

# Scaling advantages of all-to-all connectivity in physical annealers: the Coherent Ising Machine vs. D-Wave 2000Q

Ryan Hamerly<sup>1,2,†</sup>, Takahiro Inagaki<sup>3,†</sup>, Peter L. McMahon<sup>4,2,†</sup>, Davide Venturelli<sup>5,6</sup>, Alireza Marandi<sup>4</sup>, Tatsuhiro Onodera<sup>4</sup>, Edwin Ng<sup>4</sup>, Carsten Langrock<sup>4</sup>, Kensuke Inaba<sup>3</sup>, Toshimori Honjo<sup>3</sup>, Koji Enbutsu<sup>7</sup>, Takeshi Umeki<sup>7</sup>, Ryoichi Kasahara<sup>7</sup>, Shoko Utsunomiya<sup>2</sup>, Satoshi Kako<sup>2</sup>, Ken-ichi Kawarabayashi<sup>2</sup>, Robert L. Byer<sup>4</sup>, Martin M. Fejer<sup>4</sup>, Hideo Mabuchi<sup>4</sup>, Eleanor Rieffel<sup>5</sup>, Hiroki Takesue<sup>3</sup>, Yoshihisa Yamamoto<sup>4,8</sup>

May 4, 2018

---

<sup>1</sup> *Research Laboratory of Electronics, Massachusetts Institute of Technology, 50 Vassar Street, Cambridge, MA 02139, USA*

<sup>2</sup> *National Institute of Informatics, Hitotsubashi 2-1-2, Chiyoda-ku, Tokyo 101-8403, Japan*

<sup>3</sup> *NTT Basic Research Laboratories, NTT Corporation, 3-1 Morinosato Wakamiya, Atsugi, Kanagawa 243-0198, Japan*

<sup>4</sup> *E. L. Ginzton Laboratory, Stanford University, Stanford, CA 94305, USA*

<sup>5</sup> *NASA Ames Research Center Quantum Artificial Intelligence Laboratory (QuAIL), Mail Stop 269-1, Moffett Field, California 94035, USA*

<sup>6</sup> *USRA Research Institute for Advanced Computer Science (RIACS), 615 National Ave, Mountain View, California 94035, USA*

<sup>7</sup> *NTT Device Technology Laboratories, NTT Corporation, 3-1 Morinosato Wakamiya, Atsugi, Kanagawa 243-0198, Japan*

<sup>8</sup> *ImPACT Program, Japan Science and Technology Agency, Gobancho 7, Chiyoda-ku, Tokyo 102-0076, Japan*

<sup>†</sup> These authors contributed equally to this work.

## Abstract

Physical annealing systems provide a heuristic approach to solve NP-hard Ising optimization problems. It is believed that the connectivity between spins in such annealers significantly impacts the machine's computational effectiveness. In this paper we study the performance of two types of annealing machines that have very different connectivity – a commercially available quantum annealer built by D-Wave Systems, which has sparse connectivity, and coherent Ising machines based on optical parametric oscillator networks, which have all-to-all connectivity. We demonstrate an exponential ( $e^{-O(N^2)}$ ) penalty in performance for the D-Wave quantum annealer relative to coherent Ising machines when solving Ising problems on dense graphs, which is attributable to the differences in internal connectivity between the machines. This leads to a several-orders-of-magnitude time-to-solution difference between coherent Ising machines and the D-Wave system for problems with over 50 vertices. Our results provide strong experimental support to efforts to increase the connectivity of physical annealers.

## Introduction

Optimization problems are ubiquitous in science, engineering and business. Many important problems (especially combinatorial problems such as scheduling, resource allocation, route planning or community detection) belong to the NP-hard complexity class, and even for typical instances require a computation time that scales exponentially with the problem size [1]. Canonical examples such as Karp's 21 NP-complete problems [2] have attracted much attention from researchers seeking to devise new optimization methods, because by definition any NP-complete problem can be reduced to any other problem in NP with only polynomial overhead. Many approximation algorithms and heuristics (e.g., relaxations to semidefinite programs [3], simulated annealing [4], and breakout local search [5]) have been developed to search for good-quality approximate solutions as well as ground states for sufficiently small problem sizes. However, for many NP-hard optimization problems, even moderately sized problem instances can be time-consuming to solve exactly or

even approximately. Hence, there is strong motivation to find alternative approaches that can consistently beat state-of-the-art algorithms.

Despite decades of Moore’s Law scaling, large NP-hard problems remain very costly even on modern microprocessors. Thus, there is a growing interest in special-purpose machines that implement a solver directly by mapping the optimization to the underlying physical dynamics. Examples include digital CMOS annealers [6, 7], as well as analog devices such as nano-magnet arrays [8], electronic oscillators [9, 10] and laser networks [11]. Quantum adiabatic computation [12] and quantum annealing [13, 14, 15, 16] are also prominent examples, and may offer the possibility of quantum speedup [16, 17, 18] for certain NP-hard problems. However, all the non-photonic analog optimization systems realized to date suffer from limited connectivity, so that actual problems must in general first be *embedded* [19, 20] into the solver architecture native graph before they can be solved. This may add an upfront computational cost [19, 21, 22] of finding the embedding (unless previously known) and, of most relevance in this study, in general results in the use of multiple physical pseudo-spins to encode each logical spin variable, which can lead to an additional degradation of time-to-solution.

In this paper, we perform the first direct comparison between the D-Wave 2000Q quantum annealer and the Coherent Ising Machine (CIM). As we will see later, a crucial distinction between these systems is their intrinsic connectivity, which has a profound influence on their performance. Both systems are designed to solve the classical Ising problem, that is, to minimize the classical Hamiltonian:

$$H = \frac{1}{2} \sum_{ij} J_{ij} \sigma_i \sigma_j + \sum_i h_i \sigma_i \quad (1)$$

where  $\sigma_i = \pm 1$  are the Ising spins,  $J_{ij}$  are the entries of the spin-spin coupling matrix, and  $h_i$  the Zeeman

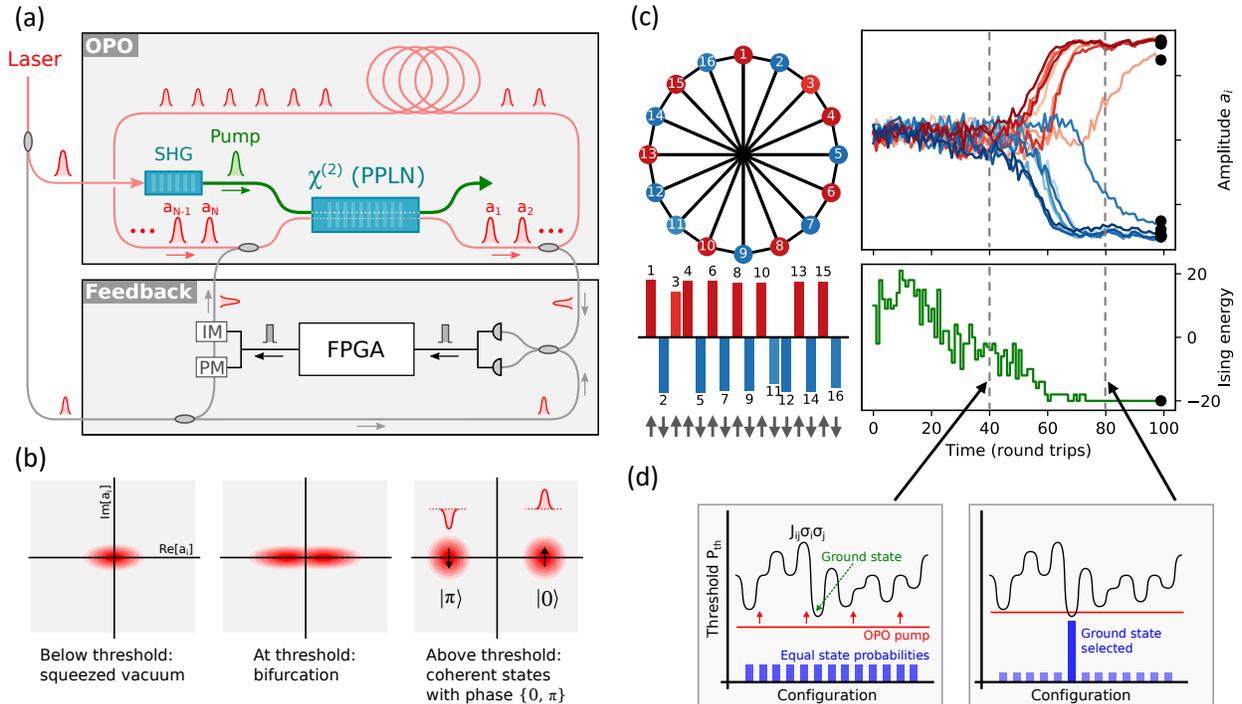


Figure 1: (a) Coherent Ising Machine design consisting of time-multiplexed OPO and measurement-feedback apparatus. See Refs. [23, 24] for details. (b) OPO state during transition from below-threshold squeezed state to (bistable) above-threshold coherent state. (c) Solution of antiferromagnetic Ising problem on the Möbius ladder with the CIM, giving measured OPO amplitudes  $a_i$  and Ising energy  $H$  as a function of time in round trips. (d) Illustration of search-from-below principle of CIM operation.

