

ARTICLE

Received 25 Dec 2015 | Accepted 24 Jun 2016 | Published 5 Aug 2016

DOI: 10.1038/ncomms12370

OPEN

Zero-temperature quantum annealing bottlenecks in the spin-glass phase

Sergey Knysh^{1,2}

A promising approach to solving hard binary optimization problems is quantum adiabatic annealing in a transverse magnetic field. An instantaneous ground state—initially a symmetric superposition of all possible assignments of N qubits—is closely tracked as it becomes more and more localized near the global minimum of the classical energy. Regions where the energy gap to excited states is small (for instance at the phase transition) are the algorithm's bottlenecks. Here I show how for large problems the complexity becomes dominated by $O(\log N)$ bottlenecks inside the spin-glass phase, where the gap scales as a stretched exponential. For smaller N , only the gap at the critical point is relevant, where it scales polynomially, as long as the phase transition is second order. This phenomenon is demonstrated rigorously for the two-pattern Gaussian Hopfield model. Qualitative comparison with the Sherrington-Kirkpatrick model leads to similar conclusions.

¹QuAIL, NASA Ames Research Center, Moffett Field, California 94035, USA. ²SGT Inc., 7701 Greenbelt Road, Suite 400, Greenbelt, Maryland 20770, USA. Correspondence and requests for materials should be addressed to S.K. (email: Sergey.I.Knysh@nasa.gov).

Quantum algorithms offer hope for tackling computer science problems that are intractable for classical computers¹. However, exponential speed-ups seen in, for example, number factoring², have not materialized for more difficult non-deterministic polynomial time (NP)-complete problems³. Those problems are targeted by the quantum adiabatic annealing algorithm (QA)^{4–6}. Any NP-hard problem can be recast as quadratic binary optimization. QA solves it by implementing a quantum Hamiltonian, written with the aid of Pauli matrices as

$$\hat{H} = -\frac{1}{2} \sum_{i,k=1}^N J_{ik} \hat{\sigma}_i^z \hat{\sigma}_k^z - \sum_{i=1}^N h_i \hat{\sigma}_i^z - \Gamma(t) \sum_{i=1}^N \hat{\sigma}_i^x. \quad (1)$$

Here the first two terms, diagonal in z -basis, encode the objective function. The last term represents the magnetic field in the transverse direction, which is decreased from $\Gamma(0) \gg 1$ to $\Gamma(T_{\text{ann}}) = 0$. The time T_{ann} needed by the algorithm is determined by a condition that the annealing rate is sufficiently low to inhibit non-adiabatic transitions:

$$d\Gamma/dt \ll \Delta E \cdot \Delta\Gamma. \quad (2)$$

These are most likely near points where the instantaneous gap to excited states ΔE attains a minimum as a function of Γ ; further, $\Delta\Gamma$ is defined as the width of the region where the gap remains comparable to its minimum value.

QA offers no worst-case guarantees on time complexity⁷, but initial assessments of typical case complexity were optimistic. Both experimental⁸ and theoretical⁹ evidence hinted at performance improvement over simulated annealing for finite-dimensional glasses; however, some empirical evidence in support of the theory has recently been called into question¹⁰. Early exact diagonalization studies also observed polynomially small gaps in the constraint satisfaction problem (CSP) on random hypergraphs¹¹, but that finding had been challenged by quantum Monte Carlo studies involving larger sizes¹². Benchmarking of a hardware implementation of QA, courtesy of D-Wave Systems, shows no improvement in the scaling of the performance^{13,14}. Whether that might be attributable to a finite temperature at which the device operates or its intrinsic noise is unclear at present^{15–17}.

Statistical physics offers some intuitive guidance: Small gaps develop at the quantum phase transition point and become exponentially small when the transition is first-order^{18–21} or when different parts of the system become critical at different times for strong-disorder continuous phase transitions²². The most promising candidates for QA are thus problems with bona fide second-order phase transitions, where the disorder is irrelevant at the quantum critical point (QCP).

The scaling analysis described here suggests a polynomially small gap at the critical point of the archetypal spin glass: the Sherrington-Kirkpatrick (SK) model^{23–25}. It has been pointed out^{9,26} that QA may still be doomed by the bottlenecks in the spin-glass phase. Exponentially small gaps away from the critical point have been observed in simulations²⁷, but adequate theoretical description of this phenomenon has proven challenging. A perturbative argument in support of this qualitative picture has been considered in ref. 26. However, the results were not borne out by more accurate analysis that took into account the extreme value statistics of energy levels²⁸.

The present manuscript sheds light on the mechanism of tunnelling bottlenecks in the spin-glass phase. Using exact, non-perturbative, methods, this is illustrated for a simple model, but the main findings are expected to be valid for quantum annealing of more realistic spin glasses. During annealing, the system must undergo a cascade of tunnellings at some specific values of

$\Gamma_1, \Gamma_2, \dots$ in an approximate geometric progression. For a finite system size, these bottlenecks are few, $O(\log N)$, and may not even appear until N is sufficiently large, highlighting the challenge of interpreting the results of numerical studies. Bottlenecks also become increasingly easier as $\Gamma \rightarrow 0$, counter to expectations that tunnellings are inhibited as the model becomes more classical. A related finding is that the time complexity of QA is exponential only in some fractional power of problem size: a mild improvement over more pessimistic estimates²⁶.

Results

Summary. The spin-glass phase, which is entered below some critical value of the transverse field Γ_c , is characterized by a large number of valleys. Often, this transition is abrupt, driven by competition between an extended state and a valley (localized state) with the lowest energy, as occurs in the random energy model^{19,20}. The exponentially small overlap between the two states then determines the gap at the phase transition. However, even if new valleys develop in a continuous manner as Γ decreases, small changes in the transverse field may result in a chaotic reordering of associated energy levels, leading to Landau-Zener avoided crossings and attendant exponentially small gaps.

Nonetheless, attempts to make this intuition exact are fraught with potential pitfalls. For increasing Γ , two randomly chosen valleys are equally likely to come either closer together or further apart in energy. In the case of the former—and further, if the sensitivity of energy levels to changes in the transverse field is so large that the levels ‘collide’ before either valley disappears—avoided crossing will occur. This may not be necessarily the case when one considers ‘collisions’ with the ground state, which are of particular concern to QA. The ground state corresponds to a valley with the lowest energy; this and other low-lying valleys obey fundamentally different statistics of the extremes.

A case in point is the analysis of ref. 26, which develops perturbation theory in Γ . The classical limit ($\Gamma = 0$) is used as a starting point; how that analysis might be extended to $\Gamma > 0$ has also been discussed²⁹. A type of CSP has been considered: classical energy levels are discrete non-negative integers (number of violated constraints) but have exponential degeneracy. ‘Zeeman splitting’ for $\Gamma > 0$ scales as \sqrt{N} , which, for large problems, may be sufficient to overcome the $O(1)$ classical gap and cause avoided crossings of levels associated with different classical energies. Yet this trend disappears if only levels with the smallest energies (after splitting) are considered; these are relevant for avoided crossings with the ground state. This about-face is not immediately apparent, only coming into play for $N \gtrsim 100$, when the exponential degeneracy of the classical ground state sets in. It has, however, been confirmed with analytical argument and numerics²⁸. Consequently, arguments based on perturbation theory cannot be used to establish the phenomenon.

