



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

Equation Planting: A Tool for Benchmarking Ising Machines

Itay Hen

Phys. Rev. Applied **12**, 011003 — Published 19 July 2019

DOI: [10.1103/PhysRevApplied.12.011003](https://doi.org/10.1103/PhysRevApplied.12.011003)

Equation Planting: A Tool for Benchmarking Ising Machines

Itay Hen^{1,2}

¹*Information Sciences Institute, University of Southern California, Marina del Rey, California 90292, USA*

²*Department of Physics and Astronomy and Center for Quantum Information Science & Technology,
University of Southern California, Los Angeles, California 90089, USA*

(Dated: July 3, 2019)

We introduce a methodology for generating benchmark problem sets for Ising machines—devices designed to solve discrete optimization problems cast as Ising models. In our approach, linear systems of equations are cast as Ising cost functions. While linear systems are easily solvable, the corresponding optimization problems are known to exhibit some of the salient features of NP-hardness, such as strong exponential scaling of heuristic solvers’ runtimes and extensive distances between ground and low-lying excited states. We show how the proposed technique, which we refer to as ‘equation planting,’ can serve as a useful tool for evaluating the utility of Ising solvers functioning either as optimizers or as ground-state samplers. We further argue that equation-planted problems can be used to probe the mechanisms underlying the operation of Ising machines.

Introduction.— Recent years have witnessed a flourishing in experimental ‘Ising machines’—special-purpose programmable devices engineered to solve discrete optimization problems [1, 2] cast as Ising models—with the tacit promise that their performance is superior to those of algorithms running on standard computers. Analog quantum devices that perform quantum annealing [3, 4] designed to find bit assignments that minimize the cost of Ising Hamiltonians have already been realized on various platforms such as arrays of superconducting flux qubits [5–11]. Other notable technologies that recently emerged are coherent Ising machines based on lasers and degenerate optical parametric oscillators [12, 13], FPGA-based quantum-inspired digital annealers [14], and mem-computing devices that operate on terminal-agnostic self-organizing logic gates [15–17].

Touting improved performance over traditional algorithms [18, 19], Ising machines have gained a large amount of interest, both in the academe as well as with the general public—and rightfully so. Many problems of theoretical and practical relevance, in areas as diverse as machine learning, materials design, software verification, portfolio management, logistics, and many more, can be cast as searching for the global minima of Ising cost functions [1, 2]. It is, however, unclear to date whether indeed any one of the aforementioned devices truly offers genuine advantages over its competitors.

One of the main bottlenecks preventing meaningful benchmarking of Ising machines is the absence of appropriate benchmark problems. Generating problem instances that are on the one hand challenging enough and whose solutions are on the other hand known in advance or are easily verifiable, so as to allow proper benchmarking, are not straightforward to construct since these two requirements are, in many respects, contradictory. To overcome this obstacle, various methods have been devised in recent years to generate problem classes with known minimizing configurations (see, e.g., Refs [20–22]). However these have generally been found to lack the hardness that characterizes NP-hard problems. At the other extreme, problem classes that are challenging to

solve but whose ground-state energies are not verifiable, have also been developed [23–25].

To bridge the above gap, in this study we advocate a methodology for generating benchmark problem sets for Ising machines that we argue may serve as a suitable tool for evaluating their performance and in turn also to indirectly probe the mechanisms underlying their operation. The proposed technique, which we refer to as ‘equation planting,’ has several desirable properties. Explicitly, the generated problem sets, while generated from easily solvable problems, possess some of the salient features of NP-hard problems—most notably extensive distances between ground and low-lying excited states.

Our approach is based on the embedding of linear systems of equations in Ising models, motivated by the observation that when linear systems of equations are cast as optimization problems, the latter form often stymies heuristic solvers [26–31]. In addition to serving as suitable benchmark problems to Ising machines functioning as optimizers, we also show that the proposed technique may be used to test the functionality of these devices as Boltzmann samplers, or more specifically as ground-state samplers [32–34], thanks to another property of these problem sets—namely, a verifiable number of ground state configurations.

We begin by describing the equation planting technique in general, moving on to focus on a specific class of problem instances using which we demonstrate the effectiveness of the method.

