

Approximate combinatorial optimization with Rydberg atoms: the barrier of interpretability

Christian de Correc,^{1,*} Thomas Ayrat,¹ and Corentin Bertrand¹

¹*Eviden Quantum Lab, 78340 Les Clayes-sous-Bois, France*

Analog quantum computing with Rydberg atoms is seen as an avenue to solve hard graph optimization problems, because they naturally encode the Maximum Independent Set (MIS) problem on Unit Disk (UD) graphs, a problem that admits rather efficient approximation schemes on classical computers. Going beyond UD-MIS to address generic graphs requires embedding schemes, typically with chains of ancilla atoms, and an interpretation algorithm to map results back to the original problem. However, interpreting approximate solutions obtained with realistic quantum computers proves to be a difficult problem. As a case study, we evaluate the ability of two interpretation strategies to correct errors in the recently introduced Crossing Lattice (CL) embedding. We find that one strategy, based on finding the closest embedding solution, leads to very high qualities, albeit at an exponential cost. The second strategy, based on ignoring defective regions of the embedding graph, is polynomial in the graph size, but it leads to a degradation of the solution quality which is prohibitive under realistic assumptions on the defect generation. Moreover, more favorable defect scalings lead to a contradiction with well-known approximability conjectures. Therefore, it is unlikely that a scalable and generic improvement in solution quality can be achieved with Rydberg platforms—thus moving the focus to heuristic algorithms.

I. INTRODUCTION

The Unit Disk (UD)-Maximum Weighted Independent Set (MWIS) problem is the native use case of Rydberg quantum computing platforms due to its intimate link with the Rydberg blockade mechanism [1–6]. However, even if it is a NP-complete problem [7, 8], finding approximate solutions to UD-MWIS is relatively easy with classical algorithms [9–13]: one can get arbitrarily close to the optimal solution in a time polynomial in the graph size. It raises the bar for reaching a putative quantum advantage over classical methods [14]. This has spurred proposals to tackle the general Maximum Independent Set (MIS) problem—which is much harder to approximate [15, 16]—with Rydberg platforms. This is done by mapping (or "embedding") the general MIS problem to a larger UD-MWIS problem [1, 17, 18]. In principle, an optimal solution of this UD-MWIS problem is then obtained with a quantum computer, for instance via an adiabatic evolution [19, 20], and can be interpreted to produce a solution of the original MIS problem as schematized in Fig. 1.

In this work, we investigate the robustness of embedding methods to imperfections in the UD-MWIS solution: imperfect UD-MWIS solutions, which are bound to appear in quantum algorithms due e.g to too short annealing times [21–23], cannot be directly mapped back to MIS solutions. We study two complementary interpretation methods, akin to error mitigation methods, that handle these imperfections to reconstruct approximate MIS solutions.

More specifically, we focus on the recently introduced embedding called the Crossing Lattice (CL) [18], which




FIG. 1. General procedure to solve MIS on a generic graph by the means of an adiabatic evolution on a UD graph. Red: selected vertices after the UD-MWIS exact resolution.

comes with a quadratic overhead in the number of vertices. While various interpretation algorithms are possible, we focus on two opposing strategies. The first one, introduced as the *distance* strategy in Sec. II, arguably yields UD-MWIS solutions closer to the optimal one, but is computationally expensive. In Sec. III, we give evidence that this strategy is limited by features of the CL solution space. The second strategy, called *deselection* strategy, proves to be inexpensive but yields lower quality results (Sec. IV). In Sec. V we use our numerical results to derive scaling laws and discuss their implications in terms of the approximability of the MIS solution with Rydberg platforms, before concluding in Sec. VI.

* c.de-correc@eviden.com

II. METHOD AND DEFINITIONS

Let $G = (E, V, W)$ be a simple finite weighted graph with set of vertices V , set of edges E and weight function $W : V \rightarrow \mathbb{R}_+^*$. A subset $S \subset V$ is said to be an Independent Set (IS), or $S \in \text{IS}(G)$, when no pair of vertices in S is adjacent in G . The MWIS problem amounts to selecting vertices that form an IS with maximal total weight. It is equivalent to finding a ground state of the Ising Hamiltonian

$$\mathcal{H}_{\text{MWIS}} = -\delta \sum_{i \in V} W(i) n_i + U \sum_{(i,j) \in E} n_i n_j \quad (1)$$

where $n_i = 1$ if the vertex i is selected and 0 otherwise. Here, $\delta > 0$ is an energy unit and $U \gg \delta$ is the interaction energy which implements the independence constraint in the limit $U \rightarrow +\infty$. On UD graphs, this Hamiltonian can be implemented with Rydberg atoms, thanks to the Rydberg blockade that acts as a repulsive interaction between neighboring atoms in the excited Rydberg state [1, 24]. For the rest of this paper we assume a perfect implementation of the independence constraint, therefore ignoring the impact of the Rydberg interaction long range tail. We also set the energy scale to $\delta = 1$ and omit the notation of δ from now on.

On arbitrary graphs, the MWIS problem can be reduced to UD-MWIS by using an *embedding*. In this work, for simplicity, we focus on MWIS embeddings of the MIS problem, that is when $W = 1$ for the embedded graph. We don't expect the generalization to MWIS to significantly change our conclusions. Formally, we define an embedding as a polynomial-time algorithm, which maps an instance G of MIS onto an instance G_{em} of UD-MWIS called *embedding graph*, and an injective function $f : \text{IS}(G) \rightarrow \text{IS}(G_{\text{em}})$ called *embedding map*. The map f should preserve the weight ordering, i.e. for every pair $S, T \in \text{IS}(G)$,

$$W(S) \geq W(T) \implies W_{\text{em}}(f(S)) \geq W_{\text{em}}(f(T)), \quad (2)$$

where W, W_{em} are the weight profiles on G, G_{em} . Being injective, f admits an inverse f^{-1} defined on $f(\text{IS}(G))$ and called the *interpretation map*. Finally, we require f^{-1} to map any MWIS of G_{em} to an MWIS of G .

