Defining and detecting quantum speedup

Troels F. Rønnow,[1] Zhihui Wang,[2] Joshua Job,[3] Sergio Boixo,[4] Sergei V. Isakov,[5]
David Wecker,[6] John M. Martinis,[7] Daniel A. Lidar,[8] and Matthias Troyer[*1]

[1] *Theoretische Physik, ETH Zurich, 8093 Zurich, Switzerland*
[2] *Department of Chemistry and Center for Quantum Information Science & Technology,*
University of Southern California, Los Angeles, California 90089, USA
[3] *Department of Physics and Center for Quantum Information Science & Technology,*
University of Southern California, Los Angeles, California 90089, USA
[4] *Google, 150 Main St, Venice Beach, CA, 90291*
[5] *Google, Brandschenkestrasse 110, 8002 Zurich, Switzerland*
[6] *Quantum Architectures and Computation Group, Microsoft Research, Redmond, WA 98052, USA*
[7] *Department of Physics, University of California, Santa Barbara, CA 93106-9530, USA*
[8] *Departments of Electrical Engineering, Chemistry and Physics,*
and Center for Quantum Information Science & Technology,
University of Southern California, Los Angeles, California 90089, USA

The development of small-scale digital and analog quantum devices raises the question of how to fairly assess and compare the computational power of classical and quantum devices, and of how to detect quantum speedup. Here we show how to define and measure quantum speedup in various scenarios, and how to avoid pitfalls that might mask or fake quantum speedup. We illustrate our discussion with data from a randomized benchmark test on a D-Wave Two device with up to 503 qubits. Comparing the performance of the device on random spin glass instances with limited precision to simulated classical and quantum annealers, we find no evidence of quantum speedup when the entire data set is considered, and obtain inconclusive results when comparing subsets of instances on an instance-by-instance basis. Our results for one particular benchmark do not rule out the possibility of speedup for other classes of problems and illustrate that quantum speedup is elusive and can depend on the question posed.

I. INTRODUCTION

The interest in quantum computing originates in the potential of a quantum computer to solve certain computational problems much faster than is possible classically. Examples are the factoring of integers [1] or the simulation of quantum systems [2]. Shor's algorithm can find the prime factors of an integer in a time that scales polynomially in the number of digits of the integer to be factored, while all known classical algorithms scale exponentially. The simulation of the time evolution of a quantum system on a classical computer also takes exponential resources, because the Hilbert space of an N particle system is exponentially large in N, while quantum hardware can simulate the same time evolution with polynomial complexity [3, 4].

In these examples the quantum algorithm is exponentially faster than the best available classical algorithm. This type of exponential quantum speedup substantially simplifies the discussion, since it renders the details of the classical or quantum hardware unimportant. According to the extended Church-Turing thesis all classical computers are equivalent up to polynomial factors [5]. Similarly, all proposed models of quantum computation are polynomially equivalent, so that a finding of exponential quantum speedup will be model-independent. In other cases, in particular on small devices, or when the quan-

tum speedup is polynomial, defining and detecting quantum speedup becomes more subtle. One such subtlety is how to properly define the hardness of a problem given prior knowledge about the answer [6].

Here we discuss how to define "quantum speedup" and show that this term may refer to different quantities depending on the goal of the study. In particular, we define what we call "limited quantum speedup"—essentially a speedup relative to a given, corresponding classical algorithm—and explain how such a speedup can be reliably detected. To illustrate these issues we compare the performance of a 503-qubit D-Wave Two (DW2) device to classical algorithms running on a standard CPU and analyze the evidence for quantum speedup on random spin glass problems. This example is particularly relevant since it is an open question whether quantum annealing [7] or the quantum adiabatic algorithm [8] can exhibit a quantum speedup for such problems. Random spin glass problems are an interesting benchmark problem, though not necessarily the most relevant for practical applications, such as machine learning. We also discuss issues that might mask or fake a quantum speedup when not considered carefully, such as comparing suboptimal algorithms or improperly accounting for the scaling of hardware resources.

II. DEFINING QUANTUM SPEEDUP

A. The classical to quantum scaling ratio

When the time to solution depends not only on the problem size N but also on the specific problem instance, then the purpose of the comparison becomes another factor in deciding how to measure performance. Specifically, when a device is used as a tool for solving problems, then the question of interest is to determine which device is better for the hardest problem, or for *almost all* possible problem instances. On the other hand, if we are interested in aspects of the underlying physics of a device then it might suffice to find *some* instances or a *subclass* of instances where a quantum device exhibits a speedup. These two questions will lead to different quantities of interest.

In all of these cases we denote the time used by a classical device to solve a problem of size N by $C(N)$ and the time used on the quantum device by $Q(N)$, defining quantum speedup as the ratio

$$S(N) = \frac{C(N)}{Q(N)}. \qquad (1)$$

Note that in both the quantum and classical case this definition includes a specific choice of algorithm and device.

The first question that arises is which classical algorithm to compare against, i.e., what is $C(N)$. This leads to different definitions of quantum speedup.

B. Five different types of quantum speedup

The optimal scenario is one of a *provable quantum speedup*, where there exists a proof that no classical algorithm can outperform a given quantum algorithm. Perhaps the best known example is Grover's search algorithm [9], which exhibits a provable quadratic speedup over the best possible classical algorithm [10], assuming an oracle.

A *strong quantum speedup* was defined in [11] by using the performance of the *best* classical algorithm for $C(N)$, whether such an algorithm is known or not. Unfortunately the performance of the best classical algorithm is unknown for many interesting problems. In the case of factoring, for example, all known classical algorithms have super-polynomial cost in the number of digits N of the number to be factored [12], while the cost of Shor's algorithm is polynomial in N. However, a proof of a classical exponential lower-bound for factorization is not known [13]. A less ambitious goal is therefore desirable, and thus one usually defines *quantum speedup* (without additional adjectives) by comparing to the *best available* classical algorithm instead of the best possible classical algorithm.

However, this notion of quantum speedup depends on there being a consensus about "best available", and this consensus may be time- and community-dependent [14]. In the absence of a consensus about what is the best classical algorithm, we define *potential (quantum) speedup* as a speedup compared to a specific classical algorithm or a set of classical algorithms. An example is the simulation of the time evolution of a quantum system, where the propagation of the wave function on a quantum computer would be exponentially faster than a direct integration of Schrödinger's equation on a classical computer. A potential quantum speedup can of course be trivially attained by deliberately choosing a poor classical algorithm (for example, factoring using classical instead of quantum period finding while ignoring known, better classical factoring algorithms), so that here too one must make a genuine attempt to compare against the best classical algorithms known, and any potential quantum speedup might be short-lived if a better classical algorithm is found.

Underlying all the above notions of quantum speedup is the availability of a fully coherent, universal quantum computer. A weaker scenario is one where the device is merely a putative or candidate quantum information processor, or where a quantum algorithm is designed to make use of quantum effects but it is not known whether these quantum effects provide an advantage over classical algorithms. To capture this scenario, which is of central interest to us in this work, we define *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. In the context of an analog quantum device this can be thought of as being the result of decohering the device. Since there is no unique way to decohere a quantum device, one may arrive at different corresponding classical algorithms. A natural example is quantum annealing implemented on a candidate physical quantum information processor *vs* either classical simulated annealing, classical spin dynamics, or simulated quantum annealing (as defined in Methods). In this comparison a limited quantum speedup would be a demonstration that quantum effects improve the annealing algorithm [15].

III. CLASSICAL AND QUANTUM ANNEALING OF A SPIN GLASS

As our primary example we will use the problem of finding the ground state of an Ising spin glass model described by a "problem Hamiltonian"

$$H_{\text{Ising}} = -\sum_{i \in \mathcal{V}} h_i \sigma_i^z - \sum_{(i,j) \in \mathcal{E}} J_{ij} \sigma_i^z \sigma_j^z , \qquad (2)$$

with N binary variables $\sigma_i^z = \pm 1$. The local fields $\{h_i\}$ and coupling $\{J_{ij}\}$ are fixed and define a *problem instance* of the Ising model. The spins occupy the vertices $\mathcal{V}$ of a graph $G = \{\mathcal{V}, \mathcal{E}\}$ with edge set $\mathcal{E}$. Solving this problem problem is NP-hard already for planar graphs [16], which means that no polynomial time algorithm to find these

ground states is known and the computational effort of all existing classical algorithms scales exponentially with problem size. NP-hardness refers only to the hardest problems, but the typical problem in our benchmarks, where the graph forms a two-dimensional (2D) lattice, is still hard since for zero local fields ($h_i = 0$) there exists a spin glass phase at zero temperature. While the critical temperature $T_c = 0$ for these 2D spin glasses makes the problem easier than 3D spin glasses with a nonzero $T_c > 0$ [17], solving the typical problem instance is nevertheless non-trivial and with all known algorithms a super-polynomial scaling is observed. While quantum mechanics is not expected to reduce this scaling to polynomial, a quantum algorithm might still scale better with problem size N than any classical algorithm.

We use simulated annealing (SA) [18], simulated quantum annealing (SQA) [19, 20], and a DW2 device to find the ground states of the Ising model above (see Methods for details). The D-Wave devices [21–24] are designed to be physical realizations of quantum annealing using superconducting flux qubits and programmable couplers. Tests on a 108-qubit D-Wave One (DW1) device [25] have shown that despite decoherence and coupling to a thermal bath, the device correlates well with SQA, which is consistent with it actually performing quantum annealing [26, 27]. It also correlates well with the predictions of a quantum master equation [28], which is consistent with it being governed by open system quantum dynamics. It is well understood that the D-Wave devices, just like any other quantum information processing device, must be error-corrected in order to overcome the effects of decoherence and control errors. While such error correction has already been demonstrated [29], our study focuses on the native performance of the device.

All annealing methods mentioned above are heuristic. They are not guaranteed to find the global optimum in a single annealing run, but only find it with a certain instance-dependent success probability $s \leq 1$. We determine the true ground state energy using an exact belief propagation algorithm [30]. We then perform at least 1000 repetitions of the annealing for each instance, count how often the ground state has been found by comparing to the exact result, and use this to estimate the success probability s for each problem instance.

