

Quantum Optimization of Fully Connected Spin Glasses

Davide Venturelli,^{1,2,*} Salvatore Mandrà,^{1,3} Sergey Knysh,^{1,4} Bryan O’Gorman,¹ Rupak Biswas,¹ and Vadim Smelyanskiy⁵

¹NASA Ames Research Center Quantum Artificial Intelligence Laboratory (QuAIL),
Mail Stop 269-1, Moffett Field, California 94035, USA

²USRA Research Institute for Advanced Computer Science (RIACS),
615 National, Mountain View, California 94043, USA

³Department of Chemistry and Chemical Biology, Harvard University,
12 Oxford Street, Cambridge, Massachusetts 02139, USA

⁴Stinger Ghaffarian Technologies Inc., 7701 Greenbelt Road, Suite 400, Greenbelt, Maryland 20770, USA

⁵Google, 150 Main Street, Venice Beach, California 90291, USA

(Received 12 March 2015; revised manuscript received 14 July 2015; published 18 September 2015)

Many NP-hard problems can be seen as the task of finding a ground state of a disordered highly connected Ising spin glass. If solutions are sought by means of quantum annealing, it is often necessary to represent those graphs in the annealer’s hardware by means of the graph-minor embedding technique, generating a final Hamiltonian consisting of coupled chains of ferromagnetically bound spins, whose binding energy is a free parameter. In order to investigate the effect of embedding on problems of interest, the fully connected Sherrington-Kirkpatrick model with random ± 1 couplings is programmed on the D-Wave Two™ annealer using up to 270 qubits interacting on a Chimera-type graph. We present the best embedding prescriptions for encoding the Sherrington-Kirkpatrick problem in the Chimera graph. The results indicate that the optimal choice of embedding parameters could be associated with the emergence of the spin-glass phase of the embedded problem, whose presence was previously uncertain. This optimal parameter setting allows the performance of the quantum annealer to compete with (and potentially outperform, in the absence of analog control errors) optimized simulated annealing algorithms.

DOI: [10.1103/PhysRevX.5.031040](https://doi.org/10.1103/PhysRevX.5.031040)

Subject Areas: Computational Physics,
Condensed Matter Physics,
Quantum Physics

I. INTRODUCTION

One tantalizing approach to solve quadratic unconstrained binary optimizations (QUBOs), such as [1,2] in their Ising formulation, is provided by programmable quantum annealing. While the founding principles of the technique have been investigated numerically [3], analytically [4,5], and experimentally [6] in the past decade, the disordered, interacting, time-dependent, and open nature of the many-body problem makes it very hard to draw universal conclusions about the power of the technique [7].

One very recent development that boosted scientific activity in this field has been the commercialization of D-Wave Two™ optimizers, which implement the annealing approach by means of a solid-state architecture consisting of hundreds of interlaced superconducting flux qubits [8]. While the manufacturing methods and the computing

technology are well documented, understanding the power of the machine is a formidable challenge for the aforementioned reasons, with the additional hindrance that the heavy integration of the circuitry entails the existence of static and dynamical sources of noise that are in part unknown. For these reasons, groups around the world have started to experimentally benchmark the machine [9–11], nurturing a lively discussion on whether the device is making functional use of quantum mechanics for computation [12,13] and how to properly measure speedups between different computational or experimental algorithms [14,15]. On a more pragmatic level, the chip was also tested to evaluate its performance on toy-application problems in the fields of network diagnostics [16], artificial intelligence [17,18], computational biology [19], and mathematics [20,21]. One typical occurrence in applied problems is when the QUBO to be solved is derived from a linear binary optimization problem with a large number of constraints, such as enforced equalities or inequalities between linear relations of variables. In this case, the resulting penalty terms in the objective function form intersecting cliques whose minimization might be a hard computational problem for classical algorithms such as simulated annealing [17].

Motivated by the great value of quantifying the power of quantum optimization on valuable applications, in this

* Author to whom correspondence should be addressed.
davide.venturelli@nasa.gov

Published by the American Physical Society under the terms of the Creative Commons Attribution 3.0 License. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

work we report on the optimal programming guidelines and performance expectation of the D-Wave TwoTM Vesuvius chip, applied to problems defined on fully connected graphs with random couplings in the absence of longitudinal local fields. This Hamiltonian corresponds to the Sherrington-Kirkpatrick (SK) model with couplings randomized from a bimodal distribution of values ± 1 [22]. The SK model is directly related to the graph partitioning problem [23], which is known to be NP hard, and supports a spin-glass phase at finite temperature with transverse fields. For these reasons, it represents one of the most interesting benchmarks to evaluate the performance of the optimizer on structured problems. Moreover, the encoding of the SK model on the D-Wave hardware has very interesting elegant symmetry properties, allowing us to investigate general procedures common to all structured optimizations on annealers, such as the parameter setting of embedding and error correction, which in the general case require heuristic numerical pre- or postprocessing.

II. THE MODEL

The D-Wave TwoTM Vesuvius chip hosted at NASA Ames Research Center features 509 working flux qubits connected by 1455 tunable composite qubits acting as Ising-interaction couplings [24], arranged in a nonplanar lattice known as a Chimera graph [25]. In order to implement general Hamiltonians that are defined on arbitrary graphs, it is customary to employ the graph-minor embedding [26] technique. This procedure consists of finding a set of connected subgraphs [logical bits (LBs), corresponding to different colors in Fig. 1] of the original graph such that each LB can be associated to a node in the original graph. This association needs to be such that for each two connected nodes there exists at least one edge between the qubits belonging to the associated LBs. While the problem of finding an optimal graph minor (i.e., minimizing the number of required nodes) is itself NP hard [27] and is typically tackled with heuristic approaches [28], for many graphs with a regular structure an efficient embedding can be found systematically.

