

What is the Computational Value of Finite Range Tunneling?

Vasil S. Denchev,¹ Sergio Boixo,¹ Sergei V. Isakov,¹ Nan Ding,¹ Ryan Babbush,¹ Vadim Smelyanskiy,¹ John Martinis,² and Hartmut Neven¹

¹Google Inc., Venice, CA 90291, USA

²Google Inc., Santa Barbara, CA 93117, USA

(Dated: December 11, 2015)

Quantum annealing (QA) has been proposed as a quantum enhanced optimization heuristic exploiting tunneling. Here, we demonstrate how finite range tunneling can provide considerable computational advantage. For a crafted problem designed to have tall and narrow energy barriers separating local minima, the D-Wave 2X quantum annealer achieves significant runtime advantages relative to Simulated Annealing (SA). For instances with 945 variables, this results in a time-to-99%-success-probability that is $\sim 10^8$ times faster than SA running on a single processor core. We also compared physical QA with Quantum Monte Carlo (QMC), an algorithm that emulates quantum tunneling on classical processors. We observe a substantial constant overhead against physical QA: D-Wave 2X again runs up to $\sim 10^8$ times faster than an optimized implementation of QMC on a single core. To investigate whether finite range tunneling will also confer an advantage for problems of practical interest, we conduct numerical studies on binary optimization problems that cannot yet be represented on quantum hardware. For random instances of the number partitioning problem, we find numerically that QMC, as well as other algorithms designed to simulate QA, scale better than SA and better than the best known classical algorithms for this problem. We discuss the implications of these findings for the design of next generation quantum annealers.

I. INTRODUCTION

Simulated annealing (SA) [1] is perhaps the most widely used algorithm for global optimization of pseudo-Boolean functions with little known structure. The objective function for this general class of problems is

$$H_P^{\text{cl}}(\mathbf{s}) = -\sum_{k=1}^K \sum_{j_1 \dots j_k=1}^N J_{j_1 \dots j_k} s_{j_1} \cdots s_{j_k}, \quad (1)$$

where N is the problem size, $s_j = \pm 1$ are spin variables and the couplings $J_{j_1 \dots j_k}$ are real scalars. In the physics literature $H_P^{\text{cl}}(\mathbf{s})$ is known as the Hamiltonian of a K -spin model. SA is a Monte Carlo algorithm designed to mimic the thermalization dynamics of a system in contact with a slowly cooling reservoir. When the temperature is high, SA induces thermal excitations which can allow the system to escape from local minima. As the temperature decreases, SA drives the system towards nearby low energy spin configurations.

Almost two decades ago, quantum annealing (QA) [2] was proposed as a heuristic technique for quantum enhanced optimization. Despite substantial academic and industrial interest [3–34], a unified understanding of the physics of quantum annealing and its potential as an optimization algorithm remains elusive. The appeal of QA relative to SA is due to the observation that quantum mechanics allows for an additional escape route from local minima. While SA must climb over energy barriers to escape false traps, QA can penetrate these barriers without any increase in energy. This effect is a hallmark of quantum mechanics, known as quantum tunneling. The

standard time-dependent Hamiltonian used for QA is

$$H(t) = -A(t) \sum_{j=1}^N \sigma_j^x + B(t) H_P, \quad (2)$$

where H_P is written as in Eq. (1) but with the spin variables s_j replaced with σ_j^z Pauli matrices acting on qubit j , and the functions $A(t)$ and $B(t)$ define the annealing schedule parameterized in terms of time $t \in [0, T_{\text{QA}}]$ (see Fig. 7). These annealing schedules can be defined in many different ways so long as the functions are smooth, $A(0) \ll B(0)$ and $A(T_{\text{QA}}) \gg B(T_{\text{QA}})$. At the beginning of the annealing, the transverse field magnitude $A(t)$ is large, and the system dynamics are dominated by quantum fluctuations due to tunneling (as opposed to the thermal fluctuations used in SA).

The question of whether D-Wave processors realize computationally relevant quantum tunneling has been a subject of substantial debate. This debate has now been settled in the affirmative with a sequence of publications [6, 8, 9, 12, 13, 18, 21] demonstrating that quantum resources are present and functional in the processors. In particular, reference [35] studied the performance of the D-Wave device on problems where eight [36] qubit co-tunneling events were employed in a functional manner to reach low-lying energy solutions.

In order to investigate the computational value of finite range tunneling in QA, we study the scaling of the exponential dependence of annealing time with the size of the tunneling domain D ,

$$T_{\text{QA}} = B_{\text{QA}} e^{\alpha D}, \quad (3)$$

where $\alpha = a_{\text{min}}/\hbar$ and a_{min} is the rescaled instanton action (see Eqs. (32) and (35)). In SA, the system escapes

from a local minimum via thermal fluctuations over an energy barrier ΔE separating the minima. The time required for such events scales as

$$T_{\text{SA}} = B_{\text{SA}} e^{\frac{\Delta E}{k_B T}}, \quad (4)$$

which is exponentially long with respect to ΔE . However, for sufficiently tall and narrow barriers

$$\frac{\Delta E}{k_B T} > \alpha D \quad (5)$$

and QA can overcome barriers exponentially faster than SA. This situation was studied in Ref. [37] and it also occurs in the benchmark problems studied in this paper.

Path integral Quantum Monte Carlo (QMC) is a method for sampling the quantum Boltzmann distribution of a d dimensional stoquastic Hamiltonian as a marginalization of a classical Boltzmann distribution of an associated $d+1$ dimensional Hamiltonian. For specific cases, it was recently shown that the exponent α for physical tunneling is identical to the corresponding quantity for QMC [38]. However, in the present work we find a very substantial computational overhead associated with the prefactor B_{QMC} in the expression for the runtime of QMC, $T_{\text{QMC}} = B_{\text{QMC}} e^{\alpha D}$. In other words, B_{QMC} can exceed B_{QA} by many orders of magnitude. The role of this prefactor becomes essential in the situations where the number of cotunneling qubits D is finite, i.e., is independent of the problem size N (or depends on N very weakly).

Because tunneling is more advantageous when energy barriers are tall and narrow, we expect this resource to be most valuable in the upper part of the energy spectrum. For instance, a random initial state is likely to have an energy well above the ground state energy for a difficult optimization problem such as the one in Eq. (1). However, the closest lower energy local minimum will often be less than a dozen spin flips away. Nevertheless, the energy barriers separating these minima may still be high. In such situations, if the transverse field is turned on to facilitate tunneling transitions, the transition rate to lower energy minima will often increase. By contrast, once the state reaches the low part of the energy spectrum, the closest lower local minimum is asymptotically N spin flips away [2, 39–41]. There, finite range tunneling may assist by effectively “chopping off” narrow energy ridges near the barrier top but the transition probability is still largely given by the Boltzmann factor. This description illustrates that finite range tunneling can be useful to quickly reach an approximate optimization, but will not necessarily significantly outperform SA when the task is to find the ground state (see Fig. 1).

The canonical QA protocol initializes the system in the symmetric superposition state, $|+\rangle^{\otimes N}$, which is the ground state at $t = 0$. By a similar argument, we expect that finite range tunneling will drive the system adiabatically across energy gaps associated with narrow barriers, preventing transitions to higher energies. However, in

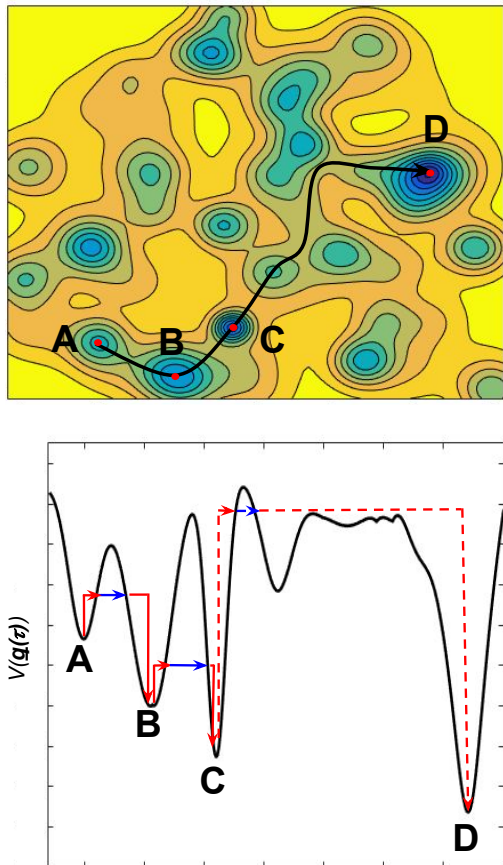


FIG. 1. **Upper:** quantum annealer dynamics are dominated by paths through the mean field energy landscape that have the highest transition probabilities. Shown is a path that connects local minimum A to local minimum D. **Lower:** the mean field energy $V(q(\tau))$ along the path from A to D, as defined by Eqs. (23) and (26). Finite range, thermally assisted tunneling can be thought of as a transition consisting of three steps: i) the system picks up thermal energy from the bath (red arrow up). ii) the system performs a tunneling transition between the entry and exit points (blue arrow). iii) the system relaxes to a local minimum by dissipating energy back into the thermal bath (red arrow down). In transitions A \rightarrow B or B \rightarrow C finite range tunneling considerably reduces the thermal activation energy needed to overcome the barrier. For long distance transitions in the lower part of the energy spectrum, such as C \rightarrow D, the transition rate is still dominated by the thermal activation energy and the increase in transition rate brought about by tunneling is negligible.

general, finite range tunneling will not be able to prevent Landau-Zener diabatic transitions for very small gaps resulting from emerging minima in the energy landscape separated by a wide barrier. This will often include the gap separating the ground state from the first excited state [2, 39–41].