This manuscript offers a fresh perspective, illustrated by studying quantum annealing of the Hopfield model. Mean-field analysis correctly describes thermodynamics if the number of random ‘patterns’ is small. The method is further extended to extract information about exact quantum energy levels. Importantly, the classical energy landscape is made much more complex by insisting that the distribution of disorder is Gaussian. Figure 1 sketches a ‘phase diagram’ obtained for this model. For decreasing Γ the gap changes as follows: (1) it is finite (does not scale with N) in the paramagnetic phase, $\Gamma > \Gamma_c$; (2) scales as $1/N^{1/3}$ in the narrow region of width $1/N^{2/3}$ around $\Gamma = \Gamma_c$; (3) increases slightly, with typical values scaling as $1/N^{1/4}$ for $\Gamma < \Gamma_c$. In addition, there are avoided crossings at isolated points $\Gamma_1, \Gamma_2, \dots$.




Figure 1 | Scaling of the gap in various regions. Sketch of the behaviour of the gap in a Hopfield model with the Gaussian distribution of disorder variables (units are arbitrary): in the paramagnetic phase ($\Gamma > \Gamma_c$), in the glassy phase ($\Gamma < \Gamma_c$) and in the critical region ($\Gamma \approx \Gamma_c$). Scaling of the typical gap in these regions is indicated in bold letters using big-O notation. The area $\Gamma < \Gamma_{\min} \sim 1/N$ is where the discrete nature of the energy landscape becomes manifest: the ground state becomes nearly completely localized. The glassy phase also contains $\log N$ isolated bottlenecks at $\Gamma_1, \Gamma_2, \Gamma_3$, and so on. (indicated with red arrows), where the gaps scale as a stretched exponential. a.u., arbitrary units.

The first non-trivial example requires two Gaussian patterns. In this case the ‘energy landscape’ is effectively one-dimensional, which greatly simplifies the analysis. The most important element of this analysis is accounting for the extreme value statistics associated with the valleys (local minima) having the lowest energies. To this end, the distribution of energies of the classical landscape must be conditioned so that they are never below the energy of the global minimum. This becomes feasible when reformulated as a continuous random process, in the limit $N \rightarrow \infty$.

This particular model is not as interesting from the computer science perspective, not least because it affords an efficient classical algorithm. It is sufficiently simple so that a complete quantitative analysis presented later on in the manuscript has been possible. Yet, the model captures the essential properties of the spin glass: its qualitative features directly apply to much more general models, including Sherrington-Kirkpatrick. The most important feature of the classical energy landscape is that it exhibits fractal properties, which both ensures that hard bottlenecks are present in the spin-glass phase and also governs their distribution. The role of the transverse field is to approximately coarse-grain it on scales determined by Γ , eliminating small barriers; thus the number of valleys decreases as Γ grows. A specific random process, corresponding to the energy landscape of the ‘infinite’-size instance, will contain every possible realization of itself at some ‘length scale’. Some realizations will contain high barriers that cannot be easily overcome; these will lead to tunnelling bottlenecks.

This intuition can be used to immediately establish the scaling of the number of tunnelling bottlenecks. Since the model contains no inherent length scale in the limit $N \rightarrow \infty$, it can be argued that the expected number of tunnelling bottlenecks in a finite interval $[\Gamma_1; \Gamma_2]$ should be a function of the ratio Γ_2/Γ_1 . The logarithm is the only function that respects additivity, that is, $\mathcal{N}_{\text{tunn}}([\Gamma_1, \Gamma_2]) = \mathcal{N}_{\text{tunn}}([\Gamma_1; \Gamma']) + \mathcal{N}_{\text{tunn}}([\Gamma', \Gamma_2])$. To obtain the total number of bottlenecks, one considers the interval $[\Gamma_{\min}; \Gamma_c]$. Here $\Gamma_c \sim 1$ is the boundary of the spin-glass phase. The lower cutoff, Γ_{\min} , corresponds to the lowest energy scale of the classical model, which scales as an inverse power of system size, for example, as $1/N$ for the Gaussian Hopfield model. In a

sense, tunnelling bottlenecks are connected to the $\Gamma = 0$ ‘fixed point’ (note that the classical gap vanishes asymptotically). To summarize, the number of tunnelling bottlenecks will grow as

$$\mathcal{N}_{\text{tunn}} \approx \alpha \ln N. \tag{3}$$

Locations (in Γ) will depend on specific disorder realization, but self-similarity ensures that the successive ratios Γ_n/Γ_{n+1} converge to a universal distribution.

This logarithmic rise is far weaker than a power law seen in some phenomenological models of temperature chaos³⁰ and, as has been argued above, likely to be a universal feature. The prefactor is model-dependent; its numerical value can be used to estimate the minimum problem size for which the mechanism becomes relevant, via $\mathcal{N}_{\text{tunn}} \approx 1$. A value of $\alpha \approx 0.15$ is obtained for the problem at hand, so that additional bottlenecks become an issue for $N \gtrsim 1,000$. Prior numerical studies similarly required large sizes before the exponentially small minimum gap was observed²⁷, and so far there has been no evidence of two or more exponentially small gaps coexisting. The slow, logarithmic increase of the expected number of bottlenecks is the most likely culprit.

A notable feature of these results is that tunnelling become progressively ‘easier’ as $\Gamma \rightarrow 0$ despite the fact that the model becomes more classical. Tunnelling gaps increase as

$$\Delta E_{\text{tunn}}^{(n)} \sim e^{-c\Gamma_n^{3/4} N^{3/4}}. \tag{4}$$

Notice that they cease to be exponentially small for $\Gamma \lesssim \Gamma_{\min}$; at that point the ground state is already localized near the correct global minimum. The power law exponent for this stretched exponential is model-dependent, related to the scaling of barrier heights. These scale as $N^{1/2}$ for the Gaussian Hopfield model, which, together with $O(N)$ scaling of the effective mass, gives rise to the $N^{3/4}$ term in the exponent.

The finding that the gaps increase for smaller Γ deserves explanation. Typically, valleys with similar energies differ by up to $N/2$ spin flips. This changes once lowest energies are considered: All spin configurations with energies less than ϵ above the global minimum are contained in a neighbourhood of radius $O[(N\epsilon)^{2/3}]$, using Hamming distance as a metric. The problem is not rendered easy by the mere fact that the global minimum is so pronounced (although theoretical analysis inspired an efficient classical algorithm for the $p=2$ Hopfield network, briefly described later on). It does imply, however, that the ground state wavefunction does not jump chaotically: Every subsequent tunnelling involves shorter distances, with $O(\Gamma N)$ spin flips, and achieves progressively better approximation to the true global minimum. Absent such a trend, annealing would be most difficult towards the end of the algorithm, when $\Gamma \sim 1/N$, and the minimum gap would exhibit less favourable exponential scaling²⁶.

In what follows, the model and its solution are described in greater detail. First, finite-size scaling of the ‘easy’ QCP bottleneck is linked to the thermodynamics of the phase transition. The next part goes beyond thermodynamics, considering small corrections to the extensive contribution to the free energy. The entire low-energy spectrum, which depends on disorder realization, is mapped onto that of a simple quantum mechanical particle in a random potential. Finally, extreme value statistics is applied to investigate the properties of that random potential near its global minimum by mapping it to a Langevin process. This yields the distribution of hard bottlenecks in a universal regime ($\Gamma \ll 1$).

Quantum Hopfield network. Consider a model with rank- p matrix of interactions and no longitudinal field ($h_i=0$):ref. 31

$$J_{ik} = \frac{1}{N} \sum_{\mu=1}^p \xi_i^{(\mu)} \xi_k^{(\mu)} \quad (5)$$

(cf. $\text{rk } J_{ik} = N$ for SK model), where $\xi_i^{(\mu)}$ are taken to be independent and identically distributed (i.i.d.) random variables of unit variance. The thermodynamics of this quantum Hopfield model has been worked out in great detail by Nishimori and Nonomura³². In fact, that study motivated the development of QA⁴.