Equation planting.— We begin by considering a linear system of m equations in n variables

$$\sum_{j=1}^n a_{ij}x_j = b_i \quad \text{for } i = 1 \dots m. \quad (1)$$

Here $\{x_1, \dots, x_n\}$ stand for variables over a given field, and $\{a_{ij}\}$ and $\{b_i\}$ are the equation coefficients. Every such linear system may be cast as an optimization prob-

lem with the corresponding cost function

$$F = \sum_{i=1}^m \left\| \sum_{j=1}^n a_{ij} x_j - b_i \right\|^2. \quad (2)$$

Since F is a sum over positive terms, any configuration $\{x_1^*, \dots, x_n^*\}$ that yields $F = 0$ is a minimizing configuration that also solves the linear system.

As already mentioned, even though the computational problem of solving a linear system of equations is easy—any given instance can always be solved in polynomial time using Gaussian elimination—the corresponding optimization problem is not necessarily easy for heuristic solvers [26, 27, 29–31].

We leverage the above setup towards constructing Ising problems that stymie heuristic solvers. To that aim, we focus henceforth on linear systems of equations modulo 2, that is,

$$\sum_{j=1}^n a_{ij} x_j = b_i \pmod{2} \quad \text{for } i = 1 \dots m, \quad (3)$$

where now both the variable set and coefficients are Booleans taking on values $\{0, 1\}$. The i th equation can therefore be written as $\bigoplus_{j:a_{ij}=1} x_j = b_i$ (here \oplus denotes the bitwise XOR operation), or in terms of Ising spins $s_i \in \{-1, 1\}$,

$$\prod_{j:a_{ij}=1} s_j = (-1)^{b_i}. \quad (4)$$

Cast in optimization form, the system of equations becomes

$$\tilde{F}_2 = \sum_{i=1}^m \left(\prod_{j:a_{ij}=1} s_j - (-1)^{b_i} \right)^2, \quad (5)$$

or, after dropping immaterial constants,

$$F_2 = - \sum_{i=1}^m (-1)^{b_i} \prod_{j:a_{ij}=1} s_j. \quad (6)$$

The cost function F_2 is a multi-spin cost function consisting of a sum of products of spin variables. The localities of the terms comprising the cost function are precisely the number of variables in the corresponding equations. The minima of F_2 correspond to the solutions of the system given in Eq. (4), provided that solutions exist. In more detail, linear systems may behave in one of three possible ways. (i) The system may have no solutions at all. This scenario corresponds to the cost function \tilde{F}_2 having a strictly positive minimal value (meaning $F_2 > -m$) and may happen if the number of equations is greater than the number of variables. (ii) The system may have a unique solution. Here, a single configuration minimizes F_2 whose minimal value in this case would be $-m$. (iii)

Last, the system may have multiple solutions if the dimension of its null space, or nullity [35], which we denote here by d_N , is nonzero. In this case, the number of minimizing configurations that yield $F_2 = -m$ is $N_{\text{GS}} = 2^{d_N}$.

Ising machines are designed to tackle two-body cost functions. The locality of the multi-spin cost F_2 may in this case be readily reduced to a two-body Ising model of the general form $\sum_{ij} J_{ij} s_i s_j + \sum_i h_i s_i$ using standard reduction techniques (see, e.g., Ref. [36]), where $\{J_{ij}\}$ and $\{h_i\}$ are integer-valued coefficients.

As we demonstrate in the next section, while the solutions of the linear systems of equations used to generate these Ising cost functions are straightforwardly obtained (if there are any), heuristic solvers will generically find this type of problems extremely difficult to solve. In what follows, we illustrate this by studying in detail one specific type of equation system—namely, random 3-regular 3-XORSAT.

3-regular 3-XORSAT as 2-body Ising.— To demonstrate the challenges that equation-planted instances present to heuristic Ising solvers, we consider as a test case linear systems $\pmod{2}$ (also known as XORSAT equations) wherein each equation contains exactly three randomly chosen Boolean variables and each variable, or bit, x_j appears in exactly three equations. This type of problem is also known as ‘3-regular 3-XORSAT’ (3R3X) and was studied previously in various contexts [28, 30, 31, 37, 38]. An n -bit instance would thus consist of n equations of the form $x_{i_1} \oplus x_{i_2} \oplus x_{i_3} = b_i$ [39]. We note that while the results presented below pertain to 3R3X, they extend in general to other classes of linear XORSAT systems.

