

ARTICLE OPEN

There are amendments to this paper

Analog errors in quantum annealing: doom and hope

Adam Pearson^{1,2}, Anurag Mishra^{1,2,6}, Itay Hen^{1,2,3} and Daniel A. Lidar^{1,2,4,5*}

Quantum annealing has the potential to provide a speedup over classical algorithms in solving optimization problems. Just as for any other quantum device, suppressing Hamiltonian control errors will be necessary before quantum annealers can achieve speedups. Such analog control errors are known to lead to J -chaos, wherein the probability of obtaining the optimal solution, encoded as the ground state of the intended Hamiltonian, varies widely depending on the control error. Here, we show that J -chaos causes a catastrophic failure of quantum annealing, in that the scaling of the time-to-solution metric becomes worse than that of a deterministic (exhaustive) classical solver. We demonstrate this empirically using random Ising spin glass problems run on the two latest generations of the D-Wave quantum annealers. We then proceed to show that this doomsday scenario can be mitigated using a simple error suppression and correction scheme known as quantum annealing correction (QAC). By using QAC, the time-to-solution scaling of the same D-Wave devices is improved to below that of the classical upper bound, thus restoring hope in the speedup prospects of quantum annealing.

npj Quantum Information (2019)5:107; <https://doi.org/10.1038/s41534-019-0210-7>

INTRODUCTION

The demonstration of scaling speedups^{1,2} using quantum hardware is the holy grail of quantum computing, and massive efforts are underway worldwide in their pursuit. A daunting obstacle is the fact that all physical implementations of quantum computers suffer from analog control errors, in which the coefficients of the Hamiltonian implemented differ from those intended (e.g., ref.³) a fact that threatens to spoil the results of computations due to the accumulation of small errors. This problem was recognized early on in the gate model of quantum computing,⁴ and soon after theoretically dealt with by error discretization via quantum error correcting codes.⁵ Moreover, the accuracy threshold theorem guarantees that if the physical gates used to implement encoded, error-corrected quantum circuits have sufficiently high fidelity, then any real, noisy quantum computation can be made arbitrarily close to the intended, noiseless computation with modest overhead.^{6–8} In contrast, the ultimate impact of analog control errors in Hamiltonian quantum computing, in particular in adiabatic quantum optimization⁹ and quantum annealing (QA),¹⁰ is not as clear. While unlike the gate-model adiabatic quantum evolution is inherently robust to path variations due to unitary control errors,¹¹ there is as of yet no equivalent mechanism of error discretization or an analogous accuracy threshold theorem in this paradigm. Yet, at the same time QA offers a currently unparalleled opportunity to explore NISQ-era (noisy intermediate-scale quantum-era)¹² quantum optimization with thousands of qubits.^{13,14} It is thus of great importance to assess the role of analog control errors in QA, and to find ways to mitigate them. Here we do so in the context of spin glass problems, which are known to exhibit a type of control error-induced bond or disorder chaos already in a purely classical setting, causing chaotic changes in the ground or equilibrium state.^{15,16}

The extent to which analog control errors present a challenge in using any physical realization of QA, such as the D-Wave processors,^{17–19} for optimization, cannot be overstated. Indeed,

an earlier study of such errors in these processors found evidence of sub-classical performance and referred to the effect as J -chaos,²⁰ a terminology we adopt here. More recently, it was shown that analog control noise causes a decrease in the probability that the implemented Hamiltonian shares a ground state with the intended Hamiltonian that scales exponentially in the size of the problem and the magnitude of the noise.²¹ This means that even if the annealer solves the implemented problem correctly, it has an exponentially shrinking probability of finding the intended ground state. In other words, subject to J -chaos an otherwise perfectly functioning quantum annealer will typically find the correct answer to the wrong problem.

To mitigate this “wrong Hamiltonian” problem and restore the prospects for a speedup in the use of QA for optimization, it is necessary to introduce techniques for error suppression and correction. This observation is not new,^{22–24} and repetition coding along with the use of energy penalties has been shown to significantly enhance the performance of quantum annealers.^{25–30} In contrast to previous work, here we directly address the impact of J -chaos on algorithmic scaling of optimization in experimental QA, while accessing a computational scale that is still far out of reach of current gate-model quantum computing devices. We employ a QA correction method to mitigate the problem, and demonstrate that while the scaling of uncorrected QA in solving random Ising spin glass problems is catastrophically affected by J -chaos—in that it is worse than even that of a deterministic (brute force) classical solver—hope for a quantum speedup³¹ is restored with error suppression and correction. This reassuring conclusion is reached here using the simplest possible error suppression and correction scheme,²⁵ so that much room for improvement remains for more advanced methods. We expect our results to apply broadly, certainly beyond the D-Wave devices to other quantum^{32–34} and semiclassical annealing implementations,^{35,36} and to other forms of analog quantum computing.³⁷

¹Department of Physics and Astronomy, University of Southern California, Los Angeles, CA 90089, USA. ²Center for Quantum Information Science and Technology, University of Southern California, Los Angeles, CA 90089, USA. ³Information Sciences Institute, University of Southern California, Los Angeles, CA 90089, USA. ⁴Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089, USA. ⁵Department of Chemistry, University of Southern California, Los Angeles, CA 90089, USA. ⁶Present address: Qulab Inc., 1642 Westwood Boulevard, Los Angeles, CA 90024, USA. *email: lidar@usc.edu

RESULTS

Inspired by classical simulated annealing, in which thermal fluctuations are used to hop over barriers, quantum annealing uses quantum fluctuations to tunnel through barriers.^{10,38–43} The D-Wave processors are physical implementations of such devices.⁴⁴ In the standard forward annealing protocol they apply a time-dependent transverse field $H_X = \sum_i \sigma_i^x$ (σ_i^a denotes the Pauli matrix of type $a \in \{x, y, z\}$ applied to qubit $i \in \{1, \dots, N\}$) and Ising Hamiltonian as follows:

$$H(s) = A(s)H_X + B(s)\tilde{H}_{\text{Ising}}, \quad (1)$$

where $\tilde{H}_{\text{Ising}} = H_{\text{Ising}} + \delta H_{\text{Ising}}$ and H_{Ising} are, respectively, the implemented (perturbed) and intended (unperturbed) “problem” Hamiltonians, while δH_{Ising} is an error term (the perturbation). Thus, \tilde{H}_{Ising} is the (wrong) Hamiltonian including analog control errors, while H_{Ising} is the Hamiltonian whose ground state we wish to find as the solution to the optimization problem specified by a set of local fields $\{h_i\}$ and couplers $\{J_{ij}\}$:

$$H_{\text{Ising}} = \sum_{i \in \mathcal{V}} h_i \sigma_i^z + \sum_{(i,j) \in \mathcal{E}} J_{ij} \sigma_i^z \sigma_j^z, \quad (2a)$$

$$\delta H_{\text{Ising}} = \sum_{i \in \mathcal{V}} \delta h_i \sigma_i^z + \sum_{(i,j) \in \mathcal{E}} \delta J_{ij} \sigma_i^z \sigma_j^z, \quad (2b)$$

where \mathcal{V} and \mathcal{E} are the vertex and edge sets of the graph \mathcal{G} , and $N = |\mathcal{G}|$. We assume that the noise is Gaussian with zero mean and standard deviation η :

$$\delta h_i, \delta J_{ij} \sim \mathcal{N}(0, \eta^2). \quad (3)$$

We note that analog errors resulting in the replacement of H_X by $\sum_i (1 + \epsilon_i) \sigma_i^x$ with random ϵ_i are expected as well, but we do not consider such transverse field errors here, since such errors are not directly programmable on the current generation of D-Wave annealers (unlike $\delta h_i, \delta J_{ij}$), and also because such errors do not contribute to J -chaos. The normalized time, $s = t/t_f$, with t_f denoting the final time, increases from 0 to 1, with $A(0) \gg B(0)$ and $B(1) \gg A(1)$, and $A(s)$ [$B(s)$] decrease (increase) monotonically. As such, the transverse field initially drives strong quantum fluctuations that eventually give way to the implemented Ising Hamiltonian. In the absence of δH_{Ising} and an environment, the adiabatic theorem guarantees that the ground state of H_{Ising} will be found if t_f is large compared to the inverse of the minimum gap of $H(s)$ and the maximum time derivative(s) of $H(s)$.^{45,46} In the presence of an environment, the adiabatic theorem instead guarantees evolution towards the steady state of the corresponding Liouvillian,^{47,48} which becomes the ground state only for sufficiently low temperature (compared to the gap of the Liouvillian). When $\delta H_{\text{Ising}} \neq 0$, the probability that the computation ends in the ground state of H_{Ising} decreases exponentially in both N and some power of η .²¹ The reason is that as more noise is added, there is an increasing probability that the spectrum changes such that the ground state is swapped with an excited state. Thus, increasing noise and problem size leads to a rapidly growing probability of failure to find the correct ground state. In experimental quantum annealing both the environment and control errors inevitably play a role.

