

Learning-Driven Annealing with Adaptive Hamiltonian Modification for Solving Large-Scale Problems on Quantum Devices

Sebastian Schulz¹, Dennis Willsch^{1,2}, and Kristel Michielsen^{1,3,4}

¹Jülich Supercomputing Centre, Institute for Advanced Simulation,
Forschungszentrum Jülich, 52425 Jülich, Germany

²Faculty of Medical Engineering and Technomathematics, University of Applied Sciences Aachen,
52428 Jülich, Germany

³AIDAS, 52425 Jülich, Germany

⁴RWTH Aachen University, 52056 Aachen, Germany
2025-10-19

We present Learning-Driven Annealing (LDA), a framework that links individual quantum annealing evolutions into a global solution strategy to mitigate hardware constraints such as short annealing times and integrated control errors. Unlike other iterative methods, LDA does not tune the annealing procedure (e.g. annealing time or annealing schedule), but instead learns about the problem structure to adaptively modify the problem Hamiltonian. By deforming the instantaneous energy spectrum, LDA suppresses transitions into high-energy states and focuses the evolution into low-energy regions of the Hilbert space. We demonstrate the efficacy of LDA by developing a hybrid quantum-classical solver for large-scale spin glasses. The hybrid solver is based on a comprehensive study of the internal structure of spin glasses, outperforming other quantum and classical algorithms (e.g., reverse annealing, cyclic annealing, simulated annealing, Gurobi, Toshiba’s SBM, VeloxQ and D-Wave hybrid) on 5580-qubit problem instances in both runtime and lowest energy. LDA is a step towards practical quantum computation that enables today’s quantum devices to compete with classical solvers.

1 Introduction

Combinatorial optimization problems (COPs) are ubiquitous in computer science, with important applications in finance [29], scheduling [52], machine learning [8, 9], computational biology [41], and operations research [11]. For many such problems, finding the optimal solution is equivalent to finding the ground

Sebastian Schulz: se.schulz@fz-juelich.de, Corresponding author:
Sebastian Schulz

state of an associated Ising spin-glass system [28]. The hardness of COPs is related to the presence of opposing spin interactions, leading to frustration and a glass phase. The latter is characterized by the presence of many low-energy local minima, that are separated by large energy barriers. This makes solving large-scale COPs often intractable on classical computers, quickly exceeding runtimes of 24 hours.

With the advent of new generations of quantum annealers from D-Wave Systems Inc. comprising more than 5000 qubits, a promising approach is the use of quantum annealing (QA) [27, 38, 55]. Inspired by the cooling of physical systems, QA uses quantum fluctuations caused by a transverse field to navigate the energy landscape of the spin system [42]. If the quantum evolution is performed adiabatically, the system arrives in the ground state of the spin glass. However, if the evolution is too fast, the system scatters into higher energy states through a series of quantum phase transitions (QPTs, used in the sense of Refs. [7, 35, 44, 45]). Hence, on real devices, finite annealing times and integrated control errors limit the performance, often making deep low-energy states (including the ground state) inaccessible [4, 50]. To address this issue, several authors recently investigated iterative procedures, such as reverse annealing [12, 31, 40, 54] and cyclic annealing [39, 51, 57], in which the annealing is repeatedly initialized from the best known classical state. In doing so, the hope is that QA gradually converges to the ground state, but in practice, these procedures often fail, as they become stuck in high-energy valleys [13, 36].

We show that systematically learned modifications to the problem Hamiltonian itself can greatly improve the performance of iterative QA on noisy hardware. With Learning-Driven Annealing (LDA), we propose a framework to link individual QA runs into a global solution strategy to mitigate hardware constraints such as finite annealing times and integrated control errors [16]. LDA works by analyzing the states

sampled from QA to adaptively modify the problem Hamiltonian to the information learned about the energy landscape. This energetically isolates low-energy valleys in the instantaneous energy spectrum, focussing subsequent annealing runs into a promising region of the Hilbert space. As a result, the QPU can reach deeper regions of the energy spectrum that are otherwise inaccessible. While we do not claim a quantum advantage per se, LDA enables current NISQ devices to tackle COPs more efficiently as long as they can be embedded onto their hardware.

We demonstrate the capabilities of LDA by designing a hybrid quantum-classical solver for large-scale spin glasses, that is based on alternating local and global search protocols. Using the D-Wave Advantage1 5.4 QPU [1] in Jülich, Germany, we benchmark our hybrid solver against leading quantum and classical algorithms, including reverse annealing [12], cyclic annealing [51, 57], simulated annealing using JUPTSA [20], Gurobi [19], Toshiba's SBM [49], VeloxQ [21, 37] and D-Wave hybrid [14], on 5580-qubit NAT-7 [24, 30] spin glasses. The results show that our hybrid solver outperforms all competing strategies in both runtime and lowest energy. LDA is a step towards practical quantum computation that enables today's quantum annealers to compete with classical solvers.

The remainder of this paper is organized as follows: In Sec. 2, we introduce the spin-glass formulation and provide insights into its underlying structures. Section 3 concerns adiabatic quantum evolutions and discusses its limitations on modern quantum annealers. In Sec. 4, we introduce the LDA framework, showcasing how learned information from sampled states can be used to guide the annealing evolution by modifying the problem Hamiltonian. In Sec. 5, we demonstrate how LDA can be used in advanced algorithms by designing a hybrid solver for spin glasses based on alternating local and global search protocols. Finally, in Sec. 6, we benchmark our hybrid solver against various classical and quantum solvers, followed by a summary of our findings in Sec. 7.

2 Spin-glass problems

Spin glasses (SG) are magnetic systems of quenched disorder, imposing conflicts between interacting magnetic moments [10]. In the Edwards-Anderson model of SG [17], N Ising spins σ_i^z , $i \in \{0, \dots, N-1\}$, are placed on the sites of a regular graph, with the Hamiltonian defined as

$$H = \sum_{i < j} J_{ij} \sigma_i^z \sigma_j^z + \sum_i h_i \sigma_i^z. \quad (1)$$

Here J_{ij} with $i < j$ denotes the coupling strength between spins i and j , and h_i represents a locally applied external magnetic field bias on spin i .

We consider both triangular lattices with periodic boundary conditions (see Fig. 1a) and the Pegasus graph used by the D-Wave Advantage1 5.4 quantum annealer [1]. In the following, we choose $|h_i| \ll |J_{ij}|$ to ensure the generation of complex SG [25]. Notably, for $h_i = 0$, H exhibits a global \mathbb{Z}_2 symmetry.

Throughout this study, we use Greek letters, e.g. $\alpha = \alpha_{N-1} \cdots \alpha_1 \alpha_0$, to represent a generic eigenstate of Eq. (1), with $\alpha_i = \pm 1$ denoting the polarization of the i -th spin. We also use the notation $\alpha^* = \alpha_{N-1}^* \cdots \alpha_1^* \alpha_0^*$ to refer to a local (in Hamming distance) minimum of Eq. (1). The eigenenergy of a state α is given by $E_\alpha = \sum_{i < j} J_{ij} \alpha_i \alpha_j + \sum_i h_i \alpha_i$ and its bit string representation is written as $\mathbf{a} = \mathbf{a}_{N-1} \cdots \mathbf{a}_1 \mathbf{a}_0$, with $\mathbf{a}_i = (1 + \alpha_i)/2$.

For a given state α , we define the set of satisfied (i.e., non-frustrated) couplers and biases as

$$\mathcal{J}^\alpha = \{(i, j) \mid J_{ij} \alpha_i \alpha_j < 0\}, \quad (2a)$$

$$\mathcal{H}^\alpha = \{i \mid h_i \alpha_i < 0\}. \quad (2b)$$

These sets consist of those couplings and biases of Eq. (1) that are satisfied in the state α , meaning that the corresponding terms in the Hamiltonian contribute to a reduction of the total energy of the system.

Intuitively, the presence of competing spin-spin interactions in Eq. (1) causes frustration, which prevents the ground state from establishing a simple long-range (anti-)ferromagnetic order. Instead, the spins align in random directions, forming a glass phase [47]. This phase is characterized by the existence of exponentially many low-energy local minima, where a *local minimum* refers to a state for which flipping a single (or a few) spins always increases the energy. The energy spectrum of these local minima typically has small gaps, with macroscopic high and wide energy barriers separating them [34]. The latter implies that $\mathcal{O}(N)$ spins would need to be flipped in order to escape a local minimum [10]. Despite the complexity of the energy landscape, the formation of local minima stems from a general structure in the interactions of the spins. Local differences in the spin-spin couplings and biases lead to the formation of spin domains $\mathcal{D}_1, \dots, \mathcal{D}_{N_{\text{dom}}} \subset \{0, \dots, N-1\}$, representing groups of highly correlated spins [5, 10].

2.1 Spin domains

We define a *spin domain* $\mathcal{D}_k \subset \{0, \dots, N-1\}$, with $k = 1, \dots, N_{\text{dom}}$, as a maximally large connected sub-spin glass of Eq. (1), such that its energy landscape is *ergodic* with respect to spin-reversal transformations. This means that the ground state of the domain is accessible from any state through a sequence of decreasing energies using solely single spin flips and state inversions (i.e. flipping all spins). In computational terms, *ergodic* means that the global minimum

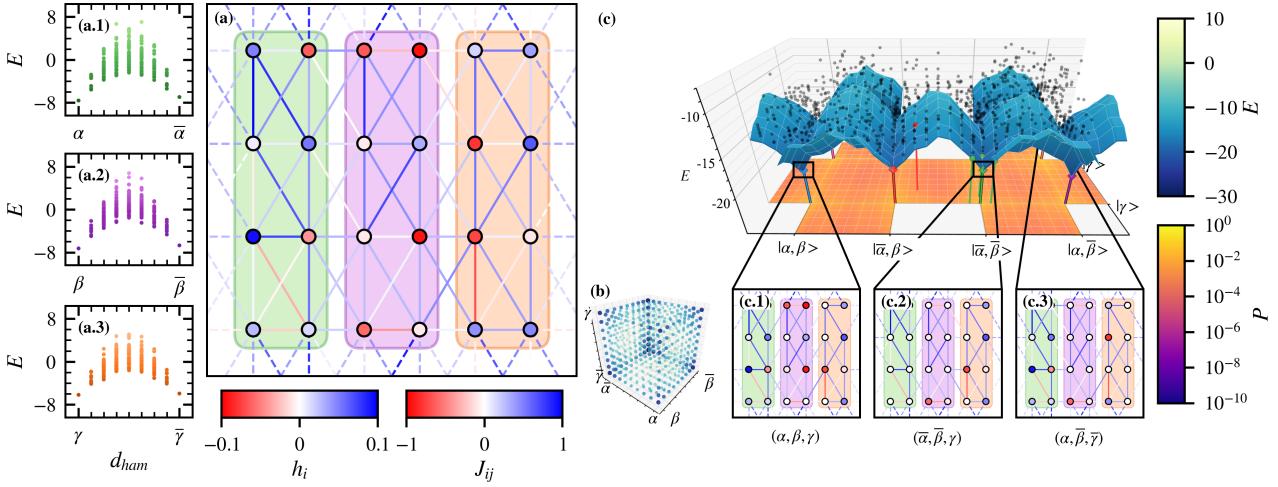


Figure 1: A 24-qubit spin-glass with $N_{\text{dom}} = 3$. (a) Spin-glass graph with the biases h_i as nodes and couplers J_{ij} as edges. Dashed edges indicate periodic boundary conditions. The colored rectangles represent the spin domains: \mathcal{D}_1 (green), \mathcal{D}_2 (purple), and \mathcal{D}_3 (orange). Each domain has two distinct minima: α and $\bar{\alpha}$ (\mathcal{D}_1), β and $\bar{\beta}$ (\mathcal{D}_2), γ and $\bar{\gamma}$ (\mathcal{D}_3). Panels (a.1-3) show the eigenenergies E of the corresponding domains as a function of the Hamming distance d_{ham} to the respective ground state α, β, γ . (b) Hypercube representation of the Hilbert space, with each axis representing the Hamming distance d_{ham} from the ground state (α, β, γ) of the respective spin domains \mathcal{D}_k . Color indicates the lowest eigenenergy at each point. The corners of the hypercube correspond to the $2^{N_{\text{dom}}} = 8$ local minima of the spin glass. (c) Energy landscape of the spin glass obtained by unfolding the hypercube facets. Each point is projected onto the closest facet (smallest Hamming distance), with the blue surface representing the lowest eigenenergies on these facets. The eight local minima are marked by colored diamonds at the facets' corners, with the states surrounding it being the energy valleys. Additionally, black dots indicate 2000 samples from a quantum annealing simulation with annealing time $T = 10$, with the probability density shown on the lower orange plane. Panels (c.1-3) highlight the specific biases h_i and couplers J_{ij} satisfied by three of the eight local minima, with the unsatisfied terms set to 0 (white). The color scheme matches that of panel (a).

of the domain can be reached via a greedy search algorithm. Note that a domain can consist of a single spin if the local magnetic field h_i dominates the interaction. An example 24-qubit spin glass with $N_{\text{dom}} = 3$ domains is shown in Fig. 1a. The respective energy spectra of the spin domains are given in panels (a.1-3).

Each spin domain \mathcal{D}_k identifies a trivial sub-spin glass, defined by the quadratic terms $i, j \in \mathcal{D}_k$ within the domain,

$$\sum_{i,j \in \mathcal{D}_k} J_{ij} \sigma_i^z \sigma_j^z. \quad (3)$$

This sub-spin glass has two degenerate ground states $\beta^{*(k)}$ and $\bar{\beta}^{*(k)}$, due to the \mathbb{Z}_2 symmetry of the spin-spin couplers, with $\bar{\beta}^{*(k)}$ denoting the inverse state of $\beta^{*(k)}$ (i.e., all spins flipped). By combining the ground states of all domains, the $2^{N_{\text{dom}}}$ local minima of the full spin glass can be constructed. Consequently, these local minima form an N_{dom} -dimensional hypercube, with the edges denoting state inversions of single spin domains (see Fig. 1b).

The complexity of the spin-glass energy landscape stems from the domain borders, where spin domains couple to each other (i.e. spin-spin couplings J_{ij} between domains) and to the local biases h_i . These interactions lift the degeneracies of the domains' ground states, assigning energy shifts to the local minima and forming a hierarchy on the hypercube [22].

For the three-dimensional hypercube shown in Fig. 1b, Fig. 1c visualizes the energy landscape obtained by unfolding its facets. Note that the hardness of spin glasses arises from the lack of knowledge about the exact domains \mathcal{D}_k . If the distribution of spin domains were known, the complexity of the problem would be reduced from 2^N to $2^{N_{\text{dom}}}$ with the number of domains $N_{\text{dom}} \ll N$.

2.2 Distance measures

For the following description of LDA, a distance measure between arbitrary spin configurations α and β is required. A commonly used metric for this purpose is the Edward-Anderson order parameter [17], measuring the overlap of the states α and β as

$$q_{EA}(\alpha, \beta) = \frac{1}{N} \sum_{i=0}^{N-1} \alpha_i \beta_i. \quad (4)$$

By definition, $q_{EA} \in [-1, 1]$, with $q_{EA} = 1 \Leftrightarrow \alpha = \beta$ and $q_{EA} = -1 \Leftrightarrow \alpha = \bar{\beta}$. Note that q_{EA} is directly related to the Hamming distance d_{ham} between the bit strings \mathbf{a} and \mathbf{b} via $d_{\text{ham}} = N(1 - q_{EA})/2$.

