

Quantum Annealers Chain Strengths: A Simple Heuristic to Set Them All

Valentin Gilbert^[0009-0001-1004-0204] and Stéphane Louise^[0000-0003-4604-6453]

Université Paris-Saclay, CEA-List, F-91120, Palaiseau, France
{valentin.gilbert, stephane.louise}@cea.fr

Abstract. Quantum annealers (QA), such as D-Wave systems, become increasingly efficient and competitive at approximating combinatorial optimization problems. However, solving problems that do not directly map the chip topology remains challenging for this type of quantum computer. The creation of logical qubits as sets of interconnected physical qubits overcomes limitations imposed by the sparsity of the chip at the expense of increasing the problem size and adding new parameters to optimize. This paper explores the advantages and drawbacks provided by the structure of the logical qubits and the impact of the rescaling of coupler strengths on the minimum spectral gap of Ising models. We show that logical qubits encoded over densely connected physical qubits require a lower chain strength to maintain the ferromagnetic coupling. We also analyze the optimal chain strength variations considering different minor embeddings of the same instance. This experimental study suggests that the chain strength can be optimized for each instance. We design a heuristic that optimizes the chain strength using a very low number of shots during the preprocessing step. This heuristic outperforms the default method used to initialize the chain strength on D-Wave systems, increasing the quality of the best solution by up to 17.2% for tested instances on the max-cut problem.

Keywords: quantum annealing · Ising chain strength · minor embedding

1 Introduction

The idea of using Quantum Annealers to solve combinatorial problems is not new and was exposed by Kadowaki et al. [13]. Despite strong theoretical proofs based on the quantum adiabatic theorem, the speed-up brought up by quantum annealers still needs to be quantified for real useful applications. It comes from the fact that quantum annealers implement a noisy version of the more general Adiabatic Quantum Computation (AQC) (the reader may refer to Albash et al. [2] for an introduction to AQC). Indeed, D-Wave systems suffer from more than five different sources of Integrated Control Errors (ICE) [1] and have limited connectivity between qubits. The maximal number of qubits available on the quantum chip limits the size of the problem that can be solved. Minor embedding

algorithms transform the initial problem into a new one that fits the sparsely connected quantum chip. The transformation consists of mapping the initial problem’s logical variables to a set of physical variables that can be straightly encoded on the physical qubits of the quantum annealer. As mentioned in Gilbert et al. [10], assessing the quality of an embedding is not trivial. The number of physical qubits used in the embedding can serve as a first quality indicator, but the physical qubit’s topology used to encode the logical qubits can also be considered. To the authors’ knowledge, a single contribution on this issue was made by Pelofske [15]. Topologies used for physical qubits are usually chains of qubits because this structure maximizes the number of potential connections of the logical qubit. E. Pelofske shows that the performance of the quantum solver is increased when logical qubits are encoded on cliques instead of chains. The minor embedding of a logical variable into a set of physical qubits requires the setting of an accurate ferromagnetic coupling. Current state-of-the-art methods often scan a static range of values to find the appropriate ferromagnetic coupler strength.

In this context, we explore the advantages and drawbacks of different logical qubit encodings. This first analysis shows that the minimum spectral gap varies depending on the qubit encoding, as well as the minimum chain strength required to maintain ferromagnetic couplings. We perform a detailed analysis of the average performance of the two main existing minor embedding methods. This analysis helps to select sets of instances on which we compute the chain break tendency w.r.t the chain strength value. In particular, we discover that the optimal value of the chain strength varies depending on the embedding choice, which suggests that a per-instance chain strength setting method is desirable. The final contribution of this article is the design of a simple algorithm used to find appropriate values for the chain strength using very few preprocessing calls to the quantum computer. This new algorithm outperforms the default method implemented by D-Wave.

The rest of the paper is organized as follows: Section 2 introduces background about quantum annealing and minor embedding methods. Section 3 surveys the related work used in the study. Section 4 describes the method and technical settings for the experiments. Section 5 analyzes and proves the minimum spectral gap reduction induced by coupler strength rescaling. Section 6 gathers experiments on embedding and chain breaks that are used to build the algorithm presented in Section 7.

2 Quantum Annealing and Minor Embedding Methods

D-Wave systems are based on the transverse field Ising model. This model can be described with a linear interpolation of two Hamiltonians: a mixing Hamiltonian H_M , which ground state (i.e., state of lower energy which can be degenerate) is easy to prepare, and a problem Hamiltonian H_P , which ground state encodes the solution to the problem. The evolution of the system at annealing fraction $s = \frac{t}{T}$ is described by the Hamiltonian:

$$H(s) = (1 - s)H_M + sH_P \quad (1)$$

where T denotes the total run time of the quantum evolution and $t \in [0, T]$. This evolution scheme is simplified compared to the real annealing schedules of D-Wave systems available at [1]. Considering a Hamiltonian based on a graph $G_s = (V_s, E_s)$ with $n = |V_s|$, the transverse field Hamiltonian $H_M = \sum_{v \in V_s} \sigma_v^x$ has its ground state defined by a uniform superposition of all the quantum states of the computational basis. The problem Hamiltonian can be fully specified by the user:

$$H_P = \sum_{v \in V_s} h_v \sigma_v^z + \sum_{(u,v) \in E_s} J_{uv} \sigma_u^z \sigma_v^z. \quad (2)$$

