

SCIENTIFIC REPORTS



OPEN

Advantages of Unfair Quantum Ground-State Sampling

Brian Hu Zhang¹, Gene Wagenbreth², Victor Martin-Mayor^{3,4} & Itay Hen^{5,6}

The debate around the potential superiority of quantum annealers over their classical counterparts has been ongoing since the inception of the field. Recent technological breakthroughs, which have led to the manufacture of experimental prototypes of quantum annealing optimizers with sizes approaching the practical regime, have reignited this discussion. However, the demonstration of quantum annealing speedups remains to this day an elusive albeit coveted goal. We examine the power of quantum annealers to provide a different type of quantum enhancement of practical relevance, namely, their ability to serve as useful samplers from the ground-state manifolds of combinatorial optimization problems. We study, both numerically by simulating stoquastic and non-stoquastic quantum annealing processes, and experimentally, using a prototypical quantum annealing processor, the ability of quantum annealers to sample the ground-states of spin glasses differently than thermal samplers. We demonstrate that (i) quantum annealers sample the ground-state manifolds of spin glasses very differently than thermal optimizers (ii) the nature of the quantum fluctuations driving the annealing process has a decisive effect on the final distribution, and (iii) the experimental quantum annealer samples ground-state manifolds significantly differently than thermal and ideal quantum annealers. We illustrate how quantum annealers may serve as powerful tools when complementing standard sampling algorithms.

Many problems of practical importance may be cast as a task of finding all the minimizing configurations, or ground-states, of a given cost function. Examples are numerous—among them are SAT filtering¹, hardware fault detection and the verification and validation (V&V) of safety-critical cyber-physical systems², to mention a few. In the V&V of safety-critical cyber-physical systems for instance, one is concerned with testing whether a given piece of software contains a bug. This problem can naturally be cast as a constraint satisfaction problem² (equivalently, as an optimization problem of the Ising-type) where finding as many of the bugs as possible is critical to the success of the mission. Similar scenarios occur in circuit fault detection where each solution corresponds to a potential discrepancy in the implementation of a circuit and where all discrepancies must be found.

The listing of all solutions of a given cost function is, however, generally an intractable task for standard algorithms (it is a problem in the complexity class #P). This is not only because of the difficulty involved in finding an optimum³, but also because of the sheer number of ground-states of the problem which may grow exponentially with input size (a property known to physicists as a non-vanishing entropy density⁴). Furthermore, the energy landscapes of certain cost functions are known to bias heuristic optimizers, as well as provable solvers, towards certain solutions and away from others⁵. Thus, the practical importance of sampling from the set of ground-states of intricate cost functions in qualitatively diverse manners is immense—both from the theoretical point of view and for practical reasons. In the context of V&V for instance, one hopes that employing a suite of qualitatively dissimilar sampling algorithms will unearth nonidentical or even disjoint sets of solutions, leading eventually to the discovery of much larger sets of bugs.

Recent technological breakthroughs that have made experimental programmable quantum annealing (QA) optimizers containing thousands of quantum bits^{6,7} available, have rekindled the interest in annealers as a revolutionary new approach to finding the minimizing assignments of discrete combinatorial cost functions. Quantum annealers^{8,9} provide a unique approach to finding the ground-states of discrete optimization problems, utilizing

¹Stanford University, Stanford, California, 94305, USA. ²Cray, Seattle, WA 98164, USA. ³Departamento de Física Teórica I, Universidad Complutense, 28040, Madrid, Spain. ⁴Instituto de Biocomputación y Física de Sistemas Complejos (BIFI), Zaragoza, Spain. ⁵Information Sciences Institute, University of Southern California, Marina del Rey, California, 90292, USA. ⁶Department of Physics and Astronomy and Center for Quantum Information Science & Technology, University of Southern California, Los Angeles, California, 90089, USA. Brian Hu Zhang, Gene Wagenbreth, Victor Martin-Mayor and Itay Hen contributed equally to this work. Correspondence and requests for materials should be addressed to I.H. (email: itayhen@isi.edu)

Received: 26 January 2017

Accepted: 21 March 2017

Published online: 21 April 2017

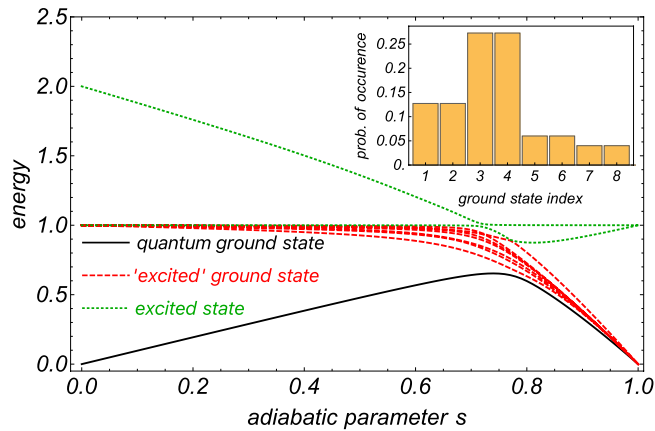


Figure 1. Ten lowest energy levels of an 8-qubit Hamiltonian interpolating between a transverse field driver Hamiltonian and a randomly generated Ising Hamiltonian. The solid black line indicates the energy of the instantaneous ground-state. The dashed red and dotted green lines indicate excited states that lead to final ground-states and final excited states, respectively. Inset: The probabilities for obtaining the various classical ground-states upon measuring the quantum ground-state at the end of the evolution in the computational basis.

gradually decreasing quantum fluctuations to traverse barriers in the energy landscape in search of global optima, a mechanism commonly believed to have no classical counterpart^{8–14}. In the context of ground-state sampling, quantum annealers thus offer the exciting possibility of discovering minimizing assignments that cannot be reached in practice with standard algorithms, potentially offering unique advantages over traditional algorithms for solving problems of practical importance.

Here, we put this hypothesis to the test by directly addressing the question of whether quantum annealers can sample the ground-state set of optimization problems differently than their classical counterparts. We further examine the potential inherent in quantum annealers to serve as useful tools in practical settings. We demonstrate, both via numerical simulations and experimentally, by testing a 512-qubit D-Wave Two quantum annealing optimizer^{6,7}, that quantum annealers not only produce different distributions over the set of ground-states than simulated thermal annealing optimizers, but that they offer an additional dimension of tunability that does not necessarily have a classical counterpart. Finally, we show that when used in conjunction with existing standard ground-state-sampling or solution counting algorithms, quantum annealers may offer certain unique advantages that may not be otherwise achievable.

Sampling the ground-state manifold of spin glasses

Similar to standard classical algorithms, quantum annealers—even ideal fully adiabatic ones held at zero temperature—when tasked with solving optimization problems will generally sample the solution space of optimization problems in a biased manner, producing certain ground-states more frequently than others. Unlike the bias exhibited by thermal algorithms, the uneven sampling of quantum annealers has its origins in the quantum nature of their dynamics: In standard quantum annealing protocols, one engineers a smoothly interpolating Hamiltonian between a simple ‘driver’ Hamiltonian H_d which provides the quantum fluctuations and a classical ‘problem’ Hamiltonian H_p that is diagonal in the computational basis and whose ground-states encode the solutions of an optimization problem

$$H(s) = (1 - s)H_d + sH_p, \quad (1)$$

where $s(t)$ is a parameter varying smoothly with time from 0 at $t=0$ to 1 at the end of the algorithm, at $t = \mathcal{T}$ [the type of problem and driver Hamiltonians we shall consider are given in Eqs (3), (4) and (5) in the next section].

