

The power of pausing: advancing understanding of thermalization in experimental quantum annealers

Jeffrey Marshall,^{1,2} Davide Venturelli,³ Itay Hen,^{2,4} and Eleanor G. Rieffel³

¹*NASA Academic Mission Services R&D Student Program, USRA Quantum Academy*

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

³*QuAIL, NASA Ames Research Center, Moffett Field, California 94035, USA*

⁴*Information Sciences Institute, University of Southern California, Marina del Rey, California 90292, USA*

We investigate alternative annealing schedules on the current generation of quantum annealing hardware (the D-Wave 2000Q), which includes the use of forward and reverse annealing with an intermediate pause. This work provides new insights into the inner workings of these devices (and quantum devices in general), particular into how thermal effects govern the system dynamics. We show that a pause mid-way through the anneal can cause a dramatic change in the output distribution, and we provide evidence suggesting thermalization is indeed occurring during such a pause. We demonstrate that upon pausing the system in a narrow region shortly after the minimum gap, the probability of successfully finding the ground state of the problem Hamiltonian can be increased by several orders of magnitude. We relate this effect to relaxation (i.e. thermalization) after diabatic and thermal excitations that occur in the region near to the minimum gap. For a set of large-scale problems of up to 500 qubits, we demonstrate that the distribution returned from the annealer very closely matches a (classical) Boltzmann distribution of the problem Hamiltonian, albeit one with a temperature at least 1.5 times higher than the (effective) temperature of the device. Moreover, we show that larger problems are more likely to thermalize to a classical Boltzmann distribution.

I. INTRODUCTION

Inspired by thermal annealing and by the adiabatic theorem of quantum mechanics, quantum annealers are designed to make use of diminishing quantum fluctuations in order to efficiently explore the solution space of particular discrete optimization problems. In the last few years, chip sizes have grown in accordance with Moore’s law, and current devices contain on the order of 2000 superconducting qubits, potentially allowing for relatively large scale problems to be solved. Though progress has been made [1, 2], it has been hard to quantitatively demonstrate a general and rigorous ‘quantum speedup’ [3].

It has been suggested that these devices may instead be useful for thermal sampling tasks [4–7], such as the NP-hard problem of sampling from a Boltzmann distribution, which has application in machine learning and artificial intelligence [5, 6, 8]. Though it has been demonstrated that (classical) thermalization may indeed be occurring in these devices [9], there are still many unanswered questions pertaining to the exact distribution they sample from.

With the introduction of the latest generation of quantum annealing hardware, the D-Wave 2000Q, which allows one to change the default annealing schedule, including the ability to pause the anneal mid-way through the evolution for extended periods of time (up to around 1ms), new insights can be gained into the operation of these devices.

In particular we use intermediate pauses, followed by a ‘quench’ (i.e., evolving the system rapidly to the read-out point) to study effects of thermalization, and how this relates to key properties of the system, including the

location of the minimum energy gap. The observations we make are relevant not only for quantum annealers, but for any quantum device which is non-negligibly coupled to a thermal environment, thus shedding light on fundamental physical processes involved across a broad range of devices.

We first go through some relevant details pertaining to annealing and the theory of thermalization in such devices, before presenting our main results. After this, we summarize and interpret our results, as well as raising some new questions which should be answered in future works. We finish with a brief discussion setting our results in a broader context.

A. Background

Transverse field quantum annealing evolves the system over rescaled time $s = t/t_a \in [0, 1]$, where t is the time, and $t_a \in [1, 2000]\mu\text{s}$ the total run-time (chosen by the user). We will occasionally refer to the rate of the anneal ds/dt which can be set to zero during the pause or take values in interval $ds/dt \in [0.0005, 1]\mu\text{s}^{-1}$ otherwise. The time-dependent Hamiltonian is of the form

$$H(s) = A(s)H_d + B(s)H_p, \quad (1)$$

where $H_p = \sum_{\langle i,j \rangle} J_{ij}\sigma_i^z\sigma_j^z + \sum_i h_i\sigma_i^z$ is the programmable Ising spin-glass problem (the final Hamiltonian) to be sampled defined by the parameters $\{J_{ij}, h_i\}$, and $H_d = -\sum_i \sigma_i^x$ is a transverse-field (or ‘driver’) Hamiltonian which provides the quantum fluctuations (the initial Hamiltonian). Here N is the total number of qubits in the problem, and $\langle i, j \rangle$ indicates the sum is only




FIG. 1. Annealing schedule in GHz (units of $h = 1$). The operating temperature ($T = 12.1$ mK, or equivalently 0.25 GHz) of the chip is also shown (black-dashed line).

over a coupled qubits, defined by the hardware ‘chimera’ graph (see Fig. 23 of Appendix D). The device we use, the D-Wave 2000Q contains 16×16 unit cells each containing 8 qubits, thus having a maximum of 2048 qubits. Note however, due to some of the qubits/couplers being faulty, the actual number of operating qubits is 2031. These ‘dead’ qubits are randomly dispersed throughout the hardware graph.

The initial state is fixed as the ground state of H_d , $|\psi(0)\rangle = |+\rangle^{\otimes N}$ where $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ (defined in the computational basis via $\sigma^z = |1\rangle\langle 1| - |0\rangle\langle 0|$). The manner in which the Hamiltonian is evolved in time is determined by the annealing schedule (i.e. the time dependence of A, B), shown in Fig. 1.

After an annealing run, the system is measured in the computational basis. Performing many such runs allows statistics about the device to be collected; useful such measures include the probability of successfully finding the ground-state of H_p (which is the solution to a classical optimization problem) which we denote as P_0 , or the average energy returned $\langle E \rangle$.

One way to provide more robust statistics, is by changing the ‘gauge’ of the problem. This is a trivial re-mapping of the problem so as to avoid certain biases which may be present for certain couplers/qubits (e.g., some couplers may have less analog control errors associated with programming in $J = +1$ as compared to $J = -1$, or certain qubits may be more likely to align with $+z$ as compared to $-z$ even in the absence of any fields). The mapping involves changing the couplings/fields as $J_{ij} \rightarrow J_{ij}r_i r_j, h_i \rightarrow h_i r_i$, where $\vec{r} = (r_1, \dots, r_N)$ is a vector of random entries $r_i \in \{-1, 1\}$. Notice any configuration $\vec{s} = (s_1, \dots, s_N)$ has a corresponding configuration of the mapped problem $\vec{s}' = (r_1 s_1, \dots, r_N s_N)$ with the same cost, thus the problem itself is exactly the same.




FIG. 2. Dimensionless annealing schedule. We plot the ratio $Q(s) := A(s)/B(s)$, and the ratio of the operating temperature ($T = 12.1$ mK) to the strength of the problem Hamiltonian, $C(s) := k_B T / B(s)$.

An important quantity identified in Ref. [9] is the ratio between the strength of the driving Hamiltonian, and the problem Hamiltonian, and we define this as $Q(s) := A(s)/B(s)$, as shown in Fig. 2. Moreover, classical thermal fluctuations are governed by the quantity $C(s) := k_B T / B(s)$, where T is the temperature of the device. Indeed, observing the relative scales of the characteristic energies associated to the driving terms (i.e. transverse field, environmental bath) with the energy of the closed system allows us to infer the existence of different regimes where a given process becomes energetically dominant. In particular, i) at early times when $Q \gg C > 1$, and the system mostly remains in the ground state of H_d , ii) when $Q \sim C \sim O(1)$ and non-trivial dynamics occur with H_d driving various transitions between computational basis states, and iii) $Q \ll C \ll 1$ when the Hamiltonian is mostly diagonal (in the computational basis) and little population transfer occurs between the eigenstates (the ‘frozen’ region) through diabatic transitions. Thermal transitions could occur but those depend also on the strength of the coupling to the thermal bath (see below).

B. Adapting the standard annealing schedule

The current generation of hardware, the D-Wave 2000Q, allows users to tweak the default schedule in various ways. In particular this gives one the ability to:

- 1) Pause the evolution at some intermediate point $s_p < 1$ in the anneal, for time t_p .
- 2) Perform reverse annealing, where the system is initialized in a classical configuration at time $s = 1$, evolved backwards to an earlier time $s_p < 1$, where




FIG. 3. Example of annealing parameter s as a function of time t for an anneal with a pause, for both forward and reverse annealing. Here a $1\mu\text{s}$ pause ($t_p = 1\mu\text{s}$) is inserted into the annealing schedule at $s_p = 0.5$ (i.e. at $t = 0.5\mu\text{s}$), which otherwise has a total anneal time of $t_a = 1\mu\text{s}$.

a pause can be inserted, before evolving the system back to time $s = 1$ where a read-out occurs.

- 3) Speed-up or slow-down the schedule during intermediate intervals of the anneal.
- 4) Provide per-qubit annealing offsets.

Based on these features, new methods of sampling from an annealer have been developed and proposed, for example exploiting reverse annealing to explore the energy landscape in a novel manner [10–14]. Moreover, performance advantages have been observed by offsetting the fields of some of the qubits, allowing one to evade spurious transitions which occur during the minimum gap [15, 16].

The work presented in this paper mostly focuses on the first capability listed above, where we embed a pause in the default annealing schedule, i.e., the Hamiltonian is fixed at $H(s_p)$ for a certain period of time chosen by the user. This will allow us to study key mechanisms determining the output of the device, such as thermalization. We also briefly study reverse annealing in a similar context as it enables one to learn about different regimes during the anneal (as discussed in Sect. II B). We show an example of an annealing schedule with a pause in Fig. 3.

C. Theory

It has been argued that the output of devices of this type corresponds to a quasi-static evolution [4]; one in which the internal system dynamics are much slower than the dynamics induced by the environment, meaning the system thermalizes instantaneously.