Effect of control noise

In order to systematically test the effect of analog control errors, we studied the performance of two D-Wave devices on random Ising instances of varying size N , to which we added artificially generated Gaussian control noise η . This noise was added to the intrinsic analog device noise η_{int} ,⁴⁹ so that the total control noise had variance $\eta_{\text{int}}^2 + \eta^2$. Adding noise in this manner allowed us to test its effect on algorithmic performance, and the efficacy of the quantum annealing correction (QAC) strategy described below.

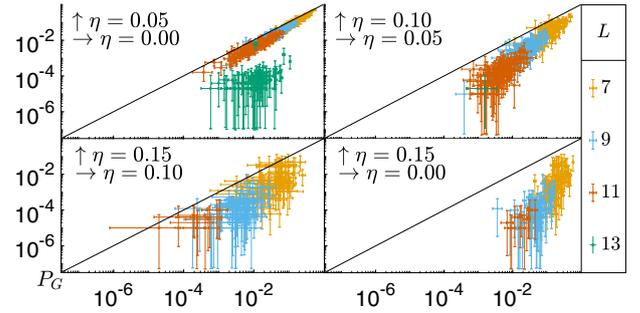


Fig. 1 Correlation plots of ground state (success) probability P_g with increasing added noise η . Results for random Ising instances are shown after addition of coupler noise $\delta J_{ij} \sim \mathcal{N}(0, \eta^2)$. In each panel every data point is P_g for the same instance with different η values. Top left: success probability decreases as we go from the no-added noise case $\eta = 0$ to $\eta = 0.05$. Note that even for $\eta = 0$ there is an intrinsic control error η_{int} . The same trend persists as we compare $\eta = 0.05$ to $\eta = 0.10$ (top right) and $\eta = 0.10$ to $\eta = 0.15$ (bottom left). The bottom right panel shows the extreme comparison of $\eta = 0$ to $\eta = 0.15$: the ground state is almost never found for $\eta = 0.15$. In all plots success probability generally decreases with increasing size L . Instances with zero success probability are not shown, hence the number of data points drops with increasing η . Unless otherwise noted, here and in later plots error bars always denote 95% confidence intervals (CI) obtained via a bootstrap. Details of our data analysis are given in Methods.

Throughout this work we used Ising instances with local fields $h_i = 0$ and couplers J_{ij} selected uniformly at random from the set $\pm\{1/6, 1/3, 1/2\}$, and chose $\eta \in \{0, 0.03, 0.05, 0.07, 0.10, 0.15\}$. Such instances have been studied before,^{26,31,50–52} but only with $\eta = 0$, that is, never subject to the systematic addition of control noise. We define “success” in all cases as finding a ground state of the unperturbed Hamiltonian H_{Ising} . See Methods for a complete description of the instances and how we verified ground states.

The number of qubits N is proportional to L^2 , the number of Chimera graph unit cells of the D-Wave devices we used (see Methods). Figure 1 displays a series of correlation plots between different levels of added noise, for different problem sizes parametrized by L . At every size, addition of noise results in a lower success probability for all instances. Increased size also results in lower success probability, as expected. Thus, Fig. 1 gives a visual confirmation of the detrimental effect of control noise; we quantify this systematically below.

Next, we test whether control noise results in J -chaos. The latter exhibits itself as large variations in the success probability of the programmed Hamiltonian across different runs. We quantify this in terms of the J -chaoticity measure σ/μ , where σ is the standard deviation of the success probability across repeated runs of a given instance, and μ is the corresponding mean (see Methods for more details). In Fig. 2, we plot the correlation of the J -chaoticity measure with increasing noise. For most instances, this quantity becomes larger with increasing noise and size, which indicates that they are becoming more chaotic. Success probability is also strongly (negatively) correlated with increasing J -chaoticity, as shown in Fig. 3. This establishes that control noise-induced J -chaos is responsible for a strong decline in performance. Before we quantify this decline in terms of the time-to-solution metric, we first address how to mitigate this problem using error suppression and correction.

Quantum annealing correction

The error suppression and correction scheme used in this work is the $[3, 1, 3]_1$ QAC code introduced in ref.²⁵ and further studied experimentally in ref.^{26–28} We refer the reader to these references for details, and to Methods for a brief summary. The $[3, 1, 3]_1$ code

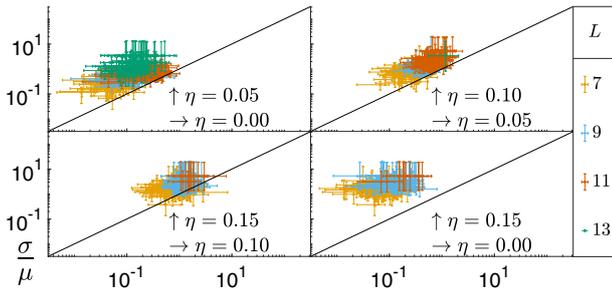


Fig. 2 Correlation plots of J -chaoticity measure σ/μ with increasing added noise η . Details are as in Fig. 1, but here we compare the J -chaoticity measure of instances with different added noise values. The measure increases as we go from the no-added noise case $\eta = 0$ to $\eta = 0.05$ (top left). The same trend persists as we compare $\eta = 0.05$ to $\eta = 0.10$ (top right) and $\eta = 0.10$ to $\eta = 0.15$ (bottom left). The bottom right panel shows the extreme comparison of $\eta = 0$ to $\eta = 0.15$, where for all instances at all sizes J -chaoticity is higher to well within the 95% CIs.

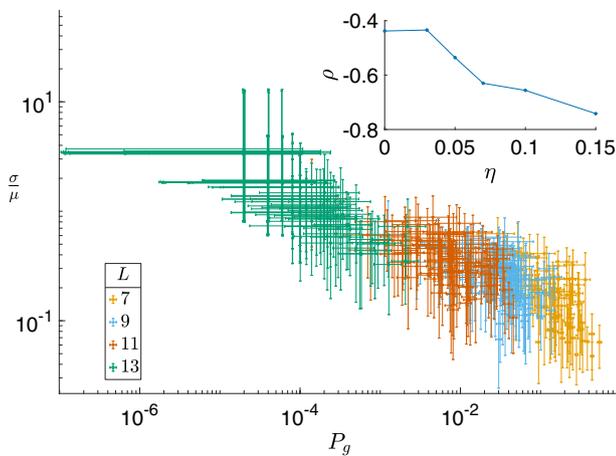


Fig. 3 Correlation between success probability and J -chaoticity. We plot the success probability and J -chaoticity measure with noise $\eta = 0.03$. There is evidently a strong correlation between the two, showing that the success probability is lowered due to J -chaos. The inset shows the Pearson's correlation coefficient ρ between success probability and J -chaoticity, as a function of added noise, for all η values tested. As η grows, success probability and J -chaoticity become more negatively correlated.

is a three-qubit repetition code that corrects bit-flip errors, with the subscript denoting one extra penalty qubit. The penalty term energetically suppresses all errors that do not commute with σ^z during the anneal.