While q_{EA} is a sufficient tool for the statistical analysis of phase transitions in disordered systems [26], it is indifferent to the distance between the energies E_α and E_β of the states. This means that two states

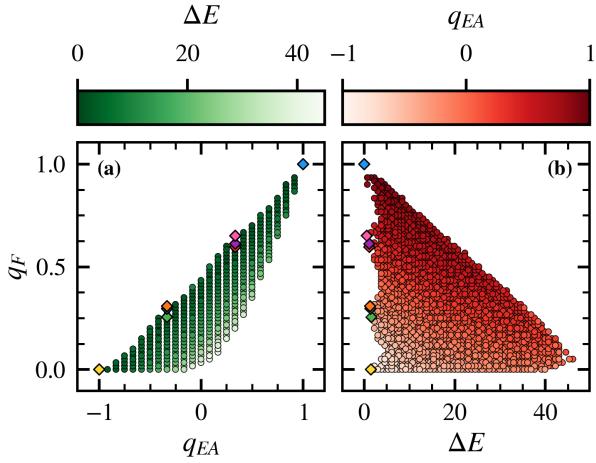


Figure 2: Comparison of the quantity $q_F(\alpha^*, \beta)$ to (a) the Edward-Anderson order parameter $q_{EA}(\alpha^*, \beta)$ and (b) the energy difference $\Delta E = E_\beta - E_{\alpha^*}$ for states β w.r.t. the global minimum α^* of the 24-qubit spin-glass shown in Fig. 1a. The colored diamonds mark the eight local minima (see Fig. 1c).

with the same order parameter q_{EA} can have significantly different energies, depending on the amount of frustration that they cause. To compensate for this, we propose a quantity $q_F(\alpha, \beta)$ that evaluates the similarity between a reference state α and a state β based on the sets of satisfied couplers and biases (see Eq. (2)). We define q_F as the fraction of couplers and biases that are simultaneously satisfied by α and β (i.e., $\mathcal{J}^\alpha \cap \mathcal{J}^\beta$ and $\mathcal{H}^\alpha \cap \mathcal{H}^\beta$) to the total set of satisfied terms by α (i.e., \mathcal{J}^α and \mathcal{H}^α),

$$q_F(\alpha, \beta) = \frac{\sum_{\substack{(i,j) \in \mathcal{J}^\alpha \cap \mathcal{J}^\beta \\ \text{with } (\alpha_i, \alpha_j) = (\beta_i, \beta_j)}} |J_{ij}| + \sum_{i \in \mathcal{H}^\alpha \cap \mathcal{H}^\beta} |h_i|}{\sum_{(i,j) \in \mathcal{J}^\alpha} |J_{ij}| + \sum_{i \in \mathcal{H}^\alpha} |h_i|}. \quad (5)$$

By definition, $q_F \in [0, 1]$, with $q_F = 1$ (0) denoting that $\mathcal{H}^\beta \subset \mathcal{H}^\alpha$ ($\mathcal{H}^\beta \cap \mathcal{H}^\alpha = \emptyset$) and $\mathcal{J}^\beta \subset \mathcal{J}^\alpha$ ($\mathcal{J}^\beta \cap \mathcal{J}^\alpha = \emptyset$). We remark that in the singular case $\mathcal{H}^\alpha = \mathcal{J}^\alpha = \emptyset$, we define $q_F = 1$. The additional condition $\alpha_i = \beta_i$ in the first summation of Eq. (5) ensures that the spins in the simultaneously satisfied couplers J_{ij} are also aligned with one another, $(\alpha_i, \alpha_j) = (\beta_i, \beta_j)$. This is necessary because $J_{ij}\alpha_i\alpha_j = J_{ij}\beta_i\beta_j < 0$ would also be satisfied if $(\alpha_i, \alpha_j) = (\beta_i, \beta_j)$. Hence, the condition lifts the spin-reversal symmetry of the spin domains, in order to account for the Hamming distance between the states.

The definition of q_F is based on the observation that each local minimum α^* can be uniquely identified (up to state inversion in case of $h_i = 0$) by the sets of satisfied couplers \mathcal{J}^{α^*} and biases \mathcal{H}^{α^*} (see Fig. 1c.1-3, where the local minima satisfy different couplers between the domains). By ranking each state according to its energy in the sub spin-glass ($\mathcal{J}^{\alpha^*}, \mathcal{H}^{\alpha^*}$), the

parameter q_F integrates both the Hamming distance (q_{EA}) and the energy distance (ΔE) from α^* into the similarity measure (see Fig. 2a). This allows q_F to distinguish the energy valley of α^* (i.e. the neighborhood of states around α^* , where some domains are in an excited state) from the other local minima in Fig. 1c. As a result, only states with both a small Hamming distance and energy distance are considered close to α^* , making q_F indifferent to low-energy states of other energy valleys (see peaks at $\Delta E \approx 0$ in Fig. 2b).

A key property of q_F is its asymmetry (i.e. $q_F(\alpha, \beta) \neq q_F(\beta, \alpha)$), since Eq. (5) considers only terms that are satisfied by the reference state α , but ignores additional terms satisfied by β . As a consequence, a state β deeper in the energy valley, with $\mathcal{J}^\beta \supset \mathcal{J}^\alpha$ and $\mathcal{H}^\beta \supset \mathcal{H}^\alpha$, typically has a larger q_F value than a higher-energy state, which violates terms in \mathcal{H}^α and \mathcal{J}^α . To make the meaning of q_F more intuitive, an explicit example of this asymmetry for a 16 spin-glass instance is shown in App. A. This means that q_F has an intrinsic bias towards lower-energy states, which will be crucial for the definition of LDA in Sec. 4.

3 Adiabatic quantum annealing

Adiabatic quantum annealing, as proposed by Farhi et al. [18], is a procedure for solving NP-hard combinatorial optimization problems (COP) through quantum fluctuations. With the search space of the COP encoded in the eigenspectrum of a problem Hamiltonian H_P (e.g. Eq. (1)), the system is complemented by a driving Hamiltonian H_D that is non-diagonal in the eigenbasis of H_P . Since the spin-glass Hamiltonian Eq. (1) is diagonal in the computational basis, we consider the transverse field Ising model

$$H_{QA}(t) = H_P + \Gamma(t) \cdot H_D \quad \text{with} \quad H_D = - \sum_{i=0}^{N-1} \sigma_i^x, \quad (6)$$

which is closely related to the design of D-Wave quantum annealers (see App. B). $\Gamma(t)$ is the annealing schedule, defining the relative strength between H_P and H_D over time $0 \leq t \leq T$. Initially, $\Gamma(t=0) \gg J$, such that the instantaneous ground state (GS) is approximately given by $|+\rangle^{\bigotimes N}$ (where $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$), which is initially separated by a large energy gap $\propto \Gamma(0)$ from the rest of the spectrum. Operating the system at a temperature that is much smaller than the scale set by $\Gamma(0)$ is supposed to initialize it in the GS of H_D . To reach the GS of H_P , $\Gamma(t)$ is decreased adiabatically to zero as $t \rightarrow T$, allowing the system to follow the instantaneous GS of $H_{QA}(t)$ [3, 18]. In order to prevent Landau-Zener transitions into excited states during this process, the adiabatic theorem requires $T^{-1} \ll \Delta_{\min}^2$ [6], where

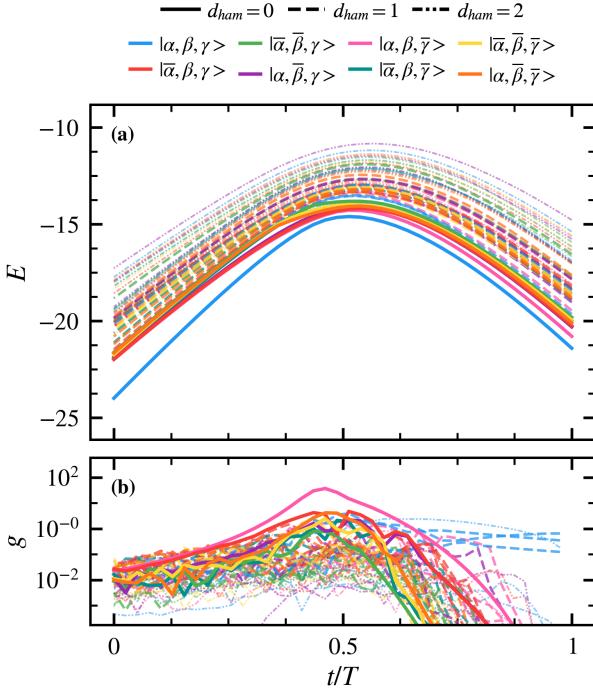


Figure 3: (a) Instantaneous energy spectrum of the transverse field Ising model (see Eq. (6)) for the 24-qubit spin-glass instance shown in Fig. 1a, using a linear annealing schedule $\Gamma(t)$. The plot presents a selection of energy levels corresponding to the 8 valleys in the energy landscape (see Fig. 1c). For each valley, the local minimum (solid lines) is shown along with three states at Hamming distance $d_{ham} = 1$ (dashed line) and three states at $d_{ham} = 2$ (dash-dotted lines), representing the lowest-energy states along the edges of the hypercube. (b) Adiabatic ratios (see Eq. (7)) for the states presented in panel (a). The 24-qubit eigenstates and eigenvalues in the vicinity of each local minimum were obtained numerically by solving the time-dependent Schrödinger equation (see e.g. [53]) for reverse-annealing starting in the given states.

Δ_{min} denotes the minimal energy gap between the two lowest instantaneous eigenstates of $H_{QA}(t)$.

The minimum annealing time T is the main measure for the computational complexity of adiabatic quantum annealing. Significant evidence exists that the presence of first-order quantum phase transitions (QPTs) during the adiabatic evolution causes an exponentially fast closure of Δ_{min} with the system size N , resulting in exponentially long annealing times [7, 44]. The authors in [5] argue that Δ_{min} scales inversely proportional to the square root of the number of competing local minima (and states near them) with energies close to the global minimum and a large Hamming distances d_{ham} . Regarding the spin-glass instances under investigation, the number of local minima scales exponentially, $2^{N_{dom}}$, with the number of domains $N_{dom} \propto N$.

Figure 3a presents the instantaneous energy spectrum of the 24-qubit spin-glass instance shown in

Fig. 1a, with the adiabatic ratio [33]

$$g(t) = \frac{|\langle m(t) | \dot{H}_{QA}(t) | GS(t) \rangle|}{|E_m(t) - E_{GS}(t)|^2} \quad (7)$$

depicted in Fig. 3b for the states marked in Fig. 1c. Here, $|m(t)\rangle$ ($|GS(t)\rangle$) denotes the instantaneous m -th excited (ground) state, $\dot{H}_{QA}(t)$ is the t derivative of the instantaneous Hamiltonian and $E_m(t)$ ($E_{GS}(t)$) corresponds to the eigenenergies of the m -th excited (ground) state.

In agreement with [2, 45], the system undergoes a second-order QPTs at $t_c/T \approx 0.45$, by evolving from the paramagnetic phase (localized in the σ_x basis) to an ordered phase (localized in the σ_z basis). As a result, the system transitions from a delocalized state, where it forms a superposition across all valleys, to a semi-localized state confined to a single valley [7]. During the QPT, all valleys are coupled, as indicated by the peaks in g for the 8 global and local minima (see solid lines in Fig. 3b). Additionally, the instantaneous energy spectrum exhibits transitions between states within the same valley. Since these states differ by small Hamming distances, they remain coupled after the second-order QPTs for $t > t_c$ (see blue dashed and dash-dotted curves in Fig. 3b).

While the second-order QPT is generally associated with a polynomial scaling of T , the presence of additional anti-crossings at $t > t_c$ between semi-localized states in large spin-glass instances induces first-order QPTs [2, 7]. These transitions correspond to quantum tunneling events, where the system relocates entirely from one valley to another, leading to energy gaps Δ that close exponentially with the Hamming distance d_{ham} between valleys. Since d_{ham} scales linearly with the system size N in our spin-glass instances, this results in an exponential scaling of T [44, 45]. Recently, locally optimized annealing procedures have been proposed, which reduce T by decelerating the annealing schedule $\Gamma(t)$ at the vicinity of the QPT [43, 48]. However, these strategies can only provide a quadratic speed up [5]. Since the required annealing times still exceed the accessible range on real devices, it is in practice impossible to solve large spin glasses using a single annealing run.

To overcome this bottleneck, several authors proposed the use of iterative protocols to incorporate classical information about the ground state into the annealing process [39, 40, 51]. The reverse annealing protocol, as proposed by Chancellor [12] and Yamashiro et al. [54], has been implemented on D-Wave QPUs. It starts by initializing the system in a computational basis state (e.g. obtained by a previous anneal or a classical heuristic), with $\Gamma = 0$. Quantum fluctuations are slowly introduced by increasing Γ to an intermediate value $\tilde{\Gamma}$ at $t = T_s$, where the process is paused for a period τ until $t = T_e$. The protocol finishes by annealing back to $\Gamma = 0$ and measuring the qubits. By repeating this cycle and initializing each

reverse anneal by the measurement of the previous, iterative reverse annealing (IRA) is hoped to move the system closer to the true ground state and overcome the above-mentioned limitation of an exponential annealing time T .

In practice, however, reverse annealing performs a dissipative local search into the neighborhood of the initial state [13]. The spread of this search is controlled by the inversion point $\tilde{\Gamma}$ and the QPTs it traverses. If no QPT is crossed, the system spreads locally (in Hamming distance) within its energy valley. However, if $\tilde{\Gamma}$ is chosen large enough, such that a QPT is crossed, the system explores the search space globally by tunneling into other energy valleys [12]. Note that T must be sufficiently short, such that the process is non-adiabatic, because otherwise the system would simply return to its initial state.

In a closed system, there is no direct mechanism in the reverse annealing protocol to favor the search of lower-energy states [13]. Instead, the system spreads uniformly in energy around the initial state, gravitating (over multiple iterations) towards an equilibrium state, where the number of states with lower and higher energies balance. This makes it exponentially hard to find the ground state for $N \gg 1$.

On real devices, however, the QPU is always in interaction with its surrounding environment. In such open systems, thermal relaxation moderates the transition between neighboring states [31, 32]. At sufficiently low temperatures, this causes a repopulation of lower-energy states during the pause, which allows reverse annealing to improve on the initial state. Passarelli et. al. [36] argue that thermal relaxation is the main mechanism driving the performance of reverse annealing. This means that the success of IRA depends mainly on the effective temperature of the QPU. This gives an intuition as to why the D-Wave Advantage1 5.4 QPU operating at $\approx 16.4mK$ [1] is not able to find the ground state of the large spin glasses under investigation through IRA (see Sec. 6).

4 Learning-Driven Annealing

With Learning-Driven Annealing (LDA) we introduce a framework for linking individual QA runs into a global solution strategy to mitigate hardware limitations, such as finite annealing times and integrated control errors. LDA differs from other iterative strategies in that it does not tune the annealing procedure (e.g. annealing time or annealing schedule), but instead learns about the problem structure (i.e. spin domains) to systematically *modify* the problem Hamiltonian. LDA acts on the instantaneous energy spectrum in order to change the strength of QPTs between the energy valleys. This focuses the annealing evolution into a low-energy region of the Hilbert space by suppressing QPTs into high-energy valleys. LDA achieves this by replacing the problem Hamiltonian

with a so-called *feature Hamiltonian* in the forward anneal. The construction of this Hamiltonian is based on the definition of the quantity q_F (see Eq. (5)) and adapts the strength of biases and couplers in H_P based on learned information about the domain structure to energetically isolate low-energy valleys in the instantaneous energy spectrum. This allows the annealing evolution to successively reach deeper states in the energy landscape.

The feature Hamiltonian, denoted as $H_F(\alpha)$, is defined w.r.t. a classical state $\alpha = \alpha_{N-1} \cdots \alpha_1 \alpha_0$ for a spin glass instance $P = (h, J)$ as

$$H_F(\alpha) = \sum_{(i,j) \in \mathcal{J}^\alpha} K_{ij} + \sum_{i \in \mathcal{H}^\alpha} h_i \sigma_i^z \quad (8)$$

$$K_{ij} = -\frac{|J_{ij}|}{2} [\alpha_i \alpha_j \sigma_i^z \sigma_j^z + \alpha_i \sigma_i^z + \alpha_j \sigma_j^z]. \quad (9)$$

\mathcal{H}^α and \mathcal{J}^α denote the sets of satisfied biases and couplers, respectively (see Eq. (2b) and (2a)). $H_F(\alpha)$ transforms the energy spectrum of P by retaining only terms that are satisfied by α in P . This *deforms* the energy landscape according to the quantity q_F (see Eq. (5)), by arranging the states based on both their Hamming distance and energy distance from α , with α becoming the new ground state (see Thm. C.1).

