



CHORUS

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

## Temperature Scaling Law for Quantum Annealing Optimizers

Tameem Albash, Victor Martin-Mayor, and Itay Hen

Phys. Rev. Lett. **119**, 110502 — Published 14 September 2017

DOI: [10.1103/PhysRevLett.119.110502](https://doi.org/10.1103/PhysRevLett.119.110502)

# Temperature scaling law for quantum annealing optimizers

Tameem Albash,<sup>1,2</sup> Victor Martin-Mayor,<sup>3,4</sup> and Itay Hen<sup>1,2</sup>

<sup>1</sup>*Information Sciences Institute, University of Southern California, Marina del Rey, California 90292, USA*

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

<sup>3</sup>*Departamento de Física Teórica I, Universidad Complutense, 28040 Madrid, Spain*

<sup>4</sup>*Instituto de Biocomputación y Física de Sistemas Complejos (BIFI), Zaragoza, Spain*

Physical implementations of quantum annealing unavoidably operate at finite temperatures. We point to a fundamental limitation of fixed finite temperature quantum annealers that prevents them from functioning as competitive scalable optimizers and show that to serve as optimizers annealer temperatures must be appropriately scaled down with problem size. We derive a temperature scaling law dictating that temperature must drop at the very least in a logarithmic manner but also possibly as a power law with problem size. We corroborate our results by experiment and simulations and discuss the implications of these to practical annealers.

**Introduction.**— Quantum computing devices are becoming sufficiently large to undertake computational tasks that are infeasible using classical computing [1–7]. The theoretical underpinning for whether such tasks exist with physically realizable quantum annealers remains lacking, despite the excitement brought on by recent technological breakthroughs that have made programmable quantum annealing (QA) [8–12] optimizers consisting of thousands of quantum bits commercially available. Thus far, no examples of practical relevance have been found to indicate a superiority of QA optimization, i.e., to find bit assignments that minimize the energy, or cost, of discrete combinatorial optimization problems, faster than possible classically [13–20]. Major ongoing efforts continue to build larger, more densely connected QA devices, in the hope that the capability to embed larger optimization problems would eventually reveal the coveted quantum speedup [21–25].

Understanding the robustness of QA optimization to errors that reduce the final ground state probability is critical. In this work, we consider perhaps the most optimistic setting where the only source of error is due to nonzero temperature. We analyze the theoretical scaling performance of *ideal fixed-temperature* quantum annealers for optimization. We show that even in the case where annealers are assumed to *thermalize instantly* (rather than only in the infinite runtime limit), the energies, or costs, of their output configurations would be computationally trivial to achieve (in a sense that we explain). We further derive a scaling law for QA optimizers and provide corroboration of our analytical findings by experimental results obtained from the commercial D-Wave 2X QA processor [26–30] as well as numerical simulations (our results equally apply to ideal thermal annealing devices). We discuss the implications of our results for both past benchmarking studies and for the engineering requirements of future QA devices.

**Fixed-temperature quantum annealers.**— In the adiabatic limit, closed-system quantum annealers are guaranteed to find a ground state of the target cost func-

tion, or final Hamiltonian  $H$ , they are to solve. The adiabatic theorem of quantum mechanics ensures that the overlap of the final state of the system with the ground state manifold of  $H$ , approaches unity as the duration of the process increases [31, 32]. For physical quantum annealers that operate at positive temperatures ( $T > 0$ ), there is no equivalent guarantee of reaching the ground state with high probability. For long runtimes, an ideal finite-temperature quantum annealer is expected to sample the Boltzmann distribution of the final Hamiltonian at the annealer temperature [33].

In what follows, we argue that even instantly-thermalizing quantum annealers [34] are severely limited as optimizers due to their finite temperature. For concreteness, we restrict to annealers for which i) the number of couplers scales linearly with the number of qubits  $N$  [35], ii) the coupling strengths are discretized and are bounded independently of problem size, and iii) the scaling of the free energy with problem size is not pathological, i.e., that our system is not tuned to a critical point. Other than the above standard assumptions, our treatment is general (we discuss the performance of quantum annealers when some of these conditions are lifted later on). For clarity, we consider optimization problems written in terms of a Hamiltonian of the Ising-type

$$H = \sum_{\langle ij \rangle} J_{ij} s_i s_j + \sum_i h_i s_i, \quad (1)$$

where  $\{s_i = \pm 1\}$  are binary Ising spin variables that are to be optimized over,  $\{J_{ij}, h_i\}$  are the coupling strengths between connected spins and external biases, respectively, and  $\langle ij \rangle$  denotes the underlying connectivity graph of the model. The discussion that follows however is not restricted to any particular model.