The *total annealing time* is defined as the time to perform R annealing runs, where R is the number of repetitions needed to find the ground state at least once with probability p:

$$R = \left\lceil \frac{\log(1-p)}{\log(1-s)} \right\rceil \quad (3)$$

In order to reduce the effect of calibration errors on the DW2, it is advantageous to repeat the annealing runs for several different encodings ("gauges") of a problem instance. See Methods for details.
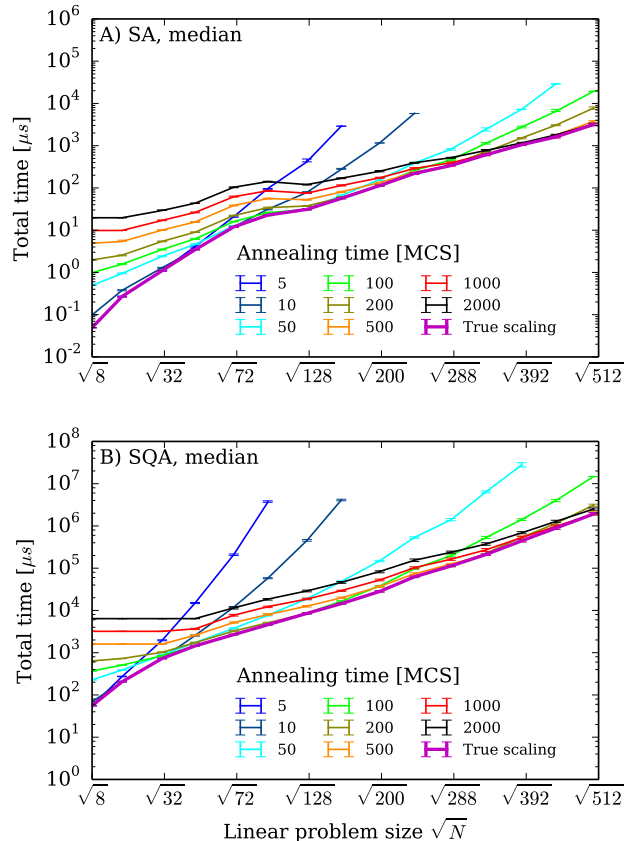

FIG. 1. **Scaling of the typical time to find a solution at constant annealing time.** Shown is the typical (median) time to find a ground state with 99% probability for spin glasses with ±1 couplings and no local field. A) for SA, B) for SQA. The envelope of the curves at constant t_a, shown in red, corresponds to the minimal time at a given problem size N and is relevant for discussion of the asymptotic scaling. Annealing times are given in units of Monte Carlo steps (MCS). One MCS corresponds to one update per spin. Note in particular that the slope for small N is much flatter at large annealing time (e.g., MCS = 4000) than that of the true scaling.

IV. CONSIDERATIONS WHEN COMPUTING QUANTUM SPEEDUP

Let us first consider the subtleties of estimating the asymptotic scaling from small problem sizes N, and inefficiencies at small problem sizes that can fake or mask a speedup. In the context of annealing methods the optimal choice of the annealing time turns out to be crucial for estimating asymptotic scaling.
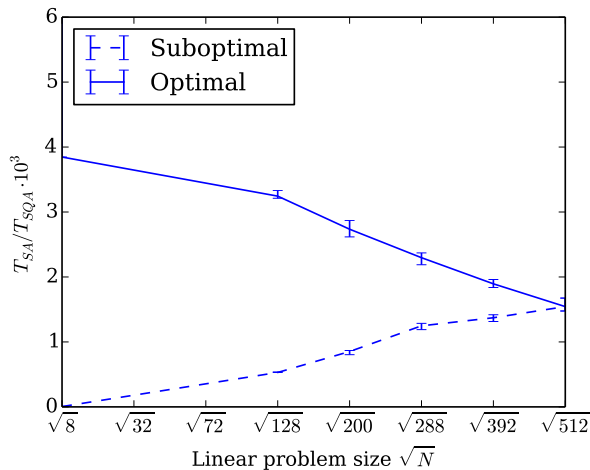
FIG. 2. **Pitfalls when detecting speedup.** Shown is the speedup of SQA over SA, defined as the ratio of median time to find a solution with 99% probability between SA and SQA. Two cases are presented: a) both SA and SQA run optimally (i.e., the ratio of the true scaling curves shown in Figure 1), and there is no asymptotic speedup (solid line). b) SQA is run suboptimally at small sizes by choosing a fixed large annealing time $t_a = 10000$ MCS (dashed line). The apparent speedup is, however, due to suboptimal performance on small sizes and not indicative of the true asymptotic behavior given by the solid line, which displays a slowdown of SQA compared to SA.

A. Asymptotic scaling: SA *vs* SQA

To illustrate these issues we consider the time to solution using SA and SQA run at different *fixed* annealing times t_a, independent of the problem size N. The problem instances we choose are random couplings of $J_{ij} = \pm 1$ on each of the edges in a perfect Chimera graph of $L \times L$ unit cells, containing $N = 8L^2$ spins (see Methods). We set the local fields $h_i = 0$. Figure 1 shows the scaling of the median total annealing time (over 1000 different random instances) for both SA and SQA to find a solution with probability $p = 0.99$. We observe that at constant t_a, as long as t_a is long enough to find the ground state almost every time, the scaling of the total effort is at first relatively flat. The total effort then rises more rapidly, once one reaches problem sizes for which the chosen annealing time is too short and the success probabilities are thus low, requiring many repetitions.

Figure 1 demonstrates that no conclusion can be drawn from annealing (simulated or in a device) about the asymptotic scaling at fixed annealing times. It is misleading to conclude about the asymptotic scaling from the initial slow increase at constant t_a, and instead the *optimal* annealing time t_a^{opt} needs to be found for each problem size N [25, 31]. The lower envelope of the scaling curves (indicated in red in Figure 1) corresponds to the total effort at an optimal size-dependent annealing time $t_a^{\text{opt}}(N)$ and can be used to infer the asymptotic scaling. In fact, the initial, relatively flat slope is due to subop-

timal performance at small problem sizes N, and should therefore not be interpreted as speedup. To illustrate this we show in Figure 2 (solid line) the true "speedup" ratio of the scaling of SA and SQA (actually a slowdown), and a misleading, fake speedup (dashed line) due to using a constant and excessively long annealing time t_a for SQA. Since the initial, slow increase of the total SQA effort at constant annealing time is a lower bound for the scaling of the true effort, the speedup *slope* obtained from this data—which depends inversely on the SQA effort—is an *upper bound*, as confirmed by Figure 2.

B. Resource usage and speedup from parallelism

A related issue is the scaling of hardware resources with problem size and parallelism in classical and quantum devices. To avoid mistaking a parallel speedup for a quantum speedup we need to scale hardware resources (computational gates and memory) in the same way for the devices we compare, and employ these resources optimally. These considerations are not universal but need to be carefully applied for each comparison of a quantum algorithm and device to a classical one.

For a problem of size N, the DW2 uses only N out of 512 qubits and $\mathcal{O}(N)$ couplers and classical logical control gates to solve a spin glass instance with N spin variables. We denote the time it needs to solve a problem by $T_{\text{DW}}(N)$. The classical simulated annealer (or simulated quantum annealer) running on a single classical CPU, on the other hand, uses fixed resources independent of problem size N, and we denote the time it requires to solve a problem by $T_{\text{SA}}(N)$. We consider here only the pure annealing times, as they are what is relevant for the asymptotic scaling rather than the readout or setup times, which scale subdominantly for large problems.

In order to avoid confusing quantum speedup with parallel speedup we thus consider as a classical counterpart to the DW2 a (hypothetical) special purpose parallel classical simulated annealing device, with the same hardware scaling as the DW2. Simulated annealing (and simulated quantum annealing) is perfectly parallelizable for the bipartite Chimera graphs realized by the DW2. The reason is that one Monte Carlo step (consisting of one attempted update per spin) can be performed in constant time, since all spins in each of the two sublattices can be updated simultaneously. The time to solve a problem on this equivalent classical device, denoted by $T_{\text{C}}(N)$, is thus related to the time $T_{\text{SA}}(N)$ taken by a simulated annealer using a fixed-size classical CPU by

$$T_{\text{C}}(N) \propto \frac{1}{N} T_{\text{SA}}(N), \qquad (4)$$

since the latter needs time $O(N)$ for one Monte Carlo step, while the former performs it in constant time.

The *quantum* part of speedup is then estimated by comparing the times required by two devices with the

same hardware scaling, giving

$$S(N) = \frac{T_C(N)}{T_{DW}(N)} \propto \frac{T_{SA}(N)}{T_{DW}(N)} \frac{1}{N}. \qquad (5)$$

The factor $1/N$ in the speedup calculation thus discounts for the intrinsic *parallel speedup* of the analog device whose hardware resources scale as N. See Methods for an alternative derivation that leads to the same results (up to subleading corrections) by using a fixed size device efficiently.

V. PERFORMANCE OF D-WAVE TWO VERSUS SA AND SQA

A. Comparing devices

If the goal is to compare the performance of devices as optimizers, then one is interested in solving *almost all* problem instances. In this case we should run the devices in such a way that all but a small fraction of the problems can be solved. This will lead to a speedup defined as the ratio of the quantiles ("R of Q") of the time to solution, with an emphasis on the high quantiles, which we discuss in Sec. V C. A complementary question is to ask whether a device exhibits better performance than another for *some* problems. To answer this question we compare the time to solution individually for each problem instance. We then consider the quantiles of the ratio ("Q of R") of the time to solution, and discuss this approach in Sec. V D.

A complementary distinction is that between *wall-clock time*, denoting the full time to solution, and the *pure annealing time*. Wall-clock time is the total time to find a solution and is the relevant quantity when one is interested in the performance of a device for applications and has been used in Ref. [32]. It includes the setup, cooling, annealing and readout times on the DW2, and the setup, annealing and measurement time for the classical annealing codes. The pure annealing time is simply Rt_a, where R is the number of repetitions and t_a the time used for a single annealing run. It is the relevant quantity when one is interested in the intrinsic physics of the annealing processes and in scaling to larger problem sizes on future devices. We discuss both wall-clock and pure annealing times below.