In Eq. (8) we use the transformation $J_{ij} \sigma_i^z \sigma_j^z \rightarrow K_{ij}$ to lift the \mathbb{Z}_2 symmetry of the spin-spin couplers J_{ij} by introducing a penalty if either of the two spins i and j does not match the reference state α , i.e.,

$$K_{ij} |\alpha_i, \alpha_j\rangle = -\frac{3}{2} |J_{ij}|, \quad K_{ij} |\alpha_i, \overline{\alpha_j}\rangle = \frac{1}{2} |J_{ij}| \quad (10)$$

$$K_{ij} |\overline{\alpha_i}, \alpha_j\rangle = \frac{1}{2} |J_{ij}|, \quad K_{ij} |\overline{\alpha_i}, \overline{\alpha_j}\rangle = \frac{1}{2} |J_{ij}|. \quad (11)$$

Note that a factor of $1/2$ is introduced in Eq. (9) to ensure that the weighting of the modified couplers remains consistent with the original problem Hamiltonian, with energy gaps of $2|J_{ij}|$. This transformation corresponds to the condition $\alpha_i = \beta_i$ in Eq. (5) and removes the degeneracy of the spin domains (i.e., each domain now has only one GS). Since the domain's degeneracy is responsible for the existence of energy valleys (see Fig. 1c), removing the \mathbb{Z}_2 symmetry flattens the energy landscape, while the restriction of $H_F(\alpha)$ to \mathcal{H}^α and \mathcal{J}^α deforms the landscape in the direction of the reference state α (see Fig. 4b.1). As a result, $H_F(\alpha)$ is frustration-free (i.e., all couplers and biases are satisfied in the GS α) and the modified energy landscape now increases monotonically with q_F from α (see Fig. 2). We remark that the monotonicity holds only for the feature Hamiltonian H_F , which will be used henceforth to construct the updated problem Hamiltonian (see Eq. 15 below).

Since the similarity measure q_F is asymmetric, a state β that lies deeper in the energy valley than α (i.e., it is closer to the local minimum α^* in terms

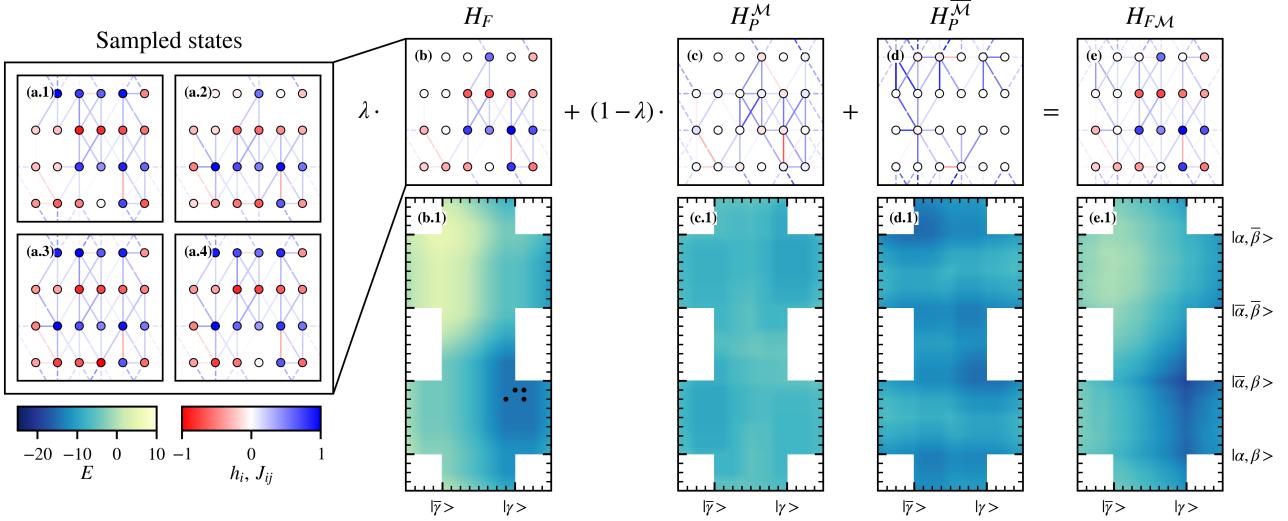


Figure 4: Construction of H_{FM} (e, Eq. (15)) for the 24-qubit spin-glass instance shown in Fig. 1, with $\lambda = 0.1$. The feature Hamiltonian H_F (b) is derived from four sampled states (a.1-4) using Eq. (12) and the bitmask \mathcal{M} , with $m_i = 1$ if the sampled states agree in the i -th bit (this is the case for all non-white nodes shown in panel (b)). Panels (c) and (d) depict subsets of the original problem Hamiltonian, constructed using Eq. (14), with (c) applying to the terms included in H_F , and (d) addressing the remaining biases and couplers. Panels (b-e.1) illustrate the corresponding energy landscapes of the Hamiltonians in (b-e), with the four sampled states shown as black dots in (b.1).

of q_F) satisfies most of the biases and couplers already satisfied by α , along with additional terms. Consequently, $\mathcal{H}^\alpha \subsetneq \mathcal{H}^\beta$ and $\mathcal{J}^\alpha \subsetneq \mathcal{J}^\beta$, such that $q_F(\alpha, \beta) \approx 1$. In this case, β is either a degenerate ground state or a weakly excited state of the feature Hamiltonian $H_F(\alpha)$. In contrast, a higher-energy state in P necessarily frustrates terms satisfied by α , making it a highly excited state of $H_F(\alpha)$ (cf. App. A for a 16-qubit example). As a result, the modified energy landscape is nearly flat in the region of Hilbert space lying deeper in the energy valley than α , termed the *search region*, while outside this region the energy increases monotonically with the Hamming distance from α (see Thm. C.2 in App. C). The size of the search region depends on the elements of \mathcal{H}^α and \mathcal{J}^α , referred to as *features*. These features represent the information provided by α about the location of a nearby local minimum α^* . If α is a high-energy state, most biases and couplers in P are frustrated, leading to a highly degenerate $H_F(\alpha)$ with a large search region, as \mathcal{H}^α and \mathcal{J}^α contain only few elements. This situation changes as α approaches a local minimum α^* . As the overlap between the satisfied terms of α and α^* increases, $H_F(\alpha)$ incorporates more features characterizing the local minimum, resulting in a progressively deformed energy landscape that narrows the search region around α^* .

The monotonicity of the modified energy landscape outside the search region suppresses all avoided level crossings in the original instantaneous spectrum that would drive the system out of the search region, as they are now separated by large energy gaps. In particular, second-order QPTs that induce delocalized-

to-localized transitions into competing valleys are removed, since these valleys are flattened in the modified landscape. Thus, the feature Hamiltonian $H_F(\alpha)$ energetically isolates an annealing path (i.e., a sequence of transitions and non-transitions at avoided level crossings) in the original instantaneous spectrum that confines the system to the search region. The length of this annealing path (i.e., the number of transitions fixed by $H_F(\alpha)$) depends on the features provided by α and is therefore related to the size of the search region.

In a forward anneal, the feature Hamiltonian confines the dynamics to the search region. However, because the energy landscape remains nearly flat within this region, it cannot resolve the original valley structure and identify the minimum α^* . As a result, the system evolves into an almost equal superposition over the search space. To address this, we reintroduce the original problem Hamiltonian H_P within the search region, yielding:

$$H_F(\alpha, \mathcal{M}) = \sum_{\substack{(i,j) \in \mathcal{J}^\alpha \\ \text{with } m_i = m_j = 1}} \lambda \cdot K_{ij} + \sum_{\substack{i \in \mathcal{H}^\alpha \\ \text{with } m_i = 1}} \lambda \cdot h_i \sigma_i^z, \quad (12)$$

$$H_P(\mathcal{M}) = \sum_{i < j} \begin{cases} (1-\lambda) \cdot J_{ij} \sigma_i^z \sigma_j^z, & \text{if } m_i, m_j = 1 \\ J_{ij} \sigma_i^z \sigma_j^z, & \text{else} \end{cases} \quad (13)$$

$$+ \sum_i \begin{cases} (1-\lambda) \cdot h_i \sigma_i^z, & \text{if } m_i = 1 \\ h_i \sigma_i^z, & \text{else} \end{cases} \quad (14)$$

$$H_{FM}(\alpha, \mathcal{M}) = H_F(\alpha, \mathcal{M}) + H_P(\mathcal{M}) \quad (15)$$

with the bitmask $\mathcal{M} = m_{N-1} \dots m_1 m_0$ controlling the mixing of H_P and $H_F(\alpha)$ and the mixing parameter $\lambda \in [0, 1]$ specifying the strength of the deformation of the energy landscape introduced by H_F . Note that for $\lambda = 0$, one recovers $H_{F\mathcal{M}}(\alpha, \mathcal{M}) = H_P$, independent of \mathcal{M} . $H_{F\mathcal{M}}$ restores the original energy relationship between the states in a local subspace of the hypercube defined by \mathcal{M} , while maintaining the deformed energy landscape outside. In particular, $\mathcal{M} = 0 \dots 0$ corresponds to using only the original problem Hamiltonian H_P , whereas $\mathcal{M} = 1 \dots 1$ indicates that only the new feature Hamiltonian H_F is used.

If we construct \mathcal{M} from a set of sampled states belonging to the same valley as α (i.e., setting $m_i = 1$ if all sampled states share the same spin orientation at site i , and $m_i = 0$ otherwise), then H_P is reintroduced only within the search region, thus recovering the minimum α^* . This is because, spins aligned across all states of a valley define the subspace of the hypercube in which the valley resides. Hence applying the feature Hamiltonian to these spins removes all competing valleys from the energy landscape. Note that α itself is not necessarily a GS of $H_{F\mathcal{M}}$. In the instantaneous spectrum, this construction restores the avoided crossings and associated dynamics within the search region, while preserving the annealing path that first brings the system into this region. Consequently, LDA enables the system to evolve to a specific point in the original instantaneous spectrum and then resume the dynamics from that point onward. This, in turn, makes it possible to retry difficult segments of the spectrum (i.e., avoided crossings with narrow gaps) until the transition is successfully traversed.

Figure 4 illustrates an example of $H_{F\mathcal{M}}$ constructed from four sampled states (black dots in Fig. 4b.1). The bitmask \mathcal{M} has a 1 at each bit positions where the sampled states agree. The corresponding feature Hamiltonian H_F is shown in Fig. 4b, where the rough energy landscape of the 24-qubit spin glass (Fig. 1c) is transformed into a more monotonic landscape (Fig. 4b.1). Since the sampled states primarily disagree on the features of the first domain, the landscape remains nearly flat between α and $\bar{\alpha}$. Incorporating the original problem Hamiltonian H_P in this domain (Fig. 4d) allows $H_{F\mathcal{M}}$ resolving the locations of the local minima $|\alpha, \beta, \gamma\rangle$ and $|\bar{\alpha}, \beta, \gamma\rangle$. As a result, all avoided crossings in the instantaneous spectrum are removed, except for the transition between these two valleys. This focuses the annealing dynamics on the relevant subspace (Fig. 4e.1), enabling repeated attempts to traverse the critical transition.

5 Hybrid optimization

LDA gives control over the exploration of the Hilbert space in quantum annealing by integrating learned in-

formation about the domain structure into the design of the problem Hamiltonian. By modifying the instantaneous energy spectrum, LDA energetically isolates local minima and creates a more monotonic energy landscape. Consequently, QPTs can be selectively suppressed, restricting the quantum evolution to a subspace of the hypercube.

In this section, we demonstrate how LDA can be used in advanced protocols to solve hard spin-glass instances. We introduce two algorithms, the local search and the global search protocol, which use an iterative application of LDA. Our hybrid optimizer is constructed as an alternating series of these two algorithms.

5.1 Local search protocol

The local search protocol is designed to converge from an initial state α_0 to a nearby local minimum α^* through an iterative application of LDA. The protocol consists of two phases, a distribution and a convergence phase.

In the distribution phase, the protocol generates states that spread evenly around the initial state α_0 (possibly with higher energies), by biasing the evolution only weakly (i.e. $\lambda \ll 1$). The goal is to identify those features from α_0 that characterize α^* , while ignoring terms that are violated in the spin domain ground states. The assumption is that the common features among the sampled states define a subset of \mathcal{H}^{α^*} and \mathcal{J}^{α^*} , providing a search area in the energy landscape where the protocol assumes α^* to be located in. In the second phase, the protocol then converges towards α^* , by focussing the annealing evolution onto the search area. This is achieved by deforming the energy landscape outside the area using H_F (i.e. $\lambda \approx 1$). Over multiple iterations of LDA, the sets of sampled states gradually reveal the elements of \mathcal{H}^{α^*} and \mathcal{J}^{α^*} , shrinking the search area until the local minimum is found.

Each iteration i of the protocol begins with the construction of $H_{F\mathcal{M}}(\alpha_i, \mathcal{M}_i)$ (see Eq. (15)) using the lowest-energy state α_i and the bitmask \mathcal{M}_i derived from the previous iteration. For the first iteration ($i = 0$, distribution phase), $\mathcal{M}_0 = 1 \dots 11$ is used. The hyperparameter λ controls the amplification of extracted features in $H_{F\mathcal{M}}$. It functions as an inverse temperature, reflecting the confidence that the extracted features form a subset of \mathcal{H}^{α^*} and \mathcal{J}^{α^*} . The goal is to maintain α^* as a ground state of $H_{F\mathcal{M}}$ throughout the iterations. Typically, a smaller λ indicates a wider spread of sampled states, as the energy landscape is less deformed (see Fig. 5a). We use a geometric schedule

$$\lambda(i) = \lambda_s \cdot \left(\frac{\lambda_f}{\lambda_s} \right)^{\frac{i}{I-1}}, \quad (16)$$

that scales between λ_s and λ_f , with I denoting the

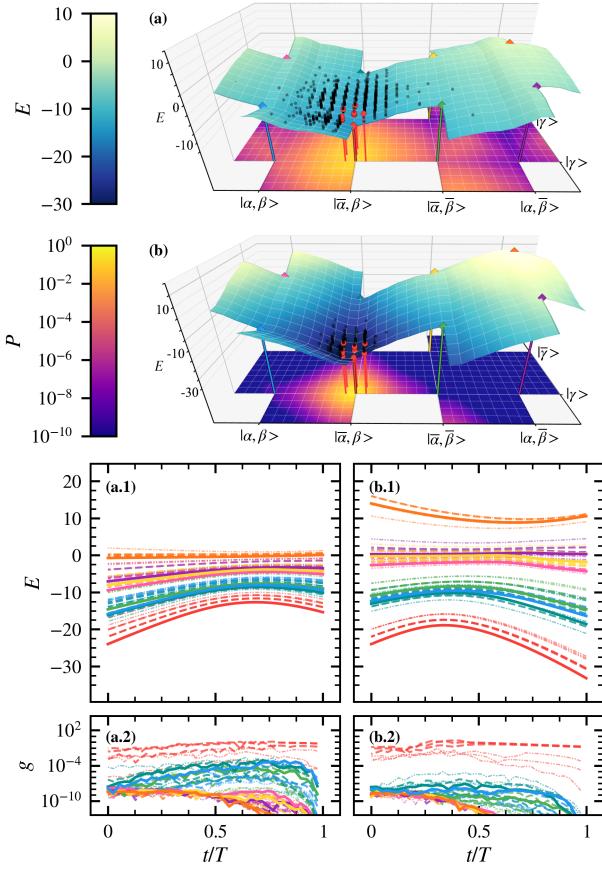


Figure 5: Two iterations of the local search protocol applied to the 24-qubit spin-glass instance shown in Fig. 1a. The protocol is executed with $Q_{\mathcal{T}} = 0.98$ and $N_{\mathcal{T}} = 5$, using (a) $\lambda = 0.2$ for the first and (b) $\lambda = 1.0$ for the second iteration. The top panels depict the energy landscape (E) of $H_{F,\mathcal{M}}$ in the first and second iteration, respectively. Black dots represent 2000 samples from a quantum annealing simulation ($T = 10$), with the probability density (P) shown on the bottom plane. Red dots indicate the subset \mathcal{T} . The initial state α_0 of the first iteration (panel a) is marked as a red dot in Fig. 1c. The color scales are consistent with those in Fig. 1c. Panels (a-b.1) display the instantaneous energy spectrum of the transverse field Ising model with $H_{F,\mathcal{M}}$ for each iteration. The energy levels correspond to those in Fig. 3, where the color of the lines indicates the association to the eight valleys in the energy landscape of H_P . A linear annealing schedule $\Gamma(t)$ is used throughout as an example. Panels (a-b.2) show the adiabatic ratios g (see Eq. (7)) for the states shown in panels (a-b.1), respectively.

total number of Iterations.

