Durham Research Online

Deposited in DRO:

31 March 2021

Version of attached file:

Published Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Chancellor, Nicholas (2020) 'Fluctuation-guided search in quantum annealing.', Physical review A., 102 (6). 062606.

Further information on publisher's website:

<https://doi.org/10.1103/PhysRevA.102.062606>

Publisher's copyright statement:

©2020 American Physical Society

Additional information:

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

Fluctuation-guided search in quantum annealingNicholas Chancellor ^{*}*Department of Physics and Durham Newcastle Joint Quantum Centre Durham University, South Road, Durham, United Kingdom*

(Received 30 September 2020; accepted 9 November 2020; published 9 December 2020)

Quantum annealing has great promise in leveraging quantum mechanics to solve combinatorial optimization problems. However, to realize this promise to its fullest extent we must appropriately leverage the underlying physics. In this spirit, we examine how the well-known tendency of quantum annealers to seek solutions where more quantum fluctuations are allowed can be used to trade off optimality of the solution to a synthetic problem for the ability to have a more flexible solution, where some variables can be changed at little or no cost. We demonstrate this tradeoff experimentally using the reverse annealing feature a D-Wave Systems quantum processing unit for both problems composed of all binary variables, and those containing some higher-than-binary discrete variables. We further demonstrate how local controls on the qubits can be used to control the levels of fluctuations and guide the search. We discuss places where leveraging this tradeoff could be practically important, namely in hybrid algorithms where some penalties cannot be directly implemented on the annealer and provide some proof-of-concept evidence of how these algorithms could work.

DOI: [10.1103/PhysRevA.102.062606](https://doi.org/10.1103/PhysRevA.102.062606)**I. INTRODUCTION AND BACKGROUND**

Quantum annealing, in which combinatorial optimization problems are mapped directly Hamiltonians and solved using sweeps of Hamiltonian parameters, has been a subject of much interest recently. This is in part due to the wide variety of potential applications, in a diverse range of subjects, including for instance air traffic control [1], hydrology [2], protein folding [3], flight gate assignment [4], finance [5–7], and even quantum field theory [8,9]. This subject has further attracted interest because of the experimental maturity of the flux qubit devices produced by D-Wave Systems Inc. which allow for large-scale experimentation.

One crucial direction in the growth of flux qubit quantum annealing is an increase in the variety of controls which users can be applied to the experimental quantum annealing process on flux qubit annealers. Traditionally formulated quantum annealing starts from an easy to prepare ground state of a so-called driver Hamiltonian and monotonically interpolates the Hamiltonian to a problem Hamiltonian with an unknown ground state. However, major advantages can be gained by using a different control pattern known as reverse annealing, which starts in a state which is a guess for the solution of the optimization problem, turns on fluctuations, and searches nearby states in Hamming distance by taking advantage of thermal dissipation [10]. Likewise, controls have been added which allow different qubits to be annealed differently [11].

These features have proven useful in a variety of ways, reverse annealing for instance is motivated by the ability to implement more complex algorithms than traditional forward annealing [12], and these algorithms have shown promising initial experimental results. For example, it was shown in

Ref. [6] that starting from the output of a simple classical algorithm can lead to a large improvement over forward annealing. References [13,14] showed that iterative methods can help over non-negative matrix factorization. The work in Ref. [15] showed experimentally that adding mutation performed using reverse annealing can aid the performance of genetic algorithms. Furthermore the simulation of the celebrated Kosterlitz-Thouless phase transition in Ref. [16] would not have been possible without reverse annealing techniques, and reverse annealing has been used to simulate quantum field theories [9].

Similarly, anneal offsets have shown promise in synchronizing the freezing of qubits [11,17]. The algorithmic use of these controls was initially motivated by numerical work which demonstrated that locally varying transverse fields can mitigate perturbative anticrossings [18]. Fluctuations which can be mitigated using these tools were shown [19] to be important in the heavy tails which were observed in early experimental work on D-Wave quantum processing units (QPUs) [20]. The demonstration that anneal offsets can mitigate these effects came later, in Ref. [17].

The tendency of quantum fluctuations to lead to uneven sampling of ground state manifolds has traditionally been viewed as a drawback for quantum annealing [21–23]. However, it has been observed that when coupled with classical techniques, this uneven sampling could be a positive feature because the states which quantum annealers find tend to be very different from those found by classical solvers, and therefore could give a more complete picture of the manifold [24] if both were used together. In all of this previous work, the uneven sampling was found to be due to quantum, rather than thermal, fluctuations. While thermal fluctuations play an unavoidable role in the experiments reported in this paper, the tendency to favor solutions with more free spins is likely due to the same quantum effects observed in this previous work.

^{*}nicholas.chancellor@gmail.com

Furthermore, the tendency of the annealer to seek out states with more free spins plays a fundamental role in the effect observed in Ref. [16]. The role of quantum fluctuations in causing uneven sampling is also important to graph isomorphism applications which have been proposed and implemented on quantum annealers [25–27]. In this paper, we explore a different advantage of this preferential search, the fact that it tends to find states which are flexible in the sense that some variables can be changed at little to no energy cost.

In this paper we experimentally investigate the role which quantum fluctuations can play in the local search which reverse annealing implements. This is done by using specialized Hamiltonians which represent hard problems for the annealer (although not necessarily hard in the computational sense) and have sets of local minima in their energy landscape where fluctuations are enhanced. We also show that the anneal offsets (controls offsets on chains or gadgets which enhance ($\delta s < 0$) or suppress ($\delta s > 0$) fluctuations) can be used to guide the search by locally enhancing fluctuations on some parts of the system. The technique of locally enhancing fluctuations is reminiscent of the methods proposed in Ref. [28] and provides some experimental validation of these concepts.

To do this, we construct problems where there is a planted solution which is guaranteed to have the lowest energy, but which does not support strong quantum fluctuations, but fluctuations can be increased by paying an energy cost. We study structures which include both “gadgets” which involve binary variables which, in some configurations, can be flipped with no energy penalty. We also investigate higher than binary variables encoded into “chains (Hamiltonian element which encodes a discrete variable)” which have a “soft” configuration where small changes have little or not effect on their energy, but start in a configuration where a small change has a large effect. We also study the effects of anneal offsets which can locally increase or suppress fluctuations.

We use these fluctuations to trade off optimality in solutions for flexibility, in other words find solutions which are a bit less optimal, but for which certain variables can be changed at little or no cost. These variables could either be binary variables directly represented by qubits, or discrete variables which can be encoded using a method we describe later. We argue that this is a property which is likely to be relevant in some real world situations and give a motivational example of how it can be used in a hybrid quantum classical algorithm to find a more optimal solution in the presence of a global penalty function which is not encoded into the annealer.

On the devices studied here (D-Wave 2000Q quantum processing units), dissipation plays an important (often positive [29]) role in the annealing process, and the reverse annealing techniques used here fundamentally rely on dissipation. Dissipation can also play a very detrimental role, as pointed out in earlier works such as Ref. [30]; however, when techniques like reverse annealing are included, even devices with more limited coherence still present opportunities for quantum advantage [31]. Recent experiments, have suggested that improvements can be made by reducing the noise on the current D-Wave devices [32,33]; however, whether these

improvements would continue until the noise is reduced to zero is an open question.

The intuition developed here, however, is likely to carry over into the more coherent protocols proposed in Refs. [34–36]. This is relevant because coherence rates can be improved through a variety of routes, in both superconducting flux qubit architectures [32,37] and trapped ion quantum annealers [38]. Furthermore, there is significant evidence that in the fully coherent regime fast, but coherent, quenches known as “diabatic” quantum computing may be a promising path to a quantum advantage [31]. This is due to both adiabatic mechanisms involving multiple energy levels [31,39] and mechanisms related to energy transfer [40–42].