B. Problem instances

The family of problem instances we use for our benchmarking tests employ couplings J_{ij} on all edges of $N = 8LL'$-vertex subgraphs of the Chimera graph of the DW2, comprising $L \times L'$ unit cells, with $L, L' \in \{1, \ldots, 8\}$. We set the fields $h_i = 0$ since nonzero values of the fields h_i destroy the spin glass phase that exists at zero field, thus making the instances easier [33]. We choose the values

of the couplings J_{ij} from $2r$ discrete values $\{n/r\}$, with $n \in \{-r, -r-1, \ldots, -1, 1, \ldots, r-1, r\}$, and call r the "range". Thus when the range $r = 1$ we only pick values $J_{ij} = \pm 1$. This choice is the least susceptible to calibration errors of the device, but the large degeneracy of the ground states in these cases makes finding a ground state somewhat easier. At the opposite end we consider $r = 7$, which is the upper limit given the four bits of accuracy of the couplings in the DW2. These problem instancess are harder since there are fewer degenerate minima, but they also suffer more from calibration errors in the device. In the Supplementary Material we present additional results for $r = 3$.

C. Performance as an optimizer: comparing the scaling of hard problem instances

1. Pure annealing time

We start our analysis by focusing on pure annealing times and show in Figure 3 the scaling of the time to find the ground state at least once with probability $p = 0.99$ for various quantiles, from the easiest instances (1%) to the hardest (99%), for two different ranges. Since we do not *a priori* know the hardness of a given problem instance we have to assume the worst case and perform a sufficient number of repetitions R to be able to solve even the hardest problem instances. Hence the scaling for the selected high quantile will apply to *all* problem instances we run on the optimizer.

In all three cases (SA, SQA, DW2) we observe, for sufficiently large N, that the total time to solution scales with $\exp(c\sqrt{N})$, as reported previously for SA and SQA [25]. The origin of the $\sqrt{N}$ exponent is well understood for exact solvers as reflecting the treewidth of the Chimera graph (see Methods and Ref. [34]), and a similar scaling is observed here for the heuristic algorithms. While the SA and SQA codes were run at an optimized annealing time for each problem size N, the DW2 has a minimal annealing time of $t_a = 20\mu s$, which is longer than the optimal time for all problem sizes (see Methods). Therefore the observed slope of the DW2 data should only be taken as a lower bound for the asymptotic scaling. Even so, we observe similar scaling for the classical codes and on DW2.

2. The ratio of quantiles

With algorithms such as SA or quantum annealing, where the time to solution depends on the problem instance, it is often not possible (and usually irrelevant) to experimentally find the hardest problem instance. It is preferable to decide instead for which fraction of problem instances one wishes to find the ground state, which then defines the relevant quantile. If we target $q\%$ of the instances then we should consider the qth percentile
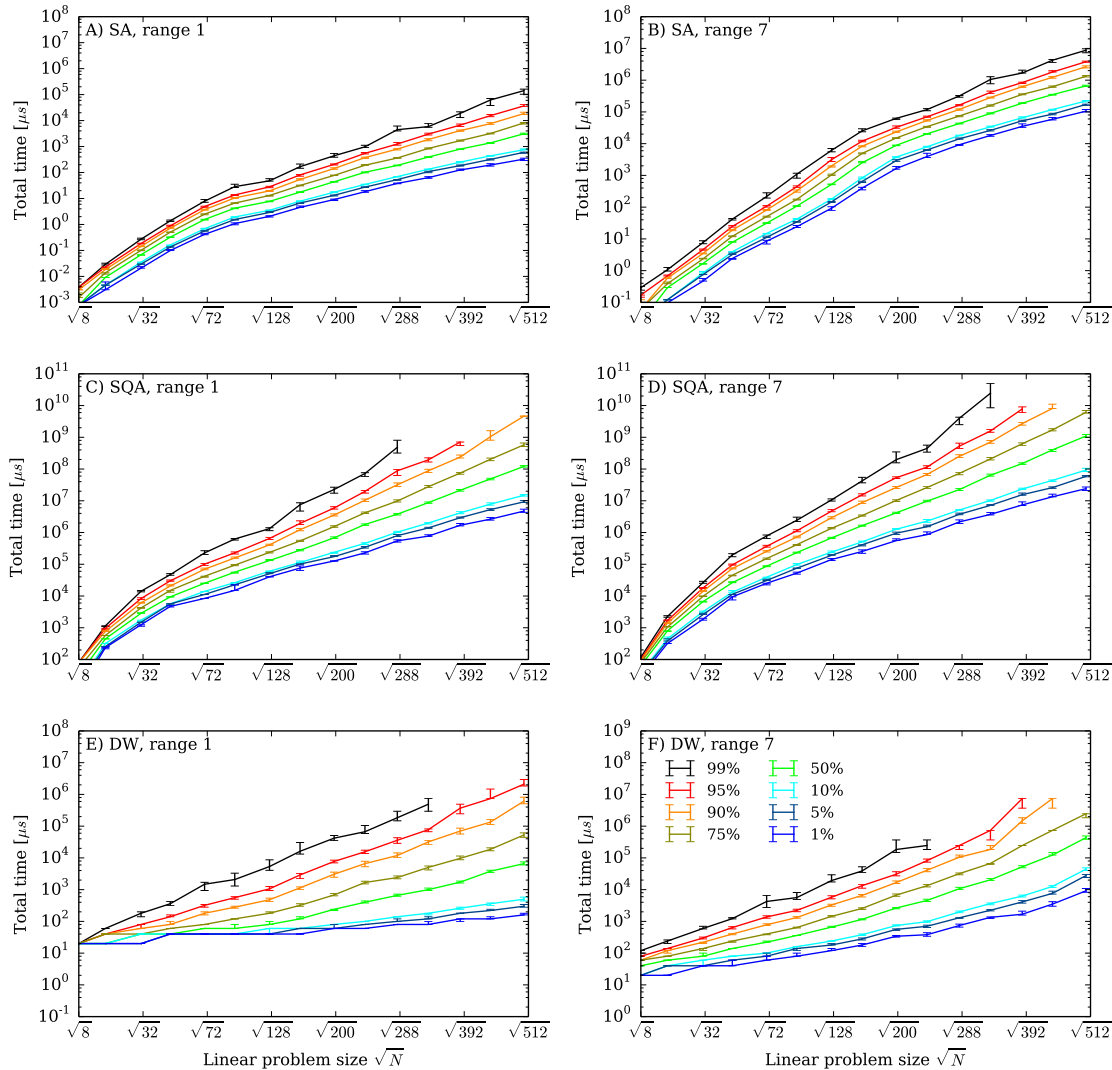
FIG. 3. **Scaling of time to solution for the ranges $r = 1$ (panels A, C and E) and $r = 7$ (panels B, D and F).** Shown is the scaling of the time to find the ground state at least once with a probability $p = 0.99$ for various quantiles of hardness, for A,B) simulated annealing (SA), C,D) simulated quantum annealing (SQA) and E,F) the DW2. The SA and SQA data is obtained by running the simulations at an optimized annealing time for each problem size. The DW2 annealing time of $20\mu s$ is the shortest possible. Note the different vertical axis scales, and that both the DW2 and SQA have trouble solving the hardest instances for the large problem sizes, as indicated by the terminating lines for the highest quantiles. More than the maximum number of of repetitions (10000 for SQA, at least 32000 for DW2) of the annealing we performed would be needed to find the ground state in those cases.

in the scaling plots shown in Figure 3. The appropriate speedup quantity is then the *ratio of these quantiles*. Denoting a quantile q of a random variable X by $[X]_q$ we can define this as

$$S_q^{\mathrm{RofQ}}(N) = \frac{[T_\mathrm{C}(N)]_q}{[T_\mathrm{DW}(N)]_q} \propto \frac{[T_\mathrm{SA}(N)]_q}{[T_\mathrm{DW}(N)]_q} \frac{1}{N}. \qquad (6)$$

Plotting this quantity for the DW2 *vs* SA in Figure 4 we find no evidence for a limited quantum speedup in

the interesting regime of large N. That is, while for all quantiles, and for both ranges (with the exception of the 50th quantile and $r = 1$), the initial slope is positive, when N becomes large enough we observe a turnaround and eventually a negative slope, showing that SA outperforms the DW2.

Taking into account that (as discussed in Sec. IV A) due to the fixed suboptimal annealing times the speedup defined in Eq. (6) is an *upper bound*, we conclude that
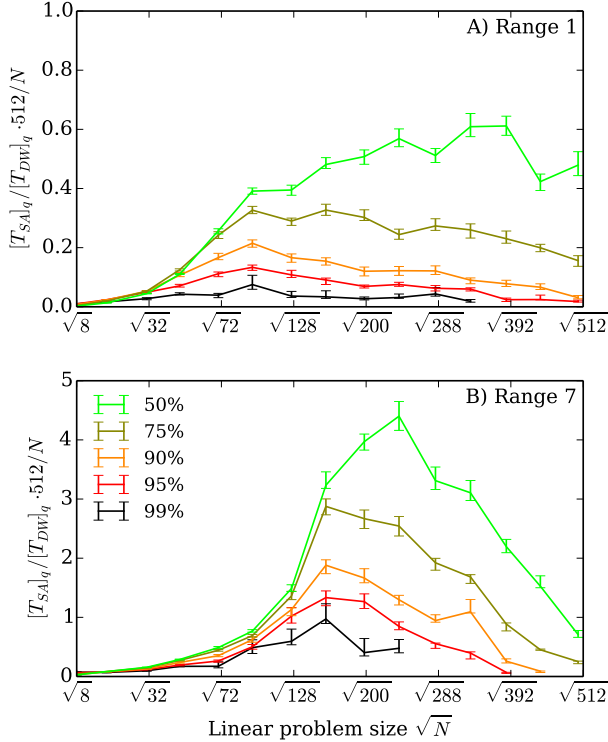
FIG. 4. **Speedup for ratio of quantiles for the DW2 compared to SA.** A) For instances with range $r = 1$. B) For instances with range $r = 7$. Shown are curves from the median (50th quantile) to the 99th quantile. 16 gauges were used. In these plots we multiplied Eq. (6) by 512 so that the speedup value at $N = 512$ directly compares one DW2 processor against one classical CPU.