Figure 1 shows an embedding of the SK model in a triangular portion [29,30] of the Vesuvius processor: each LB in the original problem of size N is represented by $(N/4) + 1$ qubits connected in a line. This means that this embedding procedure encompasses an overhead of $N^2/4 + N$ hardware qubits for encoding fully connected graphs. Note that a quadratic scaling of the embedding resources for the SK model is expected for any hardware graph with fixed degree. The embedding procedure is useful for the encoding of the problem Hamiltonian into the hardware processor as long as the qubits in each LB are collapsed on the same z value at the end of the annealing. The basic idea is to ferromagnetically couple all qubits with a negative weight J_F within a LB in such a way as to energetically penalize discordant qubit states. The remaining couplings

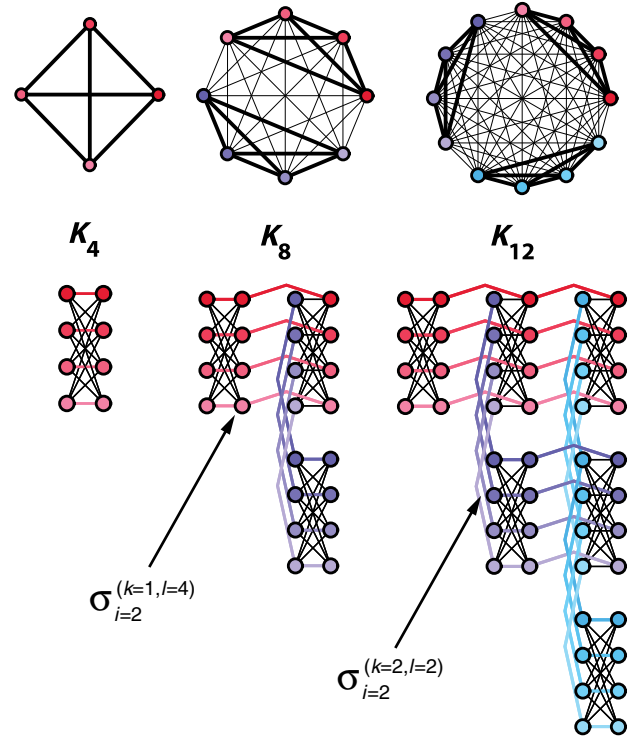


FIG. 1. Illustration of the iterative embedding procedure of the SK model in the Chimera graph. Different colors represent the N logical bits, which are arranged in $N/4$ groups of colors (reds, violets, and cyans, indexed by k). The corresponding images of fully connected graphs on top show that logical bits in the same group of colors have two different ways to be connected by a physical coupling on the Chimera graph by having a thicker edge between them. The arrows indicate two qubits with their respective indices convention with reference to the labeling of Eq. (2), where l indicates position within a given color group and i is a running label of the position of the qubit in the chain starting from the first (top left) $i = 1$ to the last $i = N/4 + 1$ (bottom right).

can be assigned to reflect the logical Hamiltonian of the problem to be solved.

With reference to the index convention illustrated in Fig. 1, the actual Hamiltonian that will be subjected to annealing is then

$$H_{\text{SK}} = \left(A(t) \sum_i S_i^X - B(t) \sum_{ij} J_{ij} S_i^Z S_j^Z \right) \quad (1)$$

$$\begin{aligned} \Rightarrow & A(t) \sum_{kli} \sigma_i^{x(kl)} - B(t) \sum_{kli} \left[J_F \sigma_i^{z(kl)} \sigma_{i+1}^{z(kl)} \right. \\ & + \delta_{ki} \sum_{l' < l} J_{(kl, k'l')} (\sigma_k^{z(kl)} \sigma_{k+1}^{z(k'l')}) \\ & \left. + \sum_{k'=k+i} \sum_{l'} J_{(kl, k'l')} (\sigma_{k+1+i}^{z(kl)} \sigma_k^{z(k'l')}) \right], \quad (2) \end{aligned}$$

where S_i^* and $\sigma_i^{*(kl)}$ are, respectively, the LB and the Pauli operators corresponding to the qubits along the $*$ direction. Since the maximum allowed energy coupling in the D-Wave Hamiltonian is 1, increasing J_F is equivalent to rescaling the logical couplings by dividing $B(t)$ by $|J_F|$ in Eq. (2). In the actual Hamiltonian of the annealing machine, the logical SK model couplings J_{ij} have been explicitly randomly chosen among the intercell couplings $J_{(kl,k'l')}$ and the couplings between different groups of colors $J_{(kl,k'l')}$, and the bounds on the summed variables are implied. $A(t)$ and $B(t)$ are the time-dependent coefficients that define the annealing schedule performed by the machine [31]. It is immediately apparent from the dynamical perspective that the optimal prescription for the value of J_F might be tricky to evaluate despite the fact that it is always possible to set its magnitude to be sufficiently high to make sure that the target ground state still lies at the bottom of the embedded classical spectrum [32].