Because the experimental details are likely to be interesting only to readers who engage with experiments on the D-Wave devices at a relatively low level, we have reserved a detailed description for the Appendix. In Sec. II we give the core experimental results, demonstrating how fluctuations can enable a tradeoff between optimality and flexibility of solutions, as well as how anneal offsets can be used to guide the search by emulating these nonengineered fluctuations. Next, in Sec. III we give a motivational example of how trading off optimality and flexibility can be useful. We then discuss some of the more detailed aspects of the experimental methods and concluded the paper with some discussion.

II. RESULTS

In this section we discuss the results of the experiments, which demonstrate how both existing and introduced fluctuations can be used to guide the search which a quantum annealer performs. First, we will introduce how the number of free gadgets (gadget in the configuration which allows more fluctuations) or soft chains (chain in a configuration where more fluctuations are allowed) can be controlled by different parameters, such as the value of the reversal parameter s^* and the anneal offsets δs applied to the chains or gadgets. The parameter s^* controls the range of reverse annealing search with $s^* = 1$ corresponds to no search, $s^* = 0$ to forward annealing. Meanwhile δs controls offsets on chains or gadgets which enhance ($\delta s < 0$) or suppress ($\delta s > 0$) fluctuations. Measures of the performance of these different control settings will be introduced in Sec. II A and further discussed in Sec. II B. A proof-of-principle example for how guided search can be useful will be discussed in Sec. III.

The first result which we find is that the number of free gadgets and soft chains both can be increased by decreasing the value of s^* , in other words by increasing the range of the search. Figure 1 shows this effect for gadgets, not only are more free gadgets the lower value of s^* , this effect is also much stronger when the gadgets are not locked, indicating that the free variables have a significant effect on the dynamics. For $s^* \gtrsim 0.45$ the dynamics are highly localized and very few if any gadgets are free, meanwhile for $s^* \lesssim 0.38$, the behavior is indistinguishable from a search with $s^* = 0.2$, effectively a global search. We have chosen a nonuniform mesh of s^* values which focuses on the regime where the reverse anneal can lead to long-range dynamics, but does not search so far that all information about the initial state is completely forgotten.

FIG. 1. Fraction of observations for different numbers of free gadgets against different values of s^* averaged over all 10 Hamiltonians (recall that s^* is a unitless quantity). Solid red (gray) lines with “X” markers represent the mean for the plot, while the dashed line with “+” markers is the mean of the other plot for comparison. (a) Unlocked gadgets with no anneal offsets. (b) Locked gadgets with no anneal offsets. The methods for creating the plots given the nonlinear mesh are explained in the Appendix.

The experiments we have conducted can be understood as probing a finite size precursor to a transition from the phase which is realized at $s^* = 1$ and the paramagnetic phase found at $s^* = 0$. The exact nature of the phase which is realized in the large system limit for $s^* = 1$ is not immediately clear so therefore neither is the nature of the transition and the factors which affect its location. The fact that annealing from the paramagnetic precursor tends not to find the ground state provides weak evidence against this phase transition being a simple transition into a ferromagnetic phase and is more consistent with a spin glass [43] or a Griffiths-type phase transition [44,45].

Figure 2 shows the same effect for embedded chains in planted solution problems. In this case, reducing the fluctuations by increasing the softness coefficient leads to fewer soft chains. As with the gadget example, nontrivial reverse annealing dynamics are seen for $0.38 \lesssim s^* \lesssim 0.45$.

Let us further observe that if we apply anneal offsets to the locked gadgets [version of gadget element which does not allow for more fluctuations (used as comparison point)], we can mimic the effect of the free variables, as Fig. 3 shows the proper choice of anneal offsets renders the distributions indistinguishable for the locked and unlocked gadgets (Hamiltonian element encoding binary variables with configurations allowing more or fewer fluctuations). We show in Sec. II B that introduced fluctuations from anneal offsets can be as effective if not more so than fluctuations due to truly free variables.

The question now becomes whether anneal offsets can similarly mimic the effect of a lower softness coefficient for chains within the planted solution Hamiltonian. Figure 4 indicates that it cannot, while a negative anneal offset parameter, $\delta s < 0$, increases the number of soft chains at intermediate

FIG. 2. Fraction of observations for different numbers of soft chains against different values of reversal parameter s^* averaged over all 10 Hamiltonians (recall that s^* is a unitless quantity). Solid red (gray) lines with “X” markers represent the mean for the plot, while the dashed line with “+” markers is the mean of the other plot for comparison. (a) Minimum softness coefficient chains with no anneal offsets. (b) Maximum softness coefficient chains with no anneal offsets. The methods for creating the plots given the nonlinear mesh are explained in the Appendix.

values of s^* , it decreases the number at low s^* . Therefore no value can be used to mimic the behavior of a lower softness coefficient simultaneously in both regimes. This is likely due to the more complicated structure of the chain encoded dis-

FIG. 3. Fraction of observations for different numbers of free gadgets against different values of reversal parameter s^* averaged over all 10 Hamiltonians (recall that s^* is a unitless quantity). Solid red (gray) lines with “X” markers represent the mean for the plot, while the dashed line with “+” markers is the mean of the other plot for comparison. (a) Unlocked gadgets with no anneal offsets. (b) Locked gadgets with anneal offsets (δs) of up to -0.04 applied to the gadgets. The methods for creating the plots given the nonlinear mesh are explained in the Appendix.

FIG. 4. Fraction of observations for different numbers of soft chains against different values of reversal parameter s^* averaged over all 10 Hamiltonians (recall that s^* is a unitless quantity). Solid red (gray) lines with “X” markers represent the mean for the plot, while the dashed line with “+” markers is the mean of the other plot for comparison. (a) Chains with maximum softness coefficient and no anneal offsets. (b) Chains with maximum softness coefficient and offsets (δs) of up to -0.04 applied to the gadgets. The methods for creating the plots given the nonlinear mesh are explained in the Appendix

crete variables. We demonstrate in Sec. II B that in contrast to the locked versus unlocked gadget example, anneal offsets cannot make up the difference between the minimum and maximum values of this parameter.

A. Conditional performance

Simply analyzing solution optimality is a losing proposition, since we have designed the experiments such that, by construction, there is no way of improving beyond the starting condition. However, there is still hope to find high-quality solutions which meet conditions which the global solution does not. We define this as *conditional performance*, the best performance attainable which also meets certain conditions. Because of how the gadgets and chains have been constructed, the condition we have chosen to analyze is how many gadgets can be in the free configuration, or chains can be in a soft configuration. This is an interesting criteria since free gadgets and soft chains both make the solution more flexible, allowing for modifications which can be made with little or no energy cost. This flexibility could be important in real world scenarios, for instance, if small changes to the solution may need to be made after the time of solving to account for unpredictable events, or if the annealer is being used as part of a hybrid solving technique where difficult to encode global constraints are not included (for an example of the latter see Ref. [6]). In Sec. III we give an example where flexible solutions can be used to gain an advantage when an additional nonlinear constraint is added.