Results however seem to suggest that the final state is not the Boltzmann distribution of the problem Hamiltonian at $s = 1$ [9]; that is, the samples are not drawn from $\rho \sim \exp(-\beta B(1)H_p)$, where $\beta = 1/k_B T$, with T the operating temperature of the device (on the order of 10mK in current devices).

A possible explanation for this observation is that there is a problem dependent freeze-out point, s^* , which occurs at some time when $Q(s^*)$ and $C(s^*)$ are ‘small’, so that the time-scale upon which the transverse field drives transitions between eigenstates of H_p is much longer compared to the system evolution time, hence little population transfer occurs. This would mean that one is sampling from a Boltzmann distribution of the Hamiltonian at time s^* , instead of at $s = 1$. If indeed this occurs when $Q(s^*) \ll 1$, the expected distribution would be close to the form $\rho \sim \exp(-\beta B(s^*)H_p)$ [4, 9].

An open-system treatment of this phenomena is described in Ref. [17] in the weak coupling regime for general problems, and in [18, 19] in the non-perturbative regime for specific problems. In the weak coupling limit, transitions between instantaneous energy levels $E_j(s) > E_i(s)$ are governed by Fermi’s Golden Rule rate Γ_{ij} :

$$\Gamma_{ij}(s) \propto \sum_{k,\alpha} |\langle E_i(s) | \sigma_k^\alpha | E_j(s) \rangle|^2 \frac{g_\alpha^2}{1 - \exp(-\beta |E_i(s) - E_j(s)|)}, \quad (2)$$

where g_α is the environment coupling strength to the $\alpha = x, y, z$ component of the system spins, and σ_k^α is the α Pauli operator acting on the k -th spin.

The explanation of freeze-out in this picture is that as $s \rightarrow 1$, energy gaps $|E_i - E_j|$ open up, as well as the matrix elements $\langle E_i(s) | \sigma^z | E_j(s) \rangle \rightarrow 0$ (which is typically the dominant environment-spin coupling [20, 21]) as the Hamiltonian becomes more diagonal in the z -basis, thus the transition rates dramatically slow down late in the anneal. Therefore, the two strongest (possibly competing) effects determining the relaxation rate Eq. (2) are the instantaneous energy gap, and $Q(s)$ (via the matrix element).

An extensive analysis of this conjecture was carried out in Ref. [9], which indeed found evidence for this phenomena, however many new questions were opened up. In particular, whilst some problems were shown to have a well defined freeze-out point late in the anneal, the majority of problems tested did not, and little could be said about these problems. Some possibilities include susceptibility to analog control error noise (errors in the analog programming of the coupling strengths J_{ij}), different parts of the problem freezing out at different points (hence no single well defined s^* exists), and different types of noise source affecting the experimental device.

With the possibility to pause the system during intermediate times, unprecedented insights into the effect of thermalization can be gained. We illustrate some of the key time-scales involved in open system annealing in Fig. 4.




FIG. 4. A cartoon example showing the various relevant time-scales involved in a quantum annealer with coupling to a thermal bath. We indicate the thermal relaxation time-scale t_r (i.e. related to the inverse of Eq. (2)) by the blue solid line, the system time-scale, t_a , by the red (horizontal) solid line, and the pause time scale, t_p , by the black dash line [22]. The relaxation time scale is shortest around the location of the minimum gap of the problem (e.g. around $s = 0.44$ in this example), and it quickly diverges exponentially as the gap opens up. We show four characteristic regions during the anneal. 1) [yellow] at early times, when the ground state is separated by a large gap from any excited state, and the population in the ground state is approximately 1. 2) [green] As the energy gap closes, and the relaxation rate increases rapidly, thermal excitations may occur. 3) [purple] As the energy gaps open up, the relaxation time scale increases (exponentially). Once $t_r > t_a$ instantaneous thermalization can no longer occur effectively. This point corresponds approximately to the freeze-out point. If however $t_r < t_p$, a pause may allow effective thermalization, which could lead to a significantly larger population in the ground state. We indicate the transition region (where $t_r \approx t_p$) by s_p^{opt} (see main text). Finally, in region 4) [blue], where t_r is the longest time scale, very little population transfer will occur (even if one pauses the system for time t_p). The dynamics are effectively frozen.

II. RESULTS

Throughout this work we typically considered problems of two different types. 1) When wanting to study large problems (such as the \mathcal{I}_{800} instance of Fig. 5) we work with problems of the planted-solution type [23], due to the ability to know in advance the general analytic form of the spectrum of H_p , including the ground state (as well as certain information about the degeneracy of the energy levels [9]). 2) When analyzing spectral properties of the full Hamiltonian, $H(s)$, we used (small) problems where the $J_{ij} \in [-1, 1]$ (uniformly random). In both cases, we used zero local fields $h_i = 0$.

While problems with local fields or large ferromagnetic structure (e.g. embedded problems) could benefit from specific analysis, we expect that the general results and

arguments presented will be generalizable to a large range of problem sets.

A. Forward annealing with a pause

We consider the following simple adaptation to the standard annealing schedule. Allow the system to run as normal to some (re-scaled) time $s_p \in [0, 1]$, upon which we pause the system for time $t_p \in [0, 2000]\mu\text{s}$, after which we continue the evolution as per normal.

We observe that this dramatically effects the samples returned from the D-Wave device, as demonstrated in Fig. 5. In particular, almost all problems we tested exhibit a strong peak in the success probability when a pause is inserted into the regular annealing schedule. The corresponding average energy returned is also significantly reduced. We define the ‘optimal pause point’, s_p^{opt} , as the point in the anneal for which a pause returns the lowest average energy returned from many samples [24] (just after $s_p = 0.4$ in this example).

Moreover, the longer the pause, the greater the increase in the success probability, as shown in Fig. 6. Here we see that the success probability, for a pause at re-scaled time $s_p = 0.44$, increases from the baseline ($\approx 0.01\%$) to over 10% for pause times longer than around $500\mu\text{s}$, and approaches 20% as the pause time approaches $2000\mu\text{s}$ (the longest allowable pause time on the D-Wave device), although saturating around 1ms (in the logarithmic regime). That is, an increase of around three orders of magnitude. This gives us new insight into the time-scales involved in these devices. It shows that even a $10\mu\text{s}$ pause (inserted within a default schedule with $t_a = 1\mu\text{s}$) can dramatically effect the nature of the samples returned from the device.

These observations are consistent with the thermalization picture mentioned in the previous section, and the cartoon in Fig. 4; we attribute the purple region in Fig. 4 to the region where the huge spike in success probability is observed, since the system can effectively relax back to the low lying energy levels on a time-scale comparable with the pause length. After this (e.g. the blue region in Fig. 4), the effect is much weaker (dropping off exponentially) as the relaxation time scale increases (notice in Fig. 5 that late in the anneal, the relative increase in success probability is much less, or non-existent, as compared to during the region around $s_p = 0.4$).

We observe similar phenomena for the second problem class we study (with $J_{ij} \in [-1, 1]$), as in Figs. 7,8. In these figures we show the effect of changing the pause time, and the anneal time respectively. In the Sect. II C we will relate this to the minimum gap (which can be computed exactly for these problems).

An effect we observe upon increasing the pause time is that the width of the peak increases, as shown in Fig. 7. Notice that in this figure all curves start to show an increase in success probability at the same pause point s_p (just after 0.4), but come back to the baseline probability




FIG. 5. Forward annealing with a pause for a single 800 qubit (planted-solution) problem instance (\mathcal{I}_{800}). The top figure shows the success probability with respect to pause time t_p , and the location of the pause s_p . The total evolution time (aside from the pause) is $1\mu\text{s}$. The corresponding bottom figure shows the average energy (in arbitrary units) returned by the device. Each data point is averaged from 10000 anneals with 5 different choice of gauge. In the absence of a pause, $P_0 \approx 10^{-4}$.

at later points for longer pause times. That is, the region of interest is slightly extended to the right. This also fits in with the model discussed in Sect. IC and the cartoon picture Fig. 4, where increasing t_p increases the size of the purple region by extending it to the right (i.e. the region where $t_a < t_r < t_p$). The location of the peak s_p^{opt} we posit to be around the point when $t_r \approx t_p$ (i.e. the interface of the blue and purple regions in Fig. 4), the last point in the anneal for which thermalization can effectively occur during a pause of length t_p . Indeed we experimentally observe (in Fig. 7) that increasing the pause time shifts the peak to later in the anneal (and also increases in size in accordance with this picture [25]).

If one instead increases the anneal time, the peak narrows (and flattens), and eventually disappears, as observed in Fig. 8. Note, in accordance with Fig. 4, the location of s_p^{opt} does not change upon increasing t_a (since this should reduce the size of the purple region from the




FIG. 6. Dependence of success probability P_0 on pause length t_p , for the same problem considered from Fig. 5 (\mathcal{I}_{800}), where we fix $s_p = 0.44$ (corresponding to the peak in Fig. 5). We see increasing the pause length corresponds to a larger success probability (although it mostly saturates around $500\mu\text{s}$). In the absence of a pause the success probability is $P_0 \approx 10^{-4}$, which increases by several orders of magnitude to around 20%. Red solid line is linear fit to tail end ($t_p > 10^{1.5}\mu\text{s}$) Inset: Same as main figure, but with linear scale on t_p -axis. Each data point is averaged from 10000 anneals with 5 different choice of gauge.




FIG. 7. Effect of changing the pause time for a 12 qubit problem instance (\mathcal{I}_{12}^0). Each data point is from 10000 annealing runs using 5 different gauges. The anneal time $t_a = 1\mu\text{s}$ for all data sets shown.

left). We also show a corresponding heat map of this effect in Appendix D (Fig. 25).

