



# CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

## Quantum annealing correction with minor embedding

Walter Vinci, Tameem Albash, Gerardo Paz-Silva, Itay Hen, and Daniel A. Lidar

Phys. Rev. A **92**, 042310 — Published 9 October 2015

DOI: [10.1103/PhysRevA.92.042310](https://doi.org/10.1103/PhysRevA.92.042310)

# Quantum Annealing Correction with Minor Embedding

Walter Vinci,<sup>1,2,3</sup> Tameem Albash,<sup>2,3,4</sup> Gerardo Paz-Silva,<sup>5,6</sup> Itay Hen,<sup>3,4</sup> and Daniel A. Lidar<sup>1,2,3,7</sup>

<sup>1</sup>*Department of Electrical Engineering, University of Southern California, Los Angeles, California 90089, USA*

<sup>2</sup>*Department of Physics and Astronomy, University of Southern California, Los Angeles, California 90089, USA*

<sup>3</sup>*Center for Quantum Information Science & Technology,*

*University of Southern California, Los Angeles, California 90089, USA*

<sup>4</sup>*Information Sciences Institute, University of Southern California, Marina del Rey, California 90292, USA*

<sup>5</sup>*Centre for Quantum Computation and Communication Technology (Australian Research Council),*

*Griffith University, Brisbane, Queensland 4111, Australia*

<sup>6</sup>*Centre for Quantum Dynamics, Griffith University, Brisbane, Queensland 4111, Australia*

<sup>7</sup>*Department of Chemistry, University of Southern California, Los Angeles, California 90089, USA*

Quantum annealing provides a promising route for the development of quantum optimization devices, but the usefulness of such devices will be limited in part by the range of implementable problems as dictated by hardware constraints. To overcome constraints imposed by restricted connectivity between qubits, a larger set of interactions can be approximated using minor embedding techniques whereby several physical qubits are used to represent a single logical qubit. However, minor embedding introduces new types of errors due to its approximate nature. We introduce and study quantum annealing correction schemes designed to improve the performance of quantum annealers in conjunction with minor embedding, thus leading to a hybrid scheme defined over an encoded graph. We argue that this scheme can be efficiently decoded using an energy minimization technique provided the density of errors does not exceed the per-site percolation threshold of the encoded graph. We test the hybrid scheme using a D-Wave Two processor on problems for which the encoded graph is a 2-level grid and the Ising model is known to be NP-hard. The problems we consider are frustrated Ising model problem instances with “planted” (a priori known) solutions. Applied in conjunction with optimized energy penalties and decoding techniques, we find that this approach enables the quantum annealer to solve minor embedded instances with significantly higher success probability than it would without error correction. Our work demonstrates that quantum annealing correction can and should be used to improve the robustness of quantum annealing not only for natively embeddable problems, but also when minor embedding is used to extend the connectivity of physical devices.

## I. INTRODUCTION

Quantum computers are in principle able to solve computational problems that are intractable for their classical counterparts [1, 2]. Remarkably, rudimentary quantum annealers—a special type of adiabatic quantum computer [3] specifically engineered to implement quantum annealing [4–9] to solve hard optimization problems [10]—are already commercially available [11–13], and their availability to the scientific community has recently stimulated a productive debate [14–24]. The question at stake is whether current experimental quantum annealing can really be considered a form of quantum computation, even when the decoherence time of the component qubits is much smaller than the overall computational time. While theory suggests that the answer to this question can be positive [25, 26], as it concerns quantum annealing experiments the answer is still open, in particular due to the “black-box” nature of quantum annealing, which makes it difficult to unambiguously determine the nature of the physical processes taking place between the initialization and readout steps. However, convincing evidence is mounting that quantum effects, including entanglement [27] and multi-qubit coherent tunneling [28], play a functional role in explaining the behavior of available annealers. Moreover, open system quantum dynamical models [28, 29] have been successful in reproducing experimental data from quantum annealing experiments [16, 21, 22, 28, 30]. A closely related question is whether quantum annealers are able to provide quantum speedup, a theoretical possibility [31] despite skepticism

based, e.g., on the appearance of exponentially small gaps [32]. While compelling evidence of a quantum speedup is still lacking with current physical annealers, it is also generally recognized that technological improvements (in particular reduced control errors and shorter annealing times) in the next generation of quantum annealers are necessary in order to unambiguously address the issue [33–37]. A careful choice of the benchmark problems is another important consideration [38, 39], in particular in light of the fragility of spin glasses to small perturbations (temperature chaos) [40, 41].

Ultimately, the usefulness of quantum annealers relies on the promise of achieving quantum information processing on a large scale. This goal can only be met through the use of some form of quantum error correction (QEC), which is necessary to protect the performance of any quantum computation against decoherence [42–44]. The scalability of quantum computing via QEC has been established in the gate model in the form of the celebrated threshold theorem (e.g, Refs. [45, 46]). Adiabatic quantum computation and quantum annealing also require some form of error correction in order to preserve their advantages over classical computation. This is true even though adiabatic dynamics is intrinsically robust to some forms of decoherence [47–49]. Several error suppression techniques suitable for adiabatic quantum computing (AQC) have been proposed [50–54]. However, despite valiant theoretical attempts [55] much less is known about how to achieve fault tolerant AQC, and some negative results have been established [56–58]. Still, it has recently been shown experimentally that it is possible to improve the performance of

physical quantum annealers with the help of tailored quantum annealing correction [59–61], a method we shall employ and generalize in this work.

Typically, realistic implementations of quantum annealing are subject to restricted (local) interactions between qubits specified by the physical hardware. Minor embedding (ME) techniques overcome this limitation by using several physical qubits to represent a single logical qubit with a larger set of interactions [62–65]. A minor embedding replaces a given optimization problem with an approximate yet implementable one. As a consequence of this approximation the performance of a quantum annealer is typically degraded. ME also adds two extra steps to the computations performed with a quantum annealer. One is an encoding procedure that determines the mapping between the logical problem and its physically embeddable representation, which in turn requires a choice of energy penalties to align the physical qubits of each logical qubit. It turns out that performance is sensitive to the mapping chosen and it is thus important to optimize it. The other is a decoding procedure that recovers the logical answer from the physical output. Decoding is nontrivial whenever the values of the physical qubits corresponding to the same logical qubit disagree, and will be discussed in detail in this work.

Quantum annealing correction (QAC) was proposed in Ref. [59], where it was demonstrated on anti-ferromagnetic chains and further discussed for random Ising problems in Ref. [60]. These works showed that QAC can improve the performance of quantum annealers. Implementing several independent copies of the same problem enables errors to be corrected using, e.g., majority vote as in classical repetition codes. Thus, importantly, the requirement of a purely adiabatic evolution can be relaxed since excited states can be tolerated as long as they can be correctly decoded. Quantum error suppression is obtained with the use of ferromagnetic energy penalties connecting the various copies, that are turned on gradually along with the final Hamiltonian. Since these can also be understood as the stabilizers of the quantum repetition code, that detect and penalize bit-flip errors during the course of the evolution, agreement is enhanced within each copy by energetically penalizing errors that would result in disagreement. The combined effect of encoding into logical qubits and quantum error suppression can be seen as an increase of the effective energy scale with which the logical problem is physically implemented. This increases the relevant minimum energy gaps that protect quantum annealing from thermal and dynamical excitations. Moreover, it improves the response of quantum annealing devices to intrinsic control errors (due to the limited accuracy with which the physical fields can be tuned) by using several physical couplings to specify a single logical interaction [59–61].

The main goal of this work is to develop and study QAC techniques for improving the performance of a quantum annealer following the expected degradation due to the implementation of ME techniques. In doing so, we explore the efficacy and applicability of several encoding and decoding strategies. Our tests are performed using a D-Wave Two (DW2) “Vesuvius” quantum annealer with 504 functional qubits.

Our paper is organized as follows. In section II, after a brief review of quantum annealing, we give an overview of ME and QAC. We explain how ME can be viewed as a map from a given logical problem to the layout dictated by the physical implementation, while QAC reverses this direction and maps from the physical level to an encoding that provides protection. Thus, while different in scope, the two approaches can be seamlessly concatenated in what we call quantum annealing correction with minor embedding (QAC-ME). We also discuss the different role of frustration effects on ME and QAC. In Sec. III we discuss different strategies for choosing the energy penalties in both ME and QAC, and with an eye to optimality, introduce a nonuniform strategy. In Sec. IV we discuss decoding strategies. These are required in order to assign a value to a logical or encoded qubit when the corresponding physical qubits have different values. A majority vote on the physical qubits is a simple choice but that may not always be justified. We thus propose energy minimization as a more effective decoding strategy and provide a general argument based on percolation theory to determine whether it can be efficiently implemented. A schematic representation of our implementation is summarized in Fig. 6. Section V introduces the “square code,” a QAC-ME scheme defined on a 2-level grid that can be implemented with the physical connectivity (“Chimera” graph) of the DW2 processor. The Ising problem on the 2-level grid is of particular interest since it was the first example of an NP-hard Ising model problem [66]. The square code is of independent interest since it is amenable to concatenation. Section VI presents our experimental results, obtained using the DW2 processor. The problem instances we study are constructed using the planted ground states method for frustrated loops introduced in Ref. [34]. This method allows for a tunable hardness parameter in the form of the effective clause density. The experimental results demonstrate the significant performance boost that can be achieved using the QAC-ME strategy, by exhibiting higher ground state success probabilities, with an improvement that is more pronounced for the harder problem instances. Section VII reports a detailed comparative experimental analysis of the different decoding strategies discussed in Sec. IV, specialized to the square code. In Sec. VIII we present the results of a simulated quantum annealing study that numerically investigates the role of temperature and control noise on the performance of the ME and QAC-ME strategies, as these are parameters that we cannot control experimentally. We conclude and discuss possible developments of our work in Sec. IX. Additional supporting information is presented in the Appendix.

## II. QUANTUM ANNEALING, MINOR EMBEDDING, AND QUANTUM ANNEALING CORRECTION

### A. Quantum annealing

The term quantum annealing was first coined (as far as we know) in 1994 by Finnila *et al.* [4], though the underlying idea was proposed at least as early as 1984 [67] as a numerical algorithm inspired by simulated annealing [68], where the

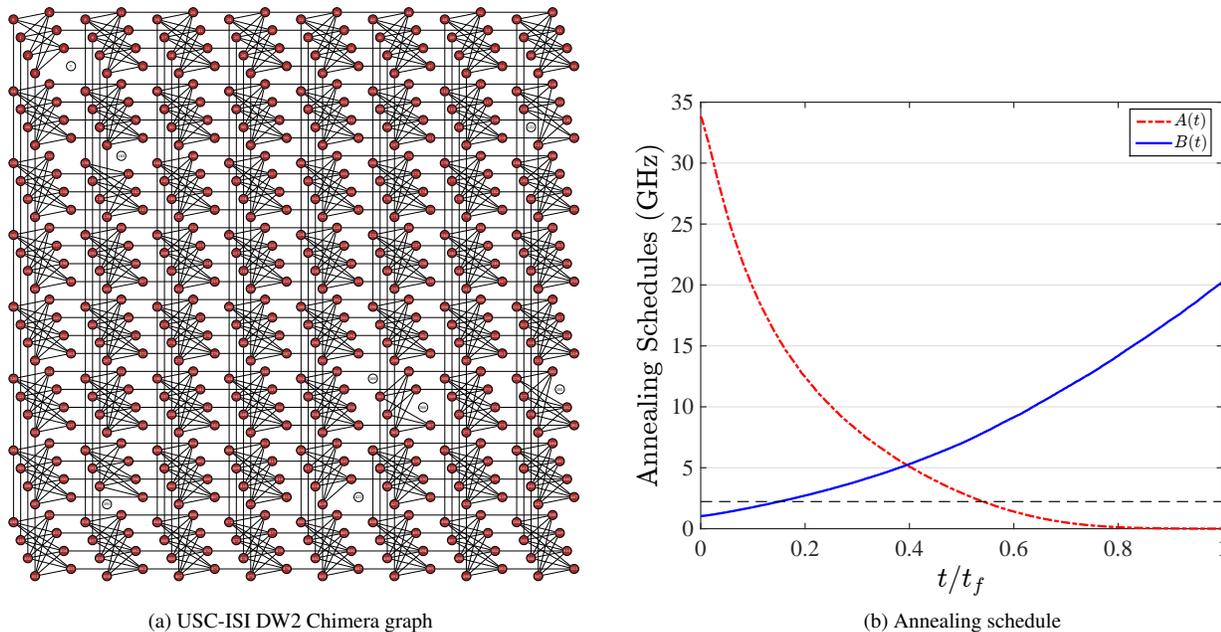


FIG. 1. (a) Physical connectivity of the “Chimera” graph of the 512-qubit DW2 processor. White circles correspond to the 8 unusable qubits in the USC-ISI DW2 processor. (b) Annealing schedule of the DW2 processor. The dashed horizontal line is the physical temperature of 17mK at which the device is kept.

role of thermal fluctuations was conjectured to be advantageously replaced by quantum tunneling. This was soon followed by numerous other studies based on a similar idea (see Refs. [5, 69–72] for some of the earliest works and Ref. [73] for a recent review). Numerical and experimental results have suggested that quantum tunneling can be more effective than thermal fluctuations for reaching the ground state [6–9]. This conclusion has recently been revisited [74], and remains the subject of active scrutiny (e.g., Ref. [75]).

Quantum annealing for spin systems is typically achieved by the following time-dependent, transverse-field Ising Hamiltonian:

$$H(t) = A(t)H_X + B(t)H_P, \quad t \in [0, t_f]. \quad (1)$$

The term  $H_X = -\sum_i \sigma_i^x$  (with  $\sigma_i^x$  being the Pauli matrix acting on qubit  $i$ ) is a transverse field whose amplitude controls the tunneling rate. The solution to an optimization problem of interest is encoded in the ground state of the problem Hamiltonian  $H_P$ , given by

$$H_P = \sum_{i \in \mathcal{V}_h} h_i \sigma_i^z + \sum_{(i,j) \in \mathcal{E}_J} J_{ij} \sigma_i^z \sigma_j^z, \quad (2)$$

where the sums run over the vertices  $\mathcal{V}$  and edges  $\mathcal{E}$  of the weighted, undirected graph  $G = (\mathcal{V}_h, \mathcal{E}_J)$ , where the local fields  $h_i$  and couplings  $J_{ij}$  are the weights. In a closed system, the adiabatic theorem ensures that if the system is initialized in the ground state of  $H(0) = A(0)H_X$  (assuming  $B(0) = 0$ ), a sufficiently slow (adiabatic) evolution results at the end of the computation in the ground state of the final Hamiltonian

$H(t_f) = B(t_f)H_P$  (assuming  $A(t_f) = 0$ ). In an open system annealing at a non-zero temperature one expects the system to evolve to a mixed final state, e.g., the Gibbs state of  $H(t_f)$ , if equilibration is possible throughout the annealing process (see, e.g., Ref. [25]).

The D-Wave devices use an array of superconducting niobium-based flux qubits in order to physically realize quantum annealing [11–13]. Each vertex of the D-Wave “Chimera” graph is a physical flux qubit and each edge is a physical inductive coupling between two qubits. Fig. 1(a) shows a pictorial representation of such a graph. In the DW2 processor used in this work the temperature is set to 17mK and the annealing schedule is given by the functions  $A(t)$  and  $B(t)$  shown in Fig. 1(b). In the following, we specify the values of the couplings  $h_i$  and  $J_{ij}$  in dimensionless units, with the maximum strength of the allowed couplings given by  $|h_i^{\max}| = 2$  and  $|J_{ij}^{\max}| = 1$ .

## B. Mapping to and from the logical problem: ME and QAC

The problem Hamiltonian defined in Eq. (2) may or may not be specified in terms that directly correspond to physical variables representing the actual degrees of freedom of a physical device. When such a specification is possible we have a *direct* (D) embedding. An example is the class of random Ising model problems studied in Refs. [17, 33], which simply require a specification of the  $h_i$  and  $J_{ij}$  on the Chimera graph of the D-Wave device.

However, often one is interested in solving problems that

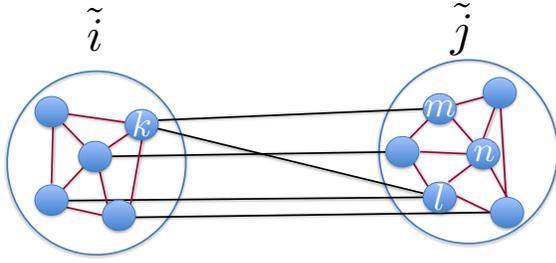


FIG. 2. (Color online) ME: Schematic of two logical groups  $\tilde{i}, \tilde{j}$  (large circles) representing two logical qubits  $\tilde{\sigma}_i^z, \tilde{\sigma}_j^z$ . Physical qubits are denoted by blue (filled) circles; physical problem couplings are denoted by black (inter logical groups) lines, penalty couplings by red (intra logical groups) lines. In the example shown the penalty coupling  $J_{j_m \tilde{j}_n}$  and the physical problem couplings  $J_{\tilde{i}_k \tilde{j}_l}, J_{\tilde{i}_k \tilde{j}_m}$  are active, among others.

do not have such a direct embedding. E.g., consider a problem defined on a complete graph, which requires a minor embedding on a graph that is not complete. In addition one might want an encoding of the Hamiltonian in order to perform error correction. To deal with this more general scenario we shall proceed to define two maps: one for ME, another for QAC. We shall also consider the concatenated map, of ME followed by QAC, which we refer to as QAC-ME.

### 1. ME: from the logical Hamiltonian to the physical Hamiltonian

When the problem Hamiltonian is defined in terms of variables that do not have a direct representation and requires a minor embedding we shall refer to it as a “logical Hamiltonian” and write

$$\tilde{H}_P = \sum_{\tilde{i} \in \tilde{\mathcal{V}}_h} \tilde{h}_{\tilde{i}} \tilde{\sigma}_i^z + \sum_{(\tilde{i}, \tilde{j}) \in \tilde{\mathcal{E}}_j} \tilde{J}_{\tilde{i}\tilde{j}} \tilde{\sigma}_i^z \tilde{\sigma}_j^z. \quad (3)$$

Quantities decorated with a tilde represent the abstract logical local fields, couplings, Pauli operators, etc. The logical Hamiltonian is defined on the vertices and edges of the logical graph  $\tilde{G} = (\tilde{\mathcal{V}}_h, \tilde{\mathcal{E}}_j)$ . Given the logical Hamiltonian which specifies the optimization problem of interest in terms of abstract variables, the task is to map it to a physical problem Hamiltonian that has a direct embedding. This is the minor embedding map

$$\tilde{H}_P \xrightarrow{\mathcal{M}_{\text{ME}}} H_P. \quad (4)$$

This map exists when  $\tilde{G}$  is a *minor* of  $G$ , i.e., it has a minor embedding in  $G$ .<sup>1</sup>

<sup>1</sup> A minor graph is obtained from  $G$  by collapsing groups of connected (physical) vertices of  $G$  into single (logical) vertices of  $\tilde{G}$  and possibly

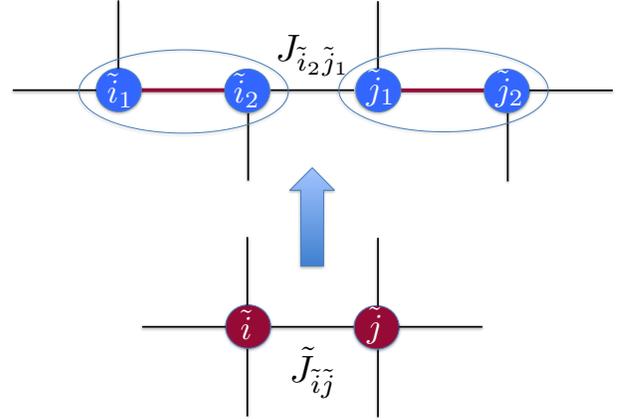


FIG. 3. (Color online) Schematic example of the map  $\mathcal{M}_{\text{ME}}$ . Two logical qubits  $\tilde{i}$  and  $\tilde{j}$  (red dots, bottom of the figure) with four logical couplings each (black, thin lines) are expanded into two physical qubits each (circled blue dots, top of the figure) with three physical couplings each. The red (thick) lines connecting the physical qubits are the penalty couplings  $J_{i_1 i_2}$  and  $J_{j_1 j_2}$ . In this example the logical coupling  $\tilde{J}_{\tilde{i}\tilde{j}}$  is replaced by a single physical coupling  $J_{i_2 j_1}$ , a special case of Eq. (6).

Each of the  $\tilde{N} = |\tilde{\mathcal{V}}_h|$  vertices of  $\tilde{G}$  is occupied by a classical binary ( $\pm 1$ ) variable, which we call a logical qubit.<sup>2</sup> The map  $\mathcal{M}_{\text{ME}}$  replaces these by  $N = |\mathcal{V}_h|$  physical qubits. A minor embedding defines a partitioning of all  $N$  physical qubits into  $\tilde{N}$  disjoint logical groups  $\{\tilde{i}\}$ , each comprising  $N_{\tilde{i}}$  physical qubits, whose union covers all physical qubits, i.e.,  $N = \sum_{\tilde{i}=1}^{\tilde{N}} N_{\tilde{i}}$ . Note that we distinguish between a *logical group* and a *logical qubit*. The latter was defined above; the former is a set of physical qubits representing the logical qubit. What identifies a logical group is a set of ferromagnetic “penalty” couplings which tie the physical qubits representing a given logical qubit together, as illustrated in Fig. 2. Penalty couplings are always “intra”, i.e., couple physical qubits of a given logical group  $\tilde{i}$ , and can thus always be written as

$$J_{\text{pen}} \in \{J_{\tilde{i}_k \tilde{i}_l}\}. \quad (5)$$

The physical qubits comprising two logical groups are coupled in a manner that depends on the problem Hamiltonian. The logical local fields and couplings are defined via

$$\tilde{h}_{\tilde{i}} = \sum_{k=1}^{N_{\tilde{i}}} h_{\tilde{i}_k}, \quad \tilde{J}_{\tilde{i}\tilde{j}} = \sum_{k=1}^{N_{\tilde{i}}} \sum_{l=1}^{N_{\tilde{j}}} J_{\tilde{i}_k \tilde{j}_l} \quad (\tilde{i} \neq \tilde{j}), \quad (6)$$

removing edges from  $G$  [76]. Note that  $\mathcal{M}_{\text{ME}}$  is a multivalued map (one-to-many).