3. Wall-clock time

While not as interesting from a complexity theory point of view, it is instructive to also compare wall-clock times for the above benchmarks, as we do in Figure 5. We observe that the DW2 performs similarly to SA run on a single classical CPU, for sufficiently large problem sizes and at high range values. Note that the large constant programming overhead of the DW2 masks the exponential increase of time to solution that is obvious in the plots of pure annealing time.
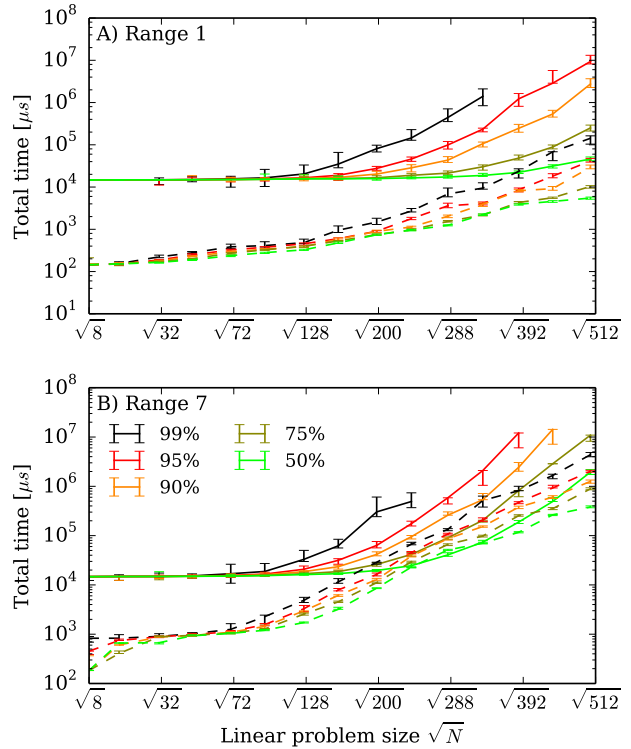

FIG. 5. **Comparing wall-clock times** A comparison of the wall-clock time to find the solution with probability $p = 0.99$ for SA running on a single CPU (dashed lines) compared to the DW2 (solid lines) using 16 gauges. A) for range $r = 1$, B) for range $r = 7$. Shown are curves from the median (50th quantile) to the 99th quantile. The large constant programming overhead of the DW2 masks the exponential increase of time to solution that is obvious in the plots of pure annealing time. Results for a single gauge are shown in the Supplementary Material.

D. Instance-by-instance comparison

1. Total time to solution

We now focus on the question of whether the DW2 exhibits a limited quantum speedup for some fraction of the instances of our benchmark set. To this end we perform individual comparisons for each instance and show in Figure 6A-B the ratios of time to solution between the DW2 and SA, considering only the pure annealing time. We find a wide scatter, which is not surprising since we previously found that DW1 performs like a simulated quantum annealer, but correlates less well with a simulated classical annealer [25]. We find that while the DW2 is sometimes up to $10\times$ faster in pure annealing time, there are many cases where it is $\geq 100\times$ slower.

Considering the wall-clock times, the advantage of the DW2 seen in Figure 6A-B for some instances tends to

the DW2 does not exhibit a speedup over SA for this particular benchmark.
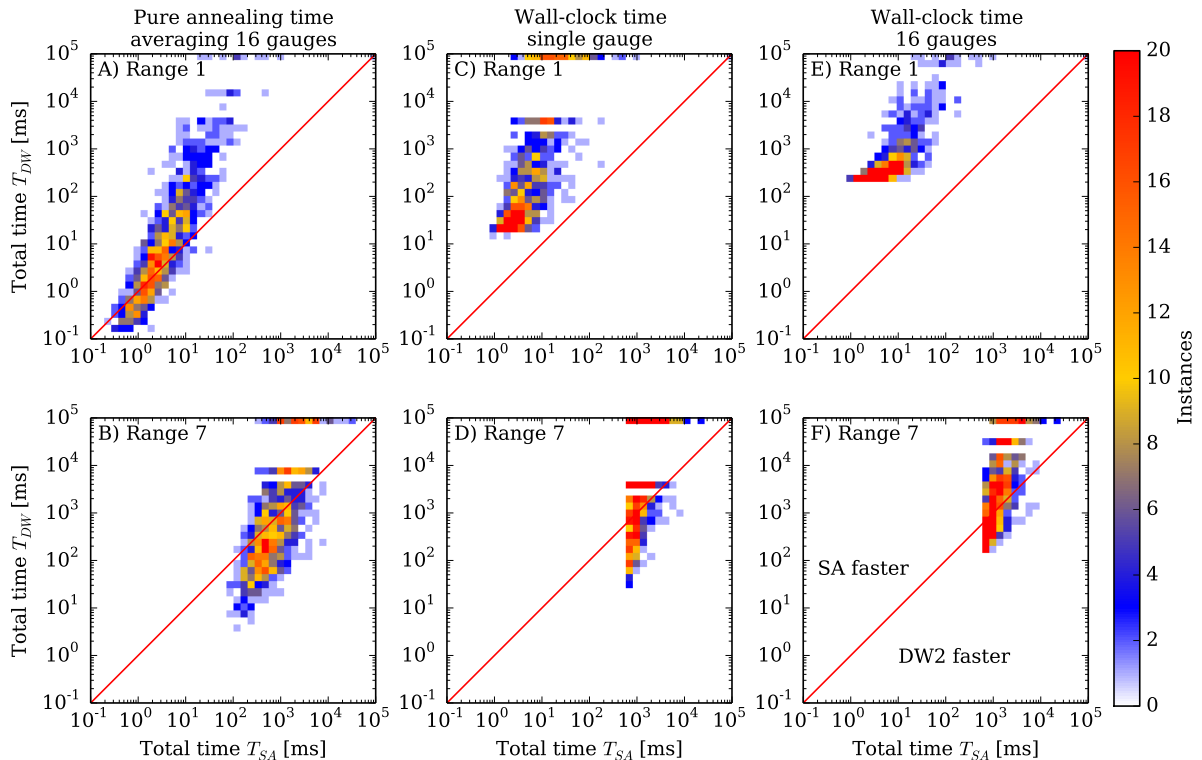
FIG. 6. **Instance-by-instance comparison of annealing times and wall-clock times.** Shown is a scatter plot of the pure annealing time for the DW2 compared to a simulated classical annealer (SA) using an average over 16 gauges on the DW2. A) DW2 compared to SA for $r = 1$, B) DW2 compared to SA for $r = 7$. The color scale indicates the number of instances in each square. Instances below the diagonal red line are faster on the DW2, those above are faster classically. Instances for which the DW2 did not find the solution with 10000 repetitions per gauge are shown at the top of the frame (no such instances were found for SA). Panels C) and D) show wall-clock times using a single gauge on the DW2. Panels E) and F) show the wall-clock time for DW2 using 16 gauges. $N = 503$ in all cases.

disappear, since it is penalized by the need for programming the device with multiple different gauge choices (see Methods). Figure 6C-D shows that for one gauge choice there are some instances, for $r = 7$, where the DW2 is faster, but many instances where it never finds a solution. Using 16 gauges the DW2 finds the solution in most cases, but is always slower than the classical annealer on a classical CPU for $r = 1$, as can be seen in Figure 6E-F. For $r = 7$ the DW2 is sometimes faster than a single classical CPU. Overall, the performance of the DW2 is better for $r = 7$ than for $r = 1$, and comparable to SA only when just the pure annealing time is considered. The difference to the results of Ref. [32] is due to the use of optimized classical codes using a full CPU in our comparison, as opposed to the use of generic optimization codes using only a single CPU core in Ref. [32].

2. Quantiles of ratio

Comparisons of the absolute time to solution are of limited importance compared to the real question of scaling,

which can give insight into the behavior of future devices that can solve larger problems. In Section V C we did not find evidence for a limited quantum speedup when considering all instances. Now we consider instead whether there is such a speedup for a subset of problem instances. To this end we study the scaling of the ratios of the time to solution for individual instances, and display in Figure 7 the scaling of various *quantiles of the ratio*

$$S_q^{\mathrm{QofR}}(N) = \left[\frac{T_{\mathrm{C}}(N)}{T_{\mathrm{DW}}(N)} \right]_q \propto \left[\frac{T_{\mathrm{SA}}(N)}{T_{\mathrm{DW}}(N)} \frac{1}{N} \right]_q. \quad (7)$$

For $r = 7$ all the quantiles bend down for sufficiently large N, so that there is no evidence of a limited quantum speedup. Yet, now there seems to be an indication of such a speedup compared to SA in the high quantiles for $r = 1$. However, for the reasons discussed in Sec. IV A, one must be careful not to overinterpret this as solid evidence for a speedup since the instances contributing here are not run at the optimal annealing time. Moreover, as discussed in the Supplementary Material, we find no evidence of a limited quantum speedup for $r = 3$. Thus, while perhaps encouraging from the per-
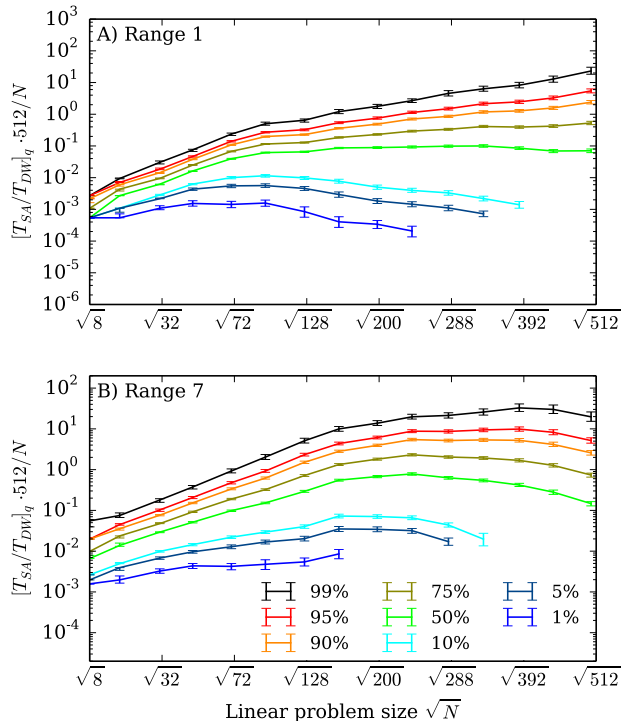
FIG. 7. **Speedup for quantiles of the ratio** of the DW2 compared to SA, for A) $r = 1$, B) $r = 7$. No asymptotic speedup is visible for any of the quantiles at $r = 7$, while some evidence of a limited quantum speedup (relative to SA) is seen for quantiles higher than the median at $r = 1$. As in Figure 4 we multiplied Eq. (7) by 512 so that the speedup value at $N = 512$ directly compares one DW2 processor against one classical CPU.

spective of a search for a (limited) quantum speedup, more work is needed to establish that the $r = 1$ result persists for those instances for which one can be sure that the annealing time is optimal.