The ground state of the Hamiltonian H_P gives the solution to the Ising cost function minimization problem :

$$\min C(\mathbf{x}) = \sum_{v \in V_s} h_v x_v + \sum_{(u,v) \in E_s} J_{uv} x_u x_v \quad (3)$$

where $x_u, x_v \in \{-1, +1\}$ and $h_v, J_{uv} \in \mathbb{R}$. The optimal solution is given by $\mathbf{x}^* = (x_0, x_1, \dots, x_{n-1})$. The Hamiltonian $H(s)$, at a fixed annealing fraction s , corresponds to a Hermitian matrix H_s , and can be decomposed in terms of its eigenvalues $E_i(s)$ and eigenvectors $|v_i(s)\rangle$:

$$H_s |v_i(s)\rangle = E_i(s) |v_i(s)\rangle \text{ with } E_0(s) < E_1(s) < \dots < E_k(s). \quad (4)$$

This decomposition is also called eigenenergies decomposition as the eigenvalues $E_i(s)$ correspond to the energy of each eigenvector $|v_i(s)\rangle$. The minimum spectral gap of the Hamiltonian Δ_{\min} is the difference between the energy of the first excited state and the energy of the ground state at any annealing fraction s :

$$\Delta_{\min} = \min_{0 \leq s \leq 1} (E_1(s) - E_0(s)). \quad (5)$$

The adiabatic theorem guarantees that a quantum state remains in its instantaneous ground state if T is chosen large enough to smooth the quantum evolution. In the best case, the time T scales as $O(1/\Delta_{\min}^2)$ [2].

When the Ising cost function of interest cannot be straightly mapped on the chip topology, one has to use a method to minor embed the problem into the quantum chip. This problem is well defined by the theory of graph minors developed by Robertson and Seymour [18]. The problem is defined as follows:

Given a source graph $G_s = (V_s, E_s)$ and a target graph $G_t = (V_t, E_t)$, the aim is to find a mapping function $\phi : V_s \rightarrow \mathcal{P}(V_t)$ such that :

1. each vertex $v \in V_s$ is mapped onto a connected subgraph $\phi(v)$ of G_t .
2. each connected subgraph must be vertex disjoint $\phi(v) \cap \phi(v') = \emptyset$, with $v \neq v'$.
3. each edge $(u, v) \in E_s$ is mapped onto at least one edge in $E_t : \forall (u, v) \in E_s, \exists u' \in \phi(u), \exists v' \in \phi(v)$, such that $(u', v') \in E_t$.

For each vertex $v \in V_s$, a ferromagnetic coupling strength $F_{\phi(v)}$ (also called chain strength) is applied to each edge of the subgraph $\phi(v)$. When this ferromagnetic coupling strength is the same for all ferromagnetic couplers, we note it F_ϕ . In the rest of the paper, we refer to V_s as the set of logical qubits and V_t as the set of physical qubits. A chain break on a logical qubit v means that at least one ferromagnetic coupling in $\phi(v)$ is corrupted. An edge of the subgraph $\phi(v)$ is a ferromagnetic coupling.

3 Related Work

Two categories of methods are used to find a minor embedding of an Ising problem. The first category comprises a set of heuristics for which the source graph G_s and target graph G_t are given as input. The state-of-the-art implementation for this category is the CMR heuristic of Cai et al. [4]. The first step of this algorithm aims to find an initial embedding of each logical qubit with possible overlaps between their associated connected subgraph of physical qubits. The second step is a refinement that tries to reduce this overlap to return a valid embedding. The second category of methods considers that the source graph G_s is a clique (i.e., a complete graph) and that the target graph G_t is static. One example is the method used to generate Clique Minor Embeddings (CME) on D-Wave systems [3], which is an iterative method that leverages the regular topology of the quantum annealer and considers inoperable qubits. Both CMR and CME methods have been extended with some pre or post-processing [21].

In [5], Choi formulates two bounds that estimate the value of the ferromagnetic coupling strength $F_{\phi(v)}$. The first bound is specific to each vertex $v \in V_s$:

$$F_{\phi(v)} < - \left(|h_v| + \sum_{u \in \text{nbr}(v)} |J_{uv}| \right) \quad (6)$$

where $\text{nbr}(v)$ gives the set of nodes connected to v . This bound is fast to compute with $O(D)$ complexity, where D is the vertex degree. V. Choi also details a more elaborated bound calculated in $O(DL)$ where L is the chain length. In the paper of Fang et al [9], the authors derive a new tighter bound that can be computed in $O(D2^L)$. The main idea is to set the chain strength with a negative strength with its absolute value greater than the maximum potential energy gain of any configuration of the physical qubits that are not part of the ground state. Thus, this method scales exponentially w.r.t the size of the chain length. Venturelli et al. [19] provided a second approach. They study the value of the optimal coupling strength by setting a global chain strength F_ϕ . They observe that the optimal coupling value F_ϕ grows as the critical point of temperature of their embedded Sherrington-Kirkpatrick model. Beyond this critical value, the success probability decreases. In the paper of Raymond et al. [17], the authors suggest that the chain strength should be tuned as:

$$\lambda = \lambda_0 \sqrt{\text{std}^2 n} = \lambda_0 \tau \sqrt{n} \quad (7)$$

where $std^2 = \frac{2}{n(n-1)} \sum_{(u,v) \in E_s} J_{uv}^2$ is the variance of the coupling strength (which is 1 for clique spin glasses). Their motivation is that a spin glass transition exists at optimal λ_0 . However, λ_0 remains to be set empirically. This paper led to the default method *uniform_torque_compensation* implemented on D-Wave systems [1], which sets the value of the chain strength:

$$F_\phi = -1.414 \times \sqrt{\bar{d}} \times \sqrt{\frac{1}{|E_s|} \sum_{(u,v) \in E_s} J_{uv}^2} \quad (8)$$