Being given an embedding scheme with some G, G_{em} and f , any configuration $S_{\text{em}} \in f(\text{IS}(G))$ is said to be *interpretable*, since $f^{-1}(S_{\text{em}})$ is well-defined and is an IS for G . Conversely, any configuration $S_{\text{em}} \notin f(\text{IS}(G))$ is said to be *non-interpretable*. Non-interpretable states arise due to the approximate resolution of UD-MWIS on G_{em} , and they require some post-processing to be interpreted as configurations of the initial problem on G .

One possible strategy, which we name *distance* strategy, consists in interpreting each configuration of the CL as the nearest interpretable configuration, with regard to the Hamming distance. This distance, denoted d ,

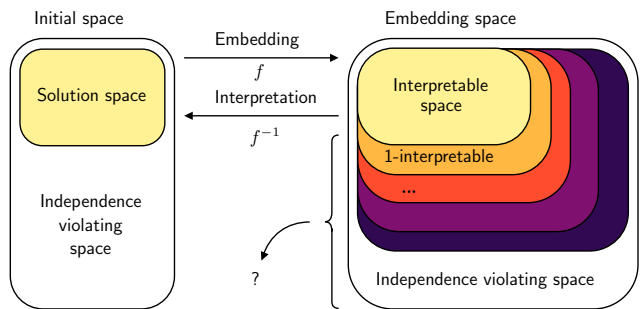


FIG. 2. Structure of the initial and embedding solution spaces for a given embedding scheme, with sub-spaces of solutions characterized by their Hamming distance $d = 0, 1, \dots$ to the nearest interpretable solution.

can be seen as the minimum number of bits to flip to reach an interpretable state. We take d both as a measure of the difficulty to apply this strategy and as a quantity of interest to study the CL solution space. A large d indicates that a larger part of the CL must be explored to find the nearest interpretable configuration, which correlates to a larger computational difficulty. As shown in Appendix E, applying this distance strategy is equivalent to solving a Quadratic Unconstrained Binary Optimization (QUBO) problem closely related to an MWIS problem on a subgraph of the original graph, showing that it can be a difficult task. Note that our definition of d and most subsequent discussions readily apply to other embeddings schemes than the CL, such as the ones discussed in App. C.

Building upon these definitions, any state with a distance (at most) d to the nearest interpretable state is thus called (at most) *d-interpretable*. With this definition, interpretable states are characterized by $d = 0$, and higher values of d are associated to states that are harder to interpret. As illustrated in Fig. 2, this gives a layered structure to the embedding space $\text{IS}(G_{\text{em}})$, with easy-to-interpret solutions at the center, and increasingly difficult configurations around.

Throughout this article, the quality of approximate solutions is evaluated by the *approximation ratio*, as is standard in optimization problems. For any $S \in \text{IS}(G)$, its approximation ratio $r \in [0, 1]$ is defined as the factor by which its weight deviates from the optimal value:

$$r = \frac{\sum_{i \in S} W(i)}{\max_{S' \in \text{IS}(G)} \sum_{i \in S'} W(i)}. \quad (3)$$

III. DISTANCE STRATEGY AND INTERPRETABILITY

We perform a numerical study of d -interpretable states at low energies, obtained thanks to an exact MWIS




FIG. 3. Path embedding of a single-vertex graph into two interpretable configurations (*green highlight*), with selected vertices in red. Due to domain walls (*red highlight*), some configurations with a low energy have a high Hamming distance d to the closest interpretable configuration.

solver [25]. In Sec. III A, we begin by applying our methodology on a simple class of graphs called *paths graphs* that can be seen as elementary components of CL graphs. In Sec. III B we extend our study of the distribution of d to CL graphs.

A. Path embeddings

A core ingredient in CL embedding is the class of *path graphs*, which are UD graphs used to embed the state of a single-vertex graph. As a preliminary, we study them separately before looking at the CL.

For any $N \geq 0$, the (unweighted) path graph P_N is defined as N vertices v_1, \dots, v_N and edges (v_i, v_{i+1}) for all $i = 1, \dots, N - 1$. Of particular interest is the case of the even-length path P_{2N} . As shown in Fig. 3, in the unweighted case, the MIS of P_{2N} has N selected vertices (in red) and a $(N + 1)$ -fold degeneracy i.e. number of configurations at the same energy. In this context, we call *domain wall* the presence of two adjacent unselected vertices. In the unweighted P_{2N} , only two MISs have no domain wall; they are said to be *antiferromagnetic*.

Antiferromagnetic configurations can be used to embed in P_{2N} a single-vertex graph $G = (V = \{v\}, E = \emptyset)$. The embedding map f is defined such that $f(\{v\})$ is the antiferromagnetic state of P_{2N} where v_1 is selected, and $f(\emptyset)$ is the other one. The fulfillment of Eq (2), and thus the optimality of the embedded MIS of G , requires to energetically penalize domain walls, thus lifting the degeneracy between the two antiferromagnetic states. To do so, we follow Ref. [18], applying a weight profile W_w , parametrized by a *weight bias* $w \in (0, 1/2)$, which reads

$$\begin{cases} W_w(v_1) &= 1/2 + w, \\ W_w(v_{2N}) &= 1/2 - w, \\ W_w(v_i) &= 1 \quad \text{for } 1 < i < 2N. \end{cases} \quad (4)$$

We now look at the corresponding Density of States (DoS). The number of states at each energy is displayed in the top panel of Fig. 4 as a function of the energy

$$\Delta E = E - E_{\text{MWIS}}, \quad (5)$$




FIG. 4. (*Top*) Low-energy DoS in the path graph P_{40} with $1/w \in \{4, 20\}$. On the left, the two interpretable states have a gray outline to enhance readability. (*Bottom*) Distribution of τ_d in increasing energy windows, with an exponential decay of $\tau_0(E)$ (dotted line).

and with a color scheme to indicate the distribution of the different values of d at each energy. Note that the energy is linearly related to the approximation ratio $r \in [0, 1]$ defined in Eq. (3), with $r \simeq 1$ for the best approximate solutions. A closed formula for this DoS is derived in App. A using graph combinatorics [26, 27], while the value of d is computed algorithmically. We observe that, even at very low energy $\Delta E \leq 1$, there is a high number of non-interpretable states with high values of d . This can be interpreted in terms of domain walls, introduced earlier. As shown in Fig. 3, the presence of only one domain wall in P_{2N} can lead to a high Hamming distance, up to $d = N$. There is thus a transition around the domain-wall energy penalty $\Delta E = 1/2 + w$, with interpretable states located at smaller energies, and non-interpretable ones right above.