It is remarkable that the peak is extremely well defined, occurs in such a concentrated region, and exists for almost all problems we studied (the only exception being some small qubit instances, which we discuss below). For problems of the planted-solution type, there




FIG. 8. Effect of changing the anneal time for \mathcal{I}_{12}^0 . Each data point is from 10000 annealing runs using 5 different gauges. The anneal time $t_p = 100\mu\text{s}$ for all data sets shown.

seems to be little dependence on problem size for the position of this peak. We demonstrate this in Fig. 9 where we see the optimal pause point, s_p^{opt} , does not vary much with problem size, and in addition, the deviation (i.e., the error bars in the figure) in the samples is more or less constant. For problems which are generated with $J_{ij} \in [-1, 1]$ (uniformly random), we see a mild effect with increasing problem size, where the optimal pause point seems to decrease, and concentrate in location (i.e. the error bars in the figure are decreasing with problem size). This effect would presumably saturate with large enough N (note, in the figure, $\text{SL}=16$ is the largest possible problem size available).

The simple observations demonstrated here show that one may be able to design more efficient annealing schedules by annealing very quickly, and pausing for a relatively short time, as compared to running the default annealing schedule for a long time. Moreover in Fig. 10, we see the width of the peak only depends very weakly (or not at all) on the problem size; this suggests that for most problems, regardless of size, there is a fairly large window in which one can pause and observe an increase in success probability.

B. Reverse annealing with a pause

Before proceeding with a statistical analysis we briefly present some relevant results using the reverse annealing protocol, with a pause, the general protocol of which is demonstrated graphically in Fig. 3. This allows us to identify some of the key regions during an anneal, which, as explained above, depend on the ratios of the various energy scales involved and associated time-scales.

We show some of our findings in Fig. 11 where we identify 4 regions of interest. 1) $s_p < s_{\text{gap}}$. The system has been evolved (from $s = 1$) past the minimum gap, and




FIG. 9. The optimal pause point s_p^{opt} as a function of problem size, for two different problem classes (see legend). The problems were generated on a square subgraph of the chimera with ‘side-length’ SL , consisting of $\text{SL} \times \text{SL}$ unit cells each containing 8 qubits (e.g. see Fig. 2 in Ref. [9]). Each SL shown [4, 8, 12, 16] corresponds to (taking account for dead qubits) $N = [127, 507, 1141, 2031]$ respectively. Each data point shown is an average over (at least) 50 instances. Error bars represent the standard deviation. Each instance (for each s_p tested) is averaged from 10000 anneals with 5 different choice of gauge, with $t_a = 1\mu\text{s}$ (not including the pause time), and $t_p = 100\mu\text{s}$.

been paused at a point where $Q > O(1)$, allowing for mixing between energy levels in the computational basis. There is no memory of the initial configuration. 2) In the region just after the minimum gap, up to around s_p^{opt} , where the lowest energy solutions are found, and corresponding to the purple region in Fig. 4. Here there is no memory of the initial state, and no clear difference between forward and reverse annealing. We expect thermalization is able to occur effectively on times scales comparable with t_p , i.e., the transition rates between energy levels $\Gamma_{ij} \geq 1/t_p$. 3) After s_p^{opt} , where there is a clear difference between forward annealing and reverse annealing, there is ‘memory’ of the initial configuration. Thus the state returned by the annealer at $s = 1$ depends heavily on the system state at the pause point s_p suggesting different time-scales and transition rates are important here. In this region, as $Q \rightarrow 0$ and $\langle E_i | \sigma^z | E_j \rangle \rightarrow 0$, the $g_{x,y}$ couplings may play more of a role, leading to qualitatively different thermalization mechanisms and time-scales. Here some transition rates Γ_{ij} may be comparable to $1/t_p$, whereas others much less. 4) Very late in the anneal, with $Q \ll C \ll 1$, almost no dynamics occur (the state returned from the annealer is almost always the same as the one initialized), i.e. $\Gamma_{ij} \ll 1/t_p$.

We mention that these general observations seem to be fairly generic, and not specific to this particular example. With this in mind, we proceed with a statistical analysis, demonstrating to what extent the picture outlined in




FIG. 10. The width of the peak on a graph of $|\langle E \rangle|$ Vs. s_p , as a function of problem size. The ‘full-width-at-half-maximum’ (FWHM) is the width of the peak in the energy curve $|\langle E \rangle|$ as a function of s_p , at $(|E(s_p^{\text{opt}})| + |E_{\text{BG}}|)/2$ where E_{BG} is the background average energy (i.e., the average energy returned from the annealer in the absence of a pause), and $E(s_p^{\text{opt}})$ the (mean) energy returned by the annealer at the optimal pause point. Note, modulus is used since all energies observed negative. The problems (same as in Fig. 9) were generated on a square subgraph of the chimera with ‘side-length’ SL, consisting of $\text{SL} \times \text{SL}$ unit cells each containing 8 qubits (e.g. see Fig. 2 in Ref. [9]). Each SL shown [4, 8, 12, 16] corresponds to (taking account for dead qubits) $N = [127, 507, 1141, 2031]$ respectively. Each data point shown is an average over (at least) 50 instances. Error bars represent the standard deviation. Each instance (for each s_p tested) is averaged from 10000 anneals with 5 different choice of gauge, with $t_a = 1\mu\text{s}$ (not including the pause time), and $t_p = 100\mu\text{s}$.

Fig. 4 holds.

C. Correlation with the minimum gap

Typical folklore of (open system) quantum annealing dictates that around the location of the minimum gap, thermal excitations from the ground state to excited energy levels may occur, and that after the gap, thermal relaxation will allow some of the excited population to fall back to the ground state [26] (of course, this is heavily dependent on the nature of the minimum gap, and hence on H_p itself, as well as the temperature and annealing schedule). This general idea is also demonstrated in Fig. 4. This framework would suggest that a pause in the annealing schedule some (problem dependent) time after the minimum gap may lead to an increase in the success probability (that is, the population in the ground state at time $s = 1$).

Working with 12 qubit problems with $J_{ij} \in [-1, 1]$ (uniformly random), we indeed find such a correlation between the location of the minimum gap, and the optimal pause point, s_p^{opt} , where for over 90 of the 100 prob-




FIG. 11. Reverse annealing with pause at s_p (solid lines), for a four different initial configurations for a 12 qubit problem \mathcal{I}_{12}^0 . We plot the average energy (arbitrary units) returned from 5000 anneals, evolved at rate $ds/dt = 1\mu\text{s}^{-1}$, and $t_p = 100\mu\text{s}$. We consider ground and first excited state configurations, as well as two highly excited energy levels. The energy level E_i is indicated on the right hand side. We also show the corresponding forward anneal curve (black-dash line). This problem has a minimum gap at $s = 0.44$ indicated in the figure. Note a sample of the spectrum for this problem is shown in Fig. 21 in Appendix A.

lems tested the best place to pause is after the minimum gap. This is demonstrated in Fig. 12, where on average $s_p^{\text{opt}} \approx s_{\text{gap}} + 0.14$.

We comment briefly on the few outliers (e.g. with $s_{\text{gap}} < s_p^{\text{opt}}$) in the data set. For some of these small 12 qubit problems, they are solved almost 100% of the time by the D-Wave (i.e. they are extremely easy optimization problems). These are problems that have large minimum gaps $\Delta_{\text{min}} > 1\text{GHz}$, and we indicate them in red in the plot (also see Fig. 24 in Appendix D). For these instances, we typically do not observe a well defined optimal pause point; since the gap is so large for all s , we expect very little thermal excitation to occur at all, hence pausing has little effect. We see these red points have a fairly random spread in the s_p^{opt} -axis.

The second set of outliers are (some of the) instances which have minimum gaps relatively late in the anneal. A possible explanation of this is that again, some of these problems do not have well defined optimal pause points. If the minimum gap is late in the anneal, when $Q(s_{\text{gap}})$ is small, the transition time scale t_r may already be too large for effective thermalization to occur during a pause (i.e. the pause has little effect). We mention this is heavily dependent on the problem itself, since the transition rates depend on the spectrum and eigenstates of the problem (see Eq. (2)), and expect this is why it seems to only be an issue for a few instances. Nevertheless, the overall trend is clear, with the majority of problem instances exhibiting an optimal pause point in a small region shortly after the location of the minimum gap.




FIG. 12. Correlation between the location of the minimum gap, s_{gap} , and the optimal pause point s_p^{opt} for 100 problems of size 12 qubits. These problems have $J_{ij} \in [-1, 1]$ (uniformly random) and $h_i = 0$. We divided the data into two groups based on their minimum gap Δ_{min} (see legend), and we also indicate by the dash-dot (vertical) line the location corresponding to $Q(s_{\text{gap}}) = 0.1$. We fixed the pause time to $t_p = 1000\mu\text{s}$, and total anneal time (excluding the pause) to $t_a = 1\mu\text{s}$. Data from the D-Wave is averaged over 10000 anneals, with 10 different choices of gauge.

We similarly study the same problems where we rescale the problem Hamiltonian by a constant factor (2,4,8). This has two effects; 1) it shifts the location of the minimum gap to later in the anneal, and it also reduces the size of the minimum gap (as an explicit example, see Fig. 26 in Appendix D). We indeed see that correspondingly, the location of the optimal pause point shifts to later in the anneal (see inset of Fig. 13). We show this explicitly for a single problem in Fig. 27 in Appendix D.

Interestingly, we also observe that the location of s_p^{opt} concentrates (thus becoming less correlated with s_{gap}) upon reducing the energy range of the problem; notice how in Fig. 13, the purple points ($H_p/8$) are almost perfectly aligned close to $s_p^{\text{opt}} = 0.8$. We also see this by noting that the error bars (standard deviation) in the s_p^{opt} -axis decrease (inset of figure).