Under the above assumptions, the ground state energies, denoted  $E_0$ , of any given problem class, scale linearly with increasing problem size (i.e., the energy is an extensive property as is generically expected from physical systems) while the classical minimal gap  $\Delta = E_1 - E_0$

remains fixed. It follows then [36] that the thermal expectation values of the intensive energy

$$\langle e \rangle_\beta = \langle H \rangle_\beta / N, \quad (2)$$

and specific heat

$$c_\beta = \partial \langle e \rangle_\beta / \partial \beta = -N [\langle e^2 \rangle_\beta - \langle e \rangle_\beta^2], \quad (3)$$

remain finite as  $N \rightarrow \infty$  for any fixed inverse-temperature  $\beta = 1/T$ . The intensive energy is discretized in steps of  $\Delta/N$ , yet its statistical dispersion  $\sigma_\beta(e) = \sqrt{-c_\beta/N}$  is much larger. Treating  $e$  as a stochastic variable, for large enough values of  $N$  it can be treated as a continuous variable as the ratio of discretization versus dispersion is  $\sqrt{-\Delta^2/(c_\beta N)}$  decaying to zero for large  $N$ . From the Boltzmann distribution it follows that the probability density of  $e$  goes as  $p_\beta(e) = Z_\beta^{-1} e^{N(s(e) - \beta e)}$ , where  $Z_\beta = \sum_n g_n e^{-\beta E_n}$  is the partition function,  $g_n$  is the degeneracy of the  $n$ -th level, i.e., the number of microstates with  $H(\{s_i\}) = E_n$ , satisfying  $2^n = \sum_{n \geq 0} g_n$ , and  $s(e)$  is the entropy density [37]. The linear combination  $\Psi_\beta(e) = s(e) - \beta e$  plays the role of a large-deviations functional for  $e$ . The most probable value of  $e$ , which we denote by  $e^*$ , is given by the maximum of  $\Psi_\beta$ . Solving  $\Psi'_\beta(e^*) = 0$ , we find [38]

$$\beta = \left. \frac{\partial s}{\partial e} \right|_{e=e^*}. \quad (4)$$

Close to  $e^*$ ,  $\Psi_\beta$  can be Taylor-expanded as  $\Psi_\beta(e) \approx \Psi_\beta(e^*) - \frac{|\Psi''_\beta(e^*)|}{2}(e - e^*)^2$ , from which it follows that

$$p_\beta(e) \approx \frac{e^{N\Psi_\beta(e^*)}}{Z_\beta} \exp \left[ -\frac{N|\Psi''_\beta(e^*)|}{2}(e - e^*)^2 \right]. \quad (5)$$

The probability density is thus approximately Gaussian in the vicinity of  $e^*$ , although deviations from the Gaussian behavior are crucial [39]. Moreover, in the limit of large  $N$ , we find

$$\langle e \rangle_\beta = e^* \quad \text{and} \quad c_\beta = \frac{-1}{|\Psi''_\beta(e^*)|}. \quad (6)$$

Therefore, the probability of finding by Boltzmann-sampling any energy  $e < e^*$  (equivalently,  $E < e^*N$ ) is exponentially suppressed in  $N$ , scaling in fact as  $\exp[-N(\Psi_\beta(e^*) - \Psi_\beta(e))]$ . We thus arrive at the conclusion that even ideal fixed temperature quantum annealers that thermalize instantaneously to the Gibbs state of the classical Hamiltonian are *exponentially unlikely* to find the ground state since  $e^* > e_0 \equiv E_0/N$ .

We now corroborate the above derivation by runs on the commercial DW2X quantum annealer [26–29]. To do so, we first generate random instances of differently sized sub-graphs of the DW2X Chimera connectivity graph [40, 41] and run them multiple times on the annealer, recording the obtained energies [42]. Figure 1

depicts typical resultant residual energy ( $E - E_0$ ) distributions. As is evident, increasing the problem size  $N$  ‘pushes’ the energy distribution farther away from  $E_0$ , as well as broadening the distribution and making it more gaussian-like. In the inset, we measure the departure of  $\langle H \rangle_\beta$  from  $E_0$  and the spread of the energies  $\sigma_\beta(H)$  over 100 ‘planted-solution’ [18] instances per sub-graph size as a function of problem size  $N$  [43]. For sufficiently large problem sizes, we find that the scaling of  $\langle H - E_0 \rangle_\beta$  is close to linear while  $\sigma_\beta(H)$  scales slightly faster than  $\sqrt{N}$ . While the slight deviations from our analytical predictions suggest that the DW2X configurations have not fully reached asymptotic behavior [44], they exhibit a trend that closely matches our assumptions with the agreement getting better with growing problem sizes.




FIG. 1. **Distributions of residual energy,  $E - E_0$ , from DW2X runs.** As problem sizes grow, the distributions become more Gaussian-like. **Inset:** Gaussians’ mean (blue) and standard deviation (red) as a function of problem size, averaged over 100 instances per size. The solid lines correspond to power-law fits of the average mean with power  $0.98 \pm 0.14$  and average standard deviation scaling with power  $0.63 \pm 0.09$ , taking into accounts all sizes but the smallest ( $1.01 \pm 0.62$  and  $0.57 \pm 0.37$  respectively if the two smallest sizes are omitted).