II. BENCHMARKING PHYSICAL QUANTUM ANNEALING ON A CRAFTED PROBLEM WITH A RUGGED ENERGY LANDSCAPE

In this section, we consider a problem designed so that finite cotunneling transitions of multiple spins strongly impacts the success probability. The previous section outlined several reasons why QA has a chance to outperform SA for problems with a rugged energy landscape.




FIG. 2. A pair of weak-strong clusters, consisting of 16 qubits in two unit cells of the Chimera graph. All qubits are ferromagnetically coupled and evolve as part of two distinct qubit clusters. At the end of the annealing evolution, the right cluster is strongly pinned downwards due to strong local fields acting on all qubits in that cell. However, the local magnetic field h_1 in the left cluster is weaker and serves as a bifurcation parameter. For $h_1 = 0.44 < 1/2$ the left cluster will reverse its orientation during the annealing sweep and eventually align itself with the right cluster.

In previous work we proposed a problem consisting of a pair of strongly connected spins (called clusters) to study the presence of functional cotunneling in D-Wave processors [35]. Each cluster coincides with a unit cell of the native hardware graph, known as the Chimera graph. The problem Hamiltonian H_P in Eq. (2) is of Ising form

$$H_P = H_P^1 + H_P^2 + H_P^{1,2} \quad (6)$$

$$H_P^k = -J \sum_{\langle j,j' \rangle \in \text{intra}} \sigma_{k,j}^z \sigma_{k,j'}^z - \sum_{j=1}^8 h_k \sigma_{k,j}^z \quad (7)$$

$$H_P^{1,2} = -J \sum_{j \in \text{inter}} \sigma_{1,j}^z \sigma_{2,j}^z. \quad (8)$$

All the couplings are ferromagnetic with $J = 1$. The index $k \in \{1, 2\}$ indexes a unit cell of the Chimera graph, the first sum in (7) goes over the *intra*-cell couplings depicted in Fig. 2, the second sum goes over the *inter*-cell couplings corresponding to $j = 5, 6, 7, 8$ in Fig. 2, and h_k denotes the local fields within each cell.

The local fields $0 < h_1 < 0.44$ (weak cluster) and $h_2 = -1$ (strong cluster) are equal for all the spins

within the cell. In this parameter regime, all spins of both clusters will point along the direction of the strong local fields in the ground state of the problem Hamiltonian H_P . However, in the initial phase of the annealing evolution, all spins in the weak cluster orient themselves by following the local field in the opposite direction. At a later stage of the annealing evolution, the pairwise coupling between clusters becomes dominant; however, in the mean field picture, the state of the weak cluster is driven into a local minimum. Using a noise model with experimentally measured parameters for the D-Wave 2X processor, we numerically verify that the most likely mechanism by which all spins arrive at the energetically more favorable configuration is multi-qubit cotunneling (see Ref. [35] and Appendix A).

Using the weak-strong cluster pairs as building blocks, larger problems are formed by connecting the strong clusters to one another in a glassy fashion. That is, the four connections between two neighboring strong clusters are all set either to $+1$ (ferromagnetic) or -1 (antiferromagnetic), at random. With this procedure, we define a large class of instances having any size that refer to as the “weak-strong cluster networks” problem. Fig. 3 shows several examples of the layout of instances that were used in these benchmark tests. Due to the fact that not all of the qubits in the D-Wave 2X processor are properly calibrated, and hence available for computation, the instances become somewhat irregular.

A. D-Wave versus Simulated Annealing

We now compare the total annealing time of SA to the total annealing time of the D-Wave 2X processor. Fig. 4 shows the time to reach the ground state with 99% success probability, as a function of problem size for different quantiles.

For D-Wave, we fix the annealing time at $20 \mu s$, the shortest time available due to engineering compromises, and estimate the probability p_i of finding the ground state for a given instance i [42]. Shorter times are optimal in this benchmark, as explained in Appendix A and Ref. [35]. The total annealing time to achieve $p_i = 0.99$ is $20 \mu s \times \log(1 - 0.99) / \log(1 - p_i)$. See Appendix A for details of the physical QA parameters (see also [43]).

We measure the computational effort of SA in units of runtime on a single core, which is natural when using server centers and convenient in order to compare to other classical approaches. Of course, the total runtime can be shortened by parallelizing the overall computation. Part of that process is embarrassingly parallel since SA finds its best solutions not in a single run, but by restarting from random bit configurations. Every restart could be executed on a different core. In fact, we used this strategy for our numerical benchmark. We find that median-case instances required 10^9 independent runs (with $945 \times 5 \cdot 10^4$ spin updates each) to find the optimal solution with 945 variables, on average.




FIG. 3. Layout of several instances of the weak-strong cluster network problem on the D-Wave 2X processor. Shown are three different sizes with 295, 490 and 945 qubits. Each cluster consists of an eight qubit unit cell of the Chimera graph. Orange dots depict qubits subject to a strong local field $h_R = -1$ while the cyan dots represent qubits with the weak field $h_L = 0.44$. Blue lines correspond to strong ferromagnetic couplings ($J = 1$) and red lines to strong antiferromagnetic couplings ($J = -1$). Note that the graphs are somewhat irregular due to the fact that not all 1152 qubits are operational.

We note that for instances with more qubits, more quantum hardware resources are brought to bear and therefore a fair comparison needs to take this into account [12, 15]. As an extreme example, one could contemplate building special purpose classical hardware that would update as many spins in parallel as possible at state-of-the-art clock rates. The sets of spins that could be updated in parallel depends on the connectivity graph. Though such considerations are reasonable, we do not explore this possibility here as we believe that future quantum annealers will have higher connectivity which will severely limit the usefulness of such parallelism.

When estimating runtimes for SA, we follow the protocol laid out by Refs. [12, 15, 44] and tuned SA for every problem size and quantile. Tuning means that the starting and end temperature, as well as the number of spin updates and the number of restarts are optimized to achieve a short overall runtime. We first measure the computational effort in units of sweeps (one sweep attempts to update all the spins). These times are plotted as $n_{\text{sweeps}} \times N \times T_{\text{su}}$, where N is the number of spins. We used a spin update time $T_{\text{su}} = 1/5$ ns (see Ref. [12]).

Our key finding in this comparison is that SA performs very poorly on the weak-strong cluster networks. The D-Wave 2X processor is $1.8 \cdot 10^8$ faster at the largest size

instances we investigated, which consisted of 945 variables. This problem was specifically engineered to cause the failure of SA: as explained above, the “weak-strong cluster networks” problem is intended to showcase the performance of annealers on a problem that benefits from finite range cotunneling. By contrast, the random Ising instances studied in Refs. [12, 15] have only low energy barriers, as explained in Ref. [45].

B. D-Wave versus Quantum Monte Carlo

Next we compared the performance of path integral Quantum Monte Carlo (QMC) with that of D-Wave for the same benchmark. QMC samples the Boltzmann distribution of a classical Hamiltonian which approximates the transverse field Ising model. In the case of a 2-spin model, the discrete imaginary time QMC classical Hamil-

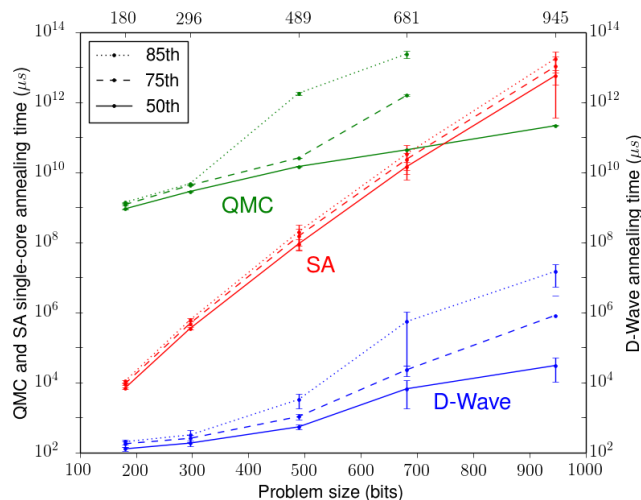


FIG. 4. Time to find the optimal solution with 99% probability for different problem sizes. We compare Simulated Annealing (SA), Quantum Monte Carlo (QMC) and the D-Wave 2X. To assign a runtime for the classical algorithms we take the number of spin updates (for SA) or worldline updates (for QMC) that are required to reach a 99% success probability and multiply that with the time to perform one update on a single state-of-the-art core. Shown are the 50th, 75th and 85th percentiles over a set of 100 instances. It occupied millions of processor cores for several days to tune and run the classical algorithms for these benchmarks. The runtimes for the higher quantiles for the largest problem size for QMC were not computed due to the high computational cost.

tonian is

$$H_{\text{cl}} = - \sum_{\tau=1}^M \left(\sum_{kj} \frac{J_{ij}}{M} \sigma_j(\tau) \sigma_j k(\tau) + J^{\perp}(s) \sum_j \sigma_j(\tau) \sigma_j(\tau + 1) \right), \quad (9)$$

where $\sigma_j(\tau) = \pm 1$ are classical spins, j, k are site indices, τ is a replica index, and M is the number of replicas. The coupling between replicas is given by

$$J^{\perp}(s) = -\frac{1}{2\beta} \ln \tanh \frac{A(s)\beta}{M}, \quad (10)$$

where β is the inverse temperature. The configurations for a given spin j across all replicas τ is called the worldline of spin j . Periodic boundary conditions are imposed between $\sigma_j(M)$ and $\sigma_j(1)$. We used continuous path integral QMC, which corresponds to the limit $\Delta\tau \rightarrow 0$ [46], and, unlike discrete path integral QMC, does not suffer from discretization errors of order $1/M$.