where \bar{d} is the average degree of the graph G_s . This formula comes from the fact that, for general Ising problems, the chain strength scales as $\tau\sqrt{n}$ where τ is the root mean square of the quadratic couplers. In practice, the value of the chain strength is mostly set with a basic chain scan method used to maximize the average expectation value of the QA. For detailed experimental studies on chain strength scanning, the reader may refer to [12,20]. The chain scan method performs well but is very expensive in the number of calls to the QA. The optimization of the chain strength is usually done based on the expectation value. A single advanced optimization method of the chain strength has been recently developed for the max-clique problem and is based on augmented lagrangian method [8]. Recent studies have benchmarked the chain break properties. Grant et al. [11] show that the chain break concentrates on specific locations of the D-Wave *2000Q* processor and design post-processing strategies to limit this bias. In another recent paper of Pelofske [16], it is seen that the approximation ratio is inversely correlated with the rate of getting chain breaks. The author shows that the optimal chain strength is also conditioned by the density of the problem and the type of quantum solver used. We reuse these conclusions as well as the first theoretical bound given in Eq. 6 to create an efficient method to set the chain strength.

4 Method

We test our methodology on two problems: the weighted Ising problem and the max-cut problem. Section 5 is based on a non-degenerate instance of the weighted Ising problem, which consists in minimizing the Ising cost function formulated in Eq. 3. The minimum spectral gap of the problem Δ_{\min} is calculated using the exact diagonalization of the Hermitian matrix H_s at each step of the annealing schedule $s \in [0, 1]$. The schedule used is the one corresponding to D-Wave *Advantage6.4* [1].

The max-cut problem is studied in Sections 6 and 7. The weights J_{uv} are set to 1 for all edges, h_v weights are set to 0. All the experiments are run at a constant annealing time of $100\mu s$. We use uniform logical weight spreading, majority vote to unembed the problem and the auto-scale method for weight rescaling. We do not use spin reversal methods or annealing offsets. The figures always show the chain strength in absolute value. The heatmaps from Fig. 2 are generated from sets of 100 instances for each size and density. Erdős-Rényi

graphs and random d -regular graphs are generated using the Python *networkx* library. The CMR heuristic runs until a valid embedding is found. The CME heuristic is the default method provided by D-Wave [1]. For all the experiments, we select the first valid embedding found. The plots from Fig. 3 are averaged over 30 instances of random Erdős-Rényi graphs of size $n = 80$ and density $p = 0.3$. The solver used is the *Advantage6.4* and the number of shots is set to 1024 for each chain strength value. The plots from Fig. 4 are obtained from a single Erdős-Rényi instance of size $n = 60$ and density $p = 0.4$. The solver used is the *Advantage2-prototype2.2* and the number of shots is set to 4096 for each chain strength and for the runs with *uniform_torque_compensation* method. The results obtained in Table 1 are obtained with solvers *Advantage6.4* and *Advantage2-prototype2.2*. The results are averaged over 30 instances of Erdős-Rényi graphs for each specified size and density. The number of shots in the preprocessing method is set to 128. The final run used to retrieve the statistics is set to 4096. The run with *uniform_torque_compensation* method is also set to 4096.

We compare the performance of CMR and CME methods in Fig. 2 using the average ratio of the number of qubits in mappings found by each method:

$$\bar{r}_{\text{emb}} = \frac{1}{N_p} \sum_{i=1}^{N_p} \frac{n_i^{\text{CMR}}}{n_i^{\text{CME}}} \quad (9)$$

where N_p is the number of instances, n_i^{CMR} and n_i^{CME} is the number of physical qubits found by the embedding method CMR resp. CME on instance i . The average breaking chain rate ϵ_b of a single shot is defined by:

$$\epsilon_b = \frac{1}{n_{\text{lq}}} \sum_{i=1}^{n_{\text{lq}}} p_b(i) \quad (10)$$

with n_{lq} the number of logical qubits. $p_b(i)$ is set to 1 if the logical qubit contains at least one broken chain and 0 otherwise. The average breaking chain rate of a serie of shots n_s is given by:

$$\bar{\epsilon}_b = \frac{1}{n_s} \sum_{i=1}^{n_s} \epsilon_b^{(i)}. \quad (11)$$

5 Logical Qubit Structure

We study the evolution of the minimum spectral gap Δ_{min} of a single weighted Ising instance by exhaustively simulating the quantum system evolution via matrix diagonalization. While Δ_{min} is not believed to directly impact the evolution of D-Wave systems that have a long annealing time [7], recent experiments have shown that these systems can operate with much shorter annealing times [14] (in the order of the nanosecond). It makes the study of the spectral gap relevant for future experiments. For this purpose, we design a single instance of an Ising