Given the scaling of the mean and standard deviation, we conclude that fixed-temperature quantum annealers will generate energies  $e$  with a fixed distance from  $e_0$ , or in terms of extensive energies, configurations obtained from fixed-temperature annealers will have energies concentrated around  $E = (1 - \epsilon)E_0$  for some  $\epsilon > 0$  and  $E_0 < 0$ .

One could now ask what the difficulty is for *classical algorithms* to generate energy values in the above range. This question has been recently answered by the discovery of a polynomial time approximation scheme (PTAS) for spin-glasses defined on a Chimera graph [45] (and which can be easily generalized to any locally connected model), where reaching such energies can be done efficiently [46]. While the scaling of the PTAS with

$\epsilon$  is not favorable, scaling as  $c^{1/\epsilon}$  for some constant  $c$ , in practice there exist algorithms (e.g., parallel tempering that we discuss later on) that are known to scale more favorably than PTAS.

**Scaling law for quantum annealing temperatures.**— In light of the above, it may seem that quantum annealers are doomed to fail as optimizers as problem sizes increase. We now argue that success may be regained if the temperature of the QA device is appropriately scaled with problem size. Specifically, we address the question of how the inverse-temperature  $\beta$  should scale with  $N$  such that there is a probability of at least  $q$  of finding the ground state.

An estimate for the required scaling can be given as follows. From the above analysis, it should be clear that the probability of finding a ground state at inverse temperature  $\beta$  will not decay exponentially with system size only if the ground state falls within the variation of the mean energy, specifically if

$$\sigma_\beta(H) = N\sigma_\beta(e) = \sqrt{-Nc_\beta}, \quad (7)$$

is comparable to

$$\langle H \rangle_\beta - E_0 = -N \int_\beta^\infty d\beta c_\beta. \quad (8)$$

The third law of thermodynamics dictates that the specific heat  $c_T \equiv d\langle e \rangle / dT$  goes to zero when  $T \rightarrow 0$ . Assuming a scaling of the form  $c_T \sim T^\alpha$ , or equivalently,  $-c_\beta \sim \beta^{-\alpha-2}$ , gives

$$\sigma_\beta(H) \sim \sqrt{\frac{N}{\beta^{\alpha+2}}} \quad \text{and} \quad \langle H \rangle_\beta - E_0 = \frac{N}{\beta^{\alpha+1}}. \quad (9)$$

For a power-law specific heat, it thus follows that the sought scaling is  $\beta \sim N^{1/\alpha}$ . If on the other hand  $c_\beta$  vanishes exponentially in  $\beta$ , the inverse-temperature scaling will be milder, of the form  $\beta \sim \log N$ .

To illustrate the above, we next present an analysis of simulations of randomly generated instances on Chimera lattices (we study several problem classes and architectures, see the Supplemental Information). To study the energy distribution generated by a thermal sampler on these instances, we use parallel tempering (PT) [47, 48], a Monte Carlo method whereby multiple copies of the system at different temperatures are simulated [49]. In Fig. 2, we show an example of how the energy distribution of a planted-solution instance changes with  $\beta$ . The qualitative behavior is similar to what we observe with increasing problem size, whereby decreasing  $\beta$  (increasing the temperature) pushes the energy distribution to larger energies and makes it more gaussian-like.

The behavior of the specific heat  $c_\beta$  as the inverse-temperature  $\beta$  becomes large is shown in Fig. 3. At large sizes, the scaling becomes  $c_\beta \propto \exp(-\Delta\beta)$  as expected




FIG. 2. **Distributions of residual energy,  $E - E_0$ , from PT simulations.** For a planted-solution instance defined on an  $L = 12$  Chimera graph, the distributions become more Gaussian-like as  $\beta$  decreases. For the case of  $\beta = 0.75$ , the mean residual energy and standard deviation are indicated. **Inset:** Scaling with problem size of the median mean energy and median standard deviation of the energy for  $\beta = 1.47$  over 100 instances.

(here,  $\Delta = 4$  is the gap). Based on our predictions above, this should mean that if for a fixed  $q$ , the minimum  $\beta^*$  such that  $p_{\beta^*}(E_0) \geq q$  falls in this exponential regime, then we should observe a scaling  $\beta^* \propto \log N$ . Indeed, the inset of Fig. 3, which shows simulation results of  $\beta^*$  versus  $N$ , exhibits the expected  $\log N$  behavior [50].




FIG. 3. **Typical specific heat with inverse-temperature.** Behavior of the median specific heat (over 100 instances) for planted-solution instances with inverse-temperature  $\beta$  for  $N = 3872$ . The behavior transitions from a polynomial scaling with  $\beta$  to an exponential scaling. **Inset:** Typical minimum inverse-temperature required for instances of size  $N$  such that the probability of the target energy  $E_T = E_0 + \delta(N)$  is at least  $q = 10^{-1}$ . Also shown are fits to  $\log N$  for all three cases and a power-law fit to  $cN^\alpha$  that finds  $\alpha = 0.19 \pm 0.05$  for the  $\delta = 0$  case, which is almost indistinguishable from the logarithmic fit.