Since the $[3, 1, 3]_1$ graph is a minor of the Chimera graph (see Methods Fig. 9), we can also implement these instances without QAC. However, to ensure a fair comparison we need to equalize the resources used with QAC and without it. The $[3, 1, 3]_1$ consumes four qubits to encode one logical qubit. Thus, we can use the same amount of resources as the encoded logical problem by running four unencoded copies in parallel, which is called the classical repetition strategy (C). To be clear, the difference between QAC and the C strategy is fourfold: QAC uses logical qubits, logical operators, and an energy penalty term, while C uses physical qubits, physical operators, and no penalty. The decoding strategy for QAC is a majority vote over the three data qubits of each logical qubit, while for C it is best of four copies of the logical problem solved by QAC. A more powerful, nested QAC strategy is known,²⁹ but it requires more physical qubits per logical qubit, and hence is less suitable for a scaling analysis of the type we perform here.

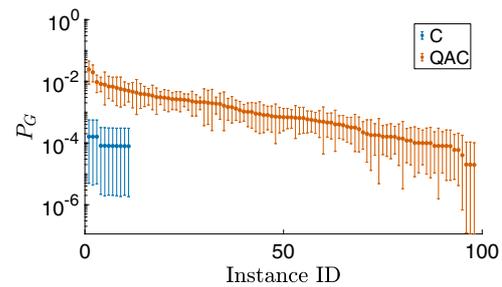


Fig. 4 Sorted success probabilities for $L = 14$ and $\eta = 0.07$, all 100 instances. C is able to find the correct ground state for only 11 instances. Meanwhile, QAC is able to find the correct ground state for all but 2 instances and increases the mean success probability over C.

We now discuss the results after the application of QAC and compare them to the C strategy. We note that this problem has also been studied using simulated quantum annealing (SQA);^{28,29} SQA is a good model for the thermalization of a quantum annealer.⁴⁰ Figure 4 illustrates that for relatively large problem sizes and strong added control noise, such as at $L = 14$ and $\eta = 0.07$, QAC is able to find nearly all ground states, while the C strategy only finds the ground state of a small fraction of instances.

More systematically, we show in Fig. 5(b) the fraction of instances where using QAC improved success probability when compared to the C strategy. If more than half of the instances exhibit better performance for the QAC strategy, applying it is useful for median instances. Evidently, QAC becomes a better strategy for large size and large noise, that is, as finding the ground state becomes harder. Similarly, we compare the J -chaoticity measure σ/μ for QAC and C, as seen in Fig. 5(a). Just as in Fig. 5(b), we see greater advantage using QAC over C with increasing size and noise. These observations are consistent with earlier results,^{25–28} where the $[3, 1, 3]_1$ performs better than the classical repetition strategy at large problem sizes. Here, we have shown that this is also true in the high control noise and J -chaos regime.

To understand the source of the improvement in the success probability, consider that application of repetition codes can decrease the effective noise on the encoded problem.²² In particular, the encoded operators of an n -qubit repetition code have an effective energy scale that scales extensively relative to the unencoded problem (see Methods), while the random control noise adds up incoherently and hence its energy contribution only scales up by a factor of \sqrt{n} . Thus, the encoding reduces the effective noise of the problem by the factor $1/\sqrt{n}$. This enhances the success probability of the encoded problem over the simple C strategy, essentially by leveraging just classical properties of the repetition code. However, there is a quantum mechanism at work as well: a mean-field analysis reveals that the penalty term reduces the tunneling barrier width and height in the QAC case.^{53,54} Indeed, we shall next see that our empirical results are inconsistent with the constant success probability enhancement that would be expected from a purely classical reduction of the effective noise by the factor $1/\sqrt{3}$ (we use an $n = 3$ repetition code).

Scaling of the time to solution

We now discuss the impact of the analog noise on the computational effort required to find a solution of the problem. This can be quantified by the time-to-solution (TTS) metric,³¹ which is the number of runs required to obtain the correct ground

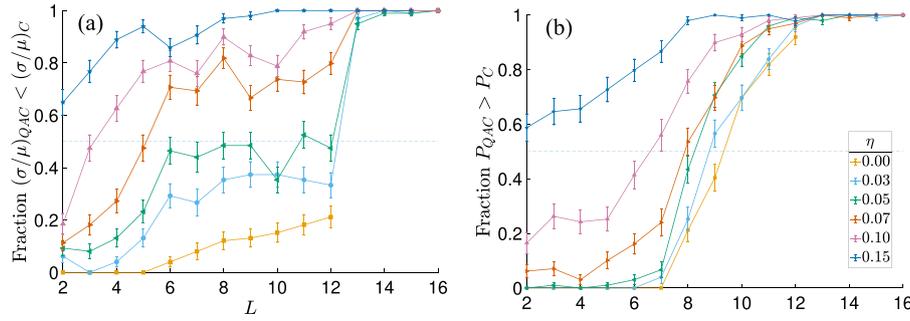


Fig. 5 Fraction of instances where QAC outperforms the C strategy. In **a**, we plot the fraction of instances where QAC found higher success probability than the C strategy after removing any instances where both failed completely (see Supplementary Information Section B.1). In **b**, we plot the fraction of instances where QAC lowered the J -chaoticity measure when compared to the C strategy. QAC becomes better with increasing size. At large noise, QAC becomes the better strategy for all instances. We note that here and the other figures below, we have omitted the $\eta = 0$ results for $L \geq 13$. The reason is a discontinuity between the DW2X and DW2000Q devices, which is discussed in detail in Supplementary Information Section C, and which is unrelated to the scaling analysis that forms the main focus of this work.

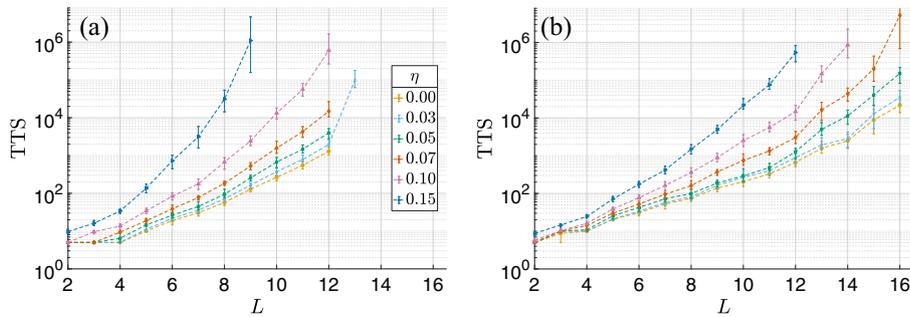


Fig. 6 TTS for QAC and the C strategy, sorted by added noise. We show the TTS to find at least one ground state of the median instance of class (L, η) . In **a**, we show the results from the classical repetition strategy. For large instances $L \geq 10$ and high noise $\eta = 0.15$, the C strategy fails to find any ground state for the median instance in our data set. In **b**, we show the results for QAC, which is always able to find the ground state of the median instance in our data set. The scaling in **b** is milder compared to **a**. For all $\eta > 0$, a missing data point indicates that the ground state was never found at that L value, for example, for all $L \geq 14$ in **a**.

state at least once with 99% success probability:

$$\text{TTS} = t_f \left\lceil \frac{\ln(1 - 0.99)}{\ln(1 - P_g)} \right\rceil, \quad (4)$$

where t_f is the total anneal time per run and P_g is the probability of finding the ground state in a given run (see Methods for details on how P_g was computed). The TTS metric is often used for benchmarking quantum annealing against classical algorithms (e.g., refs. 55–59). The TTS metric gives accurate scaling with problem size only when t_f is optimized to minimize the TTS for each size.^{13,31} In our experiments, we used a fixed $t_f = 5 \mu\text{s}$, and hence these results only place a lower bound on the true scaling,⁵⁵ but this is sufficient for our purposes. Since our anneal time was fixed, we actually report the number of runs $R = \text{TTS}/t_f$, which we still refer to as TTS. Note that $R \geq 1$ as one needs to run the annealer at least once to find the correct ground state.