To interpret this data in terms of probability of success after an adiabatic evolution, we look at the fraction $\tau_d(E)$ of at most d -interpretable states within the energy window $[E_{\text{MWIS}}, E]$. Formally,

$$\tau_d(E) = \frac{\int_{E_{\text{MWIS}}}^E d\omega \sum_{d' \leq d} \rho_{d'}(\omega)}{\int_{E_{\text{MWIS}}}^E d\omega \rho(\omega)} \quad (6)$$

where ρ is the total DoS and $\rho_{d'}$ is its restriction to d' -interpretable states, with $\rho = \sum_{d'} \rho_{d'}$. This measures the probability to get an easy-to-interpret state from random picking below energy E , and is shown in Fig. 4 (*bottom*). In the regime $\Delta E < 1/2 + w$, all states are interpretable




FIG. 5. (a) Initial graph with $N = 4$ vertices, $|E| = 3$ edges (*plain lines*) and $N(N - 1)/2 - |E| = 3$ pairs of non-adjacent vertices (*dashed lines*). (b) Corresponding CL embedding, with N even-length paths used to embed the N initial vertices and interconnected by $N(N - 1)/2$ gadgets. Colors identify the mapping between initial vertices and paths in the CL. (c) Legend, with two types of gadgets subgraphs representing initial edges (*plain diamond*) or their absence (*dotted diamond*).

since no domain walls are allowed. When E increases and $\Delta E \lesssim N/4$, $\tau_d(E)$ appears to decrease exponentially for all values $d \ll N$, before reaching a plateau near $\Delta E \sim N/2$. This observation can be accounted for by the high degeneracy of domain walls, which is expected to grow exponentially with the energy. Note that this behavior is consistent when w varies, including when $w \ll 1$ (Fig. 4, right column), for which the energy of a configuration is $\simeq l/2$ with l the number of defects in the chain.

We conclude that in the path graph embedding, states of energy above the threshold $\Delta E = 1/2 + w$ become difficult to interpret with the distance strategy, and exponentially more so as their energy increases. Quantum adiabatic computation should therefore aim at energies no much greater than this threshold. In the next section, we extend this result to CL graphs.

B. Crossing lattice embedding

Building upon the idea of the path graph embedding, the CL scheme [18] starts by embedding the N vertices of an initial graph G into N copies of P_{2N} (or P_{2N+2}). These paths are then intertwined along a square lattice grid, and interconnected by two types of subgraphs called *gadgets* to enforce constraints caused by the presence or absence of edges between pairs of vertices in G . A schematic example is displayed in Fig. 5.

In what follows, we examine the DoS of CL embeddings of unweighted random Erdős–Rényi–Gilbert graphs [28] with edge probability $p = 1/2$. We generate 20 unweighted graphs for each size $N = 3, \dots, 7$, and study their CL embedding whose vertex count $N_{\text{CL}} \sim 4N^2$ goes from 31 to 188. The value of $1/w$ is taken to be 8 or 20. These parameters are discussed in App. B, although their exact choice is not expected to change our results.

A typical example of low-energy DoS is shown in Fig. 6 (a, *top left panel*), as a function of ΔE or equivalently of the approximation ratio r , with $1 - r \propto \Delta E$. The light yellow bars with dark contour on the left are the interpretable states, and form a rescaled copy of the initial graph DoS, with the optimal solution at the lowest energy and solutions of increasing approximation ratio as energy increases. The other bars represent states with defects, which start appearing above the energy price of a domain wall $\Delta E = 1/2 - w$. Note that, unlike path embeddings, CL constraints make sure some high weight ($1/2 + w$) vertices are unselected in the ground states, for almost all initial graphs, such that the lowest energy price of a defect differs from path embeddings. The fraction of d -interpretable states $\tau_d(E)$ for increasing energy windows is shown below (a, *bottom left panel*). In good agreement with the analysis from Sec. III A, the low-energy DoS can be separated in two regimes of energy. Under the threshold $\Delta E = 1/2 - w$, every state is interpretable. Above the threshold, typical states quickly gain a high Hamming distance d . We indeed observe that, for $d \ll \mathcal{O}(N^2)$, the fraction of d -interpretable states $\tau_d(E)$ decays exponentially,

$$\tau_d(E) \sim \exp\left(-\frac{E}{E_c(d)}\right), \quad (7)$$

with a characteristic energy scale $E_c(d)$ above which states are unlikely to be better than d -interpretable. This observation is strikingly similar to the behavior of path graphs studied in Fig. 4. As displayed in Fig. 6 (c), $E_c(d)$ decreases toward zero with increasing problem size N . This is observed with a remarkably low variability between random graph instances. Similarly as for the path graph embedding, we conclude that states beyond the energy threshold $\Delta E = 1/2 - w$ require some post-processing, making them harder to interpret.

We now argue that this pattern is independent from the implementation details of the CL, and can pretend to some sort of universality. More specifically, we show that the distribution of d is reproduced by a simple gadget-less model of a product (disjoint unions) of path graphs. We compare the CL of initial size N with the graph product P_{4N}^N of N weighted copies of P_{4N} , both graphs having the same w . Using $N = 7$ and $w = 1/8$, the DoS of the path product shown in Fig. 6 (a, *top right panel*) is strikingly similar to the DoS of the CL (*top left panel*). Next, we consider a small $w = 1/20$. As in Fig. 4 (*top right panel*), in this regime, the DoS features peaks located at half-integer values of ΔE (in units of δ). Neglecting some gadget-induced defects with energy $\simeq 3/2$, there is a correspondence between energy and the number l of domain walls, i.e. $\Delta E \simeq l/2$ (see also App. B). The DoS can thus be analyzed as a collection of peaks located near $\Delta E \simeq l/2$. For $l \in \{1, 2, 3\}$, we plot in Fig. 6 (b) the normalized cumulative distribution of d within each of these peaks. With a low variability between the initial graphs (errors bars are smaller than the thickness of the line), we observe that the paths product model faithfully