(bias) terms. The Ising problem is NP-hard for non-planar couplings [25] and is one of the most widely studied problems in this complexity class. We focus on two canonical NP-hard Ising problems: unweighted MAX-CUT [2] and ground-state computation of the Sherrington-Kirkpatrick spin glass model [26].

In the CIM, the spin network is represented by a network of degenerate optical parametric oscillators (OPOs). Each OPO is a nonlinear oscillator that converts pump light to its half-harmonic [27]; it can oscillate in two identical phase states, which encode the value of the Ising spin [28]. As Fig. 1(a) shows, time multiplexing offers a straightforward way to generate many identical OPOs in a single cavity [29]. A pulsed laser with repetition time  $T$  is used to pump an optical cavity with round-trip time  $N \times T$ . Parametric amplification is provided by the  $\chi^{(2)}$  crystal; since this is an instantaneous nonlinearity, the circulating pulses in the cavity are identical and non-interacting. The approach is scalable using high repetition-rate lasers and long fiber cavities: OPO gain has been reported for up to  $N = 10^6$  pulses, and stable operation achieved for  $N = 50,000$  [30]. Each circulating pulse represents an independent OPO with a single degree of freedom  $a_i$ . Classically,  $a_i$  is a complex variable, which maps to the annihilation operator  $\hat{a}_i$  in quantum mechanics [31]. A measurement-feedback apparatus is used to apply coupling between the pulses [23, 24]. In each round trip, a small fraction of the light is extracted from the cavity and homodyned against a reference pulse (the OPO pump is created from second harmonic generation (SHG) of the reference laser, so there is good matching between the reference and the OPO signal light, which is at half the frequency of the pump). The homodyne result, in essence a measurement of  $a_i$ , is fed into an electronic circuit (consisting of an ADC, an FPGA, and a DAC) that, for each pulse, computes a feedback signal that is proportional to the matrix-vector product  $\sum_j J_{ij} a_j$ . This signal is converted back to light using an optical modulator and a reference pulse, and re-injected into the cavity. The measurement-feedback CIM has intrinsic all-to-all connectivity through its exploitation of memory in the electronic circuit (although the same effect can be obtained with optical delay-line memories in all-optical CIMs [28, 29]).

The OPO is a dissipative quantum system with a pitchfork bifurcation well adapted for modeling Ising spins: as the pump power is increased (Fig. 1(b)), the OPO state transitions from a below-threshold squeezed vacuum state [32, 33, 34] to an above-threshold coherent state [35]. Because degenerate parametric amplification is phase-sensitive, only two phase states are stable above threshold; thus the OPO functions as a classical “spin” with states  $\{|0\rangle, |\pi\rangle\}$  that can be mapped to the Ising states  $\sigma_i = \{+1, -1\}$ . The optimization process happens in the near-threshold regime where the dynamics are determined by a competition between the network loss and Ising coupling (which seek to minimize the product  $\sum_{ij} J_{ij} a_i a_j$ ), and nonlinear parametric gain (which seeks to enforce the constraints  $a_i \in \mathbb{R}, |a_i| = \text{const}$ ).

As an example, consider the Ising problem on the  $N = 16$  Möbius ladder graph with anti-ferromagnetic couplings [36]. Fig. 1(c) shows a typical run of the CIM, resulting in a solution that minimizes the Ising energy (data from Ref. [23]). The most obvious interpretation of the process is spontaneous symmetry breaking of a pitchfork bifurcation: prepared in a squeezed vacuum state and driven by shot noise, the OPO state bifurcates, during which its amplitudes  $a_i$  grow either in positive or negative value, and subsequently the system settles into the Ising ground state (or a low-lying excited state) [23, 37] (this is related to the Gaussian-state model in Ref. [38]). Another view derives from ground-state “search from below” (Fig. 1(d)). Here the Ising energy is visualized as a complicated landscape of potential oscillation thresholds, each with its own spin configuration. If the OPO pump is far below the minimum threshold, all spin configurations will be excited with near-equal probability, but once the ground-state threshold is exceeded, its probability will grow exponentially at the expense of other configurations [29]. This ground-state selection process corresponds to the  $40 \leq t \leq 60$  region in Fig. 1(c).

The D-Wave 2000Q (DW2Q) quantum annealer used in this work is installed at NASA Ames Research Center in Mountain View, California. The DW2Q has 2,048 qubits, but its “Chimera” coupling graph (i.e., the graph whose edges define the non-zero  $J_{ij}$  terms in Eq. (1)) is very sparse. Since most Ising problems are not defined on subgraphs of the Chimera, *minor embedding* is used to find a Chimera subgraph on which the corresponding Ising model has a ground state that corresponds to the classical ground state of the Ising model defined on the desired problem graph [19, 20]. Native clique embeddings [39] (Fig. 2(a)) are pre-computed embeddings that can be used for fully-connected problems or problems on dense graphs. Each *logical* qubit is associated to an L-shaped ferromagnetic chain of  $\lceil N/\kappa \rceil + 1$  *physical* qubits, where  $2\kappa$  is the number of qubits in each unit cell of the Chimera graph ( $\kappa = 4$  in the D-Wave 2000Q). Clique

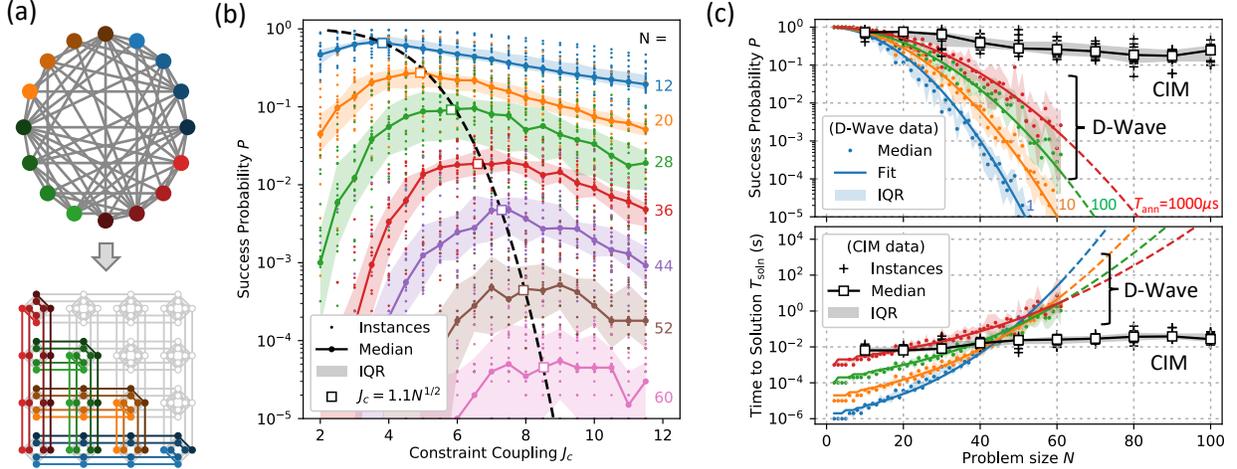


Figure 2: (a) Illustration of clique embedding: an arbitrary  $N = 16$  graph is embedded into the D-Wave chimera, each spin mapped to a ferromagnetically coupled line of physical qubits (each color is a logical qubit). (b) D-Wave ground-state probability for Sherrington-Kirkpatrick (SK) model as a function of problem size  $N$  and embedding parameter  $J_c$ . Shading indicates interquartile range (IQR, 25/75 percentile range of instances). (c) Scaling of ground-state probability and time to solution for DW2Q (with optimal  $J_c$ ) and Stanford CIM. D-Wave and CIM ran 20 and 10 instances per problem size, respectively.

embeddings are desirable because all chain lengths are equal: this simplifies the parameter setting procedure due to symmetry and it is thought to prevent desynchronized freeze-out of chains during the calculation [40]. However, the embedding introduces considerable overhead relative to the fully-connected model: for  $N$  logical qubits,  $N(\lceil N/\kappa \rceil + 1) \approx N^2/\kappa$  physical qubits are used. Due to the triangular structure of the embeddings (Fig. 2(a)), only approximately half of the annealer’s physical qubits are utilized, limiting the D-Wave 2000Q to problems with  $N \leq 64$  (the actual limit is  $N \leq 61$  due to unusable qubits on the particular machine at NASA Ames).

## Results

As a first benchmarking problem, we consider the Sherrington-Kirkpatrick (SK) spin-glass model on a fully-connected (i.e., maximally dense) graph, where the couplings  $J_{ij} = \pm 1$  are randomly chosen with equal probability [26]. Ground-state computation of the SK model is directly related to the graph partitioning problem, which is also NP-hard [41]. For each problem size  $2 \leq N \leq 61$ , 20 randomly-chosen instances were solved on the DW2Q. We consider as a performance metric the success probability  $P$ , defined as the fraction of runs on the same instance that return the ground state energy, as well as the *time to solution*  $T_{\text{soln}} = T_{\text{ann}}[\log(0.01)/\log(1 - P)]$ , which multiplies the expected number of independent runs to solve a problem with 99% probability with the time of a single run,  $T_{\text{ann}}$ .