While for problem classes with a fixed minimum gap  $\Delta$ , one may naively expect  $c_\beta$  to vanish exponentially in general, implying that a logarithmic scaling of  $\beta$  will generally be sufficient as our simulations indeed indicate, it is important to note that two-dimensional spin glasses are known to exhibit a crossover between an exponential behavior to a power law [51–54]. This crossover is characterized by a constant  $\theta \approx 1/2$ , whereby the discreteness of the gap  $\Delta$  is evident only for sizes  $N^{\theta/2} \ll \beta$ . Beyond  $N^{\theta/2} \sim \beta$ , the  $2d$  system behaves as if the coupling distribution is continuous [52, 53] at which point the system can be treated as if with continuous couplings, for which the specific heat  $c_T$  scales as  $T^\alpha$  with  $\alpha_c = 2\nu$  [51], where  $\nu = 3.53(7)$  [54]. Therefore, for an ideal quantum annealer operating beyond the crossover, a scaling of  $\beta \sim N^{1/(2\nu) \approx 0.14}$  is required. We may thus expect the same crossover to appear for instances defined on the Chimera lattice, which is  $2d$ -like. Interestingly, for the temperature scaling shown in the inset of Fig. 3, a power-law fit  $\beta \sim N^\alpha$  with  $\alpha = 0.19 \pm 0.05$  is almost indistinguishable from the logarithmic one, with a power that is consistent with the  $2d$  prediction.

**Suboptimal metrics for optimization problems.**— For many classically intractable optimization problems, when formulated as Ising models, it is crucial that solvers find a true minimizing bit assignment rather than low lying excited states. This is especially true for NP-complete/hard problems [55] where sub-optimal costs generally correspond to violated constraints that *must* be satisfied (otherwise the resultant configuration is nonsensical despite its low energy). Nonetheless, it is plausible to assume the existence of problems for which slightly sub-optimal configurations would still be of value [56]. We thus also study the necessary temperature scaling for cases where the target energies obey  $E_T \leq E_0 + \delta(N)$  with  $\delta(N)$  scaling sub-linearly with problem size. In the inset of Fig. 3, we plot the required scaling of  $\beta$  for  $\delta(N) = \text{const}$  and  $\delta(N) \propto \sqrt{N}$ . In both cases we find that a logarithmic scaling is still essential, albeit with smaller prefactors.

**Conclusions and discussion.**— We have shown that fixed temperature quantum annealers can only sample ‘easily reachable’ energies in the large problem size limit, thereby posing fundamental limitation on their performance. We derived a temperature scaling law to ensure that quantum annealing optimizers find nontrivial energy values with sub-exponential probabilities. The scaling of the specific heat with temperature controls this scaling: if  $\beta$  lies in the regime where the specific heat scales exponentially with  $\beta$ , then the inverse-temperature of the annealer must scale as  $\log N$ . However, further considerations are needed because of a possible crossover behavior in the specific heat with temperature and problem size. For Chimera graphs, because of their essentially two-dimensional structure, this may lead to a crossover to power law scaling. Little is known about this crossover

in three dimensions or for different architectures, so this concern may not be mitigated by a more complex connectivity graph.

Our results shed important light on benchmarking studies that have found no quantum speedups [17, 18, 57–59], identifying temperature as a relevant culprit for their unfavorable performance. Our analysis is particularly relevant for both the utility as well as the design of future QA devices that have been argued to sample from thermal or close-to-thermal distributions [60], calling their role as optimization devices into question.

One approach to scaling down the temperature with problem size is the (theoretically) equivalent scaling up of the overall energy scale of the Hamiltonian. However, the rescaling of the total Hamiltonian is also known to be challenging and may not represent a convenient approach for a scalable architecture. An alternative approach is to develop quantum error correction techniques to effectively increase the energy scale of the Hamiltonian by coupling multiple qubits to form a single logical qubit [61–66] in conjunction with classical post-processing [67–70] or to effectively decouple the system from the environment [71–74].

Our results reiterate the need for fault-tolerant error correction for scalable quantum annealing, however they do not preclude the utility of quantum annealing optimizers for large finite size problems, where engineering challenges may be overcome to allow the device to operate effectively at a sufficiently low temperature such that problems of interest of a finite size may be solved even in the absence of fault-tolerance. Our results only indicate that this ‘window of opportunity’ cannot be expected to continue as devices are scaled without further improvements in the device temperature or energy scale.

While our arguments above indicate that fixed-temperature quantum annealers may not be scalable as optimizers, the current study does not pertain to the usage of quantum annealers as *samplers* [60, 75, 76], where the objective is to sample from the Boltzmann distribution. The latter objective is known to be very difficult task (it is #P-hard [77–79]) and little is known about when or if quantum annealers can provide an advantage in this regard [80].

**Acknowledgements.**— TA and IH thank Daniel Lidar for useful comments on the manuscript. The computing resources were provided by the USC Center for High Performance Computing and Communications. TA was supported under ARO MURI Grant No. W911NF-11-1-0268, ARO MURI Grant No. W911NF-15-1-0582, and NSF Grant No. INSPIRE- 1551064. V.M.-M. was partly supported by MINECO (Spain), through Grant No. FIS2015-65078-C2-1-P.