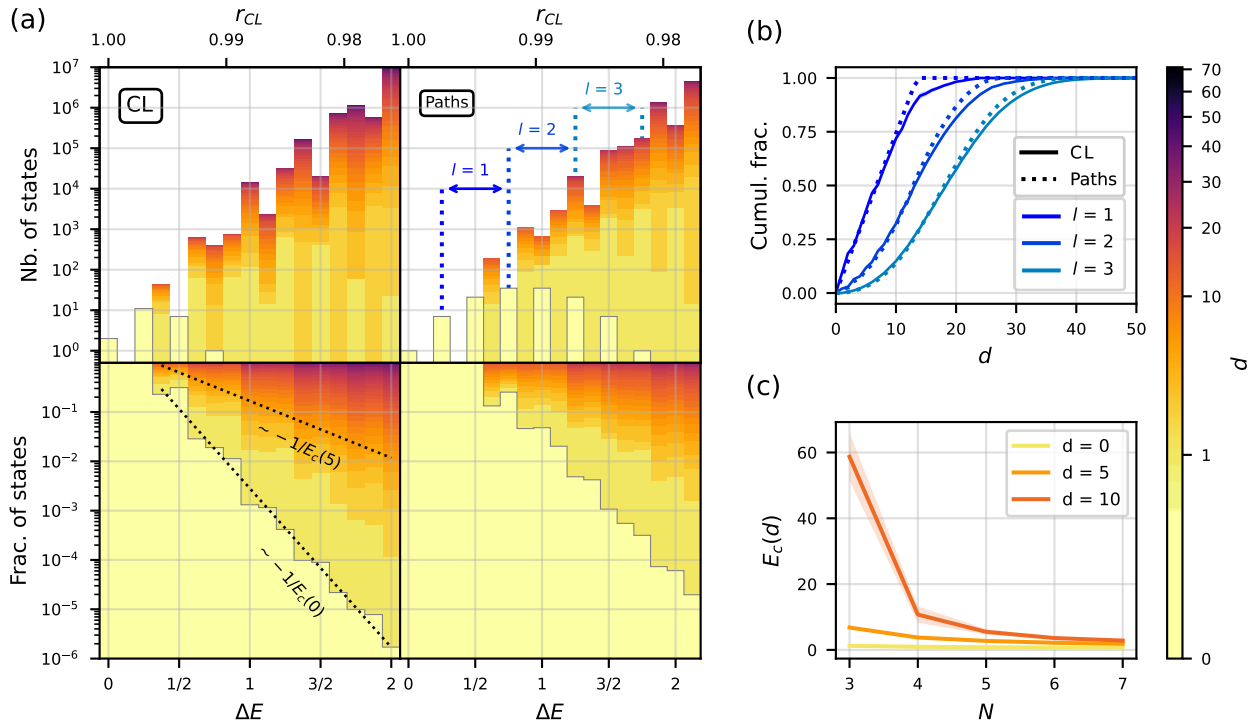


FIG. 6. (a) Low-energy sector DoS (*top*) and normalized cumulative DoS (*bottom*) for a random CL graph with initial size $N = 7$ (*left*) and for a path product P_{4N}^N (*right*), both with weight bias $w = 1/8$. Gray highlight: frontier of interpretable states. Dotted line: exponential decay $\sim \exp(-E/E_c(d))$ for $d \in \{0, 5\}$. (b) Cumulative normalized DoS restricted to energy intervals with $\simeq l \in \{1, 2, 3\}$ domain walls and $w = 1/20$, both for the CL and the path product. (c) Effect of graph size N on the characteristic energy $E_c(d)$ at which the fraction of states with distance at most d decays for the CL, with $w = 1/8$. Shaded areas indicate one standard deviation error bar.

reproduces the interpretability features of the CL. This supports the idea that the combinatorial explosion of hard-to-interpret states is an intrinsic feature of the CL scheme due to the introduction of ancilla nodes in the form of path graphs. We expect this phenomenon to be somewhat universal and thus to exist in other embedding methods, such as graph subdivisions [1, 29] (discussed in App. C) and parity encoding [30, 31], where gadgets and paths graphs introduce extensive amounts of ancillae.

From this analysis, it is natural to ask whether the choice of weight profile W_w may affect the exponential scaling Eq. (7). Indeed, the weight profile affects the distribution of defects and therefore their degeneracy. A study in this direction is presented in Appendix D, suggesting modifications to the weight profile favor defects and should be avoided.

IV. COMPARISON TO DESELECTION STRATEGY

We now compare the distance strategy with a second strategy, which is much cheaper to implement, but yields poorer quality results. Recall that a CL configuration can be seen as N paths that are either in an antiferromagnetic

ordering or have a domain-wall-like defect, as shown in Fig. 3 (a). Furthermore, a path in an antiferromagnetic ordering can be considered as selected or unselected depending on which of its two endpoints is selected. In that language, the *deselection* strategy consists in finding the defective paths and declare them unselected. Since a configuration is interpretable if and only if each path has no defect, this strategy guarantees to return an interpretable CL configuration, that is an IS of the original graph.

In terms of computational cost, finding all defects and deselection is easily done by walking through the CL and spotting succession of unselected vertices, which characterize defects. Thus, the deselection strategy costs $\mathcal{O}(N^2)$ to apply. Therefore, we consider that, unlike the distance strategy, it can be applied at any energy. The drawback of this technique is that the output configurations are of lesser quality than the distance strategy. To compare both strategies, we compute the approximation ratio r obtained after interpretation in the CL embedding of a random graph with $N = 7$ and MIS size $|\text{MIS}| = 3$. Then, for each technique and for two arbitrary thresholds $r_0 \in [0, 1]$, we show in Fig. 7 the probability that a configuration taken at random within the energy window $[E_{\text{MWIS}}, E_{\text{MWIS}} + \Delta E]$ yields $r \geq r_0$.