Figure 2 shows the median probability $P_{g.s.}$ for the analog optimizer (run at fixed annealing time $\tau = 20 \mu s$) to reach the ground state. For a given problem size N , it depends significantly on J_F and goes to zero for large and small values of J_F . For $J_F \approx 1$, the ferromagnetic couplings are not energetically stronger than the logical couplings, and we expect that the problem is not well encoded. Indeed, many chains representing LBs are found in excited states (i.e., having 1 or more kinks), as illustrated by the colored bands of the plot which display the improvement on $P_{g.s.}$ obtained by a postprocessing procedure that tries to recover logical states from broken chains by doing majority voting (similarly to error-correcting repetition codes [33,34]). Conversely, for sufficiently large J_F , defects in the LBs

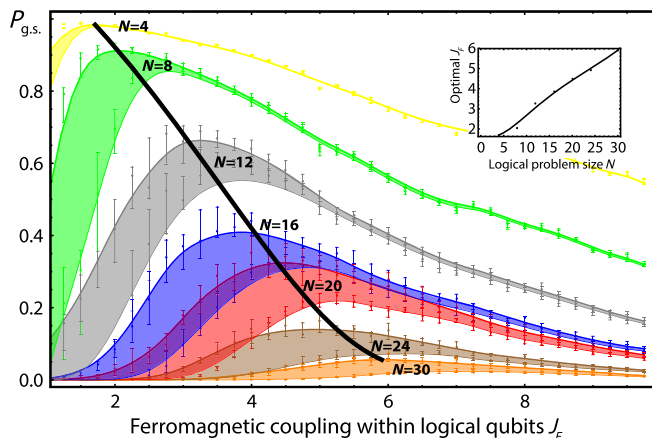


FIG. 2. Bezier fit of median probabilities of finding the ground state $P_{g.s.}$ of the encoded SK model (independently checked with exact enumeration code) after 400 000 runs on the D-Wave Two™ machine for every problem instance. The median is taken over 80 instances per size, and runs are performed using 10 random gauges. The black line and the inset indicate the optimal J_F for a given size, which increases with N as a power law close to \sqrt{N} . Error bars are obtained through resampling.

are suppressed, but the overall annealing success probability decreases after an optimal J_F . The appearance of this maximum can be connected to the expectation that the annealing dynamics is more efficient when the ferromagnetic LBs become correlated at the same time that the described SK model enters the spin-glass phase. This is because once the chains feel the ferromagnetic fixed point [for a transverse field of $A(t) \approx B(t)$], their dynamics slows down and might reasonably impede the adiabatic following of the logical ground state, while in the paramagnetic state they are more easily subjected to the formation of kinks. This argument is also supported by the scaling of the optimal coupling, which can be fit as a power law with an exponent close to $1/2$, which is consistent with the critical transverse field of the embedded SK model, which goes proportionally to $B(t)\sqrt{N}/J_F$ [35] (connected to Fig. 4 described later on). Comparisons with embedding and runs on embedded 2D lattices also support the above theory. However, other explanations that do not rely on quantum mechanics (e.g., the breaking of the ferromagnetic chains) are also plausible. [Details can be found in Supplemental Material [36].]

III. COMPARISON WITH SIMULATED ANNEALING

Figure 3 shows the median expected run time (in seconds) T_{run} for the annealing device to find the ground state with 99% probability, for different J_F and the experimentally shortest possible $\tau = 20 \mu s$. The thicker

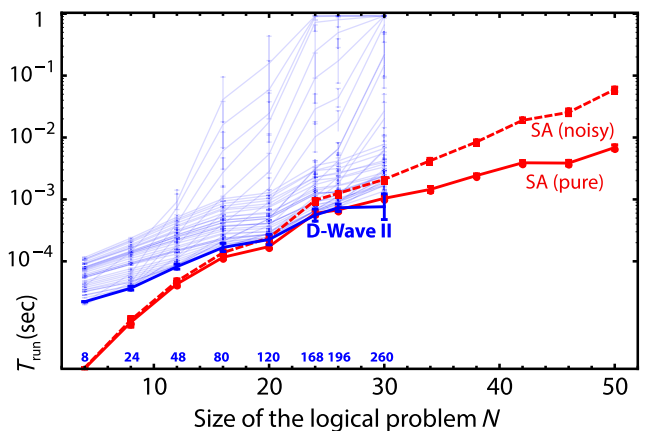


FIG. 3. Blue curves: Each curve is the expected D-Wave median run time for achieving 99% probability of finding the ground state ($P_{g.s.}$), computed as $T_{run} = \tau R$, with R being the expected number of repeated annealing cycles $R = \log(0.01)/\log(1 - P_{g.s.})$, at fixed J_F . The result shown includes the error-correction procedure, whose processing time is not considered. The best possible result irrespective of J_F for each problem size is highlighted in the thicker blue line. Red curves: Simulated annealing results for the bare logical problem (solid line) and with the introduction with the respective noise model of the D-Wave machine described in text (dashed line).