We explain this by pointing out that by dividing H_p by a large enough factor, $\beta\omega_{01}(s) < 1$ for $s > s_{\text{gap}}$ where $\omega_{01}(s) := E_1(s) - E_0(s)$ is the gap between the ground and first excited state, and β the inverse temperature of the device. This implies the system should be able to effectively thermalize until very late in the anneal, until the matrix elements in Eq. (2) become small enough (determined by Q being small enough). This means the freeze-out point (i.e. start of the purple region in Fig. 4)




FIG. 13. Effect of reducing problem energy scale (for the same 100 problems studied in Fig. 12). We divide the problem Hamiltonian H_p by 1,2,4,8 (see legend). Inset: Mean data point for each group of 100 instances upon dividing H_p by 1,2,4,8 (see legend). Dash-dot line is least squares fitting to the median data point. Error bars are the standard deviation.

is not determined so much by the problem itself (i.e. the exact spectrum as a function of s), but the annealing schedule (i.e. $Q(s)$). Thus different problems may exhibit very similar optimal pause points.

It would be worth while to explore this in more detail to confirm this hypothesis.

D. Quantum Boltzmann distribution

For a set of 10 problem instances of size 12 qubits \mathcal{T}_{12}^{0-9} (each with well defined optimal pause points, with $\Delta_{\text{min}} < 1\text{GHz}$, $Q(s_{\text{gap}}) > 0.1$) with $J_{ij} \in [-1, 1]$ (uniformly random), we compare the returned D-Wave statistics to the instantaneous quantum Boltzmann distribution $\rho \sim \exp(-\beta H(s))$ (projected to the z -eigenbasis), for various choices of β ; we vary the temperature T from $\frac{1}{4}T_{\text{DW}}$ to $4T_{\text{DW}}$, in increments of $\frac{1}{4}T_{\text{DW}}$ (where $T_{\text{DW}} = 12.1\text{mK}$ is the operating temperature of the device). We outline these calculations in Appendix B.

We study these problems, with a very long pause time of $t_p = 1000\mu\text{s}$ to allow for enough time to thermalize, and run with short anneal time (excluding pause time) $t_a = 1\mu\text{s}$ (i.e. the quickest possible time, so that we can try to isolate the effect of the pause).

An example of this analysis for a single instance is shown in Fig. 14, however we note the pattern looks much the same for all of the instances studied (see Fig. 15). Here we focus on the distribution returned from the device with a pause at the optimal pause point s_p^{opt} , and

compare this to a distribution of the form $\exp(-\beta H(s))$.

We observe that the D-Wave distribution is best described by a hotter Boltzmann distribution, at a later point in the anneal than the optimal pause point. In particular for these problems the optimal parameters (s^*, T^*) such that D_{KL} is minimized, correspond to $s^* = 0.78 \pm 0.10$ and $T^* = 26.1 \pm 8.8 \text{mK}$ (and up to 4 times the physical temperature). Moreover, $s^* - s_p^{\text{opt}} = 0.21 \pm 0.11$ (i.e. significantly later in the anneal than the optimal pause point). At values of s^* this late in the anneal (with $Q(s^*) \approx 10^{-3}$), the distribution is indeed effectively a classical Boltzmann distribution (of H_p); that is, $\exp(-\tilde{\beta} H_p)$, where $\tilde{\beta}$ is an effective inverse temperature (one which is typically hotter than T_{DW}).

This result is somewhat confusing. For one, the temperatures seem to vary wildly between instances as seen in Fig. 15, and moreover, it suggests that non-trivial dynamics can occur well past the optimal pause point. Since we expect the region around the optimal pause point to correspond approximately to the freeze-out point (i.e. the region where the system time-scale is short compared to the thermal time-scale), we would expect to see a closer match between s^* and s_p^{opt} .

We conjecture that after the optimal pause point non-trivial dynamics do indeed still occur. The intermediate pause helps the D-Wave distribution equilibrate to the instantaneous thermal distribution, and after this, although the thermal transition rate is too small to thermalize effectively (even with a pause), some dynamics will inevitably occur (recall, the minimum annealing time scale is $1 \mu\text{s}$). It is possible therefore that the distribution one eventually measures is not quite a projection of a Boltzmann distribution at a well defined (s, T) .

We partly validate this picture in Fig. 16 comparing the D-Wave distribution of a single instance to the optimal found over all (s, T) in $\exp(-\beta H(s))$. This problem fits best to a Boltzmann distribution with optimal values $(s^*, T^*) = (0.76, 18.5 \text{mK})$. We plot on a logarithmic scale to show the similarity to a classical Boltzmann distribution of H_p , for which $\ln \frac{P_i}{g_i} = -\beta E_i - \ln Z$ (i.e. a straight line on this graph). This late in the anneal we expect the Boltzmann distribution to be approximately of the form $\exp(-\tilde{\beta} H_p)$ for some appropriate $\tilde{\beta}$, since $Q(s^*) < 10^{-2}$. Also note that $Q(s_p^{\text{opt}}) = 0.07$, so this is not an unreasonable assumption even for instantaneous thermalization at the optimal pause point.

We notice that indeed, both the experimental data and the computed Boltzmann distribution seem to fit reasonably well to a linear fit (indicating classical thermalization is occurring to some extent), but also mention it is not clear whether the two distributions correspond to one and the same; there are clearly some large divergences (the y axis is a logarithmic scale).

It is also interesting to note that in the absence of a pause we do not see any clear correspondence between the D-Wave distribution and a Boltzmann distribution. For example, for the distribution returned from an anneal with pause at s_p^{opt} , the minimum KL diver-




FIG. 14. KL divergence D_{KL} between data from D-Wave P_{DW} and the Boltzmann distribution P_{QBM} (projected into the computational basis) for various choices of (s, T) , for a single 12 qubit instance (\mathcal{I}_{12}^1). The D-Wave data is sampled from a pause of length $t_p = 1000 \mu\text{s}$ at s_p^{opt} , with $t_a = 1 \mu\text{s}$ (from 10000 anneals and 10 choices of gauge). We indicate in the plot three key parameters; the physical temperature $T_{\text{DW}} = 12.1 \text{mK}$, the location of the minimum gap s_{gap} , and the optimal pause point s_p^{opt} . The white diamond corresponds to the minimum value of D_{KL} over all (s, T) , and is equal to $D_{\text{min}}^{\text{KL}} = 0.01$ bits of information. Note, to be able to distinguish the features in the plot, we set the upper limit on the plot to be $D_{\text{KL}} = 0.2$ (any value above this is mapped to this color).




FIG. 15. For the 10 problems described in the main text (\mathcal{I}_{12}^{0-9}), we plot the values (s^*, T^*) (relative to s_p^{opt}) which correspond to the minimum value of D_{KL} over all choices of (s, T) tested. We see the same trend for all of the problems, as observed directly in Fig. 14, where larger s^* implies larger T^* . Note in the plot, two of the data points lie on top of each other. The mean optimal KL-divergence found for these problems was $D_{\text{min}}^{\text{KL}} = 0.016 \pm 0.015$. Note, $T_{\text{DW}} = 12.1 \pm 1.4 \text{mK}$.




FIG. 16. Comparison of the D-Wave probability distribution and the closest fit (as measured by KL-divergence) to a quantum Boltzmann (QBM) distribution of the form $\exp(-\beta H(s))$, where the fit is over the parameters (s, T) for \mathcal{L}_{12}^2 . Optimal values for this problem are $(s^*, T^*) = (0.76, 18.5\text{mK})$, and $s_p^{\text{opt}} = 0.59$. Here p_i is the probability of observing a configuration with energy E_i , and g_i is the degeneracy of that energy level. The solid and dash-dot lines are least-squares fitting to the D-Wave data and the QBM data respectively. D-Wave data taken at the optimal pause point for this problem, from 10000 anneals (10 gauges), with $t_a = 1\mu\text{s}$, $t_p = 1000\mu\text{s}$.

gence is $D_{\min}^{\text{KL}} = 0.016 \pm 0.015$ (and $D^{\text{KL}}(s_p^{\text{opt}}, T_{\text{DW}}) = 0.076 \pm 0.065$). However, performing the same analysis with an anneal with no pause, the values vary wildly, with $D_{\min}^{\text{KL}} = 0.19 \pm 0.15$. That is (over the range of (s, T) for which we computed $\exp(-\beta H(s))$) there is typically no good choice of (s, T) , and it is not clear what the distribution is.

If the system is not able to thermalize appropriately during the anneal (e.g. because the anneal time t_a is too quick), there is no expectation for the D-Wave distribution to be close to any Boltzmann distribution.

With this in mind, though we have provided evidence suggesting thermalization is occurring to some extent in these problems, more work is required to understand precisely the distribution one is sampling from at the optimal pause point. In particular, we would like to better understand why the fitting temperature seems to vary so much between these instances, and why the dynamics seem to occur so long after the optimal pause point. The main constraint in our experiments prohibiting us to probe these details further is the maximum annealing rate $\frac{ds}{dt}$ which is limited to $1\mu\text{s}^{-1}$ on the present device. That is, even though we hope to be approximately quenching the system from after the optimal pause point (i.e. measuring in the middle of the anneal), in reality there is still plenty of time for dynamics to occur until measurement at $s = 1$.

We provide some more intriguing evidence in the next

section, where we focus on a set of (large-scale) problem instances which indeed seem to exhibit classical thermalization (i.e. thermalization to a Boltzmann distribution of H_p).

E. Classical Boltzmann distribution

Whether or not the device samples from a classical Boltzmann distribution at some point s^* late in the anneal is a hotly contested issue [4–7, 9]. If problems for which machine learning is applicable (e.g. in the restricted Boltzmann machine paradigm) freeze-out at a point late in the anneal, when $Q \ll 1$, then these devices may show an advantage over classical samplers (sampling from a Boltzmann distribution is NP-hard) [5, 6, 8].