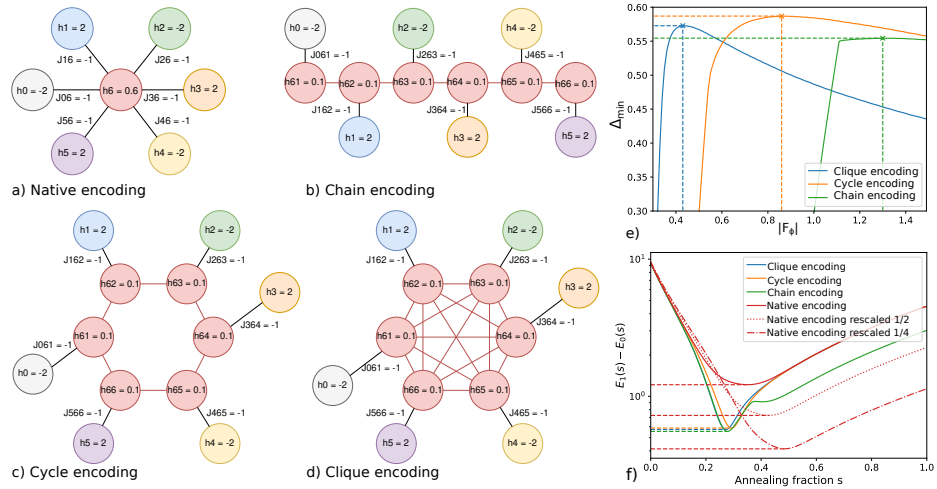


Fig. 1. Minimum spectral gap evaluations considering different types of logical qubits encoding a) Native Ising problem instance b) Same instance as in a. with the red qubit encoded as a chain of physical qubits c) Same instance as in a. with the red qubit encoded as a cycle of physical qubits d) Same instance as in a. with the red qubit encoded with a clique of physical qubits. Black edges represent logical couplers and red edges represent ferromagnetic couplers parametrized by the chain strength F_ϕ . The auto-coupler strength h_6 is uniformly spread on each physical qubit. e) Evolution of Δ_{\min} considering the whole annealing schedule for different values of the global chain strength $|F_\phi|$. f) Spectral gap evolution of each encoding type programmed with the corresponding optimal chain strength indicated by dashed lines in e.

problem with 6 nodes (see Fig. 1.a). We choose three possible encodings for the physical qubits $\phi(v_{\text{red}})$ that replace the red logical qubit v_{red} in Fig. 1.a. The selected structures are: chain, cycle and clique. Each structure is respectively shown in Fig. 1.b, 1.c and 1.d. Figure 1.e shows the evolution of the minimum spectral gap of each encoding according to the chain strength value used to maintain the ferromagnetic coupling. For each encoding, the chain strength starts at the minimum analytical value, which is sufficient to maintain the ferromagnetic coupling of the logical qubit (0.3 for the clique, 0.45 for the cycle and 0.9 for the chain). Each embedding exhibits some sweet spots that maximize the minimum spectral gap. The size of the minimum spectral gap decreases when the chain becomes too strong. The optimal value of the chain strength decreases with the density of the logical qubit encoding. This is a desirable feature as D-Wave quantum processors have a limited working range for auto-coupler (h_{range}) and coupler ($extended_j_range$) strengths. In addition, the maximum coupling range limit imposes that the total strength of the couplers linked to each qubit remains in a specific range. These two physical limitations lead to a global rescaling of coupler strengths when the values exceed these ranges.

Reusing the work of Choi [6], it is straightforward to demonstrate that rescaling the coupler strength in the problem Hamiltonian also rescales with the same factor the minimum spectral gap of the Ising Hamiltonian. Let the α -rescaled Hamiltonian with $s_2 \in [0, 1]$ and $\alpha > 1$:

$$H^{1/\alpha}(s_2) = (1 - s_2)H_M + s_2 \frac{1}{\alpha} H_P. \quad (12)$$

We take Eq. 1 as the initial Hamiltonian H with annealing fraction $s_1 \in [0, 1]$. Consider the system:

$$\begin{cases} \frac{H(s_1)}{1-s_1} = H_M + \frac{s_1}{1-s_1} H_P \\ \frac{H^{1/\alpha}(s_2)}{1-s_2} = H_M + \frac{s_2}{\alpha(1-s_2)} H_P \end{cases} \quad (13)$$

The equality $\frac{s_1}{1-s_1} = \frac{s_2}{\alpha(1-s_2)}$ can be solved with $s_1 = \frac{s_2}{(\alpha-1)(1-s_2)+1}$ and $s_2 = \frac{\frac{1}{\alpha}(1-s_1)+s_1}{1-s_1}$. Using this correspondance, we have:

$$\frac{H^{1/\alpha}(s_2)}{1-s_2} = \frac{H(s_1)}{1-s_1}. \quad (14)$$

The rescaled Hamiltonian has the form:

$$H^{1/\alpha}(s_2) = \frac{1-s_2}{1-s_1} H(s_1) = \left(1 + \left(\frac{1}{\alpha} - 1\right) s_2\right) H(s_1). \quad (15)$$

The eigenenergies of the Hamiltonian are then rescaled with the same factor. According to Eq. 4, we have:

$$E_i^{1/\alpha}(s_2) = \left(1 + \left(\frac{1}{\alpha} - 1\right) s_2\right) E_i(s_1). \quad (16)$$

Figure 1.f shows the effect of couplers strength rescaling on the spectral gap. This figure shows that rescaling the global coupling strength reduces the minimum spectral gap by the same factor. It also shifts the spectral gap to the right. Encoding logical qubits on a set of physical qubits has a detrimental effect on the spectral gap of the instance. The difference in spectral gap reduction between these encodings seems negligible compared to the rescaling of the weights. Logical qubit encodings that are dense, such as clique, require a lower coupling strength to maintain ferromagnetic coupling within the physical qubits. This type of encoding could be favored compared to chain encoding in specific cases to limit the effect of coupler strength rescaling imposed by D-Wave systems.