FIG. 7. Probabilities of obtaining upon interpretation an approximation ratio $r \geq r_0$ for two post-processing strategies, with $r_0 \in \{2/3, 1\}$. Dotted lines: deselection strategy, for which each chain with defects is considered as unselected. Plain lines: interpretation of the nearest interpretable state with respect to the Hamming distance d . Data for one random graph with initial size $N = 7$ and MIS size $|\text{MIS}| = 3$.

In the case several nearest interpretable states exist, with possibly different values of r , our definition of the distance strategy becomes ambiguous. Our convention in this situation is to take the lowest approximation ratio, so that our results still give a lower bound of the advantage of the distance strategy over the deselection strategy. We also checked these cases are rare (less than 2.2% of states displayed in Fig. 7), which legitimates our conclusions irrespective of the chosen convention.

When $\Delta E \geq 1/2$, we observe in Fig. 7 that the distance strategy tends to output better approximation ratios than the deselection strategy, and that this gap increases with the energy. This advantage is expected and we explain it as follows. First, at least in the low-energy spectrum under study, the distance strategy almost always leaves non-defective chains untouched. We checked numerically that CL vertices located in non-defective paths had to be flipped in only 0.8% of the configurations from Fig. 7. Consequently, for the vast majority of cases, the distance strategy outputs configurations of the initial graph with at least as many selected vertices than the deselection strategy, which translates into better approximation ratios for the distance strategy.

It is striking that the difference in approximation ratios for the two strategies is not very large (see also Fig. 10 in App. B3), given the huge difference in computational complexity. We explain such good results for the deselection strategy by the fact that it “corrects” some defects, while the others are energetically unfavored. Indeed, adding a domain wall to an unselected chain increases the energy by $1/2 - w$, versus $1/2 + w$ for a selected chain, but in both cases the chain is considered unselected after interpretation. This means the approximation ratio is reduced in the second case (harmful defect), but is unchanged in the first case (harmless defect), where the interpretation strategy

acts as an error correction. On top of that, the harmless defects are energetically favored compared to the harmful ones, which explains the good results of the deselection strategy at low energies, despite its crudeness. This mechanism also explains why w can, and in theory should, be chosen larger than $1/2$ when using interpretation methods, as detailed in Appendix B3.

Lower bound for r with the deselection strategy. We can find a lower bound for the approximation ratio obtained through the deselection strategy. Consider a configuration S_{CL} of the CL. Through the deselection strategy, defective chains are replaced by unselected chains, defining an interpretable CL configuration S_{des} . Since defects cost energy, deselecting defective chains lowers the energy: $E(S_{\text{des}}) \leq E(S_{\text{CL}})$.

By construction, the CL weight profile ensures that flipping any chain from selected to unselected (which is always possible without breaking the independence constraint) increases the energy by $2w$. In particular, starting from the MWIS with its $|\text{MIS}|$ selected chains, any interpretable configuration with N_{sel} selected chains has $|\text{MIS}| - N_{\text{sel}}$ less selected chains than the MWIS and thus has an energy $E_{\text{MWIS}} + 2w \cdot (|\text{MIS}| - N_{\text{sel}})$. Here we recall that $|\text{MIS}|$ is the MIS size of the initial graph. Since the approximation ratio upon interpretation is $r = N_{\text{sel}}/|\text{MIS}|$, the energy of any interpretable configuration is linearly related to its approximation ratio upon interpretation. This applies in particular to S_{des} :

$$r = 1 - \frac{E(S_{\text{des}}) - E_{\text{MWIS}}}{2w \cdot |\text{MIS}|}. \quad (8)$$

Finally, using $E(S_{\text{des}}) \leq E(S_{\text{CL}})$ we deduce a lower bound on the approximation ratio provided by the deselection strategy:

$$r \geq 1 - \frac{E(S_{\text{CL}}) - E_{\text{MWIS}}}{2w \cdot |\text{MIS}|}. \quad (9)$$

We point out that this bound cannot be improved since it is saturated whenever S_{CL} is interpretable.

V. DISCUSSION: ASYMPTOTIC BEHAVIOR

We now consider the asymptotic limit $N \rightarrow +\infty$. In this regime we expect the distance strategy to be hard to apply, as a result of App. E. In this section we link the previous results to complexity requirements on a quantum UD-MWIS solver in the case of the deselection strategy. As a first step, we prove that using the CL embedding scheme leads to a degradation of the approximation ratio.

A. Error amplification

From the previous section, a key quantity here is the energy $E(S_{\text{CL}})$ of a sampled CL configuration obtained by approximate UD-MWIS optimization. Let us assume that the quantum solver produces a number of defects (such as domain walls) that scales as N^γ for some constant $\gamma \geq 0$. For instance, under the Kibble-Zurek mechanism [3, 32–35], we would get a finite density of defects, hence $\gamma = 2$. Since the energy cost of each defect is $\Theta(1)$, this implies we get configurations of energy

$$E = E_{\text{MWIS}} + \Theta(N^\gamma) \quad (10)$$

Using Eq. (9)-(10), the deselection strategy leads to an approximation ratio lower bound

$$r \geq 1 - \frac{\Theta(N^\gamma)}{2w|\text{MIS}|} = 1 - \Theta(N^{\gamma-1}), \quad (11)$$

where we used, in the equality, the standard assumption $|\text{MIS}| = \Theta(N)$, which is verified in typical families of graphs of interest, such as bounded-degree graphs and some families of random graphs [36–38].

We distinguish between two regimes. When $\gamma < 1$, the deselection strategy ensures an approximation ratio asymptotically close to 1. When $\gamma > 1$, as expected in the current state of the art, the strategy asymptotically fails: the number of defects becomes infinitely larger than the number of chains, such that almost certainly every chain gets a defect and is deselected upon interpretation, yielding $r = 0$.