<sup>2</sup> Technically this is a misnomer since the variable is classical. However, we have in mind the entire time-dependent Hamiltonian of quantum annealing, and moreover, we can always measure the logical qubit in the logical  $\sigma^z$  basis.

where some of the  $J_{\tilde{i}_k \tilde{j}_l}$  may vanish, as illustrated in Fig. 2. Note that Eq. (6) does not uniquely specify the physical local fields and couplings (which are always “inter”), allowing for a large arbitrariness in fixing the mapping  $\mathcal{M}_{\text{ME}}$  between the logical and physical problems. This arbitrariness is in fact useful, since the efficacy of an ME implementation can be very sensitive to the specific choice of the physical couplings.

Fig. 3 shows schematically how a logical-graph vertex of degree four can be minor embedded into two physical-graph vertices of degree three, and how two logical qubits are correspondingly replaced by four physical qubits.

Ideally, a minor embedding should give a faithful representation of the original logical problem, in the sense that it should encode the solution to the logical problem in its ground state. Intuitively, this can be achieved by implementing large penalty couplings, though it is clear that a balance must be struck so as not to overwhelm the problem couplings. Consequently the optimal choice for the energy penalties is not obvious. Very large penalties make logical groups behave as independent ferromagnetic clusters. These penalties can overwhelm the relatively weaker interactions between different clusters that represent the logical problem. When the logical groups are small, theoretical and experimental analyses suggest that the optimal value of the energy penalties is of the same order of magnitude as that of the logical couplings [59, 60]:<sup>3</sup>

$$J_{\tilde{i}_k \tilde{j}_l} \sim \tilde{J}_{\tilde{i} \tilde{j}}. \quad (7)$$

We discuss how this freedom can be exploited in order to improve the performance of ME implementations in Section III.

An important consequence of the finiteness of the energy penalties within each logical group is that, informally, logical groups can “break”, i.e., the values of the physical qubits (measured in the  $\sigma^z$  eigenbasis at  $t_f$ ) representing a logical qubit can disagree. This may result in an unfaithful representation of the logical qubit by its logical group. We call logical groups with internal disagreement “broken qubits” (BQs). More formally, the logical qubit belongs to the code subspace defined by the two “all-agree” states of the logical group, and any “disagreement” is a non-code state (more on this in our discussion of QAC below). Because of Eq. (7) BQs are ubiquitous in every ME implementation. Classical post-processing (decoding) strategies that assign appropriate logical values to BQs are thus a necessary aspect of minor embedded quantum annealing. We discuss several such strategies in Section IV.

## 2. QAC: from the physical Hamiltonian to the encoded Hamiltonian

When the problem Hamiltonian  $H_P$  in Eq. (2) needs to be encoded in order to provide protection against decoherence

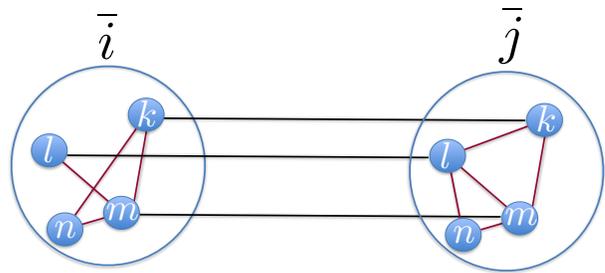


FIG. 4. (Color online) QAC: Schematic of two encoded groups  $\tilde{i}, \tilde{j}$  (large circles) representing two encoded qubits  $\bar{\sigma}_{\tilde{i}}, \bar{\sigma}_{\tilde{j}}$ . Physical qubits are denoted by blue (filled) circles; physical problem couplings are denoted by black (inter logical groups) lines, penalty couplings by red (intra logical groups) lines. In the QAC case the problem couplings are an injection (one-to-one but not necessarily onto) between the physical qubits of each encoded group. As depicted, the penalty couplings may involve all or a subset of all physical qubits of each encoded group, and may vary from group to group.

and noise we shall refer to the result as the “encoded Hamiltonian” and write

$$\bar{H}_P = \sum_{\tilde{i} \in \bar{\mathcal{V}}_h} h_{\tilde{i}} \bar{\sigma}_{\tilde{i}}^z + \sum_{(\tilde{i}, \tilde{j}) \in \bar{\mathcal{E}}_J} J_{\tilde{i} \tilde{j}} \bar{\sigma}_{\tilde{i}}^z \bar{\sigma}_{\tilde{j}}^z. \quad (8)$$

Quantities decorated with a bar represent the encoded operators, acting on the  $\bar{N}$  encoded qubits. The encoded Hamiltonian is defined on the vertices and edges of the encoded graph  $\bar{G} = (\bar{\mathcal{V}}_h, \bar{\mathcal{E}}_J)$ . Note that unlike the ME case, the same local fields and couplings are used in the encoded Hamiltonian as in the physical Hamiltonian (i.e.,  $h$  and  $J$  do not have a bar), except that they are defined on the encoded graph. In particular,  $h_{\tilde{i}} = h_i$  and  $J_{\tilde{i} \tilde{j}} = J_{ij}$ . Given the physical Hamiltonian, the task is to map it to an encoded Hamiltonian that also has a direct embedding. This is the quantum annealing correction map

$$H_P \xrightarrow{\mathcal{M}_{\text{QAC}}} \bar{H}_P. \quad (9)$$

This map exists when  $\bar{G}$  is a *subgraph* of  $G$ . I.e., QAC requires several copies of the logical graph  $\bar{G}$  to be directly embeddable into the physical graph  $G$ .

Each of the  $N = |\mathcal{V}_h|$  vertices of  $G$  is occupied by a classical binary ( $\pm 1$ ) variable. The map  $\mathcal{M}_{\text{QAC}}$  replaces these by  $\bar{N} = |\bar{\mathcal{V}}_h|$  encoded qubits. Similarly to the ME case, we identify encoded groups as sets of physical qubits tied together by ferromagnetic penalty couplings, as depicted in Fig. 4. Unlike the ME case, we impose that the problem couplings between encoded groups are injective (one-to-one) and that all encoded groups are of identical size:  $N_{\tilde{i}} = K \forall \tilde{i} \in \{1, \dots, \bar{N}\}$ . The penalty couplings may be non-uniform in order to allow for their optimization, e.g., as in Eq. (7).

In the QAC case encoding acts to provide protection. To see how this works, we briefly digress to frame our discussion in a

<sup>3</sup> When long chains are used as logical groups, the optimal value for the energy penalties seems to scale linearly with the chain length [77].

more general setting. Let us recall that one can define a quantum error correction code via a set of stabilizers,  $\{S_s\}$ , (or parity checks in classical error correction language) [44, 78]. This set of commuting operators on  $n$  qubits defines a collection of syndrome subspaces, each encoding a fixed number of encoded qubits  $k$  with corresponding encoded operators  $\{\bar{\sigma}_\kappa^x, \bar{\sigma}_\kappa^z\}_{\kappa=1}^k$  (we temporarily drop the bar from the qubit subscripts; we will sometimes return to this notation if there is no danger of confusion). These subspaces are defined according to their eigenvalues with respect to the stabilizer operators. Since our focus is on encoding the final Hamiltonian  $H_P$  we only consider bit flip errors and thus restrict our attention to quantum repetition codes.

An  $[[n, k, d]]$  quantum repetition code (encoding  $n$  physical qubits into  $k$  encoded qubits and having distance<sup>4</sup>  $d = n$  can be defined by the stabilizers  $\sigma_i^z \sigma_j^z$  where  $\{i, j\}$  can be chosen as all the nearest neighbor pairs. Such a code can correct  $t = \lfloor d/2 \rfloor$  errors, and can be reliably decoded using majority voting if the number of errors does not exceed  $t$ . The decoding is undecided in the case of a tie, i.e., for  $n/2$  errors ( $n$  even) and fails if the number of errors exceeds  $t$ .

For simplicity we consider the case of a single encoded qubit ( $k = 1$ ) per encoded group of  $n$ , and consider multiple encoded groups  $\{\ell\}$ . Such a code has the basis states

$$|\bar{0}\rangle = |0\rangle^{\otimes n} ; |\bar{1}\rangle = |1\rangle^{\otimes n} . \quad (10)$$

The encoded operators are  $\bar{\sigma}^x = (\sigma^x)^{\otimes n}$  and, e.g.,  $\bar{\sigma}^z = \sigma_1^z$ .

The QAC protection idea is to exploit the various realizations of the encoded operators in order to provide an energy boost, and to include stabilizer terms in such a way that errors causing transitions outside of the code space are penalized. In particular, writing the unencoded final Hamiltonian [Eq. (2)] in the form

$$H = \sum_{\ell} h_{\ell} \sigma_{\ell}^z + \sum_{\ell, \ell'} J_{\ell \ell'} \sigma_{\ell}^z \sigma_{\ell'}^z , \quad (11)$$

the encoded final Hamiltonian becomes

$$\bar{H} = \sum_{\ell} h_{\ell} \sum_r [\sigma_{\ell}^z]_r + \sum_{\ell, \ell'} \sum_r J_{\ell \ell'} [\sigma_{\ell}^z \sigma_{\ell'}^z]_r , \quad (12)$$

where  $[A]_r$  is the  $r$ th realization of the operator  $A$ . These multiple realizations exist since the logical Pauli operators are equivalent under multiplication by the code stabilizers, and for our repetition code any tensor product of an odd number of  $\sigma^z$  operators is an equivalent realization of  $\bar{\sigma}^z$ . This multiplicity boosts the energy scale of  $\bar{H}$  relative to  $H$ , and it increases the gap against thermal excitations. This is complemented by adding to  $\bar{H}$  a penalty term that is a sum over the available

stabilizers  $S_{\alpha}$  for each encoded qubit  $\ell$ :

$$H_{\text{pen}} = -J_{\text{pen}} \sum_{\ell, s} [S_s]_{\ell} . \quad (13)$$

Each bit-flip error anticommutes with at least one term in  $H_{\text{pen}}$  and hence incurs an energy cost of at least  $2J_{\text{pen}}$ . Such errors are thus thermally suppressed [50, 53, 54, 56, 57, 59–61].

Returning now to our specific setting and previous notation, note that whereas in the ME case the  $\bar{\sigma}_i^z$  were abstract binary variables, in the QAC case we choose the encoded Pauli operators as the total Pauli operators of each encoded qubit, i.e.,

$$\bar{\sigma}_i^{\alpha} = \sum_{k=1}^{K'} \sigma_{i_k}^{\alpha} , \quad K' \leq K, \quad (14)$$

where we allow for the possibility that not all  $K$  physical qubits in each encoded group participate in defining the encoded Pauli operators (i.e.,  $K' < K$ ). In this manner the encoded problem Hamiltonian is expressed in terms of the encoded operators of the quantum repetition code. This means that  $J_{ij} \sigma_i^z \sigma_j^z \mapsto J_{ij} \bar{\sigma}_i^z \bar{\sigma}_j^z = J_{ij} \sum_{k=1}^{K'} \sigma_{i_k}^z \sigma_{j_k}^z$ , i.e., each coupling  $J_{ij}$  is boosted by a factor of  $K'$ , since on the encoded graph each  $J_{ij}$  is represented by  $K'$  coupled pairs of physical qubits. We stress again that this boost is a useful feature since it increases the energy scale of the encoded problem Hamiltonian, a feature that works to suppress both thermal excitations and random control errors [59–61].

Broken encoded qubits appear in the QAC case just as BQs appear in the ME case, manifesting as encoded groups whose physical qubits yield different results when measured at  $t_f$ , resulting in a measured eigenvalue of  $\bar{\sigma}_i^{\alpha}$  that is smaller than  $K'$  in absolute value. Our QAC code can be viewed as a  $[[K', 1, K']]$  repetition code, i.e., a code with distance  $K'$  encoding one qubit into  $K'$ . The penalty terms  $J_{i_k, \bar{i}_l} \sigma_{i_k}^z \sigma_{\bar{i}_l}^z$  implicit in  $\bar{H}_P$  can be understood as elements of the stabilizers of the repetition code. Since these stabilizers detect bit-flip errors, as mentioned above these errors are energetically suppressed.

The transverse field is invariant under the transformation to the encoded  $\sigma^x$  operator as in Eq. (14):

$$H_X = \sum_{i=1}^N \sigma_i^x = \sum_{\bar{i}=1}^{\bar{N}} \sum_{k=1}^K \sigma_{i_k}^x = \sum_{\bar{i}=1}^{\bar{N}} \bar{\sigma}_{\bar{i}}^x = \bar{H}_X . \quad (15)$$

This observation is motivated by the fact that the D-Wave processors do not allow control over the transverse field. However, we note that for the purposes of error correction this represents a significant limitation, since (as mentioned above) the encoded  $\sigma^x$  operator of the repetition code is in fact  $\bar{\sigma}^x = \otimes_k \sigma_k^x$ , and the inability to implement this (many-body) operator compromises the efficacy of QAC [59, 60] as well as other strategies [50, 56].

<sup>4</sup> A code can detect any Pauli errors error  $E$  having weight  $\leq d - 1$ , i.e.,  $\langle \psi_a | E | \psi_b \rangle = c \delta_{ab}$  for any pair of codewords  $|\psi_a\rangle, |\psi_b\rangle$  and constant  $c$ . Equivalently, the distance is the minimum weight of a Pauli operator  $E$  such that  $\langle \psi_a | E | \psi_b \rangle \neq c \delta_{ab}$ .

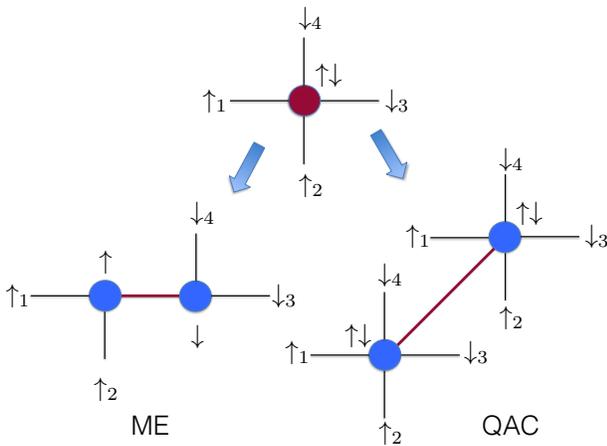


FIG. 5. (Color online) Contrasting ME and QAC via frustration effects. ME: The top configuration displays an example of a frustrated logical qubit (red) with ferromagnetic logical interactions (black). Frustration is a consequence of the opposite orientations of the outer logical qubits (represented as numbered arrows). Frustration of the logical interactions can be removed by the appearance of a broken qubit (lower left) after minor embedding. This configuration can be energetically favored if the penalty coupling (red, thick line) is sufficiently weak. Thus, in the ME case, frustration may favor the appearance of broken qubits. QAC: Now the top configuration represents a frustrated physical qubit (red) with ferromagnetic couplings (black). In this example the QAC map replaces this configuration with two copies (bottom right), coupled with a ferromagnetic penalty (red, thick). After the QAC map it is energetically favorable for the penalty-coupled qubits (blue) to be aligned for arbitrarily small values of the penalty. Thus, in the QAC case, frustration does not favor the appearance of broken qubits. Instead, configurations with broken qubits are excited states of the encoded problem for arbitrarily small values of the penalty terms.

### C. ME vs QAC: frustration effects

QAC is known to be effective in improving the performance of a quantum annealer when compared to a direct (D) embedding of optimization problems [59, 60]. On the other hand, ME typically reduces performance relative to a direct implementation of a given optimization problem. This opposite response to ME and QAC has an intuitive explanation in how ME and QAC behave under frustration, which can be useful once the transverse field is off and the evolution is effectively classical [i.e., when  $A(t) \ll kT < B(t)$ , e.g., for  $t/t_f \approx 0.8$  in Fig. 1(b)]. The problem Hamiltonian corresponding to a hard optimization problem typically represents a strongly frustrated spin system, i.e., wherein the ground state configuration cannot satisfy all logical couplings. An example of what happens when a frustrated logical qubit is minor embedded or QAC embedded is given in Fig. 5. In the ME case (where we assume that the logical problem cannot be directly embedded), if the penalty term is sufficiently weak, frustration results in a broken qubit, thus introducing new logical errors (to-be-decoded broken qubits). In the QAC case (where we assume that the problem can be directly embed-

ded), in contrast, each copy of the problem bears the same frustration independently, and it is energetically favorable to satisfy the penalty coupling, no matter how weak it is. Therefore in QAC frustration is not the reason for the appearance of logical errors (whose source can instead be, e.g., thermal). Rather, thanks to their connection via energy penalties, the different copies of the problem help each other to relax more efficiently to the logical ground state.

### D. Quantum annealing correction on minor embeddings (QAC-ME)

It is natural to consider a concatenation of the ME and QAC maps:

$$\tilde{H}_P \xrightarrow{M_{ME}} H_P \xrightarrow{M_{QAC}} \bar{H}_P. \quad (16)$$

This is the QAC-ME map. It replaces a logical problem Hamiltonian by an encoded problem Hamiltonian, after a minor embedding of the former.

The ME and QAC maps both expand the number of qubits. In the ME case a logical qubit is expanded into several physical qubits in order to enable an embedding of the logical problem Hamiltonian into a physical problem Hamiltonian that obeys the constraints of a given hardware graph. In the QAC case a physical qubit is expanded into several physical qubits in order to generate protection against decoherence and noise.

Figure 6 schematically summarizes our construction. Figure 7 shows how the various implementations are related to each other according to their expected performance when implemented on a quantum annealer. A problem formulated on a logical graph  $\bar{G}$  can be directly embedded (D) on a physical device with given hardware graph  $G$  if  $\bar{G}$  is a subgraph of  $G$ . QAC aims to improve the success probability of solving this problem as compared to a D implementation. On the other hand, for problems formulated on a logical graph that can only be minor embedded, the goal of QAC-ME is to improve the success with respect to ME. It is typically not possible to directly compare D/QAC and ME/QAC-ME, simply because direct implementations of a given optimization problem are generically unavailable. It is however of great interest to consider cases where both D and ME/QAC-ME implementations are possible. This gives us an estimate of what performance decay is due to ME, and how much of it can be recovered by quantum annealing correction schemes such as QAC-ME. We perform such comparisons below.

## III. CHOOSING THE PENALTIES

The penalty strengths should be optimized in order to maximize the performance of the embedding procedures. It is impractical to perform this optimization on an instance-by-instance basis, so it is important to find a general rule for determining the optimal choice of penalties. This is an open problem that we address below.

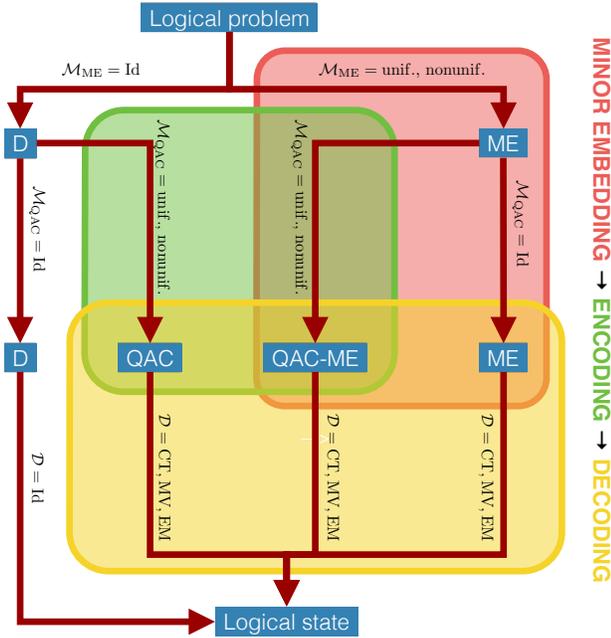


FIG. 6. The flow chart above gives a schematic representation of our implementation and summarizes the discussions of sections II, III (encoding) and IV (decoding). The minor embedding ( $\mathcal{M}_{ME}$ ) and quantum annealing correction ( $\mathcal{M}_{QAC}$ ) maps are concatenated in the encoding step. Both maps can be realized with either a uniform or nonuniform choice of the energy penalties (see section III). The final step is represented by a decoding map  $\mathcal{D}$  that reconstructs a logical state from a physical state obtained through the QAC-ME encoding. Decoding can be obtained, for example, with coin tossing (CT), majority vote (MV) or energy minimization (EM). See section IV for more details on the decoding step.

### A. Uniform penalties

Recent work has shown both experimentally and theoretically that even the simplest choice—a constant—for the penalties,

$$J_{\text{pen}} = -\gamma, \quad (17)$$

allows for successful QAC [59, 60] or ME [77]. The optimal value of  $\gamma$  is found to be roughly proportional to the value of the physical couplings that represent logical connections and typically grows with the size of the logical (or encoded) groups.

### B. Nonuniform penalties

The optimal choice of penalty strengths achieves a delicate balance of tying the cluster of physical qubits forming a logical group or encoded group together while not overwhelming the problem couplings. We can expect that a better balance can be achieved with a choice of penalties that depends on the

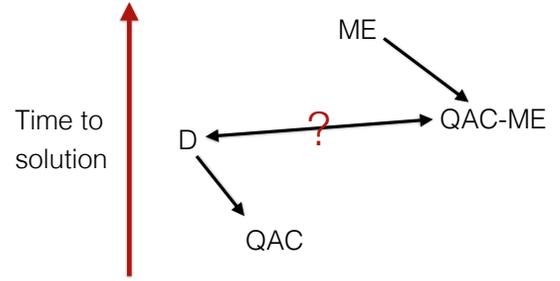


FIG. 7. Schematic comparing the performance of different embedding schemes. The time to solution (TTS) axis is not to be taken literally and represents wishful thinking. The goal of QAC is to reduce TTS (or increase the success probability) with respect to a D embedding, while QAC-ME is meant to improve over ME. Due to various tradeoffs it is not generally clear whether QAC-ME improves upon D. Of course, a D embedding may not be available, in which case QAC-ME is the only option.

local structure of the logical problem. One possibility is to choose the energy penalties as follows:

$$\text{ME: } \forall k, l \quad J_{\tilde{i}_k, \tilde{i}_l} \equiv J_{\text{pen}, \tilde{i}} = -\gamma \text{mean}_{\tilde{j}} |\tilde{J}_{\tilde{i}, \tilde{j}}| \quad (18a)$$

$$\text{QAC: } \forall k, l \quad J_{\tilde{i}_k, \tilde{i}_l} \equiv J_{\text{pen}, \tilde{i}} = -\gamma \text{mean}_{\tilde{j}} |J_{\tilde{i}, \tilde{j}}|, \quad (18b)$$

These expressions assign a value to the energy penalties associated with logical qubit  $\tilde{i}$  (encoded qubit  $\tilde{i}$ ) that is proportional to the mean of the absolute strength of the logical couplings connected to that logical qubit (encoded qubit). The overall constant factor  $\gamma$  is then optimized as in the uniform case for the best results. Typically, practical problems translate into Ising Hamiltonians with a certain amount of structure, both in the connectivity and in the value of the logical couplings. We thus expect this nonuniform choice to be generally more effective. In section VI we present a study of the efficacy of the uniform and nonuniform encoding strategies described above using the DW2 device.