In total, $N_{\mathcal{S}}$ states are generated in each iteration and stored in the set \mathcal{S} . To construct the bitmask \mathcal{M} , the states are sorted in increasing order of their energy in the original problem Hamiltonian H_P . A loop then traverses the states $\gamma \in \mathcal{S}$ and includes them in a subset \mathcal{T} if

$$|\mathcal{T}| \leq N_{\mathcal{T}} - 1 \quad \text{and} \quad q_F(\beta, \gamma) \leq Q_{\mathcal{T}} \quad \forall \beta \in \mathcal{T}, \quad (17)$$

where $N_{\mathcal{T}} < N_{\mathcal{S}}$ and $Q_{\mathcal{T}}$ are LDA parameters. This extracts a dispersed set from the sampled states,

where the second condition ensures that all states have a minimum distance w.r.t. q_F (i.e., they do not cluster together) and are uniformly distributed around α^* . The bitmask $\mathcal{M} = m_{N-1} \dots m_1 m_0$ is determined by

$$m_i = \begin{cases} 1, & \text{if } |\sum_{\beta \in \mathcal{T}} \beta_i| = N_{\mathcal{T}} \\ 0, & \text{else} \end{cases}, \quad (18)$$

such that bit m_i is only set if all states $\beta \in \mathcal{T}$ have the same bit β_i , and $\alpha_{i+1} = \mathcal{S}[0]$. The protocol continues iterating until $\alpha_i = \alpha_{i+1}$ or $i = I - 1$, with α_i being the local minimum. The full algorithm is detailed in Protocol 1.

Figure 5 shows two iterations of the local search protocol applied to the 24-qubit spin-glass instance depicted in Fig. 1a, using $Q_{\mathcal{T}} = 0.98$ and $N_{\mathcal{T}} = 5$. The initial state α_1 is marked as a red dot in Fig. 1c. Figures 5a-b.1 demonstrate how the local search progressively identifies the red curves as corresponding to the valley of the local minimum $(\bar{\alpha}, \beta, \gamma)$, separating it in the instantaneous energy spectrum and suppressing QPTs into other valleys (shown in other colors). Consequently, the evolution is increasingly confined around the local minimum, until it is sampled with high probability.

Protocol 1 Local search

Input: α_0 ▷ Initial state
 H_P ▷ Problem Hamiltonian
 $N_{\mathcal{S}}$ ▷ Size of \mathcal{S}
 $N_{\mathcal{T}}$ ▷ Size of \mathcal{T}
 $Q_{\mathcal{T}}$ ▷ Max. similarity between states in \mathcal{T}
 λ_s, λ_f ▷ Initial and final mixing strength
 I ▷ Number of iterations
 T ▷ Annealing time

Output: α_* ▷ Local minimum near α_0

```

1: procedure LOCAL SEARCH
2:    $\mathcal{M} \leftarrow 1 \dots 11$ 
3:   for  $i$  in range  $(0, I)$  do
4:      $\lambda \leftarrow \text{update\_}\lambda(i)$  ▷ Eq. 16
5:      $H_{F,\mathcal{M}} \leftarrow \text{generate\_}H_{F,\mathcal{M}}(H_P, \alpha_i, \mathcal{M}, \lambda)$  ▷ Eq. 15
6:      $\mathcal{S} \leftarrow \text{sample}(H_{F,\mathcal{M}}, N_{\mathcal{S}}, T)$ 
7:      $\mathcal{T} \leftarrow \text{select\_samples}(\mathcal{S}, H_P, N_{\mathcal{T}}, Q_{\mathcal{T}})$  ▷ Eq. 17
8:      $\mathcal{M} \leftarrow \text{update\_}\mathcal{M}(\mathcal{T})$  ▷ Eq. 18
9:      $\alpha_{i+1} \leftarrow \mathcal{T}[0]$ 
10:    if  $\alpha_i = \alpha_{i+1}$  then
11:      return  $\alpha_i$ 
12:    return  $\alpha_I$ 

```

5.2 Global search protocol

The global search protocol is designed to transition from an initial local minimum α^* to a state of a lower-energy valley. The protocol differs from the local search, as it does not directly search around α^* , but uses this state as a reference for LDA to gradually filter out higher-energy valleys from the energy landscape. This allows the protocol to overcome large energy barriers, finding lower-energy valleys even at large Hamming distances. Note that the protocol often identifies excited states of the new valley, requiring a subsequent local search to locate the corresponding local minimum.

To reach a new valley in the energy landscape, the global search transitions between the degenerate ground states of at least one domain, by flipping its $\mathcal{O}(N/N_{\text{dom}})$ spins in α^* . Notably, the procedure does not require knowledge about the domain locations, as the domain structure is inferred from the hierarchy of local minima [22]. This hierarchy arises from differences in domain sizes and inter-domain couplings, leading to different energy gaps between domain ground states. This means that a set \mathcal{S} of energetically similar sampled states (distributed across many valleys) is more likely to align in domains where the two ground states are energetically well-separated, than in domains that only weakly influence the state energies. In such cases, the optimal ground state in these domains can be partially derived from the aligned spins in \mathcal{S} . Consequently, domains where the spins in \mathcal{S} and α^* match indicate that α^* resides in the lower-energy ground state. The bitmask \mathcal{M} is then applied to focus the annealing on the remaining domains. Through successive iterations, this strategy gradually substitutes the problem Hamiltonian H_P with the feature Hamiltonian $H_F(\alpha^*)$, aligning the domains with α^* and eliminating QPTs into higher energy valleys. Eventually, H_P is only applied to the subspace of suboptimal domains, such that a state from a lower-energy valley can be sampled with high probability.

The protocol begins with the unmodified problem Hamiltonian $H_{F\mathcal{M}} = H_P$ and $\mathcal{M} = 0 \dots 00$. Running LDA $N_{\mathcal{S}}$ times generates the set of sampled states \mathcal{S} , from which the subset \mathcal{T} is constructed analogously to the local search protocol (see Eq. 17 in Sec. 5.1). Importantly, α^* is appended to \mathcal{T} , and the bitmask \mathcal{M} for the next iteration is determined using Eq. (18). The protocol continues until a state with lower energy than α^* is identified or $i = I - 1$. Notably, $\lambda \approx 1$ is used for fast convergence. The complete algorithm is presented in Protocol 2.

It is worth noting that the protocol may return to the original valley of α^* . To mitigate this, we impose the condition $q_F(\alpha^*, \beta) \leq Q_{\alpha^*} \forall \beta \in \mathcal{T}$, where the hyperparameter Q_{α^*} denotes the maximum similarity between the selected states and α^* , effectively ensuring a minimal Hamming distance between them.

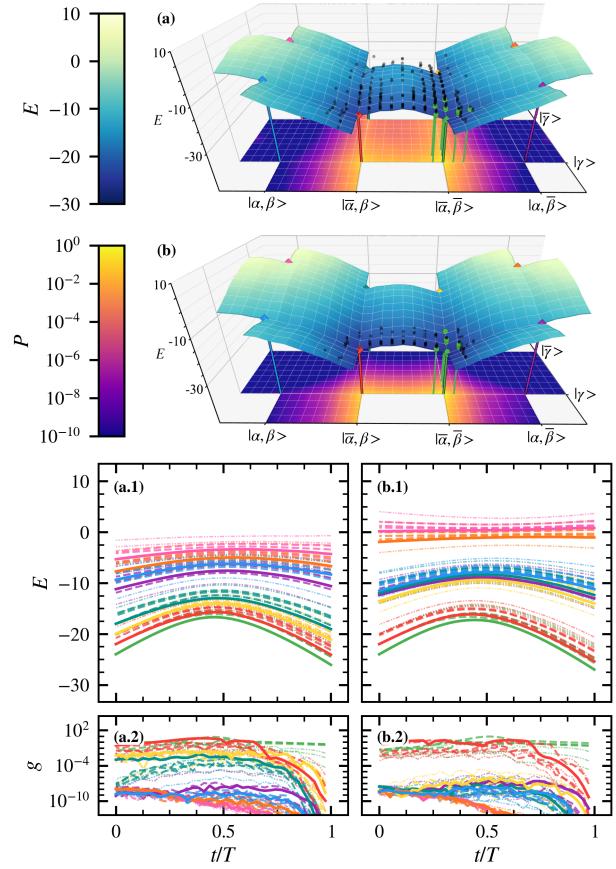


Figure 6: Two iterations of the global search protocol applied to the 24-qubit spin-glass instance shown in Fig. 1a, using $Q_{\mathcal{S}} = 0.98$, $Q_{\alpha^*} = 0.9$, $N_{\mathcal{T}} = 5$, and $\lambda = 1.0$. (a) and (b) depict the energy landscape (E) of $H_{F\mathcal{M}}$ during the first and second iterations, respectively. Black dots represent 2000 samples from a quantum annealing simulation ($T = 10$), with the probability density (P) projected onto the bottom plane. Green dots indicate the subset \mathcal{T} . Green, brown and yellow dots indicate \mathcal{T} from the first iteration with $H_{F\mathcal{M}} = H_P$ in Fig. 1c. $|\bar{\alpha}, \beta, \gamma\rangle$ is the initial local minimum α^* (red diamond). The color scales are consistent with those in Fig. 1c. Panels (a-b.1) display the instantaneous energy spectrum of the transverse field Ising model with $H_{F\mathcal{M}}$ for each iteration. The energy levels correspond to those in Fig. 3, where the color of the lines indicates the association to the eight valleys in the energy landscape of H_P . A linear annealing schedule $\Gamma(t)$ is used throughout as an example. Panels (a-b.2) show the adiabatic ratios g (see Eq. (7)) for the states shown in panels (a-b.1), respectively.

This acts as a search radius, with smaller Q_{α^*} values generally leading to valleys farther from α^* . This condition can also be applied to previously identified local minima to guarantee the exploration of an unvisited valley.

Figure 6 presents two iterations of the global search protocol applied to the 24-qubit spin-glass instance depicted in Fig. 1a, using parameters $Q_{\mathcal{T}} = 0.98$, $Q_{\alpha^*} = 0.9$, $N_{\mathcal{T}} = 5$ and $\lambda = 1.0$. The state $(\bar{\alpha}, \beta, \gamma)$ (red diamond) is selected as the initial local minimum α^* . In the first iteration (see Fig. 6a), the protocol

identifies $\bar{\alpha}$ as the optimal ground state of the first domain, followed by identifying γ as the optimal ground state of the third domain in the second iteration (see Fig. 6b). As a result, the protocol progressively isolates the valleys corresponding to the states $(\bar{\alpha}, \beta, \gamma)$ (red curves) and $(\bar{\alpha}, \bar{\beta}, \gamma)$ (green curves) in the instantaneous energy spectrum (see Fig. 6b.1), yielding states from the latter (green dots in panel b) with high probability.

Protocol 2 Global search

Input:	α^*	▷ Initial local minimum
	H_P	▷ Problem Hamiltonian
	N_S	▷ Size of \mathcal{S}
	N_T	▷ Size of \mathcal{T}
	Q_T	▷ Max. similarity between states in \mathcal{T}
	Q_{α^*}	▷ Max. similarity to α^*
	λ	▷ Mixing strength
	I	▷ Number of iterations
	T	▷ Annealing time

Output: β ▷ State from a lower-energy valley

```

1: procedure GLOBAL SEARCH
2:    $\mathcal{M} \leftarrow 0 \dots 00$ 
3:   for  $i$  in range  $(0, I)$  do
4:      $H_{FM} \leftarrow \text{generate\_}H_{FM}(H_P, \alpha^*, \mathcal{M}, \lambda)$  ▷
Eq. 15
5:      $\mathcal{S} \leftarrow \text{sample}(H_{FM}, N_S, T)$ 
6:      $\mathcal{T} \leftarrow \text{select\_samples}(\mathcal{S}, H_P, N_T, Q_T)$  ▷
Eq. 17
7:      $\mathcal{T} \leftarrow \mathcal{T} + [\alpha^*]$ 
8:      $\mathcal{M} \leftarrow \text{update\_}\mathcal{M}(\mathcal{T})$  ▷ Eq. 18
9:      $\beta \leftarrow \mathcal{T}[0]$ 
10:    if  $E_\beta < E_{\alpha^*}$  then
11:      return  $\beta$ ;
12:    return  $\beta$ 

```

6 Benchmarks

We assess the performance of the proposed hybrid solver with LDA against other quantum and classical methods on large spin-glass instances. We focus on random spin glasses that can be natively mapped onto currently available D-Wave Advantage1 QPUs. Specifically, we investigate 10 NAT-7 [24, 30] instances that almost fully utilize the Advantage1 5.4 QPU [1], using 5,580 qubits and 39,898 couplers. The instances are generated by assigning random coupling strengths J to each coupler, where J is uniformly chosen from the set $\mathcal{J} = \{\pm 1/7, \pm 2/7, \dots, \pm 1\}$. No biases h are applied to the qubits, maintaining a global \mathbb{Z}_2 symmetry. Using multiples of $1/7$ ensures that the problem and feature Hamiltonians are well within

the precision limits of the QPU [16]. Moreover, the use of a non-Sidon [46] set (i.e., the pairwise sum of two elements from the set can be an element of the set again) allows for vanishing fields and causes degeneracies between states. The problem instances are designed to feature numerous low-energy local minima, that have energies within a percent of the ground state energy but are separated by large energy barriers (i.e., $d_{ham} = \mathcal{O}(10^2)$; \mathcal{O} here denotes order of magnitude in the physics convention rather than asymptotic scaling in the computer-science convention). This enables us to investigate the capability of quantum and classical solvers to navigate rough energy landscapes.

Figure 7 presents the benchmarking results of various quantum and classical algorithms for the NAT-7 instances, with the relative energy to the best state detailed as a function of runtime for the first problem instance in panels (a)–(e). Results for the remaining instances are summarized in panels (f.1–5) (for reference, comparisons in terms of absolute energies are given in Fig. 11 in App. E). For each problem instance i , a reference state Ψ_i (black dots) is determined, which is the lowest-energy state obtained from a forward annealing run on the D-Wave Advantage1 5.4 QPU, using an annealing time of $2ms$ and generating 2000 samples. This reference state Ψ_i is used to initialize the proposed hybrid solver and other quantum strategies, to ensure a fair comparison with regard to the use of the QPU to explore the energy landscape.

Panel 7a shows results of the proposed hybrid solver executed on the D-Wave Advantage1 5.4 QPU. The two stages of the algorithm are distinguished by marker shape: squares denote 8 iterations of the local search protocol, while diamonds represent 8 iterations of the global search protocol. Starting from Ψ_1 , the solver produces 2000 samples per iteration with an annealing time of $2ms$ and spin-reversal transformations enabled (see App. B). Examples of sampled states for the two protocols are shown in App. D. The solver achieves the best energy of -9828.14 ($\Delta E = 0$) within 4 cycles, corresponding to a wall-clock runtime of ≈ 6 minutes. This runtime includes both active CPU time and QPU access time reported by D-Wave, but excludes QPU queue waiting time. The classical processing overhead of LDA amounts to $\approx 49s$, or $\approx 13.5\%$ of the total runtime, due to serial post-processing of the sampled states. Notably, this overhead could be reduced substantially through parallelization.