In the CL, let us denote by $|\text{MWIS}_{\text{CL}}|$ the MWIS total weight. The approximation ratio r_{CL} of an IS in the CL is related to its energy through

$$r_{\text{CL}} = 1 - \frac{E - E_{\text{MWIS}}}{|\text{MWIS}_{\text{CL}}|}. \quad (12)$$

With the deselection strategy, the approximation ratio after interpretation r , and before interpretation r_{CL} , are

thus related through Eq. (9) by

$$1 - r \leq \frac{(1 - r_{\text{CL}})|\text{MWIS}_{\text{CL}}|}{2w \cdot |\text{MIS}|} = \Theta(N(1 - r_{\text{CL}})), \quad (13)$$

where we used the fact that $|\text{MWIS}_{\text{CL}}| = \Theta(N^2)$. Importantly, Eq. (13) shows that solving the MWIS problem means solving the embedded MIS problem with an error up to N times higher (within a bounded factor). This amplification of the error exemplifies the limitation of embedding schemes caused by the interpretation step.

B. Approximability theory consequences

We now discuss the consequences of the error amplification in Eq. (13) due to the deselection strategy in terms of approximability theory. This theory classifies the hardness to find approximate solutions, namely the hardness of getting close to $r = 1$. We will focus on three classes of approximation: the fully-Polynomial-Time Approximation Scheme (fPTAS) class, which describes problems for which an approximation algorithm exists that guarantees an error at most ε in run time polynomial in *both* the instance size and $1/\varepsilon$; the Polynomial-Time Approximation Scheme (PTAS) class, which also ensures $r \geq 1 - \varepsilon$ but with a run time polynomial in the instance size only, and that can be worse than polynomial in $1/\varepsilon$; and the APX class, which only achieves $r \geq c$ for some constant $c < 1$. Note that the terms PTAS and fPTAS also denote the algorithms themselves (in addition to the approximation classes).

In terms of classical approximation algorithms, the UD-MWIS problem belongs to the PTAS class [9–13], while the general MIS problem is APX-complete [8, 15, 39]. MIS is thus believed to admit no PTAS. In fact, a corollary to the PCP theorem asserts that the MIS problem cannot be in the PTAS class (unless $\text{P}=\text{NP}$). Likewise, according to the quantum PCP conjecture [40], MIS should remain in the quantum equivalent of the APX class even with quantum algorithms.

Let us examine how these general statements relate to our findings. Classical PTASs exist for UD-MWIS, and due to the PCP theorem and quantum PCP conjecture, there cannot be any classical, and likely any quantum, PTAS for MIS. This means that either some degradation in the approximation ratio must happen during the interpretation, or the interpretation algorithm (or the embedding itself) is not polynomial in the problem size. The distance strategy falls in the second category (it is exponential), while the deselection strategy falls in the first category, as it is polynomial in the problem size.

For the deselection strategy, this is embodied by Eq. 13, showing that an error $\varepsilon_{\text{CL}} = \varepsilon/N$ on the UD-MWIS resolution translates to an error up to ε on the MIS problem. If an fPTAS were to exist for CL graphs, it would give an fPTAS for MIS, because

$\text{Poly}(N_{\text{CL}}, 1/\varepsilon_{\text{CL}}) = \text{Poly}(N, N/\varepsilon)$. However, this is unlikely due to the earlier PCP considerations, excluding the existence of a quantum or classical fPTAS for CL, and so for UD-MWIS as a whole. On the other hand, a PTAS for CL, which is not an fPTAS, would only provide a super-polynomial run time, e.g. $\exp(1/\varepsilon_{\text{CL}}) = \exp(N/\varepsilon)$, and therefore would not give a PTAS for MIS, in agreement with the PCP theorem. To sum up, PCP considerations rule out the existence of a PTAS for MIS and of an fPTAS for CL graphs, and so for UD-MWIS in general.

This discussion has even more practical consequences, i.e. in terms of number of defects. Going back to the assumption of Eq. (10), γ is related to the approximation ratio through Eq. (11). Consider then an hypothetical quantum solver for the CL, running polynomially in N and with a number of defects scaling as $\Theta(N^\gamma)$. With the deselection strategy, it would solve MIS with an error $\varepsilon = 1 - r \leq \Theta(N^{\gamma-1})$. If $\gamma \geq 1$, like in the standard Kibble-Zurek scenario where $\gamma = 2$, the explosion in the number of defects leads to approximation ratios so poor that no error $\varepsilon < 1$ can be guaranteed. If $\gamma < 1$, however, this algorithm would induce a PTAS for MIS, because the error can be made arbitrarily low by duplicating the initial MIS graph. Indeed, having M copies of the initial graph leads to a CL with $M \cdot N$ chains, guaranteeing an error $\leq \Theta((MN)^{\gamma-1})$ that vanishes at large M , but still solved in $\text{Poly}(MN) = \text{Poly}(N)$ time. Since this is ruled out by the previous discussion, we conclude that, very likely, there exists no quantum or classical CL solver running polynomially in N with a number of defects scaling as $o(N)$. This statement naturally extends to UD-MWIS in general.

In conclusion, no improvement in approximability class is likely to happen by using an embedding of MIS in a larger UD-MWIS problem. This represents a very practical embodiment of the PCP theorem and its quantum equivalent for embedding strategies. This of course does not prevent *heuristic* algorithms, namely algorithms with no run time guarantees, from achieving good enough approximation ratios.

VI. CONCLUSION AND OUTLOOK

In this work, we studied quantitatively the robustness to errors of embedding techniques to solve the MIS problem on Rydberg atom machines. We focused on the CL construction as a case study. We highlighted the difficulty of interpreting, or correcting, defective solutions

to the CL—which unavoidably results from realistic adiabatic procedures—into solutions to the original MIS problem, while retaining a high approximation ratio.

We studied two opposite strategies. The *distance* strategy produces high quality solutions in the presence of few defects, but is untractable with large number of defects. The *deselection* strategy, on the other hand, is straightforward to apply with any number of defects, at the expense of a worst approximation ratio. Nevertheless, asymptotically at large N , these strategies require an extremely small number of defects scaling as $\Theta(1)$ and $o(N)$ respectively, to be tractable and produce non-zero approximation ratios. With the deselection strategy, this is a consequence of the error amplification induced by the embedding, which is linked to an enlarged Hilbert space with $\Theta(N^2)$ qubits.