6 Embeddings and Chain Break Analysis

Figure 2 compares the embedding performance in terms of the number of qubits used by the CMR method [4] against the CME method [3]. The comparison is done on Erdős–Rényi and d-regular graphs. It shows that CMR performs better on sparse graphs of small size, while CME is almost always preferred for large

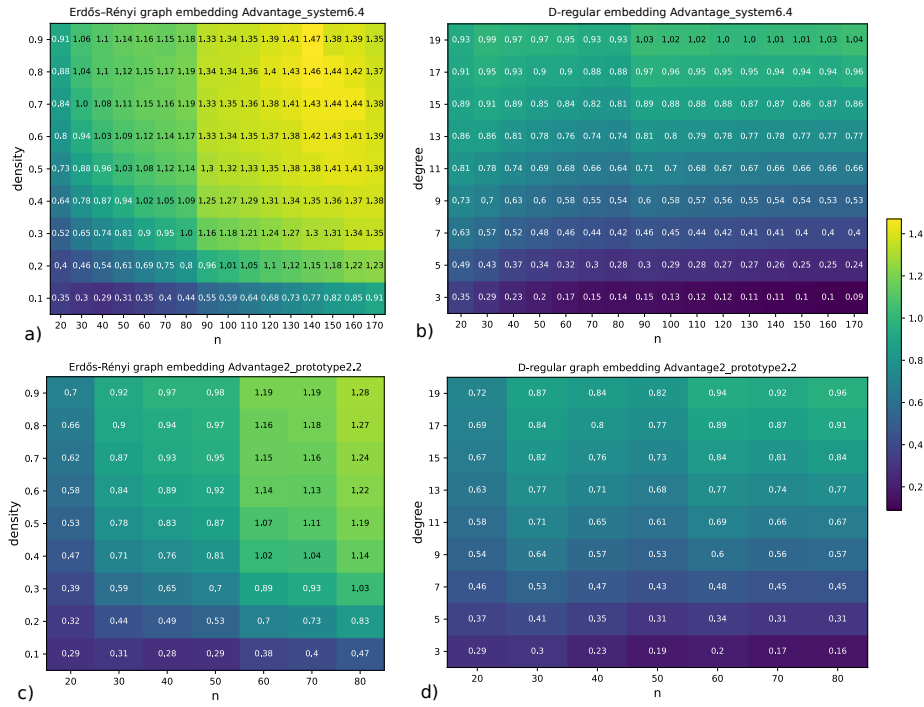


Fig. 2. Heatmaps showing the average percentage overhead of the number of qubits used to embed similar instances using either CMR or CME methods. Each score is an average done over 100 instances. The score calculation in each cell is detailed in Eq. 9. a) and b) are embeddings generated for *Advantage6.4* topology for Erdős-Rényi and d-regular graphs. b) and c) are embeddings generated for *Advantage2_prototype2.2* topology using the same instances. *Advantage6.4* and *Advantage2_prototype2.2* can embed complete graphs of maximum size 174 and 82.

graphs. We use this heatmap to select sets of instances for which CMR and CME methods produce embeddings with approximately the same number of physical qubits. We set each edge weight to 1 to build the max-cut instance associated with the graph. We choose to generate 30 random instances of the max-cut problem with size $n = 80$ and density $p = 0.3$. The aim is to analyze the impact of the embedding method on the chain strength and length. CMR and CME produce embeddings with a similar number of physical qubits for these instances. To be fair in the comparison, we force the CMR method to generate minor embeddings with $\pm 1\%$ qubits compared to the embeddings generated by CME. Figure 3.a and b. respectively show the repartition of the chain length for CMR resp. CME embeddings. While the CME method produces almost uniform chains of lengths 8 and 9, the CMR method produces chain lengths that approach a Gaussian distribution and vary from 4 to 15. The coupling strength required by the CME embeddings to obtain the same breaking chain

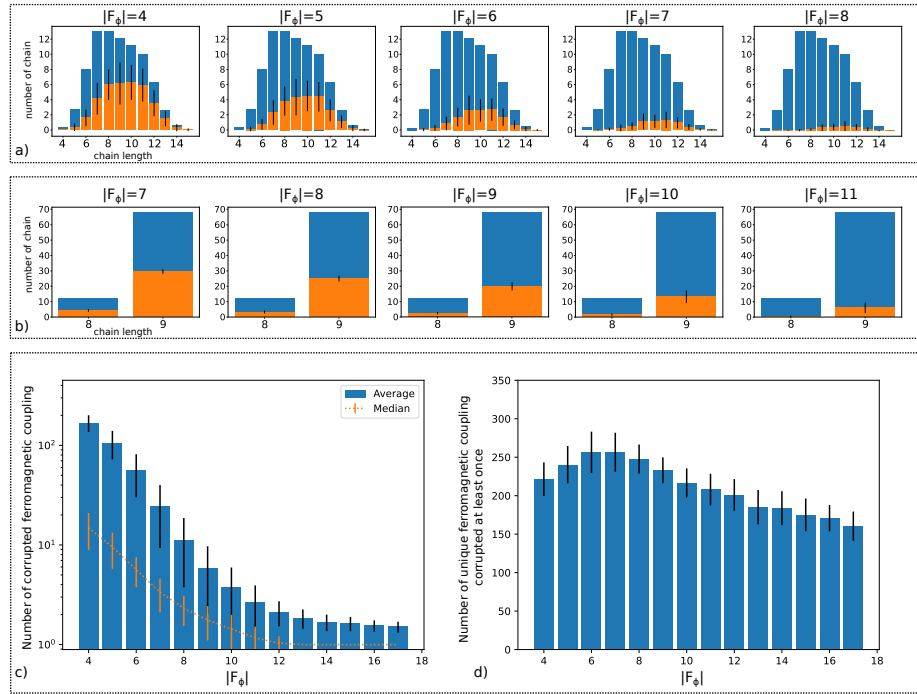


Fig. 3. Statistics on the breaking chain rate (see Eq. 11) averaged over 30 instances of Erdős-Rényi graphs of size $n = 80$ and density $p = 0.3$. a) resp. b) show the average chain length repartition of embedding obtained with CMR resp. CME methods. Blue bars show the average frequency of each chain length. Orange bars show the average breaking chain rate with a black error bar for the standard deviation. c) shows the average and median frequency of corrupted ferromagnetic couplings on CMR embeddings. d) shows the number of different ferromagnetic couplings that are corrupted at least once during the 1024 shots on CMR embeddings.