- 
- [1] Michael J. Bremner, Ashley Montanaro, and Dan J. Shepherd, “Achieving quantum supremacy with sparse and noisy commuting quantum computations,” *Quantum* **1**, 8 (2017).
- [2] Xun Gao, Sheng-Tao Wang, and L.-M. Duan, “Quantum supremacy for simulating a translation-invariant Ising spin model,” *Phys. Rev. Lett.* **118**, 040502 (2017).
- [3] Michael J. Bremner, Ashley Montanaro, and Dan J. Shepherd, “Average-case complexity versus approximate simulation of commuting quantum computations,” *Phys. Rev. Lett.* **117**, 080501 (2016).
- [4] Tomoyuki Morimae, Keisuke Fujii, and Joseph F. Fitzsimons, “Hardness of classically simulating the one-clean-qubit model,” *Phys. Rev. Lett.* **112**, 130502 (2014).
- [5] Matthew A. Broome, Alessandro Fedrizzi, Saleh Rahimi-Keshari, Justin Dove, Scott Aaronson, Timothy C. Ralph, and Andrew G. White, “Photonic boson sampling in a tunable circuit,” *Science* **339**, 794–798 (2013).
- [6] Justin B. Spring, Benjamin J. Metcalf, Peter C. Humphreys, W. Steven Kolthammer, Xian-Min Jin, Marco Barbieri, Animesh Datta, Nicholas Thomas-Peter, Nathan K. Langford, Dmytro Kundys, James C. Gates, Brian J. Smith, Peter G. R. Smith, and Ian A. Walmsley, “Boson sampling on a photonic chip,” *Science* **339**, 798–801 (2013).
- [7] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, “Characterizing Quantum Supremacy in Near-Term Devices,” ArXiv e-prints (2016), [arXiv:1608.00263 \[quant-ph\]](https://arxiv.org/abs/1608.00263).
- [8] A. B. Fimmel, M. A. Gomez, C. Sebenik, C. Stenson, and J. D. Doll, “Quantum annealing: A new method for minimizing multidimensional functions,” *Chemical Physics Letters* **219**, 343–348 (1994).
- [9] J. Brooke, D. Bitko, T. F., Rosenbaum, and G. Aeppli, “Quantum annealing of a disordered magnet,” *Science* **284**, 779–781 (1999).
- [10] Tadashi Kadowaki and Hidetoshi Nishimori, “Quantum annealing in the transverse Ising model,” *Phys. Rev. E* **58**, 5355 (1998).
- [11] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser, “Quantum Computation by Adiabatic Evolution,” [arXiv:quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106) (2000).
- [12] Giuseppe E. Santoro, Roman Martoňák, Erio Tosatti, and Roberto Car, “Theory of quantum annealing of an Ising spin glass,” *Science* **295**, 2427–2430 (2002).
- [13] A. P. Young, S. Knysh, and V. N. Smelyanskiy, “Size dependence of the minimum excitation gap in the quantum adiabatic algorithm,” *Phys. Rev. Lett.* **101**, 170503 (2008).
- [14] A. P. Young, S. Knysh, and V. N. Smelyanskiy, “First-order phase transition in the quantum adiabatic algorithm,” *Phys. Rev. Lett.* **104**, 020502 (2010).
- [15] Itay Hen and A. P. Young, “Exponential complexity of the quantum adiabatic algorithm for certain satisfiability problems,” *Phys. Rev. E* **84**, 061152 (2011).
- [16] E. Farhi, D. Gosset, I. Hen, A. W. Sandvik, P. Shor, A. P. Young, and F. Zamponi, “Performance of the quantum adiabatic algorithm on random instances of two optimization problems on regular hypergraphs,” *Phys. Rev. A* **86**, 052334 (2012), ([arXiv:1208.3757](https://arxiv.org/abs/1208.3757)).