In Fig. 6, we show the TTS required to find the ground state for the median instances at each size, for both C and QAC, sorted by different levels of added noise η . As expected for spin glasses, the TTS scales at least exponentially in L for both cases, with the scaling becoming worse for larger η . However, we note that QAC exhibits both milder scaling and lower absolute effort. The same conclusion holds when we sort instances by hardness (see Supplementary Information Section B.2). To more directly see the advantage of QAC over C, we plot the speedup ratio⁵⁵ in Fig. 7. This plot clearly shows the scaling advantage of QAC over C for sufficiently large size L and added noise η : for all added noise levels η , the slope of the speedup ratio becomes positive beyond an initial transient at small sizes L , and this happens sooner the larger η is.

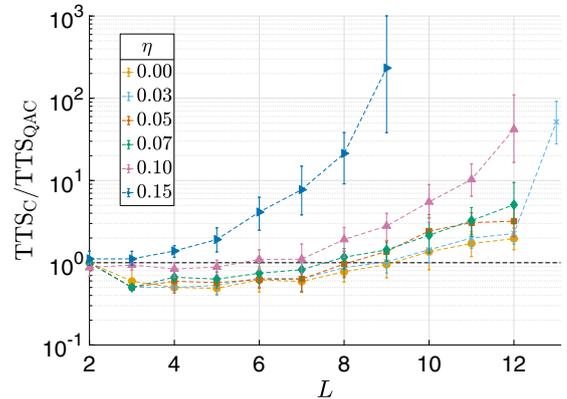


Fig. 7 Median speedup of QAC over C. We show the ratio of TTS of C and TTS of QAC, which will sit above the dashed line at speedup ratio = 1 for cases when it is advantageous to use QAC. Furthermore, a positive slope indicates a potential scaling speedup of QAC over C, since this occurs when the TTS of C is growing more quickly than of QAC.

Data collapse and scaling: doom vs. hope

We have seen that QAC outperforms the C strategy. But what is the worst-case classical cost of solving the same Ising problem instances? For a generalized Chimera graph of $L \times L$ unit cells of complete bipartite graphs $K_{r,r}$, the tree width is $w = rL + 1$; for the D-Wave devices used here $r = 4$. Dynamic programming takes time $O(L^2 2^w)$ to find the ground state of any Ising problem

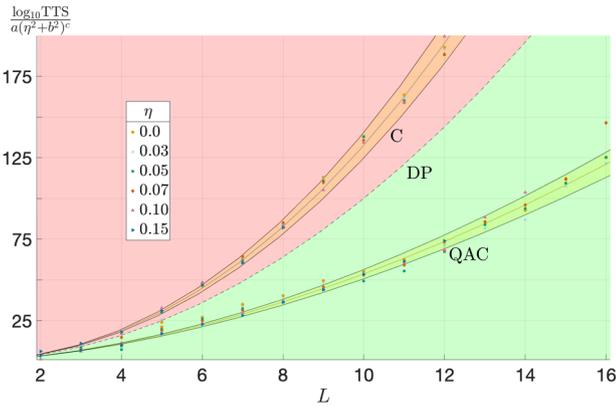


Fig. 8 Data collapse and fit. Result of the collapse of the data after fitting the TTS results shown in Fig. 6 to $10^{a(\eta^2+b^2)^c L^d}$ [Eq. (6)]. The dashed line is the asymptotic scaling of the classical dynamic programming algorithm, for which $\log_{10}(\text{TTS}) \sim L^2$. The red/green region above/below this line is where the C/QAC data lies after the data collapse. These correspond, respectively, to a guaranteed slowdown and a possible (but not guaranteed) speedup. The blue and red solid are the fits derived from the data collapse, with parameters given in Table 1. The shaded regions around the fitted lines represent the 95% CI fits as described in Supplementary Information Section B.3.

defined on such a graph.^{59,60} Here 2^w is the dimension of the exhaustive search space for each of the L^2 tree nodes of width w . Thus, in the present case, any problem can be solved exactly, deterministically, in time $\text{TTS}_{\text{DP}} = O(L^2 2^{4L})$. However, adding analog errors exponentially suppresses the probability of success. Specifically, if the errors are drawn from a Gaussian distribution with standard deviation η on an instance with N spins, then $P_g = O(e^{-\eta^2 N})$, where $\alpha \leq 1$ depends on the problem class.²¹ Thus, to find the ground state of the intended Hamiltonian H_{Ising} , running dynamic programming on the Ising instances with noise added is expected to take a time scaling as $\text{TTS}_{\text{DP}} \propto \left[\frac{\ln(1-0.99)}{\ln(1-P_g)} \right]$, which reduces, in the limit $P_g \ll 1$, to:

$$\text{TTS}_{\text{DP}}/P_g = O(L^2 2^{4L} e^{c_{\text{DP}} L^2}), \quad c_{\text{DP}} = 8\eta^2, \quad (5)$$

since the dynamic programming algorithm is only presented the intended Ising instance once every $1/P_g$ times on average. Thus, the worst-case classical cost is asymptotically $O(e^{8\eta^2 L^2})$. Note that this scaling is determined not by the intrinsic performance of the DP algorithm, but by the probability P_g that it is presented with the intended Hamiltonian, which is algorithm-independent (and is, of course, not a problem an algorithm running on classical digital computers would need to suffer from). A random guess would find the intended ground state with the same asymptotic scaling: $\text{TTS}_{\text{rand}}/P_g = 2^N e^{8\eta^2 L^2} = O(e^{c_{\text{rand}} L^2})$, but with a larger exponent: $c_{\text{rand}} = c_{\text{DP}} + 8 \ln 2$.

To compare the D-Wave device's TTS scaling to this form, we attempted a data collapse of the results shown in Fig. 6, in order to include the η dependence in the scaling function. To achieve the data collapse, we ran a comprehensive search for functions $f(L, \eta)$ that would collapse both the C and QAC data using as few fitting parameters as possible (see Methods for details of the procedure). A natural choice for such a function is a generalization of Eq. (5) with up to five free parameters, of the form $\text{TTS} = aL^2 10^{bL+c(\eta^2+d^2)^c L^d}$. However, we found it to perform poorly (see Supplementary Information Section B.3). Instead, we found that the four-parameter form

$$f(L, \eta) = 10^{a(\eta^2+b^2)^c L^d}, \quad (6)$$

where the crucial difference is the replacement of L^2 by L^d in the

Table 1. Fit parameters for Eq. (6) after data collapse of the TTS scaling data shown in Fig. 6.

Scheme	a	b	c	d
C	8.01	0.134	1.61	2.12
C lower	5.71	0.109	1.58	2.10
C upper	10.3	0.159	1.64	2.15
QAC	0.392	0.069	0.486	1.73
QAC lower	0.384	0.057	0.483	1.70
QAC upper	0.399	0.081	0.490	1.75

“Upper” and “lower” refers to the 95% CI values of the parameters, calculated as explained in detail in Supplementary Information Section B.3. Of particular note is the d parameter, which determines the asymptotic scaling. For QAC $d < 2$, while for C $d > 2$, with $d = 2$ being the scaling of an exhaustive classical solver

exponential, which works very well for both the C and QAC data (using three or fewer parameters give poor agreement). The data collapse and fit results are shown in Fig. 8, and the fit parameters along with their 95% confidence interval (CI) are given in Table 1. The relatively tight error bounds are evidence of the quality of the data collapse.

Surprisingly, we find that $d > 2$ for the C strategy, with high statistical confidence. This means that without error suppression, and even after using a majority vote among four copies of the problem, the performance of the quantum annealer is worse than that of a deterministic worst-case classical algorithm, for which $d = 2$. Hence the “doom” advertised in the title of this work. Of course, a powerful heuristic solver such as SQA⁴⁰ would outperform DP on these optimization problems, and hence the C strategy would perform even worse in comparison to SQA, which can be considered a classical algorithm that simulates the equilibration of the quantum annealer.

Fortunately, not all is lost: this disturbing finding is mitigated by QAC. As seen in Fig. 8 and Table 1, for QAC we obtain $d < 2$, again with high statistical confidence. This result restores the hope that a quantum annealer can eventually become competitive with classical optimization algorithms, but only after the incorporation of an error suppression and correction strategy such as QAC. We stress that we make no claim that QAC will outperform a sophisticated classical algorithm such as SQA, only that it gives quantum annealing a fighting chance.