When p is small ($J_{ik} \sim 1/N$), it is appropriate to replace local longitudinal fields with their mean values $h_i = \sum_k J_{ik} \langle \sigma_k^z \rangle$. The identity $\langle \sigma_i^z \rangle = h_i \sqrt{\Gamma^2 + h_i^2}$ is used to close this system of equations. For a transverse field below the critical value, $\Gamma < \Gamma_c = 1$, there appears a non-trivial ($\mathbf{m} \neq 0$) solution to the self-consistency equation for the macroscopic order parameter, a vector with p components:

$$\mathbf{m} = \frac{1}{N} \sum_{i=1}^N \xi_i \langle \hat{\sigma}_i^z \rangle. \quad (6)$$

Here, the disorder variables are also written using vector notation: $\xi_i = (\xi_i^{(1)}, \dots, \xi_i^{(p)})$.

This model is equivalent to the Curie-Weiss (quantum) ferromagnet, which has a continuous phase transition characterized by a set of mean-field critical exponents. Two of these are particularly useful in the analysis of the annealing complexity: the one for the singular component of the ground state energy

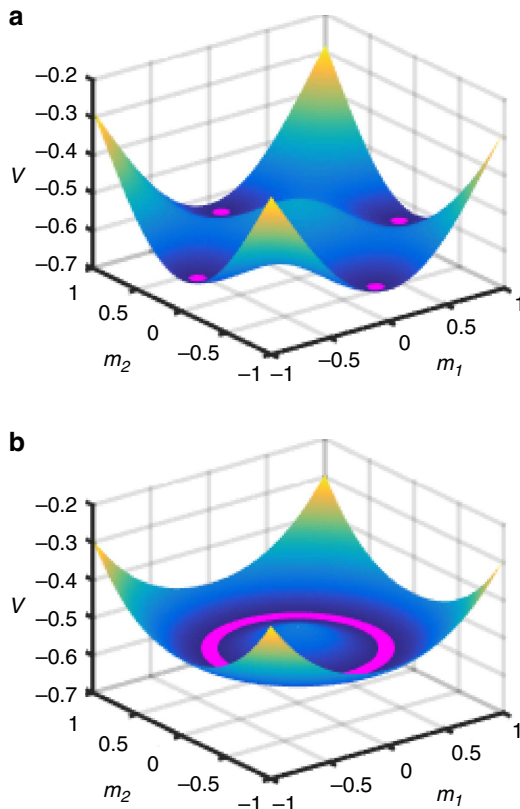


Figure 2 | The shape of the effective potential. The plots show the disorder-averaged $\mathcal{V}(\mathbf{m})$ below the critical point ($\Gamma = 0.5 < \Gamma_c$) for (a) the bimodal distribution and (b) the Gaussian distribution of disorder variables. Minima of the potential are highlighted with magenta: the $2p$ -fold degenerate global minima organized in pockets corresponding to encoded patterns in a, and a continuum of degenerate minima (connected by arbitrary rotations) in b.

($E_0^{\text{sing}}/N \propto |\gamma|^a$) as well as the dynamical exponent for the gap ($E_1 - E_0 \propto \gamma^b$), where $\gamma = \Gamma - \Gamma_c$ is the ‘distance’ from the critical point. These exponents are defined for the infinite system, yet fairly general heuristic analysis (see the Methods section) predicts finite-size scaling for the QCP bottleneck:

$$\Delta E_c \propto N^{-\frac{b}{a-b}}, \quad \Delta \Gamma_c \propto N^{-\frac{1}{a-b}}. \quad (7)$$

Substituting values $a=2$ and $b=1/2$ for the problem at hand, one may estimate the gap at the critical point and the width of the critical region to be $O(N^{-1/3})$ and $O(N^{-2/3})$, respectively.

Worse-than-any-polynomial complexity of quantum annealing might be expected for the first order phase transition, which exhibits no critical scaling (but see ref. 33). Another possibility is for the dynamical exponent to diverge at the infinite randomness QCP: the finite-size gap scales as $\exp(-c\sqrt{N})$ in a random Ising chain²². For the Hopfield model, however, this scaling is polynomial, as the disorder is irrelevant at the critical point. More intriguing is the fact that these pessimistic scenarios are not found in SK spin glass either: the model is characterized by the same set of critical exponents, albeit with logarithmic corrections^{23–25}. These corrections to scaling increase the gap and, respectively, decrease the width of the critical region by a factor of $\log^{2/9} N$. Thus, as long as $T_{\text{ann}} \gtrsim N$, non-adiabatic transitions at the critical point should be suppressed. This presents a conundrum as SK model is known to be an NP-hard problem; finding a polynomial-time (even in typical case) quantum algorithm would be a surprising development. The heuristic analysis is clearly insufficient, but ‘digging’ deeper into a problem would require a more ‘microscopic’ analysis. In the following, the problem is mapped to ordinary quantum mechanics to uncover its low-energy spectrum that explicitly depends on a particular realization of disorder, $\{\xi_i\}$.

Exact low-energy spectrum. Mean-field theory can be derived in a more systematic manner via Hubbard–Stratonovich transformation. General overview is presented below; additional details can be found in Methods and the Supplementary Note 1. Finite-temperature partition function $Z(\beta) = \text{Tr}(e^{-\beta H})$ can be written as a path integral over $\mathbf{m}(t)$, which now acquires a dependence on the imaginary time $0 \leq t \leq \beta$, with periodic boundary conditions: $\mathbf{m}(\beta) = \mathbf{m}(0)$. The value of the integral is dominated by stationary paths corresponding to the minimum of an effective potential $\mathcal{V}(\mathbf{m})$. While the discussion above has been deliberately equivocal on the distribution of disorder variables, it is now instructive to contrast bimodal ($\xi_i^{(\mu)} = \pm 1$) and Gaussian choices. The shapes of the effective potential for both scenarios are depicted in Fig. 2.

The conventional bimodal choice defines the model of associative memory: In the limit $\Gamma=0$ the ‘patterns’ can be perfectly recalled ($s_i = \pm \xi_i^{(\mu)}$) when p is small. For the Gaussian choice, the global minimum corresponds to a mixture³⁴

$$s_i = \pm \text{sgn} \left(\sum_{\mu} \alpha_{\mu} \xi_i^{(\mu)} \right), \quad (8)$$

rendering memory useless. In the bimodal case, such ‘spurious’ states only become stable once the number of patterns scales with the problem size³⁵: $p > 0.05N$. The BC_p (bimodal) or $\text{O}(p)$ (Gaussian) symmetry of the effective potential is only approximate, to leading order in N . The degeneracy of the ground state is 2 (due to global spin inversion symmetry) for almost all disorder realizations when $p \geq 3$ or $p \geq 2$ in the bimodal and Gaussian scenarios, respectively (note that that the $p=2$ bimodal case possesses an additional symmetry, which makes the ground state 4-fold degenerate). The system is in a symmetric superposition of the degenerate global minima at the end of QA. It evolves entirely in the symmetric subspace since the time-

dependent Hamiltonian commutes with $\hat{U} = \exp(i\hat{S}^x)$. Thus, it should be noted that small gaps between symmetric and antisymmetric wavefunctions are irrelevant to QA and can be ignored.