## IV. DECODING STRATEGIES

A decoding strategy is a rule that assigns a value to each logical or encoded qubit given the values of the physical qubits retrieved from the annealer in the read-out process. As already mentioned in the general discussion about ME and QAC, in the decoding process we first distinguish between “broken” and “unbroken” logical or encoded qubits. The unbroken case corresponds to a logical or encoded group of physical qubits that all have the same measured final values. In this case, decoding is trivial and consists in assigning to the logical or encoded qubit the common value of the corresponding physical qubits. The broken case, on the other hand, corresponds to a logical or encoded group where there is at least one physical qubits whose value differs from the others. Decoding strictly applies to BQs, and it can be done in different ways. We will consider two general approaches (and a

combination therefore). The first one is “local”, in the sense that the decoding assignment is performed on each logical or encoded group independently. It is the usual strategy for decoding classical repetition codes and relies on majority vote and coin tossing. The other can be considered a “global” approach, in which the value of a logical qubit is decoded by a global energy-minimization procedure that simultaneously involves all BQs.

### A. Coin tossing (CT)

The simplest decoding strategy is random decoding (“coin tossing”), wherein we assign a random  $\pm 1$  value to each logical or encoded qubit. This strategy merely serves as a baseline against which we compare all other strategies.

### B. Majority vote (MV)

In the majority vote strategy the value of each encoded qubit is chosen as the most recurrent value within an encoded group:

$$\bar{s}_i = \text{sign}\left(\sum_k s_{ik}^{\pm}\right), \quad (19)$$

where  $\bar{s}$  and  $s$  stand for the values of the encoded and physical qubits in the  $\sigma^z$  basis at  $t = t_f$ , respectively.<sup>5</sup> MV decoding is most successful under the assumption of independence of the error probability  $p_{\text{err}}$  for each physical bit. A repetition code of length  $n$  can be reliably decoded with up to  $d = \lfloor n/2 \rfloor$  errors, an event that occurs with probability  $p_{\text{err}}^d$ , which is exponentially decreasing in the code distance. An important caveat of applying MV to ME and QAC is that physical excitations above the ground state occurring during the evolution cannot be considered as generating independent errors on the physical qubits. MV itself thus does not ensure exponentially improving resiliency of the ME implementation as a function of the size of the encoded group.

In QAC schemes, where the encoded qubits are implemented as several equivalent physical copies, MV will be successful if errors in the different copies can be considered as independent. A quantum contribution in the QAC scheme arises from the energy penalties, which provide further suppression of errors. MV has indeed proven to be an effective approach to decoding QAC [59, 60].

However, MV is problematic for most ME schemes. In this case the absence of BQs is crucial for the correct representation of the logical problem as a minor embedded problem, since a logical qubit can break in several different ways, giving rise to different MV-decoded logical values. These different “breaking modes” may have similar energy cost, making

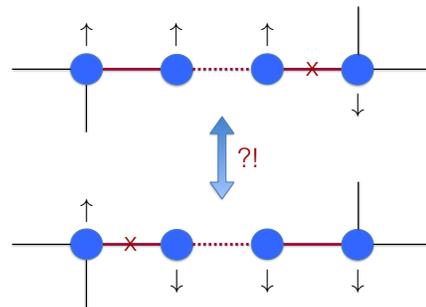


FIG. 8. Failure modes for majority vote in ME. A logical qubit is represented here by a physical chain. MV decoding on the bottom chain gives a spin-down result, on the top it gives spin-up. The chain can break iso-energetically in different places and decoding by majority vote, in this case, is likely to give no better results than coin tossing.

them similarly probable. A typical example of this is shown in Fig. 8, where a logical qubit is represented by a physical chain. The chain can break in several places, all of them costing the same energy, giving different results through MV. We should not expect, in this simple case, MV to be more effective than decoding by coin tossing.

Another issue that arises for MV is how to decode a logical qubit comprising an even number  $n$  of physical qubits. A tie ( $n/2$  errors) cannot be decoded by MV, or by any “local” decision rule other than coin tossing.

### C. Energy minimization (EM)

Since we are interested in optimization problems, a natural alternative to MV is a “global” decoding strategy wherein all BQs are decoded together in such a way as to minimize the energy of the encoded qubits. Decoding a particular state with this energy minimization strategy corresponds to solving an optimization problem on the subgraph of the encoded graph induced by the set of BQs, i.e., the qubits that require decoding. This problem is equivalent to minimizing the following “decoding” Hamiltonian

$$H_D = \sum_{\bar{i} \in \text{BQ}} \left( h_{\bar{i}} + \sum_{\bar{j} \notin \text{BQ}} J_{\bar{i}\bar{j}} \bar{s}_{\bar{j}} \right) \bar{\sigma}_{\bar{i}}^z + \sum_{\bar{i}, \bar{j} \in \text{BQ}} J_{\bar{i}\bar{j}} \bar{\sigma}_{\bar{i}}^z \bar{\sigma}_{\bar{j}}^z. \quad (20)$$

This energy minimization strategy is rather general and can be used to decode ME, QAC or QAC-MEC schemes. Moreover, it is applicable to cases where logical qubits comprise an even number of physical qubits, since it can also resolve ties.

A disadvantage of EM is that, since it requires solving a global optimization problem, it can be computationally expensive. If the number of BQs is small (in a sense we quantify below), this optimization can be done exactly; otherwise, heuristic methods are necessary. We discuss EM decoding via simulated annealing, along with experimental results, in Appendix A. In the next section we argue that for a sufficiently

<sup>5</sup> From hereon, if there is no danger of confusion, we use the QAC terminology and notation to denote both QAC and ME.

small number of BQs, EM decoding can be implemented efficiently.

#### D. Efficient optimal decoding and the percolation threshold

We now give a general criterion to determine whether energy minimization decoding can be done efficiently. We use well-known results of percolation theory [79] and rely on the simplifying assumption that an encoded qubit is broken with a qubit-independent probability  $p_{\text{BQ}}$ . This probability depends on the strength of the energy penalties used for the encoding. Eventually, the relevant value of  $p_{\text{BQ}}$  is obtained when it is evaluated at the optimal penalty value.

Let us start from two extreme cases. First, if  $p_{\text{BQ}} \ll 1$  errors are sparse, and it is unlikely that two BQs are coupled. In this case the decoding effort grows only linearly with the size of the problem because each BQ can be optimized independently. In such a regime, the EM strategy can be implemented efficiently, since the decoding time scales linearly in the number of encoded qubits

$$T_{\text{D}}(\bar{N}) \sim \bar{N}. \quad (21)$$

Second, if  $p_{\text{BQ}}$  is close to 1, most BQs are likely to be coupled. In this regime the EM decoding strategy requires one to solve an optimization problem that is as hard as the original problem:

$$T_{\text{D}}(\bar{N}) \sim T_{\text{P}}(\bar{N}), \quad (22)$$

where  $T_{\text{P}}$  is the time to solve the original problem. In this regime, EM decoding cannot be implemented efficiently.

In general, percolation theory can be used to establish a “decoding threshold”  $p_{\text{D}}$ , such that EM decoding can be implemented efficiently (in the large  $\bar{N}$  limit) if and only if  $p_{\text{BQ}} < p_{\text{D}}$  [this assumes that  $p_{\text{BQ}}$  is roughly constant across the graph, an assumption we test and confirm experimentally below (see Fig. 17)]. It turns out that  $p_{\text{D}}$  is equal to the “per-site” percolation threshold  $p_{\text{th}}$  of the encoded graph. To see this, we note that the decoding hamiltonian  $H_{\text{D}}$  is defined on a “per-site” percolation subgraph, (i.e., the subgraph induced by vertices selected with some probability  $p$ ). A fundamental result of percolation theory is that (for  $\bar{N} \gg 1$ ) there exists a threshold probability  $p_{\text{th}}$ , determined by the local structure of the graph, above which an infinite (i.e., spanning) cluster is present with probability 1. This result implies that EM decoding cannot be done efficiently if  $p_{\text{BQ}} > p_{\text{th}}$ , since it requires minimizing  $H_{\text{D}}$  for a system of BQs that is as large as the original problem.

On the other hand, below the threshold  $p_{\text{BQ}} < p_{\text{th}}$ , the probability to have an infinitely large cluster vanishes. More precisely it can be shown that the probability that a given BQ is connected to at least  $\bar{r}$  other BQs decays exponentially with the size  $\bar{r}$  of the connected domain [80]. Assuming that the  $\bar{N}$  encoded qubits grow their domains independently, we can thus express the total probability to have no domains of more

than  $\bar{r}$  connected encoded qubits as

$$P(\bar{r}) = (1 - \alpha e^{-\gamma \bar{r}})^{\bar{N}} \approx 1 - \alpha \bar{N} e^{-\gamma \bar{r}}, \quad (p_{\text{BQ}} < p_{\text{th}}) \quad (23)$$

for some constants  $\alpha$  and  $\gamma$  and for a sufficiently small exponential factor. Let us now fix the largest size of a decodable domain to be  $\bar{r}$ . If the encoded state contains a domain of larger size, we deem it “undecodable,” and we reject it. An undecodable event will happen with probability  $1 - P(\bar{r})$ . Let us again denote by  $T_{\text{P}}(\bar{r})$  the time needed to optimize over a domain of size  $\bar{r}$  (assuming that the decoding optimization is at most as hard as the original optimization problem). The total time required for decoding  $T_{\text{D}}(\bar{N})$  is proportional to the typical number  $\frac{\bar{N}}{\bar{r}}$  of independent domains of size  $\bar{r}$  in the system, multiplied by the optimization time  $T_{\text{P}}(\bar{r})$  of a region of size  $\bar{r}$ . Also, we note that the higher the probability of having no domains of more than  $r$  connected encoded qubits, the shorter is the total decoding time. Thus:

$$T_{\text{D}}(\bar{N}) \sim \frac{\bar{N}}{\bar{r}} \frac{T_{\text{P}}(\bar{r})}{P(\bar{r})}. \quad (24)$$

Equation (23) shows that if we take  $\bar{r} \sim \log(\bar{N})$ ,  $P(\bar{r})$  can be kept constant and close to 1 as  $\bar{N} \rightarrow \infty$ . Thus we have

$$T_{\text{D}}(\bar{N}) \sim \frac{\bar{N}}{\log(\bar{N})} T_{\text{P}}(\log(\bar{N})), \quad (25)$$

which is the main result of this section. Below the percolation threshold, as long as we do not attempt to decode domains whose size scales faster than  $\log(\bar{N})$ , the time required for optimal decoding through energy minimization scales sub-exponentially as compared to the time to find a solution to the original problem. In the large size limit, therefore, the decoding effort is negligible.

We note that, in principle, the constant prefactor needed to have a small  $P(\bar{r})$  can be very large and it diverges when we approach the percolation threshold. However, because the percolation transition is usually very sharp, we can expect the prefactor to be small if we are well below the threshold.

We revisit the percolation threshold experimentally below, and show (see Fig. 18) that for the problem instances we considered in this work we are well below the threshold, so that efficient decoding is possible.

## V. QAC-ME OF AN NP-HARD PROBLEM

In this section we describe an explicit QAC-ME scheme that combines a well-known NP-hard problem with a new QAC code.

### A. The two-level-grid (2LG) graph and its embedding

Consider a logical graph  $\tilde{G}$  consisting of two connected square lattices, i.e., a two-level-grid (2LG) as depicted in Fig. 9(a). The Ising spin glass problem on this graph is of

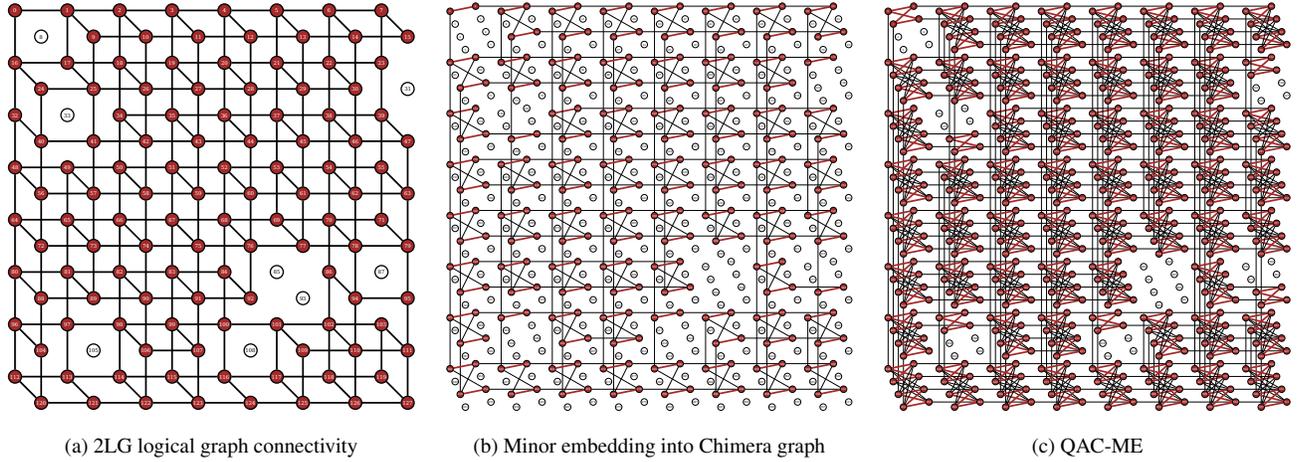


FIG. 9. (Color online) From a logical problem to a QAC-ME embedding. (a) Barahona’s “two-level-grid” (2LG) graph, for which the Ising spin glass problem with couplings in  $\{-1, 0, 1\}$  is NP-hard [66]. Disconnected vertices indicate spins with couplings set to zero, as well as unused logical qubits in the logical DW2 Chimera graph. (b) Minor embedding of the 2LG graph into the physical DW2 Chimera graph. White circles correspond to unused or unusable qubits. (c) QAC-ME embedding of the 2LG problem using the “square” code. In (b) and (c) penalties are represented by red (thick) couplings between groups of two (ME) and four (QAC-ME) physical qubits. Black (thin) links implement logical couplings. See also Fig. 11 for a more detailed description of the QAC-ME encoding.

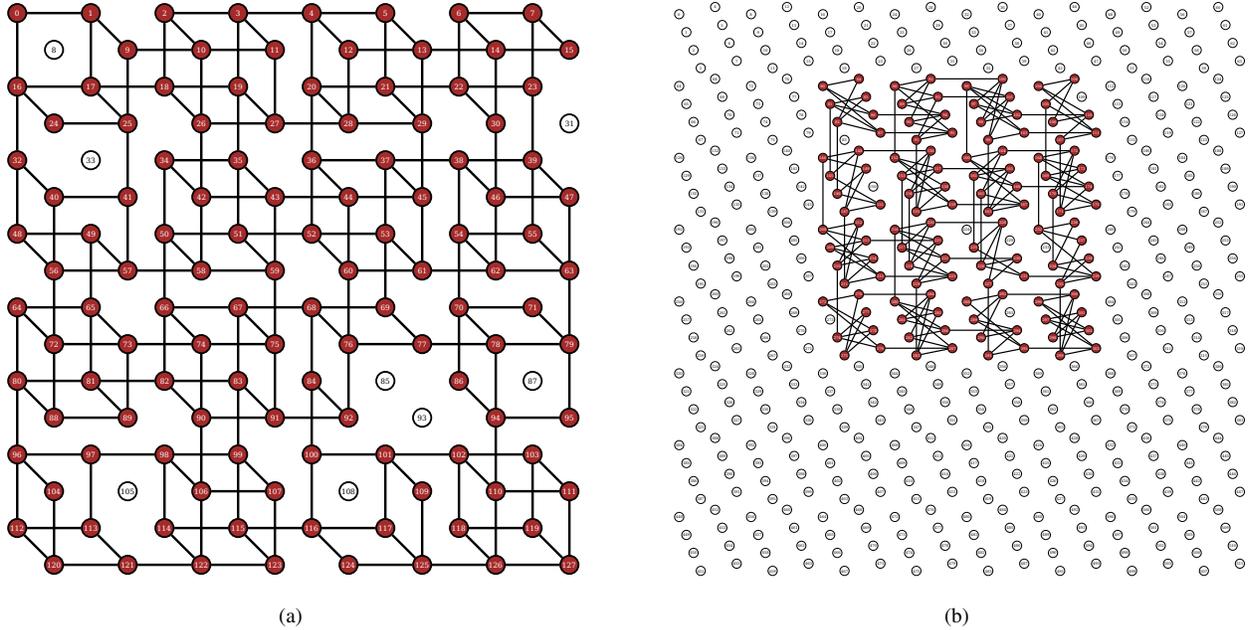


FIG. 10. (a) The largest subgraph of the 2LG (encoded graph of the square code) that can be directly embedded into the USC-ISI DW2 Chimera graph. To see that this is a subgraph, note that some of the edges depicted in Fig. 9(a) are missing, e.g., the edge between vertices 1 and 2. (b) Direct embedding of (a) into the Chimera graph. Each physical qubit in (b) corresponds to a logical qubit in (a).

interest, among other reasons, since it was perhaps the first shown to be NP-hard, for couplings in  $\{-1, 0, 1\}$  and no local fields (Barahona’s problem P3 [66]). This problem does not have a direct embedding in the Chimera graph, but it can be minor-embedded, as illustrated in Fig. 9(b).

The ME employs logical groups with pairs of physical

qubits connected by a ferromagnetic coupling. Logical couplings are represented by one or two physical couplings per physical qubit. Note that fewer than half of the physical couplings are used in this ME. The remaining physical couplings can be used to implement a QAC-ME scheme as shown in Fig. 9(c), using a new QAC “square” code we describe below.

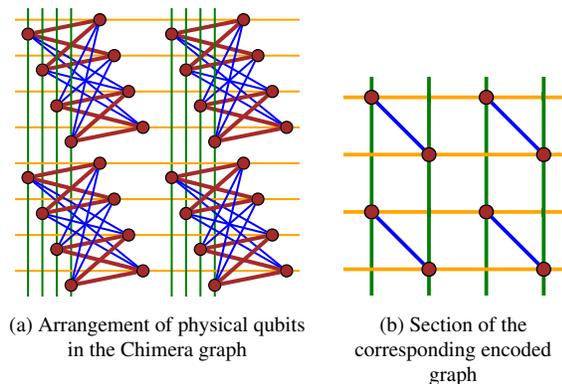


FIG. 11. (Color online) Construction of the “square”  $[[4, 1, 4]]_0$  code. Using the four physical qubits in the upper and the lower half of the Chimera unit cell, we construct two encoded qubits. In (a), the red (thick oblique) links represent the penalty couplings (stabilizers), while orange (horizontal), green (vertical) and blue (thin oblique) links form the logical Hamiltonian couplings. (b) The section of the encoded graph formed by (a). This is also the connectivity induced at level 1, where each link represents the collection of level 0 representations in (a).

### B. Direct embedding of the 2LG

Since we anticipate that minor embedding will perform poorly when compared to a direct embedding, we are also interested in studying a direct embedding of the 2LG. An interesting set of such problems includes all instances that can be defined on the subgraph of the encoded graph of the square code. An example is shown in Fig. 10(a), which depicts the largest 2LG subgraph directly embeddable into the USC-ISI DW2 Chimera graph. An example of a direct embedding of this subgraph is shown in Fig. 10(b).

### C. The square code and concatenation

We now describe the square code in detail. Using the notation  $[[n, k, d]]_p$  to denote a distance  $d$  code that uses  $n$  physical data qubits and  $p$  additional penalty qubits to encode  $k$  logical qubits, this code can be understood as a  $[[4, 1, 4]]_0$  code.<sup>6</sup>

The square code has a direct embedding into the Chimera graph, as shown in Fig. 11(a). This yields a 2LG encoded graph, as shown in Fig. 11(b), that can be extended as shown in Fig. 9(a), with extent bounded only by the size of the physical Chimera graph.

As indicated by the pairs of orange and green lines connecting pairs of encoded qubits in Fig. 11(a), each encoded qubit can couple to its neighbor via a  $\sigma^z \otimes \sigma^z$  coupling in two distinct ways, resulting in a boost by a factor of 2 at most. On

the other hand, there are 8 ways of realizing a blue logical  $\sigma^z \otimes \sigma^z$  coupling, but, since we wish to maintain the relative strength between qubit-qubit coupling terms, we use only 2 of these realizations, or alternatively we can use all 8 but with  $1/4$  the value of  $J$  per physical coupling. Similarly, there are 4 ways of realizing a logical  $\sigma^z$  operator by addressing each of the physical qubits comprising each encoded qubit, but we use only 2 of these realizations, or alternatively we can use all 4 but with half the value of  $h$  per qubit. Thus

$$\bar{h}_\ell^{(1)} = 2h_\ell, \quad \bar{J}_{\ell,\ell'}^{(1)} = 2J_{\ell\ell'}, \quad (26)$$

where we have interpreted the encoding as a level-1 code. This code induces an encoded 2LG graph of level-1 which we can exploit for concatenation. As a second layer of encoding, we use  $n \times n$  encoded level-1 qubits to form an encoded level-2 qubit on each plane of the 2LG, as depicted in Fig. 12. Note that with this encoding the 2LG graph structure of level-1 is maintained, and that an encoded level-2 qubit comprises  $4n^2$  physical qubits.

The level-2 encoding can correct  $t = 2n^2 - 1$  errors, i.e., has distance  $4n^2$  as expected from the code concatenation, and will be undecided if there are  $2n^2$  errors. To see why this is so, consider the level-2 code as a whole, and interpret it as a distance  $d = 4n^2$  qubit repetition code. A majority vote on the level-2 encoded 2LG graph can give the wrong result if more than half of the qubits are flipped, i.e., it can recover from  $2n^2 - 1$  flips, it is undecided for  $2n^2$  errors, and will give the wrong answer for more than  $2n^2$  errors. The level-2 code is thus a  $[[4n^2, 1, 4n^2]]_0$  code.

We note that an additional benefit of the encoding is the ability to fine-tune different couplings or local fields in the problem Hamiltonian. For example, suppose  $J_{ij}$  and  $h_i$  can only take values  $\{-1, 0, 1\}$ . How can we generate general problems where some couplings have fractional values? To do this we use the fact that at the level-2 encoding each coupling or local field has  $2n$  different realizations. It is now straightforward to rescale all couplings with respect to the largest coupling. E.g., suppose a given coupling  $J$  is  $9/10$  of the largest coupling  $J_{\max}$ , then with a sufficiently large  $n$ , one can turn off  $1/10$  of the realizations of  $J$  and thus achieve the desired  $J/J_{\max}$  ratio. The larger  $n$  is, the greater the accuracy of this ratio. If the quantum annealing device has other possible values of  $J_{ij}$  and  $h_i$ , then this reduces the necessary value of  $n$  for a desired ratio.