In the presence of degeneracy in the ground-state manifold of the problem Hamiltonian (which corresponds to multiple minimizing assignments for the cost function), the adiabatic theorem ensures that the state reached at the end of the adiabatic evolution in the limit $\lim_{s \rightarrow 1} H(s)$ is still uniquely defined. At the end of the QA evolution, the final state corresponds to a specific linear combination of the classical ground-states

$$|\psi_{GS}\rangle = \sum_{i=1}^D c_i |\phi_i\rangle \quad (2)$$

where $\{|\phi_1\rangle, |\phi_2\rangle, \dots, |\phi_D\rangle\}$ is the set of D classical ground states, or minimizing configurations, of the optimization problem (see Fig. 1 for an example). The $\{c_i\}^2$ are the probabilities for obtaining each of these classical ground-states upon computational-basis measurements at the end of the anneal. These define a probability distribution over the ground-state manifold and depend not only on the structure of the problem Hamiltonian but also on the nature of the quantum fluctuations provided by the driver Hamiltonian (from the point of view of quantum perturbation theory, in the simplest case where first-order perturbation theory breaks all degeneracies,

the ground-state in Equation (2) is merely the ground-state of a restricted driver Hamiltonian, specifically, the driver H_d projected onto the subspace spanned by the ground-states of H_p .

Since the distribution of minimizing configurations (henceforth, the ground-state distribution, or GSD) generated by a quantum annealer is *intrinsically quantum*, a possibility arises that some quantum GSDs cannot be efficiently generated by classical samplers. Moreover, that the choice of driver Hamiltonian H_d determines these GSDs, offers a tunable handle, or an extra knob, that potentially produces a continuum of probability distributions over the ground-state configurations. It is therefore plausible to assume that certain classically suppressed configurations, i.e., solutions that have very low probabilities of being found via thermal or other classical processes, may have high probabilities of being found or sampled with suitable choices of driver Hamiltonians^{15–18}. In such cases, quantum annealers may be used to replace or *complement* classical samplers, giving rise to a novel form of quantum enhancements.

To test whether quantum annealers indeed provide a potentially powerful platform for achieving quantum enhancements for the counting or listing of solutions of hard optimization problems, we study in detail their capabilities to sample ground-state configurations differently than their classical counterparts. As a testbed, we consider spin glasses: disordered, frustrated spin systems¹⁹ that may be viewed as prototypical classically hard (also called NP-hard) optimization problems²⁰ focusing, for reasons that will become clear later, on problems whose ground-state configurations have been computed in advance. These will be used to test the performance of classical thermal annealers, comparing the outcomes of these against the GSDs produced by ideal zero-temperature stoquastic as well as non-stoquastic quantum annealers. We shall also compare the produced GSDs to those obtained by a prototypical experimental quantum annealer—the D-Wave Two (DW2) processor^{6, 21} which consists of an array of superconducting flux qubits designed to solve Ising model instances defined on the graph hardware via a gradually decreasing transverse field Hamiltonian (further details, including a visualization of the Chimera graph and the annealing schedule used to interpolate between H_d and H_p , are provided in the Supplementary Information). These comparisons will provide insight into the potential computational power inherent in quantum devices to assist traditional algorithm in finding all (or as many as possible) minimizing configurations of discrete optimization problems.

Results: Thermal vs. quantum ground-state sampling

Setup. We generate random spin glass instances for which the enumeration of minimizing configurations is a feasible task. This a priori requirement will allow us to properly evaluate the sampling capabilities of the tested algorithms in an unbiased way. To do that, we consider optimization problems whose cost function is of the form:

$$H_p = \sum_{\langle ij \rangle} J_{ij} s_i s_j. \quad (3)$$

The Ising spins, $s_i = \pm 1$, are the variables to be optimized over, and the set of parameters $\{J_{ij}\}$ determines the cost function. To experimentally test the generated instances, we take $\langle ij \rangle$ to sum over the edges of an $N = 504$ -qubit Chimera graph—the hardware graph of the DW2 processor²¹.

The precise enumeration of all minimizing configurations of spin glass instances of more than a few dozen spins is generally an intractable task. We overcome this difficulty here by generating problem instances with *planted solutions*^{22, 23}. As we discuss in the Methods section, the structure of planted-solution instances allows for the development of a constraint solving bucket algorithm capable of enumerating all minimizing assignments. By doing so, we obtain about 2000 optimization 504-bit spin glass instances, each with less than 500 minimizing configurations, and for which we know *all* ground-state configurations. This enables the accurate evaluation of the distributions of success probabilities of *individual minimizing configurations*. We compare the GSDs as obtained by several different algorithms:

1. *Simulated annealing (SA)*—This is a well-known, powerful and generic heuristic solver²⁴. Our SA algorithm uses single spin-flip Metropolis updates with a linear profile of inverse temperatures $\beta = T^{-1}$, going from $\beta_{\min} = 0$ to $\beta_{\max} = 20$ (with β updated after every Metropolis sweep over the lattice spins). Figure 2 provides the SA results for a couple of typical instances differing in number of minimizing configurations. The figure shows the average time for each individual minimizing bit assignment plotted as a function of number of sweeps on a log-log scale. For any fixed number of sweeps, we find that the probability of obtaining certain ground-state configurations may vary by orders of magnitude. At first glance, the ‘unfair’ sampling of SA may seem to contradict the Boltzmann distribution which (in accordance with the assumption of equal a priori probability) prescribes the same probability to same-energy configurations for thermalized systems. However, we note that when SA is used as an optimizer (as it is here), the number of sweeps is not large enough for thermalization to take place. As an optimizer, the number of SA Metropolis sweeps is chosen such that on average the time to find a ground-state is minimized. The minimum time-to-solution (for individual solutions) is evident in Fig. 2, as is the convergence of all individual success probabilities into a single value in the limit of long annealing times, consistently with Boltzmann distributions.
2. *Ideal zero-temperature quantum annealer with a transverse field driver (TFQA)*—We also consider an ideal, zero-temperature, fully adiabatic quantum annealer with a transverse-field driver Hamiltonian, namely

$$H_d^{\text{TF}} = -\sum_i \sigma_i^x, \quad (4)$$

where σ_i^x is the Pauli X spin-1/2 matrix acting on spin i . The adiabatic process interpolates linearly between the above driver and the final Ising Hamiltonian as described in Eq. (1). The algorithm we devise to infer the quantum ground-state is discussed in detail in Methods.

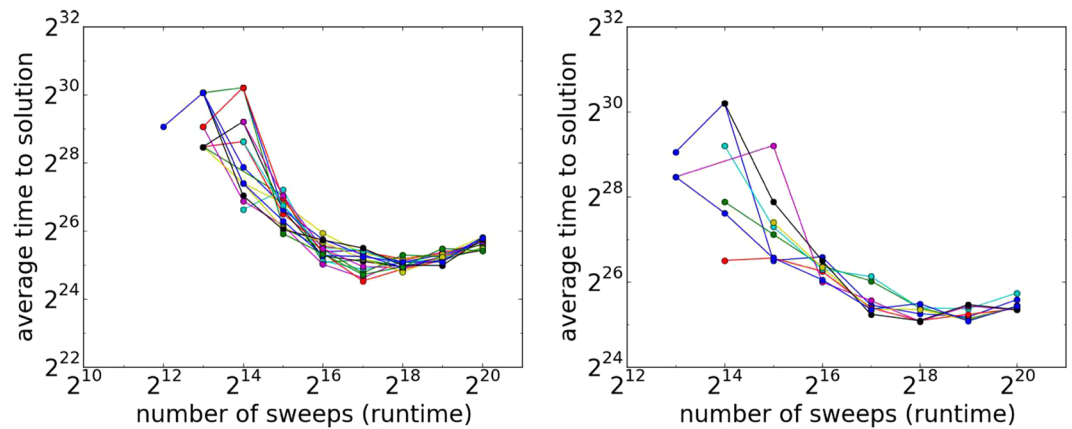


Figure 2. Simulated annealing average time to solution—the ratio of success probability to number of sweeps, vs. number of sweeps for different solutions (log-log scale). Each line represents a different solution. Left: A 504-bit instance with 16 solutions. Right: A 504-bit instance with 4 solutions. For most instances, certain solutions are reached considerably sooner than others.