We numerically compute the number of sweeps n_{sweeps} required for QMC to find the ground state with 99% probability at different quantiles. In our case, a sweep corresponds to two update attempts for each worldline.

The computational effort is $n_{\text{sweeps}} \times N \times T_{\text{worldline}}$, where N is the number of qubits and $T_{\text{worldline}}$ is the time to update a worldline. We average $T_{\text{worldline}}$ over all the steps in the quantum annealing schedule; however the value of $T_{\text{worldline}}$ depends on the particular schedule chosen. As explained above for SA, we report the total computational effort of QMC in standard units of time per single core. For the annealing schedule used in the current D-Wave 2X processor, we find

$$T_{\text{worldline}} = \beta \times 870 \text{ ns} \quad (11)$$

using an Intel(R) Xeon(R) CPU E5-1650 @ 3.20GHz.

This study is designed to explore the utility of QMC as a classical optimization routine. Accordingly, we optimize QMC by running at a low temperature, 4.8 mK. We also observe that QMC with open boundary conditions (OBC) performs better than standard QMC with periodic boundary conditions in this case [38]; therefore, OBC is used in this comparison. We further optimize the number of sweeps per run which, for a given quantile, results in the lowest total computational effort. We find that the optimal number of sweeps is 10^6 at the largest problem size. This enhances the ability of QMC to simulate quantum tunneling, and gives a very high probability of success per run in the median case, $p_{\text{success}} = 0.16$.

All the qubits in a cluster have approximately the same orientation in each local minima of the effective mean field potential. Neighboring local minima typically correspond to different orientations of a single cluster. Here, tunneling time is dominated by a single purely imaginary instanton and is described by Eq. (35) below. It was recently demonstrated that, in this situation, the exponent a_{min}/\hbar for physical tunneling is identical to that of QMC [38]. As seen in Fig. 4, we do not find a substantial difference in the scaling of QMC and D-Wave (QA). However, we find a very substantial computational overhead associated with the prefactor B in the expression $T = B e^{D a_{\text{min}}/\hbar}$ for the runtime. In other words, B_{QMC} can exceed B_{QA} by many orders of magnitude. The role of the prefactor becomes essential in situations where the number of cotunneling qubits D is finite, i.e., is independent of the problem size N (or depends on N very weakly). Between some quantiles and system sizes we observe a prefactor advantage as high as 10^8 .

C. D-Wave versus other Classical Solvers

Based on the results presented here, one cannot claim a quantum speedup for D-Wave 2X, as this would require that the quantum processor in question outperforms the best known classical algorithm. This is not the case for the weak-strong cluster networks. This is because a variety of heuristic classical algorithms can solve most instances of Chimera structured problems much faster than SA, QMC, and the D-Wave 2X [47–49] (for a possible exception see [50]) [51]. For instance, the Hamze-de Freitas-Selby algorithm [47, 48], performs a large neighborhood

search by first replacing all 8 qubit clusters with a single large spin. In the particular case of a weak-strong cluster pair this algorithm avoids the formation of the energy barrier.

However, we believe that such solvers will become ineffective for the next generation of annealers currently being designed. The primary motivation for this optimism is that the connectivity graphs of future annealers will have higher degree. For example, we believe that a 10-dimensional cubic lattice is an engineerable architecture. For such dense graphs, we expect that cluster finding will become too costly.

We have also learned from the Janus team, working with special purpose FPGAs to thermalize Ising models on a 32^3 cube, that they found cluster finding not to be helpful [52].

A Remark on Scaling

Certain quantum algorithms have asymptotically better scaling than the best known classical algorithms. While such asymptotic behavior can be of tremendous value, this is not always the case; moreover, there can be substantial value in quantum algorithms which do not show asymptotically better scaling than classical approaches. The first reason for this is that current quantum hardware is restricted to rather modest problem sizes of less than order 1000 qubits. At the same time, when performing numerical simulations of quantum dynamics, in particular when doing open system calculations, we are often limited to problem sizes smaller than 100 qubits. Extrapolating from such finite size findings can be misleading as it is often difficult to determine whether a computational approach has reached its asymptotic behavior.

When forecasting the future promise of a given hardware design, there is a tendency to focus on the qubit dimension. However, this perspective is not necessarily helpful. For instance, when looking at runtime as a function of qubit dimension, one may conclude that the measurements reported here indicate that QMC and physical annealing have a comparable slope, scale similarly and that therefore, the upside of physical annealing is bounded. However, the large and practically important prefactor depends on a number of factors such as temperature. Furthermore, we expect future hardware to have substantially better control precision, substantially richer connectivity graphs and dramatically improved T_1 and T_2 times. With such changes, next generation annealers may drastically increase the constant separation between algorithms, leading to very different performance from generation to generation. To illustrate how dramatic this effect can be, when we ran smaller instances of the weak-strong cluster networks on the older D-Wave Vesuvius chips we predicted that at 1000 variables D-Wave would be 10^4 times faster than SA. In fact, we observed a speedup of more than a factor of 10^8 . This

is because certain noise parameters were improved and the new dilution refrigerator operates at a lower temperature. Similarly, we suspect that a number of previous attempts to extrapolate the D-Wave runtimes for 1000 qubits will turn out to be of limited use in forecasting the performance of future devices. For this reason, the current study focuses on runtime ratios that were actually measured on the largest instances solvable using the current device, rather than on extrapolations of asymptotic behavior which may not be relevant once we have devices which can attempt larger problems.

III. NUMERICAL STUDIES OF QUANTUM ANNEALING FOR GENERIC PROBLEMS WITH A RUGGED ENERGY LANDSCAPE

Runtime advantages for the quantum processor described above are only valuable if they extend to problems of practical interest. While rather obvious, it may be worth delineating a criteria for problems that are suitable for treatment with a quantum annealer:

1. Solutions to the problem are valuable or interesting.
2. The problem is representable on hardware that can be built in the near future.
3. Quantum annealing offers a runtime advantage.

A. Number Partitioning

A practical problem that fulfills these criteria is the Number Partitioning Problem (NPP). The NPP is defined as follows: Given a set of N positive numbers (a_1, \dots, a_N) find a partition \mathcal{P} of this set into two groups that minimizes the partition residue $E = |\sum_{i \in \mathcal{P}} a_i - \sum_{i \notin \mathcal{P}} a_i|$. A partition \mathcal{P} can be encoded by Ising spin variables $s_j = \pm 1$: $s_j = +1$ if $j \in \mathcal{P}$ and $s_j = -1$ otherwise. Thus, the NPP cost function is

$$E(\mathbf{s}) = |\Omega_{\mathbf{s}}|, \quad \Omega_{\mathbf{s}} = \sum_{j=1}^N a_j s_j \quad (12)$$

where $\mathbf{s} = (s_1, \dots, s_N)$ is a spin configuration and $\Omega_{\mathbf{s}}$ is a *signed* partition residue. Number partitioning is also one of Garey and Johnson's six basic NP-hard problems that lie at the heart of the theory of NP-completeness [53]. In studies of the average-case computational complexity of NPP, one usually assumes that $\{a_1, \dots, a_N\}$ are independent, identically-distributed random numbers. The average-case complexity of NPP is exponential in N when the number of bits b used to represent the numbers a_j satisfies the condition: $b/N \geq \kappa_c \simeq 1 - \frac{1}{2N} \log_2 N$ [54]. Our focus is on hard instances of NPP and we will be studying the random NPP ensemble with $b = N$. The numbers a_j are uniformly sampled from the interval $[0, 2^N - 1]$. NPP has many practical applications including multiprocessor

scheduling and the minimization of VLSI circuit size and delay, public key cryptography and others (see references in [55]).

NPP is attractive for numerical studies because for $n/N > \kappa_c$, the typical runtime of all known algorithms for NPP scales exponentially with large coefficients in the exponent and often, the asymptotic behavior can already be seen at sizes as low as 20 variables. For our purposes, NPP is a useful problem to study in the context of quantum annealing because it possesses extremely rugged energy landscapes where a single bit flip can result in dramatic energy changes. Its low energy band resembles that of the random energy model as there is almost no correlation between the state and its energy [56]. The 2^N signed partition residues Ω can be thought of as drawn from the Gaussian distribution

$$p(\Omega) = \frac{2}{\sqrt{2\pi N \langle a^2 \rangle}} \exp\left(-\frac{E^2}{2N \langle a^2 \rangle}\right); \quad (13)$$

where $\langle a^2 \rangle = \frac{1}{N} \sum_{j=1}^N a_j^2$. The distribution of the cost function values $E = |\Omega| \geq 0$ is given by $2p(E)$. By picking a bit string at random, one gets an average value of the cost function $\langle E \rangle = \sqrt{2 \langle a^2 \rangle N / \pi}$. The *minimum* value of this cost function is exponentially small, with median value $E_{\min} \sim \langle E \rangle 2^{-N}$ [55].