Fig. 2(b) shows that the DW2Q performance is strongly dependent on the embedding parameter  $J_c$ , and the optimal  $J_c$  scales roughly as  $N^{1/2}$  (see methods for details). This scaling is consistent with results published on the same class of problems with the earlier D-Wave Two quantum annealer, and it is believed to be connected to the spin-glass nature of the SK Ising problem [40]. Fig. 2(c) shows that the performance on the D-Wave depends strongly on the single-run annealing time, with the values  $T_{\text{ann}} = (1, 10, 100, 1000)\mu\text{s}$  plotted here. The D-Wave annealing time is restricted to the range  $[1, 2000]\mu\text{s}$ . We observe that longer annealing times give higher success probabilities, in accordance with the expectations from the adiabatic quantum optimization approach that inspired the design of the D-Wave machine. The data fit well to a square-exponential  $P = \exp(-(N/N_0^{\text{DW}})^2)$ , where the parameter  $N_0^{\text{DW}}$  increases slowly, roughly logarithmically, with  $T_{\text{ann}}$ . The key computing figure of merit is the total annealing time to solution  $T_{\text{soln}}$ . This metric shows

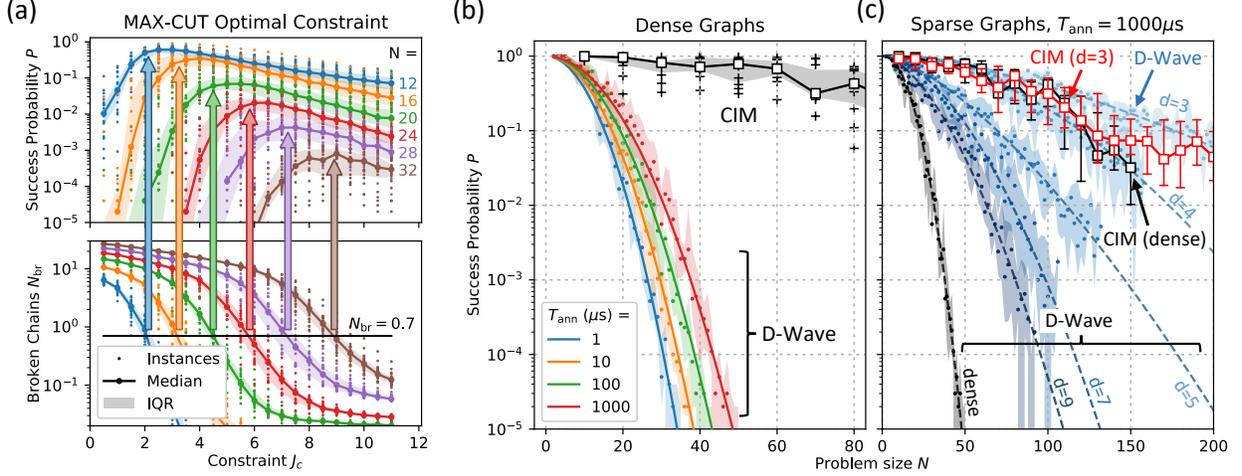


Figure 3: (a) D-Wave performance on dense MAX-CUT problems with (edge density 0.5), showing that optimal performance occurs when the  $J_c$  coupling is strong enough to make it unlikely that logical qubits (chains) become “broken” (see also Supp. Figs. S1-S2). (b) D-Wave and NTT CIM success probability for dense MAX-CUT as a function of problem size (for  $T_{\text{soln}}$  see Supp. Fig. S6). (c) D-Wave (annealing time  $T_{\text{ann}} = 1000\mu\text{s}$ ) and NTT CIM success probability for sparse graphs of degree  $d = 3, 4, 5, 6, 9$  as well as dense graphs.

a tradeoff between the annealing time of a single run and success probability: short anneals are preferred for small problems where the success probability is always close to unity and insensitive to the annealing time, and long anneals are preferred for large problems where the success probability dominates. For problem sizes  $N < 30$ , the results in Fig. 2 agree with an extrapolation of the benchmark data for  $T_{\text{ann}} = 20\mu\text{s}$  reported in Ref. [40], which used an earlier processor (the 512-qubit D-Wave, despite the engineering improvements that have been made in the last two generation chips (2X and 2000Q)).

The same SK instances for  $N = 10, 20, \dots, 60$  were solved on the CIMs hosted at Stanford University in Stanford, California and NTT Basic Research Laboratories in Atsugi, Japan [23, 24]. Additional problems with  $N \geq 60$  were also solved on the CIM, but were too large to be programmed on the DW2Q. The two Ising machines have similar performance (see Supplementary Sec. S2 for more details). Fig. 2(c) shows a plot of the success probability as a function of problem size: the exponential scaling for the CIM is shallower than the one given by the DW2Q performance. We note that the success probability  $P$  for the CIM scales approximately as  $\exp(-N/N_0^{\text{CIM}})$ , where  $N_0^{\text{CIM}}$  is a constant. The fact that for the DW2Q, success probability  $P$  scales with an  $N^2$  dependence in the exponential rather than  $N$  (as is the case for the CIM) leads to a dramatic difference in success probability between the quantum annealer and the CIM for problem sizes  $N \geq 60$ . This large difference is also reflected in the times-to-solution for the CIM versus the D-Wave quantum annealer.

We next study the DW2Q performance on MAX-CUT for both dense and sparse unweighted graphs. Unweighted MAX-CUT is the problem of finding a partition (called a cut) of the vertices  $V$  of a graph  $G = (V, E)$  where the partition is defined by two disjoint sets  $V_1$  and  $V_2$  with  $V_1 \cup V_2 = V$ , and for which the number of edges between the two sets  $|\{(v_1 \in V_1, v_2 \in V_2) \in E\}|$  is maximized. Unweighted MAX-CUT is NP-hard for general graphs [2], and can be expressed as an Ising problem by setting the anti-ferromagnetic couplings  $J_{ij} = +1$  along graph edges:  $H = \sum_{(ij) \in E} \sigma_i \sigma_j$ . Thus, the problem in Fig. 1(c) is the same as MAX-CUT on the Möbius ladder graph. Previous CIM studies have solved MAX-CUT on problems up to size  $N = 2,000$  in experiment [29, 36, 23, 24] and  $N = 20,000$  in simulation [28, 42].

Random unweighted MAX-CUT graphs of edge density 0.5 (i.e. Erdős-Rényi graphs  $G(N, \frac{1}{2})$ ) were tested on DW2Q for problems up to  $N = 61$ , and on the CIM for  $N \leq 150$ . For these graphs, clique embeddings were used, but in practice the performance did not differ from the embedding heuristic provided by the D-Wave

API [20]. In Fig. 3(a) we show that the optimal value of the embedding coupling parameter  $J_c$  appears to be correlated with the appearance of defects in the perfect polarization state expected in logical qubits at the end of the anneal. The success probability follows the same square-exponential ( $e^{-O(N^2)}$ ) trend with  $N$  as in the SK model, but the drop-off is even steeper. The CIM success probabilities are also lower than for the SK model, but are now orders of magnitude higher than the DW2Q for  $N \geq 40$ . By  $N = 50$ , the CIM success probability is  $10^5$  times larger than that for the DW2Q, and there is correspondingly a greater-than- $10^3$  times difference in the measured time to solution (Supp. Fig. S6). Extrapolated to  $N = 100$ , this difference exceeds  $10^{15}$ .

To test the effect of sparseness, Fig. 3(c) plots the performance on unweighted regular graphs of degree  $d = 3, 4, 5, 7, 9$ , where the degree of a graph is the number of edges per vertex. Despite their sparseness, MAX-CUT on these restricted graph classes is also NP-hard [43]. The CIM shows no performance difference between  $d = 3$  (cubic) and dense graphs. For DW2Q, the sparse graphs are embedded using the graph minor heuristic, which allows problems of up to size  $N = 200$  to be embedded in the DW2Q [20]. In addition, the found embeddings require significantly fewer qubits (for the sparse graphs) than the clique embeddings. For cubic graphs, the DW2Q achieves slightly better performance than the CIM, while the CIM’s advantage is noticeable for  $d \geq 5$ .

The CIM achieves similar success probabilities for cubic and dense graphs, suggesting that dense problems are not intrinsically harder than sparse ones for this class of annealer. D-Wave’s strong dependence on edge density is most likely a consequence of embedding compactness: it is known that more compact embeddings (fewer physical qubits per chain) tend to give better annealing performance, after all optimization and parameter setting is considered [20]. Since qubits on the D-Wave chimera graph have at most 6 connections, the minimum chain length is  $\ell = \lceil (d-2)/4 \rceil$ , so embeddings grow less compact with increasing graph degree (see Supp. Sec. S3). Since degree-1 and degree-2 vertices can be pruned from a graph in polynomial time (a variant of cut-set conditioning [44]),  $d = 3$  is the minimum degree required for NP-hardness. Of NP-hard MAX-CUT instances, Fig. 3(c) suggests that there is only a very narrow region ( $d = 3, 4$ ) where D-Wave matches or outperforms the CIM; for the remainder of the graphs the CIM dominates.

Fixing the problem size and varying the edge density, we see the same effect and can fill in the gap between sparse graphs and dense graphs. We constructed random unweighted graphs of degree  $d = 1, 2, \dots, (N-2)$  for each graph size  $N = 20, 30, 40, 50, 60$ . The success probabilities for DW2Q and the CIM are shown in Fig. 4(a) (for clarity only  $N = 40$  CIM data are shown). In this case, we used clique embeddings for all problems, so for a given  $N$  all the embeddings are the same. Even with the embeddings fixed, the DW2Q finds sparse problems easier to solve than dense ones. The reason is that, consistent with Ref. [40], the optimal constraint coupling is weaker for sparse problems than for dense problems (Fig. 4(b)). In general, we find that  $J_c \propto d$  for fixed  $N$ . Having a large constraint coupling could be problematic because the physical quantum annealer scales the largest coupling coefficient to the maximum coupling strength on the chip; the constraints max out this coupling and cause the logical couplings to be downscaled proportionally as  $J_c^{-1}$ . Thus dense graphs have weaker logical couplings in the embedded problem, hindering the annealer’s ability to find the ground state due to parameter misspecification or “intrinsic control errors” (ICE) [40, 45].