It is clear that the square code can be concatenated further. The resulting concatenated code can be made to have arbitrarily large distance, i.e., is capable of dealing with increasingly more errors, at the cost of more physical qubits per encoded qubit. It also exhibits a potentially beneficial growth in the encoded local fields and couplings. We describe this in Appendix B.

Decoding of the square code relies on the techniques we discussed in Sec. IV. In Appendix C we propose an alternative, recursive decoding scheme.

<sup>6</sup> The code introduced and used in Refs. [59, 60] is a  $[[3, 1, 3]]_1$  code in this notation.

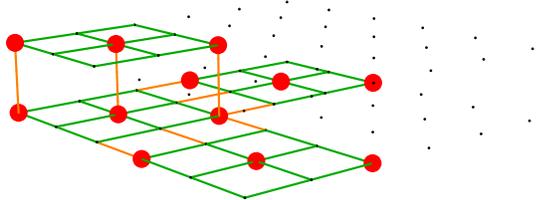


FIG. 12. (Color online) Level-2 encoding. Green (dark grey) lines represent level-1 couplings promoted to penalty couplings (stabilizers), while orange (light grey) lines represent level-1 encoded  $\sigma^z \otimes \sigma^z$  couplings (2 physical couplings). Each vertex represents a level-1 encoded qubit (4 physical qubits). Red (large) dots represent realizations of level-1  $\sigma^z$  encoded operators. The encoding uses  $n \times n$  level-1 encoded qubits, with the  $n = 3$  case depicted here.

## VI. EXPERIMENTAL RESULTS

### A. Frustrated-loop Hamiltonians with planted solutions

We now restrict our attention to a specific set of random Ising problem Hamiltonians with “planted” ground states [81, 82], which were recently introduced in the context of quantum annealing and studied in detail using the DW2 device [34], albeit without ME and QAC considerations. We briefly review how these problem Hamiltonians are constructed; for a detailed explanation see Ref. [34].

By “planting” we mean that the planted state is, by construction, a ground state of the problem Hamiltonian. The planted state need not be the unique ground state (in fact it is typically degenerate), but knowing it means that we also know the ground state energy. This advance knowledge circumvents the need to resort to computationally expensive exact solvers that would otherwise be required to determine such properties of the ground state. Without loss of generality we can always choose to plant the all-zero state:

$$|\Psi_0\rangle = |0_1, \dots, 0_N\rangle. \quad (27)$$

We generate random loops (cycles over the Chimera graph) with all ferromagnetic couplings, except one randomly chosen anti-ferromagnetic coupling. In this manner each loop is a frustrated Ising problem whose ground state is the all-zero state of the qubits covered by the loop. A problem instance is then built as a sum of such randomly generated frustrated

loops:<sup>7</sup>

$$H_P = \sum_i^L H_i^{\text{loop}}, \quad (28)$$

Similarly to the case of constrained satisfiability problems (SAT), the hardness of a problem can be tuned by varying the number of terms in the sum of Eq. (28) (i.e., the number of clauses in the terminology of SAT). Equivalently, hardness depends on a clause density  $\alpha$  defined as the ratio between the number of loops (clauses)  $L$  and the total number of qubits:

$$\alpha = \frac{L}{N}. \quad (29)$$

This turns out to generate problems with an easy-hard-easy structure as a function of the clause density. We loosely refer to the hardness peak as occurring at the critical clause density  $\alpha_c$ .<sup>8</sup> Typically, frustration and hardness peak both at the critical clause density; see Appendix D for an estimate of frustration.

While the length of each loop is arbitrary, here we only consider loops of length 4 and 6. Note that such short loops were not considered in Ref. [34] as they generated instances that were too hard (and thus resulted in insufficient solutions to be statistically meaningful) with a direct embedding on the full (then 503-qubit) Chimera graph. Instead Ref. [34] used loops of length  $\geq 8$ . Note further that recent work [39] has generated significantly harder instances for the Chimera graph, and we use the current method primarily due to its aforementioned convenient feature of advance knowledge of the ground state energy. Moreover, while the construction we have described obviously generates only a subset of the NP-hard spin glass problem instances that can be defined over the 2LG, it features the important advantage that the clause density is a tunable parameter that allows us to probe the quantum annealer on instances of varying and controllable hardness.

### B. Experimental methods

Our experiments were performed on the DW2 quantum annealing processor installed at the University of Southern California Information Sciences Institute (USC-ISI), which has been described in detail before (see, e.g., Ref. [33]). The D-Wave devices have been designed in such a way that complete logical graphs  $K_n$  of  $n$  vertices can be minor embedded into the Chimera graph representing the physical connectivity [see Fig. 1(a)] [62, 63]. The largest complete graph that can be minor embedded into the DW2 hardware graph is  $K_{32}$ .

<sup>7</sup> In somewhat confusing terminology, this Hamiltonian is “frustration-free” in the sense that the ground state of the total Hamiltonian is also the ground state of each (frustrated) term in the sum [83].

<sup>8</sup> We use the term “critical clause density” loosely here, in the sense that we do *not* mean to imply the presence of a phase transition, but merely to signal the empirically observed hardness peak.

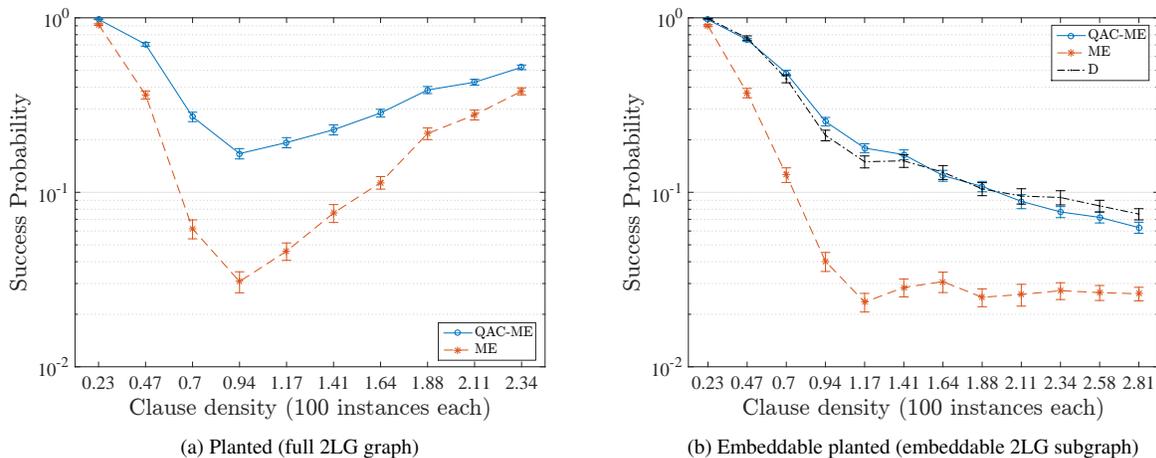


FIG. 13. (Color online) Mean success probabilities [after the renormalization of Eq. (31)] of instances generated by planting solutions on (a) the square code encoded graph, (b) on the embeddable subgraph. Decoding used the strategy of energy minimization. In (a) the comparison shown is between the ME of Fig. 9(b) and the QAC-ME of Fig. 9(c). QAC-ME is seen to improve the ME performance by almost an order of magnitude, with the largest improvement obtained at the critical clause density. In (b) the comparison shown is between D of Fig. 10(b), ME obtained by restricting Fig. 9(b) to the directly embeddable subgraph, and QAC-ME obtained by similarly restricting Fig. 9(c). It is seen that the performance loss of ME with respect to D is completely recovered by QAC-ME. Errors bars denote  $1\sigma$  confidence intervals here and in all subsequent plots.

Apart from specifying the problem Hamiltonian [Eq. (2)], the user interface enables us to choose the annealing time  $t_a$ . In all our experiments, the annealing time was set to the smallest value allowed by the hardware,  $t_a = 20\mu\text{s}$ . For each problem instance we ran  $10^4$  annealing cycles. This was done by performing 10 programming cycles of  $10^3$  runs each. For each programming cycle in which the D-Wave device is programmed with a given problem Hamiltonian, intrinsic control errors (ICE) prevent the physical couplings to be set with a precision better than  $\sim 5\%$  of the maximum allowed values  $|h_{\text{max}}|, |J_{\text{max}}|$ . This error includes both random (low frequency) noise and systematic errors (flux biases). Multiple programming cycles help to average over errors in setting the intended coupling. As an additional precaution designed to average out as much systematic error as possible, we performed random gauge transformations between each programming cycle. A gauge transformation is an invariance transformation of the Ising Hamiltonian Eq. (2) that is achieved by flipping the value of qubit  $i$  and by appropriately changing the sign of the couplings:

$$\sigma_i^z \rightarrow -\sigma_i^z \quad \Rightarrow \quad J_{ij} \rightarrow -J_{ij} \quad h_i \rightarrow -h_i, \quad \forall j. \quad (30)$$

We generated 100 random planted instances per clause density. Below we plot, for each  $\alpha$ , the mean success probability and its standard error (one  $\sigma$  deviation of the mean) for the optimal value of the penalties. Penalties were chosen and optimized according to the uniform strategy described in Sec. III A. The penalties were optimized separately for each  $\alpha$ . Decoding used the strategy of energy minimization over the broken qubits (we discuss encoding and decoding strategies in more detail in Sec. VII A).

We stress that all comparisons between the various strategies (i.e., ME, QAC-ME and D) must properly account for the physical resources required by each strategy. Specifically, ME requires twice as many physical qubits as D, while the level-1 square code (the only case we tested in this work) implementation of QAC-ME requires four times as many. This implies that one can run two ME or four D embeddings in parallel while consuming the same hardware resources (qubits) as needed for a single QAC-ME embedding. To account for this parallelism we renormalize and report the experimental D and ME ground state (“success”) probabilities as follows:

$$P_D \mapsto 1 - (1 - P_D)^4 \quad (31a)$$

$$P_{\text{ME}} \mapsto 1 - (1 - P_{\text{ME}})^2, \quad (31b)$$

which is the probability of finding the ground state at least once after four (D) or two (ME) parallel attempts, using the success probability from a single attempt (see Appendix E for a more detailed description of how the quantities above are computed from the raw data).

### C. ME vs QAC-ME for the 2LG graph

Figure 13(a) shows our experimental results for instances with planted solutions as described above and constructed on the entire 2LG graph, i.e., the encoded graph of the square code shown in Fig. 9(a). Fig. 13(a) compares the ME [corresponding to Fig. 9(b)] and QAC-ME [corresponding to Fig. 9(c)] implementations on the same set of instances. It shows that QAC-ME is effective in boosting the performance of the quantum annealer on the entire range of clause densi-

ties. Interestingly, the harder the problems, the larger is the advantage of QAC-ME over ME. We empirically identify the critical clause density as  $\alpha_{\text{crit}} \approx 0.94$ , where the success probability has a minimum. At this value QAC-ME yields success probabilities that are about one order of magnitude higher than ME's. Thus error correction becomes more effective for harder problems, as was also observed in the context of random Ising problems in Ref. [60].

#### D. D vs ME vs QAC-ME

We next compare between D and ME by constructing instances with planted solutions on the 2LG subgraph shown in Fig. 10(a). Experimental results are shown in Fig. 13(b). As expected, the ME strategy suffers a steady decay in performance compared to D, until  $\alpha \approx 1.17$  is reached. At this value, the normalized success probabilities for ME are an order of magnitude smaller than those for the D case. QAC-ME completely reverses the performance loss due to ME; the QAC-ME success probabilities are comparable to those obtained with the D implementation. This again demonstrates that QAC acts as a successful error correction strategy.

We note that there is no clear evidence of a critical clause density (or hardness peak) in this case. The ME strategy exhibits a plateau after  $\alpha \approx 1.17$ , while the D and QAC-ME success probabilities continue to decrease, with QAC-ME decreasing more slowly. Up to this value, the increase in hardness is intrinsic and can be related to the increase in frustration. It may be possible to associate the further increase in hardness for larger clause densities to a precision effect, e.g., an increase in the range (the ratio between the largest and smallest couplings in the problem Hamiltonians; see Appendix D for a quantification of this effect). This adversely affects the performance of the DW2 because a large coupling range means that some couplings are set to smaller values than can be physically implemented (the DW2 has a range of about 4 bits), due to the well-studied effect of ICE that leads to the specification of the wrong problem Hamiltonian [17, 33, 35, 36, 77]. The saturation of the ME success probabilities can be understood as being due to having arrived at this precision limit before D and QAC-ME, i.e., for  $\alpha > 1.17$  the ME results are dominated by ICE, and the improvement due to QAC is at least in part due to countering ICE. The D implementation is, of course, also susceptible to ICE, but is impacted less than ME due to the absence of broken qubits in D.

#### E. Optimized uniform vs nonuniform penalties

In Sec. III we discussed optimizing the energy penalties and proposed a new nonuniform strategy. In this section we experimentally test the effectiveness of this strategy and compare it to the previously proposed uniform strategy. Specifically, we compare the performance of the square code when penalties are chosen according to either Eq. (17) (uniform case) or

Eq. (18) (nonuniform case), where the overall penalty energy scale  $\gamma$  is a parameter that must be optimized independently.

Intuitively, a nonuniform choice of penalties should perform best when the logical couplings are distributed nonuniformly on the logical graph. The problem instances we have considered so far were generated by adding frustrated loops that are uniformly distributed over the logical graph, resulting in similarly uniformly distributed couplings. We thus consider two additional sets of nonuniform instances.

##### 1. Weighted planted frustrated-loop instances

The first set of instances is a simple modification of the construction using frustration-free Hamiltonians described above. The only difference compared to Eq. (28) is that here we consider a weighted sum of the same set of frustrated loops as used to generate Fig. 13, with the weights varying nonuniformly on the logical graph:

$$H_P = \sum_i^L w_i H_i^{\text{loop}}, \quad (32)$$

We choose the  $w_i$  as follows. We define  $\{\bar{x}_i, \bar{y}_i, \bar{z}_i\}$  to be the ‘‘center’’ of loop  $i$  on the 2LG (defined as the integer part of the mean of the coordinates of the spins comprising the loop), and set  $w_i = \bar{x}_i$ . This increases the typical strength of the couplings by a factor of 8 between the two opposite sides of the  $8 \times 8 \times 2$  2LG logical graph embedded into the DW2. This introduces a spatial non-uniformity and also reduces the magnitude of the smallest couplings of this set of instances (see Appendix D), thus making the instances significantly harder, as we shall see below. Note that the same planted solution as before ( $|\Psi_0\rangle$ ) applies, independently of the choice of the weights.

##### 2. Deformed embeddable instances

The second set are nonuniform instances we generate by ‘‘deformation’’ of the set of instances with planted solutions defined on the embeddable subgraph [Fig. 10(a)], as follows. For each instance we randomly pick a subset  $\bar{I}$  of encoded qubits (vertices of  $\bar{G}$ ) and apply a linear transformation (shifting and compression) to the encoded couplings such that:

$$J_{\bar{i}, \bar{j}} \rightarrow J_{\bar{i}, \bar{j}}/3 \quad \forall \bar{i} \in \bar{I}, \forall \bar{j} \notin \bar{I} \quad (33a)$$

$$J_{\bar{i}, \bar{j}} \rightarrow (2 \text{sign}(J_{\bar{i}, \bar{j}}) + J_{\bar{i}, \bar{j}})/3 \quad \forall \bar{i}, \bar{j} \notin \bar{I}. \quad (33b)$$

This partitions the encoded qubits into three sets  $\bar{I}, \bar{J}, \bar{K}$ , where the set  $\bar{I}$  (the one picked) is connected only through coupling that have been transformed according to Eq. (33a) (small couplings),  $\bar{K}$  is connected through couplings that have been transformed according to Eq. (33b) (large couplings), and  $\bar{J}$  is connected through couplings that have been transformed according to both Eqs. (33a) and (33b).

The average of the encoded couplings is  $\sim 0.5$ . Using

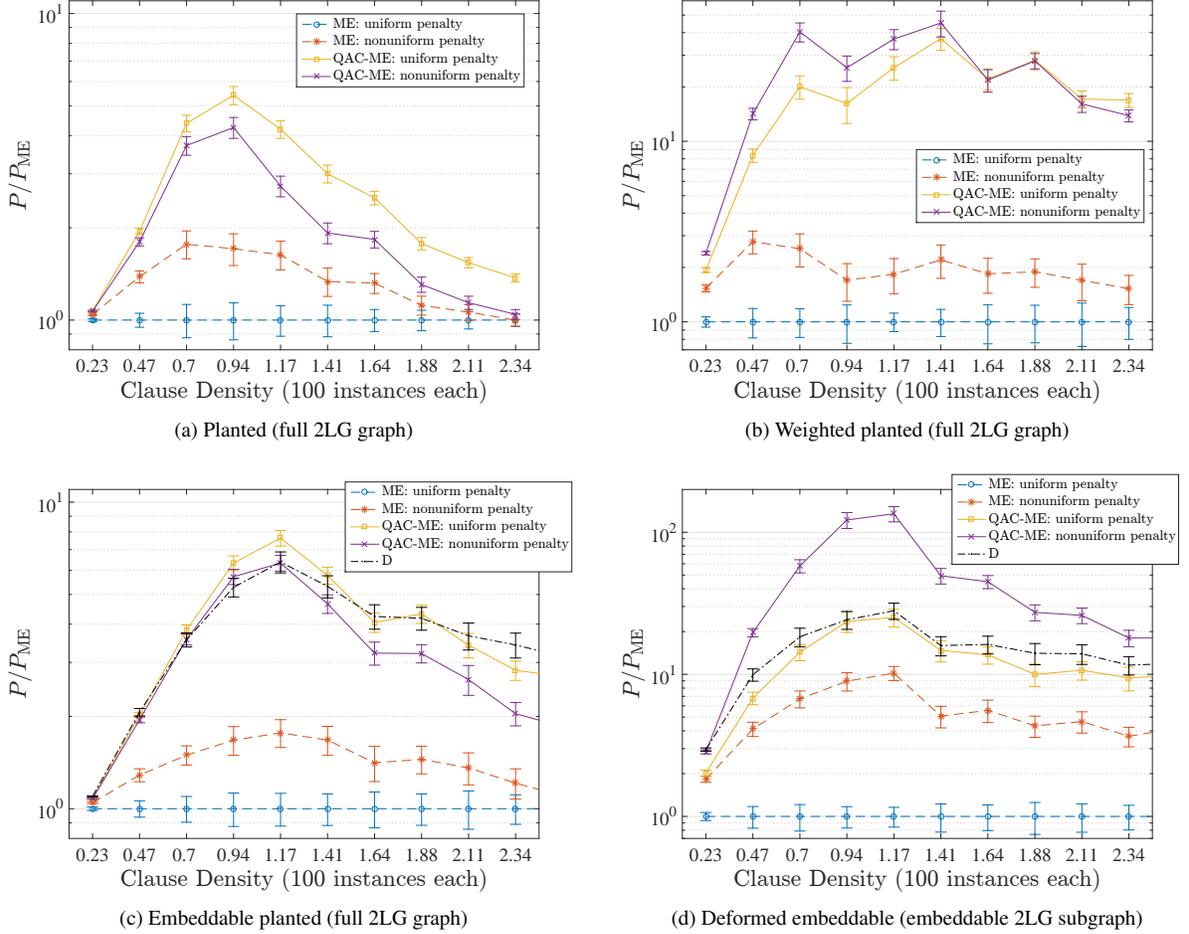


FIG. 14. (Color online) Success probabilities of ME and QAC-ME for uniform and nonuniform penalties for the four types of problem instances, normalized to the success probability of ME with uniform penalties. Decoding used the strategy of energy minimization. In all cases ME with uniform penalties performs worst, across the entire range of clause densities. ME with nonuniform penalties improves the ME performance somewhat, but always less than QAC-ME with uniform penalties. QAC-ME with nonuniform penalties does worse than QAC-ME with uniform penalties when the instances have uniform couplings [(a),(c)], but almost always does better than QAC-ME with uniform penalties when the instances have non-uniform couplings, as in (b) and (d). In the deformed embeddable case (d) QAC-ME with nonuniform couplings does significantly better even than the direct embedding.

this value in Eqs. (33a) and (33b) yields average coupling strengths of  $1/6$ ,  $1/2$  and  $5/6$  for the sets  $\bar{I}$ ,  $\bar{J}$ ,  $\bar{K}$  respectively. We randomly pick 35 qubits out of the available 120 for the  $\bar{I}$  set, which yields three similarly sized sets. Under this deformation, the originally planted configuration is no longer necessarily the ground state. In this case, we use the exact solver that uses a bucket tree elimination strategy (BTE) and is included in the D-Wave user libraries to determine the exact ground state energy. This exact solver is specifically designed to solve small, Chimera-structured problems.

Figure 14 shows the experimental results for ME and QAC-ME, with both uniform and nonuniform penalties. The panels show the four sets of instances: planted and weighted planted (both defined on the full 2LG), embeddable planted and deformed embeddable (defined on the embeddable subgraph). In all panels of Fig. 14 we see that QAC-ME outperforms

ME. The choice of nonuniform penalties improves the success probabilities of both the ME and QAC-ME in the case of nonuniform instances [panels (b) and (d)]. Moreover, we see that the nonuniform choice of penalties improves the ME implementation in the uniform case too, while it does not improve the QAC-ME scheme, suggesting that QAC-ME with a uniform penalty is already close to optimal in the case of uniform instances. (The dependence of the optimal penalty strength on clause density is shown in Appendix F.)