An obvious heuristic algorithm for NPP starts by placing the largest number in one of the two subsets. The next largest number is then placed in the set whose elements sum to the smallest value and this continues until all numbers are assigned. The idea behind this greedy heuristic is to keep the discrepancy small with every decision. This gives the scaling of the resulting partition residue as $\mathcal{O}(1/N)$. The differencing method of Karmarkar and Karp [55], also called the KK heuristics, is a polynomial time approximation algorithm. The key idea of this algorithm is to reduce the size of the numbers by replacing the two largest numbers by the absolute value of their difference. It has been proven [57] that the differencing method gives a minimum residue E_{\min}^{KK} such that

$$E_{\min}^{\text{KK}} \sim N^{-\alpha \log N}, \quad \alpha = 0.72. \quad (14)$$

The time complexity of both greedy and KK heuristics is $N \log N$ [55]. The residual energies reached by both methods are much smaller than the average partition residue $\langle E \rangle$, but far greater than the minimum residue E_{\min} . The absence of efficient heuristics is a particular feature of NPP. It is attributed to the extremely rugged energy landscape in the low part of the energy spectrum [55]. The statistics of the NPP energy landscape was studied analytically in [58] and numerically in [59].

This type of landscape leads to the exponential complexity of QA for NPP that was obtained in [58] via direct solution of the time-dependent Schrödinger equation. We observe that the particularly challenging instances of NPP violate the second condition of our “suitability” criteria as the numbers $a_i \in (0, \dots, 2^N - 1)$ should be drawn

Method	α
Branch and bound in conjunction with Karmarkar and Karp algorithm [60]	0.87
Simulated Annealing	0.98
Quantum Monte Carlo	0.81
Quantum Adiabatic Algorithm (exact diagonalization to obtain gap scaling)	0.80
Quantum Adiabatic Algorithm (time-dependent Schrödinger equation [58])	0.80
Quantum Annealing (quantum trajectory method)	0.82

TABLE I. Runtime scaling exponent for different methods to solve the Number Partitioning Problem. Values are obtained from a fit of the runtime $T \propto 2^{\alpha N}$ against numerically obtained runtimes. Notice the very poor scaling of Simulated Annealing. This is because for Simulated Annealing to work at all it has to be run at very high temperatures to overcome the enormous energy barriers present in this problem. However, at these high temperatures, SA behaves almost like random sampling and hence its scaling is almost that of exhaustive search. The scaling of Quantum Monte Carlo as well as other methods that simulate Hamiltonian evolution are comparable to each other and they all scale better than SA. This mirrors the situation we encountered for the weak-strong cluster networks. The Hamiltonian evolutions also scale better than the best known classical algorithm which employs a branch and bound method in conjunction with the Karmarkar and Karp algorithm. However, these scalings are close enough to warrant an estimation of the standard error in these values, which we have not done. The value of α for the solution of the time-dependent Schrödinger equation was initially obtained in [58].

from a set whose cardinality grows exponentially with N . This translates to a requirement that the bit precision for the coupling coefficients J_{ij} grow with N^2 if one were to express NPP as a quadratic binary optimization problem with objective function $\sum_{i,j=1}^N a_i a_j s_i s_j$ corresponding to the form (1) with $K = 2$. However, for numerical studies this is not a concern.

The runtime behavior of SA, QMC and QA, simulated using the Schrödinger equation or other methods, are shown in Table I. The relative performance of these algorithms applied to NPP is similar to their relative performance on the weak-strong cluster networks problem. QMC scales like QA and both scale better than SA. We believe this is the case because both problems are characterized by rugged energy landscapes for which tunneling transitions are a more useful way to reach low energy states than thermal transitions.

To achieve such scaling behavior it will be necessary for the size of the domains of cotunneling qubits to grow with the problem size. However it is interesting to explore a problem from a different perspective and ask a question: “*how much can the residual energy be lowered in QA until the system reaches such a state where lowering the energy further would require cotunneling of spin domains with*

sizes greater than k ?”.

To answer this question one can use an algorithm in which one starts at a random initial state and, at each step i) looks at all groups of bits of size k and ii) flips the bits of the group that results in the largest reduction of residual energy and then one iterates. We call this procedure "Algorithmic Tunneling" (AT). In other communities this would be referred to as k -opt or large neighborhood search. We emphasize that AT does not provide any information about the actual system dynamics during QA, nor the runtime of QA. We investigate AT as an upper-bound on the typical performance of QA. AT does not consider the entropic component of tunneling events arising from the statistics of different mechanisms for arriving in the same minima. Likewise, AT does not consider how the height of energy barriers effects tunneling events. However, AT allows us to develop our intuition about the value of an optimal tunneling procedure which always finds the lower energy solution within a finite Hamming distance.

To investigate the lowest residual energies that AT can reach we focus on the conditional distribution of the signed partition residues Ω' (12) over all possible spin configurations $\{\mathbf{s}'\}$ generated from a given (ancestor) configuration \mathbf{s} by simultaneous flipping of a fixed number of spins k . This conditional distribution was studied in [58] and has the form

$$P_k(\Omega|\Omega') = \frac{1}{\binom{N}{k} P(\Omega)} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{dx dx'}{8\pi^2} \zeta(x+x') \quad (15)$$

$$\times e^{\frac{i(\Omega-\Omega')x}{2}} e^{\frac{i(\Omega+\Omega')x'}{2}} \sum_{\mathbf{J}} \prod_{j \in \mathbf{J}} \cos(a_j x) \prod_{j \notin \mathbf{J}} \cos(a_j x')$$

where $P(\Omega)$ is given in (13) and $\zeta(s) = \sin(\Delta\Omega s/4)/((\Delta\Omega s/4))$. In the distribution $P_k(\Omega, \Omega')$, we averaged over the residue of the ancestor configuration within the interval $\Omega_s \in [\Omega - \Delta\Omega/2, \Omega + \Delta\Omega/2]$ where $\langle E \rangle / \binom{N}{k} \ll \Delta\Omega \ll \Omega$.

Near its maximum, the distribution $P_k(\Omega|\Omega')$ has the form

$$P_k(\Omega'|\Omega) = \frac{1}{\sqrt{2\pi N \sigma^2(q)}} \exp\left[-\frac{(\Omega' - q\Omega)^2}{2N\sigma^2(q)}\right], \quad (16)$$

where

$$\sigma(q) = \langle a^2 \rangle (1 - q^2)^{1/2}, \quad q = 1 - \frac{2k}{N}. \quad (17)$$

For small $k \ll N$ the width of the distribution is approximately $\sqrt{k\langle a^2 \rangle}$. After a single step of the AT algorithm, the average partition residue is reduced by a factor of $1 - 2k/N$. Therefore, once the the number of steps of AT far exceeds N/k , the algorithm reaches the residues $|\Omega| \ll \sqrt{k\langle a^2 \rangle}$. For those residues the distribution becomes

$$P_k(\Omega'|\Omega) \simeq P_k(0|0) = \frac{1}{\sqrt{8\pi k \langle a^2 \rangle}}. \quad (18)$$

In total we have $\binom{N}{k}$ samples of this distribution (spin configurations located on a Hamming distance k from the ancestor configuration). Therefore the minimum value is given by extreme order statistics [61]. I.e., within the set of bit configurations $\{\mathbf{s}'\}$ generated from a given (ancestor) configuration \mathbf{s} by simultaneous flipping of a fixed number of spins k , the minimum partition energy E is a random variable drawn from the exponential distribution

$$p_0(E) = \frac{1}{E_k} e^{-E/E_k}, \quad E_k = \left[\binom{N}{k} P_k(0|0) \right]^{-1}. \quad (19)$$

Therefore, average (and median) values of the minimum partition residue achieved in the AT algorithm, $E_{\min}^{\text{AT}}(k) = E_k$. For $1 \ll k \ll N$, we obtain

$$E_{\min}^{\text{AT}}(k) = 4\pi k \langle a^2 \rangle \left(\frac{k}{N} \right)^k e^{-k}. \quad (20)$$

It is instructive to compare the minimum cost values reached in AT and KK heuristics. Using (14) and (20) we get

$$\frac{E_{\min}^{\text{AT}}}{E_{\min}^{\text{KK}}} \sim N^{\alpha \log N - k} k^k e^{-k}, \quad \alpha = 0.72. \quad (21)$$

One can see from here that tunneling over barriers of length $k > \alpha \log N$ allows AT to reach cost values lower than that of the KK heuristics as N increases.

Consider AT with values of k that do not scale with N . We observe that in asymptotic limit

$$N \gg N_k = \frac{e}{k} \exp(k/\alpha), \quad (22)$$

the KK heuristic produces smaller residues than AT. If we consider tunneling with $k = 8$, corresponding to the case of the weak-strong clusters studied in this paper, then $N_k \simeq 22735$. We observe that for the high-precision ($B \geq N$) instances, NPP becomes intractable already for $N > 100$. Thus, for a broad range of problem sizes AT reaches cost values much smaller than conventional heuristics.

Motivated by above observation we conclude that asymptotic scaling behavior is not essential for this analysis. For example, one can choose AT with the barriers sizes $k = \alpha_0 \log N$ where $\alpha_0 - \alpha \gg 1/\log N$. In this case the ratio $E_{\min}^{\text{AT}}/E_{\min}^{\text{KK}}$ approaches 0 in the asymptotic limit $N \rightarrow \infty$. However the length of the barriers remains relatively small in a very broad range of N . For example, consider $\alpha_0 = 1.16$. Then for $N = 1000$ the barrier length is $k = 10$ while $E_{\min}^{\text{AT}}/E_{\min}^{\text{KK}} \sim 10^{-9}$.