The CIM has only weak dependence on the edge density  $x = d/(N-1)$ . Earlier work on  $N = 100$  graphs [23], as well as the CIM data plotted in Fig. 3(c), are consistent with this result. This suggests that the CIM has promise as a general-purpose Ising solver, achieving good performance on a large class of problems, irrespective of connectivity.

Comparing Figs. 4(a) and 3(c) we can glean some insight regarding the effect of embedding overhead on the D-Wave quantum annealer’s performance. The heuristic embeddings in Fig. 3(c) are designed to minimize the overhead factor (ratio of physical qubits to logical qubits). This ratio is much larger for the native-clique embeddings, growing linearly, i.e., as  $O(N)$  (see Supplementary Sec. S1). Fig. 4(c) compares these two D-Wave settings against the CIM at  $N = 50$ ; while the CIM out-performs on all graphs with  $d \geq 5$ , the difference between the success probabilities using clique and heuristic embeddings suggests that performance is heavily dependent on embedding overhead and the difference grows with edge density (and graph size). This illustrates an additional tradeoff in quantum annealing: poor-performing but easy-to-find embeddings vs. well-performing embeddings that require substantial pre-computation. This tradeoff is expected to favor the well-performing embeddings when the number of qubits (or connections) becomes large.

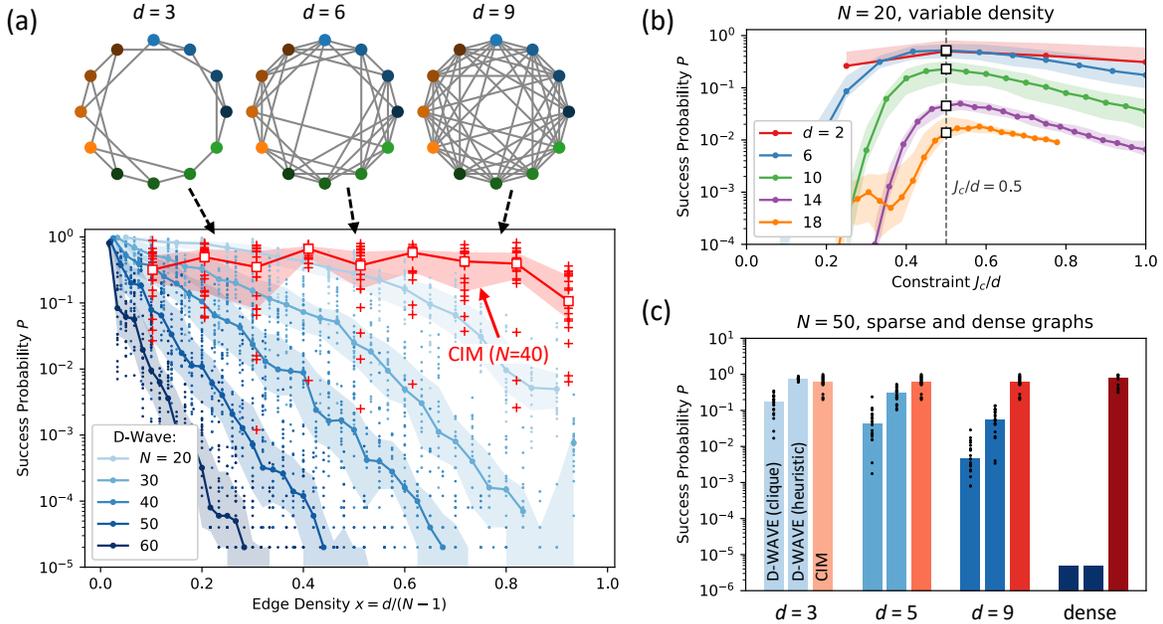


Figure 4: (a) Success probability as a function of edge density. Native clique embeddings used for D-Wave. Optimal embedding parameter (see subgraph (b)) is used, with  $T_{\text{ann}} = 1000\mu\text{s}$ . (b) D-Wave success probability as a function of graph degree, showing that the optimal  $J_c$  scales as  $J_c \propto d$  for fixed  $N$ . (For fixed edge density, the  $N$  dependence was determined previously to be  $J_c \propto N^{3/2}$ , see Fig. 3(a)). (c) Comparison of D-Wave and NTT CIM success probabilities for  $N = 50$ , using both clique embeddings and heuristically determined embeddings ( $T_{\text{ann}} = 1000\mu\text{s}$ , dense D-Wave bars are extrapolation from  $e^{-(N/N_0)^2}$  fit in Fig. 3(b))

## Discussion

In conclusion, we have benchmarked the D-Wave 2000Q system hosted at NASA Ames and measurement-feedback CIMs hosted at Stanford University and NTT Basic Research Laboratories, focusing on the MAX-CUT problems on random graphs and Sherrington-Kirkpatrick spin-glass models, and found that the merits of each machine are highly problem-dependent. Connectivity appears to be a key factor driving the performance of these machines. Problems with sparse connectivity, such as 1D chains (compare Refs. [46] and [47]) and MAX-CUT on cubic graphs (Fig. 3), can be embedded into the DW2Q with little or no overhead, resulting in similar performance from the quantum annealer and the CIMs. However, the embedding overhead for dense problems like SK is very steep, requiring  $O(N^2)$  physical qubits to represent a size- $N$  graph. This likely contributes to the CIM's performance advantage on SK models that grows exponentially with the square of the problem size. For problems of intermediate sparseness, such as MAX-CUT on graphs of degree  $d \geq 5$  the CIM is still faster by a large factor. It has been conjectured often that increased internal connectivity in quantum annealers will result in improved performance [48, 49, 50, 51, 52, 53, 54], and there are large projects underway to realize higher-connectivity quantum annealers (including efforts by D-Wave, as well as MIT-LL [55] and Google [56]). Our results provide strong experimental justification for this line of development.

## Methods

### Sample Problems

For fully-connected SK and MAX-CUT on dense graphs, 20 random instances were created of each size  $N = 2, 3, \dots, 61$  for the D-Wave. Of these, the  $N = 2, 10, 20, \dots, 60$  instances were also used for the CIM.

An additional set of random instances were created for  $N = 70, 80, \dots, 150$  for the CIM, using the same algorithm.

For the sparse-graph analysis, we computed regular graphs of size  $N = 2, 4, \dots, 300$  and degree  $d = 3, 4, \dots, 20$ , with 20 instances for each pair  $(N, d)$ . The algorithm randomly assigns edges to eligible vertices until all reach the required degree (and backtracks if it gets stuck). The same algorithm was also used for the variable-density graphs:  $d = 1, 2, \dots, (N - 2)$  for  $N = 20, 30, 40, 50, 60$ , creating 20 instances per pair  $(N, d)$ .

Exact SK ground states were found with the Spin Glass Server [57], which uses BiqMac [58], an exact branch-and-bound algorithm. For SK instances of size  $N \leq 100$ , the algorithm obtained proven ground states. For  $N > 100$  the solver timed out before exhausting all branches (runtime  $T = 3000$  s), so the result is not a guaranteed ground state; however, we believe it reaches the ground state with high probability for  $N \leq 150$  because multiple runs of the algorithm give the same state energy, and none of the CIM runs found an Ising energy lower than the Spin Glass Server result. MAX-CUT ground states for  $N \leq 30$  were found by brute-force search on a GPU; for  $20 \leq N \leq 150$  a Breakout Local Search (BLS) algorithm was used [5]. Although BLS is a heuristic solver, for  $N \leq 150$  it finds the ground state with nearly 100% probability, giving us high confidence that the BLS solutions are ground states. While the brute-force solver, D-Wave, and the CIM found states of equal energy to the BLS solution (if run long enough), they never found states of lower energy.

## D-Wave annealers

Initial D-Wave experiments were performed on the D-Wave 2X at NASA Ames Research Center and the D-Wave 2X online system at D-Wave Systems Inc. Later runs were made on the D-Wave 2000Q at NASA Ames, once that machine came online. The 2X and 2000Q systems use a C12 (12 cells  $\times$  12 cells  $\times$  4 qubits) and C16 (16  $\times$  16  $\times$  4) Chimera, respectively. For all-to-all graphs, D-Wave 2X supports  $N \leq 48$  and 2000Q supports  $N \leq 64$  (the number is slightly smaller because of broken qubits). All  $N \leq 48$  runs were consistent across the three machines as well as with extrapolation of data in Ref. [40] from runs performed on a different set of instances on the earlier generation machine D-Wave Two. All data reported in this paper came from the D-Wave 2000Q.