The cost function F_2 of an n -bit 3R3X system is therefore a sum of n three-body terms of the form $-(-1)^{b_i} s_{i_1} s_{i_2} s_{i_3}$ defined on n Ising spins. To reduce the locality of a term to two- and one-body we use a gadget that shares its four minimizing configurations [39]. For the negatively signed clause ($b_i = 0$), these are the four configurations whose product is positive, namely, $(1, 1, 1), (1, -1, -1), (-1, 1, -1)$ and $(-1, -1, 1)$ and an appropriate gadget is

$$G_{3X} = h(s_{i_1} + s_{i_2} + s_{i_3}) + \tilde{h} s_a \quad (7) \\ + J(s_{i_1} s_{i_2} + s_{i_2} s_{i_3} + s_{i_3} s_{i_1}) + \tilde{J} s_a (s_{i_1} + s_{i_2} + s_{i_3}),$$

where $(h, \tilde{h}, J, \tilde{J})$ can be either $(-1, -2, 1, 2)$ or $(-1, 2, 1, -2)$, yielding in both cases a minimal cost of -4 (we note that other gadgets exist that can equivalently be used). The above gadget is a fully connected 4-spin cost function with s_a serving as an auxiliary spin (see Fig. 1). The positively signed clauses are treated similarly.

The above gadget allows us to cast n -bit 3R3X linear systems as two-body Ising models with $2n$ spins, as each clause adds one auxiliary spin to the problem and an n -bit instance has n clauses.

At this point, we note several of the 3R3X Ising instances properties: (i) First, these instances have a




Figure 1. The two-body 3-XORSAT gadget. The gadget is a fully connected spin quadruplet with parameter values $(h, \tilde{h}, J, \tilde{J})$ as described in the main text. The four minimizing configurations of the 3-XORSAT clause are encoded into the ground states of the Ising cost of the gadget restricted to spins $s_{i_1}, s_{i_2}, s_{i_3}$. The spin s_a serves as an auxiliary spin.

bounded degree, namely 9, as each spin connects in general to three other spins for every clause in which the spin appears. (ii) The specification of Ising parameters $\{J_{ij}\}$ and $\{h_i\}$ requires only two bits of precision as both spin-spin couplings and longitudinal fields are integers in the range $[-3, 3]$. (iii) As already mentioned, these instances have known ground-state energies. (iv) Finally, these instances have a controllable ground-state degeneracy, determined by the nullity of the generating linear system. As will become evident in the next section, these properties may be leveraged to determine the utility of heuristic Ising solvers functioning either as optimizers or as ground-state samplers.

Results.— In what follows, we examine the performance of heuristic solvers on randomly generated two-body Ising 3R3X instances that have unique solutions. As a representative for state-of-the-art heuristic solvers, we use parallel tempering (PT) [40, 41] — a refinement of the celebrated yet somewhat outdated simulated annealing algorithm [42], that finds the ground-state configurations of general discrete-variable cost functions [20, 24, 37, 43–49]. In PT, multiple copies of the problem are equilibrated in parallel at different temperatures and spin configurations at adjacent temperatures are regularly swapped [39]. In addition, we consider a variant of PT (which we denote PT-H) that employs global cluster moves due to Houdayer [50] and is known to accelerate PT convergence rates for many problem classes (see, e.g., Refs. [51–53]).

We test the performance of PT and PT-H on instances of varying sizes by measuring their typical times to find minimizing configurations. Here, typical runtimes are defined as the median time to reach a minimizing configuration over 100 randomly generated instances of a given size [54].

The results are presented in Fig. 2(top) showing the exponential runtime scaling of both PT and PT-H in problem size, proportional to $e^{\alpha n}$ with $\alpha \approx 0.13$ and 0.14, respectively, indicating a strong exponential scaling. Moreover, the Houdayer cluster updates do not seem to improve the scaling. Also shown in the figure

(as shaded areas) are the interquartile ranges (the difference between 75th and 25th percentiles) of the PT and PT-H runtimes, quantitatively indicating the homogeneity in the hardness of the instances.

To better understand why these problems, while being trivial to solve using Gaussian elimination, present severe challenges to heuristic solvers, we next examine their energy landscapes. We do so by measuring the Hamming distances between local minimum configurations, as found by the solver in the course of the simulation, and the global minimum. These are plotted in Fig. 2(bottom). The low-lying excited states, characterized by small yet positive residual energies, are typically $\sim 0.6n$ spin flips away from the global minimum, implying that reaching a local minimum is no indication of the whereabouts of the ground state. The above property is a hallmark of NP-hard problems [2].