For a fair comparison, we should compare the results from the annealer with a trivial classical strategy of simply frustrat-

FIG. 5. Energy cost per free gadget for 10 different Hamiltonians using best performing value of s^* blue “+” markers are without anneal offsets, red “X” markers are best anneal offset (including the possibility of no offset). Red boxes and blue circles represent mean for without and with anneal offsets, respectively, with error bars representing standard error. Black dashed line is a guide to the eye at a cost of 2. Energy is in dimensionless coupling units, and s^* is a unitless quantity.

ing the couplings between the gadgets or chain and the rest of the problem, this “trivial” strategy leads to a cost per gadget or chain of two energy units compared to the most optimal solution. Solutions with a lower cost per gadget or chain are in principle interesting solutions, whereas those which have a higher energy than the trivial approach are not, since there is a know method which will always attain a better solution using the same starting information. Since the focus of this work is proof-of-concept rather than benchmarking, we will not explore whether or not there are other, less trivial, classical algorithms which can have better conditional performance than the annealer.

To start off, let us examine the conditional performance for the Hamiltonian with gadgets inserted without using anneal offsets. As Fig. 5 shows, even without anneal offsets the annealer is able to outperform a trivial algorithm in all but one case, in which the energy cost is more only if every gadget is made free. When different anneal offsets on the gadgets are allowed, the energy cost per free gadget never exceeds 1.5.

For discrete variables represented as domain walls [feature used to encode the information on the chains (see Appendix)], reverse annealing is also usually able to find a solution which beats the trivial approach, in fact Fig. 6 shows that even without using anneal offsets, the annealer was always able to find a solution which was better than the trivial approach when the soft region of the chain is flat [softness parameter (parameter controlling level of fluctuations within soft range) s of 0]. Even when the region of the chain which is being searched out is not flat, but a sloping minima (softness parameter s of 1), the annealer is able to beat the trivial approach in most cases, and always does both on average, and for all cases examined with less than 14 soft chains. The results for the higher softness parameter are depicted in Fig. 7. A full definition of s can be found in the Appendix.

I have now shown that reverse annealing in combination with anneal offsets can be effective at modifying solutions to meet certain conditions, but have not elucidated why or how this might happen, in the next subsection we examine potential

FIG. 6. Energy cost per soft chain for 10 different Hamiltonians using best performing value of s^* blue “+” markers are without anneal offsets, red “X” markers are best anneal offset (including the possibility of no offset). Red boxes and blue circles represent mean for without and with anneal offsets, respectively, with error bars representing standard error. Black dashed line is a guide to the eye at a cost of 2. Softness parameter used was $s = 0$ in both cases. Energy is in dimensionless coupling units, and s^* is a unitless quantity.

underlying mechanisms and discuss what the data can teach us about anneal offset strategies.

B. Performance with anneal offsets and locked gadgets

It is now worth examining more closely the role which quantum fluctuations play in conditional performance, by comparing Figs. 6 and 7 (averages directly compared in Fig. 8 (left)). We are able to see that better solutions are possible with a lower softness parameter s , the question not yet explicitly answered is whether the same is true for the fluctuations the free spins cause in the gadgets. To do this we need to compare the “free” and “locked” versions of the gadgets, where free spins are not possible, regardless of the configuration of external spins, which are described in detail in the Appendix. As Fig. 8 (right) shows, in the absence of anneal offsets having locked gadgets is very detrimental to performance, at least if more than about six gadgets are desired to be free. On the

FIG. 7. Energy cost per soft chain for 10 different Hamiltonians using best performing value of s^* black “+” markers are without anneal offsets, magenta “X” markers are best anneal offset (including the possibility of no offset). Magenta boxes and black circles represent mean for without and with anneal offsets, respectively, with error bars representing standard error. Black dashed line is a guide to the eye at a cost of 2. Softness parameter used was $s = 1$ in both cases. Energy is in dimensionless coupling units.

FIG. 8. Average energy cost per gadget (a) and chain (b) averaged over 10 different Hamiltonians. Magenta boxes and black circles represent mean of best performance without and with anneal offsets, respectively, and for locked gadgets (a) or softness parameter 1 (b). Red boxes and blue circles represent the same, but with unlocked gadgets (a) or softness parameter 0 (b). Error bars represent standard error. Black dashed line is a guide to the eye at a cost of 2. Energy is in dimensionless coupling units.

other hand there is barely any difference once anneal offsets are employed, suggesting that the offsets can enhance the fluctuations and guide the search. Conversely, the effect of anneal offsets seems to be rather minimal for discrete variables encoded in chains.

The first question to ask is what is the optimal value of s^* for given a desired number of free gadgets and soft chains, and how is this affected by factors like whether or not gadgets are locked and the softness parameter used for chains, as well as whether or not anneal offsets are used. Figure 9 shows the optimal value of s^* for both gadgets and chains under different circumstances. The first thing to notice from this figure is that, perhaps unsurprisingly, s^* decreases monotonically (within statistical uncertainty) with the desired number of free gadgets

FIG. 9. Average value of s^* (averaged over 10 Hamiltonians) used to obtain optimal conditional performance with a desired number of free gadgets (a) or soft chains (b). Red and magenta squares represent cases where no anneal offsets are used and are unlocked (softness parameter 0) and locked (softness parameter 1), respectively. Blue and black circles represent the cases where anneal offsets are used and are unlocked (softness parameter 0) and locked (softness parameter 1), respectively. Error bars represent standard error. In all cases the largest value of s^* was taken in the event of a tie. Insets are the same plots but zoomed out. Energy is in dimensionless coupling units.

FIG. 10. Average anneal offset δs taken for 10 Hamiltonians. Blue [red (lighter gray)] triangles represent Hamiltonians with unlocked (locked) gadgets. While black [magenta (darker gray)] stars represent Hamiltonians with chains with a softness parameter of 0 (1). Error bars represent standard error. Note that positive offsets indicate that fluctuations are suppressed, relative to the rest whereas negative indicates that they are enhanced. In the event of a tie, the lowest numerical value of the offsets which gave the tying energy were taken. Energy is in dimensionless coupling units.

or soft chains, this behavior makes intuitive sense, because changing more variables requires a broader search. Furthermore, constant with Fig. 8, the values of s^* based on whether or not anneal offsets are used differ much more for gadgets than for chains, indicating that allowing anneal offsets greatly changes the optimal strategy for gadgets, and does not change it as much for chains. Furthermore, except for when about 14 or more free gadgets are desired, the optimal value of s^* when anneal offsets are used is almost the same for locked and unlocked gadgets, supporting the hypothesis that increased fluctuations from anneal offsets can act as an effective proxy for truly free variables.

To better understand the role anneal offsets are playing, it is worth examining how the best choice of anneal offset depends on the number of free gadgets, or soft chains desired. As Fig. 10 shows, the best strategy is indeed to use stronger offsets in the locked gadget case, and to use them to enhance rather than suppress fluctuations on the gadgets, suggesting that there is indeed a mechanism where offsets artificially guide the search by making the locked gadgets behave as if they have free qubits.

Figure 10 further shows that domain wall encoded discrete variables show very different behavior to the gadgets, in particular, up to statistical uncertainty, the offsets used in the discrete variable case monotonically approaches zero as more soft chains are desired, while for gadgets with free binary variables, there is nonmonotonic behavior, and a trend toward locally enhancing fluctuations if more free gadgets are desired. This difference is likely due to the more complex structure of the domain wall encoded variables, leading to less tolerance to fluctuations before they no longer faithfully encode the intended variable.