rate is higher than for CMR embeddings. For the CMR method, the optimal solution probability occurs at chain strength 8, while this value is raised to 12 for CME embeddings. Figure 3.d shows that the number of different corrupted couplings only slightly changes when the chain strength is increased. Figure 3.c shows that when the chain strength is insufficient to maintain ferromagnetic couplings, the average number of corrupted ferromagnetic couplings is very high compared to the median. It suggests that the corruptions concentrate on the same few ferromagnetic couplings. When the chain strength is increased, the average number of corruptions decreases but the number of different corrupted couplings remains the same. It suggests that the corruption becomes sparse when the chain strength becomes sufficient. We computed statistics on the number of corrupted ferromagnetic couplings per logical qubit. At the optimal chain strength value, the most likely scenario is to have only a single corruption of

couplings on the connected subgraph $\phi(v)$ (99.9% cases) and very few chances of double corruption of couplings (0.01% cases).

The above experiment shows that the optimal chain strength value is related to the embedding. We select another instance that produces approximately the same number of qubits for both embedding methods for the *Advantage2_prototype2.2*. We generate four different embeddings for a single graph of 60 nodes and 0.4 density. The first embedding is generated by the CME method. We generate three other embeddings using the CMR method: one with a similar number of qubits as in the CME embedding ($\pm 1\%$), one which has 10% less qubits than the CME embedding, one which has 10% more qubits than the CME embedding. These graphs are then used to create the max-cut instances. Figure 4 shows the best cut size found with each embedding. At first, we can see that the *uniform_torque_compensation* method performs well in all the cases compared to the chain scan that requires heavy preprocessing (93% of the best cut size in the worst case with CME embedding). For each instance, the best cut size reaches a plateau when the chain strength becomes sufficient. This plateau

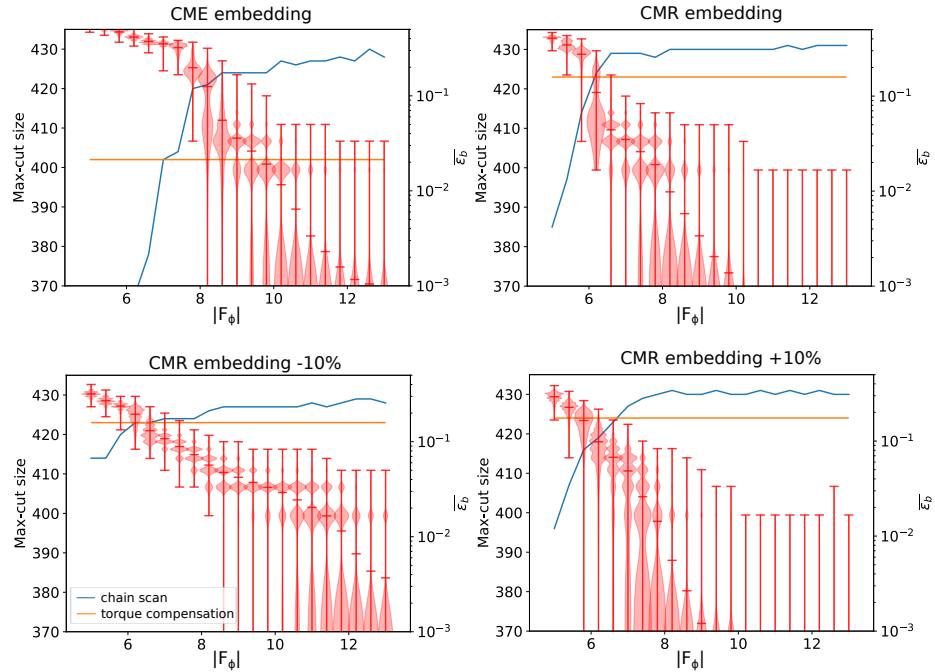


Fig. 4. Maximum cut size obtained with a chain scan on four different embeddings of the same instance. The *uniform_torque_compensation* is run once for each embedding (hence, it is independent of the chain scan). The CME and upper right CMR embedding have the same number of physical qubits ($\pm 1\%$). The two other embeddings have 10% more and 10% less qubits than the CME embedding. The red violins show the average breaking chain rate $\bar{\epsilon}_i$ related to the chain scan curve.

is reached at different chain strength values for each instance. This plateau can also be located by only considering the average breaking chain rate from Eq. 11 (for example, considering that the chain break probability should remain under 2×10^{-2}). As specified in [15], the chain break probability susceptible to provide the best result depends also on the quantum computer.