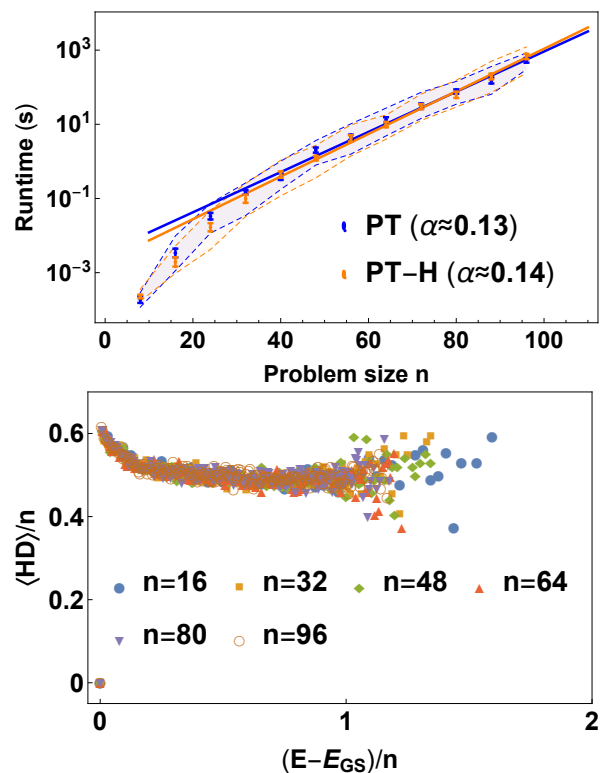


Figure 2. Top: Runtime scaling of two parallel tempering variants on random 3R3X-planted Ising instances (log-linear scale). Typical runtimes (median over 100 instances) as a function of problem size is given for parallel tempering (PT) and PT augmented with Houdayer cluster moves (PT-H). Both exhibit similarly strong exponential scaling. Fits to $e^{\alpha n}$ are shown. Error bars correspond to 1σ statistical confidence. The shaded areas correspond to the interquartile ranges for each problem size. Bottom: Typical Hamming distances (normalized by problem size) of local minima from the global minimum as a function of normalized residual energy. The figure demonstrates that low-lying excited states are typically $O(n)$ spin flips away from the global minimum.

We next illustrate the usefulness of equation planting in determining the bias of Ising machines tasked with




Figure 3. Left: Fraction of ground-state occurrences for a random 64-spin 3R3X instance with an 8-fold ground-state degeneracy as found by 1000 parallel tempering runs. The horizontal dashed line corresponds to uniform sampling. Also shown are the normalized Hamming distances between each ground state and all others. Middle: Typical p -values obtained from a one-sided χ^2 test performed on the ground state occurrences of randomly generated instances with 2-, 4- and 8-fold ground-state degeneracies as a function of problem size n . Right: Typical normalized Hamming distances between the ground states of degenerate instances as a function of problem size n for different nullities.

sampling the ground-state manifolds of Ising cost functions [24, 55, 56], which in addition to optimization is considered to be one of their main suggested uses [32–34, 49].

We first construct random Ising-encoded 3R3X instances with nullities $d_N = 1, 2, 3$ having ground-state degeneracies of 2, 4 and 8, respectively, at different problem sizes. We then study the functionality of PT-H as a ground-state sampler by simply measuring the fraction of occurrences of each ground state for any given instance [see Fig. 3(left)]. Each instance is solved 100 times, and the ground-state configurations are recorded.

We quantify the bias of PT-H on the various instances by measuring the typical p -value obtained from one-sided χ^2 tests performed on the tallied ground states of each instance under the hypothesis that the true distribution is uniform. The results are summarized in Fig. 3(middle). The smallness of the p -values indicates PT’s strong bias in sampling ground-state manifolds of these instances, implying differently sized ‘basins of attraction’ for each ground state. We also find that, similar to the ground state manifolds of NP-hard problems, the distances between the ground states of any given instance are typically extensive. These are depicted in Fig. 3(right) [57].

Summary and conclusions.— We introduced a method for generating random problem sets that we argued are optimally suitable for verifiably testing the utility of Ising machines. The proposed technique allows for the construction of random problem sets that have a number of properties that are desirable for this task: (i) verified ground-state energies; (ii) controllable ground-state degeneracy; (iii) ‘NP-hard-like’ energy landscapes; (iv) low-precision problem parameters and (v) underlying connectivity with a bounded degree.

We further illustrated the computational difficulties that heuristic state-of-the-art solvers face when tasked with finding or uniformly sampling the minimizing configurations of problem sets possessing intricate energy landscapes. We showed that equation-planted instances become difficult to solve already at small sizes, revealing rather early on the asymptotic exponential scaling expected from heuristic solvers. In addition, we noted that,

as is expected from difficult problem sets, Houdayer cluster moves [50] do not provide any additional advantages.