	SA	TFQA	NSQA	DW2
SA	—	67%	65%	93%
TFQA		—	57%	99%
NSQA			—	99%
DW2				—

Table 1. Fraction of instances with statistically significant differences in GSD between any two optimizers. Here, the p -value is set at $p=0.01$.

3. *Ideal zero-temperature quantum annealer with a non-stoquastic driver (NSQA)*—As already discussed above, quantum annealers offer more control than thermal annealers over the generated fluctuations. This is because unlike thermal fluctuations, quantum fluctuations can be *engineered* (see, e.g., refs 25, 26). One thus expects different driver Hamiltonians to yield different probability distributions over the ground-state manifold. Of particular interest are the so-called *non-stoquastic* driver Hamiltonians, which cannot be efficiently simulated by classical algorithms. As a test case, we consider an additional quantum annealing process driven by:

$$H_d^{\text{NS}} = -\sum_i \sigma_i^x + \sum_{(ij)} \tilde{J}_{ij} \sigma_i^x \sigma_j^x. \quad (5)$$

To ensure that the driver is non-stoquastic, we choose the couplings \tilde{J}_{ij} to be the same as the $\sigma_i^z \sigma_j^z$ couplings of the problem Hamiltonian, but with an arbitrarily chosen sign.

4. *Experimental D-Wave Two processor (DW2)*—We also feed the generated instances to the putative DW2 quantum annealing optimizer. This device is designed to solve optimization problems by evolving a known initial configuration—the ground-state of a transverse field towards the ground-state of the classical Ising-model Hamiltonian of Eq. (3). Each problem instance was run on the annealer for a minimum of 10^4 anneals with each anneal lasting 20 to 40 μs , for a total of more than 10^7 anneals. Each anneal ends up with a measurement in the computational basis yielding either an excited state or a classical ground-state which is subsequently recorded and which is later used to construct a GSD. To overcome the inhomogeneity of the processor as well as other systematic errors, each anneal is carried out with a randomly generated gauge (see ref. 27 for more details).

Distinguishing Probability Distributions. We test whether the different algorithms that we consider produce significantly different GSDs, or probability distributions over the ground-state manifold, on the various spin glass instances. To distinguish between two distributions generated by two methods, at least one of which is empirically estimated via experiment, we use a bootstrapped version of the Kolmogorov-Smirnov (KS) test (see Methods).

Table 1 summarizes the results of the statistical tests, listing the fraction of instances with different GSDs between any two tested optimizers. As is evident from the table, the distributions generated by the various algorithms are in general significantly different from one another—pointing to presumably different physical mechanisms generating them, namely thermal or quantum fluctuations of different sources. As we discuss later on, these pronounced differences in the GSDs can allow quantum annealers to serve as potentially powerful tools,

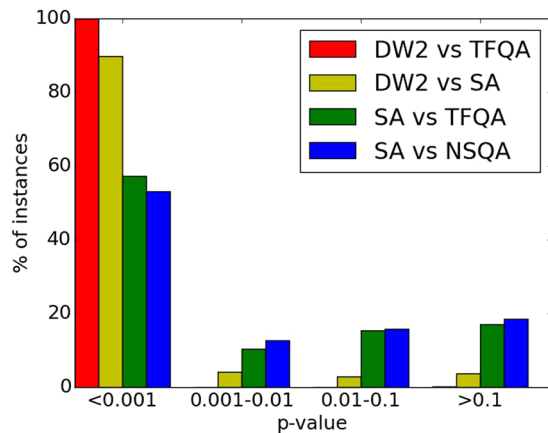


Figure 3. The p -values generated by the Kolmogorov-Smirnov tests to quantify the differences between pairs of algorithms.

when combined with standard techniques, for finding all (or as many as possible) minimizing configurations of combinatorial optimization problems.

To quantify the utility of using a combination of two (or more) methods for finding as many minimizing configurations as possible, we define the ‘bias’ $b(\mathbf{p})$ of a GSD as

$$b(\mathbf{p}) = \frac{D}{2(D-1)} \sum_{i=1}^D |p_i - 1/D|, \quad (6)$$

where we denote probability distributions by \mathbf{p} , and where p_i is the probability of obtaining ground-state i . Here, a flat distribution corresponds to $b(\mathbf{p}) = 0$, whereas an extremely biased one for which all samples are multiples of the same configuration yields $b(\mathbf{p}) = 1$. If n applications of one optimization method yield a probability distribution $\mathbf{p}^{(1)}$ and n applications of another yield $\mathbf{p}^{(2)}$, then a combined effort of $n/2$ samples from each will yield a GSD with a bias $b(\bar{\mathbf{p}})$ where $\bar{\mathbf{p}} = \frac{1}{2}(\mathbf{p}^{(1)} + \mathbf{p}^{(2)})$. We can therefore quantitatively measure the utility of a combination of two (or more) methods by comparing the bias $b(\bar{\mathbf{p}})$ against that of any one method $b(\mathbf{p}^{(1)})$ [or $b(\mathbf{p}^{(2)})$]. The smaller the bias of the combination, the greater the utility of using the two methods in conjunction. Let us next examine the differences between specific pairs of methods in more detail.

SA vs. TFQA and NSQA. We first compare the GSDs obtained by thermal simulated annealing (SA) against those generated by the transverse-field and the non-stoquastic quantum annealers. A bootstrapped KS test to decide whether they are significantly different suggests a difference that is significant at the $p < 0.01$ level in 67% of the instances with TFQA and 65% of the instances with NSQA. These results are summarized in Fig. 3. In the vast majority of cases, there is a qualitative difference between the results produced by QA and those produced by SA.

Figure 4 depicts the GSDs of several representative instances illustrating the little or no relationship between the thermal and TFQA or NSQA methods. In the left panel we find an instance for which both SA and TFQA produce similar GSDs. The middle and right instances show no clear relationship between SA and the other algorithms.

We now ask whether using a quantum annealer *together* with a simulated annealer has merit. To that aim, we compare the SA bias $b_{\text{SA}} = b(\mathbf{p}_{\text{SA}})$ against the biases of the combinations SA with TFQA and SA with NSQA. The results are summarized in Fig. 5. In these scatter plots, any data point below the $y = x$ line indicates an advantage to using ‘assisted’ ground-state sampling driven by quantum samplers. As can be clearly observed, for most of the instances the bias of the combination is significantly lower than that of using SA alone. The median SA bias of $\langle b_{\text{SA}} \rangle = 0.10(1)$ drops to $\langle b_{\text{SA}+\text{TFQA}} \rangle = 0.075(0)$ and $\langle b_{\text{SA}+\text{NSQA}} \rangle = 0.074(2)$ when assisted with TFQA and NSQA, respectively. Also noticeable is the $y = x/2$ line; data points on this line are obtained whenever SA is used in conjunction with a method that yields a flat distribution, in which case the bias is halved. We find, perhaps not surprisingly, that ideal zero-temperature annealing yields flat distributions for many of the tested instances [see also Fig. 6 (left)]. This is somewhat analogous to an ideal SA process where full thermalization is reached, in which case the generated GSDs would all be balanced.