DISCUSSION

It should be remarked that our results on optimization have no direct bearing on other tasks that quantum annealers are potentially capable of speeding up, such as approximate optimization^{61,62} and sampling.^{63–67} Nor do our results address quantum annealing slowdowns due to small gaps,^{68–71} which may be addressed via other methods, such as non-stoquastic Hamiltonians,^{72,73} reverse annealing,^{74–76} or inhomogeneous transverse field driving.^{77,78} However, none of these methods is immune to the effects of J -chaos.

To conclude, we have shown that QAC can reduce the detrimental effects of J -chaos on the performance of quantum annealers. In the regime we tested, QAC becomes more effective the higher the noise is and the larger the problem size is. The improvements seen are distinctly greater than without error suppression and correction, even after equalizing resources in terms of total qubit count, in terms of both scaling and absolute effort. Moreover, QAC undoes a catastrophic loss to an exhaustive classical algorithm by improving the scaling of the annealer's TTS to below the classical upper bound. Thus, we have demonstrated that QAC is not only an effective tool that can be used to improve

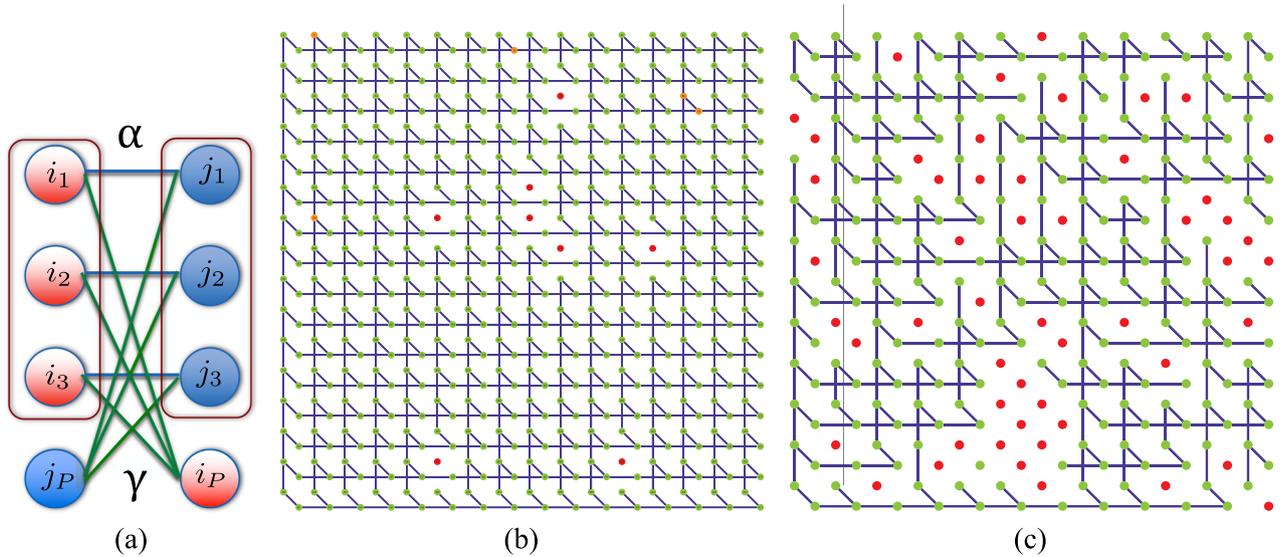


Fig. 9 The $[3, 1, 3]_1$ and the resulting logical graphs of the DW2X and DW2000Q devices. **a** A logical qubit of the $[3, 1, 3]_1$ is formed by connecting three qubits in one-half of the unit cell to a single qubit on the other half of the unit cell. The three qubits are called the “data qubits” and the single qubit plays the role of a “penalty qubit.” The data qubits are tied together with the penalty qubit by ferromagnetic couplings to ensure consistent evolution of each of the three data qubits. Each unit cell now contains two logical qubits. **b**, **c** Each logical qubit connects to its neighbor in the same cell via the intra-cell couplings and connects to its horizontal or vertical neighbor via the inter-cell couplings. This construction gives rise to a graph of degree 3 which is a minor embedding of the Chimera graph. **b** is the logical graph for the DW2X, **c** for the DW2000Q. In both graphs, green (red) circle denote operational (inactive) logical qubits. Orange circles denote logical qubits with intact data qubits, but an inactive penalty qubit. The lines denote the active couplings between the operational logical qubits. For each problem size L , we used a subgraph formed by taking an $L \times L$ square starting from the top left corner.

current quantum annealing hardware, but that error suppression and correction are essential to ensure competitive performance against classical alternatives. Further improvements using more powerful error suppression and correction strategies than the simple one we explored here are certainly expected, and undoubtedly necessary, as ultimately only a fully fault-tolerant approach is expected to be effective in the asymptotic limit of large problem sizes.

METHODS

Random Ising instances

Without noise. The set of instances used were generated randomly on the $[3, 1, 3]_1$ graph minors of the $L \times L$ Chimera graph for $L \in \{2, \dots, 12\}$ on the D-Wave 2X (DW2X) and $L \in \{13, \dots, 16\}$ on the D-Wave 2000Q (DW2000Q). There were 100 instances at each graph size such that the local fields, h_i , were 0 and the couplings J_{ij} were drawn uniformly at random from the set $\pm \frac{1}{6} \times \{1, 2, 3\}$. We found the ground state energy of these logical instances via the Hamze–Freitas–Selby (HFS) algorithm^{60,79} and parallel tempering with iso-energetic cluster (Houdayer) moves (PTICM).^{80,81} By using both, we consistently found ground state energies lower than or just as low as those found by the D-Wave devices. In a few instances, the latter found lower energies than HFS (one instance at $L = 15$ and five at $L = 16$), and these were confirmed as correct using PTICM. As such, we are confident that the ground state energies found were in fact correct.

With noise. We generated random numbers $\delta J_{ij} \sim \mathcal{N}(0, \eta^2)$. If a modified coupler value $\tilde{J}_{ij} = J_{ij} + \delta J_{ij}$ fell outside the experimentally allowed range $[-1, 1]$, we truncated it to ± 1 . Since the largest coupling in our set was 0.5 and the largest noise had a standard deviation of $\eta = 0.15$, these truncated values were only used a handful of times in our entire data collection.

The D-Wave devices used

In this study, we used the DW2X annealer installed at the USC Information Sciences Institute and the DW2000Q annealer installed at the NASA Quantum Artificial Intelligence Laboratory. The qubits of the annealer

occupy the vertices of the Chimera graph of 12×12 size for the DW2X and 16×16 for the DW2000Q (see Fig. 10 in Supplementary Information Section A). The DW2X has 1098 functional qubits, leading to a $[3, 1, 3]_1$ graph with 236 functional logical qubits. The DW2000Q has 2031 functional qubits, leading to a $[3, 1, 3]_1$ graph with 504 functional logical qubits, as shown in Fig. 9.

The annealing time t_f can be chosen in the range $[5, 2000] \mu\text{s}$ on the DW2X and $[1, 2000] \mu\text{s}$ on the DW2000Q. We used $t_f = 5 \mu\text{s}$, since this is the fastest time that could be used across both devices.

There were some differences between the structure and performance of these two devices. A discussion of these differences can be found in Supplementary Information Section C.

Data analysis

We used a Bayesian bootstrap⁸² over the underlying data (collected as described above) to compute the mean μ and the standard deviation σ of the success probabilities and their associated error bars. The ground state probability P_g used in Eq. (4) is the same as μ .