7 Chain Strength Setting Heuristic

The previous section detailed our motivations for designing a heuristic to set the chain strength for general instances of Ising models. Algorithm 1 describes this method in pseudo-code. The algorithm performs a binary search to find the optimal chain strength value within an initial interval of chain strengths *csInterval*. When the user does not specify it, the default interval's upper bound is set according to Choi's first bound (see Eq. 6 and line 3):

$$|F_\phi| = \max_{v \in V_s} |F_{\phi(v)}|. \quad (17)$$

At each iteration, the chain strength *cs* is chosen as the midpoint of the chain strength interval *csInterval* (line 6). The embedded instance *embInstance* is then sent to the quantum annealer with the specified chain strength *cs* (line 7). We then compute the breaking chain rate of the result according to Eq. 11 and check if the new breaking chain rate is in the interval specified by the user *cbInterval*. If this is the case, the algorithm converges (lines 9 and 10), and the loop breaks. If the chain break is higher than the upper bound of *cbInterval*, the chain strength is insufficient and requires an increase: we then rescale the lower

Algorithm 1 Chain strength binary search

Require: *embInstance*, *cbInterval*, *csInterval* (nullable)
Ensure: *cbInterval*[0] < *cbInterval*[1]

- 1: *hasConverged* \leftarrow *False*
- 2: **if** *csInterval* is None **then**
- 3: *csInterval* \leftarrow [0, *getUpperBound*(*embInstance*)] ▷ (see Eq. 17)
- 4: **end if**
- 5: **while** not *hasConverged* **do**
- 6: *cs* \leftarrow *csInterval*[0] + (*csInterval*[0] – *csInterval*[1])/2
- 7: *res* \leftarrow *runQA*(*embInstance*, *cs*)
- 8: $\bar{\epsilon}_b \leftarrow$ *getChainBreakRate*(*res*) ▷ (see Eq. 11)
- 9: **if** *cbInterval*[0] \leq $\bar{\epsilon}_b \leq$ *cbInterval*[1] **then**
- 10: *hasConverged* \leftarrow *True*
- 11: **else if** $\bar{\epsilon}_b >$ *cbInterval*[1] **then**
- 12: *csInterval*[0] \leftarrow *cs*
- 13: **else**
- 14: *csInterval*[1] \leftarrow *cs*
- 15: **end if**
- 16: **end while**

bound of $csInterval$ (lines 11 and 12). In the other case, the upper bound of the chain strength interval is rescaled (line 14). The subtlety of this algorithm relies in the setting of the chain break interval $cbInterval$ which stops the algorithm when the sampling produces a breaking chain rate in this interval. This interval depends on the effective noise of the quantum computer and may also depend on the size of the instance. We test our heuristic on the max-cut problem and compare the results obtained with the default method implemented by D-Wave that relies on Eq. 8. The comparison is done both on *Advantage2_prototype2.2* on small Erdős-Rényi instances ($n = 40$ and $n = 80$) and *Advantage6.4* for large instances ($n = 100$ and $n = 170$) with density $p \in \{0.1, 0.5, 0.9\}$. For each instance set, we choose the embedding heuristic that provides the best average performance based on the size and density of the instance (see Fig. 2). We empirically set the interval $cbInterval$ to $[6 \times 10^{-3}, 2 \times 10^{-2}]$ for the *Advantage2_prototype2.2* and to $[2 \times 10^{-2}, 5 \times 10^{-2}]$ for the *Advantage6.4*. This range stays untouched during the whole experiment. Each preprocessing step is done with 128 shots. The final run of our heuristic, as well as the default D-Wave method, is programmed with 4096 shots for instances run on *Advantage2_prototype2.2* and 3072 shots for instances run on *Advantage6.4*. Table 1 shows the result of this experiment. With a global setting of the value for $cbInterval$, our heuristic is able to outperform in almost every case the default *uniform_torque_compensation* method. The breaking chain rate can be obtained

Table 1. Performance comparison between D-Wave’s default method (*uniform_torque_compensation*) and our binary search heuristic on max-cut instances. The column Best cut size shows the average of the maximum cut size obtained for each instance with the *uniform_torque_compensation* method. The column Cut size improvement contains the minimum and maximum cut size improvement obtained with the chain strength found by Algorithm 1. The standard deviation is shown in column std. The column Step counts the average number of iterations required by our heuristic.

Advantage2_prototype2.2			Best cut size	Cut size improvement				Step
Instance size	Density	Embedding		min	max	mean	std	
$n = 40$	0.1	CMR	66.4	+0%	+0%	+0%	0%	5.4
	0.5	CMR	243	+0%	+2%	+0.2%	0%	5.8
	0.9	CMR	362.8	+2.1%	+8.2%	+5%	1.6%	4.6
$n = 80$	0.1	CMR	235.5	+0%	+0%	+0%	0%	4.7
	0.5	CME	804	+9.8%	+17.2%	+12.5%	0.2%	4.2
	0.9	CME	1435	+2%	+4.7%	+3.2%	0.6%	4.2
Advantage6.4			Best Cut size	Cut size improvement				Step
Instance size	Density	Embedding		min	max	mean	std	
$n = 100$	0.1	CMR	355.9	+0%	+0.3%	+0%	0%	4.5
	0.5	CME	1271.4	+5.6%	+14.5%	+8.8%	1.8%	2.7
	0.9	CME	2243	+1.4%	+3.7%	+2.5%	0.5%	3.7
$n = 170$	0.1	CMR	950.8	-2.1%	+0.6%	-0.5%	0.5%	2.1
	0.5	CME	3631.4	+2.8%	+6.2%	+4.5%	0.7%	2.1
	0.9	CME	6519.4	+0.4%	+1.4%	+0.8%	0.2%	3.2

with relatively high fidelity in few shots. It gives a considerable advantage to our method compared to the basic chain strength scans that always use the expectation value to find optimal values of the chain strength and, hence, require a very high number of shots. Our optimization method does, on average, between 2 and 6 iterations to find a suitable value for the chain strength, which corresponds to approximately 750 extra shots per instance in the worst case. This overhead is almost negligible compared to the 4096 and 3072 shots used to evaluate the final chain strength and the *uniform_torque_compensation* method. The default *uniform_torque_compensation* performs well on embeddings generated by CME for dense graphs and on embeddings generated by CMR on sparse graphs. Our heuristic seems to perform the best with mid-density instances with the increase of the cut size by up to 17.2% for the *Advantage2-prototype2.2* and 14.5% for the *Advantage6.4*.