In summary, greedy search procedure with flipping at most k bits at a time (referred above as Algorithmic Tunneling) allows to find cost values in NPP that are much below those given by KK heuristics for $k \sim 10$ at all realistic values of N .

B. K^{th} Order Binary Optimization

Our current best candidate for a problem class that fulfills all three criteria are K^{th} order binary optimization problems with $K > 2$. K^{th} order binary optimization is NP-Hard and occurs naturally in many engineering disciplines and many computational tasks. In unpublished work underway, we seek to establish that for many K -local problems, QA indeed offers a runtime advantage over SA. Currently we are focusing on $K \in \{4, 5, 6\}$. As energy landscapes get more rugged with higher K , our conjecture is that we will see larger subsets of instances for which QA runs faster as K increases. However, representing K -body terms in hardware becomes more challenging as K grows.

IV. SPIN COTUNNELING IN QA AND QMC

A. Instantons in systems with multiple spins

Cotunneling consists of system state transitions in which a group of spins simultaneously change their orientation with energy well below the energy of the (mean field) potential barriers. Tunneling is a quintessential quantum phenomenon; real time dynamics of classical trajectories cannot describe barrier penetration when the system wavefunction extends to classically forbidden regions. In such situations, the exponential decay of the wavefunction under the barrier is often captured through the path integral formalism by computing the minimum action of the trajectories in imaginary time [62, 63]. This approach was also extended to treat the tunneling of large magnetic moments with conserved total spin [64].

Tunneling in mean field spin models can be described using the path integral over spin-coherent states in imaginary time [65]. The tunneling path connects the minima of the mean field potential

$$V(\underline{\mathbf{m}}, t) = \langle \Psi_{\underline{\mathbf{m}}} | H(t) | \Psi_{\underline{\mathbf{m}}} \rangle, \quad (23)$$

where $H(t)$ is the time-dependent QA Hamiltonian from Eq. (2) and $|\Psi_{\underline{\mathbf{m}}}\rangle$ is a product state

$$|\Psi_{\underline{\mathbf{m}}}\rangle = \bigotimes_j \left[\cos \frac{\theta_j}{2} |0\rangle + e^{-i\phi_j} \sin \frac{\theta_j}{2} |1\rangle \right]. \quad (24)$$

The coherent state of the j -th spin is defined by a vector on the Bloch sphere

$$\mathbf{n}_j = (\sin \theta_j \cos \phi_j, \sin \theta_j \sin \phi_j, \cos \theta_j). \quad (25)$$

and the corresponding state of the N -spin system is defined by the vector $\underline{\mathbf{m}} = (\mathbf{n}_1, \mathbf{n}_2, \dots, \mathbf{n}_N)$.

Towards the beginning of a QA evolution, the system remains near $\underline{\mathbf{m}}_0(t)$, the global minimum of the time-dependent potential $V(\underline{\mathbf{m}}, t)$, which connects to the global minimum at the initial time. Later on, $V(\underline{\mathbf{m}}, t)$

undergoes a bifurcation, which may cause the initial minimum to become metastable. At that point, the system may be able to tunnel to the new global minimum $\underline{\mathbf{m}}_1(t)$ when $V(\underline{\mathbf{m}}_0) \simeq V(\underline{\mathbf{m}}_1)$. Here we omit the argument t , whose value corresponds approximately to the moment when the minima exchanges order. Such tunneling events are sometimes accompanied by thermal activation if QA is performed at finite temperatures (i.e. thermally assisted tunneling) [37]. The sequence of bifurcations and associated tunneling events can continue multiple times before the global minimum of H_P is reached.

Mathematically, the tunneling process can be described as an evolution of a spin system along the instanton path

$$\underline{\mathbf{q}}(\tau) = (\mathbf{n}_1(\tau), \mathbf{n}_2(\tau), \dots, \mathbf{n}_N(\tau)) \quad (26)$$

over imaginary time $\tau \in (0, \beta)$ where $\beta = \hbar/k_B T$ and T is the system temperature. The details of this analysis will be provided elsewhere [66]. Here, we simply outline the main argument. The path component for the j -th spin is described by a vector $\mathbf{n}_j(\tau)$ with periodic boundary conditions $\mathbf{n}_j(0) = \mathbf{n}_j(\beta)$. The vector $\mathbf{n}(\tau)$ is complex (see Eqs. (67),(68) in the Supplementary Material of Ref. [38]) and can be written in the form

$$\mathbf{n} = (\sin \theta_j \cosh \varphi_j, -i \sin \theta_j \sinh \varphi_j, \cos \theta_j), \quad (27)$$

corresponding to a purely imaginary azimuthal angle $\phi_j(\tau) = -i\varphi_j(\tau)$. The initial point of the instanton $\underline{\mathbf{q}}(0)$ corresponds to the initial minimum of $V(\underline{\mathbf{m}})$. The midpoint of the trajectory $\underline{\mathbf{q}}(\beta/2)$ typically corresponds to the exit point of the potential barrier in the vicinity of the final minima

$$\underline{\mathbf{q}}(0) = \underline{\mathbf{q}}(\beta) \simeq \underline{\mathbf{m}}_0, \quad \underline{\mathbf{q}}(\beta/2) \simeq \underline{\mathbf{m}}_1. \quad (28)$$

We first consider the simplified situation where the domain of D spins tunnel simultaneously. The total spin of the tunneling domain is conserved and all spins in the domain move identically through the instanton trajectory, $\mathbf{n}_j(\tau) \equiv \mathbf{n}(\tau)$ for all $j \in [1, D]$. Then, the mean field potential for the instanton can be rescaled as

$$V(\underline{\mathbf{q}}(\tau)) = Dv(\mathbf{n}(\tau)) \quad (29)$$

In the limit of low temperatures $\beta \rightarrow \infty$, the instanton action takes the form $S[\mathbf{n}(\tau)] = Da[\mathbf{n}(\tau)]$ with

$$a[\mathbf{n}(\tau)] = \frac{\hbar}{2} \omega[\mathbf{n}(\tau)] + \int_0^\infty d\tau v[\mathbf{n}(\tau)], \quad (30)$$

where ω describes a ‘‘Berry phase’’ type contribution

$$\omega[\mathbf{n}(\tau)] = \int_0^\infty d\tau (1 - \cos \theta(\tau)) \dot{\varphi}(\tau). \quad (31)$$

The exponential dependence of the tunneling rate,

$$W_{\text{QA}} \sim e^{-Da_{\min}/\hbar}, \quad a_{\min} = \min_{\mathbf{n}(\tau)} a[\mathbf{n}(\tau)], \quad (32)$$

is given by the minimum of the rescaled action.

We now consider the case when the tunneling of the spin domain is enabled by the transverse field and the Hamiltonian is of the type in Eq. (2),

$$H = -A \sum_{j=1}^N \sigma_j^x - BH_P^{\text{cl}}(\sigma_1^z, \dots, \sigma_D^z), \quad (33)$$

where $H_P^{\text{cl}}(s_1, \dots, s_D)$ is a classical cost function of binary variables $s_k = \pm 1$. Assuming that $H_P^{\text{cl}} = H_P^{\text{cl}}(\sum_{j=1}^D \sigma_j^z)$, and the system is in a state of maximum total spin, all spins will tunnel together. Thus,

$$v(\mathbf{n}(\tau)) = -A \sin \theta(\tau) \cosh \varphi(\tau) - Bg(\cos \theta(\tau)), \quad (34)$$

where the function $g(x) = D^{-1}H_P^{\text{cl}}(Dx)$ is the rescaled mean field potential energy of the spin system. At zero temperature, the minimum action can be written in a more familiar form (cf. Supplementary Material in Ref. [38] and also Ref. [67]),

$$\frac{a_{\min}}{\hbar} = \int_{\theta_0}^{\theta_1} \operatorname{arcsinh} \left(\frac{v(\theta)}{A \sin \theta} \right) \sin \theta d\theta, \quad (35)$$

where $v(\theta) = \sqrt{B^2(g(\cos \theta) - g(\cos \theta_0))^2 - A^2 \sin^2 \theta}$ is a linear velocity along the instanton path and the angles θ_i correspond to the minima of the potential $v(\mathbf{n})$ (the values of $\sinh \varphi_j = 0$ at the minima).

In general, in mean field spin models, the total spin is not conserved throughout the instanton trajectory (26). The corresponding action [65] can be written in a form that is a direct generalization of the simplified case (30)

$$S[\underline{\mathbf{q}}(\tau)] = \frac{1}{2} \sum_{i=1}^D \omega[\mathbf{n}_i(\tau)] + \int_0^\beta d\tau V[\underline{\mathbf{q}}(\tau)], \quad (36)$$

where ω is given in Eq. (31). If the azimuthal angles are purely imaginary, as discussed above, the action is real. The dominant contribution to the path integral will be given by the instanton path with the least action

$$W_{\text{QA}} \sim e^{-S_{\min}/\hbar}, \quad S_{\min} = \min_{\underline{\mathbf{q}}(\tau)} S[\underline{\mathbf{q}}(\tau)]. \quad (37)$$

In general, the action in Eq. (37) also grows linearly with the number of cotunneling spins as in the simplified case of Eq. (32) because the individual spin contributions to the action are highly correlated at the instanton trajectory. Accordingly, we may define $S_{\min} = Da_{\min}$.

