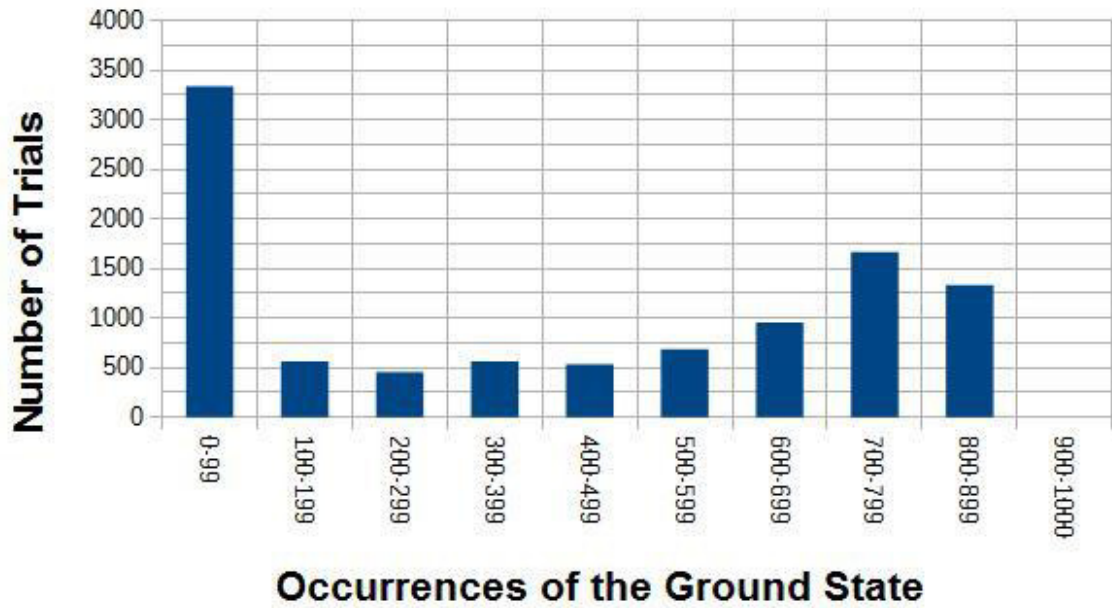


Fig. 2. Trials on the 5x5 subgraph organized by success probabilities, defined as the number of times the ground state was found. Data are from 100 spanning trees, each with 10^5 measurements.

probabilities to fluctuate widely from trial to trial. The only pattern observed was that the 3x3, 5x5, and 7x7 graphs tended to have a higher average success probability than anticipated by data from the 4x4 and 6x6 graphs. Again, only spin-up ground states were observed, and the fluctuations in these results were attributed mainly to the up-down bias.

In the final experiment, random spanning trees were again solved, but with interaction weights of -1 or +1 randomly assigned to each qubit coupler involved; again, all h_j were set to zero. This corresponds to a random gauge transformation [10]. The calculations were performed on the full 8x8 graph and on each subgraph, excluding the trivial 1x1 subgraph. All subgraphs were restricted to the upper-left corner of the Chimera lattice. Up to one hundred spanning trees were solved on each subgraph; again, each tree was submitted up to one hundred times, resulting in one thousand annealing trajectories each. The number of occurrences of the ground state energy was recorded. The mean, standard deviation, and median were calculated for each tree and across the trees of each subgraph. Additionally, for each subgraph, the trials were organized into bins according to their success probability. An example of such an analysis is shown in Fig. 2. For the 5x5 subgraph, which had 197 available qubits, one hundred different spanning trees were generated, and each one was submitted one hundred times, giving a total of ten thousand trials. On each trial, the D-Wave Two solved the tree one thousand times, and we recorded the number of times that it found the ground state energy. These trials were then organized by their success probability into the bins shown in the figure. The success probabilities of the 5x5 subgraph demonstrated a type of bimodal distribution, with a larger peak around the lower success probabilities. The success rates were highest for the 2x2 subgraph and decreased for each successive subgraph up to the full lattice. The 2x2, 3x3, and 4x4 subgraphs tended to be much more successful on the average, while the 6x6, 7x7, and 8x8 graphs showed much lower success probabilities. Additionally, the 4x4 subgraph also demonstrated a bimodal distribution, with the larger peak shifted to the right toward higher success probabilities.

In our final analysis, we looked at whether or not the ground state energy was found on each trial. If the ground state was found at least once out of the one thousand reads, this was counted as a success. The results are presented in Fig. 3. Each point represents the success probability over all trials on that subgraph. For the 2x2 and 3x3

