

This is the accepted manuscript made available via CHORUS. The article has been published as:

# Exponentially Biased Ground-State Sampling of Quantum Annealing Machines with Transverse-Field Driving Hamiltonians

Salvatore Mandrà, Zheng Zhu, and Helmut G. Katzgraber

Phys. Rev. Lett. **118**, 070502 — Published 17 February 2017

DOI: [10.1103/PhysRevLett.118.070502](https://doi.org/10.1103/PhysRevLett.118.070502)

# Exponentially-Biased Ground-State Sampling of Quantum Annealing Machines with Transverse-Field Driving Hamiltonians

Salvatore Mandrà,<sup>1,2,3,\*</sup> Zheng Zhu,<sup>4,†</sup> and Helmut G. Katzgraber<sup>4,5,‡</sup>

<sup>1</sup>*Department of Chemistry and Chemical Biology, Harvard University, 12 Oxford Street, 02138 Cambridge, Massachusetts, USA*

<sup>2</sup>*NASA Ames Research Center Quantum Artificial Intelligence Laboratory (QuAIL), Mail Stop 269-1, 94035 Moffett Field CA*

<sup>3</sup>*Stinger Ghaffarian Technologies Inc., 7701 Greenbelt Rd., Suite 400, Greenbelt, MD 20770*

<sup>4</sup>*Department of Physics and Astronomy, Texas A&M University, College Station, Texas 77843-4242, USA*

<sup>5</sup>*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501 USA*

(Dated: January 17, 2017)

We study the performance of the D-Wave 2X quantum annealing machine on systems with well-controlled ground-state degeneracy. While obtaining the ground-state of a spin-glass benchmark instance represents a difficult task, the gold standard for any optimization algorithm or machine is to sample all solutions that minimize the Hamiltonian with more or less equal probability. Our results show that while naïve transverse-field quantum annealing on the D-Wave 2X device can find the ground-state energy of the problems, it is not well suited in identifying all degenerate ground-state configurations associated to a particular instance. Even worse, some states are exponentially suppressed, in agreement with previous studies on toy model problems [New J. Phys. **11**, 073021 (2009)]. These results suggest that more complex driving Hamiltonians are needed in future quantum annealing machines to ensure a fair sampling of the ground-state manifold.

PACS numbers: 75.50.Lk, 75.40.Mg, 05.50.+q, 03.67.Lx

Optimization is ubiquitous across disciplines. Finding optimization approaches that quickly and reliably estimate the ground-state of a complex optimization problem is of great importance. While many algorithmic approaches from computer science have had a great impact in physics problems, similarly, physics-inspired optimization techniques have revolutionized optimization in fields as broad as engineering, biology, chemistry and computer science, to name a few. One physically-inspired optimization technique that has found wide-spread application is simulated annealing [1]. Similar to thermal annealing invented towards the end of the neolithic era [2], the heuristic is straightforward to implement. Initially, the system is prepared at a high temperature and it is left to thermalize. The temperature is sequentially reduced and, during the process, the system is enforced, if possible, to be in thermal equilibrium at any given temperature. At the end of the annealing (namely, when a specific target temperature is reached), the lowest energy configuration recorded during the process is returned. The procedure is repeated with different initial conditions to ensure that the obtained state is, actually, the lowest-energy state. Most importantly, it has been shown rigorously that simulated annealing can indeed obtain the ground-state of a system for sufficiently long annealing [3]; however, this is not practical. Nevertheless, it often fails to find the global minimum when the energy landscape of the problem Hamiltonian has many metastable states, such as is the case of, e.g., spin glasses [4–7]. More recently, the quantum counterpart of simulated annealing (usually called “quantum annealing”) was suggested [8–16]. In this case, quantum fluctuations are typically induced by a transverse field (instead of thermal fluctuations) to drive transitions from state to state. The advantage of quantum annealing is that the induced quantum fluctuations, in principle, could aid in the search for the optimum by allowing the system to tunnel across thin energy




FIG. 1: (Color online) Histograms of the number of times a particular ground-state configuration is found using the DW2X quantum annealer (20  $\mu$ s annealing time) sorted by rank. Data for Chimera lattice instances with  $N = 8 \cdot c^2$  sites and  $N_{GS} = 3 \cdot 2^k$  ground-states are shown. The horizontal axis is normalized by  $N_{GS}$  for easier display of the data. In all cases studied, certain ground-states are exponentially suppressed (note the vertical logarithmic axis).

barriers. To date, it remains controversial if it is able to outperform simulated annealing or other classical optimization methods.