TFQA vs. NSQA. Next, we compare the GSDs produced by the ideal zero-temperature quantum annealers, namely the transverse-field annealer (TFQA) against the non-stoquastic annealer (NSQA). As mentioned earlier, for more than half of the instances (1105 out of 1909) the quantum ground-states are found to be ‘flat’ for both annealers (the median bias for both processes was found to be $< 10^{-4}$). Discounting for those, we find in general that the chosen driver has a decisive effect on the GSDs. The GSDs of representative instances are shown in Fig. 7, showcasing the tunability of the probability distributions with respect to the ‘knob’ of quantum fluctuations. For the instance on the left, both drivers sample the ground-state manifold similarly; for the middle and right

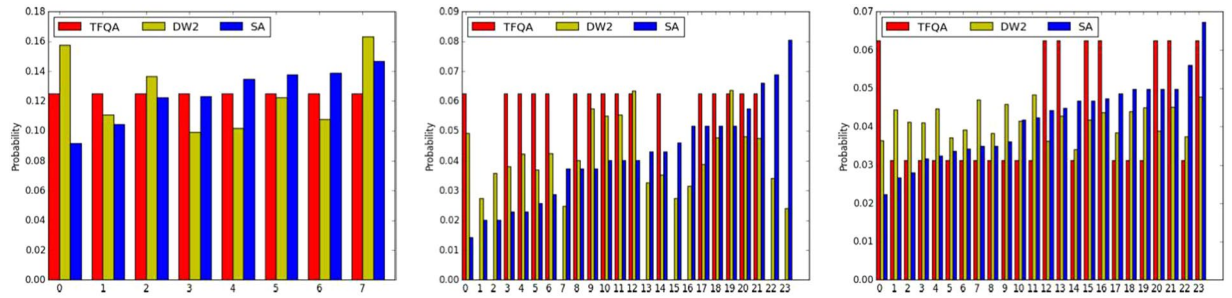


Figure 4. Three representative GSDs for simulated annealing (SA - blue), ideal transverse-field quantum annealer (TFQA - red) and the D-Wave Two experimental processor (DW2 - yellow). In the left instance, probabilities for obtaining the various ground-states predict that all solutions are about equally likely. In the middle instance, we observe that those ground-states that our analytic TFQA predicts should appear more often, and do indeed appear more often in the experimental DW2 data. In the right instance, there is no clear relationship between the various algorithms.

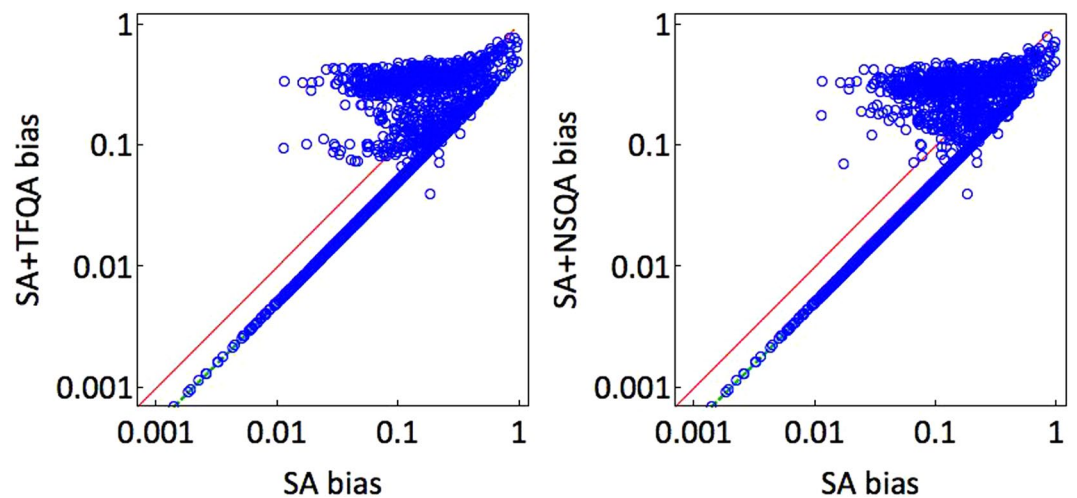


Figure 5. Scatter plots of biases of the tested instances' GSDs as obtained by SA vs. a combination of SA with ideal quantum annealers. Left: SA vs. SA + TFQA. Right: SA vs. SA + NSQA. In the majority of cases, a combination of the methods leads a smaller overall bias, i.e., a lesser degree of unfairness. The solid red line is the 'equal bias' $y = x$ line, whereas the dashed green $y = x/2$ line is the bias obtained when SA is combined with a flat, unbiased GSD, in which case the bias is halved.

instances, we observe that the driver has a profound effect on the shape of the distribution. For a fraction of the instances ($<5\%$), ground-states that had zero probability for TFQA had strictly positive probability when NSQA was used. That is, the driver often has an effect not only on the probabilities of various states but also on which states have nonzero amplitudes.

Comparing the bias associated with use of a transverse-field quantum annealer with and without the aid of a non-stoquastic driver, we find that the combination of annealers is far less biased than the transverse-field annealer alone. On average, as shown in Fig. 6 (right), NSQA distributions are slightly more biased than TFQA distributions. Nonetheless, as is indicated in the figure, which shows a scatter plot of the bias $b(\mathbf{p}_{\text{TFQA}})$ against the biases of the combination of TFQA with NSQA, there is merit in 'assisted' non-stoquastic ground-state sampling.

Experimental DW2 vs. SA and TFQA. Comparing the distributions generated by the experimental quantum annealer DW2 against SA and against an ideal transverse field quantum annealer (TFQA), we find, as in the other comparisons, only a weak relationship between the output distributions. As summarized in Fig. 3, in almost all instances, the KS test yielded a significant difference between the GSDs (also see Table 1). Figure 4 shows the GSDs for three representative instances. On the left panel, we find that DW2 as well as SA and TFQA yields an approximately flat probability distribution over the various ground-states. In the middle panel, we find those ground-states that TFQA predicts will appear more often, and indeed appeared more often in the D-Wave experimental GSD. On the right panel, no apparent relationship is found between the three GSDs.

To understand whether there is merit in using a DW2 processor alongside an SA algorithm for the purpose of producing a more balanced distribution of ground-states, we compare the bias of using SA alone vs. using SA together with DW2. The results are summarized in Fig. 8. As the left panel shows, SA and DW2 produce biases with similar distributions. Interestingly, we find that while for some instances use of both methods does reduce

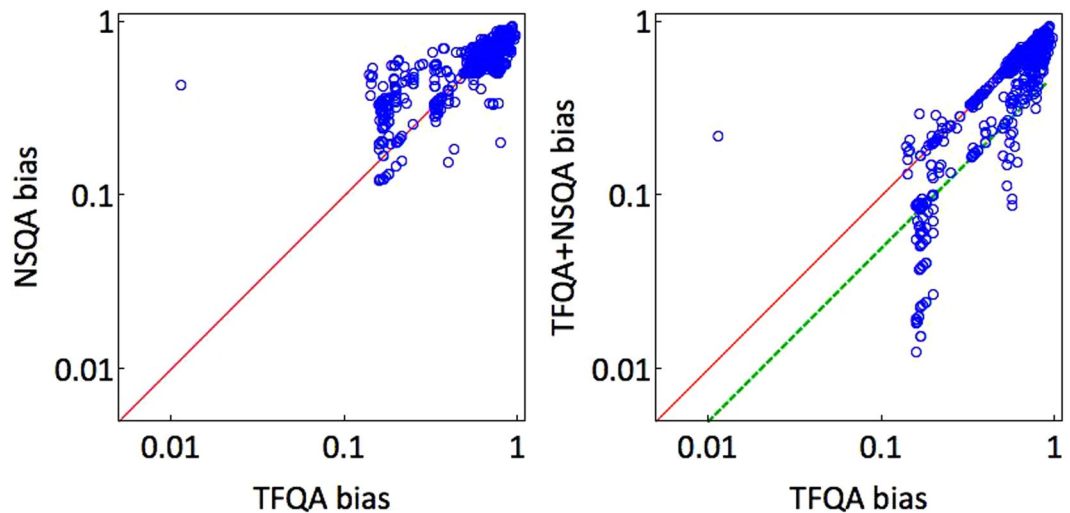


Figure 6. Scatter plots of biases of the tested instances' quantum GSDs. Left: Biases of TFQA annealing vs. NSQA annealing showing that the non-stoquastic driver is slightly more biased in general than the transverse field driver (not shown in the figure are the completely unbiased instances which constitute about 57% of the instances). Right: Biases of the TFQA GSDs vs. those of the combined application of TFQA and NSQA indicating that use of an additional non-stoquastic driver considerably reduces the bias of the GSDs.