E. Arguments for and against a speedup on the DW2

Let us consider in some more detail the speedup results discussed above. We have argued that the apparent limited quantum speedup seen in the $r = 1$ results of Figure 7 must be treated with care due to the suboptimal annealing time. It might then be tempting to argue that, strictly speaking, the comparison with suboptimal-time instances cannot be used for claiming a slowdown either, i.e., that we simply cannot infer how the DW2 will behave for optimal-time instances by basing the analysis on suboptimal times only.

However, let us make the assumption that, along with the total time, the optimal annealing time $t_a^{\mathrm{opt}}(N)$ also grows with problem size N. This assumption is sup-

ported by the SA and SQA data shown in Figure 1, and is plausible as long as the growing annealing time does not become counterproductive due to coupling to the thermal bath [35]. By definition, $T_{\mathrm{DW}}(N, t_a^{\mathrm{opt}}(N)) \leq T_{\mathrm{DW}}(N, t_a)$, where we have added the explicit dependence on the annealing time, and t_a is a fixed annealing time. Thus

$$
\begin{aligned}
S(N) &= \frac{T_{\mathrm{C}}(N)}{T_{\mathrm{DW}}(N, t_a)} \frac{1}{N} \\
&\leq \frac{T_{\mathrm{C}}(N)}{T_{\mathrm{DW}}(N, t_a^{\mathrm{opt}}(N))} \frac{1}{N} = S^{\mathrm{opt}}(N).
\end{aligned}
\tag{8}
$$

Under our assumption, $t_a^{\mathrm{opt}}(N) < t_a$ for small N, but for sufficiently large N the optimal annealing time grows so that $t_a^{\mathrm{opt}}(N) \geq t_a$. Thus there must be a problem size N^* at which $t_a^{\mathrm{opt}}(N^*) = t_a$, and hence at this special problem size we also have $S(N^*) = S^{\mathrm{opt}}(N^*)$. However, as mentioned in Section V C 1, the minimal annealing time of $20\mu s$ is longer than the optimal time for all problem sizes (see Supplementary Material), i.e., $N^* > 503$ in our case. Therefore, if $S(N)$ is a decreasing function of N for sufficiently large N, as we indeed observe in all our "R of Q" results (recall Figure 4), then since $S^{\mathrm{opt}}(N) \geq S(N)$ and $S(N^*) = S^{\mathrm{opt}}(N^*)$, it follows that $S^{\mathrm{opt}}(N)$ too must be a decreasing function for a range of N values, at least until N^*. This shows that the slowdown conclusion holds also for the case of optimal annealing times.

For the instance-by-instance comparison ("Q of R"), no such conclusion can be drawn for the subset of instances (at $r = 1$) corresponding to the high quantiles where $S_q^{\mathrm{QofR}}(N)$ is an *increasing* function of N. This limited quantum speedup may or may not persist for larger problem sizes or if optimal annealing times are used.

VI. DISCUSSION

In this work we have discussed challenges in properly defining and assessing quantum speedup, and used comparisons between a DW2 and simulated classical and quantum annealing to illustrate these challenges. *Strong* or *provable* quantum speedup, implying speedup of a quantum algorithm or device over *any* classical algorithm, is an elusive goal in most cases and one thus usually defines *quantum speedup* as a speedup compared to the best available classical algorithm. We have introduced the notion of *limited quantum speedup*, referring to a more restricted comparison to "corresponding" classical algorithms solving the same task, such as a quantum annealer compared to a classical annealing algorithm.

Quantum speedup is most easily defined and detected in the case of an exponential speedup, where the details of the quantum or classical hardware do not matter since they only contribute subdominant polynomial factors. In the case of an unknown or a polynomial quantum speedup one must be careful to fairly compare the classical and quantum devices, and, in particular, to scale hardware resources in the same manner. Otherwise *par-*

allel speedup might be mistaken for (or hide) quantum speedup.

An experimental determination of quantum speedup suffers from the problem that all measurements are limited to finite problem sizes N, while we are most interested in the asymptotic behavior for large N. To arrive at a reliable extrapolation it is advantageous to focus the scaling analysis on the part of the execution time that becomes dominant for large problem sizes N, which in our example is the pure annealing time, and not the total wall-clock time. For each problem size we furthermore need to ensure that neither the quantum device nor the classical algorithm are run suboptimally, since this might hide or fake quantum speedup.

If the time to solution depends not only on the problem size N but also on the specific problem instance, then one needs to carefully choose the relevant quantity to benchmark. We argued that in order to judge the performance over many possible inputs of a randomized benchmark test, one needs to study the high quantiles, and define speedup by considering the ratio of the quantiles of time to solution. If, on the other hand, one is interested in finding out whether there is a speedup for some subset of problem instances, then one can instead perform an instance-by-instance comparison by focusing on the quantiles of the ratio of time to solution.

We note that it is not yet known whether a quantum annealer or even a perfectly coherent adiabatic quantum optimizer can exhibit (limited) quantum speedup at all [36], although there are promising indications from simulation [20] and experiments on spin glass materials [37]. Experimental tests will thus be important. We chose to focus here on the benchmark problem of random zero-field Ising problems parametrized by the range of couplings. We did not find evidence of limited quantum speedup for the DW2 relative to simulated annealing in our particular benchmark set when we considered the ratio of quantiles of time to solution, which is the relevant quantity for the performance of a device as an optimizer. We note that random spin glass problems, while an interesting and important physics problem, may not be the most relevant benchmark for practical applications, for which other benchmarks may have to be studied.

When we focus on subsets of problem instances in an instance-by-instance comparison, we observe a possibility for a limited quantum speedup for a fraction of the instances [38]. However, since the DW2 runs at a suboptimal annealing time for most of the corresponding problem instances, the observed speedup may be an artifact of attempting to solve the smaller problem sizes using an excessively long annealing time. This difficulty can only be overcome by fixing the issue of suboptimal annealing times, e.g., by finding problem classes for which the annealing time is demonstrably already optimal.

There are several candidate explanations for the absence of a clear quantum speedup in our tests. Perhaps quantum annealing simply does not provide any advantages over simulated (quantum) annealing or other classical algorithms for the problem class we have studied [17]; or, perhaps, the noisy implementation in the DW2 cannot realize quantum speedup and is thus not better than classical devices. Alternatively, a speedup might be masked by calibration errors, improvements might arise from error correction [29], or other problem classes might exhibit a speedup [39]. Future studies will probe these alternatives and aim to determine whether one can find a class of problem instances for which an unambiguous speedup over classical hardware can be observed.

METHODS

Simulated annealing. Simulated annealing [18] performs a Monte Carlo simulation on the model of Eq. (2), starting from a random initial state at high temperature. During the course of the simulation the temperature is lowered towards zero. At the end of the annealing schedule, at low temperature, the spin configuration of the system ends up in in a local minimum. By repeating the simulation many times one may hope to find the global minimum. More specifically, SA is performed by sequentially iterating through all spins and proposing to flip them based on a Metropolis algorithm using the Boltzmann weight of the configuration at finite temperature. During the annealing schedule we linearly increase the inverse temperature over time from an initial value of $\beta = 0.1$ to a final value of $\beta = 3r$.

For the case of ± 1 couplings ($r = 1$), and for $r = 3$ we use a highly optimized multispin-coded algorithm based on Refs. [40, 41]. This algorithm performs updates on 64 copies in parallel, updating all at once. For the $r = 7$ simulations we use a code optimized for bipartite lattices [42]. Implementations of the simulated annealing codes are available in Ref. [42]. We used the code **an_ms_r1_nf** for $r = 1$, the code **an_ms_r3_nf** for $r = 3$ and the code **an_ss_ge_nf_bp** for $r = 7$.

Quantum annealing. To perform quantum annealing one maps the Ising variables σ_i^z to Pauli z-matrices and adds a transverse magnetic field in the x-direction to induce quantum fluctuations, thus obtaining the time-dependent quantum Hamiltonian

$$H(t) = -A(t) \sum_i \sigma_i^x + B(t) H_{\text{Ising}} , \quad t \in [0, t_a] . \quad (9)$$

The annealing schedule starts at time $t = 0$ with just the transverse field term (i.e., $B(0) = 0$) and $A(0) \gg k_B T$, where T is the temperature, which is kept constant. The system is then in a simple quantum state with (to an excellent approximation) all spins aligned in the x direction, corresponding to a uniform superposition over all 2^N computational basis states (products of eigenstates of the σ_i^z). During the annealing process the problem Hamiltonian magnitude $B(t)$ is increased and the transverse field $A(t)$ is decreased, ending with $A(t_a) = 0$, and couplings much larger than the temperature: $B(t_a) \max(\max_{ij} |J_{ij}|, \max_i |h_i|) \gg k_B T$. At this point the system will again be trapped in a local minimum, and by repeating the process one may hope to find the global minimum. Quantum annealing can be viewed as a finite-temperature variant of the adiabatic quantum algorithm [8].