With the advent of a new entropic sampling technique based on population annealing [27], we were able to accurately estimate the degeneracies for 225 planted-solution type instances containing 501 qubits (that is the estimated ground and first excited state degeneracies are within 5% of the known values found by planting). For more information on these techniques see Refs. [9, 27], and Appendix C. Due to the large size of these problems, we are of course not able to compute the Boltzmann distribution of the full Hamiltonian $H(s)$ as we did in the previous section.

However, having accurate values for the degeneracies allows us to calculate the classical (problem Hamiltonian) Boltzmann distribution $\rho \sim \exp(-\tilde{\beta} H_p)$, where $\tilde{\beta}$ is an effective inverse temperature, i.e., a fitting parameter, which depends on the physical temperature, and the strength problem Hamiltonian B (and in principle anything else affecting the distribution returned from the device such as various noise sources).

If the distribution returned from the D-Wave is indeed a classical Boltzmann distribution at freeze-out point s^* late in the anneal (when $Q(s^*) \ll 1$), then one would expect (ignoring any other noise sources) to find $\tilde{\beta} = \beta B(s^*)/J_{\max}$, where $\beta = 1/k_B T$ is the physical inverse temperature, $B(s^*)$ is the problem Hamiltonian strength at the freeze-out point, and $J_{\max} = \max |J_{ij}|$ is a normalization parameter (since the J_{ij} programmed into the quantum annealer are restricted to the range $[-1, 1]$).

We make two key observations. 1) all of these problems tested exhibit a strong peak in the success probability in a fairly narrow region $s_p^{\text{opt}} \in [0.35, 0.46]$ (i.e., much less varied than the small 12 qubit instances studied above). 2) The data returned from D-Wave, for all of these problems, matches very closely to a classical Boltzmann distribution, but at a higher temperature than the operating temperature of the device (at least 1.5 times higher). This is in accordance with the results of Ref. [9] where calculated freeze-out points for most large problems were very early in the annealing schedule (equivalent to a higher than expected temperature), although in this work, a pause during the anneal to more directly




FIG. 17. Fitting experimental data to linear fit for a single 500 qubit instance (\mathcal{I}_{500}). We show data for three different pause points (and from the standard annealing schedule). We see the standard schedule is almost indistinguishable from the case where the anneal is paused at $s = 0.2, 0.8$ (slope $\tilde{\beta} \approx 0.35$), however, when pausing at $s = s_p^{\text{opt}} = 0.44$ for this instance, we see a change in the distribution returned (purple, with slope $\tilde{\beta} = 0.44$). We also plot the corresponding classical Boltzmann distribution expected (black-dash line) if the system were thermalizing to H_p at $s^* = s_p^{\text{opt}}$, at temperature 12.1mK (with slope $\tilde{\beta} = 0.61$). Note, freeze-out at $s^* = 1$ would give $\tilde{\beta} = 2.52$. Inset: Average energy returned by D-Wave as a function of pause point for the same instance. The curve has a minimum value at $s_p^{\text{opt}} = 0.44$. Experimental data obtained from 10000 anneals with 5 choices of gauge, with $t_a = 1\mu\text{s}$, and $t_p = 100\mu\text{s}$.

study thermalization was not available. We demonstrate these points below.

First consider Fig. 17 where for a single instance we plot $\ln \frac{p_i}{g_i}$ against E_i , where p_i is the probability of observing a configuration with energy E_i , and g_i is the degeneracy of that energy level. One can see the data returned from D-Wave corresponds very closely to a linear fit (in fact, for all problems, and all pause points s_p , we find the R^2 (coefficient of determination) value is greater than 0.97, and up to 0.9999). That is, the data from D-Wave seems to fit to $p_i = \frac{g_i}{Z} e^{-\tilde{\beta} E_i}$, for some constants Z , and $\tilde{\beta}$ (which can be determined by least-squares fitting).

Though the results are very clear, the correct interpretation of them is not. For example, if we obtain the effective inverse temperature of the distribution $\tilde{\beta}$ from the least squares fitting, and set it equal to $\tilde{\beta} = \beta B(s^*)/J_{\text{max}}$, with knowledge of $\beta = 1/k_B T$, and J_{max} , we can calculate $B(s^*)$. If one does this however, the value returned corresponds to an extremely early point during the anneal, even earlier than s_p^{opt} (e.g. with $Q \approx 1$, or equivalently $s \approx 0.35$). If one however assumes the thermalization picture presented in Sect. IC, which suggests freeze-out should occur when the relaxation time scale is longer than the system time scale, i.e. approxi-

mately around the optimal pause point, $s_p^{\text{opt}} = 0.42 \pm 0.01$ for these problems, the temperature required for the fit is > 1.5 times higher than the physical temperature $T = 19.8 \pm 1.1\text{mK}$ (compared to $T_{\text{DW}} = 12.1 \pm 1.4\text{mK}$).

It is however not clear whether $\exp(-\beta H(s^*))$ with $s^* = 0.42$ and $Q(s^*) \approx 0.5$ would indeed correspond to a classical Boltzmann distribution (of H_p) since the off diagonal driver is still relatively strong in magnitude. Indeed, this is a somewhat similar result as from the previous section where we observed the optimal parameter value for s was in fact slightly after s_p^{opt} (and T larger than the physical temperature). If there are in-fact still dynamics after s_p^{opt} , the freeze-out point will be somewhat later in the anneal (when $Q(s)$ is smaller), and the associated temperature of the fit will be larger. We discuss some implications of this in the next section. For now we compare the samples from the optimal pause point to those from outside of it.

In Fig. 18 we plot the R^2 value found by the least squares fitting for a typical instance as a function of pause point s_p . We see that the peak corresponds closely to the optimal pause point under the pause. Moreover, we see a very similar trend for all of the problems we tested, with the pause point for which the largest R^2 value is observed differing by at most 3% of the annealing schedule from the optimal pause point; $s_{R^2}^{\text{max}} = s_p^{\text{opt}} \pm 0.03$. This indicates that the data returned from a pause at this critical point, fits better to a classical Boltzmann distribution, as compared to the rest of the data (or indeed, from the distribution returned in the absence of a pause in the schedule). If indeed the problems are thermalizing to a classical Boltzmann distribution, this work shows that by pausing the anneal at a particular (instance-dependent) point allows a more complete thermalization to occur. This result is similar to that found in the previous section.

We performed a similar analysis for three other problem sizes ($N = 31, 125, 282$ qubits), and find that increasing the problem size in general increases the mean R^2 value for a fit to a classical Boltzmann distribution. For $N = 31, 125, 282, 501$, the corresponding values are $\langle R^2 \rangle = 0.911, 0.994, 0.995, 0.997$, where the average (median) is over all instances and all pause points s_p tested. Moreover, the correlation between s_p^{opt} and $s_{R^2}^{\text{max}}$ seems to also increase with problem size, as demonstrated in Fig. 19, which shows the variation between different instances decreases with problem size, and seems to suggest that $s_p^{\text{opt}} \approx s_{R^2}^{\text{max}}$, for large N (i.e., the optimal pause point and the pause location for which the best fit to a Boltzmann distribution is observed, coincide for large problems).

In the next section we will sum up all of our results, and provide an interpretation of these slightly muddling findings.




FIG. 18. Quantifying accuracy of linear regression using R^2 for $\ln \frac{p_i}{g_i}$ as a function of E_i (as demonstrated in Fig. 17) for a single 500 qubit instance (\mathcal{I}_{500}). For each data point shown, we obtain a least-squares fitting from the distribution returned by D-Wave, from which we calculate the R^2 value. The solid line (red, right-axis) is the average energy returned from the D-Wave. We see the peak in R^2 corresponds to the region around s_p^{opt} (just after $s_p = 0.4$). Each data point is obtained from 10000 anneals with 5 choices of gauge, with $t_a = 1\mu\text{s}$, and $t_p = 100\mu\text{s}$.




FIG. 19. Average difference between the pause point s_p for which the maximal R^2 value is found (i.e., the closest to a Boltzmann distribution), and the optimal pause point, with problem size. The problems are defined on square sub-graphs of the full D-Wave chimera with side-length SL . The corresponding number of qubits is given in the legend. The data points are an average of (at least) 55 instances each. Error bars represent standard deviation. To collect our data we used 10000 anneals with 5 choices of gauge, with $t_a = 1\mu\text{s}$, and $t_p = 100\mu\text{s}$.

III. SUMMARY OF RESULTS

By using a simple adaptation to the standard annealing schedule, in which we pause the anneal for a certain time, we studied key mechanisms affecting the output of experimental quantum annealing devices.

We observed that there is very often a critical region during the anneal for which if one pauses, even for a relatively short time (say $10\mu\text{s}$), a drastically different distribution of solutions is returned from the device. In particular, it seems to more effectively sample low lying energy states in this region, resulting in a larger probability of success (probability of observing the ground state).

We are not able to explain this effect in the closed-system scenario, as discussed in Appendix A, and demonstrate that it is most likely related to thermalization, although we have raised several questions which require dedicated experiments to test further.

We studied two problem sets. The first set, small 12 qubit instances, with $J_{ij} \in [-1, 1]$, $h_i = 0$ (uniformly random), was used to demonstrate a correlation between the location of the minimum gap, and the location of the so-called optimal pause point (i.e., the point in the anneal for which a pause returns the largest success probability, or lowest average energy of the configurations). It was observed that the best place to pause the system is some time after the minimum gap (numerically we found for these instances, the best place was around 10-15% of the total anneal time later). This observation is in agreement with other experiments and theory [17, 26] which suggests that after the minimum gap – where thermal excitations may allow a significant fraction of the population to leave the ground state – population can begin to thermally relax back into the ground state. We mention this picture also explains the sharp-peaked nature of our observations; if one pauses just a little too late, since the transition rates depend exponentially on the size of the instantaneous energy gaps [17], they quickly drop off, and the pause time becomes too small for effective thermalization to occur. The exact nature of this depends heavily on the spectrum of each individual instance, since this determines the transition rates.