Embeddings were pre-computed for all problems (heuristic embeddings for sparse MAX-CUT; native clique embeddings for SK, dense MAX-CUT, and variable-density MAX-CUT) so that runs in different conditions (e.g. annealing times, constraint couplings) would use the same embeddings. The standard annealing schedule was used in all experiments, but the annealing time was tuned. Each instance was run  $10^4$ – $10^5$  times total, depending on the observed success rate. 5–10 different embeddings were used per instance and the success probability was averaged. Spin-reversal transformations were used to avoid spurious effects. After an anneal, each logical qubit value was determined by taking the majority vote of all qubits in the chain.

In all figures, the shaded regions give the [25, 75]-percentile range (inter-quartile range, or IQR) for the data. Figs. 2(b), 3(a), 4(a), show individual instances as dots and the solid line gives the median. Figs. 2(c), 3(b-c), 4(b) are too crowded to show D-Wave instances; the dots give medians and the smooth lines give analytic fits. For CIM data, medians and IQR are shown in Figs. 2(c), 3(b), while Fig. 3(c) only shows medians and IQR, due to crowding.

## CIM

CIM experiments were performed on the 100-OPO CIM at Ginzton Laboratory of Stanford University and the 2048-OPO CIM at NTT Basic Research Laboratories. The Stanford and NTT devices are described in Refs. [23] and [24], respectively. Computation time of the Stanford CIM is 1.6ms, which is the time for 1000 round-trips of the 320-m fiber ring cavity. Since the NTT CIM processes 2000-node problem in 5.0ms, which is the time for 1000 round-trips of the 1-km fiber ring cavity, we can solve up to  $\lfloor 2000/N \rfloor$  problems in parallel per the computation time.

The CIM's reliable operation depends on relative phases between the OPO pulses, injection pulses, and measurement-LO pulses being kept stable and well-calibrated. Such phase stabilization is imperfect in the experimental setups used in this study, and consequently post-selection procedures have been applied to both the Stanford and NTT CIM experimental data. This is described in detail in Supp. Sec. S2. Computation times have been reported in terms of annealing times; as with the DW2Q, these times exclude the time required to transfer data to and from the CIM

## Acknowledgements

This research was funded by the Impulsing Paradigm Change through Disruptive Technologies (ImPACT) Program of the Council of Science, Technology and Innovation (Cabinet Office, Government of Japan). R.H. is supported by an IC Postdoctoral Research Fellowship at MIT, administered by ORISE through U.S. DOE and ODNI. P.L.M. was partially supported by a Stanford Nano- and Quantum Science and Engineering Postdoctoral Fellowship. D.V. acknowledges funding from NASA Academic Mission Services, contract no. NNA16BD14C. H.M., E.N., and T.O. acknowledge funding from NSF award PHY-1648807.

Y.Y., P.L.M. and E.R. proposed the project. R.H. wrote the manuscript and performed D-Wave experiments and data analysis. T.I. and P.L.M performed NTT and Stanford CIM experiments, respectively. D.V. helped with D-Wave experiments and data analysis. A.M., C.L., R.L.B., M.M.F., and H.M. built the Stanford CIM, and T.O. and E.N. helped with theory. K.I., T.H., K.E., T.U., R.K., and H.T. built the NTT CIM. D.V., P.L.M., T.I., Y.Y., and A.M. assisted with preparation of the manuscript. S.U., S.K., and K.K. assisted with interpretation of the results.

The authors declare that they have no competing financial interests.

Correspondence and requests for materials should be addressed to R.H. (e-mail: rhamerly@mit.edu), T.I. (e-mail: inagaki.takahiro@lab.ntt.co.jp), or P.L.M. (e-mail: pmcmahon@stanford.edu).

# [Supplementary] Scaling advantages of all-to-all connectivity in physical annealers: the Coherent Ising Machine vs. D-Wave 2000Q

## S1 D-Wave embeddings and $J_c$ optimization

Native clique embeddings [39] are used for all SK problems, MAX-CUT problems on graphs with edge density 0.5, and MAX-CUT problems on varying-density graphs (Figs. 2(c), 3(b) and 4(a) respectively in main text). The code to generate the embeddings is available on GitHub [59]. Once an embedding is chosen, the embedding parameter  $J_c$  (ferromagnetic coupling between qubits in a chain) is tuned to maximize performance. In no cases does the optimal  $J_c$  depend on the annealing time.

Fig. S1 shows that the optimal  $J_c$  scales roughly as  $N^{1/2}$  for SK problems and  $N^{3/2}$  for MAX-CUT problems of edge density 0.5. In particular, the relations  $J_c = 1.1N^{1/2}$  (SK) and  $J_c = 0.047N^{3/2}$  (MAX-CUT) were used in Figs. 2(c), 3(b).

For graphs with variable edge density, it was shown in Fig. 4(b) that the optimal  $J_c$  scales as  $d$  for fixed  $N$ , with  $J_c = 0.5d = 9.5x$  for  $N = 20$  shown in the figure ( $x = d/(N - 1)$  is the edge density). Extrapolating this using the  $N^{3/2}$  relation above (which holds for constant  $x = \frac{1}{2}$ ), we used  $J_c = 9.5(N/20)^{3/2}x$ , which is very close to the  $J_c = 0.047N^{3/2}$  used for edge-density 0.5 graphs. The relation was also tested for  $N = 30$  variable edge-density graphs and found to be optimal.

Fig. 3(a) of the main text suggests that the success probability is maximized when the number of broken chains is  $N_{\text{br}} \approx 0.7$ . Plotting  $N_{\text{br}}$  as a function of  $N$  and  $J_c$  in Fig. S2, we see that  $N_{\text{br}} \approx 0.7$  for a narrow range of  $J_c$  centered around the line  $J_c = 0.047N^{3/2}$ . For a wide range of  $N$ , this value of  $J_c$  also roughly maximizes the success probability (Fig. S1).

The fact that dense MAX-CUT problems are optimally embedded when  $N_{\text{br}} = O(1)$  is an example of the general principle that  $J_c$  must neither be too strong nor too weak for a problem. If  $J_c$  is too small so that  $N_{\text{br}} \gg 1$ , the constraint is not enforced effectively and thus the embedded problem can have a ground state that is different from the logical problem. Once  $N_{\text{br}} \lesssim 1$ , increasing  $J_c$  further will not improve the computation significantly because all of the constraints are already satisfied with high probability. Rather, it degrades performance because  $J_c$  maxes out the physical coupling on the chip so that logical couplings are

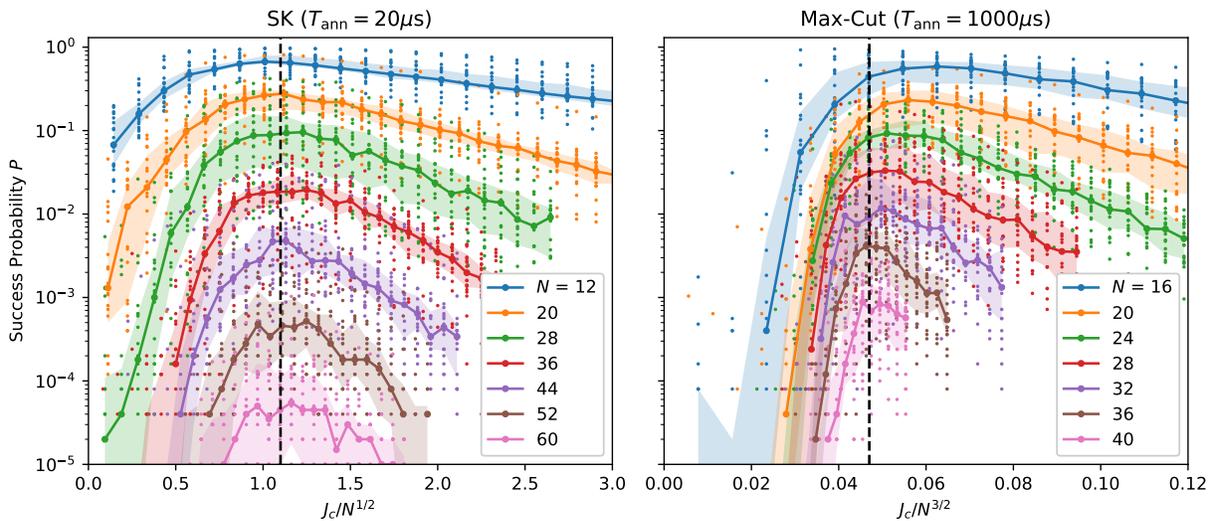


Figure S1: D-Wave success probability for SK problems and MAX-CUT problems of edge density 0.5, as a function of problem size  $N$  and embedding parameter  $J_c$ .

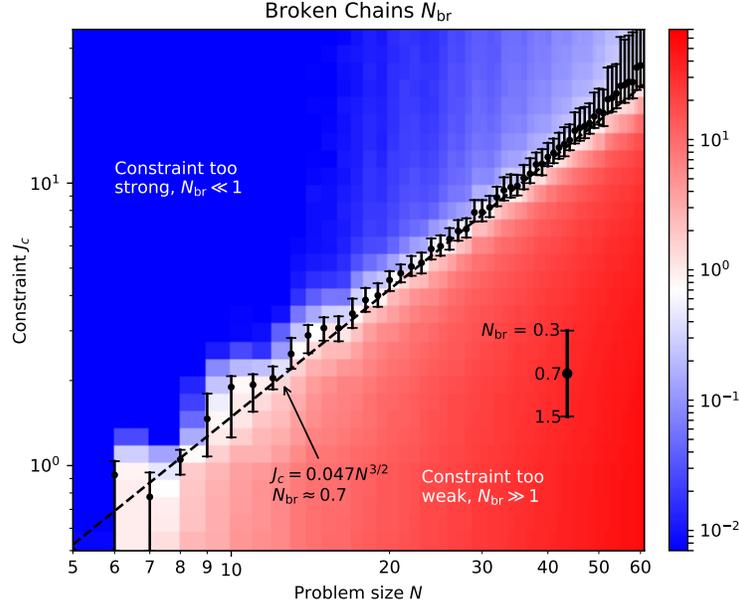


Figure S2: Max-Cut on edge-density 0.5 graphs. Broken chains as a function of problem size  $N$  and embedding parameter  $J_c$

scaled down as  $J_c^{-1}$ , which will correspondingly reduce the spectral gap of the (physical) Hamiltonian, and can also cause problems due to the finite bit precision and hardware imperfections of the D-Wave system.