Disorder fluctuations ‘nudge’ QA towards the ‘correct’ pattern as it passes the critical point in the bimodal Hopfield model. No further bottlenecks are encountered; gaps for $\Gamma < \Gamma_c$ as well as the

‘classical’ ($\Gamma = 0$) gap are $O(1)$. By contrast, the classical gap scales as $O(1/N)$ in the Gaussian Hopfield model, alerting to a ‘danger’ posed by the $\Gamma = 0$ ‘fixed point’. To find the low-energy spectrum when $\Gamma < \Gamma_c$, note that the dominant contribution to the path integral is from paths where the magnitude of the ‘magnetization’ vector remains approximately constant, close to its saddle-point value, while the angle is a slow function of time: $\mathbf{m}(t) \approx m_\Gamma \begin{pmatrix} -\sin\vartheta(t) \\ \cos\vartheta(t) \end{pmatrix}$ for $p=2$. Integrating out the amplitude degrees of freedom, the partition function is rewritten as

$$Z(\beta) \propto \int [d\vartheta(t)] e^{-\int_0^\beta \left(\frac{M}{2} \left(\frac{d\vartheta}{dt} \right)^2 + V_\Gamma(\vartheta) \right) dt}, \quad (9)$$

which describes a quantum-mechanical particle of mass $M = O(N)$ moving on a ring, subjected to a random potential

$$V_\Gamma(\vartheta) = - \sum_i \sqrt{\Gamma^2 + m_\Gamma^2 \xi_i^2 \sin^2(\vartheta - \theta_i)} + N \langle \sqrt{\dots} \rangle_\xi, \quad (10)$$

where $\xi_i \equiv \begin{pmatrix} \cos \theta_i \\ \sin \theta_i \end{pmatrix}$ and the last term, written in shorthand,

adds a constant offset so that $\langle V(\vartheta) \rangle_\xi = 0$. Notice that $V_\Gamma(\vartheta) \propto \sqrt{N}$ via central limit theorem, thereby representing a higher-order correction to the extensive part of the free energy.

Since the partition function $Z(\beta) = \sum_n e^{-\beta E_n}$ encodes information about the spectrum, low-energy (Goldstone) excitations of the many-body problem are in one-to-one correspondence with the energy levels of a quantum mechanical particle, up to a constant shift. The next step is to find the properties of this potential near a global minimum. These are relevant in a regime away from the critical point, $\Gamma \ll \Gamma_c$, where the universal behaviour characterized by the appearance of ‘hard’ bottlenecks sets in.

Evolution of the random potential. Scaling of the gap for $\Gamma < \Gamma_c$ can be obtained via semiclassical analysis. Small level splitting due to tunnelling between wells at the two degenerate global minima (this degeneracy is a consequence of the global Z_2 symmetry: $V_\Gamma(\vartheta + \pi) = V_\Gamma(\vartheta)$) is not relevant to QA. Higher degeneracies are statistically unlikely; quantization rules predict $O(N^{-1/4})$ gaps between energy levels within the symmetric subspace. But this refers to the typical gap, obtained for fixed Γ for almost all realizations of disorder. As quantum annealing sweeps the transverse field for a fixed realization of disorder, $V_\Gamma(\vartheta)$ might undergo global bifurcation. This would result in a small tunnelling gap for a specific value of Γ when the competing minima are in resonance.

Such a scenario is impossible near the QCP. Coefficients in the Fourier expansion of the random potential, $\sum_k (a_k \cos 2k\vartheta + b_k \sin 2k\vartheta)$, decrease as m^{2k}/Γ^{2k-1} so that the first harmonic dominates for $\Gamma \approx \Gamma_c$. Semiclassical analysis confirms a $O(N^{-1/3})$ gap at the critical point (where the curvature of the effective potential vanishes, leaving only the quartic part). As Γ decreases, the random potential becomes more rugged (see Fig. 3a), smooth only on scales $\Delta\vartheta \sim \Gamma$, which makes global bifurcations more likely. Furthermore, it exhibits properties that allow one to make important predictions without detailed calculations. Rescaling the potential in the vicinity of either global minimum $V_\Gamma(\vartheta^*) = V_\Gamma^*$,

$$\begin{aligned} \vartheta - \vartheta^* &\rightarrow \ell(\vartheta - \vartheta^*), \\ V_\Gamma - V_\Gamma^* &\rightarrow \ell^{3/2}(V_\Gamma - V_\Gamma^*), \end{aligned} \quad (11)$$

describes the same model but for the rescaled $\Gamma \rightarrow \ell\Gamma$ and a different realization of disorder. This scale invariance is responsible for the logarithmic scaling of the number of tunnelling bottlenecks as has been explained earlier in the text. However, it still remains necessary to demonstrate that the

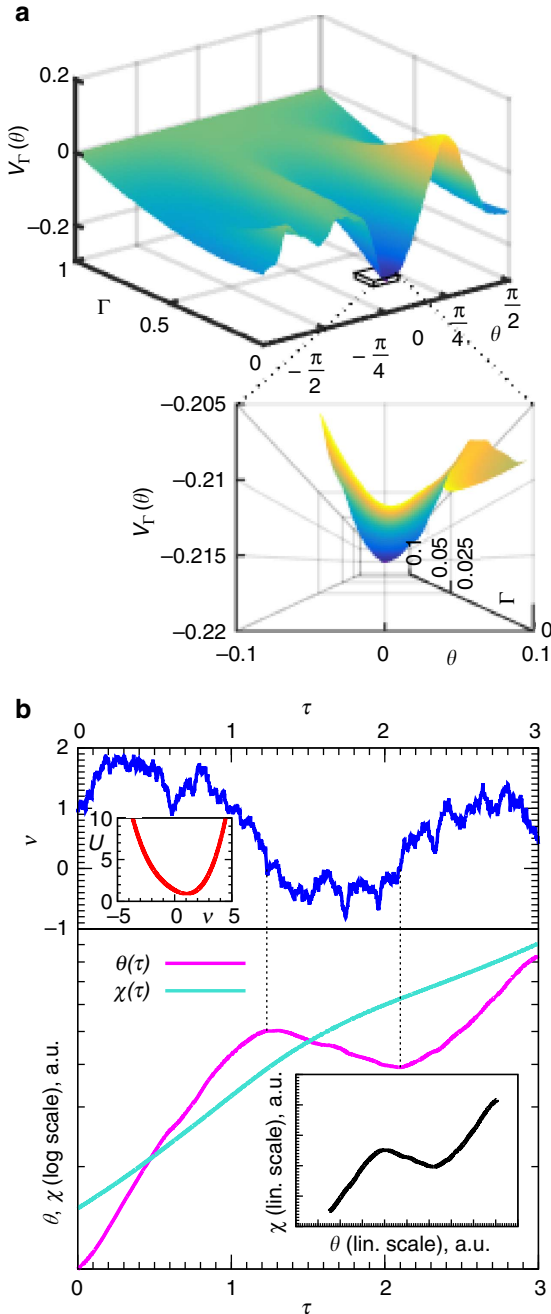


Figure 3 | Appearance of local minima. (a) Potential $V_\Gamma(\vartheta)$ for a specific realization of disorder as a function of Γ . A perspective-projection three-dimensional plot that zooms on a region near the global minimum of $V_0(\vartheta)$ is shown below. (b) The top part plots a specific realization of a stochastic process $v(\tau)$ for the Langevin potential $U(v)$ shown in the inset. The bottom part sketches $\chi(\tau)$ and $\vartheta(\tau)$ (up to a constant factor) obtained by integrating equation (12). The inset is a parametric plot $\chi(\vartheta)$ using linear scale. Fluctuations with $v(\tau) < 0$ (between the dotted vertical lines) are responsible for the appearance of local minima. a.u., arbitrary units.