FIG. 2: (Color online) Histogram of the number of times a particular ground-state configuration is found using the DW2X quantum annealer (20  $\mu$ s annealing time) sorted by rank, after adding extra Gaussian noise to the couplers with variance  $\sigma_J^2$  and biases with variance  $\sigma_h^2$ . Data for Chimera lattice instances with  $N = 8 \cdot c^2$  sites and  $N_{GS} = 3 \cdot 2^k$  ground-states are shown. The horizontal axis is normalized by  $N_{GS}$  for easier display of the data. In all cases studied, noise minimally affects the exponential bias of the sampling (note the vertical logarithmic axis).

Interest in quantum annealing has been considerably boosted by the introduction of the D-Wave quantum annealers [17]. These devices experimentally implement finite-temperature quantum annealing with a transverse field on a system of Boolean variables coupled together on a topology known as the Chimera graph [18]. Advantages in the use of the method beyond specially-crafted problems for Chimera’s architecture [19–21] remains to be found, the D-Wave 2X (DW2X) machine can be considered a huge technological feat with radically new technology. Interestingly, while many aspects of the DW2X have been scrutinized in detail, no detailed tests on its “fair sampling” [22, 23] abilities—namely the ability to sample all states of a degenerate problem with (hopefully) equal probability—have been performed. Studies on toy problems and simple Hamiltonians suggest that transverse-field driven quantum annealing does not uniformly sample all the possible ground-states resulting in some configuration being exponentially suppressed [23]. Studies on different generations of the D-Wave quantum annealer [24–26] already suggested that the sampling might be biased, but no systematic study has been performed to date. This can be seen as a noticeable shortcoming of the optimization technique.

So why is the exponential suppression of certain states, i.e., the lack of fair sampling, such a problem? First, because good optimization techniques should deliver all possible configurations that minimize the problem Hamiltonian (provided enough repetitions and using different initial conditions) in addition to being fast and reliable. This encompasses a far more stringent quality test for any optimizer. Second, and most importantly, there are many important applications for which a fair sampling of states is fundamental. In

physical applications, a fair sampling of states is imperative when estimating the ground-state entropy of a degenerate system. Similarly in computer science, for many combinatorial problems, if one can sample uniformly from the set of solutions, then one can use these different solutions to obtain a highly-accurate count of the total number of solutions [27], which is important for propositional model counting (#SAT) [28] and the knapsack solution counting problem (#Knapsack) [29]. Finally, in multiple industrial applications having many different solutions to a problem is highly desirable. For example, many uncorrelated solutions are needed to construct probabilistic membership filters using SAT formulas [30, 31]. As such, a quantum annealing machine with a transverse-field driving Hamiltonian might not be the best approach to solve these problems. On one hand, one can hope that the inherent noise found in the analog DW2X might help alleviate these biases of transverse-field quantum annealing. On the other hand, this problem could be alleviated with more complex driving Hamiltonians [23]. Unfortunately, such machines are only being constructed at the moment.

In this work we demonstrate experimentally that, for spin-glass problems with a small (known) number of ground-state configurations, the DW2X is heavily biased towards some configurations, while other minimizing configurations are exponentially suppressed. Despite applying multiple gauges, performing many runs, or increasing the annealing time, the machine is unable to sample the states fairly, i.e., it is not well suited for a wide variety of optimization applications.