blue line shows the scaling of complexity with the system size assuming the optimization on J_F with a precision of $\Delta J = 0.25$. This exponential scaling and the absolute run time seem very similar (and correlates well, see Supplemental Material [36]) to the performance of simulated annealing (SA) on the same *logical* instance set extended up to $N = 50$ whose run time is equivalently optimized over τ measured on Intel Xeon E5-2680v2 processors. However, it is well known that the Hamiltonian parameters programmed on the analog optimizer are subject to static and low-frequency noise that can be modeled as a static Gaussian disorder realization for each instance [37]. One could argue that this noise (whose presence is not fundamental but rather an engineering issue) introduces an artificial handicap in the evaluation of the performance of the D-Wave machine, as the programmed problem might significantly differ from the target objective function to be minimized (the robustness of a problem to these perturbations is sometimes called “resilience”). We introduce the noise effect in the logical instance runs with SA in order to make a fair comparison. While on the scale of the maximum physical energy programmed in the problem Hamiltonian (i.e., 3.2 GHz) this model of noise has a negligible effect [37], the rescaling of the absolute energy of the logical parameters due to the introduction of J_F proportionally amplifies the relevance of the unwanted disorders. The considered noise model spoils the $J_{(kl,k'l')}$ couplings of Eq. (2) and introduces artificial longitudinal local fields. Note that there is also an instance-dependent artificial short-range coupling between the qubits due to cross talks (background susceptibility [12]), but this contribution is in principle correctable by preprocessing and should be negligible in the classical limit. More specifically, as the logical couplings J_{ij} are chosen to be ± 1 , this implies that the problem Hamiltonian to be compared with D-Wave TwoTM runs at fixed J_F must be spoiled as follows:

$$H_{\text{dev}} = \frac{H_{\text{SK}}}{|J_F|} + \left[\sum_{ij} \xi_J^{ij} S_i S_j + \sum_i \xi_h^i S_i \right], \quad (3)$$

where ξ_J^{ij} and ξ_h^i are disorder realizations with Gaussian distribution around zero of respective standard deviations $\sigma_J^\xi = 0.035$ and $\sigma_h^\xi = 0.05$ [37]. Results are averaged over 1000 realizations for every instance, and new optimal speeds have been computed for the final scaling (see Supplemental Material [36]). What is observed is that, starting from $N = 12$, the noise significantly affects the probability for the spoiled system to find the ground state of the ideal Hamiltonian. As detailed in the Supplemental Material [36], for every fixed level of noise proportional to J_F , there is indeed a problem size above which the noise tends to shift the ground state of the noisy Hamiltonian outside the manifold of the degenerate ground states of the ideal Hamiltonian, independently from the algorithm used

to compute the ground state. We note that this effect is likely to be dominant over the slow-down of the LB dynamics conjectured to be responsible for the sharp decrease in performance of the device for large J_F observed in Fig. 2, and more analysis is needed to establish if this is the case.

Unsurprisingly, the current limitations on the number of qubits do not allow us to draw final conclusions on whether the machine has a sound limited speedup with respect to simulated annealing; the scaling results are encouraging. While it is now established that speedup might emerge artificially due to suboptimal annealing speed [14] ($\tau = 20 \mu\text{s}$ would supposedly become optimal only for larger N), as well as due to correlation between different subsequent runs [11], we show evidence that this is likely to be masked by the detrimental effect of the noise (which is expected to be significantly reduced in future generations of the device). Most importantly, our work elucidates how evaluating the comparative performance of analog optimization with respect to algorithmic methods on necessarily embedded problems is more delicate than it is on natively structured problems. This is largely because the correct representation of the target problem requires an optimal tuning of the analog optimizer, which is dependent on the hardware architecture and the programmability precision. The benchmarks [14,38] on natively structured problems need to be complemented by considering that in embedded problems the number of LBs does not reflect the number of qubits for the comparison of required resources. Moreover, the fact that the logical Hamiltonian is emergent from a coarse graining of the hardware Hamiltonian compounded by ferromagnetic chains has potentially profound consequences regarding the expected complexity of the annealing procedure on the logical problem. In the SK model this also means that the shape and the location of the critical region associated with the spin-glass phase is dependent on the internal representation parameters, such as embedding topology and optimal J_F .

IV. SPIN-GLASS PHASE OF THE EMBEDDED PROBLEM

In order to gain insights on these issues, as shown in Fig. 4, we examine, by means of Monte Carlo simulations, the emergence of the spin-glass phase of the embedded SK model [Eq. (2)], i.e., the appearance of a critical (normalized) spin-glass temperature T_{SG} as a function of $\alpha_F = J_F/\sqrt{N}$. We compute the universal spin-configuration overlap Binder ratio g [38] for different α_F , T , then estimates of the critical temperature are obtained by looking at the intersection of Binder ratios for a series of different sizes and then extrapolating the results using a scaling ansatz [39,40]. To improve equilibration, we develop a variant of a Monte Carlo procedure that combines the best properties of the Wolff cluster method and the Metropolis update. Parallel tempering




FIG. 4. The estimate of the critical temperature is obtained from finite-size scaling of intersections of Binder ratio curves $g(T) = (1/2)[3 - \langle q(T)^4 \rangle / \langle q(T)^2 \rangle^2]$ related to spin-configuration overlaps $q = (1/N) \sum_i S_i^A S_i^B$ between two replicated runs A and B at various temperatures T for each α_F in the figure. The red band around data points corresponds to the error over the intersecting region, and the dashed black line to the spin-glass temperature of the logical problem. See Supplemental Material [36] for more details.

is used to enhance it further by performing replica swaps in a two-dimensional space (in T and α_F/T). Our findings are compatible with an increase of T_{SG} with α_F towards the theoretical value of the unembedded model $1/\alpha_F$. This means that embedded problems are likely to belong to a different universality class than random Chimera problems, answering a question at the center of the current discussion in the quantum annealing community [38].