the bias, for many others it does not (right panel). When assisted with the DW2 experimental annealer, the median SA bias of $\langle b_{SA} \rangle = 0.10(1)$ remains at $\langle b_{SA+DW2} \rangle = 0.10(1)$. As we shall discuss next in more detail, while significant differences in the GSDs seem to bode well for the utility of DW2 to generate possible classically suppressed minimizing configurations with high probabilities, the origin of these differences is unclear.

Discussion

In this work we studied the capabilities of quantum annealers to sample the ground-state manifolds of degenerate spin glass optimization problems. We addressed the question of how differently ideal zero-temperature and experimental quantum annealers sample the minimizing configurations of optimization problems than the standard algorithms, specifically thermal annealers. Examining both stoquastic and non-stoquastic quantum fluctuations, we illustrated that quantum annealers produce, in general, qualitatively very different probability distributions than classical annealers; furthermore, that the final distribution depends heavily on the nature of the quantum fluctuations. Moreover, we have shown that using quantum annealers alongside thermal algorithms produces, in general, flatter distributions of ground-states; that is, a combined use is significantly more helpful in generating more ground-states than when using classical algorithms alone.

An earlier work by Matsuda *et al.*¹⁷ is worth mentioning here, where a five qubit example has been studied to show that usage of a transverse-field QA may result in the uneven sampling of a classical ground state manifold. Specifically, the authors found that some classical ground state configurations were unreachable via that driver, whereas more sophisticated drivers that ensured that any two states in the computational basis had non-vanishing matrix elements, resolved that matter. A major caveat pointed out by the authors of ref. 17, though, was that the non-locality of their enhanced driver Hamiltonian had made it very difficult to study numerically at large system sizes (in addition, non-local terms are also known to be problematic to implement physically). We also note an experimental study of the newer-generation DW2X chip¹⁸ where sampling biases have been observed as well. Our study indeed demonstrates that modifying the driver may be useful. In contrast to the above studies, we have studied non-stoquastic yet local (i.e., two-body) driver Hamiltonians, that as such are physically more reasonable. While we find that non-stoquastic local drivers do indeed produce different GSDs than transverse-field drivers, the manner in which they sample the ground state manifolds is not found to be particularly more uniform than with the standard driver.

While we have not explicitly discussed in this work the *speed* in which these distributions are obtained, we have demonstrated, for the first time to our knowledge, that annealers driven by quantum fluctuations may be used to assist existing traditional algorithms in finding all, or as many as possible, minimizing configurations of optimization problems. That quantum annealers may obtain 'classically rare' minimizing configurations has numerous immediate applications in various fields, such as k-SAT filtering, detecting hardware faults and verification and validation (V&V).

An interesting question that arises at this stage and which we believe warrants further investigation is how one should choose the strength or nature of the quantum fluctuations to boost the probabilities of classically suppressed configurations. Presumably, algorithms that adaptively control the nature of the driver terms based on repeated applications of the annealing process would be an interesting path to explore (see, e.g., ref. 25). Another question that remains open has to do with the nature of the GSDs generated by the DW2 experimental quantum annealer. As we have shown, for most instances, the experimentally generated distributions are only weakly correlated with those of the classical thermal ones, nor have these been found to correlate with the distributions

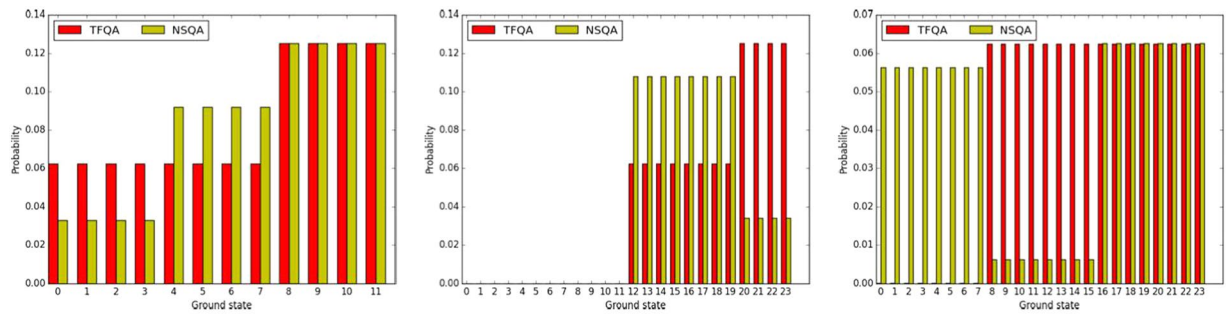


Figure 7. Three sample GSD comparisons between the two drivers: the stochastic TFQA (red) and non-stochastic NSQA (yellow). For the instance on the left both drivers sample the ground-state manifold similarly; for the middle instance less probable configurations for one method become more probable in the other and vice versa; on the right, we find an instance for which states with zero probability of occurring with one type of quantum fluctuations have distinct positive probabilities of occurring in the other.

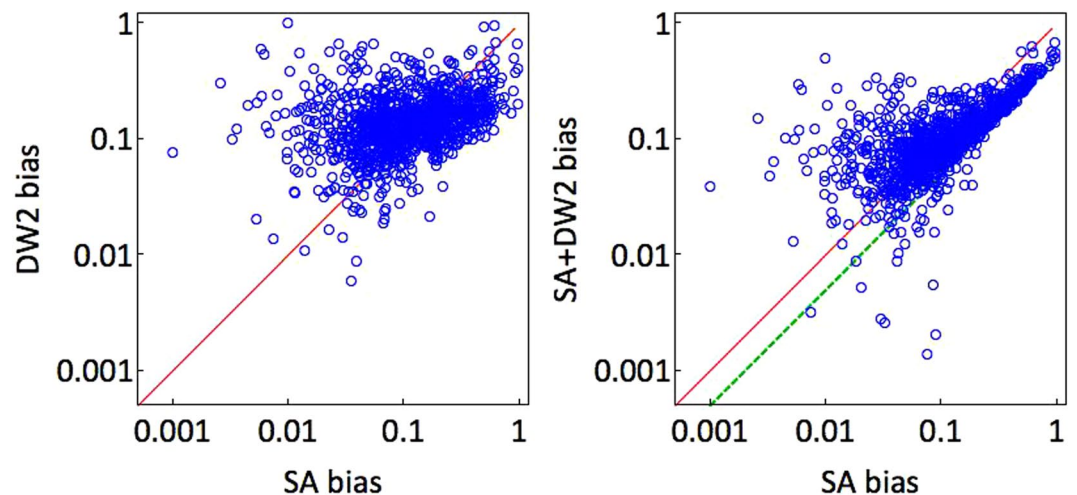


Figure 8. Scatter plots of SA biases vs. DW2 biases and a combination of SA and DW2. Left: The scatter plot indicates that the biases of the DW2 processor are similar in nature to those generated by thermal annealing. Right: Merit of using SA in conjunction with an experimental quantum annealer. Here, results are rather mixed. While for a large portion (over 50%) of the instances a combination of the methods leads to a smaller overall bias, i.e., a lesser degree of unfairness, for the rest of the instances the bias is larger. The solid red line is the ‘equal bias’ $y = x$ line, whereas the dashed green $y = x/2$ line is the bias obtained via a combination of the horizontal GSD with a flat unbiased GSD.