*Description of the benchmark instances.*— We perform the experiments on the D-Wave Systems, Inc., DW2X quantum annealing machine [17]. We use all operable qubits on the machine and encode spin-glass problems on the couplers [4, 7, 32] of the underlying Chimera topology of the system [18]. The Hamiltonian of the problem is  $\mathcal{H} = -\sum_{\{i,j\} \in \mathcal{V}} J_{ij} S_i^z S_j^z$ . The  $N$  Ising variables  $S_i^z \in \{\pm 1\}$  are defined on the vertices  $\mathcal{V}$  of the Chimera lattice of size  $N = 8 \cdot c^2$  (with  $c \in \{8, 9, 10, 11\}$ ) and do not couple to any local fields (biases). The sum is over all edges  $\mathcal{E}$  connecting vertices  $\{i, j\} \in \mathcal{V}$ . Note that some couplers and/or qubits are always inoperable. The aforementioned system sizes are for the complete lattices without taking into account any defects.

To perform a controlled study of the effects of ground-state degeneracy, we carefully choose the couplings from a Sidon set [33] with  $J_{ij} \in \{\pm 5, \pm 6, \pm 7\}$ . Furthermore, after randomly placing the couplings, we recursively traverse the lattice and shuffle the interactions randomly so that no spins have a zero local field. This prevents any additional degeneracy due to a larger number of free spins [33, 34]. Because of our choice of disorder, we find that the randomly-generated instances have typically a ground-state degeneracy of  $N_{GS} = 3 \cdot 2^k$  ( $k \in \mathbb{N}$ ). Some instances have values of  $N_{GS}$  that do not fall into the sequence  $N_{GS} = \{6, 12, 24, 48, 96, \dots\}$  because of the imperfections in the Chimera graph. We choose not to use such instances for the experiments to perform a systematic study. Note that for small sub-sections of the Chimera graph, i.e., for small  $c$ , the number of ground-states




FIG. 3: (Color online) Histogram of the number of times a particular ground-state configuration is found either using classical heuristics [using the isoenergetic cluster algorithm (ICA) [36] and the Hamze-de Freitas-Selby algorithm (HFS) [40, 41]], as well as the DW2X quantum annealer. Example data for Chimera lattice instances with  $N = 8 \cdot c^2$  sites and  $N_{GS} = 3 \cdot 2^k$  ground-states. While the classical algorithms sample the ground-state configurations fairly, the DW2X device has a bias of more than two orders of magnitude.

is typically smaller than for the largest possible lattice with  $N = 1152$  ( $c = 12$ ) sites [35]. Therefore, the available values of  $k$  are smaller.

*Experimental details.*— The number of ground states for each problem is determined classically using the isoenergetic cluster algorithm (ICA) [34, 36–39], which is known to sample the ground-state manifold fairly, especially for small numbers of ground-states (here, small means  $N_{GS} \lesssim 10^3$ ). ICA combines parallel tempering Monte Carlo with isoenergetic cluster moves (simulation parameters are shown in Table I). To ensure that the lowest energy state has been found, we independently simulate four system replicas with the same couplings. More precisely, we check that the lowest energy found by each replica (considering only the lowest temperatures) in  $N_{sw}/2$  updates, with  $N_{sw}$  the total number of updates, agree. Hence, we claim that the ground state has been found and we begin to record the ground-state configurations, and the corresponding frequencies, for the remaining  $N_{sw}/2$  updates. There is no guarantee that any solution obtained by this heuristic method is the true optimum, or that we have found all configurations that minimize the Hamiltonian. However, we ensure each configuration achieves a minimum number of 50 hits in order to increase our confidence that all accessible ground states have been found. Moreover, we also check that the lowest energy is in agreement with the Hamze-de Freitas-Selby heuristic [40, 41].

Quantum annealing experiments have been performed on the DW2X using a fixed annealing time of 20  $\mu s$ . For each instance, we used 100 distinct gauges and 1000 readouts per gauge, for a total number of  $10^5$  readouts per instance.

*Results.*— Figure 1 summarizes our results. Each panel

TABLE I: Simulation parameters for the isoenergetic cluster algorithm (ICA): for each system size  $N$ , we compute  $N_{sa}$  disorder instances.  $N_{sw} = 2^b$  is the total number of Monte Carlo sweeps for each of the  $4N_T$  replicas for a single instance,  $T_{min}$  [ $T_{max}$ ] is the lowest [highest] temperature simulated, and  $N_T$  is the number of temperatures used in the parallel tempering method. For the lowest  $N_{ICA}$  temperatures isoenergetic cluster moves are applied.