These results support the intuition that the ferromagnetic couplings need to increase as \sqrt{N} (up to logarithmic corrections) in order to properly represent the SK model for large sizes. Interestingly, the experimentally optimal α_F (see Fig. 3) in our runs is close to 1.0, meaning that the machine performs better optimizing the spectrum of an embedded representation of the SK model whose T_{SG} is significantly less than that of the logical problem, which is nevertheless faithfully represented by the embedding. Still, the critical temperature of the embedded problem at an optimal parameter setting is larger than the experimental one of the D-Wave chip, which might have profound consequences on the asymptotic computational complexity of quantum annealing on the embedded SK model.

Details on the numerics for the calculation of the classical phase diagram for the embedded SK model are provided in the following section.

V. NUMERICAL METHODS

A. Equilibration of the embedded SK Hamiltonian

The embedding of the SK model creates long ferromagnetic chains of physical spins, which correspond to a single

logical spin. When the intrachain ferromagnetic coupling J_F becomes comparatively very large, the whole chain behaves like a true logical two-level system and then we expect the embedded SK model to show the same thermodynamic properties of the logical SK model. Unfortunately, due to the presence of these long ferromagnetic chains, equilibration of the embedded SK Hamiltonian happens to be extremely long when the standard single spin-flip Metropolis-Harris procedure is used to perform Monte Carlo simulations. Indeed, for low temperature and $J_F \gg |J_{ij}|$, where J_{ij} are the couplings of the original logical SK model, spins belonging to the same chain prefer to stay aligned. Since the probability of creating a defect in a polarized chain is proportional to $\exp(-2\beta J_F)$, where $\beta = 1/T$ is the inverse of the temperature, a large part of the equilibration time is spent trying to flip a whole chain. In order to reach equilibration more quickly for the embedded SK model, we propose a variant of the Wolff cluster method [41] (a generalization of the original Swendsen-Wang cluster method [42]), which takes into account the existence of the logical superstructures.

As described in Refs. [41,42], the Wolff cluster method is effective for low temperatures, $\beta J \gg 1$, where the usual Metropolis rule has very low acceptance probability. One of the main characteristics of the cluster method is that it always flips clusters. Therefore, it works well in the presence of many domain walls. However, although it works well for low-dimensional systems, for highly connected models, clusters tend to include almost *all* the spins, so that most of the time the cluster algorithm flips the largest part of the spins back and forth, failing to equilibrate. To overcome this limitation, we design a hybrid method that combines the best aspects of the Metropolis and the Wolff methods. Because the equilibration within ferromagnetic chains is the true limit of the Metropolis update, the thermalization of the ferromagnetic chains is obtained by using the Wolff update. More specifically, the basic idea of our hybrid method is to grow a Wolff cluster only within chains and to flip those clusters by using a Metropolis update. In contrast to the standard procedure that always flips the clusters, clusters are flipped by using a probability given by the Metropolis rule. Since couplings inside the chains are already taken into account by creating the Wolff clusters, the Metropolis probability is therefore computed by using only couplings between chains. Domain wall energies inside the ferromagnetic chains are excluded: indeed, those energy have already been considered as the cluster is built.

More formally, we rewrite the embedding Hamiltonian for the SK model in the following form:

$$H(\boldsymbol{\sigma}) = - \underbrace{\sum_{\langle \alpha, \alpha' \rangle} J_{\alpha\alpha'} \sigma_{\alpha, i_{\alpha\alpha'}} \sigma_{\alpha', i_{\alpha'\alpha}}}_{H_L(\boldsymbol{\sigma})} - J_F \underbrace{\sum_{\alpha, i} \sigma_{\alpha, i} \sigma_{\alpha, i+1}}_{H_F(\boldsymbol{\sigma})}, \quad (4)$$

where the chains (which correspond to the logical spins) and the spins within each chain are indexed, respectively, by α and i . Here, $J_{\alpha\alpha'} = \pm 1$ corresponds to the original couplings of the logical model. The index $i_{\alpha\alpha'}$ is a shorthand notation to indicate the spin in chain α which is interacting with the spin in chain α' . The first term of $H(\boldsymbol{\sigma})$ describes the interaction between chains while the second term describes the ferromagnetic couplings within the chains. Given the Hamiltonian in Eq. (4), the entire procedure can be described as follows.

- (1) Choose a spin with indices (α, i) , where α labels the chain (logical spin) and i the spin's index within a chain. This can be done at random or in nested loops, with the outermost one being over i .
- (2) Grow the cluster \mathcal{C} (starting from the spin above) along the ferromagnetic chains in one direction: add spins one by one, stopping with probability $\exp(-2\beta J_F)$ (or probability 1 if the new spin fails to have the same value as $\sigma_{\alpha,i}$).
- (3) Repeat (growing the cluster) in the other direction along the chain.
- (4) Flip the cluster with probability

$$p_{\text{acc}}(\boldsymbol{\sigma}, \mathcal{C}) = \min \left\{ 1, \exp \left(-2\beta \sigma_{\alpha,i} \sum_{\alpha' | i_{\alpha\alpha'} \in \mathcal{C}} J_{(\alpha, i_{\alpha\alpha'})} (\alpha', i_{\alpha'\alpha}) \sigma_{\alpha', i_{\alpha'\alpha}} \right) \right\}. \quad (5)$$

The last expression is similar to the Metropolis formula, but, in this case, only interactions between different chains are taken into account. ‘‘Logical’’ coupling between chains α and α' is via bond-connecting spins with indices $(\alpha, i_{\alpha\alpha'})$ and $(\alpha', i_{\alpha'\alpha})$. The sum in the exponent is only over chains ‘‘interacting’’ with a cluster. Note that $\sigma_{\alpha,i}$ has been taken outside the sum as all spins within the cluster must have the same value by construction. As better detailed in the Supplemental Material [36], it is straightforward to show that the above procedure still satisfies the detailed balance; indeed, the limitation that a cluster can grow only inside a logical spin (chain) is balanced by adding the probability $p_{\text{acc}}(\boldsymbol{\sigma}, \mathcal{C})$ of flipping the cluster, which involves only interchain couplings. Equilibration analyses are reported in Supplemental Material [36].