For comparison, panel 7d presents the lowest energies (orange dots) from 30 independent runs on the QPU for the unmodified problem Hamiltonian using the same annealing parameters as before. The data shows that the QPU samples a mean energy of $\bar{\Delta E} \approx 30.51$ (see dotted orange line) with a root-mean-square deviation of $\sigma_{\Delta E} \approx 2.18$. Already the first iteration of the local search (first square in (a))

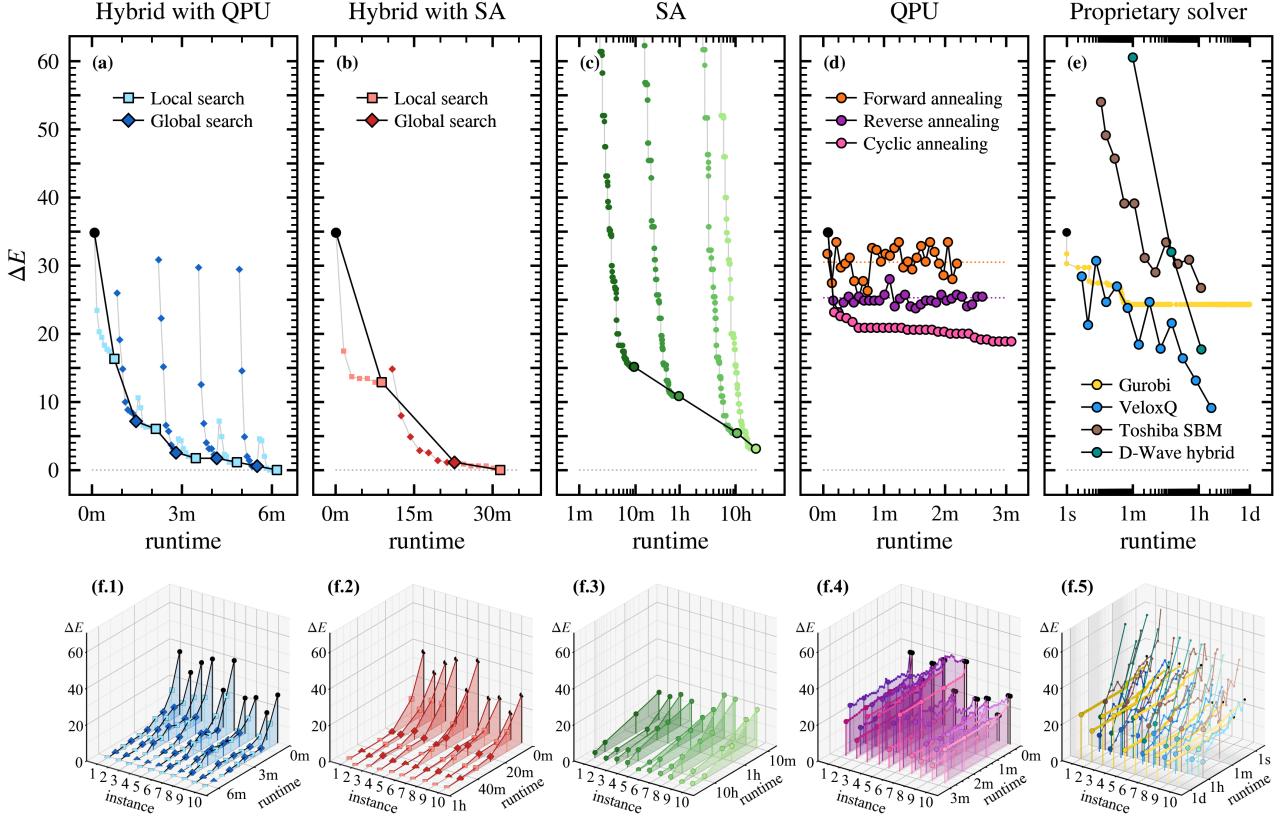


Figure 7: Performance comparison of various quantum and classical solver on 10 random 5580-qubit NAT-7 spin-glass instances. (a)–(e) Relative energy to the best solution as a function of algorithm runtime for the first problem instance. Note that only the active QPU and CPU runtimes are reported, thus excluding queue waiting times. Classical code was executed on the JUWELS Booster supercomputer [23]. (a) The proposed hybrid solver using the D-Wave Advantage1 5.4 QPU for sampling (2ms annealing time, spin-reversal transformation, 2000 samples per run). (b) The proposed hybrid solver using SA from the D-Wave Neal SDK [15] for sampling (500,000 sweeps, 200 samples, 48 threads on an Intel Xeon Platinum 8168 CPU). In both panels, squares denote the local protocol (8 iterations) and diamonds show the global protocol (8 iterations). (c) 4 executions of SA using JUPTSA [20] (Intel Xeon Platinum 8168 CPU with 48 threads) and a geometric annealing schedule from $\beta = 0.1$ to $\beta = 10$ with $600K$ ($\approx 10m$), $4M$ ($\approx 1h$), $40M$ ($\approx 10h$), and $96M$ ($\approx 24h$) spin evaluations. (d) Quantum algorithms on the D-Wave Advantage1 5.4 QPU: forward annealing (orange dots, annealing schedule: $[(0, 0), (2000, 1)]$, 2000 samples, spin-reversal transformation, mean: orange dotted line); reverse annealing [39] (purple dots, annealing schedule: $[(0, 1), (200, 0.7), (1800, 0.7), (2000, 1)]$, $\text{reinitialize_state}=\text{false}$, 2000 samples, spin-reversal transformation, mean: purple dotted line); cyclic annealing [51, 57] (pink dots, h-gain schedule: $[(0, 0), (1, 0.03), (4, 0.03), (2000, 0)]$, annealing schedule: $[(0, 1), (1, 1), (4, 0.7), (2000, 1)]$, 2000 samples, spin-reversal transformation). (e) Proprietary/classical solvers: Gurobi [19] (yellow dots), D-Wave hybrid [14] (teal dots, classified as a classical solver here as the QPU access time was $\lesssim 1\%$ of the total runtime) using 1m, 10m and 1h time limits, Toshiba's SBM algorithm [49] implemented in Ref. [37] (brown dots; 4 Nvidia H100 GPUs) and VeloxQ [21, 37] (light blue dots; 4 Nvidia H100 GPUs; instance specific tuning of parameters). (f.1–f.5) Cumulative results of all 10 NAT-7 instances and solvers. Note, that for each problem instance i the proposed hybrid solver (both QPU and SA sampling), reverse- and cyclic-annealing, and Gurobi use the same initial state Ψ_i (black dot), which is the lowest-energy state obtained from a single forward annealing run on the D-Wave Advantage1 5.4 QPU (2ms annealing time, spin-reversal transformation, 2000 samples). See App. B for a detailed description of the QPU settings.

is significantly below this energy, with $\Delta E \approx 23.43$. This demonstrates that LDA effectively guides the system through the instantaneous energy spectrum by focussing the dynamics on critical transitions, i.e., avoided crossings with small gaps, thereby allowing the system to repeatedly attempt traversing them. As a result, LDA gradually eliminates avoided crossings from the spectrum and thereby reach deeper regions of the energy landscape that are inaccessible within the annealing time of the QPU.

While LDA is primarily designed to suppress QPTs during QA, the gradual simplification of the energy landscape can also benefit classical sampling techniques such as simulated annealing (SA), TABU search, and evolutionary algorithms. To explore this potential, we replace QPU sampling in our hybrid solver with SA in Panel 7b. For convenience, we employ the SA implementation from the D-Wave Neal SDK [15], generating 200 samples with 500,000 sweeps on 48 threads of an Intel Xeon Platinum 8168 CPU, initialized from Ψ_1 (black dot). With SA, the hybrid solver identifies the best energy within a single cycle in $\approx 30\text{min}$. Although classical sampling with SA is considerably slower than QPU sampling (noting that runtime could be reduced with more computational resources), this result demonstrates that LDA is applicable beyond QA. Moreover, other classical sampling methods combined with LDA may converge even faster to the solution state, though we leave a detailed study on classical sampling for future work.

In the context of QA, the success of the SA-based hybrid solver highlights a general property of LDA: the ability to reach the GS does not depend on the quality of the sampled states. Lower-quality samples merely increase the number of cycles required to reach the GS. In this experiment, the SA-based version converged in fewer cycles than the QPU version, as it produced lower-energy samples. This indicates that LDA is broadly applicable to a wide range of NISQ devices, where QA is limited by integrated control errors and finite annealing times [1, 16], since convergence speed (i.e., number of cycles) naturally adapts to the available coherence time.

In panel 7c, we present results from conventional SA using a parallel SA software called JUPTSA [20] with 48 threads. A total of 4 runs was conducted with a geometric annealing schedule from $\beta = 0.1$ to $\beta = 10$, and $600K$ ($\approx 10\text{m}$), $4M$ ($\approx 1\text{h}$), $40M$ ($\approx 10\text{h}$), and $96M$ ($\approx 24\text{h}$) spin evaluations. Notably, with an energy of $\Delta E \approx 3.14$, SA does not reach the best energy within the 24h time limit, emphasizing the difficulty of the problem instance. Compared to panels (a) and (b), our hybrid solver finds the final energy of the 24h run in 168.13s and 961.45s using QPU and SA sampling, respectively. This speed-up is the result of the proposed local and global search protocols, and highlights their efficiency in navigating the rough energy landscape and locating low-energy

valleys. The comparison of panels (b) and (c) also reveals that short annealing runs are better for LDA, as it gives the algorithm greater control over the time-evolution to focus the dynamics into promising regions of the Hilbert space.

This observation is further supported by the results of alternative quantum algorithms shown in panel 7d, which make use of advanced programming features of D-Wave QPUs. The first ansatz we examine is reverse annealing using the IRA scheme [39] as discussed in Sec. 3 (purple dots). Initializing the system in the state Ψ_1 (black dot), we execute in total 30 runs on the QPU, with `reinitialize_state=false` and $[(0, 1), (200, 0.7), (1800, 0.7), (2000, 1)]$ as the annealing schedule (see App. B). Each dot denotes the lowest-energy state of 2000 samples from each iteration. The data shows that reverse annealing improves on the initial energy, but fails to reach the best energy with a mean of $\overline{\Delta E} \approx 25.26$ (see dotted purple line) with a root-mean-square deviation of $\sigma_{\Delta E} \approx 1.94$. As discussed in Sec. 3, this is likely because reverse annealing is constrained by the QPU’s effective temperature.

The second strategy we investigate termed cyclic annealing (pink dots) was proposed by Wang et al. [51, 57] and combines reverse annealing with the `h_gain` schedule feature (which controls the function $b(t)$ in Eq. (20), see App. B) to cycle around the tricritical point of the many-body localization transition [35]. This is achieved using a reference Hamiltonian that encodes the initial state of each cycle through local biases. The strength of this reference Hamiltonian is controlled via the `h_gain_schedule` $[(0, 0), (1, 0.05), (40, 0.05), (2000, 0)]$.

The reverse annealing schedule is $[(0, 1), (1, 1), (40, 0.4), (2000, 1)]$. Although cyclic annealing achieves lower energies than reverse annealing alone (best energy at $\Delta E \approx 20.90$), we observe that each cycle results in only minor changes to the bit string, with $\mathcal{O}(10^1)$ Hamming distance between initial and final state (for comparison, the proposed global search protocol achieves $\mathcal{O}(10^2)$ bit flips). As a result, cyclic annealing often becomes trapped in local minima. This behavior likely arises from the design of the reference Hamiltonian, which biases states based on their Hamming distance to the initial state, restricting the search to nearby states. Given that energy barriers can span $\mathcal{O}(10^2)$ bits, cyclic annealing is unlikely to transition between energy valleys, instead requiring exponentially many restarts to locate the GS.

As a result, all tested purely QPU-based strategies fail to find the best energy on the D-Wave Advantage1 5.4 QPU. While future generations of QPUs might be able to natively sample the best energy, our SA results (panels b, c) indicate that a hybrid protocol, modifying the problem Hamiltonian according to Eq. (15), might still be the best choice for practical quantum

computation, as it can guide the evolution more effectively (panel a).

As our final benchmark, we consider proprietary/classical solvers in panel 7e. These include Gurobi [19] (yellow dots), D-Wave hybrid [14] (teal dots), and two quantum-inspired-algorithms (QIA): Toshiba’s simulated bifurcation machine (SBM) [49] implemented in Ref. [37] (brown dots), and VeloxQ [21, 37] (light blue dots). Note that QIA are often regarded as a classical baseline for quantum algorithms [56]. For Gurobi, we conducted two executions: one without an initial state and one initialized with the state Ψ_1 (black dot). Both executions had a 24h time limit. Only the results of the second run are shown, as the first run failed to achieve an energy lower than $\Delta E \approx 474.57$. Using the initial state, Gurobi improves within the first minute, reaching an energy of $\Delta E \approx 24.2857$. However, subsequent progress was minimal, with the only improvements occurring to the lower-bound energy of the ground state, such that the best energy was not reached within the 24h runtime. Regarding D-Wave hybrid, we evaluated runtimes of 1min, 10min, and 1h, with the latter reaching an energy of $\Delta E \approx 17.71$. Throughout these runs, QPU access time was $\lesssim 1\%$ of the total runtime, suggesting the solver operated predominantly as a classical solver. The poor QPU utilization is likely due to the difficulty of finding a suitable embedding for the problem. Toshiba’s SBM and VeloxQ achieve energies of $\Delta E \approx 26.72$ and $\Delta E \approx 9.14$, respectively. Note that both QIA solvers were executed on substantially more powerful classical hardware (four Nvidia H100 GPUs) compared to the other classical solvers (Intel Xeon Platinum 8168 CPU). We remark that the lowest energies were found by SA and our hybrid solver, while all investigated proprietary classical solvers would likely require longer runtimes or more computing resources to reach similar energies.

In panel (f), we present the cumulative results for the 10 problem instances. Across all instances, we observe consistent qualitative behavior for the evaluated algorithms. Notably, the proposed hybrid solver (both the QPU and SA implementation) finds the lowest energy across all investigated problem instances, surpassing the other algorithms in both lowest energy and total runtime. Only for the fifth and ninth spin-glass instance, SA is able to match the energy of the proposed hybrid solver in the 24h run.

7 Conclusion

In this work, we have introduced Learning-Driven Annealing (LDA), a framework designed to link individual quantum annealing (QA) runs into a global solution strategy to mitigate hardware constraints such as finite annealing times and integrated control errors. LDA differs from other iterative strategies in that it does not change the annealing procedure (i.e,

annealing schedule or annealing time). Instead, it learns about the problem structure (i.e. the spin domains) to adaptively modify the problem Hamiltonian H_P . LDA replaces H_P with a so-called feature Hamiltonian H_F to iteratively eliminate QPTs in the instantaneous energy spectrum. H_F is defined w.r.t. a reference state α , retaining only the biases and couplers from H_P that are satisfied by α . These retained terms, called features, encode the location of a nearby local minimum α^* in the energy spectrum.

We have demonstrated that H_F deforms the energy landscape according to both the Hamming distance and the energy distance from α . When combined with a bitmask \mathcal{M} that restricts the application of the feature Hamiltonian to a subset of spins, the time evolution through the instantaneous energy spectrum is partially set, allowing LDA to concentrate the dynamics on critical transitions with small energy gaps. This is achieved by gradually removing competing valleys from the energy landscape. In this way, LDA iteratively eliminates avoided crossings to access states deep within the energy valley that would otherwise remain inaccessible within the QPU’s annealing time.

To demonstrate the efficacy of LDA, we have developed a hybrid quantum-classical solver for large-scale spin-glass problems. The solver alternates between a local and global search protocol, both using an iterative application of LDA.

The local search protocol converges from an initial state α_0 to a nearby local minimum α^* . It begins with a distribution phase that scatters the system around the initial state to identify characteristic features of the energy valley, followed by a convergence phase that gradually narrows the search area to isolate α^* . Numerical simulations on a 24-qubit spin-glass confirmed that this approach isolates the energy valley of α^* in the instantaneous energy spectrum, ensuring that the search is restricted to transitions within the valley.