B. Tunneling simulation in QMC

In the path integral QMC algorithm one introduces an extra dimension associated with the imaginary time axis in order to simulate the multispin tunneling phenomenon on a classical computer, as seen above. This

is done by using a Suzuki-Trotter decomposition and representing the partition function of the system Z in terms of the path integral over the spin trajectories $\underline{\sigma}(\tau) = \{\sigma_j(\tau)\}_{j=1}^N$, see Eq. (9). These trajectories are periodic along the imaginary time axis $\sigma_j(0) = \sigma_j(\beta)$. For each spin j , the set of values $\sigma_j(\tau) = \pm 1$ form a path component referred to as a worldline. The time step along the worldline is $\Delta\tau = \beta/M$ where M is the number of Trotter slices (i.e. the number of spin replicas in the worldline). Sampling the system states along this extra dimension introduces an additional overhead in classical computation that does not exist for the corresponding quantum dynamics.

The runtime T_{QMC} of QMC can be thought of as a product of three factors,

$$T_{\text{QMC}} = N n_{\text{sweeps}} T_{\text{worldline}} \quad (38)$$

where N is the problem size that is equal to the number of worldlines. The number of sweeps n_{sweeps} , in general, depends exponentially on the typical size D of the cotunneling domain, $n_{\text{sweeps}} \propto e^{\alpha D}$, where $\alpha = \alpha(\beta)$ also depends on the inverse temperature β . In cases where $D = \mathcal{O}(N)$ the growth of this factor with N reflects a major computational bottleneck of QMC and QA. As was shown recently by some of the authors [38], the exponent in QMC and QA is the same for a broad class of problems.

According to the findings reported in this paper, the prefactor in n_{sweeps} , along with the factors N and $T_{\text{worldline}}$ in T_{QMC} , are significantly different for QMC and QA. The value of $T_{\text{worldline}}$ found in our simulations is given in Eq. (B1). We also find that

$$n_{\text{sweeps}} \gg 1 \text{ ns}/T_{\text{QA}}, \quad (39)$$

where T_{QA} is the duration of QA and we use a normalization factor of one nanosecond, corresponding to the typical time scale of superconducting QA devices. We expect that the above relation will remain true even if the scaling of both quantities with Da_{\min}/\hbar is the same.

C. Comparison of QA and QMC for the “weak-strong cluster pair” problem

We use the modeling considerations described above to compare QA and QMC in the system corresponding to the weak-strong cluster pair problem (see discussion in Sec. II, Fig. 2, and Eqs. (6),(7),(8)). Tunneling in this system corresponds to an avoided crossing between the two lowest energy levels of the Hamiltonian, shown in Fig. 5. All other levels lie high above the first two and are never excited. During tunneling, the total spin of the left cluster reverses orientation. The number of cotunneling spins is $D = 8$ in this case while the total number of spins $N = 16$.

Proper analysis of the tunneling probabilities and related QA success rates should also account for coupling to the environment. We studied the success probabilities




FIG. 5. Gap of the quantum Hamiltonian for $h_1 = 0.44$ as a function of the annealing parameter. The solid line is the energy difference between the ground state and the first excited state. The avoided crossing at $t/T_{QA} = 0.62$ corresponds to a minimum gap of 248 MHz. The next excited state is separated by a gap in excess of 2 GHz.

in QA for the 2-cluster problem using the approach developed in [35]. The results are summarized in Fig. 11 of App. A. Decreasing the temperature of QA compared to the temperature of the D-Wave device suppresses the effect of thermal excitation from the ground state and leads to an increase of the final success probability p_0 to find the ground state. Once the temperature reaches 5 mK, the success probability stays above 90%, even at QA schedules that are faster than $T_{QA}=300$ nanoseconds. On the other hand, adiabatic transitions near the avoided crossing are suppressed even at $T_{QA} \simeq 71$ nanoseconds, as can be seen from solutions to the time-dependent Schrödinger equation, shown in Fig. 12.

In the previous studies involving D-Wave devices (see [35] for references) it was inferred from the data that the low frequency noise components of the spectral density, providing a leading contribution to the qubit line-width W (A1),(A2), have effective frequency cutoffs much below 314 MHz (15 mK). In the current QA schedule of 20 μ s, the system spends only a small fraction of this time in the vicinity of the avoided crossing where thermal excitations from the ground state are possible. For a QA schedule duration of ~ 100 ns, we expect that the effective noise strength will be weaker than at the current schedule. This would lead to the suppression of the thermal excitations from the ground state.

To compare simulations of QA at a temperature of 5 mK with the QMC performance, we choose a duration of QA such that the probability to reach the ground state at the end of QA equals 0.95. As we discussed above this can be achieved at

$$T_{QA} = 71 \text{ ns} . \quad (40)$$

In setting up the QMC simulations our objective is to select the two parameters, number of sweeps per qubit n_{sweeps} and β , to minimize T_{QMC} for a given probability

of success p_0 to find the system at the end of the QA in the ground state where all spins point down. Essentially we need to minimize the product of βn_{sweeps} keeping p_0 fixed.




FIG. 6. Probability of success versus β for QMC with the DWave 2X schedule. Different colors correspond to different number of sweeps (see legend). The embedded plot shows the probability of success at saturation for different number of sweeps. We use periodic boundary conditions, which performed better than open boundary conditions in this case.

In Fig. 6 we plot the success probability of QMC $p_0 = p_0(\beta, n_{\text{sweeps}})$ as a function of β (inverse temperature) for different number of sweeps. We see that increasing β increases p_0 ; however, the success probability saturates at some value

$$p_0(\beta, n_{\text{sweeps}}) \leq p_0^{\text{sat}}(n_{\text{sweeps}}), \quad (41)$$

which itself depends on the number of sweeps. The saturation probability $p_0^{\text{sat}}(n_{\text{sweeps}})$ is plotted in the insert to Fig. 6. By fixing the success probability $p_0 = 0.95$ we select the optimal number of sweeps. Then, by looking on the main plot we determine the value of β_{sat} where saturation occurs. A detailed study shows that this procedure gives the minimum value of the product βn_{sweeps} . The optimal values are

$$n_{\text{sweeps}} = 23,000, \quad \beta = 32.5 . \quad (42)$$

The total time to update one worldline with the D-Wave 2X schedule is

$$T_{\text{worldline}} = 28.3 \mu\text{s} \quad (43)$$

and the total runtime QMC per qubit, according to Eq. (38), is

$$\frac{T_{QMC}}{N} = 0.75 \text{ seconds} . \quad (44)$$

By comparing this with Eq. (40), we estimate that

$$\frac{T_{QMC}/N}{T_{QA}} \sim 10^7 \quad (T_{QA} = 71 \text{ ns}, \quad T = 5 \text{ mK}). \quad (45)$$

This ratio will need to be multiplied by the number of qubits to obtain the overall speed up factor (e.g., $\sim 10^{10}$ for 1000 qubits).

Implementing fast QA schedules or operating flux qubits at 5 mK will require improvements in the control electronics and other elements of the design. Furthermore, readout will need to be made much faster than in the current D-Wave devices. However, the estimates presented above serve to emphasize the significant promise of QA, as compared to QMC when the system adiabatic evolution “under the gap” becomes coherent and thermal excitations are suppressed.

V. DESIGNING FUTURE ANNEALERS OF PRACTICAL RELEVANCE

Should numerical studies confirm that QA offers a substantial runtime advantage, there is still one more hurdle to overcome. We need to ensure that K -local problems can be economically represented in hardware. We would like to be able to tell a user: “if you have a binary optimization problem with N variables and L terms, and the many-body order of the highest term is K , then you can send this problem to the quantum annealing co-processor”. However, annealers built to-date only support pairwise qubit couplings, i.e. $K = 2$. Two routes have been proposed to increase the locality.

One route is to build physical K -body couplers. However, it may prove difficult to layout K -local couplers on a two dimensional chip or even in layered architectures. Of course, the general case in which one aims to implement all possible $\sum_{k=1}^K \binom{N}{k}$ couplings will be infeasible. While many applications will only necessitate $L = O(N)$ coupling terms, this could still prove challenging. Furthermore, economically embedding problems in a fixed graph with only a limited number of specific K -local terms may prove difficult.

Another possibility is that we could use logical embeddings that map K -local problems to 2-local problems. A new proposal on how to accomplish such embeddings has been put forth [68], which has reinvigorated interest in this direction. Our main worry regarding any reduction to quadratic problems is that this will involve ancillary qubits. As argued above, it is crucial that the problem features tall and narrow barriers for tunneling transitions to contribute and the introduction of additional variables may cause these barriers to become wider. This makes purely thermal annealing more competitive and may negate gains seen in the numerics prior to embedding.

VI. SUMMARY

It is often quipped that Simulated Annealing is only for the “ignorant or desperate”. Yet, in our experience we find that lean stochastic local search techniques such as SA are often very competitive and they continue to be one of the most widely used optimization schemes. This is because for sufficiently complex optimization tasks with little structure to exploit (such as instances of K^{th} order binary optimization) it often takes considerable expert effort to devise a faster method. Therefore we regard SA as the generic classical competition that Quantum Annealing needs to beat.

Here we showed that for carefully crafted proof-of-principle problems with rugged energy landscapes that are dominated by large and tall barriers QA can have a significant runtime advantage over SA. We found that for problem sizes involving nearly 1000 binary variables, quantum annealing is more than 10^8 times faster than SA running on a single core. We also compared the quantum hardware to QMC. While the scaling of runtimes with size between these two methods is comparable, they are again separated by a large factor sometimes as high as 10^8 .