With the above stated, it should be noted that since the problem sets are generated from trivially solvable systems, special-purpose algorithms could conceivably be tailored to identify the underlying equations within the Ising cost function and utilize the equation structure to solve these Ising costs polynomially fast. On the other hand, the equation structure can be easily obfuscated if one (i) employs a wider range of gadgets to embed the various XORSAT clauses, (ii) constructs instances using a randomly weighted sum of terms, (iii) generates non-regular linear systems of equations with potentially different numbers of variables per equation or (iii) adds to existing XORSAT cost functions non-XORSAT gadgets whose ground states coincide with those of the instance, as these are known in advance. An interesting question that arises in this context is whether one can obfuscate such easily solvable problems in a manner that is provably intractable to ‘reverse engineer’ back to easily solvable form.

In the generation of the random XORSAT instances considered in this study, no specific connectivity constraints were assumed, which in turn led to the generation of Ising instances with random connections among the various spins. These instances could not be directly programmed into Ising machines that have rigid and sparse architectures [34]. This limitation may be addressed in one of two ways. The first is to convert the randomly connected instances to match the Ising machine hardware graph via embedding schemes at the price of adding auxiliary spins to the problem [58]. Alternatively, one can directly generate equation-planted instances that are native to the device connectivity, which may require devising novel 3-XORSAT gadgets.

Another aspect of equation planting that has not been explored here is the relative hardness of r -regular k -XORSAT instances for different k, r greater than three — problem classes that are expected to yield yet harder instances. Varying k, r would presumably allow for ‘hardness tunability’, which is often a desired property in the benchmarking of heuristic solvers. Another way for

achieving hardness tunability that was not pursued here is the use of non-regular random graphs, such as Poisson graphs, rather than random regular graphs [27].

The NP-hard-like energy landscapes characterizing equation-planted instances coupled with the fact that their ground state manifolds can be known in advance, may also be used as a tool for gaining insight into the mechanisms underlying the operation of Ising machines. This is because the machines' performance on these problems serves as an indicator of the locality of the heuristic being employed. As we demonstrated above, local search approaches are not very successful whereas global approaches such as Gaussian elimination are able to decipher their inherent structure very easily. Another possible use of this type of problems is perhaps also the study of spin glasses, specifically the onset of temperature chaos [43, 59, 60].

One intriguing future research direction would be the development of heuristic solvers that identify XORSAT-type relations between the spins of Ising cost functions and utilize the ease with which such sub-problems can be minimized in order to efficiently find global minimal

costs.

ACKNOWLEDGMENTS

We thank Tameem Albash and Federico Ricci-Tersenghi for useful comments and discussions. The research is based upon work (partially) supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via the U.S. Army Research Office contract W911NF-17-C-0050. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes, notwithstanding any copyright annotation thereon. This material is based on research sponsored by AFRL under agreement number FA8750-18-1-0109. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