The global search protocol transitions from a local minimum α^* to a state in a lower-energy valley by leveraging the hierarchy of local minima. Using α^* as a reference, the protocol identifies spin domains in suboptimal configurations and gradually replaces the problem Hamiltonian H_P with the feature Hamiltonian H_F . Through numerical simulation, we verified that this process isolates lower-energy valleys in the instantaneous energy spectrum, allowing the system to cross high and wide energy barriers and reach a lower-energy valley even at large Hamming distances.

We have benchmarked the proposed hybrid solver against leading quantum and classical methods, including reverse annealing [12], cyclic annealing [51, 57], simulated annealing [20], Gurobi [19], Toshiba’s SBM [49], D-Wave hybrid [14] and VeloxQ [21, 37], on 5580-qubit NAT-7 [24, 30] spin-glass instances. The results show that the proposed hybrid solver outperforms all competing algorithms in both runtime and

lowest energy using both QPU sampling from the D-Wave Advantage1 5.4 and SA sampling. Remarkably, none of the other methods matched the best energies found by the proposed hybrid solver within a 24h time limit (except for SA in the fifth and ninth problem instance), effectiveness of our protocols in exploring rugged energy landscapes and identifying low-energy valleys, even for large-scale instances.

It is important to emphasize that these results do not demonstrate a quantum advantage, as the same energies obtained with QPU sampling can also be reproduced by combining LDA with classical SA. Nevertheless, our experiments on real quantum hardware show that LDA effectively mitigates hardware limitations of NISQ devices (particularly short coherence times) by systematically guiding the time evolution through the instantaneous spectrum. This indicates that, in principle, NISQ devices could solve any COP that can be embedded onto the hardware using LDA, although further studies are needed to assess the applicability of our local and global search protocols to other problem classes. Moreover, LDA could be integrated with hybrid quantum-classical sampling strategies (e.g., combining classical local search with QPU-based global search), potentially enabling even faster convergence. Overall, LDA represents a step toward practical quantum optimization by providing classical control over the quantum time evolution, allowing current quantum annealers to compete with state-of-the-art classical solvers in both runtime and solution quality.

Acknowledgments

The authors thank Hans De Raedt, Madita Willsch, Vrinda Mehta, Fengping Jin, Jaka Vodeb, and Paul Warburton for comments and discussions. The authors thank Jakub Pawłowski and Bartłomiej Gargas for sharing performance results of Toshiba’s SBM and VeloxQ on the investigated problem instances. The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time on the GCS Supercomputer JUWELS [23] at Jülich Supercomputing Centre (JSC). The authors gratefully acknowledge the Jülich Supercomputing Centre (<https://www.fz-juelich.de/ias/jsc>) for funding this project by providing computing time on the D-Wave Advantage™ System JUPSI through the Jülich Unified Infrastructure for Quantum computing (JUNIQ). D.W. acknowledges support from the project JUNIQ that has received funding from the German Federal Ministry of Education and Research (BMBF) and the Ministry of Culture and Science of the State of North Rhine-Westphalia.

References

- [1] Qpu-specific physical properties: Advantagesystem5.4. Technical report, D-Wave Systems Inc., 2025. URL https://docs.dwavesys.com/docs/latest/_downloads/b566a7b6e9f535194611f51a3d890cd6/09-1264A-G_QPU_Properties_Advantage_system5_4.pdf. D-Wave User Manual 09-1264A-G.
- [2] Tameem Albash and Daniel A. Lidar. Adiabatic quantum computation. *Rev. Mod. Phys.*, 90:015002, Jan 2018. DOI: [10.1103/RevModPhys.90.015002](https://doi.org/10.1103/RevModPhys.90.015002). URL <https://link.aps.org/doi/10.1103/RevModPhys.90.015002>.
- [3] Tameem Albash and Daniel A. Lidar. Adiabatic quantum computation. *Reviews of Modern Physics*, 90(1), 2018. ISSN 0034-6861. DOI: [10.1103/RevModPhys.90.015002](https://doi.org/10.1103/RevModPhys.90.015002). URL <https://journals.aps.org/rmp/pdf/10.1103/RevModPhys.90.015002>.
- [4] Boris Altshuler, Hari Krovi, and Jérémie Roland. Anderson localization makes adiabatic quantum optimization fail. *Proceedings of the National Academy of Sciences of the United States of America*, 107(28):12446–12450, 2010. DOI: [10.1073/pnas.1002116107](https://doi.org/10.1073/pnas.1002116107). URL <https://www.pnas.org/doi/10.1073/pnas.1002116107>.
- [5] M. H. S. Amin. Effect of local minima on adiabatic quantum optimization. *Physical Review Letters*, 100(13):130503, 2008. ISSN 0031-9007. DOI: [10.1103/PhysRevLett.100.130503](https://doi.org/10.1103/PhysRevLett.100.130503). URL <https://journals.aps.org/prl/pdf/10.1103/PhysRevLett.100.130503>.
- [6] M. H. S. Amin. Consistency of the adiabatic theorem. *Physical Review Letters*, 102(22):220401, 2009. ISSN 0031-9007. DOI: [10.1103/PhysRevLett.102.220401](https://doi.org/10.1103/PhysRevLett.102.220401). URL <https://journals.aps.org/prl/pdf/10.1103/PhysRevLett.102.220401>.
- [7] M. H. S. Amin and V. Choi. First-order quantum phase transition in adiabatic quantum computation. *Physical Review A*, 80(6), 2009. DOI: [10.1103/PhysRevA.80.062326](https://doi.org/10.1103/PhysRevA.80.062326). URL <https://arxiv.org/pdf/0904.1387.pdf>.
- [8] Mohammad H. Amin, Evgeny Andriyash, Jason Rolfe, Bohdan Kulchytskyy, and Roger Melko. Quantum boltzmann machine. *Phys. Rev. X*, 8(2):021050, may 2018. DOI: [10.1103/physrevx.8.021050](https://doi.org/10.1103/physrevx.8.021050). URL <https://doi.org/10.1103/physrevx.8.021050>.
- [9] Marcello Benedetti, John Realpe-Gómez, Rupak Biswas, and Alejandro Perdomo-Ortiz. Estimation of effective temperatures in quantum annealers for sampling applications: A case study with possible applications in deep learning. *Phys. Rev. A*, 94:022308, Aug 2016. DOI: [10.1103/PhysRevA.94.022308](https://doi.org/10.1103/PhysRevA.94.022308). URL <https://link.aps.org/doi/10.1103/PhysRevA.94.022308>.

[10] K. Binder and A. P. Young. Spin glasses: Experimental facts, theoretical concepts, and open questions. *Rev. Mod. Phys.*, 58:801–976, Oct 1986. DOI: 10.1103/RevModPhys.58.801. URL <https://link.aps.org/doi/10.1103/RevModPhys.58.801>.

[11] Richard J Boucherie, Aleida Braaksma, and Henk Tijms. *Operations Research*. World Scientific, 2021. DOI: 10.1142/12343. URL <https://www.worldscientific.com/doi/abs/10.1142/12343>.

[12] Nicholas Chancellor. Modernizing quantum annealing using local searches. *New Journal of Physics*, 19(2):023024, 2017. ISSN 1367-2630. DOI: 10.1088/1367-2630/aa59c4. URL <https://iopscience.iop.org/article/10.1088/1367-2630/aa59c4/meta>.

[13] Nicholas Chancellor and Viv Kendon. Experimental test of search range in quantum annealing. *Physical Review A*, 104(1), 2021. DOI: 10.1103/PhysRevA.104.012604. URL <https://journals.aps.org/pra/pdf/10.1103/PhysRevA.104.012604>.

[14] D-Wave Systems Inc. D-wave hybrid solver service: An overview, 2025. URL https://www.dwavesys.com/media/4bnpi53x/14-1039a-b_d-wave_hybrid_solver_service_an_overview.pdf.

[15] D-Wave Systems Inc. D-wave neal, 2025. URL <https://github.com/dwavesystems/dwave-neal?tab=readme-ov-file>.

[16] D-Wave Systems Inc. Error sources for problem representation, 2025. URL https://docs.dwavesys.com/docs/latest/c_qpu_ice.html.

[17] S. F. Edwards and P. W. Anderson. Theory of spin glasses. *Journal of Physics F: Metal Physics*, 5(5):965–974, 1975. ISSN 0305-4608. DOI: 10.1088/0305-4608/5/5/017. URL <https://iopscience.iop.org/article/10.1088/0305-4608/5/5/017>.

[18] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. Quantum computation by adiabatic evolution, 2000. URL <https://arxiv.org/abs/quant-ph/0001106>.

[19] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2024. URL <https://www.gurobi.com>.

[20] Fengping Jin, Hans De Raedt, Vrinda Mehta. Juptysa: Jülich parallel tempering simulated annealer, in house software, 2024.

[21] P. Hanussek, J. Pawłowski, Z. Mzaouali, and B. Gardas. Solving quantum-inspired dynamics on quantum and classical annealers. 2025. DOI: 10.48550/arXiv.2509.03952. URL <https://arxiv.org/abs/2509.03952>.

[22] G. Hed, A. K. Hartmann, D. Stauffer, and E. Domany. Spin domains generate hierarchical ground state structure in $j = +/-1$ spin glasses. *Physical Review Letters*, 86(14):3148–3151, 2001. ISSN 0031-9007. DOI: 10.1103/PhysRevLett.86.3148. URL https://www.researchgate.net/publication/12044094_Spin_Domains_Generate_Hierarchical_Ground_State_Structure_in_J_1_Spin_Glasses.

[23] Jülich Supercomputing Centre. JUWELS Cluster and Booster: Exascale Pathfinder with Modular Supercomputing Architecture at Juelich Supercomputing Centre. *J. of Large-Scale Res. Facil.*, 7:A183, 2021. DOI: 10.17815/jlsrf-7-183. URL <http://dx.doi.org/10.17815/jlsrf-7-183>.

[24] Helmut G. Katzgraber, Firas Hamze, and Ruben S. Andrist. Glassy chimeras could be blind to quantum speedup: Designing better benchmarks for quantum annealing machines. *Physical Review X*, 4(2):021008, 2014. DOI: 10.1103/PhysRevX.4.021008.

[25] Helmut G. Katzgraber, Firas Hamze, Zheng Zhu, Andrew J. Ochoa, and H. Munoz-Bauza. Seeking quantum speedup through spin glasses: The good, the bad, and the ugly. *Physical Review X*, 5(3), 2015. DOI: 10.1103/PhysRevX.5.031026. URL <https://journals.aps.org/prx/pdf/10.1103/PhysRevX.5.031026>.

[26] Andrew D. King, Jack Raymond, Trevor Lanting, Richard Harris, Alex Zucca, Fabio Altomare, Andrew J. Berkley, Kelly Boothby, Sara Ejtemaei, Colin Enderud, Emile Hoskinson, Shuiyuan Huang, Eric Ladizinsky, Allison J. R. MacDonald, Gaelen Marsden, Reza Molavi, Travis Oh, Gabriel Poulin-Lamarre, Mauricio Reis, Chris Rich, Yuki Sato, Nicholas Tsai, Mark Volkmann, Jed D. Whittaker, Jason Yao, Anders W. Sandvik, and Mohammad H. Amin. Quantum critical dynamics in a 5,000-qubit programmable spin glass. *Nature*, 617(7959):61–66, 2023. DOI: 10.1038/s41586-023-05867-2.

[27] Andrew D. King, Alberto Nocera, Marek M. Rams, Jacek Dziarmaga, Roeland Wiersema, William Bernoudy, Jack Raymond, Nitin Kaushal, Niclas Heinsdorf, Richard Harris, Kelly Boothby, Fabio Altomare, Andrew J. Berkley, Martin Boschnak, Kevin Chern, Holly Christiani, Samantha Cibere, Jake Connor, Martin H. Dehn, Rahul Deshpande, Sara Ejtemaei, Pau Farré, Kelsey Hamer, Emile Hoskinson, Shuiyuan Huang, Mark W. Johnson, Samuel Kortas, Eric Ladizinsky, Tony Lai, Trevor Lanting, Ryan Li, Allison J. R. MacDonald, Gaelen Marsden, Catherine C. McGeoch, Reza Molavi, Richard Neufeld, Mana Norouzpour, Travis Oh, Joel Pasvolsky, Patrick Poitras, Gabriel Poulin-Lamarre, Thomas Prescott, Mauricio Reis, Chris Rich, Mohammad Samani, Benjamin Sheldon, Anatoly Smirnov, Edward Sterpka, Berta Trullas Clavera, Nicholas Tsai, Mark

Volkmann, Alexander Whiticar, Jed D. Whitaker, Warren Wilkinson, Jason Yao, T. J. Yi, Anders W. Sandvik, Gonzalo Alvarez, Roger G. Melko, Juan Carrasquilla, Marcel Franz, and Mohammad H. Amin. Computational supremacy in quantum simulation. *arXiv e-prints*, art. arXiv:2403.00910, March 2024. DOI: 10.48550/arXiv.2403.00910.

[28] Andrew Lucas. Ising formulations of many np problems. *Frontiers in Physics*, 2, 2014. DOI: 10.3389/fphy.2014.00005. URL <http://arxiv.org/pdf/1302.5843>.

[29] Michael Marzec. Portfolio optimization: Applications in quantum computing. *Handbook of High-Frequency Trading and Modeling in Finance*, pages 73–106, 2016. DOI: 10.1002/9781118593486.ch4. URL <https://doi.org/10.1002/9781118593486.ch4>.

[30] Catherine C. McGeoch and Pau Farré. Milestones on the quantum utility highway: Quantum annealing case study. *ACM Transactions on Quantum Computing*, 5(1):1–30, 2024. ISSN 2643-6809. DOI: 10.1145/3625307.

[31] Vrinda Mehta, Hans De Raedt, Kristel Michielsen, and Fengping Jin. Performance of quantum annealing for 2-sat problems with multiple satisfying assignments. 2025. DOI: 10.48550/arXiv.2502.01423. URL <http://arxiv.org/pdf/2502.01423>.

[32] Vrinda Mehta, Hans De Raedt, Kristel Michielsen, and Fengping Jin. Unraveling reverse annealing: A study of d-wave quantum annealers. 2025. DOI: 10.48550/arXiv.2502.08575. URL <http://arxiv.org/pdf/2502.08575>.

[33] Yuichiro Mori, Shiro Kawabata, and Yuichiro Matsuzaki. How to experimentally evaluate the adiabatic condition for quantum annealing. *Scientific Reports*, 14(1):8177, 2024. ISSN 2045-2322. DOI: 10.1038/s41598-024-58286-2. URL <https://www.nature.com/articles/s41598-024-58286-2>.

[34] Sudip Mukherjee and Bikas K. Chakrabarti. On the question of ergodicity in quantum spin glass phase and its role in quantum annealing. *Journal of the Physical Society of Japan*, 88(6):061004, 2019. ISSN 0031-9015. DOI: 10.7566/JPSJ.88.061004. URL <https://arxiv.org/pdf/1809.01944.pdf>.

[35] Hidetoshi Nishimori and Gerardo Ortiz. *Elements of phase transitions and critical phenomena*. Oxford University Press, Oxford, first published in paperback edition, 2015. ISBN 9780199577224. DOI: 10.1093/acprof:oso/9780199577224.001.0001. URL <https://library.oapen.org/handle/20.500.12657/76440>.

[36] Gianluca Passarelli, Ka-Wa Yip, Daniel A. Lidar, Hidetoshi Nishimori, and Procolo Lucig-
nano. Reverse quantum annealing of the p -spin model with relaxation. *Physical Review A*, 101(2), 2020. DOI: 10.1103/PhysRevA.101.022331. URL <https://link.aps.org/accepted/10.1103/PhysRevA.101.022331>.

[37] J. Pawłowski, J. Tuziński, P. Tarasiuk, A. Przybysz, R. Adamski, K. Hendzel, L. Pawela, and B. Gardas. Veloxq: A fast and efficient qubo solver. 2025. DOI: 10.48550/arXiv.2501.19221. URL <http://arxiv.org/pdf/2501.19221>.