For the sparse graphs, embeddings were found using the heuristic of Cai et al. [20], which is available as part of the D-Wave API toolkit. For each sparse graph instance, we attempt to generate 10 embeddings using the heuristic with a time-out of 60 seconds. The probability of finding an embedding is shown in Fig. S3(a) (the  $d = 3$  case is in agreement with [20, Fig. 7]). The time required to find an embedding (on average) and the number of physical qubits  $N_{\text{emb}}$  are also plotted in Fig. S3(a).

Fig. S3(b) shows the number of physical qubits for graphs of degree  $d = 3, 4, 5, 7, 9$  embedded using the heuristic, as well as the average chain length  $L = N_{\text{emb}}/N$ . This is compared against the clique embeddings described above.

Because the heuristic embeddings differ markedly from clique embeddings, we do not use the formula  $J_c = 9.5(N/20)^{3/2}x$  derived above. Rather, the optimal  $J_c$  is found by hand, running the quantum annealer for a range of  $N$ ,  $d$  and  $J_c$  (Fig. S4). We find that the optimal  $J_c$  is independent of  $N$  for sufficiently large  $N$ , while it increases slightly for small  $N$  for  $d = 7, 9$ . We interpolate using the curves of Fig. S4 to find the embedding parameter used in the main text (Fig. 3(c)).

## S2 CIM data and post-selection

The CIM is based on an OPO network, which is sensitive to optical phase fluctuations. During the course of operation, the phase of the injection beam will drift. This drift is slow compared to experimental timescales, but can become large if a calculation is run thousands of times.

To filter out out-of-phase computations (which always lead to the wrong answer), each CIM includes a phase-checking mechanism, albeit somewhat different for the NTT and the Stanford CIMs. We summarize both here.

In the NTT system, phase stability and calibration is implemented with a phase-check graph: the 2,048 spins in the CIM are partitioned into a 16-spin (unused) header, a 32-spin bipartite graph for phase checking, and a

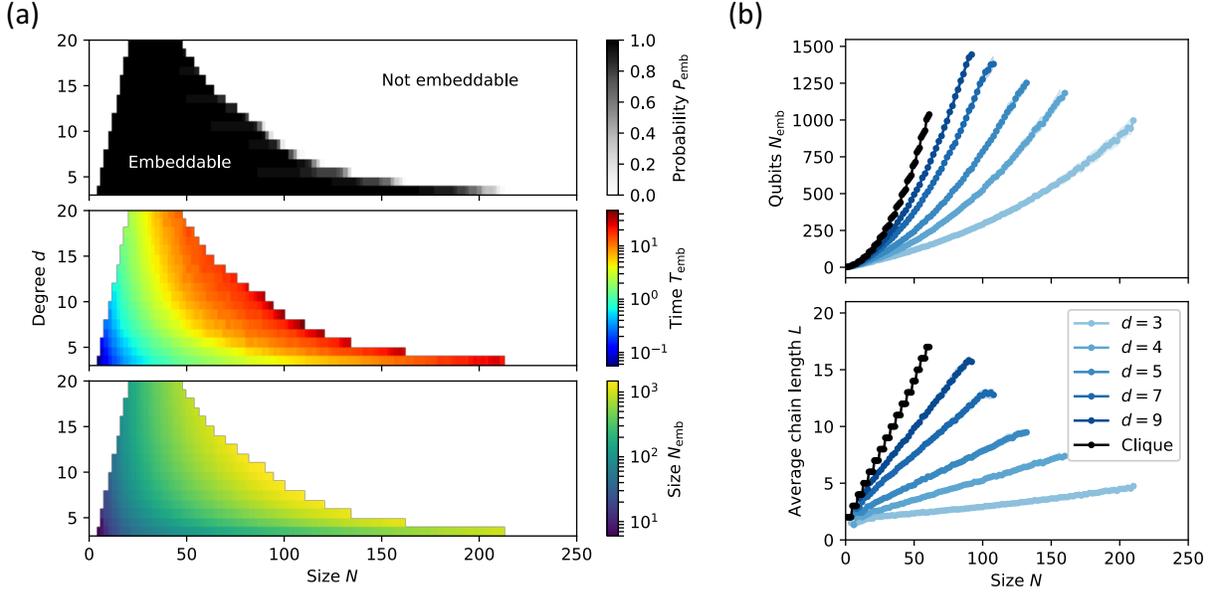


Figure S3: (a) Probability of finding an embedding using the heuristic, average time required to find an embedding, and number of physical qubits as a function of graph parameters  $(N, d)$  for fixed-degree graphs. (b) Number of qubits and average embedding chain length as functions of  $N$ .

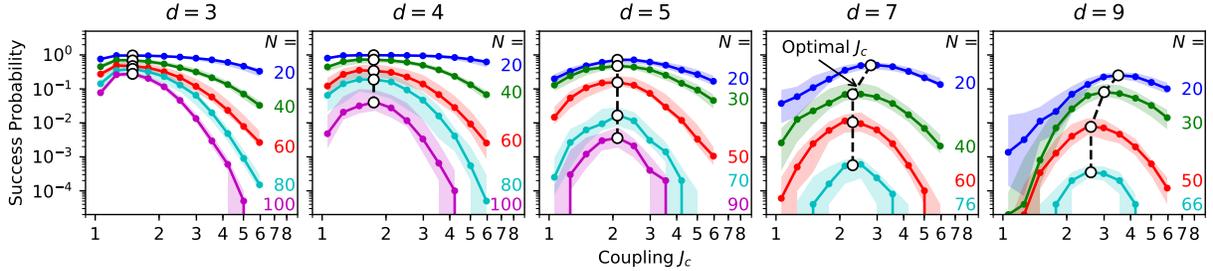


Figure S4: Choice of optimal coupling for sparse graphs using the heuristic embedding.

“frame” of 2,000 spins for the desired problem. Since  $N \ll 2000$  for the problems in this paper, we can solve up to  $\lfloor 2000/N \rfloor \approx 2000/N$  problems in parallel per frame. The coupling matrix  $J_{ij}$  has a block-diagonal structure (Fig. S5(a)).

The couplings of the bipartite graph for phase-check are randomly set to  $+1$  or  $-1$  and the value of the phase-check Hamiltonian  $H_{PC} = \frac{1}{2} \sum_{ij} J_{ij} \sigma_i \sigma_j$  is computed after each run. If the optical phase is incorrect, we find  $H_{PC} > 0$  because the system couplings are reversed and the machine is trying to minimize  $-H_{PC}$ . The top plot of Fig. S5(b) shows the phase-check  $H_{PC}$  value (normalized to the maximum) as a function of time.  $H_{PC}$  drops sharply to a negative value when the CIM is in phase, making it a good proxy for the CIM phase.

In the bottom plots of Fig. S5(b), three data-filtering techniques are shown. Here we plot the free-running success probability (fraction of instances per frame in the ground state) for an  $N = 50$  problem (40 trials running in parallel per frame). Averaging over all frames requires no post-processing, but gives a low success probability because we are including many trials when the machine is out of phase. Filtering on the phase-check graph (green curve) does significantly better; however, we are still averaging over the edges of the phase-check region where the system is only marginally in phase. Still better success probabilities can be

found by looking for the best batch of 1,000 consecutive trials (20 consecutive frames) in the series (red curve). This generally corresponds to the the CIM working in its best condition: when the feedback signal is well in phase. This is the success probability we could expect from a well-engineered CIM where the optical phase, pump power, and other optical degrees of freedom have been sufficiently stabilized.

We compare the three post-selection methods in Fig. S5(c) to show that our post-selection techniques give only a constant improvement in success probability, and this constant is never more than an order of magnitude. Thus, we can safely conclude that the CIM’s performance advantage does not arise from cherry-picking good samples from the data. The “best batch” method (red curves in Fig. S5) is used to process all CIM data reported in the main text.

The data collected from the Stanford CIM was also post-processed to select only the runs on the machine for which the optical setup was optimally stable. However, the procedure for post-selection was slightly different to that used for the data from the NTT CIM. In the case of the Stanford CIM, a recording of the homodyne measurement of the output pulses immediately before a run began was stored. During this recording phase, constant-amplitude pulses were injected into the cavity. If the entire system is phase-stable, then the recorded homodyne measurement results should not show large fluctuations from pulse to pulse. Furthermore, the particular value of the phase of the injected light is also relevant (not just that it is ideally constant), since the computation mechanism relies on interference of injected pulses with pulses in the cavity, and how much interference is obtained is partially determined by the phase of the injection pulses. We therefore post-selected not only for stability, but also for a particular mean value of the homodyne measurement results, which was determined on an instance-by-instance basis. The net effect of this post-selection procedure is to produce success probabilities that represent the probabilities one would obtain if the CIM was always phase-stable whenever a computation was run, and the phase was correctly calibrated for each problem instance.