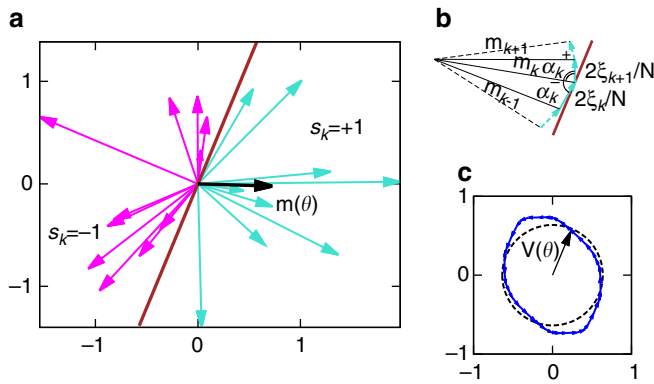


Figure 4 | Geometric interpretation of the random potential. Illustration for a random instance with $N = 20$. **(a)** A vector $\mathbf{m}(\vartheta)$ (black arrow) is defined by drawing a separating line at angle ϑ (in brown) and assigning spins with ξ_k on either side of separating line values $s_k = +1$ (turquoise arrows) and $s_k = -1$ (magenta arrows), respectively. **(b)** Cartoon shows how $\mathbf{m}(\vartheta)$ changes as the separating line rotates counterclockwise for increasing ϑ : the vector $\mathbf{m}(\vartheta)$ is incremented/decremented by $(2/N)\xi_k$ (turquoise arrows) whenever s_k changes sign. **(c)** These increments form a closed path, approximated by a circle; fluctuations around the ideal circle (dashed black line) define a random potential (in blue).

density of bottlenecks is non-zero, which entails an examination of the properties of the random potential in the limit $\Gamma = 0$.

The classical optimization problem corresponds to maximizing the magnitude of $\sum_i \xi_i s_i$. A necessary condition for a local minimum is that two sets of vectors, $\{\xi_i | s_i = +1\}$ and $\{\xi_i | s_i = -1\}$, can be separated by a line. As the angle of this line changes, fluctuations of the amplitude give rise to a random potential $V_0(\vartheta)$ (see Fig. 4). This suggests a linear-time algorithm for finding a global minimum: Sort all vectors by angle (this may introduce an extra $\log N$ factor due to sorting overhead) and exhaustively check all possible angles. Of course, QA algorithm is too generic to exploit the specific structure of the problem; moreover *ad hoc* efficient algorithms are unlikely to exist for more general spin-glass problems.

On short intervals, the random process is described as an undamped Langevin process^{36,37} in the continuous ($N \rightarrow \infty$) limit (hence the exponent of 3/2 in equation (11), corresponding to its fractal dimension). Properly taking into account the statistics of the extremal properties, the process must be conditioned on the fact that $V_0(\vartheta^*) > V_0(\vartheta_0^*) = V_0^*$ away from the global minimum. As described in Methods and the Supplementary Notes 2 and 3, such a conditioned process consists of two uncorrelated branches,

$$V_0(\vartheta) - V_0^* \propto \sqrt{N} \begin{cases} \chi_+, & \vartheta > \vartheta_0^* \\ \chi_-, & \vartheta < \vartheta_0^* \end{cases}, \text{ Integrating equations}$$

$$\begin{aligned} d(\ln \chi_{\pm})/d\tau &= v_{\pm}(\tau), \\ d\vartheta/d\tau &= \pm \chi_{\pm}^{2/3}, \end{aligned} \quad (12)$$

defines $\chi_+(\vartheta)$ and $\chi_-(\vartheta)$ parametrically, in terms of random processes $v_{\pm}(\tau)$ that correspond to a Brownian motion in a non-linear potential depicted in Fig. 3b. This potential is biased towards positive ‘velocities’ v so that $V_0 - V_0^* \sim \sqrt{N}(\vartheta - \vartheta_0^*)^{3/2}$ from equation (12). It will, however, experience arbitrary percentage drops due to subpaths with $v < 0$ (albeit with decreasing probability).

For small but finite Γ , the ‘separating line’ becomes blurred. The random potential adds a ‘quantum correction’ (see Methods and the Supplementary Note 2), $V_{\Gamma}(\vartheta) - V_0(\vartheta) = O(\sqrt{N}\Gamma^{3/2})$.




Figure 5 | Numerical results. **(a-c)** The L-shaped figure formed by the three panels shows the distribution of the tunnelling jumps and the tunnelling exponents. Here $\Delta\vartheta/\Gamma$ is proportional to the number of spin flips for the tunnelling event, normalized by ΓN . The distribution of the tunnelling gaps is quantified using a prefactor in the exponent of equation (4), namely $c = |\ln \Delta E_{\text{tunn}}/(\Gamma N)^{3/4}$. The panel **a** is a scatterplot of c and $\Delta\vartheta/\Gamma$. The histograms of univariate distributions of c and $\Delta\vartheta/\Gamma$ (‘projections’ of the scatterplot) are plotted in panels **b** and **c**, respectively. **(d)** A histogram of distribution of Γ_n/Γ_{n+1} for the successive tunnelling events.

For two minima to come into resonance, they cannot be more than $\Delta\vartheta \sim \Gamma$ apart (that is, $O(N\Gamma)$ spin flips). The tunnelling exponent is given by under-the-barrier action $\mathcal{A} \sim \sqrt{M\Delta V\Delta\vartheta}$, where the effective mass $M \sim N/\Gamma^2$ and $\Delta V \sim \sqrt{N}\Gamma^{3/2}$, leading to equation (4). Numerical results for the universal distribution of tunnelling jumps and exponents are exhibited in Fig. 5.

Discussion

Poor scalability of classical annealing of spin-glass models had been linked to the phenomenon of temperature chaos³⁸. Interestingly, its existence in mean-field glasses had been debated³⁹⁻⁴¹, although it is uncontroversial in finite-dimensional models^{42,43}. Similarly, failures of quantum annealing might be attributed to transverse field chaos. The phenomenon described in this manuscript represents a much stronger finding. The mere fact that ground states at Γ and $\Gamma + \Delta\Gamma$ will have vanishingly small overlap as $N \rightarrow \infty$ is not inconsistent with the continuously evolving ground state and poses no ‘threat’ to QA. By contrast, ‘hard’ bottlenecks correspond to isolated discontinuities that persist as $\Delta\Gamma \rightarrow 0$.

To dwell upon the generality of these results, first note that scaling of the tunnelling exponent will depend on the universality class of the model. The SK model, for instance, exhibits different scaling of barrier heights, believed to be $\propto N^{1/3}$ (see for example, ref. 44). In the model studied, the decrease in the number of spins involved in tunnelling offsets the divergence of the effective mass in the classical limit $\Gamma \rightarrow 0$. As the height of the barriers also decreases, the tunnelling gaps widen towards the end of the algorithm. One might expect qualitatively similar behaviour in realistic spin glasses.

The logarithmic scaling of the number of bottlenecks is due to self-similar properties of the free energy landscape in the interval $[\Gamma_{\text{min}}, \Gamma_c]$. The lower cutoff should correspond to the smallest energy scale in the classical limit, which for the SK model is also a negative power of N , namely $1/\sqrt{N}$. This is related to the linear vanishing of the density of distribution of effective fields as $h \rightarrow 0$ at zero temperature⁴⁵ (since $\int_0^{h_{\text{min}}} P(h)dh \sim 1/N$). The picture is

less clear for CSPs, where the energies are constrained to be non-negative integers; that is, the classical gap is $O(1)$. These energy levels have exponential degeneracy, which is lifted by the transverse field. A value sufficient to make the spectrum effectively quasi-continuous might serve as an appropriate lower cutoff Γ_{\min} in problems of this type. Lack of ‘hard’ bottlenecks in the Hopfield model with the bimodal distribution of disorder and $p = O(1)$ could be attributed to the fact that the number of valleys is finite, which is not representative of a true spin glass.