III. MOTIVATIONAL EXAMPLE FOR FLEXIBLE SOLUTIONS

Now that it has been shown that the underlying dynamics of quantum annealers can be used to find solutions which are

more flexible, it is worth demonstrating an example where such solutions could be useful. To do this, let us consider a problem which natively fits onto the chimera graph, but is also subject to global nonlinear penalty. Such global penalties are likely to be encountered in realistic problems, and for example may arise when a shared resource is being used for different purposes and there is a penalty which depends on the *total* amount required. A simple example of how such a constraint could arise in the real world is minimizing the total cost of a project if a company owns X number of a piece of equipment, so there is no penalty for a solution which uses any number up to X ; however, there is a cost associated with renting every additional piece of equipment beyond the original X .

While techniques are known to implement global nonlinear penalties on quantum annealers, for example, those proposed in Refs. [46,47], these techniques require a fully connected graph and number of auxiliary qubits equal to the number of original qubits, and such an encoding is not practical for large problems on existing quantum annealers. Consider first an alternative strategy for solving such problems: first, encode the entire problem except for the global penalty onto the annealer and use reverse annealing techniques to find solutions with various levels of trade-off between flexibility (for example, measured by the number of free gadgets) and optimality. We then perform greedy optimization as described in the Methods section starting from the best solution found at each level of flexibility. This greedy optimization is performed against the entire problem including the nonlinear penalty.

Before considering the results for the QPU-sized problems used in earlier demonstrations, it is worth demonstrating this approach with a simpler 16-qubit example. To do this, consider the Hamiltonian used in Ref. [29], which is in turn similar to the Hamiltonian considered in Ref. [48]. This Hamiltonian has both a local minimum where eight of the 16 qubits are “free,” able to exist in either the zero or one state without incurring an energy penalty, and a global minimum where none of the qubits are free; the (unique) ground state and first excited state manifold of this Hamiltonian are depicted in Fig. 11. At least for short run times, the close avoided crossings in these devices mean that quantum annealers will typically find the false minimum with more free qubits due to a close avoided crossing relatively late in the annealing schedule [29].

We now consider the ability of the solution to adjust to nonlinear penalties of different strength. The global nonlinear penalty we elect to use is nonlinear function of the Hamming distance $\mathcal{D}$ from a random state

$$E(\mathcal{D}) = 1 - \exp \left\{ - \frac{[\mathcal{D} - (\frac{n}{2} + \sqrt{n+1})]^2}{n+1} \right\}, \quad (1)$$

where n is the number of qubits involved in the Hamiltonian. The states which the annealer returns will be a Hamming distance $\mathcal{D} = \frac{n}{2}$ away from most random states, therefore this penalty offsets the Gaussian from the point where a typical solution will sit by its standard deviation, $\sqrt{n+1}$. This will guarantee that the nonlinear penalty will have a substantial gradient for typical starting states.

Equipped with this definition let us consider the results of adding a nonlinear penalty followed by a greedy search for

FIG. 11. The 16-qubit gadget used in Ref. [29]. Edges represent ferromagnetic coupling of unit strength, and circles represent qubits. Red (and dotted borders) indicates that a qubit is subject to a field of $+1$, while magenta (and dashed borders) indicates -1 and gray (solid borders) indicates no field. In the top figure diagram the arrows indicate the unique ground state which satisfies $+1$ fields on the outer qubits but frustrates the -1 field. The bottom diagram is the first excited manifold, where a superimposed 0 and 1 indicates that a qubit is “free” and can take either value without affecting the energy.

the 16-qubit problem mentioned earlier. As Fig. 12 shows, it is much easier for the greedy search heuristic to compensate for the global nonlinear penalty starting from the higher energy but more flexible solution which the annealer finds as compared to the true minimum; the result is that for moderate penalty strength, the more flexible state is a superior choice for a starting configuration.

FIG. 12. Energy in dimensionless coupling units after applying a nonlinear penalty with a strength given by the x axis and performing greedy search. The gold (lighter gray) line shows results for starting from the true minimum, while the blue (darker gray) line shows the result starting from the higher energy, but more flexible false minimum the annealer typically finds. The dashed red line is the energy of the true lowest energy state. Ten thousand samples were taken for each point on this plot, and statistical error bars are smaller than the depicted lines.

FIG. 13. Energy in dimensionless coupling units of the best solution (where planted solution energy is defined to be zero) with a different number of free gadgets versus nonlinear penalty strength. The color encoding of the number of gadgets is depicted in the inset. This plot is for Hamiltonian number 7 and for the best solution found including the use of anneal offsets, although it is typical of the behavior seen in both cases. The green (dark gray) line is included as a visual aid and follows the state with zero free gadgets. These data were averaged over 300 choices of random states, and in cases where multiple states were tied for the lowest energy for a given number of free gadgets, a new state was chosen at random for each sample.

A. Synthetic use case: Optimizing with global nonlinear penalties

We now consider what happens when we apply a nonlinear penalty followed by greedy search to states with different numbers of free gadgets found for QPU scale problems. While neither the original problem nor the nonlinear penalty is based on anything which one might encounter in the real world, recall that situations where a problem containing a nonlinear penalty must be solved are realistic, and this can therefore be considered a “synthetic” use case for a quantum annealer, not directly based on an application, but with a structure which is likely to be encountered in the real world. We start by considering the best solutions the annealer could find with different free gadget numbers for a single Hamiltonian, in this case Hamiltonian number 7. As Fig. 13 shows, as the penalty strength is increased to a moderate value, the best solution is no longer obtained from starting a greedy search at the true energy minimum, but from starting with a more flexible state with more free gadgets. For these experiments we consider only the best solution found with each number of free gadgets, choosing at random in the event of a tie. All greedy searches are performed on the unlocked gadgets, where having free variables is likely to improve the solution quality when the nonlinear penalty is added.

From Fig. 14 we can see that the behavior seen in Fig. 13 is indeed typical of results found both with and without anneal offsets, although, unsurprisingly, the cases where anneal offsets are used perform better on average since lower energy solutions can be found by using anneal offsets.

Finally, consider the optimal number of free gadgets in the starting state for different Hamiltonians and penalty strengths. Figure 15 shows that, for both the strategy using anneal offsets and the one which does not, the typical number of free gadgets in the best performing state increases for a while with penalty

FIG. 14. Energy difference in dimensionless coupling units between greedy search performed with a nonlinear penalty starting in planted solution (with no free gadgets) and the best performing state found via reverse annealing. Panel (a) considers only solutions found without using anneal offsets, while panel (b) is the same but including offsets. The colored lines represent the average over all 10 Hamiltonians, and the gray squares represent individual Hamiltonians. Panel (c) shows only the averages, with the blue (dark gray) circles representing the method inducing offsets and red (lighter gray than the circles but still darker than background) squares without. These data were averaged over 300 choices of random states, and in cases where multiple states were tied for the lowest energy for a given number of free gadgets, a new state was chosen at random for each sample. Penalty strength is in dimensionless coupling units.

strength and then settles to an average across all Hamiltonians of around seven free gadgets. While it is possible that the average number of free gadgets is slightly higher for the strategy using offsets, the difference is relatively small. It is however clear that for the solutions which used anneal offsets, there is a much wider variety of solutions and, in particular, a tendency to use some solutions with many more free gadgets.

Recall that the gadgets and chains were observed to behave qualitatively differently in the previous section, so the results from gadgets may or may not carry over to chains. The purpose of this section is to provide proof-of-concept for the usefulness of fluctuation guided behaviors, not to provide exhaustive evidence of how they can be used, so we will not analyze these in detail. The data associated with this paper however are publicly available [49], and analysis of the chain data from the perspective of hybrid algorithms, similar to what we have done here for gadgets, is likely to produce interesting results.