B. Replica exchange Monte Carlo method in ‘‘two dimensions’’

The standard replica exchange method [43] (more commonly known as parallel tempering) improves equilibration by simultaneously simulating different replicas of the system, each of them at a different temperature. More specifically, the replica exchange method is based on ‘‘local’’ updates, namely, single-spin Metropolis updates of each replica, and ‘‘global’’ updates, namely, swaps between two replicas at different temperatures. We implement a natural

extension of this approach by considering a two-dimensional parameter space. It is instructive to write the embedded SK Hamiltonian [see Eq. (2)] in the following form:

$$-\beta H(\boldsymbol{\sigma}) = K_L \underbrace{\sum_{\langle \alpha, \alpha' \rangle} J_{\alpha\alpha'} \sigma_{\alpha, i_{\alpha\alpha'}} \sigma_{\alpha', i_{\alpha'\alpha}}}_{H_L(\boldsymbol{\sigma})} + K_F \underbrace{\sum_{\alpha, i} \sigma_{\alpha, i} \sigma_{\alpha, i+1}}_{H_F(\boldsymbol{\sigma})}, \quad (6)$$

where the notation is the same as used in Eq. (4).

The ‘‘dimensionless’’ coupling constants are $K_L \equiv \beta$ and $K_F = \beta J_F$. The convention of using integer values $J_{\alpha\alpha'} = \pm 1$ for the logical couplings implies that the energy of the logical Hamiltonian scales like $E \sim N^{3/2}$. In the limit of large J_F , when the embedding is traced out and the thermodynamics of the whole system corresponds to the thermodynamics of the logical system, the relevant inverse temperature must scale as $\beta = \beta' / \sqrt{N}$, with $\beta' > 0$. Consequently, to have the correct scaling of the embedding Hamiltonian in Eq. (6) in the paramagnetic phase, J_F must properly scale as $J_F = \alpha_F \sqrt{N}$, where $\alpha_F > 0$. Observe that the scalings of β and J_F are consistent with the choice of the grid; indeed, at the critical temperature $\beta_{\text{SG}} = 1/T_{\text{SG}} = 1/\sqrt{N}$, $K_F = \beta_{\text{SG}} J_F = \alpha_F$ maintains a finite value, independently of the system size.

Replica swaps are attempted between nearest neighbors on a square lattice spanned by (K_L, K_F) and are accepted with probability

$$p(K_L \leftrightarrow K'_L) = \min \{ 1, e^{(K_L - K'_L)[H_L(\boldsymbol{\sigma}) - H_L(\boldsymbol{\sigma}')] } \} \quad (7a)$$

or

$$p(K_F \leftrightarrow K'_F) = \min \{ 1, e^{(K_F - K'_F)[H_F(\boldsymbol{\sigma}) - H_F(\boldsymbol{\sigma}')] } \}, \quad (7b)$$

depending on the axis. Here, $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma}'$ represent spin configurations of the replicas. It is straightforward to verify that this maintains Gibbs equilibrium for the combined system.

Figure 5 shows an example of the intersection of two Binder ratios for the embedded SK model with, respectively, $N = 64$ and $N = 32$ logical spins $[T_{\text{SG}}^{64}(\alpha_F)]$. As expected, the logical SK model is recovered for sufficiently large α_F . Further details on the calculation of the finite-size critical temperatures by using Binder ratio crossings and the parameters we use in the simulations are presented in Supplemental Material [36].

C. Calculation of the critical temperature in the thermodynamic limit

The calculation of the critical temperature T_{SG} in the thermodynamic limit can be done by fitting the finite-size estimations of the critical temperature T_{SG}^N while increasing the system size N . More specifically, since T_{SG}^N scales like




FIG. 5. Example of the intersection between two Binder ratio surfaces with, respectively, $N = 32$ and $N = 64$. The intersecting manifold is $T_{SG}^{64}(\alpha_F)$ (see Supplemental Material [36]). For $\alpha_F \gg 1$, the logical SK model is recovered.

$$T_{SG}^N = T_{SG} + aN^{-1/\nu} + o(N^{-1/\nu}), \quad (8)$$

where ν is the critical exponent that identifies the universality class of the system, the analysis of the scaling of T_{SG}^N by varying N allows for the extrapolation of both T_{SG} and ν of the system (α_F dependence is implied). However, the expected scaling behavior as in Eq. (8) can be observed only for very large systems [44]. In particular, for the logical SK model, it is necessary to study systems with more than $N = 512$ spins to correctly identify both T_{SG} and ν . Therefore, for smaller system sizes, subleading terms cannot be neglected and the functional form of the scaling of T_{SG}^N must be properly adjusted. Unfortunately, the addition of subleading terms also introduces new fitting parameters, which makes the proper fit of both T_{SG} and ν hard to find.