An interesting observation is that since ‘hard’ bottlenecks correspond to Landau–Zener crossings, annealing times need not be exponentially small. The probability that QA fails to follow the ground state every single time in n repeated experiments is

$$p_{\text{fail}}^{(n)} = \left(e^{-\frac{\pi \Delta E \Delta T}{4(dT/dt)}} \right)^n, \quad (13)$$

which exactly matches the probability of failure for the annealing rate that is n times slower. Using shorter annealing cycles with many repetitions can minimize the effects of decoherence. It suffices that non-adiabatic transitions are suppressed at the critical point only, $\frac{dT}{dt} \lesssim 1/N$.

Even with polynomial annealing rates, coherent evolution would require much better isolation from the environment than what is currently feasible. The only commercial implementation of QA (D-Wave) must contend with a fairly strong coupling to a thermal bath. On the positive side, it accommodates faster annealing cycles, acting as a ‘safety valve’ to dissipate any heat generated during the non-adiabatic process. At the same time, it all but ensures that the system is always in thermal equilibrium with the environment.

The theory presented here breaks down when $\Gamma < T$ so that quantum bottlenecks described here may not be a limiting factor if $\ln(\Gamma/T) \lesssim 1/\alpha$. Main source of errors will be from exponentially many thermally occupied excited states. If the annealing profile were adjusted so that the energy spacing increased beyond T towards the end, this would effectively implement classical annealing. An intricate relationship between temperature, problem size, and the properties of the spin-glass model might determine which mechanism (quantum or classical annealing) will be dominant.

Yet another tradeoff in the design of D-Wave chip is a quasi-two-dimensional (2D) topology of interactions J_{ik} due to fabrication constraints, which incurs significant performance penalty when mean-field models are ‘embedded’ into a ‘Chimera’ graph⁴⁶. So-called ‘native’ problems corresponding to uniformly random instances on this quasi-2D lattice have been argued to be poor candidates for QA due to the lack of a finite-temperature classical phase transition⁴⁷; at the same time, a quantum phase transition at $\Gamma_c > 0$ is expected in 2D glasses⁴⁸.

Whereas first-order phase transition immediately implies exponential complexity, even for small sizes, problems having a continuous phase transition may remain tractable up to a threshold, N_c , beyond which tunnelling bottlenecks become dominant ($\alpha \ln N_c \sim 1$). This ‘tractability threshold’ serves as a silver lining for this otherwise negative result. Moreover, the picture of ‘hard’ bottlenecks may be equally applicable to classical annealing. A recent study demonstrated that classically ‘hard’ instances that exhibit temperature chaos also take much longer time to solve on a D-Wave machine⁴⁹. While in some crafted examples classical annealing is at a unique disadvantage due to first-order phase transition⁵⁰, for most interesting problems both classical and quantum transitions are second-order. In such a scenario, the density of bottlenecks becomes a tie-breaker for evaluating relative performance. Whether quantum annealing can be advantageous in terms of this metric and determining which models will benefit is a practically important question for follow-up work.

There remains a question whether the failure mechanism described here can be somehow circumvented. Such a feat is feasible, for example, for a disordered Ising chain, where an exponentially small gap develops at the critical point, which is a manifestation of Griffiths singularities²². Modification of QA that requires controlling the transverse field for each spin individually can suppress these singularities and restore a polynomial gap.

Based on comparison with another exactly solvable model, it seems that frustration—in addition to disorder—is essential for the appearance of small gaps in the spin-glass phase proper. The seemingly random profile of the energy landscape for finite Γ heralds difficulties in avoiding these bottlenecks in generic spin glasses. Although any prospects of exponential speedup should be met with skepticism, heuristics inspired by spin-glass theory revolutionized branch-and-bound algorithms for CSPs⁵¹. One can remain hopeful that theoretical advances can similarly aid quantum optimization.

Methods

Scaling analysis. Finite-size scaling of the gap at QCP is best understood using an example of a finite-dimensional system. In thermodynamic limit, both correlation length and characteristic time diverge near the phase transition as

$$\xi_c \propto (\Gamma - \Gamma_c)^{-\nu}, \quad \tau_c \propto (\Gamma - \Gamma_c)^{-z\nu}. \quad (14)$$

In a finite system this divergence is smoothed out as soon as the correlation length becomes comparable to lattice size ($\xi_c \sim N^{1/d}$). The minimum gap (the reciprocal of τ) and the width of the critical region should scale as $N^{-z/d}$ and $N^{-1/(d\nu)}$, respectively. In this paper, the product $z\nu$ has been labelled as exponent b . Singular behaviour of normalized ground state energy (free energy) is related to the specific heat exponent ($a = 2 - \alpha$). Dimensionality d can be eliminated with the aid of hyperscaling relation $2 - \alpha = (d + z)\nu$ to yield equation (7) of the main text. Independent estimates of the specific heat exponent can be obtained from the exponents for the order parameter and the susceptibility ($2 - \alpha = 2\beta + \gamma$).

Mapping to quantum mechanics. Finite-temperature partition function can be written as a sum over a set of paths $\{s_i(t)\}$ with $0 \leq t \leq \beta$, where $s_i(t)$ alternates between the values ± 1 . Hubbard–Stratonovich transformation can be used to rewrite it as a path integral

$$\begin{aligned} Z &\sim \sum_{\{\{s_i(t)\}\}} e^{\frac{\beta}{2\Gamma} \int_0^\beta (\sum_i \xi_i s_i(t))^2 dt + \sum_i \mathcal{K}_\Gamma[s_i(t)]} \\ &\sim \int [d\mathbf{m}(t)] e^{-\frac{\beta}{2} \int_0^\beta \mathbf{m}^2(t) dt + \sum_i \ln Z_i} \end{aligned} \quad (15)$$

The ‘kinetic’ term $\mathcal{K}_\Gamma[s(t)]$ in the first equation penalises kinks, representing Γ -dependent ferromagnetic couplings between Trotter slices. As the interaction term is decoupled, the problem reduces to that of evaluating the single-site partition function Z_i for a spin subjected to a magnetic field with the transverse component Γ and the ‘time-dependent’ longitudinal component $h_i(t) = \xi_i \mathbf{m}(t)$. To leading order in N , the saddle-point of the path integral (15) is a solution of mean-field equations. This becomes a degenerate manifold for Gaussian disorder distribution; to determine higher-order contribution all paths such that $|m(t)| = m_\Gamma$ are considered. Evaluating Z_i is best performed in the adiabatic basis that diagonalizes the 2-level Hamiltonian $\hat{H}_i = -h_i(t)\hat{\sigma}^z - \Gamma\hat{\sigma}^x$,

$$Z_i = \text{Tr} \mathcal{T} - e^{\int_0^\beta \hat{H}_i(t) dt} = \text{Tr} \mathcal{T} e^{-\int_0^\beta (\hat{E}_i(t) + \frac{\beta}{2\Gamma} \hat{V}_i(t)) dt}. \quad (16)$$

Here $\hat{E}_i(t)$ is diagonal with eigenvalues $\pm \sqrt{\Gamma^2 + h_i^2(t)}$. Its fluctuations around the mean give rise to the random potential $V_\Gamma(\vartheta)$. Non-adiabatic terms due to rotation of the basis ($\hat{U} = e^{-i\hat{V}_i(t)}$) are treated using second order perturbation theory, giving rise to a kinetic term $\propto (d\vartheta/dt)^2$. Note a simple form of the perturbative term in equation (16) owing to the fact that $\hat{V}_i(t)$ commute for all t . The details of this calculation are given in Supplementary Note 1.