Our observations on the CL graphs seem independent from the connectivity between chains, suggesting some form of universality, when approximating embedding graphs. As a consequence, we expect similar results when extending our approach to integer factorization [18, 41], QUBO problems [18, 42], and related embedding techniques [1–4, 29–31]. We also expect the $o(1/N)$ requirement in the density of defects to apply to other strategies which output better approximation ratios than the deselection strategy.

Using approximability theory, we demonstrated that getting a $o(1/N)$ density of defects in polynomial time is unlikely to be feasible, due to the quantum PCP conjecture. Otherwise, the combination with e.g. the deselection strategy would produce an fPTAS for MIS, which is ruled out by this conjecture. Further work would be required to study more precisely how this limitation is embodied in adiabatic algorithms, in particular in terms of energy gaps. A more realistic goal could thus be to outperform existing approximate classical PTAS or heuristics for UD-MWIS.

ACKNOWLEDGMENTS

We thank F. Meneses, A. Oliveira and H. Pichler for useful exchanges. This work was supported by the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101018180 (HPC-QS). This publication has received funding under Horizon Europe programme HORIZON-CL4-2022-QUANTUM-02-SGA via the project 101113690 (PASQuanS2.1).

[1] H. Pichler, S.-T. Wang, L. Zhou, S. Choi, and M. D. Lukin, arXiv preprint arXiv:1809.04954 (2018).
 [2] H. Pichler, S.-T. Wang, L. Zhou, S. Choi, and M. D. Lukin, arXiv preprint arXiv:1808.10816 (2018).

[3] S. Ebadi, A. Keesling, M. Cain, T. T. Wang, H. Levine, D. Bluvstein, G. Semeghini, A. Omran, J.-G. Liu, R. Samajdar, X.-Z. Luo, B. Nash, X. Gao, B. Barak, E. Farhi, S. Sachdev, N. Gemelke, L. Zhou, S. Choi,

- H. Pichler, S.-T. Wang, M. Greiner, V. Vuletić, and M. D. Lukin, *Science* **376**, 1209 (2022), <https://www.science.org/doi/pdf/10.1126/science.abo6587>.
- [4] M. Kim, K. Kim, J. Hwang, E.-G. Moon, and J. Ahn, *Nature Physics* **18**, 755 (2022).
- [5] A. G. de Oliveira, E. Diamond-Hitchcock, D. M. Walker, M. T. Wells-Pestell, G. Pelegrí, C. J. Picken, G. P. A. Malcolm, A. J. Daley, J. Bass, and J. D. Pritchard, *Demonstration of weighted graph optimization on a rydberg atom array using local light-shifts* (2024), [arXiv:2404.02658 \[quant-ph\]](https://arxiv.org/abs/2404.02658).
- [6] P. Cazals, A. François, L. Henriët, L. Leclerc, M. Marin, Y. Naghmouchi, W. d. S. Coelho, F. Sikora, V. Vitale, R. Watrigant, *et al.*, *arXiv preprint arXiv:2502.04291* (2025).
- [7] R. M. Karp, *Complexity of computer computations* **1**, 85 (1972).
- [8] M. R. Garey and D. S. Johnson, *J. ACM* **25**, 499–508 (1978).
- [9] T. Erlebach, K. Jansen, and E. Seidel, in *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '01* (Society for Industrial and Applied Mathematics, USA, 2001) p. 671–679.
- [10] T. Matsui, in *Japanese Conference on Discrete and Computational Geometry* (Springer, 1998) pp. 194–200.
- [11] G. K. Das, M. De, S. Kolay, S. C. Nandy, and S. Sur-Kolay, *Information Processing Letters* **115**, 439 (2015).
- [12] S. C. Nandy, S. Pandit, and S. Roy, *Information Processing Letters* **127**, 58 (2017).
- [13] M. Bonamy, E. Bonnet, N. Bousquet, P. Charbit, and S. Thomassé, in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, 2018) pp. 568–579.
- [14] M. F. Serret, B. Marchand, and T. Ayril, *Physical Review A* **102**, 052617 (2020).
- [15] J. Håstad, *Acta Mathematica* **182**, 105 (1999).
- [16] D. Zuckerman, in *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2006) p. 681–690.
- [17] C. Dalyac, *Quantum many-body dynamics for combinatorial optimisation and machine learning*, Ph.D. thesis, Sorbonne Université (2023).
- [18] M.-T. Nguyen, J.-G. Liu, J. Wurtz, M. D. Lukin, S.-T. Wang, and H. Pichler, *PRX Quantum* **4**, 010316 (2023).
- [19] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, *Science* **292**, 472 (2001).
- [20] T. Albash and D. A. Lidar, *Rev. Mod. Phys.* **90**, 015002 (2018).
- [21] B. F. Schiffer, D. S. Wild, N. Maskara, M. Cain, M. D. Lukin, and R. Samajdar, *Phys. Rev. Res.* **6**, 013271 (2024).
- [22] A. Miessen, D. J. Egger, I. Tavernelli, and G. Mazzola, *PRX Quantum* **5**, 040320 (2024).
- [23] L. Bombieri, Z. Zeng, R. Tricarico, R. Lin, S. Notarnicola, M. Cain, M. D. Lukin, and H. Pichler, *arXiv preprint arXiv:2411.04645* (2024).
- [24] E. Urban, T. A. Johnson, T. Henage, L. Isenhower, D. Yavuz, T. Walker, and M. Saffman, *Nature Physics* **5**, 110 (2009).
- [25] J.-G. Liu, X. Gao, M. Cain, M. D. Lukin, and S.-T. Wang, *SIAM Journal on Scientific Computing* **45**, A1239–A1270 (2023).
- [26] V. E. Levit and E. Mandrescu, in *Discrete Mathematics and Theoretical Computer Science*, edited by C. S. Calude, M. J. Dinneen, and V. Vajnovszki (Springer Berlin Heidelberg, Berlin, Heidelberg, 2003) pp. 237–256.
- [27] J. L. Arocha, *Revista Ciencias Matemáticas* **5**, 103 (1984).
- [28] E. N. Gilbert, *The Annals of Mathematical Statistics* **30**, 1141 (1959).
- [29] C. Dalyac, L.-P. Henry, M. Kim, J. Ahn, and L. Henriët, *Phys. Rev. A* **108**, 052423 (2023).
- [30] W. Lechner, P. Hauke, and P. Zoller, *Science advances* **1**, e1500838 (2015).
- [31] M. Lanthaler, C. Dłaska, K. Ender, and W. Lechner, *Physical Review Letters* **130**, 220601 (2023).
- [32] T. W. Kibble, *Journal of Physics A: Mathematical and General* **9**, 1387 (1976).
- [33] W. H. Zurek, *Nature* **317**, 505 (1985).
- [34] W. H. Zurek, *Physics Reports* **276**, 177 (1996).
- [35] Y. M. Bunkov and H. Godfrin, *Topological defects and the non-equilibrium dynamics of symmetry breaking phase transitions*, Vol. 549 (Springer Science & Business Media, 2000).
- [36] Y. Caro, *New results on the independence number*, Tech. Rep. (Technical Report, Tel-Aviv University, 1979).
- [37] V. K. Wei, A lower bound on the stability number of a simple graph (1981).
- [38] A. Coja-Oghlan and C. Efthymiou, in *Society for Industrial and Applied Mathematics and Association for Computing Machinery. Proceeding of the ACM-SIAM Symposium on Discrete Algorithms* (Society for Industrial and Applied Mathematics, 2011) p. 136.
- [39] C. Bazgan, B. Escoffier, and V. T. Paschos, *Theoretical Computer Science* **339**, 272 (2005).
- [40] D. Aharonov, I. Arad, and T. Vidick, *The quantum pcp conjecture* (2013), [arXiv:1309.7495 \[quant-ph\]](https://arxiv.org/abs/1309.7495).
- [41] J. Park, S. Jeong, M. Kim, K. Kim, A. Byun, L. Vignoli, L.-P. Henry, L. Henriët, and J. Ahn, *Physical Review Research* **6**, 023241 (2024).
- [42] A. Byun, J. Jung, K. Kim, M. Kim, S. Jeong, H. Jeong, and J. Ahn, *Advanced Quantum Technologies* **7**, 2300398 (2024).