[38] Elijah Pelofske. Comparing three generations of D-Wave quantum annealers for minor embedded combinatorial optimization problems. *Quantum Sci. Technol.*, 10(2):025025, February 2025. ISSN 2058-9565. DOI: 10.1088/2058-9565/adb029.

[39] Elijah Pelofske, Georg Hahn, and Hristo Djidjev. Initial state encoding via reverse quantum annealing and h-gain features. *IEEE Transactions on Quantum Engineering*, 4:1–21, 2023. DOI: 10.1109/TQE.2023.3319586. URL <https://arxiv.org/pdf/2303.13748.pdf>.

[40] Alejandro Perdomo-Ortiz, Salvador E. Venegas-Andraca, and Alán Aspuru-Guzik. A study of heuristic guesses for adiabatic quantum computation. *Quantum Information Processing*, 10(1):33–52, 2011. ISSN 1573-1332. DOI: 10.1007/s11128-010-0168-z.

[41] Alejandro Perdomo-Ortiz, Neil Dickson, Marshall Drew-Brook, Geordie Rose, and Alán Aspuru-Guzik. Finding low-energy conformations of lattice protein models by quantum annealing. *Scientific Reports*, 2012. DOI: 10.1038/srep00571. URL <https://doi.org/10.1038/srep00571>.

[42] Atanu Rajak, Sei Suzuki, Amit Dutta, and Bikas K. Chakrabarti. Quantum annealing: an overview. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, 381(2241):20210417, 2023. DOI: 10.1098/rsta.2021.0417. URL <https://arxiv.org/pdf/2207.01827.pdf>.

[43] Jérémie Roland and Nicolas J. Cerf. Quantum search by local adiabatic evolution. *Physical Review A*, 65(4), 2002. DOI: 10.1103/PhysRevA.65.042308. URL <https://journals.aps.org/pra/pdf/10.1103/PhysRevA.65.042308>.

[44] Gernot Schaller and Ralf Schützhold. The role of symmetries in adiabatic quantum algorithms. *Quantum Information & Computation* 10, 2010. DOI: 10.48550/arXiv.0708.1882. URL <https://arxiv.org/pdf/0708.1882>.

[45] Ralf Schützhold and Gernot Schaller. Adiabatic quantum algorithms as quantum phase transitions: First versus second order. *Phys. Rev. A*, 74:060304, Dec 2006. DOI: 10.1103/PhysRevA.74.060304. URL <https://link.aps.org/doi/10.1103/PhysRevA.74.060304>.

[46] S. Sidon. Ein satz ber trigonometrische poly-

nome und seine anwendung in der theorie der fourier-reihen. *Mathematische Annalen*, 106(1):536–539, 1932. ISSN 1432-1807. DOI: 10.1007/BF01455900. URL <https://link.springer.com/article/10.1007/BF01455900>.

[47] D. L. Stein and C. M. Newman. Spin Glasses: Old and New Complexity. *arXiv e-prints*, art. arXiv:1205.3432, 2012. DOI: 10.48550/arXiv.1205.3432. URL <https://arxiv.org/pdf/1205.3432.pdf>.

[48] Yuki Susa, Yu Yamashiro, Masayuki Yamamoto, Itay Hen, Daniel A. Lidar, and Hidetoshi Nishimori. Quantum annealing of the p -spin model under inhomogeneous transverse field driving. *Physical Review A*, 98(4), 2018. DOI: 10.1103/PhysRevA.98.042326. URL <https://journals.aps.org/pra/pdf/10.1103/PhysRevA.98.042326>.

[49] Toshiba Digital Solutions Corporation. Sqbm+: (commercial service) quantum-inspired optimization solution based on sbm technology, 2025. URL <https://www.global.toshiba/ww/products-solutions/ai-iot/sbm.html>.

[50] Aarón Villanueva, Peyman Najafi, and Hilbert J. Kappen. Why adiabatic quantum annealing is unlikely to yield speed-up. *Journal of Physics A: Mathematical and Theoretical*, 56(46):465304, 2023. ISSN 1751-8113. DOI: 10.1088/1751-8121/ad0439. URL <https://iopscience.iop.org/article/10.1088/1751-8121/ad0439/pdf>.

[51] Hanteng Wang, Hsiu-Chung Yeh, and Alex Kamenev. Many-body localization enables iterative quantum optimization. *Nature Communications*, 13(1):5503, 2022. ISSN 2041-1723. DOI: 10.1038/s41467-022-33179-y. URL <https://www.nature.com/articles/s41467-022-33179-y>.

[52] Dennis Willsch, Madita Willsch, Carlos D. Gonzalez Calaza, Fengping Jin, Hans De Raedt, Marika Svensson, and Kristel Michielsen. Benchmarking advantage and d-wave 2000q quantum annealers with exact cover problems. *Quantum Information Processing*, 21(4), apr 2022. DOI: 10.1007/s11128-022-03476-y. URL <https://doi.org/10.1007%2Fs11128-022-03476-y>.

[53] Dennis Willsch, Madita Willsch, Fengping Jin, Kristel Michielsen, and Hans De Raedt. Gpu-accelerated simulations of quantum annealing and the quantum approximate optimization algorithm. *Computer Physics Communications*, 278:108411, 2022. ISSN 0010-4655. DOI: 10.1016/j.cpc.2022.108411. URL <https://www.sciencedirect.com/science/article/pii/S0010465522001308>.

[54] Yu Yamashiro, Masaki Ohkuwa, Hidetoshi Nishimori, and Daniel A. Lidar. Dynamics of reverse annealing for the fully connected p -spin model. *Physical Review A*, 100(5), 2019. DOI: 10.1103/PhysRevA.100.052321. URL <https://journals.aps.org/pra/pdf/10.1103/PhysRevA.100.052321>.

[55] Sheir Yarkoni, Elena Raponi, Thomas Bäck, and Sebastian Schmitt. Quantum annealing for industry applications: introduction and review. *Reports on progress in physics. Physical Society (Great Britain)*, 85(10), 2022. DOI: 10.1088/1361-6633/ac8c54. URL <https://iopscience.iop.org/article/10.1088/1361-6633/ac8c54/pdf>.

[56] Qing-Guo Zeng, Xiao-Peng Cui, Bowen Liu, Yao Wang, Pavel Mosharev, and Man-Hong Yung. Performance of quantum annealing inspired algorithms for combinatorial optimization problems. *Commun. Phys.*, 7(249):1–10, 2024. ISSN 2399-3650. DOI: 10.1038/s42005-024-01705-7.

[57] Hao Zhang, Kelly Boothby, and Alex Kamenev. Cyclic quantum annealing: searching for deep low-energy states in 5000-qubit spin glass. *Scientific Reports*, 14(1):30784, 2024. ISSN 2045-2322. DOI: 10.1038/s41598-024-80761-z. URL <https://www.nature.com/articles/s41598-024-80761-z>.

A Asymmetry of q_F

Figure 8 compares the similarity measure q_F for two reference states α and β on a 16-qubit spin-glass instance with four energy valleys. Here, α corresponds to the global minimum, while β is a state located in the same energy valley. Both states have a Hamming distance of $d_{\text{ham}} = 2$. The problem graph shown in panel (a) reveals that the satisfied terms $\mathcal{J}^\alpha, \mathcal{J}^\beta$ and $\mathcal{H}^\alpha, \mathcal{H}^\beta$ coincide for most couplers and biases, indicated in green. State α satisfies five couplers not satisfied by β (shown in blue), whereas β satisfies three additional couplers and two biases (shown in red). Panel (b) depicts q_F for both reference states across all computational basis states γ , highlighting the asymmetry of the similarity measure. Notably, while the global minimum (at $q_F(\alpha, \gamma) = 1$) also yields $q_F(\beta, \gamma) \approx 1$, the reverse overlap $q_F(\alpha, \beta)$ is significantly smaller. This demonstrates that q_F is intrinsically biased toward lower-energy states. As a result, the feature Hamiltonian on average guides the system down an energy valley.

B QPU settings

The time-dependent Hamiltonian implemented on the D-Wave Advantage1 5.4 QPU is defined as

$$H_{QA}(t) = -\frac{A(s(t))}{2} \left[\sum_i \sigma_x^i \right] \quad (19)$$

$$+ \frac{B(s(t))}{2} \left[b(t) \cdot \sum_i h_i \sigma_z^i + \sum_{i < j} J_{ij} \sigma_z^i \sigma_z^j \right], \quad (20)$$

where h_i and J_{ij} represent the linear and quadratic terms of the problem Hamiltonian H_P , respectively. The function $s(t) \in [0, 1]$ parameterizes the annealing schedule for which A and B are the strengths of the driving and problem Hamiltonians, respectively (the particular functions $A(s)$ and $B(s)$ for the QPU used are shown in [1]). The function $b(t)$ is called `h_gain` schedule.

For forward annealing, we use $s(t) = t/T$ and $b(t) = 1$, where T is the total annealing time. Additionally, we apply a spin-reversal transformation to H_P for every batch of 400 samples from the QPU. This transformation involves generating a random binary string $r \in \{\pm 1\}^N$ and modifying the problem terms via $h_i \rightarrow h_i r_i$ and $J_{ij} \rightarrow J_{ij} r_i r_j$. Sampled states α are transformed back using $\alpha_i \rightarrow \alpha_i r_i$. We use the spin-reversal transformation to mitigate memory effects on the QPU when rapidly submitting similar problems.

Reverse annealing is realized by changing the annealing schedule via $s(t)$ to begin at $s(0) = 1$ where $B(s) \gg A(s)$. The schedule is specified as a series of tuples $[t/T, s]$ (see Sec. 6), with linear interpolation

between points. Inspired by [13], we use a schedule with a symmetric pause, by first annealing to an intermediate value $0 < s' < 1$, holding s' , and then returning to $s = 1$. For the large spin glasses under investigation, we found $s' = 0.7$ to be optimal. For all reverse annealing runs, we set `reinitialize_state=false`, initializing each run from the previously sampled state.

Cyclic annealing uses both reverse annealing and the QPU's `h_gain` feature, which introduces a time-dependent gain $b(t)$ for the linear terms h_i in H_P . Similar to $s(t)$, $b(t)$ is defined as series of tuples $[t/T, b]$, with linear interpolation between points. Since the large spin glasses under investigation have no linear biases (i.e., $h_i = 0$), we use these terms to encode the initial state α of each cycle, setting $h_i = -\alpha_i$. The schedules for $s(t)$ and $b(t)$ are inspired by [57].

C Properties of H_F

Theorem C.1. *Given the feature Hamiltonian $H_F(\alpha)$ (see Eq. (8)) to a reference state α , then α is a ground state of $H_F(\alpha)$.*

Proof. Let $H_F(\alpha)$ be the feature Hamiltonian to a reference state $\alpha = \alpha_{N-1} \dots \alpha_1 \alpha_0$ and let $\beta = \beta_{N-1} \dots \beta_1 \beta_0$ be an arbitrary state, with $\alpha_i, \beta_i \in \{-1, 1\}$. Then

$$E_\beta = \langle \beta | H_F(\alpha) | \beta \rangle \quad (21)$$

$$= \sum_{(i,j) \in \mathcal{J}^\alpha} \frac{-|J_{ij}|}{2} \cdot [\alpha_i \alpha_j \beta_i \beta_j + \alpha_i \beta_i + \alpha_j \beta_j] \quad (22)$$

$$+ \sum_{i \in \mathcal{H}^\alpha} h_i \beta_i. \quad (23)$$

For each term in the first sum, we have

$$\alpha_i \alpha_j \beta_i \beta_j + \alpha_i \beta_i + \alpha_j \beta_j = \begin{cases} 3 & \text{if } (\beta_i, \beta_j) = (\alpha_i, \alpha_j) \\ -1 & \text{if } (\beta_i, \beta_j) = (\overline{\alpha_i}, \alpha_j) \\ -1 & \text{if } (\beta_i, \beta_j) = (\alpha_i, \overline{\alpha_j}) \\ -1 & \text{if } (\beta_i, \beta_j) = (\overline{\alpha_i}, \overline{\alpha_j}) \end{cases}, \quad (24)$$

so Eq. (22) is minimal if $(\beta_i, \beta_j) = (\alpha_i, \alpha_j)$ for all $(i, j) \in \mathcal{J}^\alpha$. Similarly, for the second sum in Eq. (23), we have

$$h_i \beta_i < 0 \Leftrightarrow \beta_i = \alpha_i, \quad (25)$$

by the definition of \mathcal{H}^α (see Eq. (2b)) Thus, the energy E_β is minimal if $\beta = \alpha$. Therefore, α is a ground state of $H_F(\alpha)$. \square

Theorem C.2. *Given the feature Hamiltonian $H_F(\alpha)$ (see Eq. (8)) to a reference state α , then the energy E_β of an arbitrary state $\beta \neq \alpha$ (see Eq. (21)) is monotonic as a function of Hamming distance*

$d_{ham}(\alpha, \beta)$ from α . This means that there exists a path from β to α via $d_{ham}(\alpha, \beta)$ successive single bit flips, each of which either maintains or decreases the energy.

Proof. Let $H_F(\alpha)$ be the feature Hamiltonian to a reference state $\alpha = \alpha_{N-1} \dots \alpha_1 \alpha_0$, and let $\beta = \beta_{N-1} \dots \beta_1 \beta_0$ be an arbitrary state such that $\alpha \neq \beta$, with $\alpha_i, \beta_i \in \{-1, 1\}$. To prove the statement, we construct a path $P = (\alpha = \gamma^{(0)}, \gamma^{(1)}, \dots, \gamma^{(d_{ham}(\alpha, \beta)-1)} = \beta)$ with monotonically increasing energies.

Let $\gamma^{(0)} = \gamma_{N-1}^{(0)} \dots \gamma_1^{(0)} \gamma_0^{(0)} = \alpha$. Then by Lem. C.1, $\gamma^{(0)}$ is a ground state of $H_F(\alpha)$, with all terms in $E_{\gamma^{(0)}}$ giving only negative contributions. To construct the path P , we define the step from $\gamma^{(n)}$ to $\gamma^{(n+1)}$ for $n = 0, \dots, d_{ham}(\alpha, \beta) - 2$. The energy of $\gamma^{(n)}$ is, according to Eq. (24),

$$E_{\gamma^{(n)}} = \sum_{\substack{(i,j) \in \mathcal{J}^\alpha \\ (\gamma_i^{(n)}, \gamma_j^{(n)}) = (\alpha_i, \alpha_j)}} \frac{-3|J_{ij}|}{2} \quad (26)$$

$$+ \sum_{\substack{(i,j) \in \mathcal{J}^\alpha \\ (\gamma_i^{(n)}, \gamma_j^{(n)}) = (\alpha_i, \alpha_j)}} \frac{|J_{ij}|}{2} \quad (27)$$

$$+ \sum_{\substack{(i,j) \in \mathcal{J}^\alpha \\ (\gamma_i^{(n)}, \gamma_j^{(n)}) = (\alpha_i, \alpha_j)}} \frac{|J_{ij}|}{2} \quad (28)$$

$$+ \sum_{\substack{(i,j) \in \mathcal{J}^\alpha \\ (\gamma_i^{(n)}, \gamma_j^{(n)}) = (\alpha_i, \alpha_j)}} \frac{|J_{ij}|}{2} \quad (29)$$

$$+ \sum_{i \in \mathcal{H}^\alpha} h_i \gamma_i^{(n)}. \quad (30)$$

Let $i \in \{0, \dots, N-1\}$ be the first bit for which $\gamma_i^{(n)} \neq \beta_i$ and set $\gamma^{(n+1)}$ to be the state with the single bit i flipped, i.e., $\gamma^{(n+1)} = \gamma_{N-1}^{(n)} \dots \gamma_{i+1}^{(n)} \beta_i \gamma_{i-1}^{(n)} \dots \gamma_0^{(n)}$. If bit i occurs in the set \mathcal{J}^α (see Eq. (2a)), this single bit flip will either increase the energy of the coupling contribution by moving terms from Eq. (26) to Eq. (27) or Eq. (28)—thereby increasing the energy by multiples of $2|J_{ij}|$ —or it will maintain the energy by moving terms from Eq. (27) or Eq. (28) to Eq. (29). Similarly, if bit i occurs in the set \mathcal{H}^α (see Eq. (2b)), the bit flip will increase the energy of the bias contribution in Eq. (30) by $2|h_i|$. Thus each step from $\gamma^{(n)}$ to $\gamma^{(n+1)}$ for $n = 0, \dots, d_{ham}(\alpha, \beta) - 2$ will either increase or maintain the energy. After N iterations, $\gamma^{(d_{ham}(\alpha, \beta)-1)} = \beta$ by the definition of the Hamming distance. Therefore, P is a path of states with monotonically increasing energies. \square

D Illustration of the local and global search protocol

Figures 9 and 10 show representative illustrations of the local and global search protocols applied to the first 5580-qubit NAT-7 spin-glass instance (see Fig. 7), respectively. They depict the set of sampled states \mathcal{S} , the set of selected states \mathcal{T} and the feature Hamiltonian $H_{\mathcal{M}}$ across four iterations of the protocols.