Continuous limit. In the limit $N \rightarrow \infty$, the random potential $V_\Gamma(\vartheta)$ defined in equation (10) in the main text converges to a Gaussian process that can be specified completely by its covariance matrix $\langle V_\Gamma(\vartheta) V_\Gamma(\vartheta') \rangle$. This can be ‘diagonalized’, alternatively expressing the random potential as a linear combination of white-noise processes $\{\zeta_n(\vartheta)\}$. One representation, as a convolution series $\sum_{n=0}^\infty (f_\Gamma^{(n)} * \zeta_n)(\vartheta)$ with kernels

$$f_\Gamma^{(n)}(\vartheta) \propto \int_0^\infty \sqrt{\Gamma^2 + \zeta^2 m_\Gamma^2 \sin^2 \vartheta} \zeta^{\alpha+1} e^{-\zeta^2/2} L_n^{(\alpha)}\left(\frac{\zeta^2}{2}\right) d\zeta, \quad (17)$$

relies on orthogonal properties of associated Laguerre polynomials. The choice

$\alpha = 1$ ensures that only $n = 0$ term survives in the limit $\Gamma = 0$. The random potential should satisfy a stochastic equation

$$\frac{d^2 V_0}{d\vartheta^2} + V_0 = \frac{2\sqrt{N}}{\pi} \zeta_0. \quad (18)$$

As a side note observe that a similar equation is obtained by taking a continuous limit in the identities that follow from elementary geometry (see Fig. 4 in the main text, $V_k = \frac{1}{2}m_k^2$):

$$\frac{V_{k+1} - V_k}{\zeta_{k+1}} - \frac{V_k - V_{k-1}}{\zeta_k} = \frac{2}{N} (\zeta_{k+1} + \zeta_k) + 4\sqrt{2V_k/N} \cos \frac{\alpha_k^+ - \alpha_k^-}{2} \cos \frac{\alpha_k^+ + \alpha_k^-}{2}. \quad (19)$$

For finite Γ , the form of the potential is modified as follows: It is convolved with a smoothing kernel of width $\Delta\vartheta \sim \Gamma$, which raises the global minimum by $O(\sqrt{N}\Gamma^{3/2})$. Additional random contributions (from $n \geq 1$) have similar scaling. This derivation is presented in greater detail in Supplementary Note 2.

Extreme value statistics. In the vicinity of the global minimum, the statistics of the classical potential is fundamentally different. The ‘returning force’ in equation (18) can be neglected; additionally the process should be conditioned on the fact that it stays above its value at $\vartheta = 0$ (no generality is lost by choosing the global minimum as the origin). This problem has been a subject of a considerable body of work^{36,37}, although important aspects ought to be revisited. Here, I present a particularly simple self-contained derivation.

A pair (χ, v) —where $\chi \propto V_0 - V_0^*$ is interpreted as the ‘coordinate’ and $v = d\chi/d\vartheta$ as the ‘velocity’ (ϑ being ‘time’)—forms a Markov process. The probability distribution $p(\vartheta; \chi, v)$ satisfies, for $\vartheta > 0$,

$$\frac{\partial p}{\partial \vartheta} = -v \frac{\partial p}{\partial \chi} + \frac{1}{2} \frac{\partial^2 p}{\partial v^2}, \quad \text{while } p(\vartheta; +0, v) = 0 (v > 0) \quad (20)$$

serves as a boundary condition for the absorbing boundary. This probability is ‘renormalized’ to condition on the fact that it survives until some $\Theta \gg \vartheta$. It becomes a conserved quantity, but the diffusion equation adds a drift, $-\partial(\log P_\Theta)/\partial v$, repelling from the boundary. The ‘survival’ probability P_Θ in the limit $\Theta \rightarrow \infty$ is, up to ‘time-reversal’, the universal asymptotic solution of (20), reduced to ordinary differential equation using scaling ansatz:

$$p(\vartheta, \chi, v) \propto \frac{\chi^{2\alpha/3} p_*(v/\chi^{1/3})}{\vartheta^\alpha}. \quad (21)$$

This exploits a fact that fractal dimensions of ‘velocity’ and ‘coordinate’ are $[v] = [\vartheta]^{1/2}$ and $[\chi] = [\vartheta]^{3/2}$. The asymptotics is dominated by solutions with the smallest possible exponent, $\alpha = 1/4$ out of the infinite set of eigenvalues for the ordinary differential equation. This matches a known value obtained with a different method³⁶.

The next step performs a change of variables, introducing ‘dimensionless’ velocity $v = v/\chi^{1/3}$, and ‘logarithmic’ coordinate $\mu = \ln \chi$. With the ‘time’ variable redefined via $d\tau = d\vartheta/\chi^{2/3}$, Markov process is described by a tuple (ϑ, μ, v) . Marginalizing out μ and ϑ in the equation for $p(\tau; \vartheta, \mu, v)$ produces Fokker–Planck equation for a stochastic motion of particle in a potential

$$U(v) = -\ln \left(\left(v/6^{1/3} \right) \text{Ai} \left(v^2/6^{2/3} \right) + \text{Ai}' \left(v^2/6^{2/3} \right) \right). \quad (22)$$

Given a particular realization of $v(\tau)$, the full process (μ, v, ϑ) is determined deterministically, by integration (see equation (12) in the main text). The construction of a realization of a random process is performed independently for $\vartheta > 0$ and $\vartheta < 0$. More detailed analysis is presented in Supplementary Note 3.

Numerical simulations rescale this random potential instead of evolving the transverse field: $\vartheta \rightarrow e^{-\frac{\Delta\Gamma}{\Gamma}} \vartheta$ and $\chi \rightarrow e^{-\frac{3\Delta\Gamma}{2\Gamma}} \chi$. The process is extended to larger values of τ as needed (details of the process for small τ , where they fall below the numerical precision, are ‘forgotten’). A fairly large range of τ is required to gather the sufficient statistics.

Data availability. The data that support the findings of this study are available from the corresponding author upon request.

References

- Nielsen, M. & Chuang, I. *Quantum Computation and Quantum Information* 10th edn (Cambridge University Press, 2011).
- Shor, P. W. Algorithms for quantum computation: discrete logarithms and factoring. *Proc. 35th Ann. Symp. Foundations of Computer Science*. (ed. Goldwasser, S.) 124–134 (IEEE Computer Society Press, 1994).
- Garey, M. R. & Johnson, D. S. *Computers and Intractability: A Guide to the Theory of NP-Completeness* (W. H. Freeman, 1979).
- Kadowaki, T. & Nishimori, H. Quantum annealing in the transverse Ising model. *Phys. Rev. E* **58**, 5355 (1998).
- Farhi, E., Goldstone, J., Gutmann, S. & Sipser, M. Quantum computation by adiabatic evolution. Preprint at <http://arxiv.org/abs/quant-ph/0001106> (2000).