Consider g gauges where for each gauge i we find s_i successful readouts out of the total M readouts. We used the β -function $\beta(s_i, M - s_i)$ as our posterior probability distribution of success, that is, it is our best guess of the distribution of the success probability, given the observation of s_i successful hits in M attempts. To draw one sample of our bootstrap distribution, we did the following:

- First, sample from each of the β -distributions. Let B_i be a sample from the distribution $\beta(s_i, M - s_i)$ and let $\vec{B} = \{B_1, B_2, \dots, B_g\}$.
- Then, sample a point from the g -dimensional uniform Dirichlet distribution. This is a g -dimensional vector \vec{D} .
- The estimate of the success probability for this bootstrap sample is given by

$$\mu = \vec{D} \cdot \vec{B} = \sum_{i=1}^g D_i B_i \quad (7)$$

and the estimate of the standard deviation of this sample is the square root of the weighted variance of \vec{B} where the weights are given by \vec{D} ,

$$\sigma^2 = \sum_{i=1}^g D_i (B_i - \mu)^2. \quad (8)$$

In Eqs. (7) and (8), we have used the fact that the samples of the uniform Dirichlet distribution sum to 1: $\sum_i D_i = 1$.

- From these two quantities, we computed σ/μ . Other quantities of interest can be similarly derived from a combination of \bar{B} and \bar{D} .

We repeated these steps a large number of times to obtain a bootstrap distribution over our quantity of interest (in this case, μ and σ/μ). Our best estimate of the quantity and its associated error bars are given by the mean and the spread of the bootstrap distribution, respectively.

Data for C was collected as follows. For each instance and each noise value, we ran the annealer with five random gauges^{83,84} with 10,000 readouts each. If p is the success probability of one unencoded copy and each copy is statistically independent, then the C strategy will have at least one successful copy with probability $1 - (1 - p)^4$. Here, we only collected data for a single copy, and then used this combinatorial formula to get an estimate on the success probability of the classical repetition case, as in earlier work.²⁶ In fact, due to crosstalk this provides an upper bound on the actual performance of the C strategy (see Supplementary Information Section C.2), so that our results favor QAC over C even more than our plots indicate.

Data for QAC (see below) was collected as follows. Every data collection run for problem instance i of size L can be labeled by two additional parameters; the strength of artificial injected noise η and the strength of the penalty value γ . The penalty strength was chosen from the set $\gamma \in \{0.1, 0.2, \dots, 0.5\}$. For each data collection run (i, L, η, γ) , we ran the annealer with five random gauges with 10,000 readouts each, for a total of 50,000 readouts. From these data we estimated the mean success probability over the gauges, $P_g(i, L, \eta, \gamma)$. The optimal penalty value

$$\gamma_{\text{opt}} = \arg \max_{\gamma} P_g(i, L, \eta, \gamma) \quad (9)$$

maximizes the success probability within the chosen range of penalty values. The results shown in the Results section were picked to be at this optimal value for each instance. Histograms of the optimal strengths for each problem size are shown in Supplementary Information Section D.

Quantum annealing correction

In QAC we encode each logical Pauli-Z operator as a sum of n such physical operators, that is,

$$\bar{\sigma}_i^z = \sum_{l=1}^n \sigma_{i_l}^z, \quad \bar{\sigma}_i^z \bar{\sigma}_j^z = \sum_{l=1}^n \sigma_{i_l}^z \sigma_{j_l}^z. \quad (10)$$

Furthermore, we add a term to ferromagnetically couple the physical copies through an auxiliary qubit, that is,

$$H_p = - \sum_{i=1}^N \sigma_{i_p}^z \bar{\sigma}_i^z, \quad (11)$$

where N is the number of logical variables in the original optimization problem. We refer to the physical copies, $\sigma_{i_l}^z$, as data qubits and the auxiliary qubit, $\sigma_{i_p}^z$, as a penalty qubit. Thus, we arrive at the following encoding of our logical problem:

$$\bar{H}(s) = A(s)H_X + B(s)(a\bar{H}_{\text{Ising}} + \gamma H_p), \quad (12)$$

where \bar{H}_{Ising} is the encoded version of H_{Ising} in Eq. (2a), that is, $\sigma_i^z \mapsto \bar{\sigma}_i^z$ and $\sigma_i^z \sigma_j^z \mapsto \bar{\sigma}_i^z \bar{\sigma}_j^z$, and a is an overall energy scale for the problem Hamiltonian (not used in this work, but complementary to adding control noise.²⁹) When we add control noise to the QAC Hamiltonian, we replace h_i by $\tilde{h}_i = h_i + \delta h_i$ and J_{ij} by $\tilde{J}_{ij} = J_{ij} + \delta J_{ij}$. The perturbations δh_i and δJ_{ij} are chosen independently for each qubit i and coupler (i, j) from the distribution specified in Eq. (3).

The current generations of D-Wave devices allow a direct implementation of this code in the Ising Hamiltonian for $n = 3$, as shown in Fig. 9, but are unable to encode the driver Hamiltonian H_X , as this requires many-body terms of the form $(\sigma^x)^{\otimes n}$. Thus, increasing the penalty strength, γ , begins to diminish the effect of the quantum fluctuations that drive quantum annealing. On the other hand, larger γ values are more able to suppress bit-flip errors. Thus, there exists an optimal value of γ , which depends on the spectrum of the problem instance.^{25–28}

After annealing, we obtain a state vector where each data qubit is measured in the computational basis. From this, we can obtain a state vector of logical qubits via a variety of decoding strategies.²⁸ In this work, we exclusively used the strategy in which each logical qubit is decoded by a majority vote of its constituent data qubits.

Data collapse

Here we explain our procedure for identifying the optimal fit and data collapse function, and for extracting CIs and error bars. We considered trial TTS functions of the form $f(L, \eta) = 10^{g(L, \eta)}$, with:

$$g_1 = a(\eta^2 + b^2)^c L^d, \quad g_2 = g_1 + \log_{10}(e), \quad (13a)$$

$$g_3 = aL + c(\eta^2 + b^2)^{d_1} L^{d_2} + \log_{10}(eL^2). \quad (13b)$$

For g_3 we focused on the three cases $\{d_1 = 1/2, d_2 = 2\}$, $\{d_1 = d, d_2 = 2\}$, and $\{d_1 = 1/2, d_2 = d\}$. Thus, our trial functions had either four ($\{a, b, c, d\}$) or five ($\{a, b, c, d, e\}$) free fitting parameters. For each trial function we computed non-linear least-squares fits to the median TTS data for C on $L \in \{2, \dots, 12\}$ and QAC on $L \in \{2, \dots, 16\}$. The fitting parameters were initially allowed to take any values. However, we only accepted fits with $a \geq 0$ in g_3 , since $a < 0$ (scaling that decreases with L) would have to reflect overfitting. Thus, we also computed fits where we squared all the fitting parameters [i.e., replaced a by a^2 in Eqs. (13a) and (13b), etc.] in order to enforce positivity. Furthermore, we tested if the discrepancy between the ideal and actual number of Chimera graph couplers made a difference by fitting with an effective L ; see Supplementary Information Section C for details. Thus, for each trial function there were four different methods: unconstrained/squared fitting parameters with L /effective L . Lastly, all fits were attempted with each of the optimization methods possible in Mathematica: SimulatedAnnealing, RandomSearch, NelderMead, and DifferentialEvolution.

Across the different methods and optimization algorithms used, g_1 was consistently the best of the 4-parameter fits and was always very close to g_2 , which is its 5-parameter generalization. The g_3 functions always resulted either in $a < 0$ or otherwise a very poor fit. Parameter squaring also improved the fit quality, and of the optimization methods only NelderMead tended to give inferior results.

After determining the three parameters $\{a, b, c\}$ for the median TTS data for g_1 , we found least-squares fits to the upper and lower bounds determined by the 95% CIs for the median TTS, by using the same set $\{a, b, c\}$ and letting only d be a free parameter. In this manner we found d_- and d_+ , the exponents that provide respective lower and upper bounds on d for the median TTS data. In turn, d_- and d_+ have associated 95% CIs, denoted Δd_- and Δd_+ . The reported range of d in Table 1 is then $[d_- - \Delta d_-, d_+ + \Delta d_+]$. The resulting fits for each η are shown in Supplementary Information Section B.3.