Simulated quantum annealing. Simulated quantum annealing (SQA) [19, 20] is an annealing algorithm based on discrete-time path-integral quantum Monte Carlo simulations of the transverse field Ising model, following the above annealing schedule at a constant low temperature, but using Monte Carlo dynamics instead of the open system evolution of a quantum system. This amounts to sampling the world line configurations of the quantum Hamiltonian (9) while slowly changing the couplings. The algorithm we used is similar to that of Ref. [43], but uses cluster updates along the imaginary time direction, typically with 64 time slices. Our annealing schedule is linear, as shown in Figure 9B): the Ising couplings are ramped up linearly while the transverse field is ramped down linearly over time.

The layout of the D-Wave Two Vesuvius chip. The Chimera graph of the DW2 used in our tests is shown in Figure 8. Each unit cell is a balanced $K_{4,4}$ bipartite graph. In the ideal Chimera graph (of 512 qubits) the degree of each vertex is 6. For the scaling analysis we considered $L \times L'$ rectangular sub-lattices of the Chimera graph, and restricted our simulations and tests on the DW2 to the subset of functional qubits within these subgraphs. More generally the $N = 2cL^2$-vertex Chimera graph comprises an $L \times L$ grid of $K_{c,c}$ unit cells, and the (so-called TRIAD) construction of Ref. [34] can be used to embed the complete L-vertex graph K_L, where $L = 4c$. The treewidth of the $N = 2cL^2$-vertex Chimera graph comprising an $L \times L$ grid of $K_{c,c}$ unit cells is $cL + 1 \sim \mathcal{O}(\sqrt{N})$ [34]. The treewidth of the 512-vertex Chimera graph shown in Figure 8 is 33. Dynamic programming can always find the true ground state of the corresponding Ising model in a time that is exponential in the treewidth, i.e., that scales as $\exp(c\sqrt{N})$.

Annealing schedule of the D-Wave Two device. Nominally, the DW2 performs annealing by implementing the time-dependent Hamiltonian $H(t) = -A(t)\sum_i \sigma_i^x + B(t)H_{\text{Ising}}$, where $t \in [0, t_a]$. However, in reality the transverse field varies somewhat and the actual Hamiltonian realized is described more accurately as

$$H(t) = -\sum_i A_i(t)\sigma_i^x + B(t)H_{\text{Ising}} \qquad (10)$$

The annealing schedules $A_i(t)$ and $B(t)$ used in the device are shown in figure 9A). We used the minimal annealing time of $t_a = 20\mu s$ provided by the device since this always gave us the shortest total time to solution (see below). The source of this minimal annealing time is engineering restrictions in the DW2. There are four annealing lines, and their synchronization becomes harder for faster annealers. The filtering of the input control lines introduces some additional distortion in the annealing control.

The DW2 is programmed by providing the sets of couplings J_{ij} and local longitudinal fields h_i (which together define a problem instance by specifying H_{Ising}), the number of repetitions R of the annealing to be performed, the annealing time t_a, and a number of other parameters which were included in our wall-clock results and which are described below.

Gauge averaging on the D-Wave device. Calibration inaccuracies cause the couplings J_{ij} and h_i that are realized in the DW2 to be slightly different from the intended and
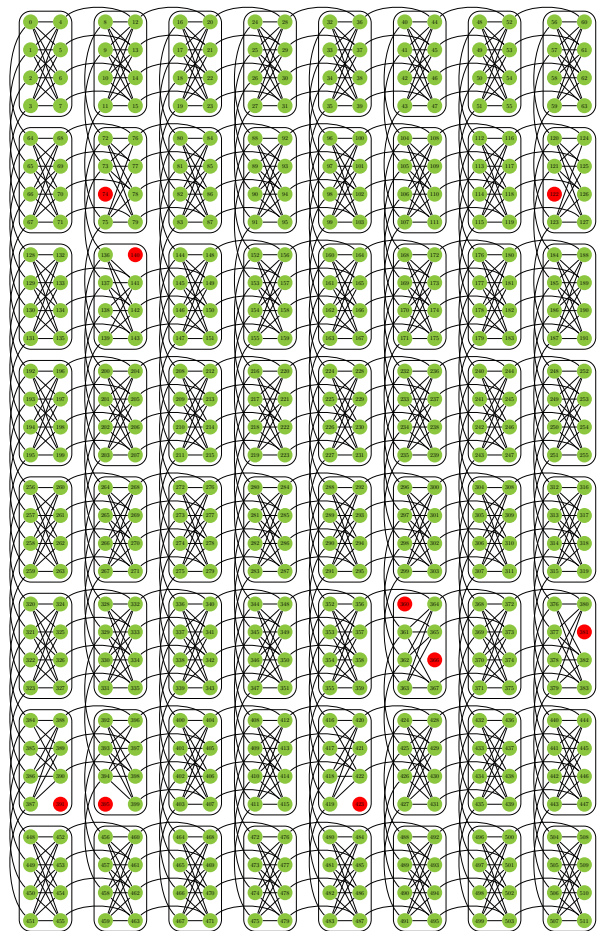

FIG. 8. **Qubits and couplers in the D-Wave Two device.** The DW2 "Vesuvius" chip consists of an 8×8 two-dimensional square lattice of eight-qubit unit cells, with open boundary conditions. The qubits are each denoted by circles, connected by programmable inductive couplers as shown by the lines between the qubits. Of the 512 qubits of the device located at the University of Southern California used in this work, the 503 qubits marked in green and the couplers connecting them are functional.

programmed values ($\sim 5\%$ variation). These calibration errors can sometimes lead to the ground states of the model realized in the device being different from the perfect model. To overcome these problems it is advantageous to perform annealing on the device with multiple encodings of a problem instance into the couplers of the device [25]. To realize these different encodings we use a gauge freedom in realizing the Ising spin glass: for each qubit we can freely define which of the two qubits states corresponds to $\sigma_i = +1$ and $\sigma_i = -1$. More formally this corresponds to a gauge transformation that changes spins $\sigma_i^z \to a_i\sigma_i^z$, with $a_i = \pm 1$ and the couplings as $J_{ij} \to a_i a_j J_{ij}$ and $h_i \to a_i h_i$. The simulations are invariant under such a gauge transformation, but (due to calibration errors which break the gauge symmetry) the results returned by the DW2 are not.

If the success probability of one annealing run is denoted by s, then the probability of failing to find the ground state after R independent repetitions (annealing runs) each having
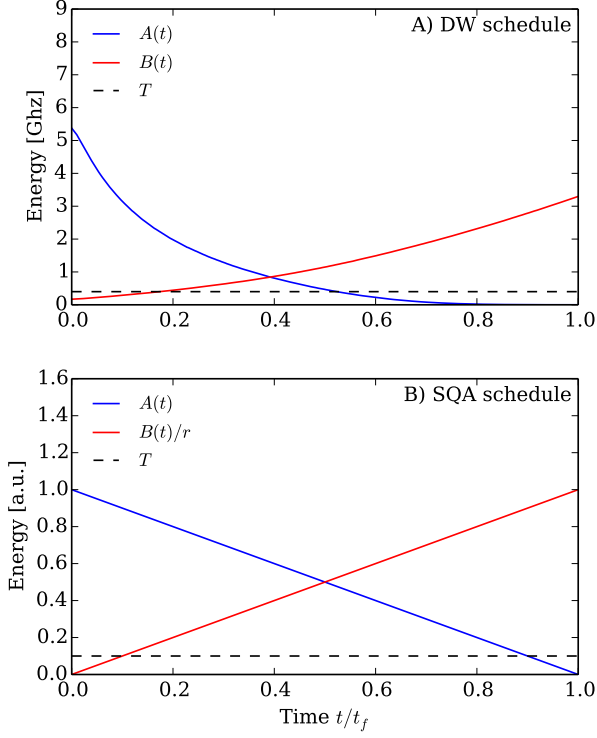
N	t_p [ms]	t_r [μs]
8	14.7 ± 0.3	51.0 ± 0.2
16	14.8 ± 0.3	53.0 ± 0.2
31	14.8 ± 0.3	57.9 ± 0.2
47	14.9 ± 0.4	60.6 ± 0.2
70	15.0 ± 0.4	64.5 ± 0.2
94	15.2 ± 0.3	68.3 ± 0.2
126	15.6 ± 0.2	73.1 ± 0.2
158	15.5 ± 0.2	78.0 ± 0.2
198	15.5 ± 0.2	80.8 ± 0.2
238	15.7 ± 0.2	83.5 ± 0.2
284	15.8 ± 0.2	83.6 ± 0.1
332	16.0 ± 0.3	87.1 ± 0.2
385	16.6 ± 1.0	87.1 ± 0.6
439	16.6 ± 0.1	90.4 ± 0.1
503	16.6 ± 0.2	90.5 ± 0.1

TABLE I. **Wallclock times on the DW2**. Listed are measured programming times t_p and annealing plus readout times t_r (for a pure annealing time of $20\mu s$) on the DW2 for various problem sizes.

$N \le 238$	$N = 284, 332$	$N = 385, 439$	$N = 503$
1000	2000	5000	10000

TABLE II. **Repetitions of annealing runs used on the DW2**. This table summarizes the total number of repetitions used to estimate the success probabilities on the DW2 for various system sizes.

FIG. 9. **Annealing schedules.** A) The amplitude of the median transverse field $A(t)$ (decreasing, blue) and the longitudinal couplings $B(t)$ (increasing, red) as a function of time. The device temperature of $T = 18$mK is indicated by the black horizontal dashed line. B) The linear annealing schedule used in simulated quantum annealing.

success probability s is $(1-s)^R$, and the total success probability of finding the ground state at least once in R repetitions is

$$P = 1 - (1-s)^R. \tag{11}$$

Thus the number of repetitions needed to find the ground state at least once with probability p is found by solving $p = 1 - (1-s)^R$, i.e., Eq. (3).