To avoid the introduction of new fitting parameters, we incorporate the subleading terms as an effective critical exponent, leaving the critical temperature unchanged, namely, using the ansatz

$$T_{SG}^N = T_{SG} + a'N^{-b'} + o(N^{-b'}), \quad (9)$$

where $b' > 0$. The above approximation, which converges to the exact extrapolation for large N , allows the calculation of the spin-glass critical temperature T_{SG} without introducing any further fitting parameter. Scalings of T_{SG}^N by varying the system size can be found in Supplemental Material [36].

VI. CONCLUSIONS

To summarize, we show that random fully connected spin glasses are solvable on Chimera-graph-based

annealers through the embedding procedure. The resulting Hamiltonian belongs to a different universality class with finite spin-glass temperature, unlike the random problems on the Chimera graph. Moreover, when the embedding parameters are scaled with \sqrt{N} , an optimal performance is achieved for the device we test. This prescription allows the D-Wave device to match the performance of optimized simulated annealing on the same problems, despite the presence of analog control errors. Should these precision errors be mitigated in future hardware, it is expected that the comparative performance of the quantum annealer with simulated annealing will improve significantly.

ACKNOWLEDGMENTS

We acknowledge useful discussions with Dr. Trevor Lanting, Dr. Andrew King, Dr. Alejandro Perdomo-Ortiz, Dr. Eleanor G. Rieffel, Professor M. Troyer, and Professor H. Katzgraber. Part of the work was supported by NASA (Sponsor Awards No. NNX14AF62G and No. NNX12A K33A). The numerical work was supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IAA 145483, and by the AFRL Information Directorate under Grant No. F4HBKC4162G001.

- [1] A. Lucas, *Ising Formulations of Many NP Problems*, *Front. Phys.* **2**, 5 (2014).
- [2] V. N. Smelyanskiy, E. G. Rieffel, S. I. Knysh, C. P. Williams, M. W. Johnson, M. C. Thom, W. G. Macready, and K. L. Pudenz, *A Near-Term Quantum Computing Approach for Hard Computational Problems in Space Exploration*, [arXiv:1204.2821](https://arxiv.org/abs/1204.2821).
- [3] G. E. Santoro, R. Martoňák, E. Tosatti, and R. Car, *Theory of Quantum Annealing of an Ising Spin Glass*, *Science* **295**, 2427 (2002).
- [4] T. Kadowaki and H. Nishimori, *Quantum Annealing in the Transverse Ising Model*, *Phys. Rev. E* **58**, 5355 (1998).
- [5] S. Morita and H. Nishimori, *Mathematical Foundation of Quantum Annealing*, *J. Math. Phys. (N.Y.)* **49**, 125210 (2008).
- [6] J. Brooke, D. Bitko, G. Aeppli *et al.*, *Quantum Annealing of a Disordered Magnet*, *Science* **284**, 779 (1999).
- [7] A. M. Zagoskin, E. Ilichev, M. Grajcar, J. J. Betouras, and F. Nori, *How to Test the “Quantumness” of a Quantum Computer?*, *Front. Phys.* **2**, 33 (2014).
- [8] P. I. Bunyk, E. Hoskinson, M. W. Johnson, E. Tolkacheva, F. Altomare, A. J. Berkley, R. Harris, J. P. Hilton, T. Lanting, and J. Whittaker, *Architectural Considerations in the Design of a Superconducting Quantum Annealing Processor*, *IEEE Trans. Appl. Supercond.* **24**, 1700110 (2014).
- [9] C. C. McGeoch and C. Wang, *Experimental Evaluation of an Adiabatic Quantum System for Combinatorial Optimization*, in *Proceedings of the ACM International Conference on Computing Frontiers, 2013* (Ischia, Italy, 2013), p. 23.
- [10] S. Dash, *A Note on QUBO Instances Defined on Chimera Graphs*, [arXiv:1306.1202v2](https://arxiv.org/abs/1306.1202v2).