For higher order optimization problems, rugged energy landscapes will become typical. As we saw in our experiments with the D-Wave 2X, problems with such landscapes stand to benefit from quantum optimization because quantum tunneling makes it easier to traverse tall and narrow energy barriers. Therefore, we expect that quantum annealing might also deliver runtime advantages for problems of practical interest such as K^{th} order binary optimization with larger K .

More work is needed to turn quantum enhanced optimization into a practical technology. The design of next generation annealers must facilitate the embedding of problems of practical relevance. For instance, we would like to increase the density and control precision of the connections between the qubits as well as their coherence. Another enhancement we wish to engineer is to support the representation not only of quadratic optimization but of higher order optimization as well. This necessitates that not only pairs of qubits can interact directly but also larger sets of qubits. Such improvements will also make it easier for end users to input hard optimization problems.

The work presented here focused on the computational resource that is experimentally most accessible for quantum annealers: finite range tunneling. However this analysis is far from complete. A coherent annealer could accelerate the transition through saddle points, an issue slowing down the training of deep neural networks, for reasons similar to those that make a quantum walk spread faster than a classical random walker [69–71]. It could also dramatically accelerate sampling from a probability distribution via the mechanism of many body delocalization [72]. The computational value of such physics still needs to be properly understood and modeled.




FIG. 7. Schedule functions for D-Wave 2X chip. The annealing parameter is $s = t/T_{\text{QA}}$ for time t and total annealing time T_{QA} .

Acknowledgements – We would like to thank Matthias Troyer for discussions and Edward Farhi and Masoud Mohseni for helping to review the manuscript.

Appendix A: Quantum Annealing results for Weak-Strong cluster problem with 16 qubits

We developed a detailed modeling of the quantum annealing process and incoherent multi-qubit cotunneling for the weak-strong cluster problem in Ref. [35]. Using a noise model with experimentally measured parameters for the D-Wave 2X processor, we numerically verified that the spins arrive at the energetically more favorable configuration via multi-qubit tunneling. In what following we will refer to the modeling in this reference for the details.

In the present study we apply this detailed model for the new schedule functions $A(s)$, $B(s)$ (see Fig. 7) and for the new values of the noise parameters, line-width W , Ohmic coefficient η , and temperature of the device T for the D-Wave 2X processor. The noise parameters were measured near the end of the quantum annealing schedule, $s = 1$. The values of the noise parameters at a point during the annealing can be related to the measured ones (see Appendix A 5 in [35])

$$\begin{aligned} (W(s)/W_{\text{MRT}})^2 &= \eta/\eta_{\text{mRT}} = B(s)/B(1) \\ W_{\text{MRT}} &= 530\text{MHz}, \quad \eta_{\text{MRT}} = 0.12, \quad T \simeq 12\text{mK} \end{aligned} \quad (\text{A1})$$

The population $p_0(t)$ of the ground state during the QA process obeys the equations

$$\begin{aligned} \frac{dp_0}{dt} &= -(W_{01}(s) + W_{10}(s))p_0(s) + W_{10}(s), \quad (\text{A2}) \\ \frac{W_{01}(s)}{W_{10}(s)} &= e^{-\Delta_{10}(s)/k_B T}, \quad \Delta_{10}(s) = E_1(s) - E_0(s), \end{aligned}$$




FIG. 8. Logarithmic plot of the transition rate $W_{10}(s)$ vs s after the avoided crossing. Plot corresponds to $T = 12\text{mK}$ and $h_1 = 0.44$.

where $W_{ij}(s)$ is a transition rate from the state i to the state j whose explicit form is given in Ref. [35]. The transition rate $W_{10}(s)$ decays very fast after the avoided crossing (see Fig. 8) because the weak cluster (left cluster in Fig. 2) becomes progressively more polarized along the z -axis and the effective size of the tunneling domain $D = D(s)$ grows. This gives rise to the multi-qubit “freezing phenomenon” where the system population gets partially trapped in the excited state after certain value of s in later stages of the QA Ref. [35]. Fig. A2 shows the the ground state population given by the solution of (A2). The success probability to be at the ground state at the end of the annealing is $p_0 = 0.88$, which is very close to the experimentally observed mean value of 0.9. It is seen that the equilibrium population of the ground state exceeds the actual population for $s \gtrsim 0.64$, corresponding to the onset of freezing of the transition rates.

Appendix B: D-Wave versus Quantum Monte Carlo with Linear Schedule

There is ongoing work directed to optimized the QMC parameters further. In preliminary results, we compared QMC with a linear schedule against D-Wave [73]. The transverse field in this case is lower, resulting in a faster time $T_{\text{worldline}}$ to update a worldline ($T_{\text{worldline}}$ scales linearly with the transverse field). We measured this time to be

$$T_{\text{worldline}} = \beta \times 115 \text{ ns}. \quad (\text{B1})$$

This time is consistent with the one reported in Ref. [12].




FIG. 9. Probability of occupation of the ground state during the QA process is shown with solid black line. Red dashed line gives the equilibrium population of the ground state. Results are given at $T = 12\text{mK}$ and $h_1 = 0.44$. Dashed red line gives the equilibrium population of the ground state. The quantum annealing time is $t_{\text{QA}} = 20\ \mu\text{sec}$.

We also took a different approach when optimizing β and the number of sweeps per run to minimize the total computational effort. In the case reported in Sec. II B we optimized the number of sweeps for each quantile at fixed $\beta = 10$. In the case of a linear schedule, we used our knowledge of the structure of the weak-strong cluster networks problem to optimize β and the number of sweeps n_{sweeps} concurrently. We first measure the probability of success $p(n_{\text{sweeps}}, \beta)$ for a single weak-strong cluster pair. Then we estimate the performance for the cluster network problems taking into account the number of cluster pairs c for each size. The estimate is

$$\text{total time} \propto n_{\text{sweeps}} \times \beta \times \left[\frac{\log(1 - 0.99)}{\log(1 - p(n_{\text{sweeps}}, \beta)^c)} \right]. \quad (\text{B2})$$

Here we run with periodic boundary conditions (PBC) in imaginary time. We estimate an optimal $\beta = 130$ for all sizes. We then optimize the number of sweeps for each quantile and size running the actual benchmark.

The results, following the same methodology as in Sec. II B, are plotted in Fig. 13. We obtain a prefactor $O(10^7)$ for the median and up to $O(10^9)$ in the 85th quantile.

When optimizing QMC to the extent performed here a methodological concern arises. Since QMC has many parameters and modes of execution (e.g. temperature, number of sweeps, annealing schedule, open versus closed boundary conditions in imaginary time, discrete or con-




FIG. 10. Logarithmic plots of the transition rate $W_{10}(s)$ vs s after the avoided crossing for different temperatures. Red, green, mustard and blue colors corresponds to $T = 12\text{mK}$, 8 mK, 6 mK and 4 mK respectively. All plots correspond to $h_1 = 0.44$ and $T_{\text{QA}} = 20\ \mu\text{sec}$.




FIG. 11. Success probability of QA p_0 vs T . Different colors corresponds to different duration of the QA process T_{QA} . Plots correspond to $h_1 = 0.44$.

tinuous time Monte Carlo), overlearning can become an issue when working with just 100 instances. Moreover, optimizations over many parameters will become computationally prohibitive as the problem size increases. By contrast, the current quantum hardware has only a single parameter that can be tuned, the number of annealing sweeps.




FIG. 12. Plots show the results of the solution of the time-dependent Schrödinger equation for the closed system QA. Probabilities of occupation of ground state and first excited state as a function of QA time are plotted with blue and red points, respectively. The QA time to reach the probability of success 0.95 equals to 70.9 ns. All plots correspond to $h_1 = 0.44$.




FIG. 13. Time to find the optimal solution with 99% probability for different problem sizes. We compare Quantum Monte Carlo (QMC) with a linear schedule and the D-Wave 2X. To assign a runtime for QMC we take the number of worldline updates that are required to reach a 99% success probability and multiply that with the time to perform one update on a single state-of-the-art core. Shown are the 50th, 75th and 85th percentiles over a set of 100 instances. The runtimes for the higher quantiles for the largest problem size for QMC were not computed because the computational cost was too high.