The post-selected success probabilities were only on average  $5\times$  higher than the success probabilities obtained when no post-selection was applied. This implies that even if one is pessimistic about the prospects of improvement to the optical phase stabilization of the CIM, and one assumes that the most stable the

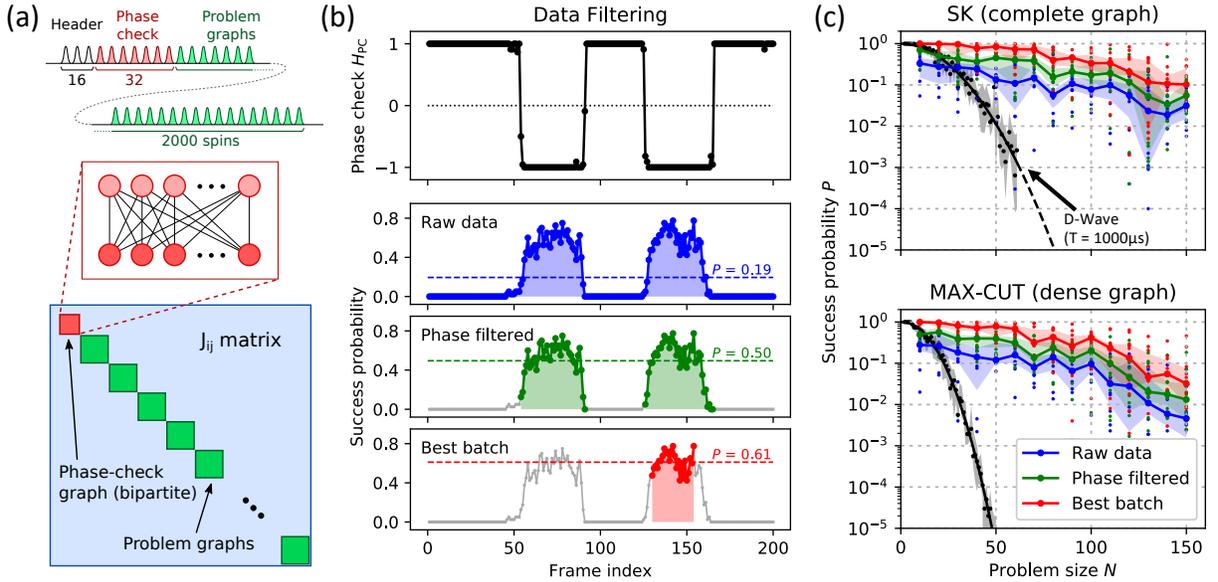


Figure S5: (a) Partitioning of NTT CIM spins into a 16-spin header, a 32-spin phase-check graph, and 2,000 spins for problem graphs, and the resulting  $J_{ij}$  matrix. (b) Phase-check Hamiltonian  $H_{PC}$  as a function of time (frame index), and three post-selection techniques for inferring the success probability. (c) NTT CIM success probability for SK and  $x = 0.5$  MAX-CUT problems as a function of post-selection method.

machine will ever be is as it was during the experiments reported in this paper, then at worst one should divide the success probabilities for the Stanford CIM reported in this paper by  $5\times$ . This gives the estimate for the expected success probabilities for a machine that has the same fundamental operating principle as the currently implemented CIM at Stanford, as well as the same experimental imperfections (including phase noise) that the current setup has.

The CIMs at Stanford and NTT were run on the same (randomly-chosen) Ising problems for  $N \leq 100$  MAX-CUT (edge density  $x = 0.5$ ) and SK (fully connected). The average success probabilities of the two machines agree to within a factor of 5 (Fig. S6).

In order to compare the solution time  $T_{\text{soln}}$  with D-Wave, we need the physical annealing time for the CIM. A strict minimum for the annealing time is given by the product of the time between pulses (equal to  $1/f$  where  $f$  is the pump repetition frequency), the size of the problem  $N$ , and the number of round trips per run  $R$ :

$$T_{\text{ann}}^{(\text{min})} = \frac{NR}{f} \quad (\text{S1})$$

This is the effective annealing time if perfect parallelization is achieved and all spins are used for logic (i.e. a negligible fraction of phase-check and dummy spins). Both Stanford and NTT CIMs use  $R = 1000$  round trips.

However, the annealing time is generally longer than  $T_{\text{ann}}^{(\text{min})}$  because dummy spins are added to the cavity to compensate for the delays due to the DAC / ADC electronics in the feedback circuit and to give the FPGA more time to finish the coupling computation. This increases the cavity round-trip time and thus the annealing time.

In the NTT CIM, we used 5056 pulses in a 1-km fiber ring cavity as: 16-spin (header), 32-spin (phase check), 2000-spin (solve problem), 100-spin (blank), 2808-spin (free running in FPGA calculation time), 100-spin

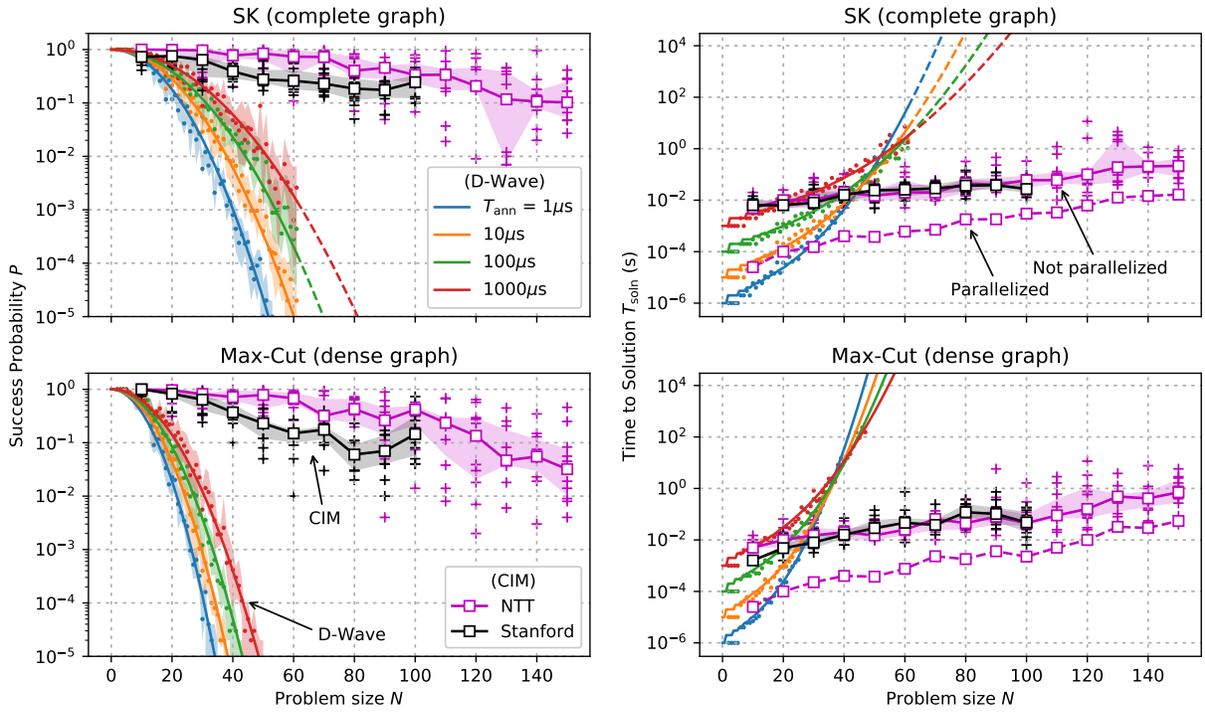


Figure S6: Comparison of Stanford and NTT CIM performance for SK and dense MAX-CUT problems. D-Wave data for  $T_{\text{ann}} = 1, 10, 100, \text{ and } 1000\mu\text{s}$  are also plotted.

(blank). The pump repetition rate is 1 GHz and the round-trip time is  $5\mu\text{s}$ . As only 2000 of 5056 pulses are used, even if perfect parallelism is employed, the annealing time is approximately  $2.5\times$  longer than Eq. (S1), or  $T_{\text{ann}} = (2.5N)\mu\text{s}$ , where  $N$  is the problem size. Fig. S6 plots the NTT CIM time-to-solution both with and without parallelism, to enable a fair comparison with the D-Wave annealer (we did not attempt to parallelize D-Wave to run multiple problems per anneal).

In the Stanford CIM, which did not employ parallelism due to its smaller number of spins, the annealing time is  $T_{\text{ann}} = 1.6\text{ ms}$  for all problems. The Stanford CIM [23] features a 320-m fiber ring cavity that contains 160 optical pulses (repetition rate 100 MHz), of which up to 100 can be used to encode Ising problems. The data in Fig. 2(c) come from the Stanford CIM, where the above annealing time combined with the formula  $T_{\text{soln}} = T_{\text{ann}}[\log(0.01)/\log(1 - P)]$  is used to calculate the time to solution.