-
- [1] B Korte and J. Vygen, *Combinatorial Optimization: Theory and Algorithms* (Springer, 2012).
 - [2] C.H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*, Dover Books on Computer Science (Dover Publications, 2013).
 - [3] T. Kadowaki and H. Nishimori, "Quantum annealing in the transverse Ising model," *Phys. Rev. E* **58**, 5355 (1998).
 - [4] 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).
 - [5] S. K. Tolpygo, V. Bolkhovsky, T. J. Weir, L. M. Johnson, M. A. Gouker, and W. D. Oliver, "Fabrication process and properties of fully-planarized deep-submicron nb/al-alox/nb josephson junctions for vlsi circuits," *IEEE Transactions on Applied Superconductivity* **25**, 1–12 (2015).
 - [6] S. K. Tolpygo, V. Bolkhovsky, T. J. Weir, C. J. Galtbraith, L. M. Johnson, M. A. Gouker, and V. K. Semenov, "Inductance of circuit structures for mit ll superconductor electronics fabrication process with 8 niobium layers," *IEEE Transactions on Applied Superconductivity* **25**, 1–5 (2015).
 - [7] X. Y. Jin, A. Kamal, A. P. Sears, T. Gudmundsen, D. Hover, J. Miloshi, R. Slattery, F. Yan, J. Yoder, T. P. Orlando, S. Gustavsson, and W. D. Oliver, "Thermal and residual excited-state population in a 3d transmon qubit," *Phys. Rev. Lett.* **114**, 240501 (2015).
 - [8] M W Johnson, P Bunyk, F Maibaum, E Tolkacheva, A J Berkley, E M Chapple, R Harris, J Johansson, T Lanting, I Perminov, E Ladizinsky, T Oh, and G Rose, "A scalable control system for a superconducting adiabatic quantum optimization processor," *Superconductor Science and Technology* **23**, 065004 (2010).
 - [9] A J Berkley, M W Johnson, P Bunyk, R Harris, J Johansson, T Lanting, E Ladizinsky, E Tolkacheva, M H S Amin, and G Rose, "A scalable readout system for a superconducting adiabatic quantum optimization system," *Superconductor Science and Technology* **23**, 105014 (2010).
 - [10] R. Harris, M. W. Johnson, T. Lanting, A. J. Berkley, J. Johansson, P. Bunyk, E. Tolkacheva, E. Ladizinsky, N. Ladizinsky, T. Oh, F. Cioata, I. Perminov, P. Spear, C. Enderud, C. Rich, S. Uchaikin, M. C. Thom, E. M. Chapple, J. Wang, B. Wilson, M. H. S. Amin, N. Dickson, K. Karimi, B. Macready, C. J. S. Truncik, and G. Rose, "Experimental investigation of an eight-qubit unit cell in a superconducting optimization processor," *Phys. Rev. B* **82**, 024511 (2010).
 - [11] P. I Bunyk, E. M. Hoskinson, M. W. Johnson, E. Tolkacheva, F. Altomare, A.J. Berkley, R. Harris, J. P. Hilton, T. Lanting, A.J. Przybysz, and J. Whittaker, "Architectural considerations in the design of a superconducting quantum annealing processor," *IEEE Transactions on Applied Superconductivity* **24**, 1–10 (Aug. 2014).
 - [12] Yoshihisa Yamamoto, Kazuyuki Aihara, Timothee Leleu, Ken-ichi Kawarabayashi, Satoshi Kako, Martin Fejer, Kyo Inoue, and Hiroki Takesue, "Coherent ising machines—optical neural networks operating at the quantum limit," *npj Quantum Information* **3**, 49 (2017).
 - [13] Peter L. McMahon, Alireza Marandi, Yoshitaka Haribara, Ryan Hamerly, Carsten Langrock, Shuhei Tamate, Takahiro Inagaki, Hiroki Takesue, Shoko Utsunomiya, Kazuyuki Aihara, Robert L. Byer, M. M. Fejer, Hideo Mabuchi, and Yoshihisa Yamamoto, "A fully-