- Das, A. & Chakrabarti, B. K. Quantum annealing and analog quantum computation. *Rev. Mod. Phys.* **80**, 1061–1081 (2008).
- van Dam, W., Mosca, M. & Vazirani, U. How powerful is adiabatic quantum computation? *Proc. 42nd IEEE Symp. FOCS* 279–287 (2001).
- Brooke, J., Bitko, D., Rosenbaum, F.T. & Aeppli, G. Quantum annealing of a disordered magnet. *Science* **284**, 779–781 (1999).
- Santoro, G., Martoňák, R., Tosatti, E. & Car, R. Theory of quantum annealing of spin glasses. *Science* **295**, 2427–2430 (2002).
- Heim, B., Ronnow, T. F., Isakov, S. V. & Troyer, M. Quantum versus classical annealing of Ising spin glasses. *Science* **348**, 215–217 (2015).
- Farhi, E. *et al.* A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science* **292**, 472–475 (2001).
- Young, A. P., Knysh, S. & Smelyanskiy, V. N. First order phase transition in the quantum adiabatic algorithm. *Phys. Rev. Lett.* **104**, 020502 (2010).
- Boixo, S. *et al.* Evidence for quantum annealing with more than one hundred qubits. *Nat. Phys.* **10**, 218–224 (2014).
- Ronnow, T. F. *et al.* Defining and detecting quantum speedup. *Science* **345**, 420–424 (2014).
- Boixo, S. *et al.* Computational role of multiqubit tunneling in a quantum annealer. *Nat. Commun.* **7**, 10327 (2016).
- Katzgraber, H. G., Hamze, F., Zhu, Z., Ochoa, A. J. & Munos-Bauza, H. Seeking quantum speedup through spin glasses: the good, the bad, and the ugly. *Phys. Rev. X* **5**, 031026 (2015).
- Zhu, Z., Ochoa, A. J., Schnabel, S., Hamze, F. & Katzgraber, H. G. Best-case performance of quantum annealers on native spin-glass benchmarks: How chaos can affect success probabilities. *Phys. Rev. A* **93**, 012317 (2016).
- Smelyanskiy, V. N., von Toussaint, U. & Timucin, D. A. Dynamics of quantum adiabatic evolution algorithm for number partitioning. Preprint at <http://arxiv.org/abs/quant-ph/0202155> (2002).
- Goldschmidt, Y. Y. Solvable model of the quantum spin glass in a transverse field. *Phys. Rev. B* **41**, 4858–4861 (1990).
- Jörg, T., Krzaka, F., Kurchan, J. & Maggs, A. C. Simple glass models and their quantum annealing. *Phys. Rev. Lett.* **101**, 147204 (2008).
- Jörg, T., Krzaka, F., Semerjian, G. & Zamponi, F. First-order transitions and the performance of quantum algorithms in random optimization problems. *Phys. Rev. Lett.* **104**, 207206 (2010).
- Fisher, D. S. Critical behavior of random transverse-field Ising spin chains. *Phys. Rev. B* **51**, 6411–6461 (1995).
- Miller, J. & Huse, D. Zero-temperature critical behavior of the infinite-range quantum Ising spin glass. *Phys. Rev. Lett.* **70**, 3147–3150 (1993).
- Ye, J., Sachdev, S. & Read, N. Solvable spin glass of quantum rotors. *Phys. Rev. Lett.* **70**, 4011–4014 (1993).
- Read, N., Sachdev, S. & Ye, J. Landau theory of quantum spin glasses of rotors and Ising spins. *Phys. Rev. B* **52**, 384–410 (1995).
- Altshuler, B., Krovi, H. & Roland, J. Anderson localization casts clouds over adiabatic quantum optimization. *Proc. Natl Acad. Sci. USA* **107**, 12446–12450 (2010).
- Farhi, E. *et al.* The performance of the quantum adiabatic algorithm on random instances of two optimization problems on regular hypergraphs. *Phys. Rev. A* **86**, 052334 (2012).
- Knysh, S. & Smelyanskiy, V. N. On the relevance of avoided crossings away from quantum critical point to the complexity of quantum adiabatic algorithm. Preprint at <http://arxiv.org/abs/1005.3011> (2010).
- Laumann, C. R., Moessner, R., Scardicchio, A. & Sondhi, S. L. Quantum annealing: the fastest route to quantum computation? *Eur. Phys. J. J.* **224**, 75–88 (2015).
- Krzaka, F. & Martin, O. C. Chaotic temperature dependence in a model of spin glasses. *Eur. Phys. J. B* **28**, 199–209 (2002).
- Hopfield, J. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl Acad. Sci.* **79**, 2554–2558 (1982).
- Nishimori, H. & Nonomura, Y. Quantum effects in neural networks. *J. Phys. Soc. Jpn.* **65**, 3780–3796 (1996).
- Laumann, C. R., Moessner, R., Scardicchio, A. & Sondhi, S. L. The quantum adiabatic algorithm and scaling of gaps at first order quantum phase transitions. *Phys. Rev. Lett.* **109**, 030502 (2012).
- Bovier, A., van Enter, A. C. D. & Nelderhauser, B. Stochastic symmetry-breaking in a gaussian Hopfield model. *J. Stat. Phys.* **95**, 181–213 (1999).
- Hertz, J., Krogh, A. & Palmer, R. G. *Introduction to the Theory of Neural Computation* (Addison-Wesley, 1995).
- Sinai, Y. G. On the distribution of some functions of the integral of a random walk. *Theor. Math. Phys.* **90**, 219–241 (1992).
- Groeneboom, P., Jongbloed, G. & Wellner, J. A. Integrated brownian motion, conditioned to be positive. *Ann. Prob.* **27**, 1283–1303 (1999).
- Bray, A. J. & Moore, M. A. Chaotic nature of the spin-glass phase. *Phys. Rev. Lett.* **58**, 57–60 (1987).
- Mulet, R., Pagnani, A. & Parisi, G. Against temperature chaos in naive Thouless-Anderson-Palmer equations. *Phys. Rev. B* **63**, 184438 (2001).

40. Rizzo, T. & Crisanti, A. Chaos in temperature in the Sherrington-Kirkpatrick model. *Phys. Rev. Lett.* **90**, 137201 (2003).
41. Billoire, A. Rare events analysis of temperature chaos in the Sherrington-Kirkpatrick model. *J. Stat. Mech.* **2014**, P040016 (2014).
42. Kondor, I. On chaos in spin glasses. *J. Phys. A* **22**, L163–L168 (1989).
43. Katzgraber, H. G. & Krzakala, F. Temperature and disorder chaos in three-dimensional Ising spin glasses. *Phys. Rev. Lett.* **98**, 017201 (2007).
44. Vertechi, D. & Virasoro, M. A. Energy barriers in SK spin-glass model. *J. Phys. France* **50**, 2325–2332 (1989).
45. Sommers, H. J. & Dupont, W. Distribution of frozen fields in the mean-field theory of spin glasses. *J. Phys. C* **17**, 5785–5793 (1984).
46. Venturelli, D. *et al.* Quantum annealing of fully-connected spin glass. *Phys. Rev. X* **5**, 031040 (2015).
47. Katzgraber, H. G., Hamze, F. & Andrist, R. S. Glassy chimeras may be blind to quantum speedup. *Phys. Rev. X* **4**, 021008 (2014).
48. Rieger, H. & Young, A. P. Zero-temperature quantum phase transition of a two-dimensional Ising spin glass. *Phys. Rev. Lett.* **72**, 4141–4144 (1994).
49. Martin-Mayor, V. & Hen, I. Unraveling quantum annealers using classical hardness. *Sci. Rep.* **5**, 15324 (2015).
50. Farhi, E., Goldstone, J. & Gutmann, S. Quantum adiabatic annealing algorithms versus simulated annealing. Preprint at <http://arxiv.org/abs/quant-ph/0201031> (2002).
51. Mézard, M., Parisi, G. & Zecchina, R. Analytic and algorithmic solution of random satisfiability problems. *Science* **297**, 812–815 (2002).

Acknowledgements

I would like to thank Vadim Smelyanskiy for useful discussions. This work was supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IAA 145483, and by the Air Force Research Laboratory (AFRL) Information Directorate under grant F4HBKC4162G001. The views and conclusions contained herein are those of the author 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.

Author contributions

S.K. performed the research and wrote the manuscript.

Additional information

Supplementary Information accompanies this paper at <http://www.nature.com/naturecommunications>

Competing financial interests: The authors declare no competing financial interests.

Reprints and permission information is available online at <http://npublishing.com/reprintsandpermissions/>

How to cite this article: Knysh, S. Zero-temperature quantum annealing bottlenecks in the spin-glass phase. *Nat. Commun.* **7**:12370 doi: 10.1038/ncomms12370 (2016).



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>

© The Author(s) 2016