Following [25], after splitting these repetitions into R/G repetitions for each of G gauge choices with success probabilities s_g, the total success probability becomes

$$P^{(G)} = 1 - \prod_{g=1}^{G} (1-s_g)^{R/G}. \tag{12}$$

If we denote by $\overline{s}$ the geometric mean of the success probabilities of the individual gauges

$$\overline{s} = 1 - \prod_{g=1}^{G} (1-s_g)^{1/G}, \tag{13}$$

then Eq. (12) can be written in the same form as Eq. (11):

$$P^{(G)} = 1 - (1-\overline{s})^R. \tag{14}$$

We thus use the geometric mean $\overline{s}$ in our scaling analysis.

Wall-clock and annealing times. We show results mainly for pure annealing times, but also for wall clock times. The pure annealing time for R repetitions is straightforwardly defined as

$$t_{\text{anneal}} = Rt_a. \tag{15}$$

Wall clock times include the time for programming, cooling, annealing, readout and communication. We have performed tests on the DW2 with varying numbers of repetitions R and performed a linear regression analysis to fit the total wall clock time for each problem size to the form $t_p(N) + Rt_r(N)$, where $t_p(N)$ is the total preprocessing time and $t_r(N)$ is the total run time per repetition for an N-spin problem. The values of t_p and t_r are summarized in Table I. With these numbers we obtain the total wall clock time for R annealing runs split over G gauges (with R/G annealing runs each) as

$$t_{\text{total}}(N) = Gt_p(N) + Rt_r(N). \tag{16}$$

To calculate pure annealing times for the simulated annealer we determine the total effort in units of of Monte Carlo updates (attempted spin flips), and then convert to time by dividing by the number of updates that the codes can perform per second [42]. Our classical reference CPU is an 8-core Intel Xeon E5-2670 CPU, which was introduced around the same time as the DW2.

To obtain wall-clock times we measure the actual time needed to perform a simulation on the same Intel Xeon
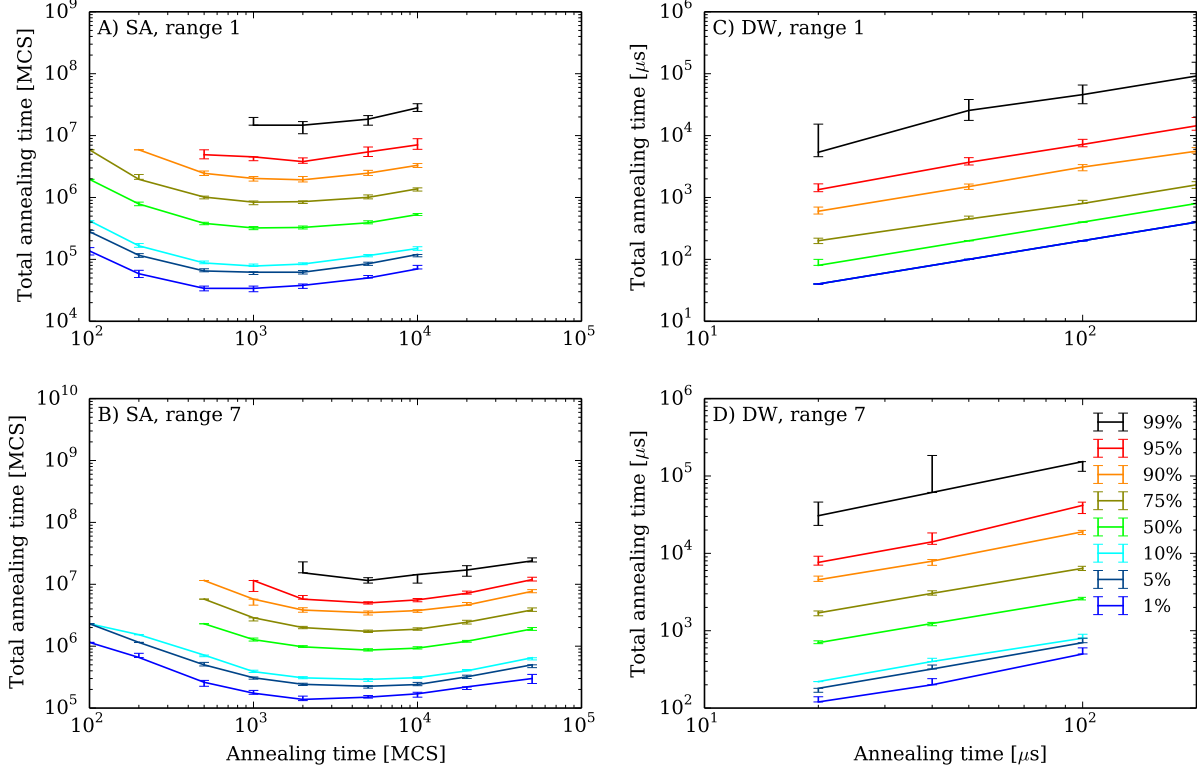
FIG. 10. **Optimal annealing times for the simulated annealer and for the D-Wave device.** Shown is the total effort $R(t_a)t_a$ as a function of annealing time t_a for various quantiles of problems with $r = 1$ and $r = 7$ (see Supplementary Information for $r = 3$). A) and B) SA, where the minimum of the total effort determines the optimal annealing time $t_a^{\rm opt}$. C) and D) DW2, where we find a monotonically increasing total effort, meaning that the optimal time $t_a^{\rm opt}$ is always shorter than the minimal annealing time of 20μs.

E5-2670 CPU. Since the multi-spin codes perform at 64 repetitions in parallel, we always make at least 1024 repetitions when running 16 threads on 8 cores. This causes the initially flatter scaling in wall-clock times as compared to pure annealing times. The measured initialization time includes all preparations needed for the algorithm to run, and the spin flip rate was computed for the 99% quantile for 503 qubits. For smaller system sizes or lower quantiles, the spin flip rate is lower since the problems are not hard enough to benefit from parallelization over several cores.

Optimal annealing times. As discussed in the main text we need to determine the optimal annealing time $t_a^{\rm opt}$ for every problem size N in order to make meaningful extrapolations of the time to find a solution. To determine $t_a^{\rm opt}$ we perform annealing runs at different annealing times t_a, determine the success probabilities $s(t_a)$ of 1000 instances, and from them the required number of repetitions $R(t_a)$ to find the ground state with a probability of 99%. The total effort $R(t_a)t_a$ diverges for $t_a \to 0$ and $t_a \to \infty$ and has a minimum at an optimal annealing time $t_a^{\rm opt}$. The reason is that for short t_a the success probability goes to zero, which leads to a diverging total effort, while for large t_a the time also grows since one always needs to perform at least one annealing run and the total effort is thus bounded from below by t_a.

In Figure 10 (left) we plot various quantiles of the total effort $R(t_a)t_a$ for the simulated annealer as a function of

t_a to determine the optimal annealing time $t_a^{\rm opt}$. For the DW2 we find, as shown in Figure 10 (right) that the minimal annealing time of $20\mu s$ is always longer than the optimal time and we thus always use the device in a suboptimal mode. As a consequence the scaling of time to solution is underestimated, as explained in detail in the main text.

Alternative consideration of parallel versus quantum speedup. In the consideration of how to disentangle parallel and quantum speedup it may seem more natural to assume fixed computational resources of a given device. We will show that this leads to the same scaling as Eq. (5). We might be tempted to define the speedup in this case as $S(N) = T_{\rm SA}(N)/T_{\rm DW}(N)$. However, in this manner only a fraction $N/512$ of the qubits are used while the classical code uses the available CPU fully, independently of problem size. This suboptimal use of the DW2 may again be incorrectly interpreted as speedup. The same issue would appear when comparing a classical analog annealer against a classical simulated annealer.

As in the discussion of optimal annealing times above, we need to ensure an optimal implementation to correctly assess speedup. For the DW2 (or a similarly constructed classical analog annealer) this means that one should always attempt to make use of the entire device: we should perform as many annealing runs in parallel as possible. Let us denote the machine size by M (e.g., $M = 512$ in the DW2 case). With this

we define a new, optimized, annealing time

$$T_{\text{DW}}^{\text{opt}}(N) = T_{\text{DW}}(N) \frac{1}{\lfloor M/N \rfloor}, \qquad (17)$$

and the correct speedup in our case is then

$$S(N) = \frac{T_{\text{SA}}(N)}{T_{\text{DW}}^{\text{opt}}(N)} = \frac{T_{\text{SA}}(N)}{T_{\text{DW}}(N)} \left\lfloor \frac{M}{N} \right\rfloor. \qquad (18)$$

Omitting the floor function ($\lfloor \ \rfloor$), which only gives subdominant corrections in the limit $M \to \infty$ we recover Eq. (5).

The conclusion that the speedup function includes a factor proportional to $1/N$ is validated from yet another perspective, that focuses on the annealing time. Instead of embedding $C \equiv \lfloor M/N \rfloor$ different instances in parallel, we can embed C replicas of a given instance. Each replica r (where $r \in \{1, \ldots, C\}$) results in a guess $E_{r,i}$ of the ground state energy for the ith run, and we can take $E_i = \min_r E_{r,i}$ as the proposed solution for that run. If the replicas are independent and each has equal probability s of finding the ground state, then using C replicas the probability that at least one will find the ground state is $s' = 1 - (1-s)^C$, which is also the probability that E_i is the ground state energy for the ith

run. Repeating the argument leading to Eq. (3), the number of repetitions required to find the ground state at least once with probability p is then:

$$R' = \left\lceil \frac{\log(1-p)}{\log(1-s')} \right\rceil = \left\lceil \frac{\log(1-p)}{C \log(1-s)} \right\rceil = \frac{R}{C}. \qquad (19)$$

Focusing on the pure annealing time we have $T_{\text{DW}}^{\text{opt}}(N) = t_a R'$ and $T_{\text{DW}}(N) = t_a R$, which yields Eq. (17).