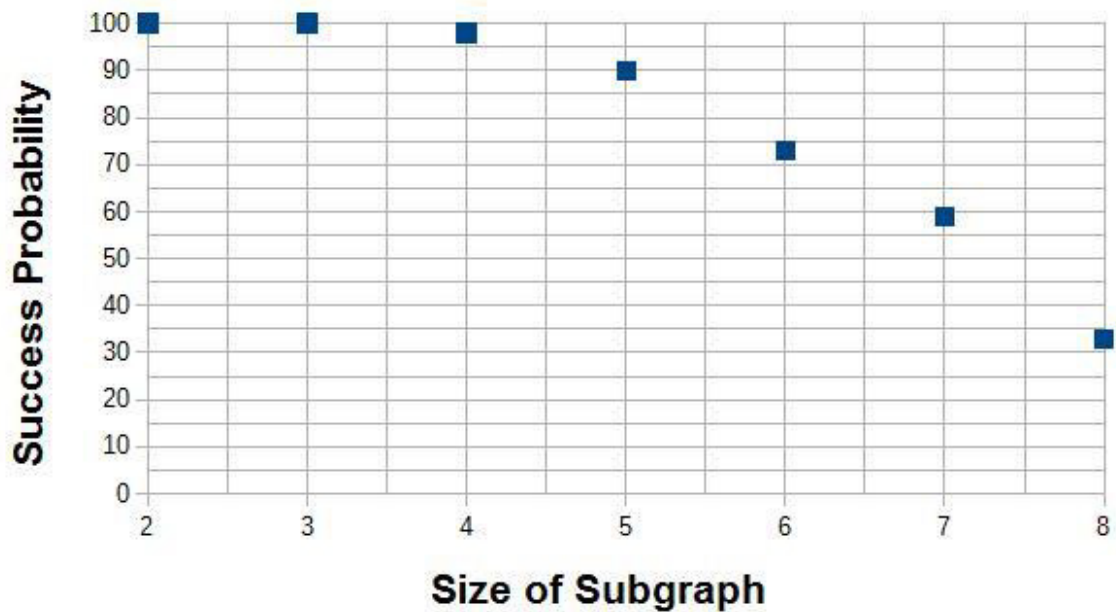


Fig. 3. Success probabilities for each subgraph.

subgraphs, the D-Wave Two found the ground state energy 100% of the time; this means that the ground state energy was found at least once on each trial. The success probability decreases for each larger subgraph; the success probability for the full 8x8 graph was 33%. The success of an adiabatic quantum annealer, such as the D-Wave Two, is a function of the annealing time and the temperature [11]. For infinite annealing time and operation at absolute zero temperature, the ground state is guaranteed, theoretically, to be found [12]. The chip used in our study had an annealing time of 20 μ s and operated at 20 mK. The next generation D-Wave computer, a 1000-qubit chip, is expected to operate at a lower temperature, and this should increase the success probability for each subgraph.

In summary, we concluded in our first study of the ferromagnet and the antiferromagnet that the probability of finding the ground state correlates with the strength of the interaction, as expected; higher success rates are seen for stronger interaction weights and decrease at lower interaction weights. In our second study of ferromagnetic spanning trees, we determined that the up-down bias was influencing our results to an unacceptable degree, prompting us to study spanning trees with randomly assigned interaction weights of -1 or +1. We determined this to sufficiently mitigate the bias. The final results show that the probability of finding the ground state energy decreases on average with the number of qubits for fixed annealing time and temperature. The philosophy of our study is similar to a recent preprint [13].

Acknowledgements

Supported in part by NSF grants DGE-0947419 and DMR-1206233. Computer time supplied by D-Wave Systems, Inc.

References

- [1] <http://www.dwavesys.com>
- [2] Broder A. Generating random spanning trees. *30th Annual Symposium on Foundations of Computer Science* 1989; 442-447.
- [3] Pemantle R. Choosing a spanning tree for the integer lattice uniformly. *Ann. Probab.* 1991; 19:1559-1574.

- [4] Wilson DB. Generating random spanning trees more quickly than the cover time. *Proceedings of the Twenty-eighth Annual ACM Symposium on the Theory of Computing* 1996; 296-303.
- [5] Ronnow TF, Wang Z, Job J, Isakov SV, Wecker D, Martinis JM, Lidar DA, Troyer M. Defining and detecting quantum speedup. *Science* 2014; 345:420-424.
- [6] Johnson MW, et al. Quantum annealing with manufactured spins. *Nature* 2011; 473:194-198.
- [7] Lanting T, et al. Entanglement in a quantum annealing processor. *Phys. Rev. X* 2014; 4.
- [8] Bian Z, Chudak F, Israel R, Lackey B, Macready WG, Roy A. Discrete optimization using quantum annealing on sparse Ising models. *Front. Phys.* 2014; 2.
- [9] <https://oeis.org>. The authors thank Dr. Edward D. Dahl of D-Wave for pointing this out.
- [10] Boixo S, Ronnow TF, Isakov SV, Wang Z, Wecker D, Lidar DA, Martinis JM, and Troyer M. Evidence for quantum annealing with more than one hundred qubits. *Nature Physics* 2014; 10:218.
- [11] Rieffel EG, Venturelli D, O’Gorman B, Do MB, Pristay EM, Smelyanskiy VN. A case study in programming a quantum annealer for hard operational planning problems. *Quant. Inf. Process.* 2014; 14:1-36.
- [12] Rieffel EG, Polak W. *Quantum Computing: A Gentle Introduction*. MIT Press; 2011.
- [13] Hen I, Job J, Albash T, Ronnow TF, Troyer M, Lidar DA. Probing for quantum speedup in spin glass problems with planted solutions. arXiv:1502.01663v2.