As a final remark of this section, note the large boost of nearly two orders in magnitude in the success probabilities that QAC-ME achieves over ME in the weighted case shown in Fig. 14(b). In this case QAC-ME is particularly successful. A similar advantage of QAC-ME over ME is observed in Fig. 14(d), where the advantage of QAC-ME with nonuniform penalties is the most pronounced, easily outperforming

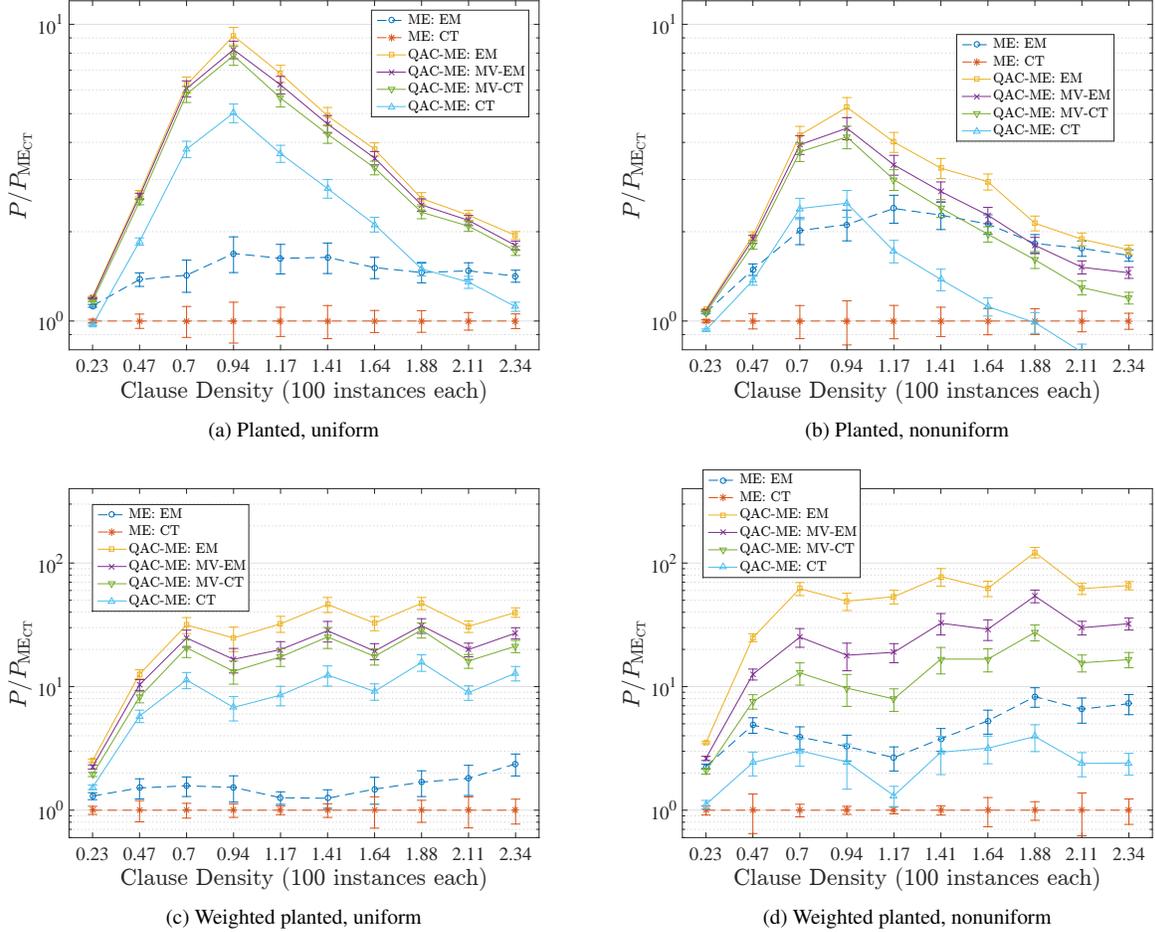


FIG. 15. (Color online) Comparison between different decoding strategies for the ME and QAC-ME implementations, for planted [(a) and (b)] and weighted planted [(c) and (d)] instances. Panel (a) and (c) [(b) and (d)] is for the case of uniform [nonuniform] penalties. The success probabilities plotted are normalized to the ME case with random decoding ( $P_{\text{ME}_{\text{CT}}}$ ). Panels demonstrate the improvement of the decoded results when going from CT to MV and EM. In the ME case, the EM strategy resoundingly beats CT. In the QAC-ME case the ordering is  $\text{EM} > \text{MV-EM} > \text{MV-CT} > \text{CT}$ , with the advantage of the best strategy (EM) being more significant in the nonuniform case.

the direct embedding as well [in contrast to the tie in the embeddable planted case, seen in Fig. 13(b)].

## VII. IMPACT OF THE DECODING STRATEGY

We now consider the impact of various decoding strategies defined in Sec. IV on performance as measured in terms of success probabilities. As already discussed in Sec. IV B, because both the ME and QAC-ME implementations of the square code involve an even number of physical qubits (two and four respectively), the possibility of ties means that broken qubits can be undecodable by majority vote. In the ME case, in particular, every broken qubit is indeed a tie. In the QAC-ME case there are two possibilities: “partially” broken qubits (decodable through majority vote) and ties.

We thus consider the following decoding strategies:

### 1. ME:

- (a) Decode ties randomly via coin tossing (CT);
- (b) Decode ties via energy minimization (EM).

### 2. QAC-ME:

- (a) Decode all broken qubits via CT;
- (b) Decode all broken qubits via energy minimization (EM);
- (c) Decode partially broken qubits via MV and ties via CT (MV-CT);
- (d) Decode partially broken qubits via MV and ties via EM (MV-EM);
- (e) Decode partially broken qubits via MV and ties via EM. Randomly select a number of unbroken qubits equal to the number of partially broken qubits and decode via EM (MV-EM(R)).

### A. Performance of the various decoding strategies

While we do not expect random decoding (the CT strategy) to work well, it sets a useful baseline against which we can compare our other decoding strategies. The comparison between all strategies in the case of instances defined on the full 2LG graph is shown in Figs. 15 and 16. Figure 15 is for the planted and weighted planted instances and reports results for both the ME and QAC-ME cases. Figure 16 focuses on the MV-EM(R) decoding strategy (QAC-ME only). Complementary results for the embeddable planted and the deformed embeddable cases are shown in Appendix G.

As expected, CT is the least effective decoding strategy, and is substantially outperformed by MV. The optimal choice is always the EM strategy (recall Sec. IV C), which yields improvements over both MV and MV-EM. The benefit of EM is larger in the case of nonuniform penalties and weighted instances. In Fig. 15(d) we see that EM decoding improves the success probabilities by a factor of about 3 as compared to the second best decoding strategy, MV-EM. The advantage of EM over MV and CT is directly related to the number of broken qubits at the optimal penalty value. The larger this number, the more bit flip errors we expect to correct during the decoding process by doing energy minimization. As we shall see in the next section, the typical number of broken qubits is usually larger for the nonuniform penalties and for the nonuniform instances we have constructed.

It is interesting to check whether unbroken qubits (where all physical qubits agree), which cannot be decoded via MV, can be further improved by EM. The MV-EM(R) decoding strategy is designed to probe this, by singling out a fraction of the unbroken qubits and applying the EM strategy to them. We thus compare MV-EM(R) to MV-EM and EM decoding in Fig. 16, again for the planted and weighted planted instances defined on the full 2LG graph, with nonuniform penalties. We find that the MV-EM(R) strategy offers no advantage over MV-EM, leaving EM as the superior strategy. Since MV-EM(R) is essentially indistinguishable from MV-EM, we may conclude that the unbroken qubits make a negligible contribution to the error rate.

### B. Broken qubit probability and the percolation threshold for efficient decoding

Following the discussion in Sec. IV D, efficient optimal decoding via the EM strategy is possible if the probability  $p_{\text{BQ}}$  to find a broken encoded qubit is smaller than the percolation threshold  $p_{2\text{LG}}$  of the 2LG. This ensures that no large clusters of broken qubits are formed that would make energy minimization too computationally costly. While we do not know the value of the percolation threshold of the (infinite) 2LG, it must certainly be bounded between the thresholds  $p_{\text{cubic}}$  and  $p_{\text{square}}$  for cubic and square lattices, respectively [84, 85]:

$$0.3116 = p_{\text{cubic}} < p_{2\text{LG}} < p_{\text{square}} = 0.5927. \quad (34)$$

We may further expect  $p_{2\text{LG}}$  to be closer to  $p_{\text{square}}$  than to  $p_{\text{cubic}}$  since the 2LG is geometrically more similar to a square lattice. Recall that the 2LG is the encoded graph of the square code, and hence we identify the latter's percolation threshold with  $p_{2\text{LG}}$ .

We first check that the experimental probability  $p_{\text{BQ}}$  to find broken qubits in the square code depends only weakly on the position of the encoded qubits. This is confirmed (for the QAC-ME case) in Fig. 17 at the optimal penalty value for a fixed clause density, and justifies the assumption underlying our percolation threshold argument about the efficiency of the EM decoding strategy in Sec. IV D.

Figure 18 shows the probability of a broken qubit as a function of the clause density, for both ME and QAC-ME, in the case of instances with planted and weighted planted solutions on the full 2LG, for both uniform and nonuniform penalties. The most important finding is that in all cases, the experimental probability of a broken logical or encoded qubit is well below  $p_{\text{cubic}}$ , and hence below the percolation threshold of the square code. This suggests that scalable and efficient decoding of the square code through energy minimization is possible, at least for the (frustrated) class of instances considered here. Additional supporting arguments are given in Appendix H using a worst-case analysis.

There are a few other interesting observations we can make about Fig. 18. First, the very fact that the observed  $p_{\text{BQ}}$  is non-vanishing is the reason that a decoding of ME and QAC-ME is necessary. Second, note that in both the ME and QAC-ME cases, the optimal penalty values tend to keep the probability  $p_{\text{BQ}}$  of broken qubits fairly constant over the range of clause densities  $\alpha$  we have studied. The optimal penalty values can thus be converted into an optimal broken qubits probability that depends on the set of instances chosen and on the implementation (ME or QAC-ME) considered. Third, note that  $p_{\text{BQ}}$  is typically smaller for ME than for QAC-ME, which is unsurprising given that the QAC step in QAC-ME adds a second way for qubits to break, in addition to ME. Fourth, note that  $p_{\text{BQ}}$  is typically larger with nonuniform penalties or weighted instances. The intuitive reason for this is that nonuniform penalties are by design weaker for qubits that are less strongly coupled. This facilitates the appearance of more broken qubits. This is consistent with the experimental finding of the previous section that EM decoding results in a larger advantage in these cases, simply because there are more broken qubits to be optimized. Finally, note that the differences in performance of the various decoding strategies considered in Sec. VII A can be related to  $p_{\text{BQ}}$ .

## VIII. TESTING TEMPERATURE AND NOISE DEPENDENCE VIA SIMULATED QUANTUM ANNEALING

To gain a better understanding of the mechanisms that are most responsible for the experimental success of QAC, we conducted numerical simulations that illuminate the role of temperature and ICE. A master equation simulation treating the device as an open quantum system [29] is unfeasible due to the large number of qubits involved in our experiments.

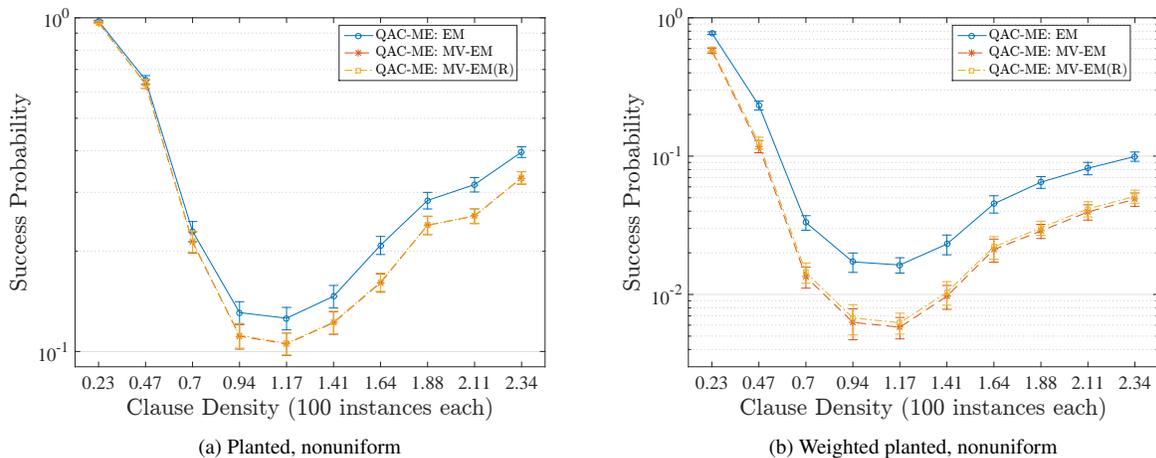


FIG. 16. (Color online) Comparison between different decoding strategies for the QAC-ME implementation with nonuniform penalties, for (a) planted and (b) weighted planted instances. The MV-EM(R) strategy, where a random set of unbroken qubits is subjected to energy minimization, is shown along with MV-EM and EM. It can be seen that MV-EM(R) does no better than MV-EM. This implies that logical errors are mostly concentrated on broken qubits. The EM data shown here is the same as for the QAC-ME case in Figs. 15(b) and 15(d).

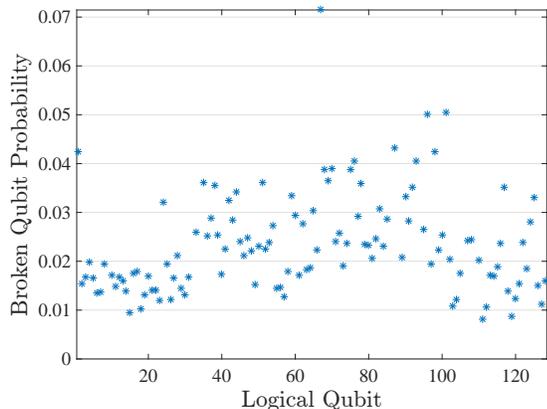


FIG. 17. Broken qubit median probabilities for the “planted” instance set at the critical clause density  $\alpha = 0.94$ . Results are shown for the full set of 100 instances implemented with the QAC-ME scheme, for the optimal value of the (uniform) penalty  $\gamma = 0.2$ . The abscissa is the index of the encoded qubits on the 2LG graph shown in Fig. 9(a).

Hence we resort to a semi-dynamical approach based on quantum Monte Carlo. The approach, called simulated quantum annealing (SQA) [17, 86, 87], attempts to mimic the evolution of an open quantum system where stochastic processes (typical of a noisy regime) dominate the dynamics, and essentially provides the instantaneous Gibbs state along the annealing evolution (see Appendix I for details). This approach has successfully reproduced the success probability distribution of the D-Wave devices for random Ising instances [17], although it does not provide a complete description when one also accounts for ground state degeneracy and excited states

[21], or open system dynamics [88].

Our strategy is not to use SQA as a faithful model for the DW2 device. We thus do not attempt to tune the SQA parameters to find a best fit for the experimental results. Rather, we use SQA as a numerical tool to explore the response of our QAC-ME scheme to changes of various important quantities that cannot be controlled experimentally, specifically temperature and noise. The noise level  $\chi$  is the standard deviation of statistically independent Gaussian noise added to all physical couplings  $J_{ij}$  and local fields  $h_i$ , and is introduced to reproduce ICE. We performed simulations for different values of the number of sweeps (a sweep is a single update of all spins), but found that the results have a very weak dependence on this number, so it is not a crucial parameter in our analysis. On the other hand, the temperature and the noise level heavily affect our numerical results. Figure 19 summarizes our numerical results for the “embeddable planted” instances [compare with the experimental results shown in Fig. 13(b)]. It shows that QAC-ME becomes competitive with respect to ME and D only for sufficiently large temperature and/or noise. We regard this result as strong evidence that the QAC schemes are particularly efficient in reducing the amount of errors due to thermal noise and imprecisions in parameter settings.

## IX. CONCLUSIONS AND DISCUSSION

In order to exploit the potential advantages of quantum annealing over classical solvers in solving “non-native” optimization problems they must accommodate a minor embedding (ME) step that maps the given problem to the given device hardware graph. In this work we introduced a new quantum annealing correction (QAC) scheme that is compatible with the minor embedding step—the “square code”—and studied its efficacy experimentally using the DW2 device. We

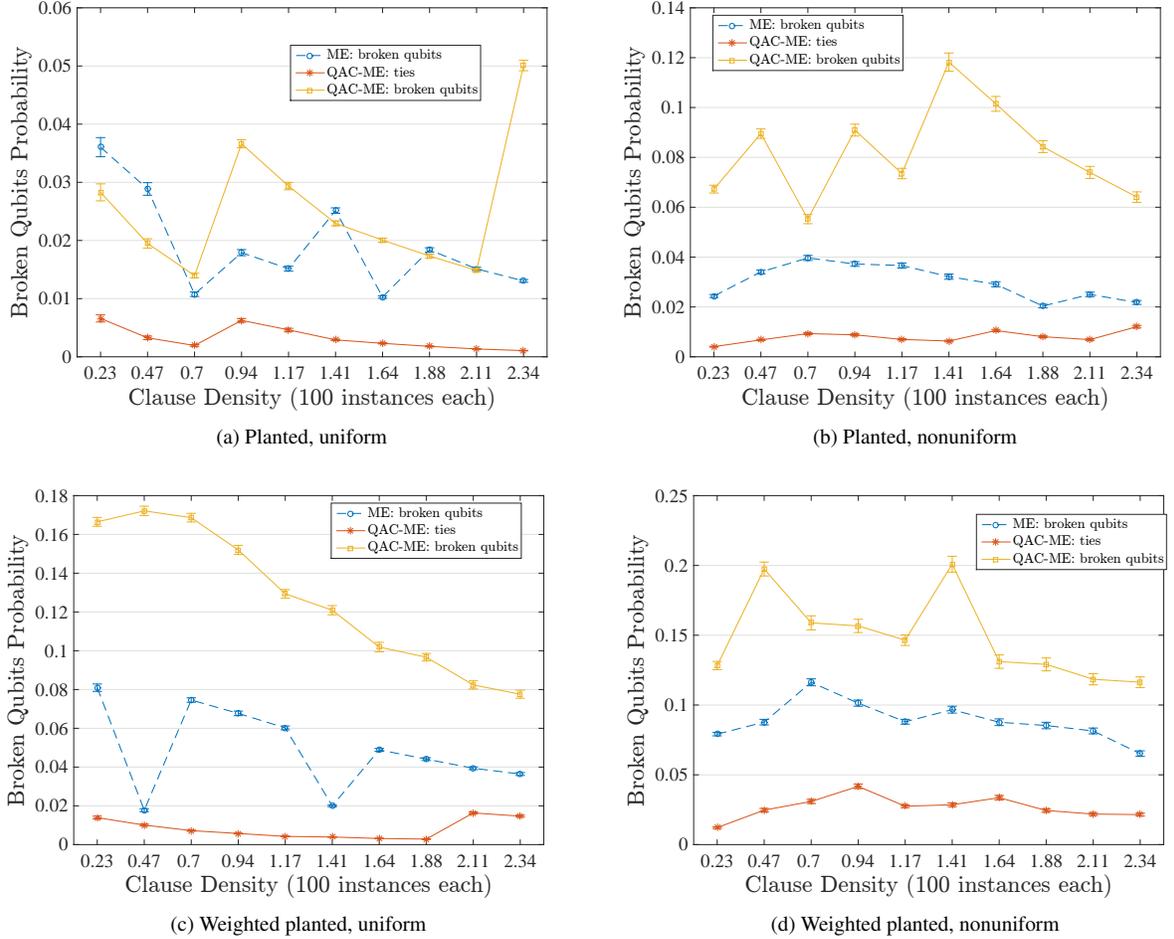


FIG. 18. (Color online) Probability  $p_{BQ}$  of broken qubits in ME and QAC-ME, as a function of the clause density, for planted (top) and weighted planted (bottom) problem instances defined on the full 2LG graph. Left: uniform penalties. Right: nonuniform penalties. ME almost always has a lower  $p_{BQ}$  than QAC-ME. Also shown is the probability of tied qubits in the QAC-ME case (which is already included in the broken qubits probability).

demonstrated that this hybrid QAC-ME scheme significantly boosts performance by applying it to frustrated Ising problems with known (“planted”) ground states. In the construction of such instances, a controllable clause density parameter  $\alpha$  tunes both hardness and frustration. The largest performance boost (measured in terms of the success probability of finding the ground state) via QAC-ME occurs at the critical clause density  $\alpha_c$ , where hardness peaks.

We have also considered problem instances that are not only embeddable into the square code but also directly into the physical “Chimera” hardware graph of the DW2 device. This allows for a comparison of the direct, ME and QAC-ME embeddings for the same problem instance. As expected ME behaves poorly with respect to the direct embedding. The performance improvement achieved with QAC-ME is particularly significant at  $\alpha_c$  and suffices to match that of the direct embedding. This is a proof-of-concept that QAC-ME can (and should) be used to overcome the limitations of minor embedding on quantum annealers. QAC-ME thus has the potential

to extend the competitiveness of these devices beyond natively embeddable optimization problems.

Several questions remain to be addressed in future work, requiring the use of next-generation quantum annealers. The most important of these is arguably the scaling of error-corrected quantum annealing. To properly address this question we need to be able to embed larger logical problems (the square code considered in this work implements up to 120 encoded qubits) on devices with a larger number of physical qubits. An important aspect for any scaling study is the determination of the “optimal” implementation. In the case of a direct embedding, this boils down to the question of determining the optimal annealing time (see, e.g., Ref. [34]). Optimality in minor embedded quantum annealing is a broader question. This is due to the freedom in choosing the minor embedding and, once this is given, the exact map between the logical and the physical couplings. In particular, the performance of minor embedded quantum annealing is sensitive to the choice of energy penalties, and determining the optimal choice of penal-

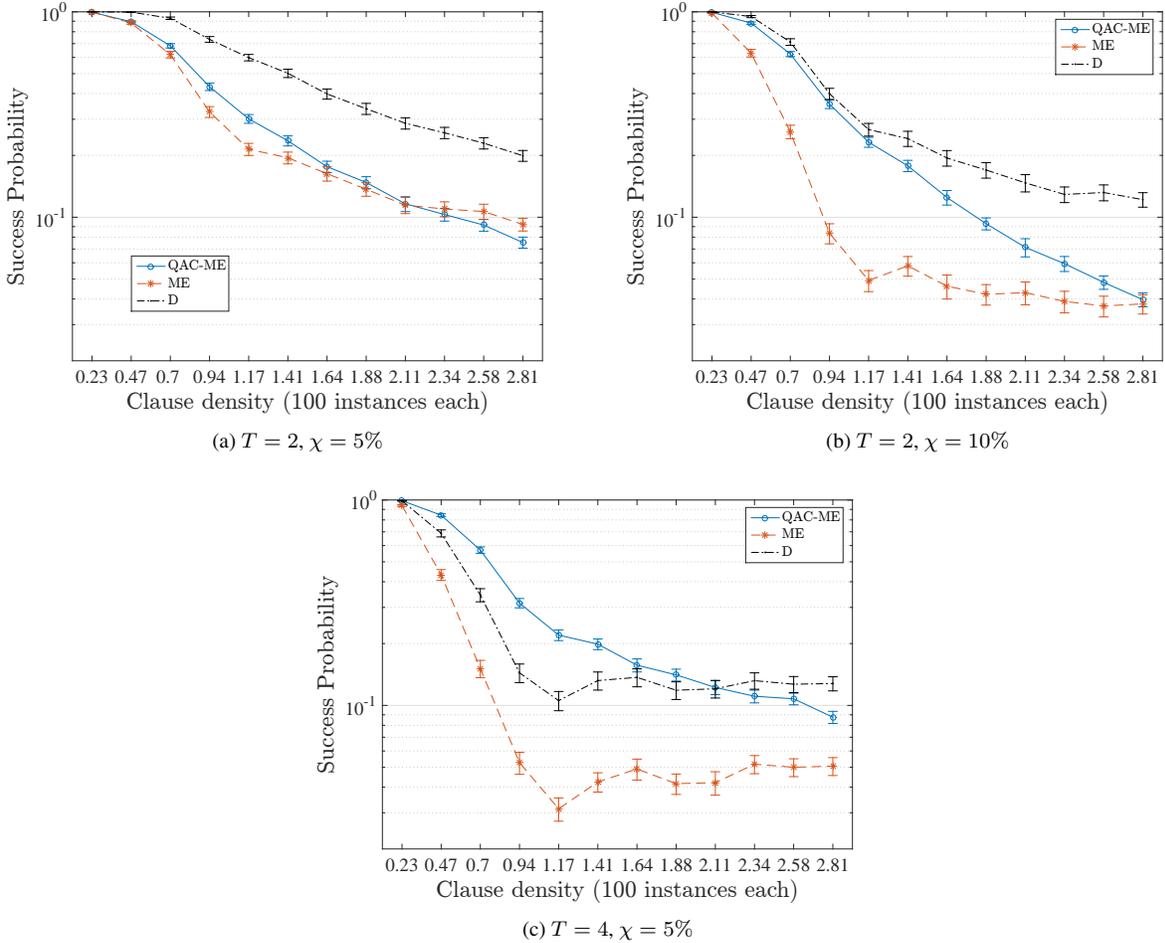


FIG. 19. (Color online) SQA simulation results for the “embeddable planted” instances class. (a) Relatively low temperature  $T$  and low noise strength  $\chi$ : QAC-ME does not provide an advantage over the uncorrected ME approach. Note how success probabilities are typically larger for SQA in this case with respect to the experimental results [Fig. 13(b)]. (b) Relatively low temperature and high noise strength: QAC-ME gains a clear advantage over ME but is still inferior to D. (c) Relatively high temperature and low noise strength: QAC-ME gives a consistent advantage over ME, and it is competitive with D. This situation is the most similar to the experimental results. We used 20,000 sweeps in all cases shown. Temperature is in units of  $B(t_f)/k_B$ .

ties is important both for determining the scaling of physical quantum annealing and for optimizing the performance on practical applications. We have shown that choosing nonuniform penalties that depend on the local strength of the logical couplings can give a significant advantage with respect to a uniform choice.