IV. METHODS

All reverse annealing experiments were performed using the maximum allowed annealing rate on both the forward and reverse anneal, and at this rate the entire (forward) anneal would be completed in $5\mu\text{s}$. All experiments used a hold time τ of $20\mu\text{s}$. All annealer calls were set to perform 1000 individual runs. The reverse annealing experiments presented here were performed using the D-Wave Matlab API between 27 October 2018 and 30 October 2018 on a commercially

FIG. 15. Top: number of free gadgets in optimal solution for all 10 Hamiltonians versus penalty strength, shading indicates number of Hamiltonians where the same number of free gadgets are optimal for the same penalty strength; a legend for the shading levels appears in the middle of the figure. The colored lines are the mean. Panel (a) is the best solutions not including anneal offsets, and (b) is including them. (c) The mean from the top two plots shown on the same axis to compare them. Penalty strength is in dimensionless coupling units.

available D-Wave 2000Q QPU with QPU time purchased by BP plc. Data are publicly available at Ref. [49].

Greedy optimization was performed by checking all single bit flips and performing the one which reduces the energy the most, choosing at random in the event of a tie. The greedy procedure is repeated until no single bit flip will reduce the energy.

All plots were produced in the Python language [50] and the matplotlib plotting package [51] code used to produce the plots and perform the experiments is available from the same public repository as the experimental data. Heat-map plots with nonlinear grids were plotted such that the center of each cell aligns with the value of each axis. The NumPy [52] and SciPy [53] packages were also used as well as jupyter notebooks [54] and the IPython interpreter [55].

V. DISCUSSION AND CONCLUSIONS

In this paper, we have demonstrated how fluctuations can guide quantum annealers to trade off optimality for more flexible solutions, as well as motivated cases where such a tradeoff could be useful. The particular useful case we focus on is when a problem involved global penalties which cannot practically be implemented on the annealer. While in the past the tendency of quantum annealers to find solutions where fluctuations are stronger has been seen as a weakness, for instance, in inhibiting the ability to uniformly sample ground states, we demonstrate ways in which it could be useful.

In addition to demonstrating that the existing fluctuations on the annealer can help guide searches toward more flexible states, we show that locally offsetting the annealing schedule of the qubits can be used to guide the search. This provides experimental motivation for methods like those proposed in

Ref. [28], which incorporate bitwise uncertainty into algorithms.

Realistic problems may have chain like structures if the encode discrete variables, but will not contain gadgets like the ones discussed here. Even for problems which encode discrete variables, it will not necessarily be obvious which chains should have the offsets applied. Identifying areas where there is more potential for flexibility to be increased using anneal offsets in the setting of either binary or higher-than-binary discrete variables is an interesting problem, but one which is beyond the scope of the current work. Techniques could include experimental analysis to see where soft chains or free qubit variables appear and offset those parts of the problem Hamiltonian, but testing would need to be done to determine whether this approach is useful or practical.

While not explored here, it is likely that analogous effects could be seen in quantum inspired algorithms based on spin-like systems, for example, quantum Monte Carlo techniques [56] which should show analogous effects to the fluctuations observed here. In fact, the proof-of-concept numerics in Ref. [28] exhibited that fluctuations can attract quantum Monte Carlo dynamics preferentially to some minima over others. This work has introduced alternative ways in which quantum annealers and related algorithms can be used, beyond directly finding the most optimal solution, an important direction in hybrid quantum-classical computing. By laying the groundwork for how modifying fluctuations locally can be used algorithmically to guide a search, the work here opens a path to using these modified fluctuation strengths algorithmically, in a similar vein to currently used reverse annealing techniques, but guiding the direction of the search, rather than the starting point.

ACKNOWLEDGMENTS

This work was supported by BP plc, by EPSRC fellowship EP/S00114X/1, EPSRC grant EP/L022303/1, and EPSRC Hub EP/M013243/1. The author thanks Adam Callison for pointing out several typos in the original arXiv version, as well as Andrew King for providing key references and perspective for the discussion of previous work.

APPENDIX: EXPERIMENTAL SETUP

These experiments involve both specially engineered Hamiltonians to construct a search space with the necessary properties and the use of advanced control features of the QPU, both anneal offsets and reverse annealing, which are used in combination. We first describe how the Hamiltonians are constructed, and then how they are used in the actual experimental protocols. Before we do this, it is useful to provide some background on the operation of the quantum annealer. In the simplest implementation of its protocol, QPU realizes a transverse field Ising Hamiltonian

$$H = -A(s) \sum_i X_i + B(s) H_{\text{prob}}, \quad (\text{A1})$$

where $A(s)$ and $B(s)$ are functions of a control parameter $0 \leq s \leq 1$, which are nonlinear in general but could also be linear in principle, X_i is a Pauli X acting on qubit i , and H_{prob} is a

programmable Ising problem Hamiltonian,

$$H_{\text{prob}} = \sum_{ij} J_{ij} Z_i Z_j + \sum_i h_i Z_i, \quad (\text{A2})$$

where Z_i is a Pauli Z acting on qubit i ; the details of how J_{ij} and h_i are chosen is discussed later.

The QPU is designed so that $|\frac{A(0)}{B(0)}| \gg 1$ and $|\frac{A(1)}{B(1)}| \ll 1$, the ratio $|\frac{A(s)}{B(s)}|$ decreases monotonically with s , and neither A nor B change sign during the protocol. We do employ a more advanced feature known as anneal offsets, which slightly changes the form of Eq. (A1) and will be discussed in due course.

1. Hamiltonian construction

The goal of the experiments in this paper are to study the ability of a quantum annealer to use fluctuations to find high-quality solutions which are flexible in the sense that changing some elements of the solution will not affect the energy of the solution or will affect it only very little. Since we are not developing this study as a benchmark against classical methods, we have focused on designing Hamiltonians which are difficult to solve for the annealer and have a known solution, but which are not necessarily computationally hard problems. To this end, the problems used here build on the planted solution construction from Ref. [57], which yields limited computational hardness [58,59] (for state-of-the-art solution planting techniques, see Ref. [60]). Furthermore, we use many more clauses than would be desirable to construct the hardest problems in the interest of ensuring that the problem graph is connected and to reduce the degeneracy of the ground state manifold.

The methods which we use, proposed in Ref. [57], construct problems with planted solutions by generating overlapped frustrated loops on the edges of the underlying graph via random walks which terminate when they intersect their own path. We use planted solution problems with loop size less than six and 8000 loops on a QPU with approximately 2000 qubits (some of which are reserved for specialized features as discussed later in this section), with a coupling arrangement known as a chimera graph. The details of this coupling graph are not important to understand this study, but a full description of the graph can be found in Ref. [61]. Figures 17 and 18 depict chimera graphs with a 3×3 grid of eight qubit unit cells, and the 2000Q has the same eight qubit unit cells arranged in a 16×16 grid.

In addition to having a known planted solution, the experimental Hamiltonians also need features which can explore the ability of the annealer to use fluctuations to find more flexible solutions. Since we intend to study the ability to find more flexible solutions in both a binary and discrete setting, in other words both in the setting where a variable can take two values and in the setting where it still takes a finite number of discrete values, but can take more than two, two different strategies need to be employed: gadgets where variables are allowed to become “free” should be embedded, henceforth referred to as “gadgets,” as well as chains of qubits which encode discrete variables using the domain wall encoding described in Ref. [62], henceforth referred to as “chains.”

FIG. 16. A free spin gadget inserted into a planted solution Hamiltonian. Green circles (dark gray) and dashed edges indicate qubits which are not part of the gadget, while black qubits indicate the active qubits within the gadget, with black edges indicating ferromagnetic couplers of unit strength and red edges indicating antiferromagnetic couplers of unit strength. Gray circles and edges indicate the unused couplers and qubits within the gadget. Pink (gray) is used as a guide to the eye.