For a subset of these problems, we compared the distribution returned from the device with a pause at s_p^{opt} to a quantum Boltzmann distribution of the form $\exp(-\beta H(s))$, treating $T = k_B/\beta$ and s as fitting parameters. We found that the optimal values of (s, T) were such that s was actually significantly later (e.g. 20% of t_a) in the schedule compared to s_p^{opt} , and that the temperature fluctuated wildly between instances, and was much higher than the physical temperature (up to 4 times higher). The interpretation of this is not quite clear, since, according to the thermalization picture painted above, the region of the optimal pause point should correspond approximately to the freezing region, after which transitions between energy levels stop. Moreover, the fluctuating temperatures do not seem consistent with the

physical temperature fluctuations on the device.

It is therefore still not clear whether the sampling is from a Boltzmann distribution of this form when pausing at the optimal pause point. We mention it is also not clear whether we would even expect it to, since, even ignoring possible analog errors which are known to effect the output of such devices, the minimum annealing rate of $1\mu s^{-1}$ in principle allows for the distribution to change non-negligibly from directly after the pause until measurement at the end of the anneal (even if the thermal time-scale after the pause is large, if $Q(s)$ is non-negligible, internal dynamics may still play a role). So whilst our data does seem to fit to a Boltzmann distribution, the fact that some dynamics still occur may cause the parameters of the fit to be slightly different from what one would otherwise expect, making it tricky to correctly interpret. It is clear more work is required to explain these observations.

With the advent of a new entropic sampling algorithm [27], we were able to obtain accurate degeneracy estimates for large – 500 qubit – problems of the planted-solution type. This allowed us to compare the samples returned from the device to a classical Boltzmann distribution (i.e., arising from thermal sampling of the problem Hamiltonian). Remarkably, we found that for all samples tested, that the data does fit extremely well to a Boltzmann distribution of the problem Hamiltonian, with R^2 values up to 0.9999 (and always larger than 0.97). Moreover, we observed that for all problems, the distribution returned by an anneal with a pause at the optimal pause point corresponds to the best fit to a Boltzmann distribution, indicating these samples are thermalizing more completely.

We also provide evidence which suggests that the larger a problem is, the more likely it is to have a ‘good’ fit to a classical Boltzmann distribution. According to the thermalization picture presented in Sect. IC, this suggests that larger problems in fact freeze-out later in the anneal.

This is an intriguing result, and goes beyond what was found in Ref. [9], where it was shown that if one attempts to fit the data to a Boltzmann distribution at the physical temperature of the device, impossibly early freeze-out points are observed. We find the same here, however, if one reverses this argument, and fits the data, leaving the temperature as a free parameter, we find exceptionally close fits to a Boltzmann distribution (which get better with problem size).

The difficulty however in interpreting this data is that the optimal pause point seems slightly too early in the anneal to observe a classical Boltzmann distribution (since $Q(s_p^{\text{opt}}) \approx 0.5$ for these problems), and moreover, the temperature was at least 1.5 times the physical temperature (however, the fluctuations in the fitted temperature were fairly consistent with those on the actual device). This also suggests there are likely non-negligible dynamics occurring after s_p^{opt} (i.e. up to a point when Q is smaller). Taking this picture would then imply the fitted

temperature is even larger than 1.5 times the physical.

Though we can explain to some extent the higher than observed temperatures (and in principle the fluctuations), some possibilities of which include 1) the measured temperature of the device not corresponding to the actual temperature of the qubits, 2) the temperature of the device changing even during a single anneal, 3) analog errors (programming errors in the J_{ij}) which masquerade as a higher observed temperature, it is hard to reconcile, without further analysis, the fact that we seem to observe data which fits to a Boltzmann distribution, but one at a much later point than the optimal pause point.

As we have stressed previously, it would be extremely useful to be able to perform a more adequate quenching of the dynamics, that is, allowing for an annealing rate of much quicker than $\frac{ds}{dt} = 1\mu s^{-1}$. Moreover, it would be useful to have more accurate temperature data (indeed, the temperature data is sampled just once every few hours from the device meaning we can not perform a ‘real-time’ analysis).

With this in mind we summarize by noting that our results do indicate thermalization is occurring effectively in the region of the optimal pause point, however without addressing the two above points, it is challenging to fully consolidate this picture with the freeze-out conjecture (e.g. from Fig. 4).

IV. CONCLUSION AND OUTLOOK

The work presented here has several practical implications for quantum annealing. Firstly, it is clear that one may be able to design more effective annealing schedules by including a short pause at some intermediate point in the anneal, possibly allowing for orders of magnitude improvement in sampling from the low lying energy states (i.e., those which are of interest from the point of optimization). We demonstrated that across a particular problem set (even for different problem sizes), the optimal pause point does not seem to change much and that the size of the region in which one should pause remains fairly constant, hence good general heuristics may be derived.

In fact, this type of annealing schedule may allow us to actually learn fundamental properties about the nature of certain problem classes, such as the location of the minimum gap, which is extremely hard to determine numerically.

By varying the pause point (at least for ‘large’ problems of the planted solution type), one can effectively vary the temperature at which one samples from the classical Boltzmann distribution. This has implications for machine learning (for example in the use of restricted Boltzmann machines) [5, 6, 8]. Of course, one would need to perform more analysis for different problem types before this type of sampling can be understood in general.

Having said this, more work is needed to pin down the exact cause of the phenomena reported here. In particu-

lar there are two main lingering points. 1) We would like to more precisely understand the distribution the device is sampling from at the optimal pause point. To do this, one would need to be able to more effectively quench the system; that is, measure the device exactly at an intermediate point in the anneal, instead of having to finish the anneal in a relatively long time of $O(1\mu s)$, during which non-trivial dynamics can still occur. This is a hardware limitation imposed on the current device we performed our experiments on. 2) It is critical to understand the role of temperature in these devices. On the one hand, the device does seem to be sampling from a Boltzmann distribution (with a pause), however the temperature at which the samples are returned from are higher than the physical temperature of the device. Understanding this will allow more effective quantum devices (not just annealers) to be built, as well as allow for more effective sampling.

V. DATA ACQUISITION

For all instances labeled as \mathcal{I}_N^s (where N is the qubit number, and s a serial number when applicable), we pro-

vide the problem instance itself as part of the ancillary files associated with this article.

ACKNOWLEDGMENTS

We thank Lev Barash for computing the degeneracies of the 500 qubit problems in Sect. II E. E.G.R. and D.V. would like to acknowledge support from the NASA Transformative Aeronautic Concepts program and the NASA Ames Research Center. D.V. was supported by NASA Academic Mission Services, contract number NNA16BD14C. E.G.R. and D.V. were also supported in part by the AFRL Information Directorate under grant F4HBKC4162G001 and the Office of the Director of National Intelligence (ODNI) and the Intelligence Advanced Research Projects Activity (IARPA), via IAA 145483. 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 ODNI, IARPA, AFRL, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purpose notwithstanding any copyright annotation thereon.

-
- [1] T. Albash and D. A. Lidar, “Demonstration of a Scaling Advantage for a Quantum Annealer over Simulated Annealing,” *Phys. Rev. X* **8**, 031016 (2018).
- [2] V. S. Denchev, S. Boixo, S. V. Isakov, N. Ding, R. Babbush, V. Smelyanskiy, J. Martinis, and H. Neven, “What is the Computational Value of Finite-Range Tunneling?” *Phys. Rev. X* **6**, 031015 (2016).
- [3] T. F. Ronnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer, “Defining and detecting quantum speedup,” *Science* **345**, 420 (2014).
- [4] M. H. Amin, “Searching for quantum speedup in quasistatic quantum annealers,” *Phys. Rev. A* **92**, 052323 (2015).
- [5] M. H. Amin, E. Andriyash, J. Rolfe, B. Kulchitsky, and R. Melko, “Quantum Boltzmann Machine,” *Phys. Rev. X* **8**, 021050 (2018).
- [6] M. Benedetti, J. Realpe-Gómez, R. Biswas, and A. Perdomo-Ortiz, “Estimation of effective temperatures in quantum annealers for sampling applications: A case study with possible applications in deep learning,” *Phys. Rev. A* **94**, 022308 (2016).
- [7] B. H. Zhang, G. Wagenbreth, V. Martin-Mayor, and I. Hen, “Advantages of Unfair Quantum Ground-State Sampling,” *Scientific Reports* **7**, 1044 (2017).
- [8] S. H. Adachi and M. P. Henderson, “Application of Quantum Annealing to Training of Deep Neural Networks,” ArXiv e-prints (2015), [arXiv:1510.06356 \[quant-ph\]](https://arxiv.org/abs/1510.06356).
- [9] J. Marshall, E. G. Rieffel, and I. Hen, “Thermalization, freeze-out, and noise: Deciphering experimental quantum annealers,” *Phys. Rev. Applied* **8**, 064025 (2017).
- [10] A. D. King, J. Carrasquilla, J. Raymond, I. Ozfidan, E. Andriyash, A. Berkley, M. Reis, T. Lanting, R. Harris, F. Altomare, K. Boothby, P. I. Bunyk, C. Enderud, A. Fréchet, E. Hoskinson, N. Ladizinsky, T. Oh, G. Poulin-Lamarre, C. Rich, Y. Sato, A. Y. Smirnov, L. J. Swenson, M. H. Volkmann, J. Whittaker, J. Yao, E. Ladizinsky, M. W. Johnson, J. Hilton, and M. H. Amin, “Observation of topological phenomena in a programmable lattice of 1,800 qubits,” *Nature* **560**, 456 (2018).
- [11] M. Ohkuwa, H. Nishimori, and D. A. Lidar, “Reverse annealing for the fully connected p -spin model,” ArXiv e-prints (2018), [arXiv:1806.02542 \[quant-ph\]](https://arxiv.org/abs/1806.02542).
- [12] D. Ottaviani and A. Amendola, “Low Rank Non-Negative Matrix Factorization with D-Wave 2000Q,” *arXiv preprint arXiv:1808.08721* (2018).
- [13] D. Venturelli and A. Kondratyev, “Reverse Quantum Annealing Approach to Portfolio Optimization Problems,” unpublished (ACQ18) (2018).
- [14] D-Wave Technical Whitepaper, *Reverse Quantum Annealing for Local Refinement of Solutions* (2018), [dwaveReverse, https://www.dwavesys.com/sites/default/files/14-1018A-A_Reverse_Quantum_Annealing_for_Local_Refinement_of_Solutions.pdf](https://www.dwavesys.com/sites/default/files/14-1018A-A_Reverse_Quantum_Annealing_for_Local_Refinement_of_Solutions.pdf).
- [15] T. Lanting, A. D. King, B. Evert, and E. Hoskinson, “Experimental demonstration of perturbative anticrossing mitigation using nonuniform driver Hamiltonians,” *Phys. Rev. A* **96**, 042322 (2017).
- [16] J. I. Adame and P. L. McMahon, “Inhomogeneous driving in quantum annealers can result in orders-of-magnitude improvements in performance,” *arXiv preprint arXiv:1806.11091* (2018).
- [17] T. Albash, S. Boixo, D. A. Lidar, and P. Zanardi, “Quantum adiabatic Markovian master equations,” *New Journal of Physics* **14**, 123016 (2012).