The optimal choice of the decoding procedure is another important question. We found global energy minimization over broken qubits to be the optimal choice for our problem instances. We showed that in order for this decoding technique to be performed at a negligible computational cost, the typical probability to have broken logical or encoded qubits should be lower than the per-site percolation threshold of the logical or encoded graph. It turns out that the square code, for the frustrated problems we have considered, is efficiently and optimally decodable through energy minimization even when its size is scaled. Our results demonstrate that the square code

on the tested DW2 lies in a “decodable” regime, with a sufficiently small number of broken qubits when the penalties are optimized. It is important to study whether optimal decoding can be performed efficiently in other ME and QAC-ME schemes. Whenever this is not possible, efficient, albeit suboptimal, strategies should be developed. To this end, it is likely that techniques developed in the decoding of Low Density Parity Check (LDPC) codes [89] will find important applications in quantum annealing too. In particular, suboptimal decoding through belief propagation [90] is a promising approach to successfully decode ME quantum annealing when the broken qubit probability is larger than the decoding threshold.

More generally, it is important to understand how the performance of minor embedded quantum annealing degrades with the increase of the size of the logical or encoded qubits. It is likely that the larger the logical qubits required for a given

minor embedding are (this will be the case for increasingly larger logical problems) the more important QAC-ME will be as a tool to prevent the degradation of the performance of quantum annealers on minor embedded implementations.

While our experimental conclusions are naturally limited to the particular set of problem instances we studied, we may extrapolate that our results more generally support the effectiveness of the QAC-ME strategy. Confirming this can have important implications in scaling quantum annealing to solve larger and harder optimization problems of practical interest, and we expect QAC-ME techniques to play a central role in achieving this goal.

## X. ACKNOWLEDGEMENTS

We thank Anurag Mishra and Leonid Pryadko for helpful discussions. Access to the D-Wave Two was made avail-

able by the USC-Lockheed Martin Quantum Computing Center. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. Part of the computing resources were provided by the USC Center for High Performance Computing and Communications. I.H. and D.A.L. acknowledge support under ARO grant number W911NF-12-1-0523. The work of W.V., T.A. and D.A.L. was supported under ARO MURI Grant No. W911NF-11-1-0268.

- 
- [1] D. Bacon and W. van Dam, *Commun. ACM* **53**, 84 (2010).
- [2] A. W. Harrow, arXiv:1501.00011 (2014).
- [3] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, arXiv:quant-ph/0001106 (2000).
- [4] A. B. Finnila, M. A. Gomez, C. Sebenik, C. Stenson, and J. D. Doll, *Chemical Physics Letters* **219**, 343 (1994).
- [5] T. Kadowaki and H. Nishimori, *Phys. Rev. E* **58**, 5355 (1998).
- [6] J. Brooke, D. Bitko, T. F. Rosenbaum, and G. Aeppli, *Science* **284**, 779 (1999).
- [7] J. Brooke, T. F. Rosenbaum, and G. Aeppli, *Nature* **413**, 610 (2001).
- [8] W. M. Kaminsky and S. Lloyd, in *Quantum Computing and Quantum Bits in Mesoscopic Systems*, edited by A. Leggett, B. Ruggiero, and P. Silvestrini (Kluwer Academic/Plenum Publ., 2004) arXiv:quant-ph/0211152.
- [9] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, C. J. S. Truncik, S. Uchaikin, J. Wang, B. Wilson, and G. Rose, *Nature* **473**, 194 (2011).
- [10] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, *Science* **292**, 472 (2001).
- [11] M. W. Johnson, P. Bunyk, F. Maibaum, E. Tolkacheva, A. J. Berkley, E. M. Chapple, R. Harris, J. Johansson, T. Lanting, I. Perminov, E. Ladizinsky, T. Oh, and G. Rose, *Superconductor Science and Technology* **23**, 065004 (2010).
- [12] A. J. Berkley, M. W. Johnson, P. Bunyk, R. Harris, J. Johansson, T. Lanting, E. Ladizinsky, E. Tolkacheva, M. H. S. Amin, and G. Rose, *Superconductor Science and Technology* **23**, 105014 (2010).
- [13] R. Harris, M. W. Johnson, T. Lanting, A. J. Berkley, J. Johansson, P. Bunyk, E. Tolkacheva, E. Ladizinsky, N. Ladizinsky, T. Oh, F. Cioata, I. Perminov, P. Spear, C. Enderud, C. Rich, S. Uchaikin, M. C. Thom, E. M. Chapple, J. Wang, B. Wilson, M. H. S. Amin, N. Dickson, K. Karimi, B. Macready, C. J. S. Truncik, and G. Rose, *Phys. Rev. B* **82**, 024511 (2010).
- [14] S. Suzuki and A. Das (guest eds.), *Eur. Phys. J. Spec. Top.* **224**, 1 (2015).
- [15] S. Aaronson, “D-Wave: Truth finally starts to emerge,”
- [16] S. Boixo, T. Albash, F. M. Spedalieri, N. Chancellor, and D. A. Lidar, *Nat. Commun.* **4**, 2067 (2013).
- [17] S. Boixo, T. F. Ronnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer, *Nat. Phys.* **10**, 218 (2014).
- [18] J. A. Smolin and G. Smith, arXiv:1305.4904 (2013).
- [19] L. Wang, T. F. Ronnow, S. Boixo, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer, arXiv:1305.5837 (2013).
- [20] S. W. Shin, G. Smith, J. A. Smolin, and U. Vazirani, arXiv:1401.7087 (2014).
- [21] T. Albash, T. F. Ronnow, M. Troyer, and D. A. Lidar, *Eur. Phys. J. Spec. Top.* **224**, 111 (2015).
- [22] T. Albash, W. Vinci, A. Mishra, P. A. Warburton, and D. A. Lidar, *Physical Review A* **91**, 042314 (2015).
- [23] S. W. Shin, G. Smith, J. A. Smolin, and U. Vazirani, arXiv:1404.6499 (2014).
- [24] P. J. D. Crowley, T. urić, W. Vinci, P. A. Warburton, and A. G. Green, *Physical Review A* **90**, 042317 (2014).
- [25] T. Albash and D. A. Lidar, arXiv:1503.08767 (2015).
- [26] K. Kechedzhi and V. N. Smelyanskiy, arXiv:1505.05878 (2015).
- [27] T. Lanting, A. J. Przybysz, A. Y. Smirnov, F. M. Spedalieri, M. H. Amin, A. J. Berkley, R. Harris, F. Altomare, S. Boixo, P. Bunyk, N. Dickson, C. Enderud, J. P. Hilton, E. Hoskinson, M. W. Johnson, E. Ladizinsky, N. Ladizinsky, R. Neufeld, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, S. Uchaikin, A. B. Wilson, and G. Rose, *Phys. Rev. X* **4**, 021041 (2014).
- [28] S. Boixo, V. N. Smelyanskiy, A. Shabani, S. V. Isakov, M. Dykman, V. S. Denchev, M. Amin, A. Smirnov, M. Mohseni, and H. Neven, arXiv:1411.4036 (2014).
- [29] T. Albash, S. Boixo, D. A. Lidar, and P. Zanardi, *New J. of Phys.* **14**, 123016 (2012).
- [30] N. G. Dickson, M. W. Johnson, M. H. Amin, R. Harris, F. Altomare, A. J. Berkley, P. Bunyk, J. Cai, E. M. Chapple, P. Chavez, F. Cioata, T. Cirip, P. deBuen, M. Drew-Brook, C. Enderud, S. Gildert, F. Hamze, J. P. Hilton, E. Hoskinson, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Lanting, T. Mahon, R. Neufeld, T. Oh, I. Perminov, C. Petroff, A. Przybysz, C. Rich, P. Spear, A. Tcaciuc, M. C. Thom, E. Tolkacheva,

- S. Uchaikin, J. Wang, A. B. Wilson, Z. Merali, and G. Rose, Nat. Commun. **4**, 1903 (2013).
- [31] R. D. Somma, D. Nagaj, and M. Kieferová, Phys. Rev. Lett. **109**, 050501 (2012).
- [32] T. Jörg, F. Krzakala, J. Kurchan, A. C. Maggs, and J. Pujos, EPL (Europhysics Letters) **89**, 40004 (2010).
- [33] T. F. Rønnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer, Science **345**, 420 (2014).
- [34] I. Hen, J. Job, T. Albash, T. F. Rønnow, M. Troyer, and D. Lidar, arXiv:1502.01663 (2015).
- [35] A. D. King and C. C. McGeoch, arXiv:1410.2628 (2014).
- [36] A. D. King, arXiv:1502.02098 (2015).
- [37] M. Amin, arXiv:1503.04216 (2015).
- [38] H. G. Katzgraber, F. Hamze, and R. S. Andrist, Phys. Rev. X **4**, 021008 (2014).
- [39] H. G. Katzgraber, F. Hamze, Z. Zhu, A. J. Ochoa, and H. Munoz-Bauza, arXiv:1505.01545 (2015).
- [40] V. Martin-Mayor and I. Hen, arXiv:1502.02494 (2015).
- [41] Z. Zhu, A. J. Ochoa, S. Schnabel, F. Hamze, and H. G. Katzgraber, arXiv:1505.02278 (2015).
- [42] D. Gottesman, in *Quantum Information Science and Its Contributions to Mathematics*, Proceedings of Symposia in Applied Mathematics, Vol. 68, edited by S. Lomonaco (Amer. Math. Soc., Rhode Island, 2010) p. 13, eprint arXiv:0904.2557.
- [43] F. Gaitan, *Quantum Error Correction and Fault Tolerant Quantum Computing* (Taylor & Francis Group, Boca Raton, 2008).
- [44] D. Lidar and T. Brun, eds., *Quantum Error Correction* (Cambridge University Press, Cambridge, UK, 2013).
- [45] E. Knill, R. Laflamme, and W. H. Zurek, Science **279**, 342 (1998).
- [46] D. Aharonov, A. Kitaev, and J. Preskill, Phys. Rev. Lett. **96**, 050504 (2006).
- [47] A. M. Childs, E. Farhi, and J. Preskill, Phys. Rev. A **65**, 012322 (2001).
- [48] M. S. Sarandy and D. A. Lidar, Phys. Rev. Lett. **95**, 250503 (2005).
- [49] M. H. S. Amin, P. J. Love, and C. J. S. Truncik, Phys. Rev. Lett. **100**, 060503 (2008).
- [50] S. P. Jordan, E. Farhi, and P. W. Shor, Phys. Rev. A **74**, 052322 (2006).
- [51] D. A. Lidar, Phys. Rev. Lett. **100**, 160506 (2008).
- [52] G. Quiroz and D. A. Lidar, Phys. Rev. A **86**, 042333 (2012).
- [53] A. Ganti, U. Onunkwo, and K. Young, Phys. Rev. A **89**, 042313 (2014).
- [54] A. D. Bookatz, E. Farhi, and L. Zhou, arXiv:1407.1485 (2014).
- [55] A. Mizel, arXiv:1403.7694 (2014).
- [56] K. C. Young, M. Sarovar, and R. Blume-Kohout, Phys. Rev. X **3**, 041013 (2013).
- [57] M. Sarovar and K. C. Young, New J. of Phys. **15**, 125032 (2013).
- [58] I. Marvian and D. A. Lidar, Phys. Rev. Lett. **113**, 260504 (2014).
- [59] K. L. Pudenz, T. Albash, and D. A. Lidar, Nat. Commun. **5**, 3243 (2014).
- [60] K. L. Pudenz, T. Albash, and D. A. Lidar, Phys. Rev. A **91**, 042302 (2015).
- [61] K. C. Young, R. Blume-Kohout, and D. A. Lidar, Phys. Rev. A **88**, 062314 (2013).
- [62] V. Choi, Quant. Inf. Proc. **7**, 193 (2008).
- [63] V. Choi, Quant. Inf. Proc. **10**, 343 (2011).
- [64] C. Klymko, B. D. Sullivan, and T. S. Humble, Quantum Information Processing **13**, 709 (2014).
- [65] J. Cai, W. G. Macready, and A. Roy, arXiv:1406.2741 (2014).
- [66] F. Barahona, J. Phys. A: Math. Gen **15**, 3241 (1982).
- [67] W. Ebeling, A. Engel, B. Esser, and R. Feistel, Journal of Statistical Physics **37**, 369 (1984).
- [68] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, Science **220**, 671 (1983).
- [69] P. Ruján, Zeitschrift für Physik B Condensed Matter **73**, 391 (1988).
- [70] R. L. Somorjai, The Journal of Physical Chemistry **95**, 4141 (1991).
- [71] K. A. Olszewski, L. Piela, and H. A. Scheraga, The Journal of Physical Chemistry **96**, 4672 (1992).
- [72] P. Ray, B. K. Chakrabarti, and A. Chakrabarti, Phys. Rev. B **39**, 11828 (1989).
- [73] A. Das and B. K. Chakrabarti, Rev. Mod. Phys. **80**, 1061 (2008).
- [74] B. Heim, T. F. Rønnow, S. V. Isakov, and M. Troyer, Science **348**, 215 (2015).
- [75] S. Muthukrishnan, T. Albash, and D. A. Lidar, arXiv:1505.01249 (2015).
- [76] R.J. Wilson, *Introduction to Graph Theory*, 4th ed. (Addison-Wesley, Reading, Massachusetts, 1997).
- [77] D. Venturelli, S. Mandrà, S. Knysh, B. O’Gorman, R. Biswas, and V. Smelyanskiy, arXiv:1406.7553 (2014).
- [78] D. Gottesman, Phys. Rev. A **54**, 1862 (1996).
- [79] S. R. Broadbent and J. M. Hammersley, Proceedings of the Cambridge Philosophical Society **53**, 629 (1957).
- [80] M. Menshikov, Soviet Mathematics Doklady **33**, 856 (1986).
- [81] W. Barthel, A. K. Hartmann, M. Leone, F. Ricci-Tersenghi, M. Weigt, and R. Zecchina, Phys. Rev. Lett. **88**, 188701 (2002).
- [82] F. Krzakala and L. Zdeborová, Phys. Rev. Lett. **102**, 238701 (2009).
- [83] S. Bravyi and B. Terhal, *SIAM Journal on Computing*, SIAM Journal on Computing **39**, 1462 (2009).
- [84] M. Levinshtein, B. Shklovskii, M. Shur, and A. Efros, Zh. Eksp. Theor. Fiz **69**, 386 (1975).
- [85] P. Grassberger, Journal of Physics A Mathematical General **25**, 5867 (1992).
- [86] R. Martoňák, G. E. Santoro, and E. Tosatti, Phys. Rev. B **66**, 094203 (2002).
- [87] G. E. Santoro, R. Martoňák, E. Tosatti, and R. Car, Science **295**, 2427 (2002).
- [88] T. Albash, I. Hen, F. M. Spedalieri, and D. A. Lidar, ArXiv e-prints (2015), arXiv:1506.03539 [quant-ph].
- [89] R. G. Gallager, *IRE Transactions on Information Theory*, IRE Transactions on Information Theory **8**, 21 (January 1962).
- [90] D. MacKay, *Information Theory, Inference & Learning Algorithms* (Cambridge University Press, New York, NY, 2003).
- [91] D. Bacon, Physical Review A **78**, 042324 (2008).
- [92] U. Wolff, Phys. Rev. Lett. **62**, 361 (1989).

## Appendix A: Energy Minimization Decoding

We performed decoding through energy minimization by implementing simulated annealing on the “decoding” Hamiltonian defined in Eq. (20). We used a linear schedule for the temperature, with the initial and final temperatures equal to  $T_{\text{in}} = 4 \max_{i,j}(J_{ij})$  and  $T_{\text{fin}} = 0.1 \min_{i,j}(J_{ij})$  respectively. We chose the computational effort in the decoding step to be constant for all the decoding attempts. For each readout after an annealing cycle, decoding consisted of repeating 10 simulated annealing runs with 10 sweeps. For each decoding attempt, we defined a decoding success probability  $P_{\text{dec}}$  as the fraction of times the lowest decoded configuration was found. We consider this quantity as a measure of our decoding quality, with a  $P_{\text{dec}}$  close to 1 meaning that the best decoded configuration is likely to be the optimally decoded state. A small value of  $P_{\text{dec}}$ , on the other hand, means that a better decoded configuration could likely be found with increased computational effort (a larger number of sweeps and/or repetitions in the simulated annealing procedure).

$P_{\text{dec}}$  is shown in Fig. 20 for a few representative examples. Recall that the number of broken qubits is typically larger for QAC-ME (Fig. 18), which therefore requires a larger decoding effort in order to obtain a similar decoding quality. This explains why we observe lower decoding success probabilities for the QAC-ME case with energy minimization. Nonuniform penalties are also responsible for a larger number of broken qubits. Improving  $P_{\text{dec}}$  for QAC-ME by expending more effort on the simulated annealing step would thus further increase the advantage of QAC-ME with respect to ME, as well as the advantages of nonuniform penalties with respect to the uniform case.

## Appendix B: Concatenating the square code

In Sec. V C of the main text we mentioned that the square code can be concatenated, and described the first and second encoding levels. To see how this can be continued, consider, e.g., combining  $n \times n$  encoded level-2 qubits to form an encoded level-3 qubit on each plane of the 2LG, to obtain a  $[[4n^4, 1, 4n^4]]_0$  code. In general we have  $[n, r] \equiv [[4n^{2(r-1)}, 1, 4n^{2(r-1)}]]_0$  for a level- $r$  code with integer  $r \geq 1$ . The realizable stabilizer terms included in the penalty Hamiltonian suffice to penalize every possible bit-flip error, at every level of concatenation.

It is interesting to consider the growth in the encoded fields and couplings. As illustrated in Fig. 12, for levels  $r \geq 2$  the encoded couplings grow in proportion to the side length of the square of encoded qubits, and the encoded fields must grow in the same proportion in order to maintain their relative strengths. E.g., the level-2 local fields correspond to  $n$  level-1 realizations, each with two level-0 realizations. Analogously, every encoded level-2 coupling comprises  $n$  level-1 couplings, which are in turn equivalent to 2 level-0 couplings. Thus  $\bar{h}_\ell^{(2)} = n\bar{h}_\ell^{(1)} = 2nh_\ell$  and  $\bar{J}_{\ell\ell'}^{(2)} = n\bar{J}_{\ell\ell'}^{(1)} = 2nJ_{\ell\ell'}$ . At level  $r$  each square has side length  $n^{r-1}$ , so the encoded

level- $r$  fields and couplings are

$$\bar{h}_\ell^{(r)} = 2n^{r-1}h_\ell, \quad \bar{J}_{\ell\ell'}^{(r)} = 2n^{r-1}J_{\ell\ell'}, \quad (\text{B1})$$

i.e., the energy scale of the level- $r$  final Hamiltonian has been boosted by a factor of  $2n^{r-1}$ . Note that the encoded penalty strength grows in proportion to the area of each square if all the penalties are utilized. Since this would result in overly rigid qubits a simple remedy is to use only a linear number of the total available penalties.

Since the number of physical qubits required grows quadratically faster than the strength of the encoded local fields and couplings, it is not a priori clear that a higher concatenation level  $r$  is beneficial. However, in a Markovian model where the probability of a thermal excitation at inverse temperature  $\beta$  is proportional to the Gibbs factor  $e^{-\beta\Delta}$  (e.g., Ref. [25]), with  $\Delta$  the gap from the ground state, it is possible to realize a gain by increasing  $r$ , provided the gap grows along with the energy scale of the level- $r$  final Hamiltonian. Determining the scaling of the gap with  $r$  is a non-trivial problem since the energy scale of the initial transverse field Hamiltonian does not also grow under the present encoding, as remarked earlier. Our proposal for concatenating the square code for QAC can be viewed as a physically realizable (and hence severely compromised) version of Bacon’s proposal for quantum concatenated-code Hamiltonians [91].

## Appendix C: Recursive decoding of the square code

In this section we propose an alternative (untested) recursive decoding scheme of the square code, that is similar to the energy minimization strategy.