E Benchmarking results of NAT-7 instances

Figure 11 compares the benchmarking results of the investigated quantum and classical solvers in Sec. 6 within the 10 NAT-7 spin-glass instances. The figure is based on the same data shown in Fig. 7 but highlights the absolute energies found by each solver. Note that the problem instances are designed to feature numerous low-energy local minima, that have energies within a percent of the ground state energy but are separated by large energy barriers ($d_{ham} = (100)$). This allows us to study the solvers' ability to navigate rough energy landscapes. Consequently, the energy differences between solvers are typically small relative to the overall energy scale of the problems, but their corresponding states exhibit a large Hamming distance d_{ham} from one another. The proposed hybrid solver (both QPU and SA sampling), reverse-and cyclic-annealing, and Gurobi use the same initial state Ψ_i (black dot), which is the lowest-energy state obtained from a single forward annealing run on the D-Wave Advantage1 5.4 QPU (2ms annealing time, spin-reversal transformation, 2000 samples).

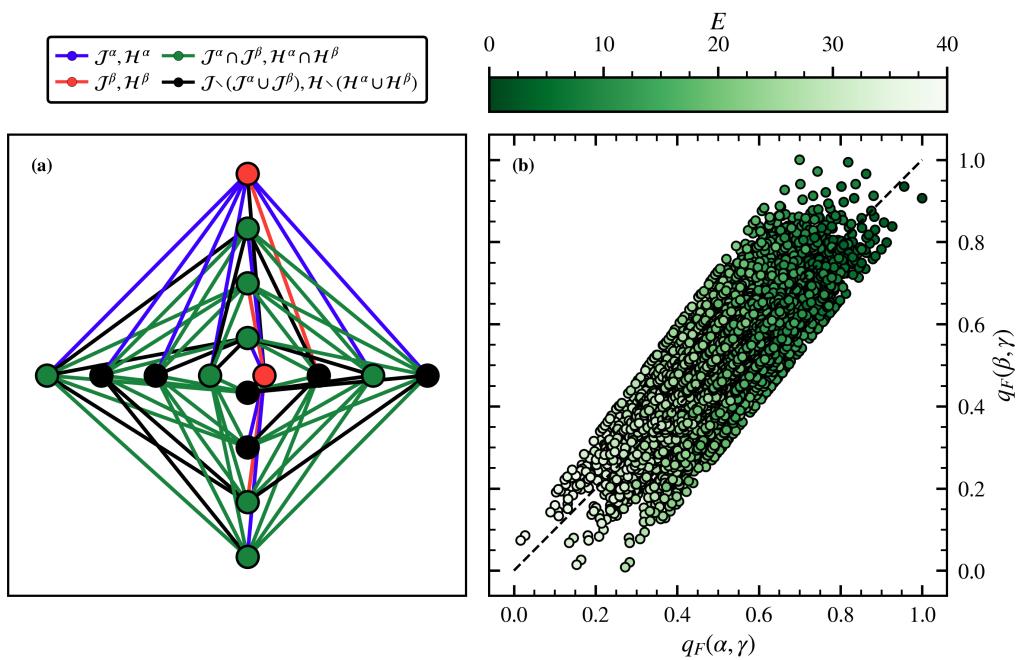


Figure 8: Asymmetry of the similarity measure q_F w.r.t. the global minimum α and a state β located in the global minimum well. (a) 16-qubit spin-glass graph, where nodes represent biases h_i and edges represent couplers J_{ij} . Colors indicate whether h_i and J_{ij} belong to the sets of couplers $\mathcal{J}^\alpha, \mathcal{J}^\beta$ and biases $\mathcal{H}^\alpha, \mathcal{H}^\beta$ satisfied by α and/or β . (b) Comparison of $q_F(\alpha, \gamma)$ and $q_F(\beta, \gamma)$ across all computational basis states γ , illustrating the asymmetry of the similarity measure. While $q_F(\beta, \alpha)$ is close to one, while $q_F(\alpha, \beta)$ is significantly lower, revealing an intrinsic bias toward the lower-energy state.

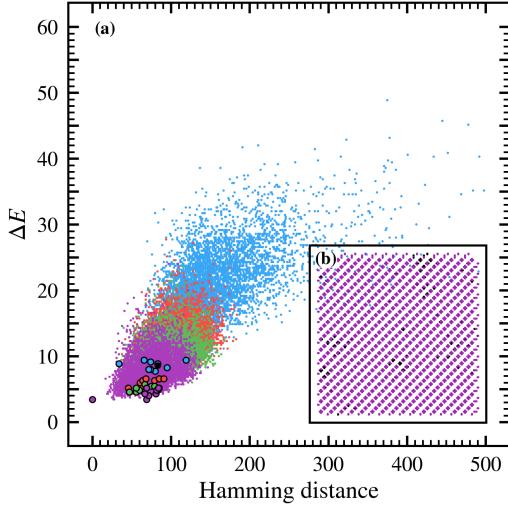


Figure 9: Four iterations of the local search protocol applied to the first NAT-7 spin-glass instance, shown in Fig. 7. (a) Set \mathcal{S} of 6000 sampled states for each iteration, presented as a function of their Hamming distance from the protocol's final state (purple circle) and their energy difference ΔE relative to the best-known state. The energy axis is consistent with Fig. 7a. The black circle denotes the reference state for the protocol, while blue, red, green, and purple dots represent sampled states from the first, second, third, and fourth iterations, respectively. Colored circles depict the elements of the subset \mathcal{S}_C used for the calculation of the bitmask \mathcal{M} (see Eq. (18)) of each iteration. (b) Embedding of the feature Hamiltonian on the D-Wave Advantage1 5.4 QPU [1], with qubits represented as nodes. Purple nodes indicate qubits for which a coupler and/or bias was modified in the feature Hamiltonian in the fourth iteration relative to the original problem Hamiltonian. Black nodes represent qubits for which the original problem Hamiltonian was used.

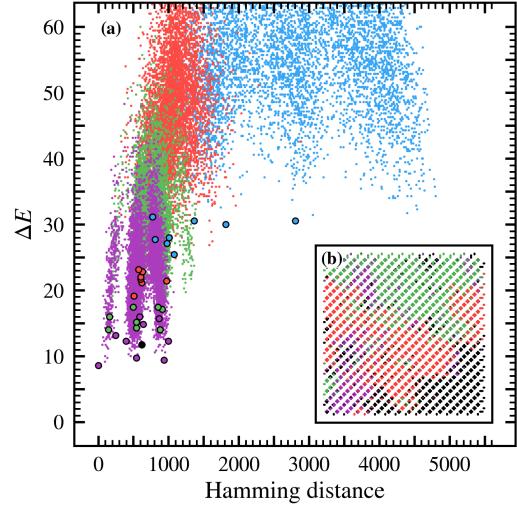


Figure 10: Four iterations of the global search protocol applied to the first NAT-7 spin-glass instance, shown in Fig. 7. (a) Set \mathcal{S} of 6000 sampled states for each iteration, presented as a function of their Hamming distance from the protocol's final state (purple circle) and their energy difference ΔE relative to the best-known state. The energy axis is consistent with Fig. 7a. The black circle denotes the reference state for the protocol, while blue, red, green, and purple dots represent sampled states from the first, second, third, and fourth iterations, respectively. Colored circles depict the elements of the subset \mathcal{S}_C used for the calculation of the bitmask \mathcal{M} (see Eq. (18)) of each iteration. (b) Embedding of the feature Hamiltonian on the D-Wave Advantage1 5.4 QPU [1], with qubits represented as nodes. Colored nodes indicate qubits for which a coupler and/or bias was modified in the feature Hamiltonian relative to the original problem Hamiltonian, with the color encoding the iteration in which the modification occurred. Black nodes represent qubits for which the original problem Hamiltonian was used. Since connected qubits are typically spatially close on the QPU, the spin-glass domain structure is reflected in the clustering of colored nodes.

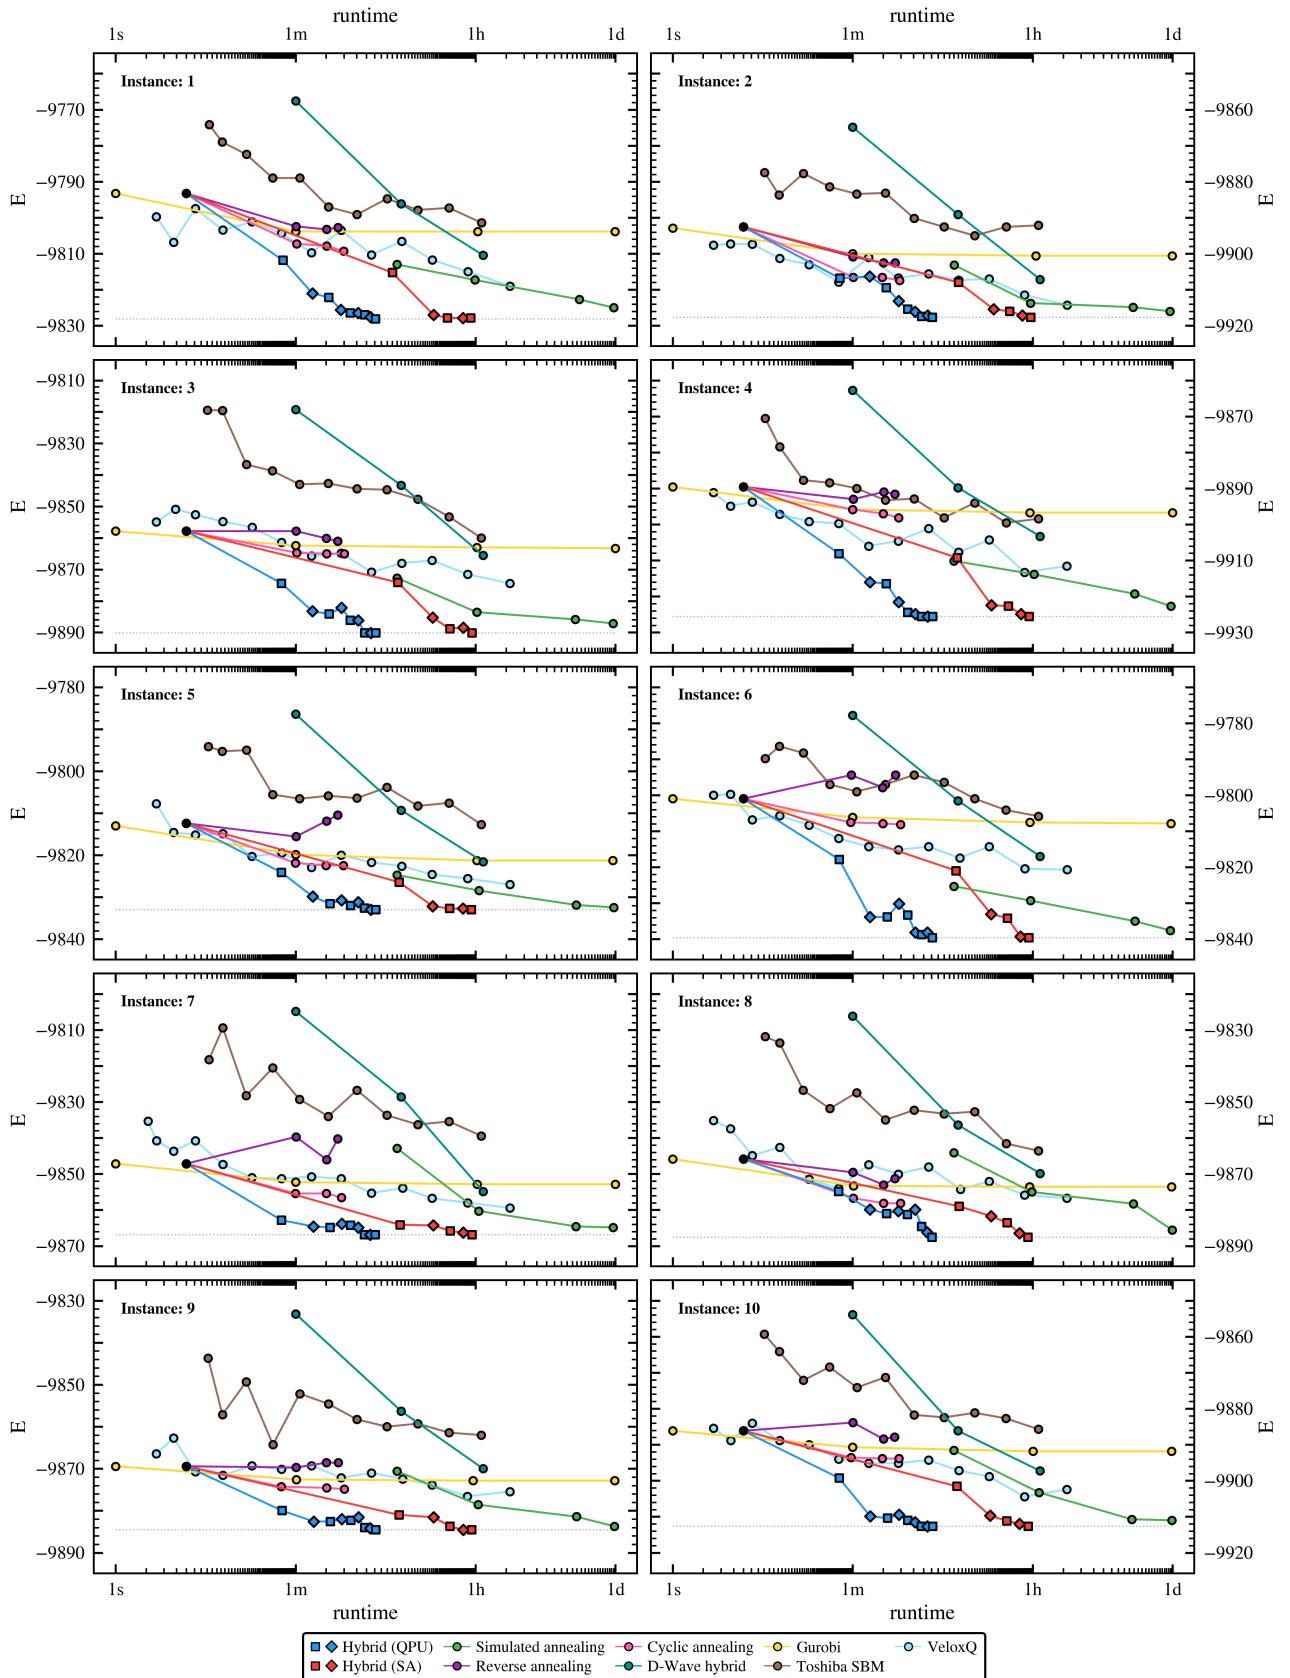


Figure 11: Performance comparison of various quantum and classical solver on 10 random 5580-qubit NAT-7 spin-glass instances. The data corresponds to Fig. 7, but depicts the absolute energies found by each solver.