DATA AVAILABILITY

All raw data is available upon reasonable request from the corresponding author. A Mathematica notebook containing the TTS results, analysis scripts, error analysis, and our detailed fitting and data collapse results is available.⁸⁵

Received: 15 August 2019; Accepted: 15 October 2019;

Published online: 28 November 2019

REFERENCES

- Shor, P. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.* **26**, 1484–1509 (1997).
- Bravyi, S., Gosset, D. & König, R. Quantum advantage with shallow circuits. *Science* **362**, 308 (2018).
- Barends, R. et al. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature* **508**, 500–503 (2014).
- Landauer, R. Is quantum mechanics useful? *Proc. R. Soc. Lond. Ser. A* **353**, 367 (1995).
- Shor, P.W. Fault-tolerant quantum computation. In *Proc. 37th Conference on Foundations of Computer Science* 56–65 (1996). <http://ieeexplore.ieee.org/document/548464/>.
- Aliferis, P., Gottesman, D. & Preskill, J. Quantum accuracy threshold for concatenated distance-3 codes. *Quant. Inf. Comput.* **6**, 97 (2006).
- Chao, R. & Reichardt, B. W. Fault-tolerant quantum computation with few qubits. *npj Quant. Inform.* **4**, 42 (2018).
- Lidar, D. & Brun, T. (eds). *Quantum Error Correction* (Cambridge University Press, Cambridge, 2013).
- Farhi, E., Goldstone, J., Gutmann, S. & Sipser, M. Quantum computation by adiabatic evolution. Preprint at <http://arxiv.org/abs/quant-ph/0001106> (2000).
- Kadowaki, T. & Nishimori, H. Quantum annealing in the transverse Ising model. *Phys. Rev. E* **58**, 5355 (1998).

11. Childs, A. M., Farhi, E. & Preskill, J. Robustness of adiabatic quantum computation. *Phys. Rev. A* **65**, 012322 (2001).
12. Preskill, J. Quantum computing in the nisq era and beyond. *Quantum* **2**, 79 (2018).
13. Albash, T. & Lidar, D. A. Demonstration of a scaling advantage for a quantum annealer over simulated annealing. *Phys. Rev. X* **8**, 031016 (2018).
14. Mandrà, S. & Katzgraber, H. G. A deceptive step towards quantum speedup detection. *Quant. Sci. Technol.* **3**, 04LT01 (2018).
15. Bray, A. J. & Moore, M. A. Chaotic nature of the spin-glass phase. *Phys. Rev. Lett.* **58**, 57–60 (1987).
16. Katzgraber, H. G. & Krzakala, F. Temperature and disorder chaos in three-dimensional Ising spin glasses. *Phys. Rev. Lett.* **98**, 017201 (2007).
17. Bunyk, P. I. et al. Architectural considerations in the design of a superconducting quantum annealing processor. *IEEE Trans. Appl. Superconductivity* **24**, 1–10 (2014).
18. Harris, R. et al. Experimental demonstration of a robust and scalable flux qubit. *Phys. Rev. B* **81**, 134510 (2010).
19. Harris, R. et al. Experimental investigation of an eight-qubit unit cell in a superconducting optimization processor. *Phys. Rev. B* **82**, 024511 (2010).
20. Martin-Mayor, V. & Hen, I. Unraveling quantum annealers using classical hardness. *Scientific Rep.* **5**, 15324 EP – (2015).
21. Albash, T., Martin-Mayor, V. & Hen, I. Analog errors in Ising machines. *Quant. Sci. Technol.* **4**, 02LT03 (2019).
22. Young, K. C., Blume-Kohout, R. & Lidar, D. A. Adiabatic quantum optimization with the wrong hamiltonian. *Phys. Rev. A* **88**, 062314 (2013).
23. Bian, Z. et al. Discrete optimization using quantum annealing on sparse Ising models. *Front. Phys.* **2**, 56 (2014).
24. Zhu, Z., Ochoa, A. J., Schnabel, S., Hamze, F. & Katzgraber, H. G. Best-case performance of quantum annealers on native spin-glass benchmarks: How chaos can affect success probabilities. *Phys. Rev. A* **93**, 012317 (2016).
25. Pudenz, K. L., Albash, T. & Lidar, D. A. Error-corrected quantum annealing with hundreds of qubits. *Nat. Commun.* **5**, 3243 (2014).
26. Pudenz, K. L., Albash, T. & Lidar, D. A. Quantum annealing correction for random Ising problems. *Phys. Rev. A* **91**, 042302 (2015).
27. Mishra, A., Albash, T. & Lidar, D. A. Performance of two different quantum annealing correction codes. *Quant. Inf. Proc.* **15**, 609–636 (2015).
28. Vinci, W., Albash, T., Paz-Silva, G., Hen, I. & Lidar, D. A. Quantum annealing correction with minor embedding. *Phys. Rev. A* **92**, 042310 (2015).
29. Vinci, W., Albash, T. & Lidar, D. A. Nested quantum annealing correction. *npj Quant. Inf.* **2**, 16017 (2016).
30. Vinci, W. & Lidar, D. A. Scalable effective-temperature reduction for quantum annealers via nested quantum annealing correction. *Phys. Rev. A* **97**, 022308 (2018).
31. Rønnow, T. F. et al. Defining and detecting quantum speedup. *Science* **345**, 420–424 (2014).
32. Weber, S. J. et al. Coherent coupled qubits for quantum annealing. *Phys. Rev. Appl.* **8**, 014004 (2017).
33. Novikov, S. et al. Exploring more-coherent quantum annealing. (2018) *IEEE International Conference on Rebooting Computing (ICRC)*, McLean, VA, USA, 2018, pp. 1–7. <https://doi.org/10.1109/ICRC.2018.8638625>.
34. Goto, H. Quantum computation based on quantum adiabatic bifurcations of Kerr-nonlinear parametric oscillators. *J. Phys. Soc. Jpn* **88**, 061015 (2019).
35. Inagaki, T. et al. A coherent Ising machine for 2000-node optimization problems. *Science* **354**, 603 (2016).
36. Goto, H., Tatumura, K. & Dixon, A. R. Combinatorial optimization by simulating adiabatic bifurcations in nonlinear Hamiltonian systems. *Sci. Adv.* **5**, eaav2372 (2019).
37. Das, A. & Chakrabarti, B. K. Colloquium: quantum annealing and analog quantum computation. *Rev. Mod. Phys.* **80**, 1061–1081 (2008).
38. 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 (1989).
39. Brooke, J., Bitko, D., Rosenbaum, F. T. & Aeppli, G. Quantum annealing of a disordered magnet. *Science* **284**, 779–781 (1999).
40. Santoro, G. E., Martoňák, R., Tosatti, E. & Car, R. Theory of quantum annealing of an Ising spin glass. *Science* **295**, 2427–2430 (2002).
41. Boixo, S. et al. Computational multiqubit tunnelling in programmable quantum annealers. *Nat. Commun.* **7**, 10327 (2016).
42. Muthukrishnan, S., Albash, T. & Lidar, D. A. Tunneling and speedup in quantum optimization for permutation-symmetric problems. *Phys. Rev. X* **6**, 031010 (2016).
43. Denchev, V. S. et al. What is the computational value of finite-range tunneling? *Phys. Rev. X* **6**, 031015 (2016).
44. Johnson, M. W. et al. Quantum annealing with manufactured spins. *Nature* **473**, 194–198 (2011).
45. Jansen, S., Ruskai, M.-B. & Seiler, R. Bounds for the adiabatic approximation with applications to quantum computation. *J. Math. Phys.* **48**, 102111 (2007).
46. Lidar, D. A., Rezaekhani, A. T. & Hama, A. Adiabatic approximation with exponential accuracy for many-body systems and quantum computation. *J. Math. Phys.* **50**, 102106 (2009).
47. Avron, J. E., Fraas, M., Graf, G. M. & Grech, P. Adiabatic theorems for generators of contracting evolutions. *Comm. Math. Phys.* **314**, 163–191 (2012).
48. Venuti, L. C., Albash, T., Lidar, D. A. & Zanardi, P. Adiabaticity in open quantum systems. *Phys. Rev. A* **93**, 032118 (2016).
49. The D-Wave documentation refers to this as integrated control errors (ICEs). According to https://docs.dwavesys.com/docs/latest/c_qpu_1.html, $\eta_{\text{int}} \leq 0.015$ for the DW2000Q device).
50. Boixo, S. et al. Evidence for quantum annealing with more than one hundred qubits. *Nat. Phys.* **10**, 218–224 (2014).
51. Katzgraber, H. G., Hamze, F. & Andrist, R. S. Glassy chimeras could be blind to quantum speedup: designing better benchmarks for quantum annealing machines. *Phys. Rev. X* **4**, 021008 (2014).
52. King, A. D. & McGeoch, C. C. Algorithm engineering for a quantum annealing platform. Preprint at <http://arxiv.org/abs/1410.2628> (2014).
53. Matsuura, S., Nishimori, H., Albash, T. & Lidar, D. A. Mean field analysis of quantum annealing correction. *Phys. Rev. Lett.* **116**, 220501 (2016).
54. Matsuura, S., Nishimori, H., Vinci, W., Albash, T. & Lidar, D. A. Quantum-annealing correction at finite temperature: ferromagnetic p -spin models. *Phys. Rev. A* **95**, 022308 (2017).
55. Hen, I. et al. Probing for quantum speedup in spin-glass problems with planted solutions. *Phys. Rev. A* **92**, 042325 (2015).
56. Mandrà, S., Zhu, Z., Wang, W., Perdomo-Ortiz, A. & Katzgraber, H. G. Strengths and weaknesses of weak-strong cluster problems: a detailed overview of state-of-the-art classical heuristics versus quantum approaches. *Phys. Rev. A* **94**, 022337 (2016).
57. King, J. et al. Quantum annealing amid local ruggedness and global frustration. Preprint at <http://arxiv.org/abs/1701.04579> (2017).
58. Hamerly, R. et al. Experimental investigation of performance differences between coherent Ising machines and a quantum annealer. *Sci. Adv.* **5**, eaau0823 (2019).
59. Jünger, M. et al. Performance of a quantum annealer for Ising ground state computations on chimera graphs. Preprint at <https://arxiv.org/abs/1904.11965> (2019).
60. Selby, A. Efficient subgraph-based sampling of Ising-type models with frustration. Preprint at <http://arxiv.org/abs/1409.3934> (2014).
61. King, J., Yarkoni, S., Nevisi, M. M., Hilton, J. P. & McGeoch, C. C. Benchmarking a quantum annealing processor with the time-to-target metric. Preprint at <http://arxiv.org/abs/1508.05087> (2015).
62. Vinci, W. & Lidar, D. A. Optimally stopped optimization. *Phys. Rev. Appl.* **6**, 054016 (2016).
63. Adachi, S. H. & Henderson, M. P. Application of quantum annealing to training of deep neural networks. Preprint at <http://arxiv.org/abs/1510.06356> (2015).
64. Amin, M. H., Andriyash, E., Rolfe, J., Kulchitsky, B. & Melko, R. Quantum Boltzmann machine. *Phys. Rev. X* **8**, 021050 (2018).
65. Mott, A., Job, J., Vlimant, J.-R., Lidar, D. & Spiropulu, M. Solving a Higgs optimization problem with quantum annealing for machine learning. *Nature* **550**, 375 EP – (2017).
66. Li, R. Y., DiFelice, R., Rohs, R. & Lidar, D. A. Quantum annealing versus classical machine learning applied to a simplified computational biology problem. *npj Quantum Inf.* **4**, 14 (2018).
67. Perdomo-Ortiz, A., Benedetti, M., Realpe-Gómez, J. & Biswas, R. Opportunities and challenges for quantum-assisted machine learning in near-term quantum computers. *Quantum Sci. Technol.* **3**, 030502 (2018).
68. van Dam, W., Mosca, M. & Vazirani, U. How powerful is adiabatic quantum computation? In *Proc. 42nd IEEE Symposium on Foundations of Computer Science, 2001* 279–287 (8–11 Oct 2001).
69. Reichardt, B. W. The quantum adiabatic optimization algorithm and local minima. In *Proc. Thirty-sixth Annual ACM Symposium on Theory of Computing, STOC '04*, 502–510 (ACM, New York, NY, 2004). Erratum: <http://www-bcf.usc.edu/breichar/Correction.txt>.
70. Jörg, T., Krzakala, F., Kurchan, J., Maggs, A. C. & Pujos, J. Energy gaps in quantum first-order mean-field-like transitions: the problems that quantum annealing cannot solve. *Europhys. Lett.* **89**, 40004 (2010).
71. Laumann, C. R., Moessner, R., Scardicchio, A. & Sondhi, S. L. Quantum adiabatic algorithm and scaling of gaps at first-order quantum phase transitions. *Phys. Rev. Lett.* **109**, 030502 (2012).
72. Nishimori, H. & Takada, K. Exponential enhancement of the efficiency of quantum annealing by non-stoquastic Hamiltonians. *Front. ICT* **4**, 2 (2017).
73. Albash, T. Role of nonstoquastic catalysts in quantum adiabatic optimization. *Phys. Rev. A* **99**, 042334 (2019).