- [11] S. Boixo, T. F. Rønnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer, *Evidence for Quantum Annealing with More Than One Hundred Qubits*, *Nat. Phys.* **10**, 218 (2014).
- [12] T. Albash, W. Vinci, A. Mishra, P. A. Warburton, and D. A. Lidar, *Consistency Tests of Classical and Quantum Models for a Quantum Annealer*, *Phys. Rev. A* **91**, 042314 (2015).
- [13] S. Boixo, V. N. Smelyanskiy, A. Shabani, S. V. Isakov, M. Dykman, V. S. Denchev, M. Amin, A. Smirnov, M. Mohseni, and H. Neven, *Computational Role of Collective Tunneling in a Quantum Annealer*, arXiv:1411.4036.
- [14] T. F. Rønnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer, *Defining and Detecting Quantum Speedup*, *Science* **345**, 420 (2014).
- [15] J. A. Smolin and G. Smith, *Classical Signature of Quantum Annealing*, arXiv:1305.4904.
- [16] A. Perdomo-Ortiz, J. Fluegemann, S. Narasimhan, V. Smelyanskiy, and R. Biswas, *A Quantum Annealing Approach for Fault Detection and Diagnosis of Graph-Based Systems* (to be published).
- [17] E. G. Rieffel, D. Venturelli, B. O’Gorman, M. B. Do, E. M. Prystay, and V. N. Smelyanskiy, *A Case Study in Programming a Quantum Annealer for Hard Operational Planning Problems*, *Quantum Inf. Process.* **14**, 1 (2015).
- [18] B. O’Gorman, R. Babbush, A. Perdomo-Ortiz, A. Aspuru-Guzik, and V. Smelyanskiy, *Bayesian Network Structure Learning Using Quantum Annealing*, *Eur. Phys. J Special Topics* **225**, 163 (2015).
- [19] A. Perdomo-Ortiz, N. Dickson, M. Drew-Brook, G. Rose, and A. Aspuru-Guzik, *Finding Low-Energy Conformations of Lattice Protein Models by Quantum Annealing*, *Sci. Rep.* **2**, 571 (2012).
- [20] Z. Bian, F. Chudak, W. G. Macready, L. Clark, and F. Gaitan, *Experimental Determination of Ramsey Numbers*, *Phys. Rev. Lett.* **111**, 130505 (2013).
- [21] S. Santra, G. Quiroz, G. Ver Steeg, and D. A. Lidar, *MAX 2-SAT with Up to 108 Qubits*, *New J. Phys.* **16**, 045006 (2014).
- [22] D. Sherrington and S. Kirkpatrick, *Solvable Model of a Spin-Glass*, *Phys. Rev. Lett.* **35**, 1792 (1975).
- [23] Y. Fu and P. W. Anderson, *Application of Statistical Mechanics to NP-Complete Problems in Combinatorial Optimisation*, *J. Phys. A* **19**, 1605 (1986).
- [24] R. Harris, M. W. Johnson, T. Lanting, A. J. Berkley, J. Johansson, P. Bunyk, E. Tolkacheva, E. Ladizinsky, N. Ladizinsky, T. Oh *et al.*, *Experimental Investigation of an Eight-Qubit Unit Cell in a Superconducting Optimization Processor*, *Phys. Rev. B* **82**, 024511 (2010).
- [25] M. W. Johnson, P. Bunyk, F. Maibaum, E. Tolkacheva, A. J. Berkley, E. M. Chapple, R. Harris, J. Johansson, T. Lanting, I. Perminov *et al.*, *A Scalable Control System for a Superconducting Adiabatic Quantum Optimization Processor*, *Supercond. Sci. Technol.* **23**, 065004 (2010).
- [26] W. M. Kaminsky and S. Lloyd, *Quantum Computing and Quantum Bits in Mesoscopic Systems* (Springer, New York, 2004), pp. 229–236.
- [27] D. Eppstein, *Finding Large Clique Minors is Hard*, *J. Graph Algorithms Appl.* **13**, 197 (2009).
- [28] J. Cai, B. Macready, and A. Roy, *A Practical Heuristic for Finding Graph Minors*, arXiv:1406.2741.
- [29] V. Choi, *Minor-Embedding in Adiabatic Quantum Computation: II. Minor-Universal Graph Design*, *Quantum Inf. Process.* **10**, 343 (2011).
- [30] C. Klymko, B. D. Sullivan, and T. S. Humble, *Adiabatic Quantum Programming: Minor Embedding with Hard Faults*, *Quantum Inf. Process.* **13**, 709 (2014).
- [31] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk *et al.*, *Quantum Annealing with Manufactured Spins*, *Nature (London)* **473**, 194 (2011).
- [32] V. Choi, *Minor-Embedding in Adiabatic Quantum Computation: I. The Parameter Setting Problem*, *Quantum Inf. Process.* **7**, 193 (2008).
- [33] K. L. Pudenz, T. Albash, and D. A. Lidar, *Error-Corrected Quantum Annealing with Hundreds of Qubits*, *Nat. Commun.* **5**, 3243 (2014).
- [34] K. C. Young, R. Blume-Kohout, and D. A. Lidar, *Adiabatic Quantum Optimization with the Wrong Hamiltonian*, *Phys. Rev. A* **88**, 062314 (2013).
- [35] S. Suzuki, J.-i. Inoue, and B. K. Chakrabarti, *Quantum Ising Phases and Transitions in Transverse Ising Models* (2013).
- [36] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevX.5.031040> for further details on numerics and additional benchmarkings.
- [37] A. D. King and C. C. McGeoch, *Algorithm Engineering for a Quantum Annealing Platform*, arXiv:1410.2628v2.
- [38] H. G. Katzgraber, F. Hamze, and R. S. Andrist, *Glassy Chimeras Could be Blind to Quantum Speedup: Designing Better Benchmarks for Quantum Annealing Machines*, *Phys. Rev. X* **4**, 021008 (2014).
- [39] H. G. Katzgraber, *Introduction to Monte Carlo Methods*, arXiv:0905.1629.
- [40] R. N. Bhatt and A. P. Young, *Search for a Transition in the Three-Dimensional J Ising Spin-Glass*, *Phys. Rev. Lett.* **54**, 924 (1985).
- [41] U. Wolff, *Collective Monte Carlo Updating for Spin Systems*, *Phys. Rev. Lett.* **62**, 361 (1989).
- [42] J.-S. Wang and R. H. Swendsen, *Cluster Monte Carlo Algorithms*, *Physica (Amsterdam)* **167A**, 565 (1990).
- [43] K. Hukushima and K. Nemoto, *Exchange Monte Carlo Method and Application to Spin Glass Simulations*, *J. Phys. Soc. Jpn.* **65**, 1604 (1996).
- [44] A. Billoire, L. A. Fernandez, A. Maiorano, E. Marinari, V. Martin-Mayor, and D. Yllanes, *Finite-Size Scaling Analysis of the Distributions of Pseudo-Critical Temperatures in Spin Glasses*, *J. Stat. Mech.* (2011) P10019.