- [17] Troels F. Rønnow, Zhihui Wang, Joshua Job, Sergio Boixo, Sergei V. Isakov, David Wecker, John M. Martinis, Daniel A. Lidar, and Matthias Troyer, “Defining and detecting quantum speedup,” *Science* **345**, 420–424 (2014).
- [18] Itay Hen, Joshua Job, Tameem Albash, Troels F. Rønnow, Matthias Troyer, and Daniel A. Lidar, “Probing for quantum speedup in spin-glass problems with planted solutions,” *Phys. Rev. A* **92**, 042325– (2015).
- [19] Sergio Boixo, Tameem Albash, Federico M. Spedalieri, Nicholas Chancellor, and Daniel A. Lidar, “Experimental signature of programmable quantum annealing,” *Nat. Commun.* **4**, 2067 (2013).
- [20] Tameem Albash, Walter Vinci, Anurag Mishra, Paul A. Warburton, and Daniel A. Lidar, “Consistency tests of classical and quantum models for a quantum annealer,” *Phys. Rev. A* **91**, 042314– (2015).
- [21] “Quantum enhanced optimization (qeo),” .
- [22] S. K. Tolpygo, V. Bolkhovskiy, T. J. Weir, L. M. Johnson, M. A. Gouker, and W. D. Oliver, “Fabrication process and properties of fully-planarized deep-submicron Nb/Al-AlOx/Nb Josephson junctions for VLSI circuits,” *IEEE Transactions on Applied Superconductivity* **25**, 1–12 (2015).
- [23] S. K. Tolpygo, V. Bolkhovskiy, T. J. Weir, C. J. Galbraith, L. M. Johnson, M. A. Gouker, and V. K. Semenov, “Inductance of circuit structures for MIT II superconductor electronics fabrication process with 8 niobium layers,” *IEEE Transactions on Applied Superconductivity* **25**, 1–5 (2015).
- [24] X. Y. Jin, A. Kamal, A. P. Sears, T. Gudmundsen, D. Hover, J. Miloshi, R. Slattery, F. Yan, J. Yoder, T. P. Orlando, S. Gustavsson, and W. D. Oliver, “Thermal and residual excited-state population in a 3D transmon qubit,” *Phys. Rev. Lett.* **114**, 240501 (2015).
- [25] “D-wave systems previews 2000-qubit quantum system,” .
- [26] M W Johnson, P Bunyk, F Maibaum, E Tolkacheva, A J Berkley, E M Chapple, R Harris, J Johansson, T Lanting, I Perminov, E Ladizinsky, T Oh, and G Rose, “A scalable control system for a superconducting adiabatic quantum optimization processor,” *Superconductor Science and Technology* **23**, 065004 (2010).
- [27] A J Berkley, M W Johnson, P Bunyk, R Harris, J Johansson, T Lanting, E Ladizinsky, E Tolkacheva, M H S Amin, and G Rose, “A scalable readout system for a superconducting adiabatic quantum optimization system,” *Superconductor Science and Technology* **23**, 105014 (2010).
- [28] R. Harris, M. W. Johnson, T. Lanting, A. J. Berkley, J. Johansson, P. Bunyk, E. Tolkacheva, E. Ladizinsky, N. Ladizinsky, T. Oh, F. Cioata, I. Perminov, P. Spear, C. Enderud, C. Rich, S. Uchaikin, M. C. Thom, E. M. Chapple, J. Wang, B. Wilson, M. H. S. Amin, N. Dickson, K. Karimi, B. Macready, C. J. S. Truncik, and G. Rose, “Experimental investigation of an eight-qubit unit cell in a superconducting optimization processor,” *Phys. Rev. B* **82**, 024511 (2010).
- [29] P. I. Bunyk, E. M. Hoskinson, M. W. Johnson, E. Tolkacheva, F. Altomare, A. J. Berkley, R. Harris, J. P. Hilton, T. Lanting, A. J. Przybysz, and J. Whittaker, “Architectural considerations in the design of a superconducting quantum annealing processor,” *IEEE Transactions on Applied Superconductivity* **24**, 1–10 (Aug. 2014).