74. Perdomo-Ortiz, A., Venegas-Andraca, S. E. & Aspuru-Guzik, A. A study of heuristic guesses for adiabatic quantum computation. *Quantum Inf. Process.* **10**, 33–52 (2011).
75. Chancellor, N. Modernizing quantum annealing using local searches. *N. J. Phys.* **19**, 023024 (2017).
76. Ohkuwa, M., Nishimori, H. & Lidar, D. A. Reverse annealing for the fully connected SpS -spin model. *Phys. Rev. A* **98**, 022314 (2018).
77. Susa, Y., Yamashiro, Y., Yamamoto, M. & Nishimori, H. Exponential speedup of quantum annealing by inhomogeneous driving of the transverse field. *J. Phys. Soc. Jpn* **87**, 023002 (2018).
78. Susa, Y. et al. Quantum annealing of the SpS -spin model under inhomogeneous transverse field driving. *Phys. Rev. A* **98**, 042326 (2018).
79. Hamze, F. & de Freitas, N. in *UAI* (eds Chickering, D. M. & Halpern, J. Y.) 243–250 (AUAI Press, Arlington, 2004). <http://dl.acm.org/citation.cfm?id=1036873>.
80. Houdayer, J. A cluster Monte Carlo algorithm for 2-dimensional spin glasses. *Eur. Phys. J. B* **22**, 479–484 (2001).
81. Zhu, Z., Ochoa, A. J. & Katzgraber, H. G. Efficient cluster algorithm for spin glasses in any space dimension. *Phys. Rev. Lett.* **115**, 077201 (2015).
82. Rubin, D. B. The Bayesian bootstrap. *Ann. Statist.* **9**, 130–134 (1981).
83. Boixo, S., Albash, T., Spedalieri, F. M., Chancellor, N. & Lidar, D. A. Experimental signature of programmable quantum annealing. *Nat. Commun.* **4**, 2067 (2013).
84. Job, J. & Lidar, D. Test-driving 1000 qubits. *Quantum Sci. Technol.* **3**, 030501 (2018).
85. Pearson, A. & Lidar, D. datacollapse.nb. Mathematica notebook for data collapse and fitting (2019). <https://www.notebookarchive.org/data-collapse-from-analog-errors-in-quantum-annealing-doom-and-hope--2019-10-ana4eit/>.

ACKNOWLEDGEMENTS

We are grateful to Tameem Albash for useful discussions. The computing resources used for this work were provided by the USC Center for High Performance Computing and Communications. 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 US 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 US Government. The US Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. This work was partially supported by Oracle.

AUTHOR CONTRIBUTIONS

A.P. and A.M. are co-first authors. D.L. and I.H. conceived the project. A.M. performed the DW2X data collection and initial data analysis. A.P. performed the DW2000Q data collection and subsequent data analysis. D.L. guided the entire project. D.L. and A.P. wrote the manuscript with input from A.M. and I.H.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary information is available for this paper at <https://doi.org/10.1038/s41534-019-0210-7>.

Correspondence and requests for materials should be addressed to D.A.L.

Reprints and permission information is available at <http://www.nature.com/reprints>

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) 2019