obtained by the zero-temperature ideal transverse-field quantum annealers. The dominant mechanisms that determine the distributions emerging from the D-Wave devices remain unknown, however. Specifically, the question of whether the yielded distributions are ‘quantum’ in nature—i.e., dominated by the quantum fluctuations, or classical, i.e., set by thermal fluctuations, or simply by intrinsic control errors (ICE)—is still left unanswered. One plausible explanation is that the generated distributions are a result of a combination of thermal and quantum fluctuations, given the finite-temperature of the device. A recent conjecture²⁸ suggests that these experimental devices in fact sample from the quantum ground-state at a point midway through the anneal at a so-called ‘freeze-out’ point (or region) where thermal fluctuations become too weak to generate any changes to the state. Another plausible scenario is that the generated distributions are determined by intrinsic control errors—those errors that stem from the analog nature of the processor and as such may have a decisive effect on the resulting ground-state distributions²⁹. If the latter conjecture is found to be true, then it could be that these differences in GSDs are easily reproducible by classical means by simply injecting random errors to the problem parameters.

Finally, this study raises an interesting question concerning the computational complexity of faithfully sampling the quantum ground-states of spin glass Hamiltonians. While for (ideal) quantum annealers the successful sampling of the quantum ground-state consists solely of performing an adiabatic anneal followed by a computational basis measurement, classical algorithms must follow a different path to accomplish the same task. Since non-stochastic Hamiltonians cannot be efficiently simulated, the prevailing algorithm to date for sampling the quantum ground-state of such Hamiltonians is one which uses degenerate perturbation theory. The latter technique consists of first diagonalizing the problem Hamiltonian to find all its (classical) ground-states, followed by an application of degenerate perturbation theory (up to N -th order in the worst case). This last step consists

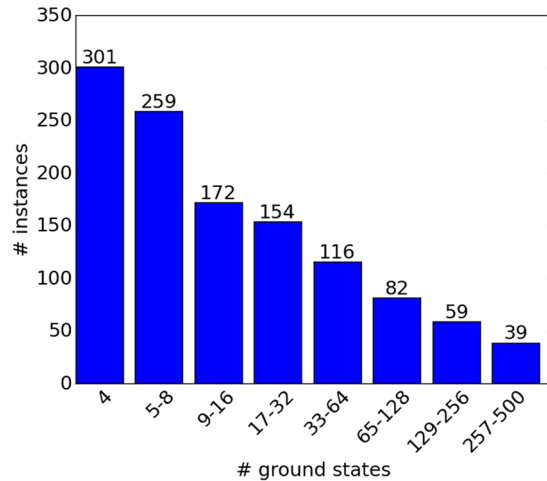


Figure 9. Histogram of number of problem instances vs. number of ground-states. The distribution of number of solutions is a function of clause density after discarding instances with more than 500 solutions.

of tracking, in the worst case, all N -spin paths from any one classical ground-state to any other. At least naively, the computational complexity involved in doing so scales as $N!$. Thus, the problem of sampling from quantum ground-states generated by non-stoquastic driver Hamiltonians may eventually prove to be a path to quantum enhancements of a novel kind.

Methods

To test the capabilities of quantum annealers tasked with the identification of the individual ground-state configurations of a given problem versus classical algorithms, we first generate random spin glass instances for which the enumeration of the minimizing configurations is feasible. This a priori requirement allows for the proper evaluation of the sampling capabilities of the tested algorithms in an unbiased and consistent manner.

Generation of Instances. For the generation of instances, in this work we have chosen to study problems constructed around ‘planted solutions’—an idea borrowed from constraint satisfaction (SAT) problems^{30,31}. In these problems, the planted solution represents a ground-state configuration of Eq. (3) that minimizes the energy and is known in advance. The Hamiltonian of a planted-solution spin glass is a sum of terms, each of which consists of a small number of connected spins, namely, $H = \sum_j H_j$ ²². Each term H_j is chosen such that one of its ground-states is the planted solution. It follows then that the planted solution is also a ground-state of the total Hamiltonian, and its energy is the ground-state energy of the Hamiltonian. Knowing the ground-state energy in advance circumvents the need to verify the ground-state energy using exact (provable) solvers, which rapidly become too expensive computationally as the number of variables grows. The interested reader will find a more detailed discussion of planted Ising problems in ref. 22. As we show next, studying this type of problem also allows us to devise a practical algorithm capable of finding all minimizing configurations of the generated instances.

Enumeration of Minimizing Configurations. The enumeration of all minimizing configurations of a given optimization problem is a difficult task in the general case, belonging to the complexity class $\#P$. The exhaustive search for all solutions of an N -spin Ising problem becomes unfeasible for $N > 40$ bit problems as the search space grows exponentially with the size of the problem. To successfully address this difficulty, we use the fact that our generated instances consist of a sum of terms—each of which has all minimizing configurations as its ground-state (these are commonly referred to as frustration free Hamiltonians). To enumerate all minimizing configurations, we implement a form of the ‘bucket’ algorithm³² designed to eliminate variables one at a time to perform an exhaustive search efficiently (for a detailed description of the algorithm, see the Supplementary Information). Implementing the bucket algorithm and running it on the randomly generated planted-solution instances, discarding instances with more than 500 solutions, has yielded the histogram depicted in Fig. 9 which provides the number of problem instances used in this study versus number of ground-states.

Calculation of the Quantum Ground-state. Here we review the algorithm used to compute the ‘quantum ground-state’ of H_p , namely, the s -dependent ground-state in the limit $s \rightarrow 1$ (from below) of the s -dependent Hamiltonian,

$$H(s) = (1 - s)H_d + sH_p. \quad (7)$$

We consider an ideal zero-temperature adiabatic evolution that starts with the driver H_d (e.g., a transverse-field Hamiltonian) at $s = 0$, and ends with the problem Hamiltonian H_p at $s = 1$, where the problem Hamiltonian encodes a classical cost function with D degenerate ground-states.

Calculating the quantum ground-state in the presence of a degenerate ground subspace of an $N \approx 500$ -qubit problem Hamiltonian is a nontrivial task. In general, it requires the diagonalization of $H(s)$ which scales exponentially with the number of spins. Our workaround combines the Rayleigh-Ritz variational principle, taking

advantage of the fact the sought state is that of minimal energy, combined with degenerate perturbation theory (see e.g., ref. 33), in which H_d is considered a perturbation to H_p .

We begin by observing that perturbation theory separates the quantum ground-state from the other states in the limit $s \rightarrow 1$ (as shown for example in Fig. 1) using a growing sequence of sets of states S_k and subspaces V_k such that:

1. The set $S_1 = \{|\phi_1\rangle, |\phi_2\rangle, \dots, |\phi_D\rangle\}$ contains all the classical ground-states.
2. Subspace V_1 is the linear span of S_1 (i.e., the set of all linear combinations of vectors in S_1).
3. The states $|\phi\rangle \in S_k$ meet three requirements, (a) $|\phi\rangle$ is orthogonal to V_{k-1} , (b) $|\phi\rangle$ belongs to the computational basis and hence is an eigenstate of H_p , and (c) $|\phi\rangle$ has a non-vanishing matrix element $\langle \phi | H_d^{k-1} | \phi_i \rangle$ for some classical ground-state $|\phi_i\rangle \in S_1$.
4. V_k is the linear span of the union $S_1 \cup S_2 \cup \dots \cup S_k$.