- [30] J. King, S. Yarkoni, M. M. Nevisi, J. P. Hilton, and C. C. McGeoch, “Benchmarking a quantum annealing processor with the time-to-target metric,” ArXiv e-prints (2015), [arXiv:1508.05087 \[quant-ph\]](#).
- [31] T. Kato, “On the adiabatic theorem of quantum mechanics,” *J. Phys. Soc. Jap.* **5**, 435 (1950).
- [32] Sabine Jansen, Mary-Beth Ruskai, and Ruedi Seiler, “Bounds for the adiabatic approximation with applications to quantum computation,” *J. Math. Phys.* **48**, – (2007).
- [33] Lorenzo Campos Venuti, Tameem Albash, Daniel A. Lidar, and Paolo Zanardi, “Adiabaticity in open quantum systems,” [arXiv:1508.05558](#) (2015).
- [34] The existence of small minimum gaps prior to the end of the anneal suggests that it is extremely unlikely that the Gibbs state before crossing these gaps would have a larger overlap with the ground state manifold of the final Hamiltonian than after. Therefore, measuring the system midway through a quantum annealing process will generically yield lower success probabilities than measurements taking place at the end of it [16, 81, 82]. We give two analytical examples in the Supplemental Information.
- [35] This is equivalent to having a bounded degree connectivity graph.
- [36] The analysis is based on the equivalence between the Canonical and the Microcanonical Ensembles of Statistical Mechanics. This equivalence is reviewed in many places, see e.g., Ref. [83].
- [37] Equivalently,  $Ns(e)$  is the logarithm of the number of microstates with intensive energy  $e$ .
- [38] A relation best known as the second law of thermodynamics  $de = Tds$ .
- [39] An energy probability density that is precisely Gaussian implies that the energy density is a linear function of the inverse temperature  $\beta$  and hence the specific heat is a constant. We elaborate on this point in the Supplemental Information.
- [40] Vicky Choi, “Minor-embedding in adiabatic quantum computation: I. The parameter setting problem,” *Quant. Inf. Proc.* **7**, 193–209 (2008).
- [41] Vicky Choi, “Minor-embedding in adiabatic quantum computation: II. Minor-universal graph design,” *Quant. Inf. Proc.* **10**, 343–353 (2011).
- [42] The reader is referred to the Supplemental Information for further details.
- [43] Details of these instances as well as similar results obtained for other problem classes are given in the Supplemental Information.
- [44] The D-Wave processors are known to suffer from additional sources of error such as problem specification errors [58, 59] and freeze-out before the end of the anneal [84] that prevent thermalization to the programmed problem Hamiltonian.
- [45] R. Saket, “A PTAS for the Classical Ising Spin Glass Problem on the Chimera Graph Structure,” ArXiv e-prints (2013), [arXiv:1306.6943 \[cs.DS\]](#).
- [46] By no means however, is it meant that PTAS is able to thermally sample from a Boltzmann distribution of the input problem. In fact, it should be clear that PTAS does not.
- [47] C. J. Geyer, “Parallel tempering,” in *Computing Science and Statistics Proceedings of the 23rd Symposium on the Interface*, edited by E. M. Keramidas (American Statistical Association, New York, 1991) p. 156.
- [48] Koji Hukushima and Koji Nemoto, “Exchange monte carlo method and application to spin glass simulations,” *Journal of the Physical Society of Japan* **65**, 1604–1608 (1996).
- [49] Details of our PT implementation can be found in the Supplemental Information.
- [50] Similar scaling behavior for other classes of Hamiltonians, specifically 3-regular 3-XORSAT instances and random  $\pm 1$  instances, is also observed, and we give the results for these instances in the Supplemental Information.
- [51] T. Jörg, J. Lukic, E. Marinari, and O. C. Martin, “Strong universality and algebraic scaling in two-dimensional ising spin glasses,” *Phys. Rev. Lett.* **96**, 237205 (2006).
- [52] C. K. Thomas, D. A. Huse, and A. A. Middleton, “Zero and low temperature behavior of the two-dimensional  $\pm j$  ising spin glass,” *Phys. Rev. Lett.* **107**, 047203 (2011).
- [53] Francesco Parisen Toldin, Andrea Pelissetto, and Ettore Vicari, “Finite-size scaling in two-dimensional ising spin-glass models,” *Phys. Rev. E* **84**, 051116 (2011).
- [54] L. A. Fernandez, E. Marinari, V. Martin-Mayor, G. Parisi, and J. J. Ruiz-Lorenzo, “Universal critical behavior of the two-dimensional ising spin glass,” *Phys. Rev. B* **94**, 024402 (2016).
- [55] A. Lucas, “Ising formulations of many NP problems,” *Front. Phys.* **2**, 5 (2014).
- [56] Walter Vinci and Daniel A. Lidar, “Optimally stopped optimization,” *Phys. Rev. Applied* **6**, 054016 (2016).
- [57] Sergio Boixo, Troels F. Ronnow, Sergei V. Isakov, Zhihui Wang, David Wecker, Daniel A. Lidar, John M. Martinis, and Matthias Troyer, “Evidence for quantum annealing with more than one hundred qubits,” *Nat. Phys.* **10**, 218–224 (2014).
- [58] Victor Martin-Mayor and Itay Hen, “Unraveling quantum annealers using classical hardness,” *Scientific Reports* **5**, 15324 EP – (2015).
- [59] Zheng Zhu, Andrew J. Ochoa, Stefan Schnabel, Firas Hamze, and Helmut G. Katzgraber, “Best-case performance of quantum annealers on native spin-glass benchmarks: How chaos can affect success probabilities,” *Phys. Rev. A* **93**, 012317 (2016).
- [60] Mohammad H. Amin, “Searching for quantum speedup in quasistatic quantum annealers,” *Phys. Rev. A* **92**, 052323 (2015).
- [61] Kristen L Pudenz, Tameem Albash, and Daniel A Lidar, “Error-corrected quantum annealing with hundreds of qubits,” *Nat. Commun.* **5**, 3243 (2014).
- [62] Kristen L. Pudenz, Tameem Albash, and Daniel A. Lidar, “Quantum annealing correction for random Ising problems,” *Phys. Rev. A* **91**, 042302 (2015).
- [63] Walter Vinci, Tameem Albash, Gerardo Paz-Silva, Itay Hen, and Daniel A. Lidar, “Quantum annealing correction with minor embedding,” *Phys. Rev. A* **92**, 042310– (2015).
- [64] Shunji Matsuura, Hidetoshi Nishimori, Tameem Albash, and Daniel A. Lidar, “Mean field analysis of quantum annealing correction,” [arXiv:1510.07709](#) (2015).
- [65] Walter Vinci, Tameem Albash, and Daniel A Lidar, “Nested quantum annealing correction,” *Npj Quantum Information* **2**, 16017 EP – (2016).
- [66] Shunji Matsuura, Hidetoshi Nishimori, Walter Vinci, Tameem Albash, and Daniel A. Lidar, “Quantum-annealing correction at finite temperature: Ferromagnetic  $p$ -spin models,” *Phys. Rev. A* **95**, 022308 (2017).