$N$	$N_{sa}$	$b$	$T_{min}$	$T_{max}$	$N_T$	$N_{ICA}$
512	4164	19	0.06	3.05	33	18
648	6970	19	0.06	3.05	33	18
800	11199	19	0.06	3.05	33	18
968	16739	19	0.06	3.05	33	18




FIG. 4: (Color online) Comparison of  $\Theta_{max}$  (maximum absolute difference of the empiric cumulative distribution with respect to the expected distribution) for the distribution of ground states found by the DW2X against a uniform random distribution. Each point corresponds to a specific instance and the error bars are computed by bootstrapping the data. The closer the points are to the diagonal, the more the ground-states for that specific instance have been uniformly sampled. As one can see, DW2X does not uniformly sample the ground-states. The example data are for Chimera lattices with  $N = 8 \cdot c^2$  sites and  $N_{GS} = 3 \cdot 2^k$  ground-states. For the analysis, we only considered those instances for which DW2X found at least 50 solutions (independently of the ground-state configuration).

shows a histogram with the number of times a particular ground-state configuration is found by the DW2X. The horizontal axis represents the index of a given ground-state configurations, normalized by  $3 \cdot 2^k$  for a better readability. For each instance, indexes of ground states are ordered so that the ground-states with the largest probability have the largest index. Each panel represents different experiments at a fixed system size  $N = 8 \cdot c^2$ , while each line considers only experiments with a fixed number of ground-state configurations  $N_{GS} = 3 \cdot 2^k$ . Error bars are computed by averaging each bin over a given number of samples. In all cases stud-

ied, some ground-state configurations are exponentially suppressed (note the vertical logarithmic axis). We obtain similar results by increasing the annealing time to  $200 \mu s$ . It is important to stress that the exponential bias is minimally affected by introducing additional artificial noise to the target Hamiltonian, as shown in Figure 2. Both random biases and coupler noise are drawn from a Gaussian distribution with variances  $\sigma_h^2$  and  $\sigma_j^2$ , respectively. In addition, we compare the sampling of the DW2X to the two most efficient classical heuristics in Figure 3. While the bias is minimal for the classical approaches (due to Poissonian fluctuations [22]), a bias of approximately two orders of magnitude persists for the DW2X device.

Finally, to better appreciate the exponential suppression of some ground states of the DW2X, we introduce the observable  $\Theta_{\max}$  defined as the maximum absolute difference of the empiric cumulative distribution  $\bar{F}(x)$  with respect to the cumulative of a uniform distribution  $F(x)$ , namely  $\Theta_{\max} = \max_x |\bar{F}(x) - F(x)|$ , with  $x$  the ground state index. The test (which is similar in the purpose of the Kolmogorov-Smirnov test) is useful to understand how close an empiric distribution is to the expected distribution. More precisely, the smaller  $\Theta_{\max}$  is, the more similar the distributions are. In Figure 4, we show the comparison of  $\Theta_{\max}$  for the distribution of ground-states found by the DW2X against random numbers uniformly chosen in the set  $\{1, 2, \dots, N_{GS}\}$ . In general, the number of ground states that DW2X can find widely varies from instance to instance. Therefore, to perform a fairer analysis, we extract an amount of random numbers which is equal to the number of solutions (regardless of the ground-state configurations) that the DW2X has found for the given instance. Each point in the plot corresponds to a specific instance and the error bars are computed by bootstrapping the data after the randomization of the ground-state indexes. The diagonal line represents the best value that  $\Theta_{\max}$  can assume: the closer the points are to the diagonal, the more uniformly the ground-states for that specific instance have been sampled. For the analysis, we only considered those instances for which DW2X has found at least 50 solutions (independently of the ground-state configuration). As one can see, the results show that all the considered instances are far from the optimal diagonal, which confirms that the DW2X using a transverse-field driving Hamiltonian does not sample uniformly. In addition, results from instances with fixed  $c$  and different  $k$  suggest that the DW2X slightly improves its sampling by increasing the total number of ground-states. An intuitive understanding of how degeneracy of ground-states changes sampling can be obtained by considering level crossings between ground-states and low-energy excited-states: instances with less degeneracy tend to be harder and more likely to have level crossings [33, 42], therefore a longer annealing time is required to reach a stationary distribution of ground states. Instances with larger degeneracy, however, have a slightly better fair sampling for the same amount of annealing time.