For instance, for the transverse-field driver Eq. (4), the set S_2 is composed of the states obtained from each classical ground-state by flipping only one bit (these new states must not be ground-states themselves).

Additionally, we take advantage of the symmetry of our Hamiltonians $H(s)$ under global bit flip [see Eqs (3), (4) and (5)]. For any vector in the computational basis $|\phi\rangle$, we denote the state obtained by flipping all bits in $|\phi\rangle$ as $|\bar{\phi}\rangle$, and the symmetric (+) and antisymmetric (−) components of $|\phi\rangle$ as $(|\phi\rangle \pm |\bar{\phi}\rangle)/\sqrt{2}$. Similarly, since H_d is also symmetric under global bit-flip, the subspaces V_k split into symmetric and antisymmetric subspaces $V_k = V_k^S \oplus V_k^A$. We shall therefore only consider V_k^S from now on.

Perturbation theory prescribes in a well-defined manner the specific linear combination of states in V_k^S that compose the ground-state to k -th order. However, this prescription is contrived. To simplify the procedure, we turn to the Rayleigh-Ritz variational principle. Restricting our test functions to the subspace V_k^S , the ground-state can also be defined as

$$|\psi_{GS}(s)\rangle = \arg \min_{\psi \in V_k^S} \frac{\langle \psi | H(s) | \psi \rangle}{\langle \psi | \psi \rangle}, \quad (8)$$

i.e., the state with least energy. Our variational estimate is consistent with the k -th order perturbation theory because V_k^S contains all the states relevant to k -th order perturbation. The minimizing state $|\psi_{GS}(s)\rangle$ can be easily found using a conjugate gradient method (for real Hamiltonians, such as ours, the search for the minimum in Equation (8) can be restricted to real states $|\Psi\rangle$; our conjugate gradient follows Appendix A of ref. 34).

Our algorithm proceeds as follows. Considering initially only the order $k=1$ in Eq. (8), we set $s = s_*$, a small number such that the ground-state of $H(s_*)$ is close to the ground-state of H_d and for which H_p serves as a perturbation (we set $s_* = 0.1$) and minimize $H(s_*)$. We then verify that the minimizing state is not degenerate. To do that, it is sufficient to run the conjugate gradient minimization routine twice, starting from two independent random states and checking that the two obtained ground-states are linearly dependent. Finding the same ground-state twice in the presence of degeneracies is a zero-measure event as the minimization is that of a quadratic function, for which there are no local maxima separating different minima. If the minimizing state is found to be unique, we proceed to slowly increase s up to $s \rightarrow 1$ minimizing along the way. If on the other hand the degeneracies are not broken at $s = s_*$, we expand the variational search to V_2^S and then proceed with the numerical annealing all the way to $s \rightarrow 1$. The state obtained at the end of this procedure is the quantum ground state of H_p as driven by H_d .

In the case of H_d^{TF} we found that out of about 1900 randomly generated spin glass instances, considering V_1^S sufficed to find the quantum ground-state for 912 of them, and the inclusion of V_2^S sufficed for the rest. In the non-stoquastic H_d^{NS} case, 1508 were solved with V_1^S , and the rest with V_2^S . We note that for those instances in which all degeneracies are removed at the level of V_1^S , the numerical annealing can be dispensed of altogether. This is because the problem Hamiltonian H_p , when restricted to V_1^S , is merely the classical ground-state energy times the identity.

Comparing Empirical Distribution Functions. Here we discuss the statistical comparisons of pairs of GSDs over the ground-state manifold of a spin glass with D solutions, where at least one of the distributions is empirically estimated via experiment or simulation. Let the two probability distributions be denoted $\mathbf{p}^{(1)}$, $\mathbf{p}^{(2)}$ where $p_i^{(1)}$ and $p_i^{(2)}$ are the probabilities of obtaining ground-state i with the first and second method, respectively. Let N_1 and N_2 be the sample sizes (i.e., number of anneals) for these two experiments (if one of the methods generates an analytic probability distribution, this corresponds to taking its number of anneals to infinity). Hence, if ground-state i was found $n_i^{(1)}$ times with method one, then $p_i^{(1)} = n_i^{(1)}/N_1$ (with similar definitions for method two). It follows that the square of the χ^2 -distance between distributions $\mathbf{p}^{(1)}$ and $\mathbf{p}^{(2)}$ is ref. 35

$$\|\mathbf{p}^{(1)} - \mathbf{p}^{(2)}\|_{\chi^2}^2 = \sum_{i=1}^D \frac{\left(\sqrt{\frac{N_1}{N_2}} n_i^{(2)} - \sqrt{\frac{N_2}{N_1}} n_i^{(1)} \right)^2}{n_i^{(1)} + n_i^{(2)}}. \quad (9)$$

If one of the two probabilities, say $\mathbf{p}^{(2)}$, is known analytically, we may just take the limit $N_2 \rightarrow \infty$ in Eq. (9), obtaining

$$\|\mathbf{p}^{(1)} - \mathbf{p}^{(2)}\|_{\chi^2, \text{one sided}}^2 = \sum_{i=1}^D \frac{(N_i p_i^{(2)} - n_i^{(1)})^2}{N_i p_i^{(2)}}. \quad (10)$$

The Kolmogorov-Smirnov test proceeds as follows. Here, for the null hypothesis, one assumes that both $\mathbf{p}^{(1)}$ and $\mathbf{p}^{(2)}$ are drawn from the same underlying probability distribution. If the null hypothesis holds, then for large N_1 , N_2 and D , the probability distribution for $\|\mathbf{p}^{(1)} - \mathbf{p}^{(2)}\|_{\chi^2}$ can be computed assuming Gaussian statistics. Hence, one would just compare the $\|\mathbf{p}^{(1)} - \mathbf{p}^{(2)}\|_{\chi^2}$ computed from empirical data with the known χ^2 -distribution. This comparison yields the probability (or p -value) that the null hypothesis holds³⁵.

In our case, neither N_1 , N_2 nor D are exceedingly large. Therefore, rather than relying on a precomputed χ^2 -probability, we resort to using a bootstrapped version of the Kolmogorov-Smirnov (KS) test: assuming for the null hypothesis that the probability distributions associated with the two methods are in fact identical, the underlying distribution p is given by: $p_i = (N_1 p_i^{(1)} + N_2 p_i^{(2)}) / (N_1 + N_2)$ (this p_i is our best guess for the probability of getting the i -th ground-state given the null hypothesis). As a next step, we simulate synthetic experiments based on this underlying distribution. Each synthetic experiment consists of extracting N_1 ground-states to form a synthetic probability distribution \mathbf{P}_1 , and N_2 ground-states to form a probability distribution \mathbf{P}_2 (if a method is analytic then we set $\mathbf{P}_i = \mathbf{p}_i$ by law of large numbers). We simulate a large number of such synthetic experiments, and measure the proportion of experiments for which we find $\|\mathbf{P}_1 - \mathbf{P}_2\| > \|\mathbf{p}^{(1)} - \mathbf{p}^{(2)}\|$. This proportion corresponds to the p -value for our test. As a check, we have compared our p -values with the ones obtained using the tabulated χ^2 probability verifying that the p -values of the two tests agree to within a few percent.