- programmable 100-spin coherent ising machine with all-to-all connections,” *Science* **354**, 614–617 (2016).
- [14] Maliheh Aramon, Gili Rosenberg, Elisabetta Valiante, Toshiyuki Miyazawa, Hirotaka Tamura, and Helmut G. Katzgraber, “Physics-Inspired Optimization for Quadratic Unconstrained Problems Using a Digital Annealer,” *arXiv e-prints*, arXiv:1806.08815 (2018), arXiv:1806.08815 [physics.comp-ph].
 - [15] Massimiliano Di Ventra and Fabio L. Traversa, “Perspective: Memcomputing: Leveraging memory and physics to compute efficiently,” *Journal of Applied Physics* **123**, 180901 (2018).
 - [16] Fabio L. Traversa and Massimiliano Di Ventra, “Mem-Computing Integer Linear Programming,” *arXiv e-prints*, arXiv:1808.09999 (2018), arXiv:1808.09999 [cs.ET].
 - [17] Fabio Lorenzo Traversa, Chiara Ramella, Fabrizio Bonani, and Massimiliano Di Ventra, “Memcomputing np-complete problems in polynomial time using polynomial resources and collective states,” *Science Advances* **1**, e1500031 (2015).
 - [18] Yoshitaka Haribara, Hitoshi Ishikawa, Shoko Utsunomiya, Kazuyuki Aihara, and Yoshihisa Yamamoto, “Performance evaluation of coherent ising machines against classical neural networks,” *Quantum Science and Technology* **2**, 044002 (2017).
 - [19] Takahiro Inagaki, Yoshitaka Haribara, Koji Igarashi, Tomohiro Sonobe, Shuhei Tamate, Toshimori Honjo, Alireza Marandi, Peter L. McMahon, Takeshi Umeki, Koji Enbutsu, Osamu Tadanaga, Hirokazu Takenouchi, Kazuyuki Aihara, Ken-ichi Kawarabayashi, Kyo Inoue, Shoko Utsunomiya, and Hiroki Takesue, “A coherent ising machine for 2000-node optimization problems,” *Science* **354**, 603–606 (2016).
 - [20] Itay Hen, Joshua Job, Tameem Albash, Troels F. Rønnow, Matthias Troyer, and Daniel A. Lidar, “Probing for quantum speedup in spin-glass problems with planted solutions,” *Phys. Rev. A* **92**, 042325 (2015).
 - [21] James King, Sheir Yarkoni, Jack Raymond, Isil Ozfidan, Andrew D. King, Mayssam Mohammadi Nevisi, Jeremy P. Hilton, and Catherine C. McGeoch, “Quantum annealing amid local ruggedness and global frustration,” *Journal of the Physical Society of Japan* **88**, 061007 (2019).
 - [22] Jeffrey Marshall, Victor Martin-Mayor, and Itay Hen, “Practical engineering of hard spin-glass instances,” *Phys. Rev. A* **94**, 012320 (2016).
 - [23] Salvatore Mandrà and Helmut G. Katzgraber, “A deceptive step towards quantum speedup detection,” *Quantum Science and Technology* **3**, 04LT01 (2018).
 - [24] V. Martín-Mayor and I. Hen, “Unraveling quantum annealers using classical hardness,” *Scientific Reports* **5**, 15324 (2015).
 - [25] Davide Venturelli, Salvatore Mandrà, Sergey Knysh, Bryan O’Gorman, Rupak Biswas, and Vadim Smelyanskiy, “Quantum optimization of fully connected spin glasses,” *Phys. Rev. X* **5**, 031040 (2015).
 - [26] S. Franz, M. Mézard, F. Ricci-Tersenghi, M. Weight, and R. Zecchina, “A ferromagnet with a glass transition,” *Europhys. Lett.* **55**, 465 (2001).
 - [27] W. Barthel, A. K. Hartmann, M. Leone, F. Ricci-Tersenghi, M. Weight, and R. Zecchina, “Hiding solutions in random satisfiability problems: A statistical mechanics approach,” *Phys. Rev. Lett.* **88**, 188701 (2002).
 - [28] 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.* **104**, 207206 (2010), (arXiv:0911.3438).
 - [29] Federico Ricci-Tersenghi, “Being glassy without being hard to solve,” *Science* **330**, 1639–1640 (2010).
 - [30] Marco Guidetti and A. P. Young, “Complexity of several constraint satisfaction problems using the heuristic, classical, algorithm, WalkSAT,” *Phys. Rev. E* **84**, 011102 (2011), (arXiv:1102.5152).
 - [31] Edward Farhi, David Gosset, Itay Hen, A. W. Sandvik, Peter Shor, A. P. Young, and Francesco Zamponi, “Performance of the quantum adiabatic algorithm on random instances of two optimization problems on regular hypergraphs,” *Phys. Rev. A* **86**, 052334 (2012).
 - [32] S. H. Adachi and M. P. Henderson, “Application of Quantum Annealing to Training of Deep Neural Networks,” *ArXiv e-prints* (2015), arXiv:1510.06356 [quant-ph].
 - [33] Marcello Benedetti, John Realpe-Gómez, Rupak Biswas, and Alejandro Perdomo-Ortiz, “Estimation of effective temperatures in quantum annealers for sampling applications: A case study with possible applications in deep learning,” *Phys. Rev. A* **94**, 022308 (2016).
 - [34] Jeffrey Marshall, Eleanor G. Rieffel, and Itay Hen, “Thermalization, freeze-out, and noise: Deciphering experimental quantum annealers,” *Phys. Rev. Applied* **8**, 064025 (2017).
 - [35] David C. Lay, *Linear Algebra and Its Applications* (3rd ed.) (Addison Wesley, Boston, MA, 2005).
 - [36] Ryan Babbush, Bryan O’Gorman, and Alán Aspuru-Guzik, “Resource efficient gadgets for compiling adiabatic quantum optimization problems,” *Annalen der Physik* **525**, 877–888 (2013).
 - [37] Tameem Albash, Victor Martin-Mayor, and Itay Hen, “Temperature scaling law for quantum annealing optimizers,” *Phys. Rev. Lett.* **119**, 110502 (2017).
 - [38] Tameem Albash, Victor Martin-Mayor, and Itay Hen, “Analog errors in ising machines,” *Quantum Science and Technology* **4**, 02LT03 (2019).
 - [39] See Supplemental Material at [URL] for additional details.
 - [40] K. Hukushima and K. Nemoto, “Exchange Monte Carlo method and application to spin glass simulations,” *J. Phys. Soc. Japan* **65**, 1604 (1996), (arXiv:cond-mat/9512035).
 - [41] E. Marinari, “Optimized Monte Carlo methods,” in *Advances in Computer Simulation*, edited by J. Kertész and I. Kondor (Springer-Verlag, 1998) p. 50, (arXiv:cond-mat/9612010).
 - [42] S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi, “Optimization by simulated annealing,” *Science* **220**, 671 (1983).
 - [43] L. A. Fernandez, V. Martin-Mayor, G. Parisi, and B. Seoane, “Temperature chaos in 3d ising spin glasses is driven by rare events,” *EPL (Europhysics Letters)* **103**, 67003 (2013).
 - [44] L. A. Fernandez and V. Martin-Mayor, “Exact identities as an opportunity for Monte Carlo error reduction,” *Phys. Rev. E* **79**, 051109 (2009), (arXiv:0904.3415).
 - [45] L. A. Fernandez, V. Martin-Mayor, S. Perez-Gaviro, A. Tarancon, and A. P. Young, “Phase transition in the three dimensional Heisenberg spin glass: Finite-size scaling analysis,” *Phys. Rev. B* **80**, 024422 (2009), (arXiv:0905.0322).
 - [46] H. G. Ballesteros, A. Cruz, L. A. Fernández, V. Martín-