The abysmal fair-sampling performance of transverse-field quantum annealing on the DW2X suggests that the machine is

not well suited for applications where many uncorrelated optimal states are needed. Surprisingly, neither the intrinsic thermal fluctuations nor the application of multiple gauges seem to affect these results [43]. Attempting to run the machine for longer annealing times (see Fig. 5) has a negligible effect on the poor sampling of the machine. This is in agreement with simulations on simple toy models [23]. There, simulations showed that the use of more complex driving Hamiltonians might alleviate this problem. Finally, changing the energy scale of the Hamiltonian in the device, as well as adding additional artificial noise, does not affect the poor sampling (see Figure 2). As such, and in agreement with the aforementioned analytical results, the transverse-field driver is likely the source of the bias. Unfortunately, at the moment neither quantum annealing machines with more complex driving Hamiltonians nor quantum Monte Carlo simulations to emulate these are readily available. However, the aforementioned results strongly argue for more complex annealing architectures in future devices.

*Summary.*— We have demonstrated experimentally that the D-Wave 2X quantum annealer is unable to fairly sample states of degenerate random spin-glass problems. In fact, some states are exponentially suppressed compared to others. This means that transverse-field quantum annealing might not be well suited for applications where many uncorrelated solutions are needed. This could also explain the poor performance of the implementation of probabilistic membership filters on the D-Wave device [31]. Our results are in agreement with previous theoretical and numerical studies [23] on toy models and suggest that the ever-growing quantum annealing community should put more emphasis on mitigating this problem by, e.g., using more complex driving Hamiltonians [23] or developing hybrid architectures that encourage thermal fluctuations [44]. We do emphasize, however, that degenerate embedded problems might be affected differently by this problem. For example, different embeddings might influence the sampling differently, a problem that should be studied in the future.

*Acknowledgements.*— We would like to thank A. Aspuru-Guzik, F. Hamze, A. King, A. J. Ochoa and A. Perdomo-Ortiz for fruitful discussions. We also thank E. G. Rieffel and D. Venturelli for carefully reading the manuscript. H.G.K. acknowledges support from the NSF (Grant No. DMR-1151387) and would like to thank Zaya for inspiration to initiate this project. S.M. was supported by NASA (Sponsor Award Number: NNX14AF62G). We thank the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing HPC resources (Stampede Cluster) and Texas A&M University for access to their Ada and Lonestar clusters. The research of H.G.K. and Z.Z. is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via MIT Lincoln Laboratory Air Force Contract No. FA8721-05-C-0002. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official poli-






FIG. 5: (Color online) Comparison of  $\Theta_{\max}$  (defined as the maximum absolute difference of the empiric cumulative distribution with respect to the expected one) for the distribution of ground states found by the DW2X for two different annealing times,  $20 \mu\text{s}$  and  $200 \mu\text{s}$ . Each point in the plot corresponds to a specific instance and the error bars are computed by bootstrapping the data. Only instances for which DW2X has found a solution for both annealing times are represented. The data are for Chimera lattice instances with  $N = 8 \cdot c^2$  sites and  $N_{\text{GS}} = 3 \cdot 2^k$  ground-states. Even though there is a slight improvement by increasing the total annealing time,  $\Theta_{\max}$  remains quite large, thus showing that sampling remains strongly biased even when the annealing time is increased tenfold.

cies or endorsements, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purpose notwithstanding any copyright annotation thereon. All authors contributed equally to the project.

\* Electronic address: [smandra@fas.harvard.edu](mailto:smandra@fas.harvard.edu)

† Electronic address: [zzwtgts@tamu.edu](mailto:zzwtgts@tamu.edu)

‡ Electronic address: [hgk@tamu.edu](mailto:hgk@tamu.edu)

- [1] S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi, *Optimization by simulated annealing*, Science **220**, 671 (1983).
- [2] For example, one of the objects found with Ötzi the iceman was a copper axe.
- [3] S. Geman and D. Geman, IEEE Trans. Pattern. Anal. Mach. Intell. **PAMI-6**, 721 (1984).
- [4] K. Binder and A. P. Young, *Spin Glasses: Experimental Facts, Theoretical Concepts and Open Questions*, Rev. Mod. Phys. **58**, 801 (1986).
- [5] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
- [6] A. P. Young, ed., *Spin Glasses and Random Fields* (World Scientific, Singapore, 1998).
- [7] D. L. Stein and C. M. Newman, *Spin Glasses and Complex-*

ity, Primers in Complex Systems (Princeton University Press, 2013).