References

- Douglass, A., King, A. D. & Raymond, J. *Theory and Applications of Satisfiability Testing – SAT 2015: 18th International Conference, Austin, TX, USA, September 24–27, 2015, Proceedings*, chap. Constructing SAT Filters with a Quantum Annealer, 104–120 (Springer International Publishing, Cham, 2015).
- Pudenz, K. L. & Lidar, D. A. Quantum adiabatic machine learning. *Quantum Information Processing* **12**, 2027–2070, doi:10.1038/ncomms4243 (2013).
- Papadimitriou, C. & Steiglitz, K. *Combinatorial Optimization: Algorithms and Complexity*. Dover Books on Computer Science (Dover Publications, 2013).
- Pauling, L. The structure and entropy of ice and of other crystals with some randomness of atomic arrangement. *Journal of the American Chemical Society* **57**, 2680–2684, doi:10.1021/ja01315a102 (1935).
- Bastea, S. & Duxbury, P. M. Ground state structure of random magnets. *Phys. Rev. E* **58**, 4261–4265, doi:10.1103/PhysRevE.58.4261 (1998).
- Johnson, M. W. *et al.* Quantum annealing with manufactured spins. *Nature* **473**, 194–198, doi:10.1038/nature10012 (2011).
- Berkley, A. J. *et al.* Tunneling spectroscopy using a probe qubit. *Phys. Rev. B* **87**, 020502(R), doi:10.1103/PhysRevB.87.020502 (2013).
- Kadowaki, T. & Nishimori, H. Quantum annealing in the transverse Ising model. *Phys. Rev. E* **58**, 5355–5363, doi:10.1103/PhysRevE.58.5355 (1998).
- Farhi, E. *et al.* A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science* **292**, 472–475, doi:10.1126/science.1057726 (2001).
- Finnila, A., Gomez, M., Sebenik, C., Stenson, C. & Doll, J. Quantum annealing: A new method for minimizing multidimensional functions. *Chemical Physics Letters* **219**, 343–348, doi:10.1016/0009-2614(94)00117-0 (1994).
- Brooke, J., Bitko, D., Rosenbaum, F. T. & Aeppli, G. Quantum annealing of a disordered magnet. *Science* **284**, 779–781, doi:10.1126/science.284.5415.779 (1999).
- Santoro, G., Martoňák, R., Tosatti, E. & Car, R. Theory of quantum annealing of an Ising spin glass. *Science* **295**, 2427–2430, doi:10.1126/science.1068774 (2002).
- Das, A. & Chakrabarti, B. K. Colloquium. *Rev. Mod. Phys.* **80**, 1061–1081, doi:10.1103/RevModPhys.80.1061 (2008).
- Ray, P., Chakrabarti, B. K. & Chakrabarti, A. Sherrington-kirkpatrick model in a transverse field: Absence of replica symmetry breaking due to quantum fluctuations. *Phys. Rev. B* **39**, 11828–11832, doi:10.1103/PhysRevB.39.11828 (1989).
- Boixo, S., Albash, T., Spedalieri, F. M., Chancellor, N. & Lidar, D. A. Experimental signature of programmable quantum annealing. *Nat Commun* **4**, doi:10.1038/ncomms3067 (2013).
- Vinci, W., Albash, T., Mishra, A., Warburton, P. A. & Lidar, D. A. Distinguishing classical and quantum models for the D-Wave device. *arXiv:1403.4228* (2014).
- Matsuda, Y., Nishimori, H. & Katzgraber, H. G. Quantum annealing for problems with ground-state degeneracy. *Journal of Physics: Conference Series* **143**, 012003, doi:10.1088/1742-6596/143/1/012003 (2009).
- Mandrà, S., Zhu, Z. & Katzgraber, H. G. Exponentially biased ground-state sampling of quantum annealing machines with transverse-field driving hamiltonians. *Phys. Rev. Lett.* **118**, 070502, doi:10.1103/PhysRevLett.118.070502 (2017).
- Young, A. P. (ed.) *Spin Glasses and Random Fields* (World Scientific, Singapore, 1998).
- Barahona, F. On the computational complexity of Ising spin glass models. *J. Phys. A* **15**, 3241–3253, doi:10.1088/0305-4470/15/10/028 (1982).
- Bunyk, P. I. *et al.* Architectural considerations in the design of a superconducting quantum annealing processor. *Applied Superconductivity, IEEE Transactions on* **24**, 1–10, doi:10.1109/TASC.2014.2318294 (Aug. 2014).
- Hen, I. *et al.* Probing for quantum speedup in spin-glass problems with planted solutions. *Phys. Rev. A* **92**, 042325, doi:10.1103/PhysRevA.92.042325 (2015).
- King, J., Yarkoni, S., Nevisi, M. M., Hilton, J. P. & McGeoch, C. C. Benchmarking a quantum annealing processor with the time-to-target metric. *ArXiv e-prints 1508.05087* (2015).
- Kirkpatrick, S., Gelatt, C. D. Jr. & Vecchi, M. P. Optimization by simulated annealing. *Science* **220**, 671–680, doi:10.1126/science.220.4598.671 (1983).
- Dickson, N. G. & Amin, M. H. Algorithmic approach to adiabatic quantum optimization. *Phys. Rev. A* **85**, 032303, doi:10.1103/PhysRevA.85.032303 (2012).
- Crosson, E., Farhi, E., Yen-Yu Lin, C., Lin, H.-H. & Shor, P. Different Strategies for Optimization Using the Quantum Adiabatic Algorithm. *ArXiv e-prints 1401.7320* (2014).
- Boixo, S. *et al.* Evidence for quantum annealing with more than one hundred qubits. *Nat. Phys.* **10**, 218–224, doi:10.1038/nphys2900 (2014).
- Amin, M. H., Andriyash, E., Rolfe, J., Kulchitsky, B. & Melko, R. Quantum Boltzmann Machine. *ArXiv e-prints 1601.02036* (2016).
- Martin-Mayor, V. & Hen, I. Unraveling quantum annealers using classical hardness. *Scientific Reports* **5**, 15324 EP– (2015).

30. Barthelemy, W. *et al.* Hiding solutions in random satisfiability problems: A statistical mechanics approach. *Physical Review Letters* **88**, 188701, doi:[10.1103/PhysRevLett.88.188701](https://doi.org/10.1103/PhysRevLett.88.188701) (2002).
31. Krzakala, F. & Zdeborová, L. Hiding quiet solutions in random constraint satisfaction problems. *Physical Review Letters* **102**, 238701, doi:[10.1103/PhysRevLett.102.238701](https://doi.org/10.1103/PhysRevLett.102.238701) (2009).
32. Dechter, R. Bucket elimination: a unifying framework for processing hard and soft constraints. *Constraints* **2**, 51–55, doi:[10.1023/A:1009796922698](https://doi.org/10.1023/A:1009796922698) (1997).
33. Galindo, A. & Pascual, P. *Quantum Mechanics II* (Springer-Verlag, Berlin Heidelberg, 1991).
34. Hen, I. Excitation gap from optimized correlation functions in quantum Monte Carlo simulations. *Phys. Rev. E* **85**, 036705, doi:[10.1103/PhysRevE.85.036705](https://doi.org/10.1103/PhysRevE.85.036705) (2012).
35. Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. *Numerical Recipes in C* (Cambridge University Press, Cambridge, 1995).

Acknowledgements

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. Computation for the work described here was also supported by the University of Southern California's Center for High-Performance Computing (<http://hpcc.usc.edu>).

Author Contributions

I.H. and V.M.M. conceived the ideas behind this work. V.M.M., B.H.Z. and G.W. carried out the numerical simulations and statistical tests. I.H. performed the tests on the experimental quantum annealer. All authors reviewed the manuscript.

Additional Information

Supplementary information accompanies this paper at doi:[10.1038/s41598-017-01096-6](https://doi.org/10.1038/s41598-017-01096-6)

Competing Interests: The authors declare that they have no competing interests.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2017