[1] Shor, P. W. Algorithms for quantum computation: discrete logarithms and factoring. *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on* 124–134 (20-22 Nov 1994). URL http://dx.doi.org/10.1109/SFCS.1994.365700.

[2] Feynman, R. Simulating physics with computers. *International Journal of Theoretical Physics* **21**, 467–488 (1982). URL http://dx.doi.org/10.1007/BF02650179.

[3] Lloyd, S. Universal quantum simulators. *Science* **273**, 1073–1078 (1996). URL http://www.sciencemag.org/content/273/5278/1073.abstract.

[4] Berry, D. W., Childs, A. M., Cleve, R., Kothari, R. & Somma, R. D. Exponential improvement in precision for simulating sparse hamiltonians. *arXiv:1312.1414* (2013). URL http://arXiv.org/abs/1312.1414.

[5] Parberry, I. Parallel speedup of sequential machines: a defense of parallel computation thesis. *SIGACT News* **18**, 54–67 (1986).

[6] Smolin, J. A., Smith, G. & Vargo, A. Oversimplifying quantum factoring. *Nature* **499**, 163–165 (2013). URL http://dx.doi.org/10.1038/nature12290.

[7] Kadowaki, T. & Nishimori, H. Quantum annealing in the transverse Ising model. *Phys. Rev. E* **58**, 5355–5363 (1998). URL http://link.aps.org/doi/10.1103/PhysRevE.58.5355.

[8] Farhi, E. *et al.* A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science* **292**, 472–475 (2001). URL http://dx.doi.org/10.1126/science.1057726.

[9] Grover, L. K. Quantum mechanics helps in searching for a needle in a haystack. *Physical Review Letters* **79**, 325–328 (1997). URL http://link.aps.org/doi/10.1103/PhysRevLett.79.325.

[10] Bennett, C., Bernstein, E., Brassard, G. & Vazirani, U. Strengths and weaknesses of quantum computing. *SIAM Journal on Computing* **26**, 1510–1523 (1997). URL http://dx.doi.org/10.1137/S0097539796300933.

[11] Papageorgiou, A. & Traub, J. F. Measures of quantum computing speedup. *Phys. Rev. A* **88**, 022316 (2013). URL http://link.aps.org/doi/10.1103/PhysRevA.88.022316.

[12] Pomerance, C. A tale of two sieves. *Notices of the Amer. Math. Soc.* **43**, 1473–1485 (1996). URL http://www.ams.org/notices/199612/pomerance.pdf.

[13] Such a proof seems unlikely to be found any time soon since it would imply that factoring is not in the complexity class P (polynomial) and thus P and NP (non-deterministic polynomial) are distinct, solving the longstanding P versus NP question.

[14] For example, it may be the case, though it seems unlikely, that a classified polynomial-time factoring algorithm is available to parts of the intelligence community.

[15] We compare quantum annealing only to classical simulated annealing and simulated quantum annealing in this study. Another example of a limited quantum speedup would be Shor's factoring algorithm running on a fully coherent quantum computer *vs* a classical computer where the period finding using a quantum circuit has been replaced by a classical period finding algorithm.

[16] Barahona, F. On the computational complexity of Ising spin glass models. *Journal of Physics A: Mathematical and General* **15**, 3241 (1982). URL http://stacks.iop.org/0305-4470/15/i=10/a=028.

[17] Helmut G. Katzgraber, R. S. A., Firas Hamze. Glassy Chimeras are blind to quantum speedup: Designing better benchmarks for quantum annealing machines (2014). URL http://arxiv.org/abs/1401.1546. 1401.1546.

[18] Kirkpatrick, S., Gelatt, C. D. & Vecchi, M. P. Optimization by simulated annealing. *Science* **220**, 671–680 (1983). URL http://dx.doi.org/10.1126/science.

220.4598.671.

[19] Martoňák, R., Santoro, G. E. & Tosatti, E. Quantum annealing by the path-integral Monte Carlo method: The two-dimensional random Ising model. *Phys. Rev. B* **66**, 094203 (2002). URL http://link.aps.org/doi/10.1103/PhysRevB.66.094203.

[20] Santoro, G. E., Martoňák, R., Tosatti, E. & Car, R. Theory of quantum annealing of an Ising spin glass. *Science* **295**, 2427–2430 (2002). URL http://dx.doi.org/10.1126/science.1068774.

[21] Harris, R. *et al.* Experimental investigation of an eight-qubit unit cell in a superconducting optimization processor. *Phys. Rev. B* **82**, 024511 (2010). URL http://link.aps.org/doi/10.1103/PhysRevB.82.024511.

[22] Johnson, M. W. *et al.* A scalable control system for a superconducting adiabatic quantum optimization processor. *Superconductor Science and Technology* **23**, 065004 (2010). URL http://stacks.iop.org/0953-2048/23/i=6/a=065004.

[23] Berkley, A. J. *et al.* A scalable readout system for a superconducting adiabatic quantum optimization system. *Superconductor Science and Technology* **23**, 105014 (2010). URL http://stacks.iop.org/0953-2048/23/i=10/a=105014.

[24] Johnson, M. W. *et al.* Quantum annealing with manufactured spins. *Nature* **473**, 194–198 (2011). URL http://dx.doi.org/10.1038/nature10012.

[25] Boixo, S. *et al.* Quantum annealing with more than one hundred qubits (2013). URL http://arxiv.org/abs/1304.4595. 1304.4595.

[26] Smolin, J. A. & Smith, G. Classical signature of quantum annealing (2013). URL http://arxiv.org/abs/1305.4904. arXiv:1305.4904.

[27] Wang, L. *et al.* Comment on: 'Classical signature of quantum annealing' (2013). URL http://arxiv.org/abs/1305.5837. arXiv:1305.5837.

[28] Boixo, S., Albash, T., Spedalieri, F. M., Chancellor, N. & Lidar, D. A. Experimental signature of programmable quantum annealing (2012). URL http://arxiv.org/abs/1212.1739. 1212.1739.

[29] Pudenz, K. L., Albash, T. & Lidar, D. A. Error corrected quantum annealing with hundreds of qubits. *arXiv:1307.8190* (2013). URL http://arXiv.org/abs/1307.8190.

[30] Dechter, R. Bucket elimination: A unifying framework for reasoning. *Artificial Intelligence* **113**, 41–85 (1999).

[31] While one can in principle look for an optimal annealing time for each individual problem instance, we instead determine an averaged optimal annealing time $t_a^{\mathrm{opt}}(N)$ for each problem size N by annealing many instances at various annealing times t_a, and then use these for all future problems of that size.

[32] McGeoch, C. C. & Wang, C. Experimental evaluation of an adiabatic quantum system for combinatorial optimization. In *Proceedings of the 2013 ACM Conference on Computing Frontiers* (2013).

[33] Young, A. P. & Katzgraber, H. G. Absence of an Almeida-Thouless Line in Three-Dimensional Spin Glasses. *Phys. Rev. Lett.* **93**, 207203 (2004). URL http://link.aps.org/doi/10.1103/PhysRevLett.93.207203.

[34] Choi, V. Minor-embedding in adiabatic quantum computation: II. Minor-universal graph design. *Quantum Information Processing* **10**, 343–353 (2011). URL http://dx.doi.org/10.1007/s11128-010-0200-3.

[35] Sarandy, M. S. & Lidar, D. A. Adiabatic quantum computation in open systems. *Physical Review Letters* **95**, 250503– (2005). URL http://link.aps.org/doi/10.1103/PhysRevLett.95.250503.

[36] An important exception is Ref. [44], but this example does not (yet) admit implementation on a physical device.

[37] Brooke, J., Bitko, D., Rosenbaum, T. & Aeppli, G. Quantum annealing of a disordered magnet. *Science* **284**, 779–781 (1999). URL http://www.sciencemag.org/content/284/5415/779.abstract. http://www.sciencemag.org/content/284/5415/779.full.pdf.

[38] The fact that some problems are solved faster on classical hardware and some on the DW2 raises the possibility of a hybrid approach that benefits by solving a problem instance using both, and then selecting the best solution found.

[39] For example, it is possible that a randomized benchmark with a tunable "hardness knob", such as the clause density in the MAX-2SAT problem [45], will allow for a more fine-tuned exploration of the performance potential of a quantum annealer than a purely randomized benchmark such as we have used here. It was also proposed that embedding 3D spin-glass problems into the Chimera graph, with a finite critical temperature, might yield problem classes that are better suited as benchmarks for quantum annealing [17].

[40] Bhanot, G., Duke, D. & Salvador, R. A fast algorithm for the Cyber 205 to simulate the 3d Ising model. *J. Stat. Phys.* **44**, 985–1002 (1986). URL link.springer.com/content/pdf/10.1007/BF01011918.pdf.

[41] Heuer, H.-O. A fast vectorized Fortran 77 program for the Monte Carlo simulation of the three-dimensional Ising system. *Comput. Phys. Commun.* **59**, 387 – 398 (1990). URL http://www.sciencedirect.com/science/article/pii/0010465590901865.

[42] Isakov, S., Zintchenko, I., Rønnow, T. & Troyer, M. Optimized simulated annealing for Ising spin glasses (2014). URL http://arxiv.org/abs/1401.1084. 1401.1084.

[43] Martoňák, R., Santoro, G. E. & Tosatti, E. Quantum annealing by the path-integral Monte Carlo method: The two-dimensional random Ising model. *Phys. Rev. B* **66**, 094203 (2002). URL http://link.aps.org/doi/10.1103/PhysRevB.66.094203.

[44] Somma, R. D., Nagaj, D. & Kieferová, M. Quantum speedup by quantum annealing. *Physical Review Letters* **109**, 050501– (2012). URL http://link.aps.org/doi/10.1103/PhysRevLett.109.050501.

[45] Santra, S., Quiroz, G., Steeg, G. V. & Lidar, D. MAX 2-SAT with up to 108 qubits. *arXiv:1307.3931* (2013). URL http://arXiv.org/abs/1307.3931.