8 Conclusion

This paper presented a detailed analysis of the minimum spectral gap evolution considering different topologies for the set of physical qubits representing a single logical qubit. Each encoding required a specific chain strength value to maintain ferromagnetic couplings between the physical qubits. The denser the logical qubit encoding was, the lower the chain strength had to be. This feature is desirable as the coupler strength rescaling, usually driven by chain strength values in combinatorial problems, reduces the minimum spectral gap of the problem, which may have an impact on very short annealing time evolutions. This consideration could be included in future embedding heuristics designs to enhance the quality of the generated minor embeddings.

The analysis of the breaking chain rate considering different embeddings of the same instance has shown that the optimal chain strength varies depending on the embedding used. This experiment led to the design of a simple but fast heuristic used to optimize the chain strength for each instance. The heuristic converged in most cases in less than 5 preprocessing steps. This number has to be considered cautiously as it strongly depends on the chain break interval, which acts as the breaking condition of the heuristic. We used large intervals in this experiment. The precise identification of the breaking chain rate that produces the best results for each quantum computer could help to refine these bounds. Even if this new method does not provide better solutions than the basic chain scan, the original use of the breaking chain rate in the optimization process drastically reduces the preprocessing time required by a simple chain scan and does not require any assumption over the scanning range of chain strengths.

The performance of our heuristic can be questioned due to the relatively low gain on the max-cut size (17% in the best case). However, even a small percentage gain in the cut size can reduce the Time To Solution of several orders of magnitude, which can justify the use of our heuristic.

A detailed study on the determination of refined breaking chain intervals for each solver is a relevant perspective. This heuristic could also be used as a fast

preprocessing routine to find relatively good global chain strengths, followed by a refinement step optimizing each ferromagnetic coupling strength individually.

References

1. D-wave system. solver properties and parameters (2024), https://docs.dwavesys.com/docs/latest/doc_solver_ref.html
2. Albash, T., Lidar, D.A.: Adiabatic quantum computation. *Reviews of Modern Physics* **90**(1), 015002 (2018)
3. Boothby, T., King, A.D., Roy, A.: Fast clique minor generation in chimera qubit connectivity graphs. *Quantum Information Processing* **15**, 495–508 (2016)
4. Cai, J., Macready, W.G., Roy, A.: A practical heuristic for finding graph minors. arXiv preprint arXiv:1406.2741 (2014)
5. Choi, V.: Minor-embedding in adiabatic quantum computation: I. the parameter setting problem. *Quantum Information Processing* **7**, 193–209 (2008)
6. Choi, V.: The effects of the problem hamiltonian parameters on the minimum spectral gap in adiabatic quantum optimization. *QIP* **19**(3), 90 (2020)
7. Dickson, N.G., Johnson, M.W., et al.: Thermally assisted quantum annealing of a 16-qubit problem. *Nature Communications* **4**(1) (2013)
8. Djidjev, H.N.: Logical qubit implementation for quantum annealing: augmented lagrangian approach. *Quantum Science and Technology* **8**(3), 035013 (2023)
9. Fang, Y.L., Warburton, P.: Minimizing minor embedding energy: an application in quantum annealing. *Quantum Information Processing* **19**(7), 191 (2020)
10. Gilbert, V., Rodriguez, J.: Discussions about high-quality embeddings on Quantum Annealers. In: EU/ME meeting. Troyes, France (2023)
11. Grant, E., Humble, T.S.: Benchmarking embedded chain breaking in quantum annealing. *Quantum Science and Technology* **7**(2), 025029 (2022)
12. Hamerly, R., Inagaki, T., McMahon, et al.: Experimental investigation of performance differences between coherent ising machines and a quantum annealer. *Science Advances* **5**(5) (2019)
13. Kadowaki, T., Nishimori, H.: Quantum annealing in the transverse ising model. *Phys. Rev. E* **58**, 5355–5363 (Nov 1998)
14. King, A.D., Nocera, A., et al.: Computational supremacy in quantum simulation. arXiv preprint arXiv:2403.00910 (2024)
15. Pelofske, E.: 4-clique network minor embedding for quantum annealers. arXiv preprint arXiv:2301.08807 (2023)
16. Pelofske, E.: Comparing three generations of d-wave quantum annealers for minor embedded combinatorial optimization problems. arXiv:2301.03009 (2023)
17. Raymond, J., Ndiaye, N., et al.: Improving performance of logical qubits by parameter tuning and topology compensation. In: 2020 IEEE International Conference on Quantum Computing and Engineering (QCE). pp. 295–305. IEEE (2020)
18. Robertson, N., Seymour, P.: Graph minors .XIII. the disjoint paths problem. *Journal of Combinatorial Theory, Series B* **63**(1), 65–110 (1995)
19. Venturelli, D., Mandrà, S., et al.: Quantum optimization of fully connected spin glasses. *Physical Review X* **5**(3), 031040 (2015)
20. Willsch, D., Willsch, M., et al.: Benchmarking advantage and d-wave 2000q quantum annealers with exact cover problems. *QIP* **21**(4) (2022)
21. Zbinden, S., Bärttschi, A., et al.: Embedding Algorithms for Quantum Annealers with Chimera and Pegasus Connection Topologies, p. 187–206. Springer International Publishing (2020)