- Mayor, J. Pech, J. J. Ruiz-Lorenzo, A. Tarancón, P. Téllez, C. L. Ullod, and C. Ungil, “Critical behavior of the three-dimensional ising spin glass,” *Phys. Rev. B* **62**, 14237–14245 (2000).
- [47] Helmut G Katzgraber, Simon Trebst, David A Huse, and Matthias Troyer, “Feedback-optimized parallel tempering monte carlo,” *Journal of Statistical Mechanics: Theory and Experiment* **2006**, P03018–P03018 (2006).
- [48] David J. Earl and Michael W. Deem, “Parallel tempering: Theory, applications, and new perspectives,” *Phys. Chem. Chem. Phys.* **7**, 3910–3916 (2005).
- [49] Brian Hu Zhang, Gene Wagenbreth, Victor Martin-Mayor, and Itay Hen, “Advantages of unfair quantum ground-state sampling,” *Scientific Reports* **7**, 1044 (2017).
- [50] J. Houdayer, “The wormhole move: A new algorithm for polymer simulations,” *The Journal of Chemical Physics* **116**, 1783–1787 (2002).
- [51] Helmut G. Katzgraber, L. W. Lee, and A. P. Young, “Correlation length of the two-dimensional ising spin glass with gaussian interactions,” *Phys. Rev. B* **70**, 014417 (2004).
- [52] Jérôme Houdayer and Alexander K. Hartmann, “Low-temperature behavior of two-dimensional gaussian ising spin glasses,” *Phys. Rev. B* **70**, 014418 (2004).
- [53] Tameem Albash and Daniel A. Lidar, “Demonstration of a scaling advantage for a quantum annealer over simulated annealing,” *Phys. Rev. X* **8**, 031016 (2018).
- [54] Both algorithms were run on one core of a 3.5 GHz 6-Core Intel Xeon E5 processor.
- [55] S. Bastea and P. M. Duxbury, “Ground state structure of random magnets,” *Phys. Rev. E* **58**, 4261–4265 (1998).
- [56] Yoshiki Matsuda, Hidetoshi Nishimori, and Helmut G Katzgraber, “Quantum annealing for problems with ground-state degeneracy,” *Journal of Physics: Conference Series* **143**, 012003 (2009).
- [57] The bias exhibited by PT may seem at first to contradict the Boltzmann distribution which (in accordance with the assumption of equal a priori probability) prescribes the same probability to same-cost configurations for thermalized systems. However, we note that when PT is used as an optimizer (as it is here), the number of sweeps is not large enough for thermalization to take place.
- [58] Jun Cai, William G. Macready, and Aidan Roy, “A practical heuristic for finding graph minors,” *arXiv e-prints*, arXiv:1406.2741 (2014), arXiv:1406.2741 [quant-ph].
- [59] T. Aspelmeier, A. J. Bray, and M. A. Moore, “Why temperature chaos in spin glasses is hard to observe,” *Phys. Rev. Lett.* **89**, 197202 (2002), (arXiv:cond-mat/0207300).
- [60] L. Berthier and A. P. Young, “Temperature cycles in the Heisenberg spin glass,” *Phys. Rev. B* **71**, 214429 (2005), (arXiv:cond-mat/0503012).