- [8] A. B. Finnila, M. A. Gomez, C. Sebenik, C. Stenson, and J. D. Doll, *Quantum annealing: A new method for minimizing multi-dimensional functions*, Chem. Phys. Lett. **219**, 343 (1994).
- [9] T. Kadowaki and H. Nishimori, *Quantum annealing in the transverse Ising model*, Phys. Rev. E **58**, 5355 (1998).
- [10] J. Brooke, D. Bitko, T. F. Rosenbaum, and G. Aeppli, *Quantum annealing of a disordered magnet*, Science **284**, 779 (1999).
- [11] 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 **292**, 472 (2001).
- [12] G. Santoro, E. Martoňák, R. Tosatti, and R. Car, *Theory of quantum annealing of an Ising spin glass*, Science **295**, 2427 (2002).
- [13] A. Das and B. K. Chakrabarti, *Quantum Annealing and Related Optimization Methods* (Edited by A. Das and B.K. Chakrabarti, Lecture Notes in Physics 679, Berlin: Springer, 2005).
- [14] G. E. Santoro and E. Tosatti, *TOPICAL REVIEW: Optimization using quantum mechanics: quantum annealing through adiabatic evolution*, J. Phys. A **39**, R393 (2006).
- [15] A. Das and B. K. Chakrabarti, *Quantum Annealing and Analog Quantum Computation*, Rev. Mod. Phys. **80**, 1061 (2008).
- [16] S. Morita and H. Nishimori, *Mathematical Foundation of Quantum Annealing*, J. Math. Phys. **49**, 125210 (2008).
- [17] See <http://www.dwavesys.com>.
- [18] P. Bunyk, E. Hoskinson, M. W. Johnson, E. Tolkacheva, F. Altomare, A. J. Berkley, R. Harris, J. P. Hilton, T. Lanting, and J. Whittaker, *Architectural Considerations in the Design of a Superconducting Quantum Annealing Processor*, IEEE Trans. Appl. Supercond. **24**, 1 (2014).
- [19] D. Venturelli, S. Mandrà, S. Knysh, B. O’Gorman, R. Biswas, and V. Smelyanskiy, *Quantum Optimization of Fully-Connected Spin Glasses* (2014), (arXiv:cond-mat/1406.7553).
- [20] V. S. Denchev, S. Boixo, S. V. Isakov, N. Ding, R. Babush, V. Smelyanskiy, J. Martinis, and H. Neven, *What is the Computational Value of Finite Range Tunneling?* (2015), (arXiv:quant-ph/1512.02206).
- [21] S. Mandrà, Z. Zhu, W. Wang, A. Perdomo-Ortiz, and H. G. Katzgraber, *Strengths and Weaknesses of Weak-Strong Cluster Problems: A Detailed Overview of State-of-the-art Classical Heuristics vs Quantum Approaches* (2016), (arXiv:1604.01746).
- [22] J. J. Moreno, H. G. Katzgraber, and A. K. Hartmann, *Finding low-temperature states with parallel tempering, simulated annealing and simple Monte Carlo*, Int. J. Mod. Phys. C **14**, 285 (2003).
- [23] Y. Matsuda, H. Nishimori, and H. G. Katzgraber, *Ground-state statistics from annealing algorithms: quantum versus classical approaches*, New J. Phys. **11**, 073021 (2009).
- [24] S. Boixo, T. Albash, F. M. Spedalieri, N. Chancellor, and D. A. Lidar, *Experimental signature of programmable quantum annealing*, Nat. Commun. **4**, 2067 (2013).
- [25] T. Albash, T. F. Rønnow, M. Troyer, and D. A. Lidar, *Reexamining classical and quantum models for the D-Wave One processor*, Eur. Phys. J. Spec. Top. **224**, 111 (2015).
- [26] A. D. King, E. Hoskinson, T. Lanting, E. Andriyash, and M. H. Amin, *Degeneracy, degree, and heavy tails in quantum annealing*, Phys. Rev. A **93**, 052320 (2016).
- [27] M. R. Jerrum, L. G. Valiant, and V. V. Vazirani, *Random generation of combinatorial structures from a uniform distribution*, Theoretical Computer Science **43**, 169 (1986).
- [28] C. P. Gomes, A. Sabharwal, and B. Selman, *Model counting*, in