Fortunately, the planted solution construction does not require a full chimera graph to be effective. This means that the construction can be performed with some qubits reserved for either gadgets or chains, and these features can be added in later. The gadgets are constructed with the following properties:

- (1) Couplings to neighboring qubits within the planted solution construction
- (2) A unique ground state when all external qubits which the gadget couples to are in the $|0\rangle$ (or $|1\rangle$) configuration
- (3) Degenerate ground state with free variables in cases where the external couplings do not agree, and therefore the planted solution components may be frustrated, and these have equal energy to the unique state
- (4) Occupy a single chimera unit cell

Figure 16 depicts a gadget which obeys these properties embedded into a larger problem Hamiltonian. The top row of Fig. 17 depicts the lowest energy states of this gadget when external qubits either agree or disagree. To be able to separate the effects of fluctuations due to free spins from other effects, we have also developed a “locked” version of the free variable gadget; in this version some of the couplings are reduced to half the strength of the others so that the lowest energy state when the external variables do not agree no longer contains free variables; these are depicted on the bottom row of Fig. 17.

To embed discrete variables, we use the domain wall encoding from Ref. [62] to encode a variable with 16 possible values within a 15-qubit chain with unit ferromagnetic coupling strength. For completeness, we review the domain wall encoding at the end of the Appendix. We use the field controls of the annealer to control the potential on this chain such

FIG. 17. Top: Minimum energy state of free variable gadget with different configurations of external variables; left is where all external variables agree, right is where one disagrees. Bottom: Same but for locked gadget. Gray edges and circles indicate unused couplers and qubits, green indicates (dark gray) external qubits, red (dark gray, between pairs of dark nodes), and black edges indicate anti-ferromagnetic and ferromagnetic coupling, respectively, while thick edges indicate coupling of unit strength and thin indicate coupling with a strength of 0.5. Superimposed 1 and 0 characters indicate free variables. Slashes indicate frustrated couplings, with gray slashes indicating multiple possibilities depending on the values of the free variables.

that the 0 value of the variable (all qubits are in the $|0\rangle$ configuration) has the same energy as the minimum energy in a “soft” region which corresponds to seven consecutive values of the discrete variable which are randomly chosen to start anywhere from two to six (recall that we use a convention where the allowed values run from 0 to 15). The chain is coupled to the rest of the problem Hamiltonian on the first and last qubit of the soft region, such that the planted solution must be frustrated if the domain wall is in the soft region. All other values of this variable have an energy which is two energy units higher than either the minimum of the soft range (range of values where more fluctuations are allowed) or the 0 state of the variable. Henceforth we refer to a chain where the domain wall is in the soft region as a “soft chain.” The qubit chain used to encode the domain wall variable is randomly placed within the planted solution problem by performing 15 steps of a non-self-intersecting random walk on the hardware graph; an example of a chain within the larger Hamiltonian is depicted in Fig. 18.

The potential within the soft range is always equal to the 0 value of the variable at the midpoint m of the range. Away from the midpoint the potential increases such that $E(m+j) = E(m) + \varsigma|j/2|$. Lower values of ς allow for more fluctuations since it costs less energy for the domain wall to move away from the center of the chain.

For both the gadget and chain versions of the problem, 10 Hamiltonians were created at random. Other than the reduced strength of the gadget couplings, the free and locked gadget

FIG. 18. Chain encoding domain wall variable placed within a planted solution problem. Black circles indicate qubits which are part of the domain wall encoding but not within the soft range, blue (darker gray) is a guide to the eye to indicate the soft range and the externally coupled qubits at each end of the range. Green (lighter gray) indicates those outside of the domain wall variable. Solid edges indicate ferromagnetic coupling of unit strength, while dashed edges indicate couplings which are part of the planted solution encoding. Pink (gray) is used as a guide to the eye.

runs use the same 10 Hamiltonians. Each Hamiltonian includes either 15 gadgets or 15 chains.

2. Annealing protocol

The key feature of the Hamiltonians constructed for these experiments is that they have known planted solutions. This is crucial for the purpose of this study: to explore the ability of the device to trade off between optimality and flexibility, we need to start off in a state which is known to be optimal. Fortunately the reverse annealing feature [10] allows for a search around the planted (or any other classical) state. The reverse annealing feature uses a protocol which starts the QPU in a state determined by the user at $s = 1$, anneals to a value s^* held for a time τ , and then anneals back to $s = 1$ as depicted in Fig. 19. Thermal dissipation allows the device to seek out lower energy states during the reverse annealing protocol.

In addition to the reverse annealing feature, we also make use of another feature called anneal offsets [11]. The function of this feature is to offset the annealing parameter on different qubits. In particular, we offset the parameter values of either the chains or the gadgets (a subset of qubits we call g), which makes the Hamiltonian

$$H = -A(s) \sum_{i \notin g} X_i + B(s) \left(\sum_{i \notin g, j \notin g} J_{ij} Z_i Z_j + \sum_{i \notin g} h_i \right) - A(s + \delta s) \sum_{i \in g} X_i + B(s + \delta s) \times \left(\sum_{i \in g, j \in g, j \notin g} J_{ij} Z_i Z_j + \sum_{i \in g} h_i \right), \quad (\text{A3})$$

FIG. 19. Schematic of reverse annealing protocol; solid line is the value of s for qubits not in gadgets or chains (or all qubits in no offset case). Dot-dashed line is the schedule for chains and gadgets with positive δs , the dashed line is the same for negative δs . The quantity s is unitless, and time is typically in units of μs .

where $i \in g$ means that qubit i belongs to a gadget or chain and $i \notin g$ means that it does not. Effectively, for a positive value of δs the strength of the couplings and longitudinal field terms within g as well as coupling between g and the rest of the qubits is increased, while the transverse field within g is decreased. For negative g the opposite happens. In simplified terms, positive δs leads to weaker fluctuations with g , while negative δs leads to stronger. Since the original purpose of the anneal offset purpose was to suppress fluctuations, previous work has focused most strongly on positive δs .

For all experiments reported here, the anneals in the reverse annealing protocol were performed at the maximum allowed rate, which traverses from $s = 0$ to $s = 1$ in $5 \mu s$ and a hold time τ of $20 \mu s$ was used. As a comparison point, the energy scale of these devices leads to natural frequencies in the GHz range. The same parameters were used for chains and gadgets. For all values of s^* , we used a linearly spaced grid of 11 values of δs evenly spaced between -0.2 to 0.2 , inclusive of the end points. Since not all qubits are capable of the full range of offset values, the maximum magnitude allowed (positive or negative) value was used when the desired value fell outside of the range. Because we wanted to study extreme values of s^* as well as studying more values within a region of interest where the data were observed to change rapidly with s^* , we chose the nonuniform grid of 19 values of s^* depicted in Fig. 20.

3. Review of the domain wall variable encoding strategy

The domain wall encoding strategy used here was originally developed to undertake the research described in this paper, since it can encode discrete variables on a chimera graph without requiring minor embedding unlike the more traditional one-hot strategy. Because the domain wall encoding strategy has been observed to significantly outperform the one-hot strategy on several key metrics related to embedding on realistic hardware graphs, a full description of this technique as well as numerical evidence of its superior performance has been published elsewhere [62], and an experimental study of its comparative performance to one-hot is

FIG. 20. Grid of values of s^* used in this study; dotted lines are guides to the eyes to show the region of higher interest, between about $s^* = 0.4$ and $s^* = 0.5$. Both quantities in this plot are unitless.

forthcoming [63]. This technique has also been used in solving quantum field theory problems using quantum annealers [8].