- [1] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, *Science* **220**, 671 (1983).
- [2] T. Kadowaki and H. Nishimori, *Physical Review E* **58**, 5355 (1998); J. Brooke, D. Bitko, T. F. Rosenbaum, and G. Aeppli, *Science* **284**, 779 (1999); Y. H. Lee and B. J. Berne, *The Journal of Physical Chemistry A* **104**, 86 (2000); E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, *Science* **292**, 472 (2001); G. E. Santoro, R. Martonak, E. Tosatti, and R. Car, *Science* **295**, 2427 (2002).
- [3] A. Das and B. K. Chakrabarti, *Quantum annealing and related optimization methods*, Vol. 679 (Springer Science & Business Media, 2005).
- [4] A. Das and B. K. Chakrabarti, *Reviews of Modern Physics* **80**, 1061 (2008).
- [5] R. Harris, M. W. Johnson, T. Lanting, A. J. Berkley, J. Johansson, P. Bunyk, E. Tolkacheva, E. Ladizinsky, N. Ladizinsky, T. Oh, F. Cioata, I. Perminov, P. Spear, C. Enderud, C. Rich, S. Uchaikin, M. C. Thom, E. M. Chapple, J. Wang, B. Wilson, M. H. S. Amin, N. Dickson, K. Karimi, B. Macready, C. J. S. Truncik, and G. Rose, *Phys. Rev. B* **82**, 024511 (2010).
- [6] M. Johnson, M. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. Berkley, J. Johansson, P. Bunyk, *et al.*, *Nature* **473**, 194 (2011).
- [7] V. Bapst, L. Foini, F. Krzakala, G. Semerjian, and F. Zamponi, *Physics Reports* **523**, 127 (2013).
- [8] S. Boixo, T. Albash, F. M. Spedalieri, N. Chancellor, and D. A. Lidar, *Nat. Commun.* **4** (2013).
- [9] N. Dickson, M. Johnson, M. Amin, R. Harris, F. Altomare, A. Berkley, P. Bunyk, J. Cai, E. Chapple, P. Chavez, *et al.*, *Nat. Commun.* **4**, 1903 (2013).
- [10] C. C. McGeoch and C. Wang, in *Proceedings of the ACM International Conference on Computing Frontiers (ACM, 2013)* p. 23.
- [11] S. Dash, arXiv:1306.1202.
- [12] S. Boixo, T. F. Rønnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer, *Nature Physics* **10**, 218 (2014).
- [13] T. Lanting, A. Przybysz, A. Y. Smirnov, F. Spedalieri, M. Amin, A. Berkley, R. Harris, F. Altomare, S. Boixo, P. Bunyk, *et al.*, *Phys. Rev. X* **4**, 021041 (2014).
- [14] S. Santra, G. Quiroz, G. Ver Steeg, and D. A. Lidar, *New J. Phys.* **16**, 045006 (2014).
- [15] T. F. Rønnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer, *Science* **345**, 420 (2014).
- [16] W. Vinci, K. Markström, S. Boixo, A. Roy, F. M. Spedalieri, P. A. Warburton, and S. Severini, *Sci. Rep.* **4** (2014).
- [17] S. W. Shin, G. Smith, J. A. Smolin, and U. Vazirani, arXiv:1401.7087.
- [18] W. Vinci, T. Albash, A. Mishra, P. A. Warburton, and D. A. Lidar, arXiv:1403.4228.
- [19] C. C. McGeoch, *Synthesis Lectures on Quantum Computing* **5**, 1 (2014).
- [20] D. Venturelli, S. Mandrà, S. Knysh, B. O’Gorman, R. Biswas, and V. Smelyanskiy, arXiv:1406.7553.
- [21] T. Albash, T. F. Rønnow, M. Troyer, and D. A. Lidar, *The European Physical Journal Special Topics* **224**, 111 (2015).
- [22] A. D. King and C. C. McGeoch, arXiv:1410.2628 (2014).
- [23] P. Crowley, T. Duric, W. Vinci, P. Warburton, and A. Green, *Physical Review A* **90**, 042317 (2014).
- [24] W. Vinci, T. Albash, G. Paz-Silva, I. Hen, and D. A. Lidar, *Physical Review A* **92**, 042310 (2015).
- [25] I. Hen, J. Job, T. Albash, T. F. Rønnow, M. Troyer, and D. Lidar, arXiv preprint arXiv:1502.01663 (2015).
- [26] D. S. Steiger, T. F. Rønnow, and M. Troyer, arXiv preprint arXiv:1504.07991 (2015).
- [27] D. Venturelli, D. J. Marchand, and G. Rojo, arXiv preprint arXiv:1506.08479 (2015).
- [28] B. Bauer, L. Wang, I. Pizorn, and M. Troyer, arXiv:1501.06914 (2015).
- [29] T. Albash, I. Hen, F. M. Spedalieri, and D. A. Lidar, arXiv preprint arXiv:1506.03539 (2015).
- [30] H. G. Katzgraber, F. Hamze, Z. Zhu, A. J. Ochoa, and H. Munoz-Bauza, arXiv:1505.01545 (2015).
- [31] N. Chancellor, S. Szoke, W. Vinci, G. Aeppli, and P. A. Warburton, arXiv preprint arXiv:1506.08140 (2015).
- [32] A. Perdomo-Ortiz, J. Fluegemann, R. Biswas, and V. N. Smelyanskiy, arXiv preprint arXiv:1503.01083 (2015).
- [33] A. Perdomo-Ortiz, B. O’Gorman, J. Fluegemann, R. Biswas, and V. N. Smelyanskiy, arXiv preprint arXiv:1503.05679 (2015).
- [34] W. Vinci, T. Albash, and D. A. Lidar, arXiv preprint arXiv:1511.07084 (2015).
- [35] S. Boixo, V. N. Smelyanskiy, A. Shabani, S. V. Isakov, M. Dykman, V. S. Denchev, M. Amin, A. Smirnov, M. Mohseni, and H. Neven, arXiv:1411.4036 .
- [36] In still unpublished experiments we saw evidence of tunneling events involving up to 12 qubits. However it is difficult to exclude higher order tunneling processes that may break up the cotunneling group.
- [37] K. Kechedzhi and V. N. Smelyanskiy, arXiv:1505.05878.
- [38] S. V. Isakov, G. Mazzola, V. N. Smelyanskiy, Z. Jiang, S. Boixo, H. Neven, and M. Troyer, arXiv:1510.08057.
- [39] B. Altshuler, H. Krovi, and J. Roland, *Proceedings of the National Academy of Sciences* **107**, 12446 (2010).
- [40] E. Farhi, D. Gosset, I. Hen, A. W. Sandvik, P. Shor, A. P. Young, and F. Zamponi, *Phys. Rev. A* **86**, 052334 (2012).
- [41] S. Knysh, arXiv:1506.08608.
- [42] We take the mean over 16 gauges.
- [43] R. Harris, J. Johansson, A. J. Berkley, M. W. Johnson, T. Lanting, S. Han, P. Bunyk, E. Ladizinsky, T. Oh, I. Perminov, E. Tolkacheva, S. Uchaikin, E. M. Chapple, C. Enderud, C. Rich, M. Thom, J. Wang, B. Wilson, and G. Rose, *Phys. Rev. B* **81**, 134510 (2010).
- [44] S. V. Isakov, I. N. Zintchenko, T. F. Rønnow, and M. Troyer, *Computer Physics Communications* **192**, 265 (2015).
- [45] H. G. Katzgraber, F. Hamze, and R. S. Andrist, *Phys. Rev. X* **4**, 021008 (2014).
- [46] H. Rieger and N. Kawashima, cond-mat/9802104.
- [47] F. Hamze and N. de Freitas, arXiv:1207.4149 .
- [48] A. Selby, arXiv:1409.3934 (2014).
- [49] I. Zintchenko, M. B. Hastings, and M. Troyer, *Physical Review B* **91**, 024201 (2015).
- [50] J. King, S. Yarkoni, M. M. Nevisi, J. P. Hilton, and C. C. McGeoch, arXiv:1508.05087 .
- [51] Nevertheless, it is interesting to note that Ref. [15] defines

a *limited quantum speedup* as “a speedup obtained when comparing specifically with classical algorithms that correspond to the quantum algorithm in the sense that they implement the same algorithmic approach, but on classical hardware ... a natural example is quantum annealing implemented on a candidate physical quantum information processor versus either classical simulated annealing, classical spin dynamics, or simulated quantum annealing.” In the weak-strong cluster networks benchmark we show that D-Wave 2X outperforms these three algorithms. We find a substantial advantage against QMC in the prefactor, not the scaling. For the comparison with classical spin dynamics please see Ref. [35].

- [52] Victor Martin-Mayor, personal communication.
- [53] M. R. Garey and D. S. Johnson, San Francisco, LA: Freeman (1979).
- [54] S. Mertens, Physical Review Letters **81**, 4281 (1998).
- [55] S. Mertens, Theoretical Computer Science **265**, 79 (2001).
- [56] S. Mertens, Physical Review Letters **84**, 1347 (2000).
- [57] B. Yakir, Mathematics of Operations Research **21**, 85 (1996).
- [58] V. N. Smelyanskiy, U. v. Toussaint, and D. A. Timucin, quant-ph/0202155.
- [59] P. F. Stadler, W. Hordijk, and J. F. Fontanari, Physical Review E **67**, 056701 (2003).
- [60] R. E. Korf, Artificial Intelligence **106**, 181 (1998).
- [61] J. E. Angus, Technometrics **32**, 110 (1990).
- [62] S. Coleman, Physical Review D **15**, 2929 (1977).
- [63] A. I. Vainshtein, V. I. Zakharov, V. A. Novikov, and M. A. Shifman, Soviet Physics Uspekhi **25**, 195 (1982).
- [64] E. M. Chudnovsky and J. Tejada, *Macroscopic Quantum Tunneling of the Magnetic Moment* (Cambridge, UK: Cambridge University Press, 1998).
- [65] N. Nagaosa, *Quantum field theory in condensed matter physics* (Springer Science & Business Media, 2013).
- [66] V. N. Smelyanskiy, S. Isakov, S. Boixo, and H. Neven, (to be published).
- [67] A. Garg, J. Math. Phys **39**, 5166 (1998).
- [68] W. Lechner, P. Hauke, and P. Zoller, Science Advances **1**, e1500838 (2015).
- [69] A. Ambainis, International Journal of Quantum Information **1**, 507 (2003).
- [70] A. M. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D. A. Spielman, in *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing* (ACM, 2003) pp. 59–68.
- [71] J. Kempe, Contemporary Physics **44**, 307 (2003).
- [72] C. L. Baldwin, C. R. Laumann, A. Pal, and A. Scardicchio, arXiv:1509.08926 .
- [73] The annealing function $A(s)$ decreases linearly from 1 to 0, while $B(s)$ increases linearly from 0 to 1.