- Handbook of Satisfiability*, edited by A. Biere, M. Heule, H. van Maaren, and T. Walsch (IOS Press, 2008).
- [29] P. Gopalan, A. Klivans, R. Meka, D. Stefankovic, S. Vempala, and E. Vigoda, in *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on* (IEEE, 2011), p. 817.
  - [30] S. A. Weaver, K. J. Ray, V. W. Marek, A. J. Mayer, and A. K. Walker, *Satisfiability-based set membership filters*, Journal on Satisfiability, Boolean Modeling and Computation (JSAT) **8**, 129 (2014).
  - [31] A. Douglass, A. D. King, and J. Raymond, *Constructing SAT Filters with a Quantum Annealer*, in *Theory and Applications of Satisfiability Testing – SAT 2015* (Springer, 2015), pp. 104–120.
  - [32] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing: An Introduction* (Oxford University Press, New York, 2001).
  - [33] H. G. Katzgraber, F. Hamze, Z. Zhu, A. J. Ochoa, and H. Munoz-Bauza, *Seeking Quantum Speedup Through Spin Glasses: The Good, the Bad, and the Ugly*, Phys. Rev. X **5**, 031026 (2015).
  - [34] Z. Zhu, A. J. Ochoa, F. Hamze, S. Schnabel, and H. G. Katzgraber, *Best-case performance of quantum annealers on native spin-glass benchmarks: How chaos can affect success probabilities*, Phys. Rev. A **93**, 012317 (2016).
  - [35]  $N = 1152$  is the theoretical maximal size of the chip. However, trapped fluxes and manufacturing imperfections reduce the problem size to approximately 100 sites.
  - [36] Z. Zhu, A. J. Ochoa, and H. G. Katzgraber, *Efficient Cluster Algorithm for Spin Glasses in Any Space Dimension*, Phys. Rev. Lett. **115**, 077201 (2015).
  - [37] Z. Zhu, A. J. Ochoa, and H. G. Katzgraber, *Efficient sampling of ground-state configurations for quasi two-dimensional ising spin glasses*, in preparation (2016).
  - [38] Z. Zhu, C. Fang, and H. G. Katzgraber, *borealis - A generalized global update algorithm for Boolean optimization problems* (2016), (arXiv:1605.09399).
  - [39] J. Houdayer, *A cluster Monte Carlo algorithm for 2-dimensional spin glasses*, Eur. Phys. J. B. **22**, 479 (2001).
  - [40] F. Hamze and N. de Freitas, in *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence* (AUAI Press, Arlington, Virginia, United States, 2004), UAI '04, p. 243, ISBN 0-9749039-0-6.
  - [41] A. Selby, *Efficient subgraph-based sampling of Ising-type models with frustration* (2014), (arXiv:cond-mat/1409.3934).
  - [42] S. Boixo, T. F. Rønnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer, *Evidence for quantum annealing with more than one hundred qubits*, Nat. Phys. **10**, 218 (2014).
  - [43] At this time it is unclear of the intrinsic noise of the D-Wave device either helps with the fair sampling or makes it worse. However, given the results presented in Ref. [23], we expect that noise should not be the sole source of the lack of fair sampling.
  - [44] N. G. Dickson, M. W. Johnson, M. H. Amin, R. Harris, F. Altomare, A. J. Berkley, P. Bunyk, J. Cai, E. M. Chapple, P. Chavez, et al., *Thermally assisted quantum annealing of a 16-qubit problem*, Nat. Commun. **4**, 1903 (2013).