In the interest of making the current paper self-contained, we review the basics of the domain wall encoding and some of its key features. As Fig. 21 (top), shows, the domain wall encoding is produced by creating a linearly connected chain of $n - 1$ qubit with frustrating fields on each end such that there are n total possible domain wall locations, including frustrating the the fields at either end, which can be thought of as couplings between the terminal qubits and “virtual” qubits which are constrained to take either the 1 or 0 value. As was discussed in detail in Ref. [62], any interaction between two discrete variables can be realized using two body couplings

FIG. 21. Top: Encoding of a discrete variable as a domain wall position, where the domain wall is depicted in blue, real qubits in green, and “virtual” qubits which are fixed are depicted in pink with dotted borders. Bottom: A binary variable in the domain wall encoding reduces to the standard qubit representation.

between the qubits in the domain wall encoding, and arbitrary penalties can be realized by putting fields (single body terms) on the chain. Moreover, the domain wall encoding of a binary variable simply reduced to a normal qubit representation, as depicted in Fig. 21 (bottom).

I am interested in simple couplings which force frustration in the planted solution problem if the domain wall variable takes one of its soft values, while simultaneously avoiding the need for minor embedding. To do this, we place a single ferromagnetic coupler between the qubits encoding the discrete variables and the other qubits at each end of the soft region. For the additional energy penalties on the chain, we make use of the fact that a single (nonextreme) value of the discrete variable can be penalized using a term of the form $\delta_i = \frac{1}{2}(Z_i - Z_{i-1})$. For the extreme values can be penalized in the same way, but omitting terms which correspond to virtual qubits. This method is described in more detail in Ref. [62], and software for realizing these encodings can be found at Ref. [64].

- [1] T. Stollenwerk, B. O’Gorman, D. Venturelli, S. Mandrà, O. Rodionova, H. Ng, B. Sridhar, E. G. Rieffel, and R. Biswas, Quantum annealing applied to de-conflicting optimal trajectories for air traffic management, *IEEE Trans. Intell. Transp. Syst.* **21**, 285 (2019).
- [2] D. O’Malley, An approach to quantum-computational hydrologic inverse analysis, *Sci. Rep.* **8**, 6919 (2018).
- [3] A. Perdomo-Ortiz, N. Dickson, M. Drew-Brook, G. Rose, and A. Aspuru-Guzik, Finding low-energy conformations of lattice protein models by quantum annealing, *Sci. Rep.* **2**, 571 (2012).
- [4] T. Stollenwerk, E. Lobe, and M. Jung, Flight gate assignment with a quantum annealer, in *Quantum Technology and Optimization Problems*, edited by S. Feld and C. Linnhoff-Popien (Springer, Cham, 2019), pp. 99–110.
- [5] M. Marzec, Portfolio optimization: Applications in quantum computing, in *Handbook of High-Frequency Trading and Modeling in Finance*, edited by I. Florescu, M. C. Mariani, H. E. Stanley, and F. G. Viens (John Wiley & Sons, 2016), Chap. 4, pp. 73–106.
- [6] D. Venturelli and A. Kondratyev, Reverse quantum annealing approach to portfolio optimization problems, [arXiv:1810.08584](https://arxiv.org/abs/1810.08584).
- [7] R. Orus, S. Mugel, and E. Lizaso, Forecasting financial crashes with quantum computing, *Phys. Rev. A* **99**, 060301 (2019).
- [8] S. Abel, N. Chancellor, and M. Spannowsky, Quantum computing for quantum tunnelling, [arXiv:2003.07374](https://arxiv.org/abs/2003.07374).
- [9] S. Abel and M. Spannowsky, Observing the fate of the false vacuum with a quantum laboratory, [arXiv:2006.06003](https://arxiv.org/abs/2006.06003).
- [10] Reverse quantum annealing for local refinement of solutions, https://www.dwavesys.com/sites/default/files/14-1018A-A_Reverse_Quantum_Annealing_for_Local_Refinement_of_Solutions.pdf (2019).
- [11] E. Andriyash, Z. Bian, F. Chudak, M. Drew-Brook, A. D. King, W. G. Macready, and A. Roy, Boosting integer factoring performance via quantum annealing offsets, Technical Report 14-1002A-B, D-Wave Systems Inc, Burnaby BC, Canada, 2016, https://www.dwavesys.com/sites/default/files/14-1002A_B_tr_Boosting_integer_factorization_via_quantum_annealing_offsets.pdf.
- [12] N. Chancellor, Modernizing quantum annealing using local searches, *New J. Phys.* **19**, 023024 (2017).
- [13] D. Ottaviani and A. Amendola, Low rank non-negative matrix factorization with D-wave 2000Q, [arXiv:1808.08721](https://arxiv.org/abs/1808.08721).
- [14] J. Golden and D. O’Malley, Reverse annealing for nonnegative/binary matrix factorization, [arXiv:2007.05565](https://arxiv.org/abs/2007.05565).
- [15] J. King, M. Mohseni, W. Bernoudy, A. Fréchette, H. Sadeghi, S. V. Isakov, H. Neven, and M. H. Amin, Quantum-assisted genetic algorithm, [arXiv:1907.00707](https://arxiv.org/abs/1907.00707).