- [18] V. N. Smelyanskiy, D. Venturelli, A. Perdomo-Ortiz, S. Knysh, and M. I. Dykman, “Quantum annealing via environment-mediated quantum diffusion,” *Phys. Rev. Lett.* **118**, 066802 (2017).
- [19] S. Boixo, V. N. Smelyanskiy, A. Shabani, S.V. Isakov, M. Dykman, V. S. Denchev, M. H. Amin, A. Y. Smirnov, M. Mohseni, and H. Neven, “Computational multiqubit tunnelling in programmable quantum annealers,” *Nat. Commun.* **7**, 10327 (2016).
- [20] A. J. Leggett, S. Chakravarty, A. T. Dorsey, Matthew P. A. Fisher, Anupam Garg, and W. Zwerger, “Dynamics of the dissipative two-state system,” *Rev. Mod. Phys.* **59**, 1 (1987).
- [21] R. Hanson, V. V. Dobrovitski, A. E. Feiguin, O. Gywat, and D. D. Awschalom, “Coherent Dynamics of a Single Spin Interacting with an Adjustable Spin Bath,” *Science* **320**, 352 (2008).
- [22] (), The transition time t_r shown in Fig. 4 is calculated using the spectrum for a 12 qubit problem instance \mathcal{T}_{12}^0 (spectrum shown in Fig. 21 of Appendix D). In this example, we do not calculate the full transition rates, but just use $t_r(s) \propto 1 - \exp(-\beta\omega(s))$ for demonstrative purposes. Here $\omega(s) = E_1(s) - E_0(s)$ and $\beta = 1/k_B T$ with $T = 12.1\text{mK}$. The other time scales shown $t_{a,p}$ are arbitrarily chosen.
- [23] 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).
- [24] (), Throughout this work the optimal pause point s_p^{opt} is evaluated by finding the lowest average energy (from many thousands of anneals) returned from the D-Wave device. In principle one could use the ground state success probability as a metric, however this is often not practical since 1) some hard problems are not solved even once by the D-Wave, and 2) for randomly generated instances (i.e., not of the planted-solution type), the ground state energy is typically unknown.
- [25] For a given problem, the later one can thermalize, the greater the ground state success probability, since the gap ω_{01} opens up so e.g. the ratio $P_1/P_0 = \exp(-\beta\omega_{01})$ is reduced.
- [26] N. G. Dickson, M. W. Johnson, M. H. Amin, R. Harris, F. Altomare, A. J. Berkley, P. Bunyk, J. Cai, E. M. Chapple, P. Chavez, F. Cioata, T. Cirip, P. deBuen, M. Drew-Brook, C. Enderud, S. Gildert, F. Hamze, J. P. Hilton, E. Hoskinson, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Lanting, T. Mahon, R. Neufeld, T. Oh, I. Perminov, C. Petroff, A. Przybysz, C. Rich, P. Spear, A. Tcaciuc, M. C. Thom, E. Tolkacheva, S. Uchaikin, J. Wang, A. B. Wilson, Z. Merali, and G. Rose, “Thermally assisted quantum annealing of a 16-qubit problem,” *Nat. Commun.* **4**, 1903 (2013).
- [27] L. Barash, J. Marshall, M. Weigel, and I. Hen, “Estimating the Density of States of Frustrated Spin Systems,” ArXiv e-prints (2018), [arXiv:1808.04340 \[quant-ph\]](https://arxiv.org/abs/1808.04340).
- [28] F. Wang and D. P. Landau, “Efficient, Multiple-Range Random Walk Algorithm to Calculate the Density of States,” *Phys. Rev. Lett.* **86**, 2050 (2001).
- [29] F. Wang and D. P. Landau, “Determining the density of states for classical statistical models: A random walk algorithm to produce a flat histogram,” *Phys. Rev. E* **64**, 056101 (2001).

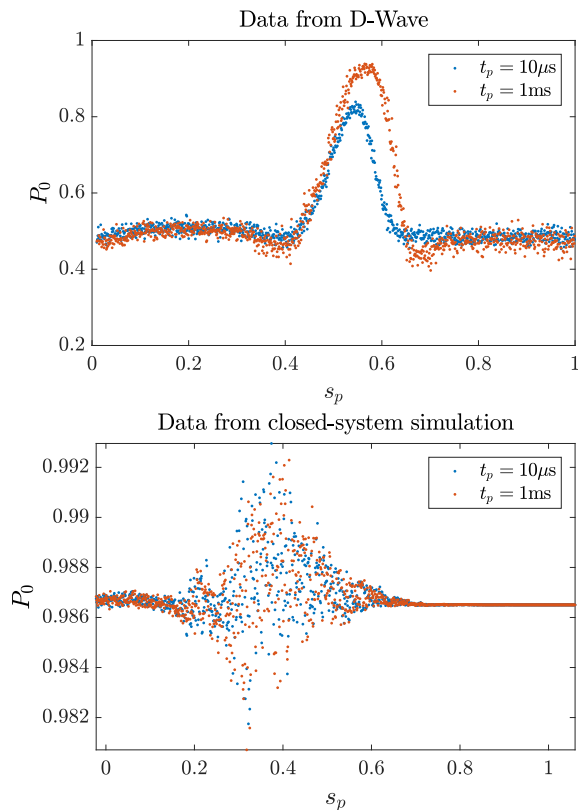


FIG. 20. Comparison between results from the experimental D-Wave device (top), and a closed-system Schrödinger evolution (bottom) for a single 12 qubit problem instance (\mathcal{I}_{12}^0). Both plots show the success probability P_0 against the pause point, for two different pause times ($t_p = 10, 1000\mu\text{s}$) as shown by the legend. The D-Wave data is from 10000 annealing runs, using 5 different gauges. Each plot contains 1000 data points evenly distributed in $s_p \in [0, 1]$. Both have an annealing time of $t_a = 1\mu\text{s}$ (in addition to the pause time). The simulation uses 1000 time steps in the calculation of the evolution operator.

Appendix A: Expected closed-system dynamics

We consider the effect of an intermediate pause under the closed system (Schrödinger) evolution alone, studying small 12 qubit problems, with $J_{ij} \in [-1, 1]$ (uniformly random) and $h_i = 0$.

It is interesting to note that the effect of pausing the anneal does have a noticeable effect, as demonstrated in Fig. 20, even in the closed system case. We believe this is essentially caused by Rabi oscillations during the pause, and we note it does not match the observed output from the D-Wave device. We explain below.

We consider the three regions in Fig. 20 (bottom). 1) During the evolution, when $s < 0.2$ the state is almost entirely in the instantaneous ground state, $|\psi(s)\rangle \approx |E_0(s)\rangle$. Thus, when the system is paused, and evolved under $H(s_p)$, very little happens since just an overall global phase is acquired. 2) A little later on, when the




FIG. 21. Spectrum (lowest 10 energy levels) of the problem considered in Figs. 7, 8, 11, 20 (\mathcal{I}_{12}^0). This problem has a small, well defined minimum gap of 0.15 GHz located at $s = 0.44$. The units are defined with $h = 1$.

energy gap starts to close between $s \in [0.2, 0.4]$ (see Fig. 21), diabatic transitions to excited energy levels may occur. Once a non-negligible amount of the population has been transferred to excited states, a pause will give rise to Rabi oscillations between the eigen-states of H_p , hence directly affecting the success probability at the end of the anneal. 3) Late in the anneal, after around $s = 0.7$, the driver Hamiltonian is essentially negligible, so can not drive any transitions between energy levels, hence a pause, will only change the relative phases of eigen-states of H_p , but not affect the probabilities upon measurement in the computational basis.