Suppose we already decoded all the untied encoded qubits with only the tied encoded qubits remaining to be decoded. Consider a problem Hamiltonian  $H^{(0)} = \sum_{(i,j) \in G} J_{ij} \sigma_i^z \sigma_j^z$ , where the sum is over all connected pairs of logical qubits  $i$  and  $j$  in the encoded graph  $\bar{G}$  (the graph obtained after, say, the first level of encoding using the square code), and for notational simplicity here we use undecorated Pauli operators to denote the encoded operators. At the end of a (quantum) annealing run some encoded qubits will be tied. The untied qubits now have fixed values  $\{s_i\}$ , so they act as local fields on the tied qubits, weighted by the couplings. We can rewrite the problem Hamiltonian after this run as

$$H^{(1)} = \sum_{j' \in \text{ties}} h_{j'} \sigma_{j'}^z + \sum_{(i',j') \in \bar{G}'} J_{i'j'} \sigma_{i'}^z \sigma_{j'}^z, \quad (\text{C1})$$

where the sum runs only over the tied qubits, which occupy the vertices of a subgraph  $\bar{G}'$  of  $\bar{G}$ , and where the local fields are

$$h_{j'} = \sum_i J_{ij'} s_i, \quad (\text{C2})$$

with the sum including all the untied qubits coupled to the given tied qubit  $j'$ .

The renormalized Hamiltonian  $H^{(1)}$  represents a new Ising problem. To find the values of the tied qubits we can run  $H^{(1)}$  on the (quantum) annealer, i.e., implement the EM strategy by bootstrapping the problem. If after such a run there are still tied qubits we can repeat the same process, defining a new Hamiltonian  $H^{(2)} = \sum_{j'' \in \text{ties}} h_{j''} \sigma_{j''}^z + \sum_{(i'', j'') \in \bar{G}''} J_{i'' j''} \sigma_{i''}^z \sigma_{j''}^z$ , with  $h_{j''} = \sum_{i'} J_{i' j''} s_{i'}$ , etc. Note that successive runs of the quantum annealer need not necessarily resolve the tied qubits, but since  $\bar{G}''$  is by construction a subgraph of  $\bar{G}'$ , this process will converge. Any ties that remain after the process has converged can be decoded using the previously described EM strategy. Accounting for wall-clock times, it may be preferable to start running this sequence of recursive optimizations on a classical solver once the subgraph size has become small enough.

#### Appendix D: Broken Qubit Probability, Range and Frustration

In this section we present the complete data set in Figs. 21 (planted), 22 (embeddable planted), 23 (weighted planted), and 24 (deformed embeddable) for the broken qubit (and ties) probabilities for the instances studied for all values of the penalties, for both the ME and QAC-ME cases. Note how, as expected, the broken qubit probability ( $p_{\text{BQ}}$ ) strongly depends on the strength of the energy penalties. One more feature that is evident from the same figures is the different behavior of  $p_{\text{BQ}}$  in the uniform and nonuniform cases, at a given value of the penalty strength  $\gamma$ . In the uniform case,  $p_{\text{BQ}}$  is mostly a monotonically decreasing function of the clause density  $\alpha$ , while in the nonuniform case there is a clear peak that is located close to the critical clause density.

There is a simple and intuitive explanation of this behavior in terms of range and frustration, that is consistent with the general discussion of Sec. IID. Let us first give a quantitative definition of both frustration and range. We define frustration as the percentage of  $J_{ij}$  couplings that are not satisfied on the planted ground state solution. The inverse range is the smallest (non-zero) absolute value of the  $J_{ij}$  couplings of a given instance. Each time we implement a given instance, we always set the value of the largest coupling equal to 1. This exploits the largest energy scale achievable with the DW2 processor. The range thus gives us an estimate of the required precision needed to correctly implement the given instance.

Figure 25 shows the computed values of the required (inverse) range and the frustration of the various classes of instances. We can see that the qualitative behavior of  $p_{\text{BQ}}$  in the uniform case follows that of the inverse range. This is related to the fact that a smaller inverse range implies the presence of a larger number of weak couplings that are dominated by the uniform and relatively larger energy penalty, making the appearance of a broken qubit less likely. On the other hand, the value of the energy penalties in the nonuniform case is adjusted to the typical value of the logical couplings. Typically, a smaller inverse range results in weaker energy penalties, balancing the range effect. Thus, in the nonuniform case  $p_{\text{BQ}}$  is more positively correlated with frustration. This is consistent with the general expectation that  $p_{\text{BQ}}$  is large for highly frus-

trated instances.

#### Appendix E: Data Collection and Analysis

As already mentioned in Section VIB, the presence of control errors (ICE) due to low frequency noise and systematic biases prevents the physical couplings to be set with a precision better than  $\sim 5\%$ . To average out the effects of such errors, we run 10 programming cycles for all the encoding schemes (D, ME, QAC-ME) and penalty values for every instance considered. Each programming cycle implements a gauge transformation and consists of  $10^3$  annealing runs. For each instance, encoding strategy and penalty value, (altogether labeled here by the index  $i$ ) we extract an ‘‘intrinsic’’ success probability  $P^i$  as the overall probability to find the ground state over the 10 programming cycles. In other words, this is the mean of the success probabilities extracted from each programming cycle, and it serves as our best guess for the intrinsic (without ICE) success probability of each encoding of a given instance

$$P^i = \sum_g P_g^i, \quad g = 1, \dots, 10. \quad (\text{E1})$$

We then use these gauge-averaged probabilities to compute the renormalized success probabilities for the D and ME cases:

$$P_{\text{D}}^i \mapsto 1 - (1 - P_{\text{D}}^i)^4, \quad (\text{E2a})$$

$$P_{\text{ME}}^i \mapsto 1 - (1 - P_{\text{ME}}^i)^2, \quad (\text{E2b})$$

$$P_{\text{QAC-ME}}^i \mapsto P_{\text{QAC-ME}}^i. \quad (\text{E2c})$$

As explained in the main text, this renormalization is necessary to correctly take into account the amount of hardware resources that each implementation requires. We then use these renormalized probabilities to compute 5000 means, obtained from 5000 bootstrapped samples of 100 instances (picked with replacement among the 100 total number of instances) for each clause density  $\alpha$  and penalty value  $\gamma$ . The quantities plotted in Fig. 13 (as well as in all other figures plotting experimental data) are the mean and the standard error of such 5000 samples of bootstrapped means.

The same approach was followed to compute means and error bars of other similar plots. See, e.g., the plot showing broken qubits probabilities (Fig. 18).

#### Appendix F: Energy Penalties

We show explicitly in Fig. 26 the value of the optimal  $\gamma$  for all the cases considered in the main text. As can be seen, this value is fairly constant as a function of the clause density  $\alpha$ . There is a slight tendency for the optimal  $\gamma$  to decrease with  $\alpha$  in the case of uniform penalties. This is most pronounced in the ME implementation of the deformed embeddable instances. When choosing nonuniform penalties, however, the optimal  $\gamma$  again becomes almost constant. This shows that the nonuniform choice of penalties has the potential of being

closer to optimality than the uniform case with a single choice of the optimal penalty strength. Since optimization of the energy penalties can be time consuming, this property is important for any practical encoding scheme. Ideally, we would like an encoding scheme (for the choice of penalties) that is close to optimality and does not depend on any quantity that has to be optimized on an instance-by-instance case.

### Appendix G: Decoding Strategies

We show in Figs. 27 and 28 the comparison between various decoding strategies for the embeddable planted and the deformed embeddable cases. This complements the results shown in Figs. 15 and 16 of the main text, which showed the results for the planted and weighted planted instances. Results for the embeddable sets are very similar to the sets defined on the full 2LG. We notice here too that the use of the EM strategy is crucial for the QAC-ME encoding to give comparable (for the planted embeddable set) or better (for the deformed embeddable set) results than the direct implementation (D).

### Appendix H: Worst-case Scenario Decodability of the Square Code

We can set an upper limit (worst-case scenario) for the probability to obtain ties and broken qubits in general by assuming that all four physical qubits comprising an encoded qubit are uncorrelated. This corresponds to assuming that the output of the annealer corresponds to a thermal distribution with infinite temperature. Within a single encoded qubit there are 6 physical qubit states that correspond to a tie (out of a total of 16 states), so we obtain  $p_{wc}^{tie} = 6/16 = 0.375$  for the worst case probability to obtain a tie. The worst case probability to obtain a broken qubit is  $p_{wc}^{BQ} = 14/16 = 0.875$ . While  $p_{wc}^{BQ}$  is definitely above the percolation threshold of the square code,  $p_{wc}^{tie}$  is almost certainly below it since the device output is better approximated by a finite temperature thermal distribution (as opposed to an infinite temperature distribution). Moreover, the non-vanishing penalty term and (if present) longitudinal field disfavor broken qubits and ties. This strongly suggests that in the more realistic scenario the number of broken qubits is well below the worst-case value. Fig. 29 shows the mean size of the connected domains over  $10^4$  randomly generated configurations for increasingly large  $N \times N$  2LG graphs. The fact that this mean size saturates instead of growing with  $N$  is evidence that  $p_{wc}^{tie} < p_{2LG}$ . This means that the square code is efficiently decodable (at least using using the MV-EM strategy) even in the worst-case scenario.

### Appendix I: Simulated Quantum Annealing

We reported SQA results in Sec. VIII of the main text. Here we briefly review this technique. SQA is a quantum Monte Carlo based algorithm whereby Monte Carlo dynamics are used to sample from the instantaneous Gibbs state associated

with the Hamiltonian in Eq. (1). The state at the end of the quantum Monte Carlo simulation of the quantum Hamiltonian  $H(t)$  is introduced as the initial state of the Monte Carlo simulation with Hamiltonian  $H(t + \Delta t)$ . This proceeds until  $H(t_f)$  is reached. Originally proposed as an optimization algorithm [86, 87], it has gained some traction as a classically efficient model for finite temperature quantum annealers [17, 21, 33, 34]. SQA does not capture the unitary dynamics of the quantum system, but it is hoped that the sampling of the instantaneous Gibbs state captures thermal processes in the quantum annealer, which may be the dominant dynamics if the evolution is sufficiently slow. Although there is strong evidence that SQA does not completely capture the final-time statistics of the D-Wave processors [21, 28], at present it is the only viable means to simulate large open quantum annealing systems.

We used discrete-time quantum Monte Carlo in our simulations, which we briefly review. For Hamiltonians of the form of the transverse Ising model, such as in Eqs. (1) and (2), the sampling from the instantaneous Gibbs state is done by sampling from the dual classical system with Hamiltonian  $\mathcal{H}_C$  and inverse temperature  $\beta$ :

$$\beta\mathcal{H}_C(t) = J_{\perp}(t) \sum_{i,\tau} s_{i,\tau} s_{i,\tau+1} \quad (I1)$$

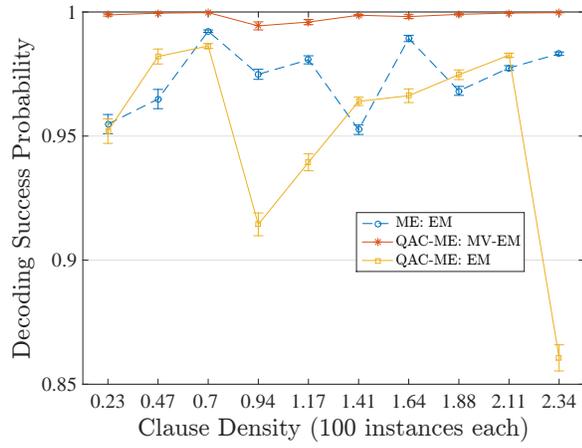
$$+ \frac{\beta}{N_{\tau}} B(t) \sum_{\tau} \left[ \sum_i h_i s_{i,\tau} + \sum_{i<j} J_{ij} s_{i,\tau} s_{j,\tau} \right],$$

where  $N_{\tau}$  is the number of Trotter slices used along the Trotter direction,  $s_{i,\tau}$  denotes the  $i$ th classical spin on the  $\tau$ th Trotter slice, and

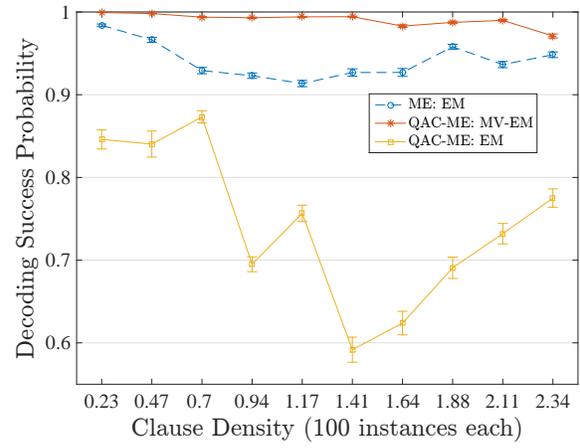
$$J_{\perp}(t) \equiv \frac{1}{2} \ln [\tanh(\beta A(t)/N_{\tau})] < 0 \quad (I2)$$

is the nearest-neighbor (ferromagnetic) coupling strength along the Trotter direction.

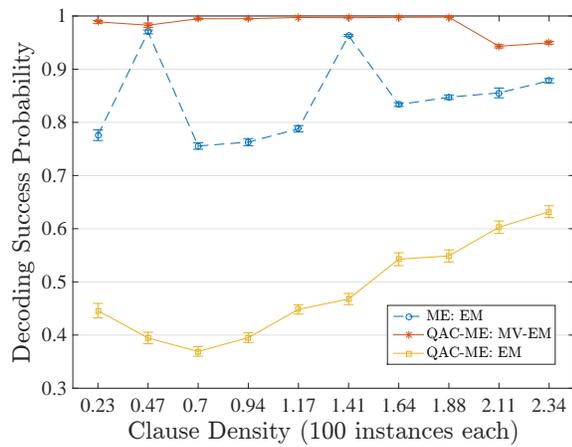
In the simulations presented in the main text, the number of Trotter slices was fixed to 64. Spin updates were performed via Wolff-cluster updates [92] along the Trotter direction only.



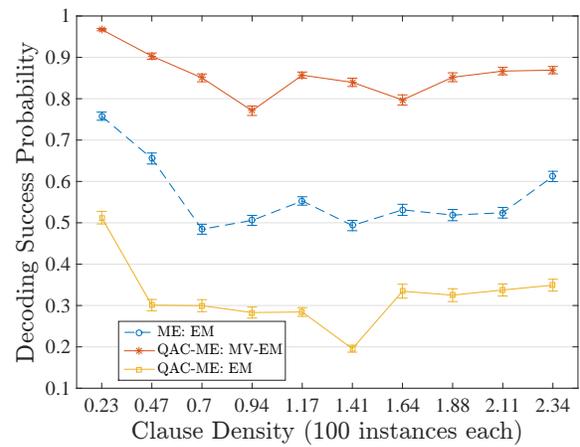
(a) Planted, uniform



(b) Planted, nonuniform

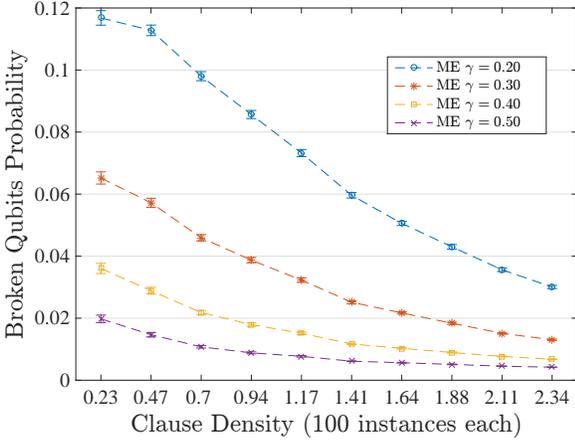


(c) Weighted planted, uniform

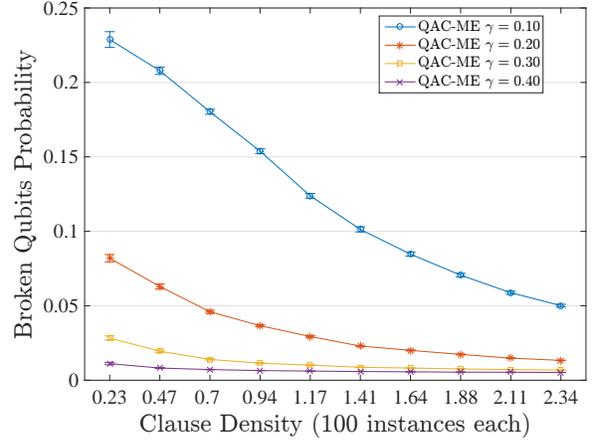


(d) Weighted planted, nonuniform

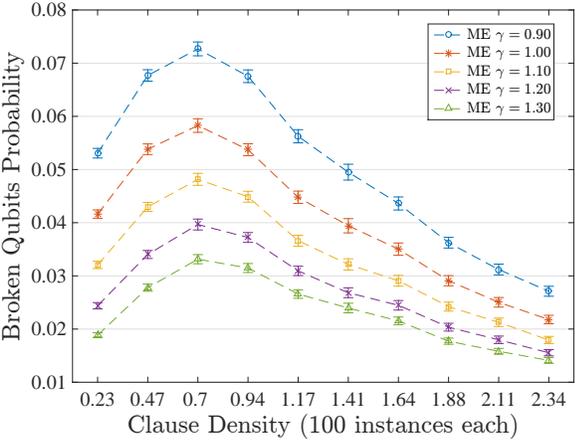
FIG. 20. (Color online) Energy minimization success probabilities of decoding using simulated annealing.



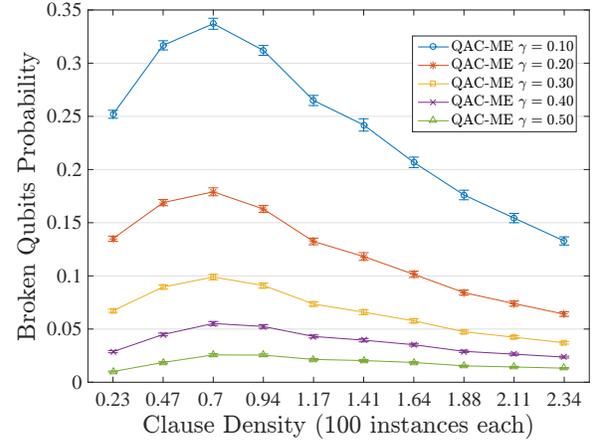
(a) Planted, uniform



(b) Planted, uniform

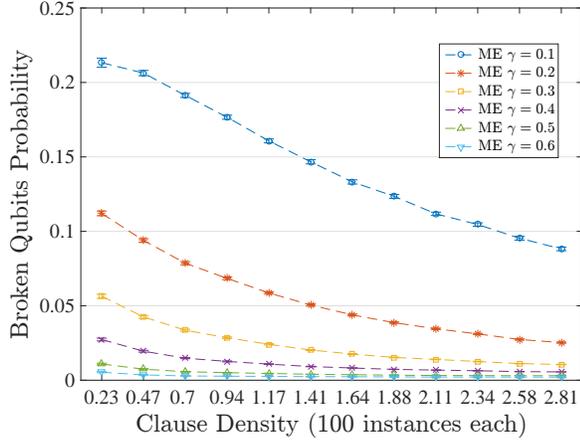


(c) Planted, nonuniform

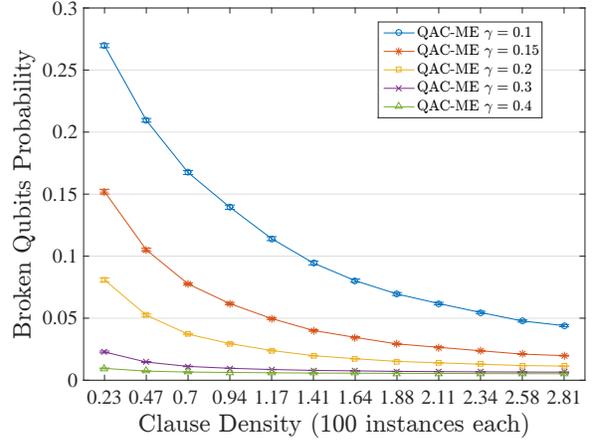


(d) Planted, nonuniform

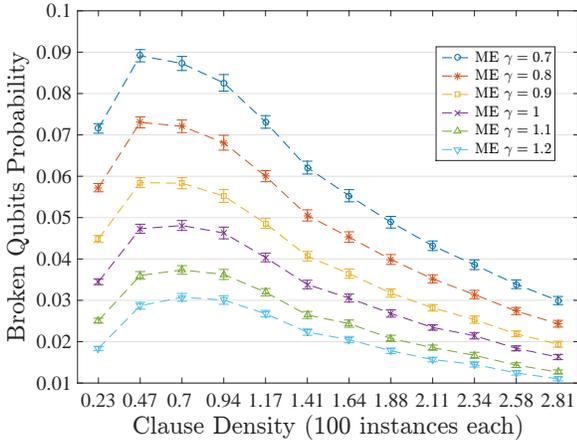
FIG. 21. (Color online) Probability of broken qubits for ME (left) and QAC-ME (right) at different penalty values, for the planted and embeddable planted instances. Top row: uniform. Bottom row: nonuniform. The dependence on the clause density is monotonic in the uniform case and exhibits a maximum in the nonuniform case.



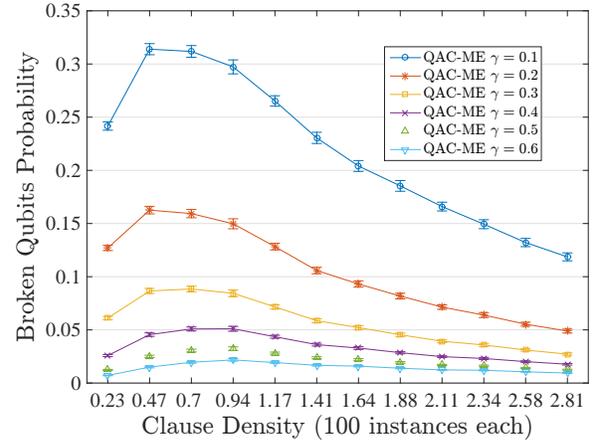
(a) Embeddable planted, uniform



(b) Embeddable planted, uniform

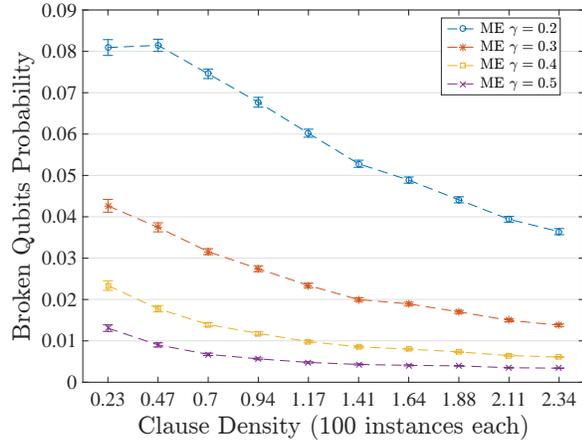


(c) Embeddable planted, nonuniform

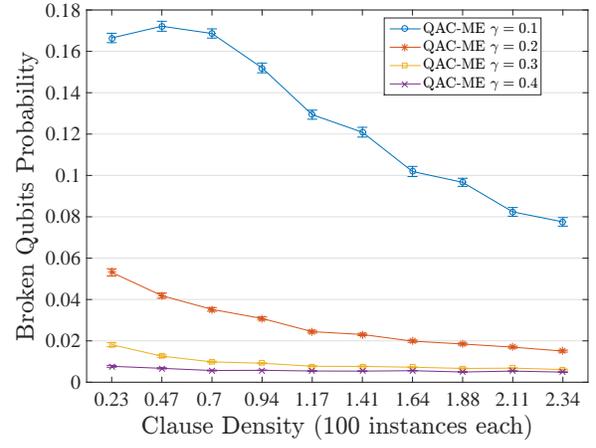


(d) Embeddable planted, nonuniform

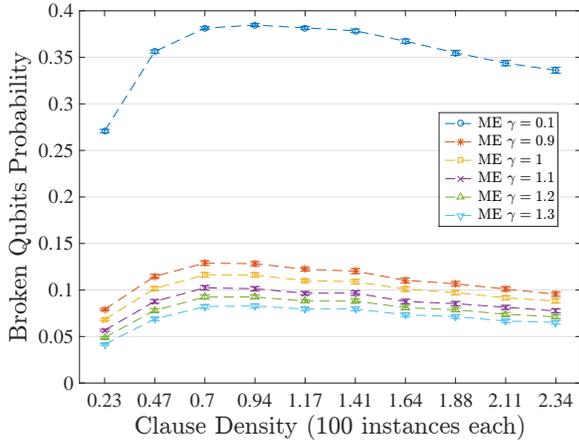
FIG. 22. (Color online) Probability of broken qubits for ME (left) and QAC-ME (right) at different penalty values, for the embeddable planted instances. Top row: uniform. Bottom row: nonuniform. The dependence on the clause density is monotonic in the uniform case and exhibits a maximum in the nonuniform case.