## References

- [1] Rudich, S. & Wigderson, A. *Computational Complexity Theory*, vol. 10 (American Mathematical Soc., 2004).
- [2] Karp, R. M. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, 85–103 (Springer, 1972).
- [3] Goemans, M. X. & Williamson, D. P. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)* **42**, 1115–1145 (1995).
- [4] Kirkpatrick, S., Gelatt, C. D. & Vecchi, M. P. Optimization by simulated annealing. *Science* **220**, 671–680 (1983).
- [5] Benlic, U. & Hao, J.-K. Breakout local search for the max-cut problem. *Engineering Applications of Artificial Intelligence* **26**, 1162–1173 (2013).
- [6] Yoshimura, C. *et al.* Uncertain behaviours of integrated circuits improve computational performance. *Scientific Reports* **5** (2015).
- [7] Tsukamoto, S., Takatsu, M., Matsubara, S. & Tamura, H. An accelerator architecture for combinatorial optimization problems. *FUJITSU Science and Technology Journal* **53**, 8–13 (2017).
- [8] Sutton, B., Camsari, K. Y., Behin-Aein, B. & Datta, S. Intrinsic optimization using stochastic nanomagnets. *Scientific Reports* **7** (2017).
- [9] Parihar, A., Shukla, N., Jerry, M., Datta, S. & Raychowdhury, A. Vertex coloring of graphs via phase dynamics of coupled oscillatory networks. *Scientific Reports* **7**, 911 (2017).
- [10] Yin, X. *et al.* Efficient analog circuits for boolean satisfiability. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* **26**, 155–167 (2018).
- [11] Tait, A. N. *et al.* Neuromorphic photonic networks using silicon photonic weight banks. *Scientific Reports* **7**, 7430 (2017).
- [12] Farhi, E. *et al.* A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science* **292**, 472–475 (2001).
- [13] Kadowaki, T. & Nishimori, H. Quantum annealing in the transverse Ising model. *Physical Review E* **58**, 5355 (1998).
- [14] Santoro, G. E. & Tosatti, E. Optimization using quantum mechanics: quantum annealing through adiabatic evolution. *Journal of Physics A: Mathematical and General* **39**, R393 (2006).
- [15] Boixo, S. *et al.* Evidence for quantum annealing with more than one hundred qubits. *Nature Physics* **10**, 218–224 (2014).

- [16] Santoro, G. E., Martoňák, R., Tosatti, E. & Car, R. Theory of quantum annealing of an Ising spin glass. *Science* **295**, 2427–2430 (2002).
- [17] Somma, R. D., Nagaj, D. & Kieferová, M. Quantum speedup by quantum annealing. *Physical Review Letters* **109**, 050501 (2012).
- [18] Rønnow, T. F. *et al.* Defining and detecting quantum speedup. *Science* **345**, 420–424 (2014).
- [19] Choi, V. Minor-embedding in adiabatic quantum computation: I. the parameter setting problem. *Quantum Information Processing* **7**, 193–209 (2008).
- [20] Cai, J., Mcready, W. G. & Roy, A. A practical heuristic for finding graph minors. *arXiv preprint arXiv:1406.2741* (2014).
- [21] Choi, V. Minor-embedding in adiabatic quantum computation: II. minor-universal graph design. *Quantum Information Processing* **10**, 343–353 (2011).
- [22] Klymko, C., Sullivan, B. D. & Humble, T. S. Adiabatic quantum programming: minor embedding with hard faults. *Quantum Information Processing* **13**, 709–729 (2014).
- [23] McMahon, P. L. *et al.* A fully-programmable 100-spin coherent Ising machine with all-to-all connections. *Science* **354**, 614–617 (2016).
- [24] Inagaki, T. *et al.* A coherent Ising machine for 2000-node optimization problems. *Science* **354**, 603–606 (2016).
- [25] Barahona, F. On the computational complexity of Ising spin glass models. *Journal of Physics A: Mathematical and General* **15**, 3241 (1982).
- [26] Kirkpatrick, S. & Sherrington, D. Solvable model of a spin-glass. *Physical Review Letters* **35**, 1792–1796 (1975).
- [27] Boyd, R. W. *Nonlinear optics* (Academic Press, 2003).
- [28] Wang, Z., Marandi, A., Wen, K., Byer, R. L. & Yamamoto, Y. Coherent Ising machine based on degenerate optical parametric oscillators. *Physical Review A* **88**, 063853 (2013).
- [29] Marandi, A., Wang, Z., Takata, K., Byer, R. L. & Yamamoto, Y. Network of time-multiplexed optical parametric oscillators as a coherent Ising machine. *Nature Photonics* **8**, 937–942 (2014).
- [30] Takesue, H. & Inagaki, T. 10-GHz clock time-multiplexed degenerate optical parametric oscillators for a photonic Ising spin network. *Optics Letters* **41**, 4273–4276 (2016).
- [31] Walls, D. F. & Milburn, G. J. *Quantum Optics* (Springer Science & Business Media, 2007).
- [32] Milburn, G. & Walls, D. Production of squeezed states in a degenerate parametric amplifier. *Optics Communications* **39**, 401–404 (1981).
- [33] Drummond, P., McNeil, K. & Walls, D. Non-equilibrium transitions in sub/second harmonic generation. *Journal of Modern Optics* **28**, 211–225 (1981).
- [34] Wu, L.-A., Kimble, H., Hall, J. & Wu, H. Generation of squeezed states by parametric down conversion. *Physical Review Letters* **57**, 2520 (1986).
- [35] Kinsler, P. & Drummond, P. D. Quantum dynamics of the parametric oscillator. *Physical Review A* **43**, 6194 (1991).
- [36] Takata, K. *et al.* A 16-bit coherent Ising machine for one-dimensional ring and cubic graph problems. *Scientific Reports* **6**, 34089 (2016).
- [37] Hamerly, R. *et al.* Topological defect formation in 1D and 2D spin chains realized by network of optical parametric oscillators. *International Journal of Modern Physics B* **30**, 1630014 (2016).
- [38] Clements, W. R. *et al.* Gaussian optical Ising machines. *Physical Review A* **96**, 043850 (2017).

- [39] Boothby, T., King, A. D. & Roy, A. Fast clique minor generation in chimera qubit connectivity graphs. *Quantum Information Processing* **15**, 495–508 (2016).
- [40] Venturelli, D. *et al.* Quantum optimization of fully connected spin glasses. *Physical Review X* **5**, 031040 (2015).
- [41] Fu, Y. & Anderson, P. W. Application of statistical mechanics to NP-complete problems in combinatorial optimisation. *Journal of Physics A: Mathematical and General* **19**, 1605 (1986).
- [42] Haribara, Y., Utsunomiya, S. & Yamamoto, Y. Computational principle and performance evaluation of coherent Ising machine based on degenerate optical parametric oscillator network. *Entropy* **18**, 151 (2016).
- [43] Alimonti, P. & Kann, V. Hardness of approximating problems on cubic graphs. *Algorithms and Complexity* 288–298 (1997).
- [44] Dechter, R. & Pearl, J. *The cycle-cutset method for improving search performance in AI applications* (University of California, Computer Science Department, 1986).
- [45] D-Wave Systems, Inc. Technical description of the D-Wave quantum processing unit (2016). Part of D-Wave API documentation.
- [46] Gardas, B., Dziarmaga, J. & Zurek, W. H. Defects in quantum computers. *arXiv preprint arXiv:1707.09463* (2017).
- [47] Inagaki, T. *et al.* Large-scale Ising spin network based on degenerate optical parametric oscillators. *Nature Photonics* **10**, 415–419 (2016).
- [48] Dumoulin, V., Goodfellow, I. J., Courville, A. C. & Bengio, Y. On the challenges of physical implementations of RBMs. In *AAAI*, vol. 2014, 1199–1205 (2014).
- [49] Katzgraber, H. G., Hamze, F. & Andrist, R. S. Glassy chimeras could be blind to quantum speedup: Designing better benchmarks for quantum annealing machines. *Physical Review X* **4**, 021008 (2014).
- [50] Biswas, R. *et al.* A NASA perspective on quantum computing: Opportunities and challenges. *Parallel Computing* **64**, 81–98 (2017).
- [51] Mott, A., Job, J., Vlimant, J.-R., Lidar, D. & Spiropulu, M. Solving a Higgs optimization problem with quantum annealing for machine learning. *Nature* **550**, 375 (2017).
- [52] Rosenberg, D. *et al.* 3D integrated superconducting qubits. *NPJ Quantum Information* **3**, 42 (2017).
- [53] Weber, S. J. *et al.* Coherent coupled qubits for quantum annealing. *Physical Review Applied* **8**, 014004 (2017).
- [54] Rieffel, E. G. *et al.* A case study in programming a quantum annealer for hard operational planning problems. *Quantum Information Processing* **14**, 1–36 (2015).
- [55] Weber, S. *et al.* Hardware considerations for high-connectivity quantum annealers. *Bulletin of the American Physical Society* (2018).
- [56] Chen, Y. *et al.* Progress towards a small-scale quantum annealer I: Architecture. In *APS Meeting Abstracts* (2017).
- [57] Liers, F. & Jünger, M. Spin glass server. URL <http://informatik.uni-koeln.de/spinglass/>.
- [58] Rendl, F., Rinaldi, G. & Wiegele, A. Solving Max-Cut to optimality by intersecting semidefinite and polyhedral relaxations. *Math. Programming* **121**, 307 (2010).
- [59] D-Wave Systems, Inc. Chimera embedding. <https://github.com/dwavesystems/chimera-embedding> (2016).