We describe three fundamental differences between the simulation, and the results from the experimental device (in addition to the large difference in success probability). Firstly it is evident there is much less structure in the closed system case; although the success probability does seem to increase on average, there is much more variability. This is due to the sensitivity of the period of the Rabi oscillations to the energy gaps (and hence to the location of the pause s_p). Second, there is seemingly no qualitative difference between a short and long pause in the closed system case, as compared to the observed phenomena which has an increasing peak with pause time. We believe this is due again to the nature of the Rabi oscillations. If the gap between energy levels is of the order of 1 GHz (see Fig. 21), the time-scale of the Rabi oscillations is much shorter than the pause times considered here (e.g. 1ns compared with $100\mu\text{s}$). Lastly, the location at which this Rabi dynamics occurs does not correspond precisely to that observed in the experiment (it is seemingly shifted slightly earlier in the schedule).

We also plot both data sets on the same axis in Fig. 22 for reference. This shows the Rabi induced oscillations are effectively negligible compared to the effect observed




FIG. 22. Same as Fig. 20, but with both data sets on the same plot. Both data sets correspond to $t_p = 1000\mu\text{s}$, and $t_a = 1\mu\text{s}$.

on the physical device.

Appendix B: Computing the quantum Boltzmann distribution

We wish to compare the distribution returned from the D-Wave device, i.e., the probability that a configuration with energy E is returned $P_{\text{DW}}(E)$, to what would be expected if the device were instantaneously thermalizing to the quantum Boltzmann distribution $\rho(s, T) := \frac{1}{Z} e^{-\beta H(s)}$, where $Z = \text{Tr} e^{-\beta H(s)}$, and $\beta = 1/k_B T$ with T a temperature parameter.

We note that the D-Wave device can only measure in the computational (z) basis, thus to compare the probability distributions we compute

$$P_{\text{QBM}}^{(s, T)}(E) = \sum_{z: H_p|z\rangle = E|z\rangle} \langle z | \rho(s, T) | z \rangle \quad (\text{B1})$$

for $s \in [0, 1]$ (steps of 0.01), and $T \in [\frac{1}{4}, 4]T_{\text{DW}}$ (steps of $\frac{1}{4}T_{\text{DW}}$).

Appendix C: Computing the classical Boltzmann distribution for large problems

In Sect. II E we analysed planted-solution type problems [23], of four different problem sizes, $N \in \{31, 125, 282, 501\}$. These were (a subset of) the same instances as studied in Ref. [9], and are generated on subgraphs of the full chimera of side-length $\text{SL} \in \{2, 4, 6, 8\}$ respectively. Each N group tested consisted of (at least) 55 problem instances, with a random number of sub-Hamiltonian loops chosen (as described in more detail in Ref. [9]).




FIG. 23. The D-Wave 2000Q ‘chimera’ graph which we conducted our experiments on (the device is housed at NASA Ames Research Center). There are 16×16 unit cells each containing 8 qubits. Dead (malfunctioning) qubits are not shown on the graph. Top left bordered in red is an example of a square subgraph of side-length $\text{SL}=4$.

The benefit to using this problem type is that one knows in advance the spectrum of H_p , and one can calculate exactly the degeneracy of the ground and first excited states. The description of this algorithm is outlined in Ref. [7].

Knowledge of the exact ground and first excited state degeneracies is extremely powerful as it allows one to verify any estimated degeneracy values from entropic sampling techniques (such as the well known Wang-Landau method [28, 29]).

We used a newly devised algorithm for estimating the density of states based on population-annealing [27] to obtain accurate estimates of the degeneracies for the largest instances tested (501 qubits), for which traditional (e.g. Wang-Landau) approaches failed. Here ‘accurate’ implies neither the ground nor first excited state degeneracy estimate differed by more than 5% of the exact values. For the 125 and 282 qubit problems we were able to use the Wang-Landau algorithm to obtain accurate estimates of the degeneracies. For the 31 qubit instances we used exact enumeration to compute the degeneracies.

Appendix D: Supplemental figures

In Fig. 23 we show the full (working) D-Wave 2000Q hardware graph. All of our experiments were conducted on this graph.




FIG. 24. Success probability P_0 (under the default annealing schedule with $t_a = 1\mu\text{s}$) as a function of minimum gap Δ_{\min} for the 100 problem instances of size 12 qubits reported on in the main text. We also plot the operating temperature of the device (black-dash line). The units are defined with $\hbar = 1$. The data is from 10000 anneals with 10 choices of gauge.

In Sect. II C we analysed 100 problems of size 12 qubits

for which $J_{ij} \in [-1, 1]$ (uniformly random), and $h_i = 0$. In Fig. 24 we show how the success probability of these problems depends on the minimum gap Δ_{\min} . For problems for which Δ_{\min} is larger than around 1GHz, the problems are solved with nearly 100% success probability.

In the main text (Sect. II A), one effect we studied was varying the total anneal time t_a , but keeping the pause time t_p constant (e.g. Fig. 8). Here, in Fig. 25 we show the corresponding heat map (i.e. where t_p is also varied for each choice of t_a). One sees that the pause is essentially an efficient way of increasing the success probability (lowering average energy) without increasing the anneal time; notice for a short anneal time, $t_a = 1\mu\text{s}$, with a pause of around $20\mu\text{s}$ at $s \approx 0.4$ gives approximately the same average energy as an anneal for time of 1ms.

One interesting observation in Sect. II C was that by dividing the energy scale of the problem Hamiltonian, $H_p \rightarrow H_p/C$ for (e.g.) $C = 1, 2, 4, 8$, was that the peak in the success probability shifts to later in the anneal. This was partly due to the minimum gap which shifts to later in the anneal (see Fig. 26), but we also related this in the main text to the diminishing quantum fluctuations Q . In Fig. 27 we show the corresponding heat map for a single problem instance upon dividing the problem energy scale.

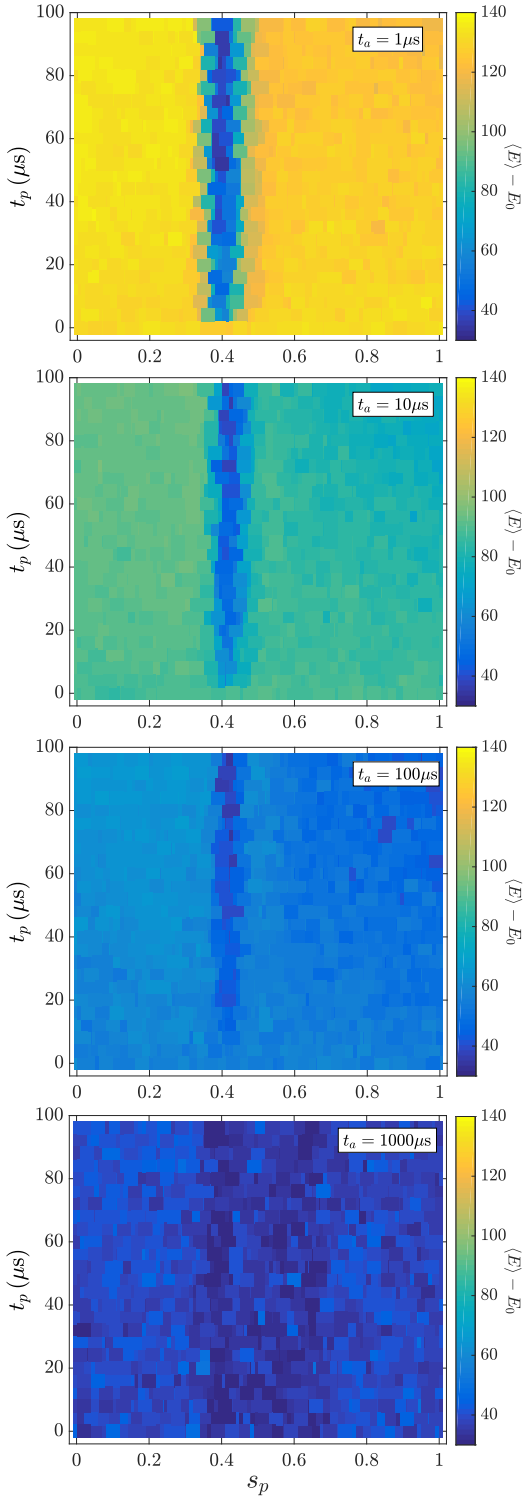


FIG. 25. Effect of changing the total annealing time (not including the pause time), t_a , for a 500 qubit planted problem instance (\mathcal{I}_{500}^g). The heat map color corresponds to the average energy, $\langle E \rangle - E_0$ (arbitrary units) returned from the annealer. From top to bottom the total annealing time $t_a = 1, 10, 100, 1000 \mu\text{s}$ (see legend). In the bottom figure, with longest annealing time, the pause has little to no effect. Each data point is averaged from 5000 anneals with 5 different choice of gauge.




FIG. 26. Changing of the spectral properties upon re-scaling the problem. This is the same 12 qubit problem studied in Fig. 27 (\mathcal{I}_{12}^0). The location of the minimum gap changes as $s_{\text{gap}} = [0.438, 0.508, 0.558, 0.608]$, and the minimum gap itself changes accordingly as $\Delta_{\text{min}} = [0.15, 0.10, 0.06, 0.03]$ GHz, when the problem Hamiltonian is re-scaled by $[1, 2, 4, 8]$ respectively (see legend). Energy units defined via $\hbar = 1$.




FIG. 27. Success probability P_0 heat map for a single 12 qubit instance (\mathcal{I}_{12}^0) where the annealing schedule has a pause of length t_p inserted at s_p , where from top to bottom the problem Hamiltonian H_p has been re-scaled by a factor of 1,2,4,8 (that is, $H_p \rightarrow H_p/C$, where $C = 1, 2, 4, 8$). See Fig. 26 for the corresponding minimum gap plot. Each data point is an average from 10000 anneals with 5 gauges. Notice the change in the color bar scale between the different images.