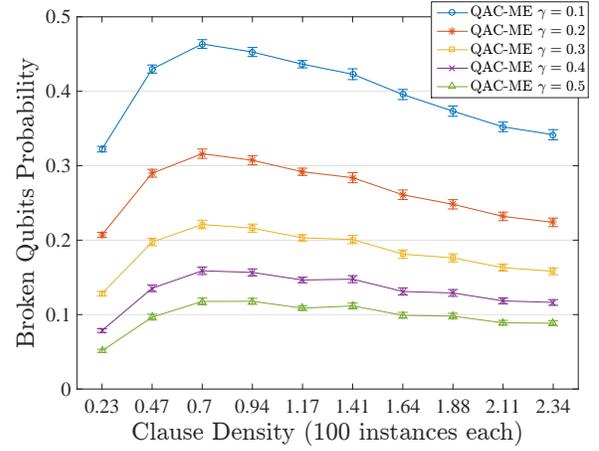
(a) Weighted planted, uniform



(b) Weighted planted, uniform



(c) Weighted planted, nonuniform



(d) Weighted planted, nonuniform

FIG. 23. (Color online) Probability of broken qubits for ME (left) and QAC-ME (right) at different penalty values, for the weighted planted and deformed embeddable instances. Top row: uniform. Bottom row: nonuniform. The dependence on the clause density is monotonic in the uniform case and exhibits a maximum in the nonuniform case.

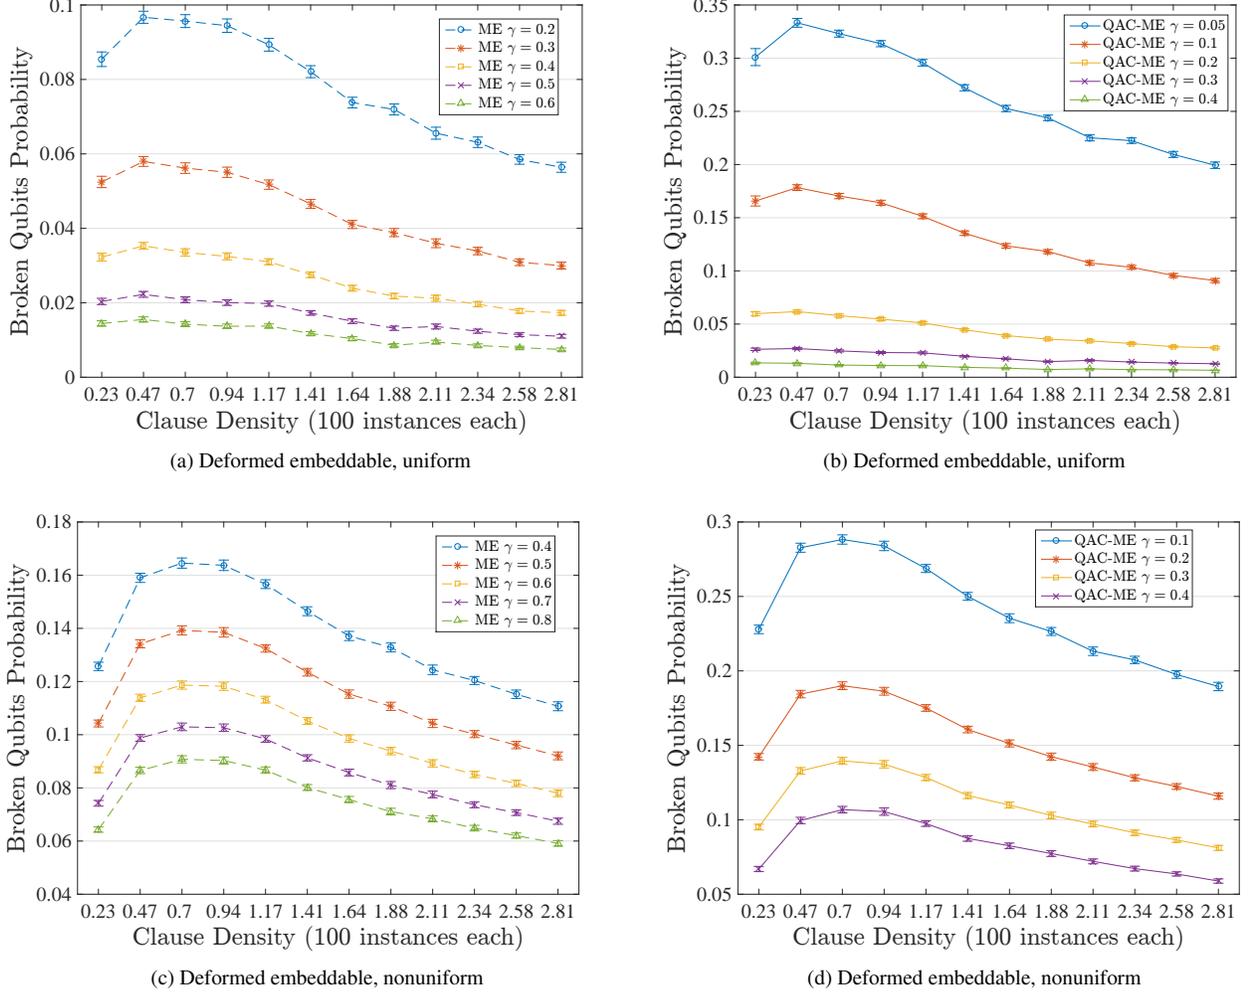


FIG. 24. (Color online) Probability of broken qubits for ME (left) and QAC-ME (right) at different penalty values, for the deformed embeddable instances. Top row: uniform. Bottom row: nonuniform. The dependence on the clause density is monotonic in the uniform case and exhibits a maximum in the nonuniform case.

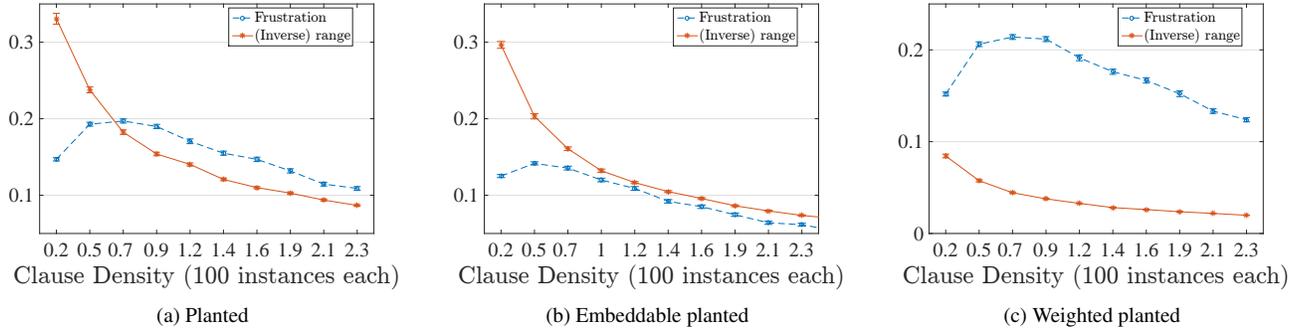


FIG. 25. (Color online) Range and frustration as a function of the clause density. Note the correlation of frustration (range) to the number of broken qubits obtained with the nonuniform (uniform) encoding choices of Figs. 21-24.

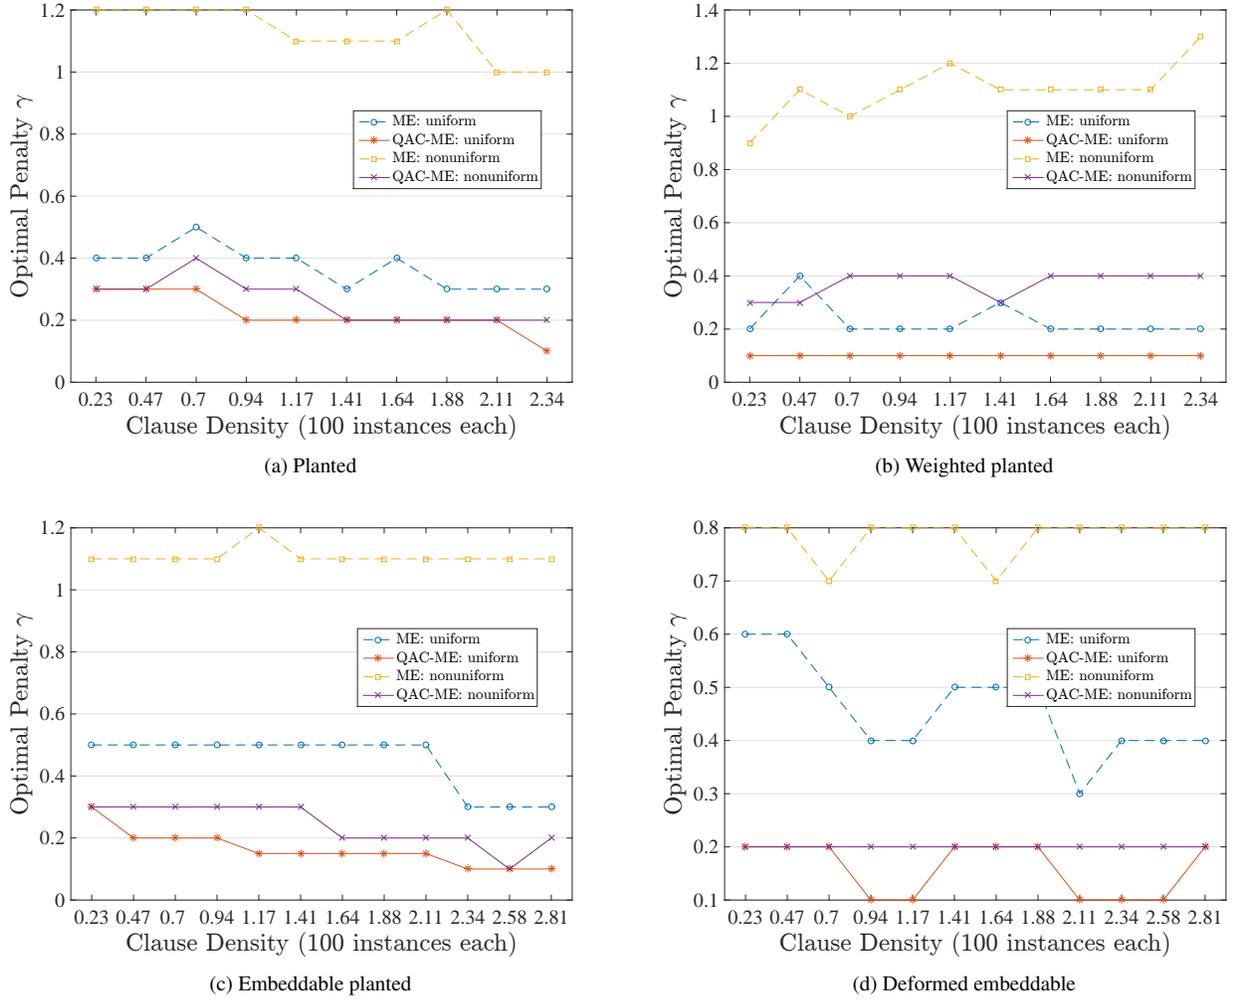


FIG. 26. (Color online) The four panels show the optimal penalty strength  $\gamma$  for each set of instances and encoding choice (uniform/nonuniform). The optimal  $\gamma$  is fairly constant over the different values of the clause density  $\alpha$ .

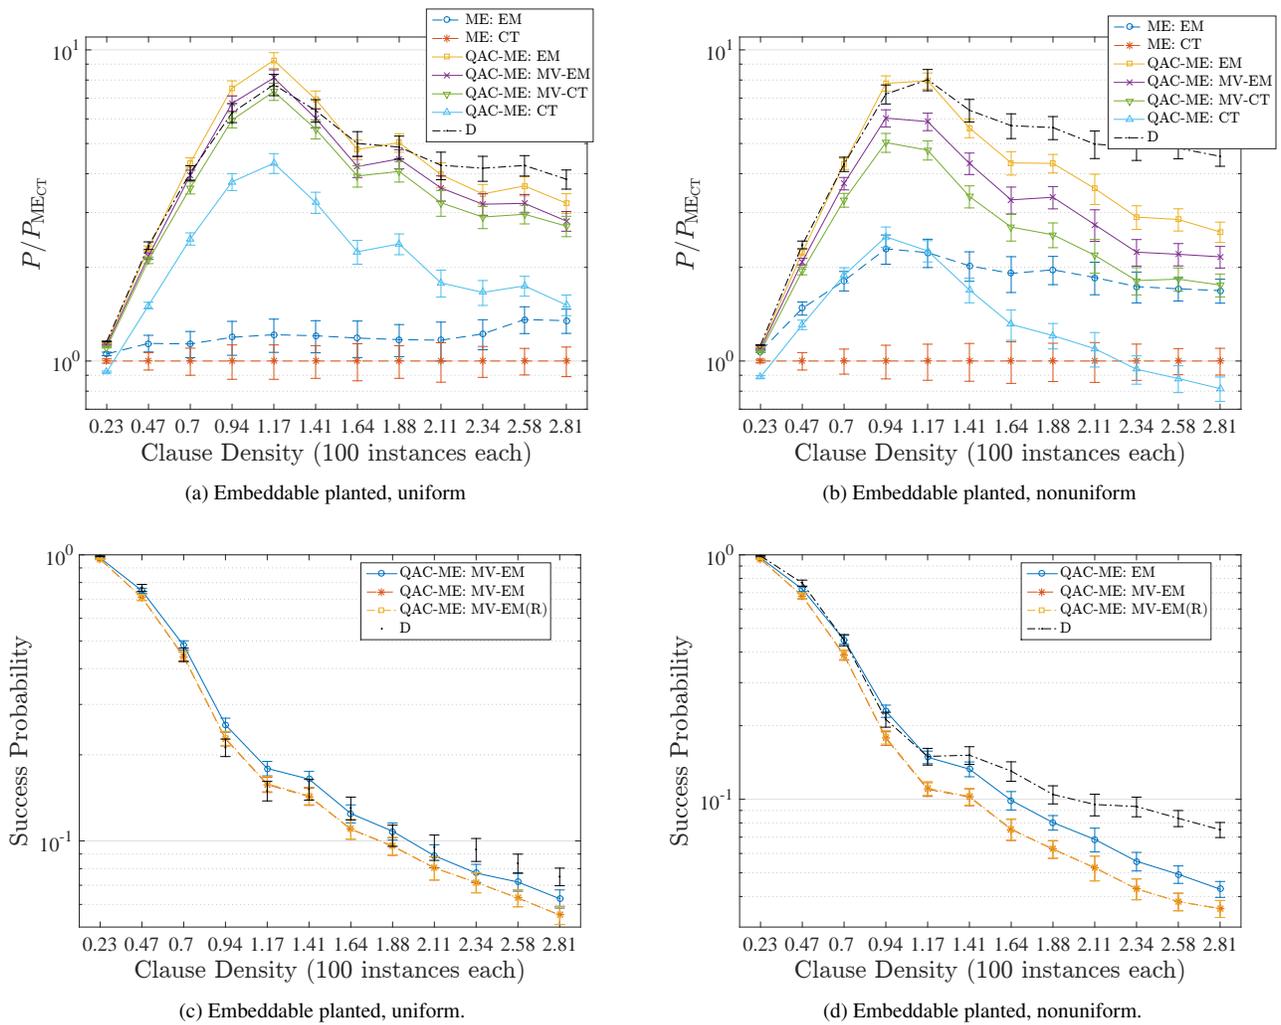


FIG. 27. (Color online) Comparison between the decoding strategies for the embeddable planted instances.

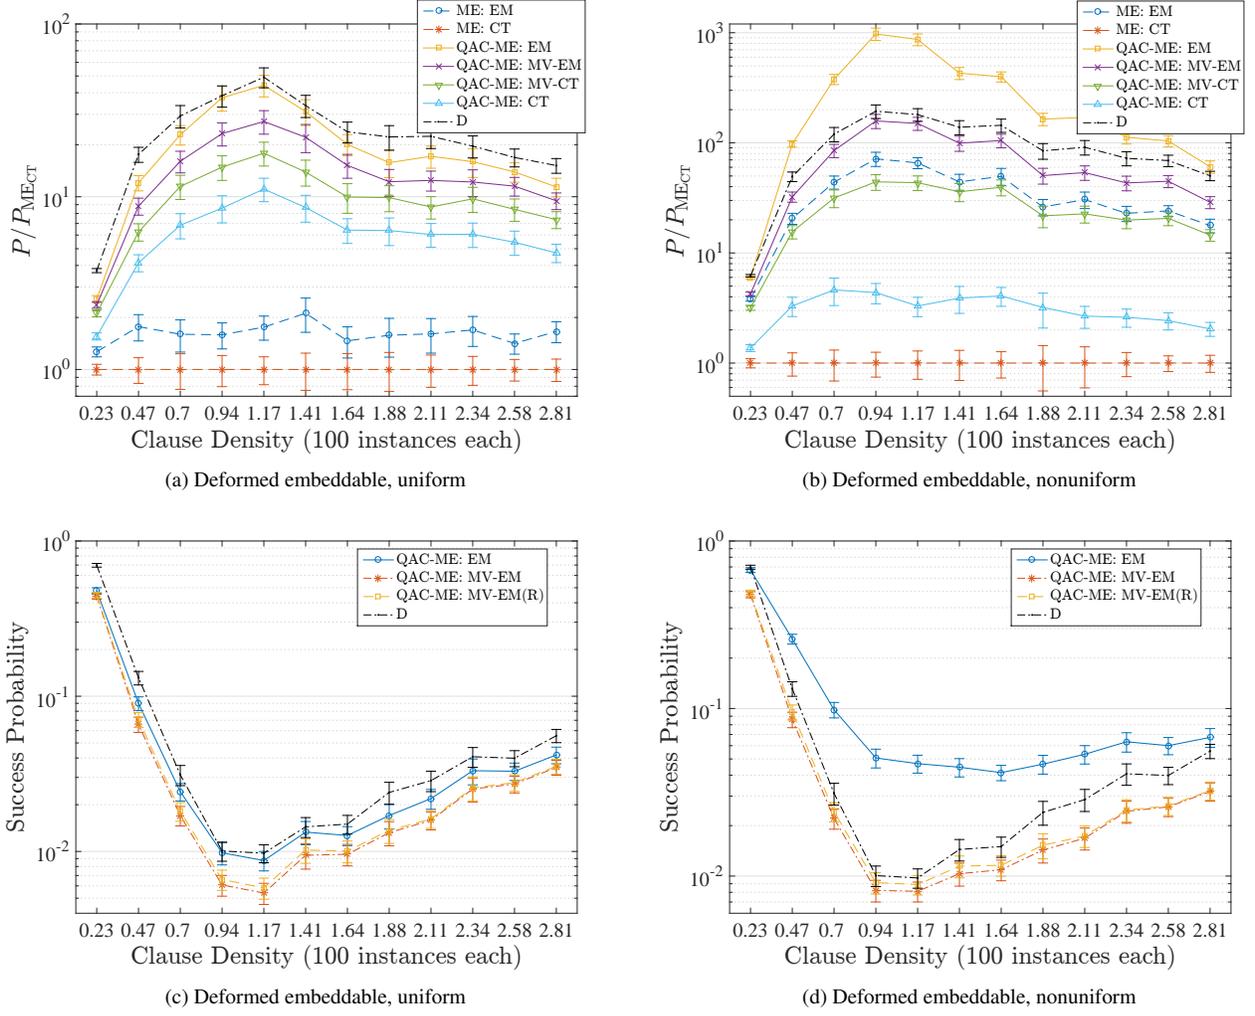


FIG. 28. (Color online) Comparison between the decoding strategies for the deformed embeddable instances.

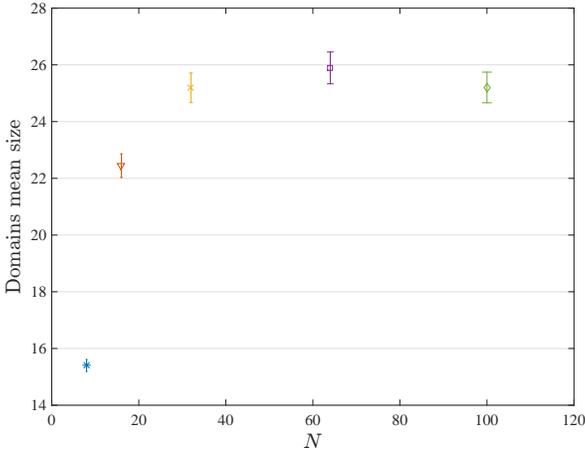


FIG. 29. (Color online) Mean size of domains connected to a central spin for various  $N \times N \times 2$  2LG graph sizes for the worst-case percolation probability  $p_{wc}^{tie} = 0.375$ . The mean size of connected domains saturates at a size of about 25 spins. We used  $10^4$  random configurations per point.