- [67] Nicholas Chancellor, “Modernizing quantum annealing using local searches,” *New Journal of Physics* **19**, 023024 (2017).
- [68] N. Chancellor, “Modernizing Quantum Annealing II: Genetic Algorithms and Inference,” ArXiv e-prints (2016), [arXiv:1609.05875 \[quant-ph\]](#).
- [69] Hamed Karimi and Gili Rosenberg, “Boosting quantum annealer performance via sample persistence,” *Quantum Information Processing* **16**, 166 (2017).
- [70] H. Karimi, G. Rosenberg, and H. G. Katzgraber, “Effective optimization using sample persistence: A case study on quantum annealers and various Monte Carlo optimization methods,” ArXiv e-prints (2017), [arXiv:1706.07826 \[cs.DM\]](#).
- [71] S. P. Jordan, E. Farhi, and P. W. Shor, “Error-correcting codes for adiabatic quantum computation,” *Phys. Rev. A* **74**, 052322 (2006).
- [72] Adam D. Bookatz, Edward Farhi, and Leo Zhou, “Error suppression in hamiltonian-based quantum computation using energy penalties,” *Physical Review A* **92**, 022317–(2015).
- [73] Zhang Jiang and Eleanor G. Rieffel, “Non-commuting two-local hamiltonians for quantum error suppression,” *Quantum Information Processing* **16**, 89 (2017).
- [74] Milad Marvian and Daniel A. Lidar, “Error suppression for hamiltonian-based quantum computation using subsystem codes,” *Phys. Rev. Lett.* **118**, 030504 (2017).
- [75] Steven H. Adachi and Maxwell P. Henderson, “Application of quantum annealing to training of deep neural networks,” [arXiv:1510.06356 \(2015\)](#).
- [76] M. Benedetti, J. Realpe-Gómez, R. Biswas, and A. Perdomo-Ortiz, “Quantum-assisted learning of graphical models with arbitrary pairwise connectivity,” ArXiv e-prints (2016), [arXiv:1609.02542 \[quant-ph\]](#).
- [77] C.H. Papadimitriou, *Computational Complexity* (Addison Wesley Longman, Reading, Massachusetts, 1995).
- [78] Leslie G. Valiant, “The complexity of enumeration and reliability problems,” *SIAM Journal on Computing* **8**, 410–421 (1979).
- [79] Vilhelm Dahllöf, Peter Jonsson, and Magnus Wahlström, “Counting models for 2sat and 3sat formulae,” *Theoretical Computer Science* **332**, 265 – 291 (2005).
- [80] L. Campos Venuti, T. Albash, M. Marvian, D. Lidar, and P. Zanardi, “Relaxation vs. adiabatic quantum steady state preparation: which wins?” ArXiv e-prints (2016), [arXiv:1612.07979 \[quant-ph\]](#).
- [81] Boris Altshuler, Hari Krovi, and Jrmie Roland, “Anderson localization makes adiabatic quantum optimization fail,” *Proceedings of the National Academy of Sciences* **107**, 12446–12450 (2010).
- [82] Sergey Knysh, “Zero-temperature quantum annealing bottlenecks in the spin-glass phase,” *Nature Communications* **7**, 12370 EP – (2016).
- [83] V. Martín-Mayor, “Microcanonical approach to the simulation of first-order phase transitions,” *Phys. Rev. Lett.* **98**, 137207 (2007).
- [84] Mohammad H. Amin, “Searching for quantum speedup in quasistatic quantum annealers,” *Phys. Rev. A* **92**, 052323 (2015).
- [85] See Supplemental Material for additional details about the derivations and simulations, which includes Refs. [86–92].
- [86] M. Falcioni, E. Marinari, M.L. Paciello, G. Parisi, and B. Taglienti, “Complex zeros in the partition function of the four-dimensional su(2) lattice gauge model,” *Physics Letters B* **108**, 331 – 332 (1982).
- [87] Alan M. Ferrenberg and Robert H. Swendsen, “Optimized monte carlo data analysis,” *Phys. Rev. Lett.* **63**, 1195–1198 (1989).
- [88] Andrew D. King, Trevor Lanting, and Richard Harris, “Performance of a quantum annealer on range-limited constraint satisfaction problems,” [arXiv:1502.02098 \(2015\)](#).
- [89] Firas Hamze and Nando de Freitas, “From fields to trees,” in *UAI*, edited by David Maxwell Chickering and Joseph Y. Halpern (AUAI Press, Arlington, Virginia, 2004) pp. 243–250.
- [90] Alex Selby, “Efficient subgraph-based sampling of ising-type models with frustration,” [arXiv:1409.3934 \(2014\)](#).
- [91] Lov K. Grover, “Quantum mechanics helps in searching for a needle in a haystack,” *Phys. Rev. Lett.* **79**, 325–328 (1997).
- [92] Jérémie Roland and Nicolas J. Cerf, “Quantum search by local adiabatic evolution,” *Phys. Rev. A* **65**, 042308–(2002).