- [16] A. D. King *et al.*, Observation of topological phenomena in a programmable lattice of 1, 800 qubits, *Nature (London)* **560**, 456 (2018).
- [17] T. Lanting, Andrew D. King, B. Evert, and E. Hoskinson, Experimental demonstration of perturbative anticrossing mitigation using nonuniform driver Hamiltonians, *Phys. Rev. A* **96**, 042322 (2017).
- [18] N. G. Dickson and M. H. Amin, Algorithmic approach to adiabatic quantum optimization, *Phys. Rev. A* **85**, 032303 (2012).
- [19] A. D. King, E. Hoskinson, T. Lanting, E. Andriyash, and M. H. Amin, Degeneracy, degree, and heavy tails in quantum annealing, *Phys. Rev. A* **93**, 052320 (2016).
- [20] D. S. Steiger, T. F. Rønnow, and M. Troyer, Heavy Tails in the Distribution of Time to Solution for Classical and Quantum Annealing, *Phys. Rev. Lett.* **115**, 230501 (2015).
- [21] Y. Matsuda, H. Nishimori, and H. G. Katzgraber, Quantum annealing for problems with ground-state degeneracy, *J. Phys.: Conf. Ser.* **143**, 012003 (2009).
- [22] S. Mandrà, Z. Zhu, and Helmut G. Katzgraber, Exponentially Biased Ground-State Sampling of Quantum Annealing Machines with Transverse-Field Driving Hamiltonians, *Phys. Rev. Lett.* **118**, 070502 (2017).
- [23] M. S. Könz, G. Mazzola, A. J. Ochoa, H. G. Katzgraber, and M. Troyer, Uncertain fate of fair sampling in quantum annealing, *Phys. Rev. A* **100**, 030303(R) (2019).
- [24] B. H. Zhang, G. Wagenbreth, V. Martin-Mayor, and I. Hen, Advantages of unfair quantum ground-state sampling, *Sci. Rep.* **7**, 1044 (2017).
- [25] I. Hen and A. P. Young, Solving the graph-isomorphism problem with a quantum annealer, *Phys. Rev. A* **86**, 042310 (2012).
- [26] W. Vinci, K. Markström, S. Boixo, A. Roy, F. M. Spedalieri, P. A. Warburton, and S. Severini, Hearing the shape of the Ising model with a programmable superconducting-flux annealer, *Sci. Rep.* **4**, 5703 (2014).
- [27] Z. G. Izquierdo, R. Zhou, K. Markström, and I. Hen, Discriminating non-isomorphic graphs with an experimental quantum annealer, *Phys. Rev. A* **102**, 032622 (2020).
- [28] N. Chancellor, Modernizing quantum annealing II: Genetic algorithms with the inference primitive formalism, *arXiv:1609.05875*.
- [29] N. G. Dickson *et al.*, Thermally assisted quantum annealing of a 16-qubit problem, *Nat. Commun.* **4**, 1903 (2013).
- [30] S. Ashhab, J. R. Johansson, and F. Nori, Decoherence in a scalable adiabatic quantum computer, *Phys. Rev. A* **74**, 052330 (2006).
- [31] E. J. Crosson and D. A. Lidar, Prospects for quantum enhancement with diabatic quantum annealing, *arXiv:quant-ph/2008.09913*.
- [32] D-Wave Systems Inc., Improved coherence leads to gains in quantum annealing performance, https://www.dwavesys.com/sites/default/files/14-1037A-A_Improved_coherence_leads_to_gains_QA_performance.pdf (2019).
- [33] N. Chancellor and V. Kendon, Search range in experimental quantum annealing, *arXiv:2008.11054*.
- [34] A. Perdomo-Ortiz, S. E. Venegas-Andraca, and A. Aspuru-Guzik, A study of heuristic guesses for adiabatic quantum computation, *Quant. Inform. Proc.* **10**, 33 (2011).
- [35] Q.-H. Duan, S. Zhang, W. Wu, and P.-X. Chen, An alternative approach to construct the initial Hamiltonian of the adiabatic quantum computation, *Chin. Phys. Lett.* **30**, 010302 (2013).
- [36] T. Graß, Quantum Annealing with Longitudinal Bias Fields, *Phys. Rev. Lett.* **123**, 120501 (2019).
- [37] D-Wave Systems Inc., Probing mid-band and broad-band noise in lower-noise D-wave 2000Q fabrication stacks, https://www.dwavesys.com/sites/default/files/14-1034A-A_Probing_noise_in_LN_2000Q_fabrication_stacks.pdf (2019).
- [38] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, Probing many-body dynamics on a 51-atom quantum simulator, *Nature (London)* **551**, 579 (2017).
- [39] S. Jansen, M.-B. Ruskai, and R. Seiler, Bounds for the adiabatic approximation with applications to quantum computation, *J. Math. Phys.* **48**, 102111 (2007).
- [40] A. Callison, N. Chancellor, F. Mintert, and V. Kendon, Finding spin glass ground states using quantum walks, *New J. Phys.* **21**, 123022 (2019).
- [41] M. B. Hastings, Duality in quantum quenches and classical approximation algorithms: Pretty good or very bad, *Quantum* **3**, 201 (2019).
- [42] A. Callison, M. Festenstein, J. Chen, L. Nita, V. Kendon, and N. Chancellor, An energetic perspective on rapid quenches in quantum annealing, *arXiv:2007.11599*.
- [43] E. Vincent and V. Dupuis, Spin glasses: Experimental signatures and salient outcomes, in *Frustrated Materials and Ferroic Glasses*, edited by T. Lookman and X. Ren, Springer Series in Material Science Vol. 275 (Springer Nature, Cham, 2018), pp 31–56.
- [44] T. Vojta, Quantum Griffiths effects and smeared phase transitions in metals: Theory and experiment, *J. Low Temp. Phys.* **161**, 299 (2010).
- [45] K. Nishimura, H. Nishimori, and H. G. Katzgraber, Griffiths-McCoy singularity on the diluted chimera graph: Monte Carlo simulations and experiments on quantum hardware, *Phys. Rev. A* **102**, 042403 (2020).
- [46] N. Chancellor, S. Zohren, P. A. Warburton, S. Benjamin, and S. Roberts, A direct mapping of Max k-SAT and high order parity checks to a chimera graph, *Sci. Rep.* **6**, 37107 (2016).
- [47] N. Chancellor, S. Zohren, and P. A. Warburton, Circuit design for multi-body interactions in superconducting quantum annealing systems with applications to a scalable architecture, *npj Quant. Inform.* **3**, 21 (2017).
- [48] S. Boixo, T. Albash, F. M. Spedalieri, N. Chancellor, and D. A. Lidar, Experimental signature of programmable quantum annealing, *Nat. Commun.* **4**, 2067 (2013).
- [49] N. Chancellor, Fluctuation guided search in quantum annealing (data set), data archive at Durham University, doi:10.15128/r1c534fn95w (2019).
- [50] G. Van Rossum and F. L. Drake, Python language reference manual, Network Theory United Kingdom (2003).
- [51] J. D. Hunter, Matplotlib: A 2D graphics environment, *Comput. Sci. Eng.* **9**, 90 (2007).
- [52] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith *et al.*, Array programming with NumPy, *Nature* **585**, 357 (2020).

- [53] Scipy 0.17.1, 2016, <https://www.scipy.org/> (2016).
- [54] T. Kluyver, B. Ragan-Kelley, F. Pérez, B. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. Hamrick, J. Grout, S. Corlay *et al.*, Jupyter notebooks—A publishing format for reproducible computational workflows, in *Positioning and Power in Academic Publishing: Players, Agents and Agendas* (IOS Press, Southampton, UK, 2016), pp. 87–90.
- [55] F. Pérez and B. E. Granger, IPython: A system for interactive scientific computing, *Comput. Sci. Eng.* **9**, 21 (2007).
- [56] R. Martoňák, G. E. Santoro, and E. Tosatti, Quantum annealing by the path-integral monte carlo method: The two-dimensional random Ising model, *Phys. Rev. B* **66**, 094203 (2002).
- [57] I. Hen, J. Job, T. Albash, T. F. Rønnow, M. Troyer, and D. A. Lidar, Probing for quantum speedup in spin-glass problems with planted solutions, *Phys. Rev. A* **92**, 042325 (2015).
- [58] A. D. King, T. Lanting, and R. Harris, Performance of a quantum annealer on range-limited constraint satisfaction problems, [arXiv:1502.02098](https://arxiv.org/abs/1502.02098).
- [59] F. Hamze, D. C. Jacob, A. J. Ochoa, D. Perera, W. Wang, and H. G. Katzgraber, From near to eternity: Spin-glass planting, tiling puzzles, and constraint-satisfaction problems, *Phys. Rev. E* **97**, 043303 (2018).
- [60] D. Perera, I. Akpabio, F. Hamze, S. Mandra, N. Rose, M. Aramon, and H. G. Katzgraber, Chook—A comprehensive suite for generating binary optimization problems with planted solutions, [arXiv:2005.14344](https://arxiv.org/abs/2005.14344).
- [61] D-Wave Systems Inc, D-wave QPU architecture: Topologies, https://docs.dwavesys.com/docs/latest/c_gs_4.html.
- [62] N. Chancellor, Domain wall encoding of discrete variables for quantum annealing and QAOA, *Quantum Sci. Technol.* **4**, 045004 (2019).
- [63] J. Chen, N. Chancellor, and T. Stollenwerk, Experimental tests of a domain wall encoding of discrete variables (unpublished).
- [64] N. Chancellor, Code associated with paper Domain wall encoding of integer variables for quantum annealing and QAOA, <https://doi.org/10.15128/r27d278t029>.