Appendix A: DoS in paths graphs

This section introduces tools from graph theory that provide a concise notation for the DoS in MWIS problems. In App. A 2, this is used to characterize the DoS for path graphs. The following sections expand on this formalism to treat the case of weighted paths and CL graphs.

1. Independence polynomials

To study the non-interpretable states that arise from suboptimal UD-MWIS optimization, it is convenient to borrow several notions from the fields of combinatorics and graph theory. The *independence polynomial* $\text{IP}(G, W; x)$ is a generalized polynomial, whose coefficient at any (possibly non-integer) degree $w \in \mathbb{R}_+$ is the number of ISs whose vertices have a total weight w :

$$\text{IP}(G, W; x) = \sum_{S \in \text{IS}(G)} x^{\sum_{i \in S} W(i)} = \text{Tr}(x^{-\mathcal{H}_{\text{MWIS}}}). \quad (\text{A1})$$

As a reminder, the energy scale $\delta = 1$ is omitted. We will frequently drop the notation of the weight profile W in $\text{IP}(G, W)$ if the graph is unweighted, as well as the notation of the indeterminate x . One can relate the independence polynomial with the canonical partition function $\mathcal{Z}(\beta)$, introducing a dimensioning factor β which represents the thermodynamical inverse temperature:

$$\mathcal{Z}(\beta) = \text{Tr}(e^{-\beta \mathcal{H}_{\text{MWIS}}}) = \text{IP}(G, W; x = e^{-\beta}). \quad (\text{A2})$$

Conversely, the independence polynomial characterizes the DoS ρ seen as a distribution.

The partition function depends on the DoS as

$$\mathcal{Z}(\beta) = \int dE e^{-\beta E} \rho(E). \quad (\text{A3})$$

Inverting this relation and replacing $\mathcal{Z}(\beta)$ by $\text{IP}(z)$ with the change of variable $z = e^{-\beta}$, we get

$$\rho(E) = \int_{-i\infty}^{+i\infty} \frac{d\beta}{2i\pi} \mathcal{Z}(\beta) e^{\beta E} = \frac{1}{2i\pi} \oint_{\mathcal{C}} \text{IP}(z) z^{E-1} dz. \quad (\text{A4})$$

where \mathcal{C} is an anticlockwise contour in the complex plane around the origin.

Both formulations being equivalent, we choose to work with independence polynomials due to the simplicity of notations.

2. Unweighted paths

Let P_N be the path graph with $N \in \mathbb{N}$ vertices. A previous result from Refs [26, 27] is that

$$\forall N \geq 0, \quad \text{IP}(P_N) = \sum_{k \geq 0} \binom{N-k+1}{k} x^k, \quad (\text{A5})$$

where the summand is zero whenever $k > \lfloor (N+1)/2 \rfloor$. In particular the two first path polynomials are obtained from a combinatorial enumeration:

$$\text{IP}(P_0) = 1, \quad \text{IP}(P_1) = 1 + x, \quad (\text{A6})$$

which is quickly verified by a combinatorial enumeration of the ISs in P_0 and P_1 . For the sake of completeness, we mention that Eq. (A5) can be proved by induction with a well-known *Fibonacci recurrence relation* from Ref. [27]:

$$\forall N \geq 1, \quad \text{IP}(P_{N+1}) = \text{IP}(P_N) + x \cdot \text{IP}(P_{N-1}). \quad (\text{A7})$$

Moreover, we provide an apparently new derivation of Eq. (A5) based on combinatorics. Given N, k , an IS in P_N with k selected vertices can be built by choosing how many unselected vertices are on the right and left of each selected vertex. This amounts to choosing two non-negative integers l, r and $(k - 1)$ positive integers m_1, \dots, m_{k-1} such that

$$l + m_1 + \dots + m_{k-1} + r = N - k. \quad (\text{A8})$$

Here $(l, m_1, \dots, m_{k-1}, r)$ characterizes the IS in P_N made by the succession of $l \geq 0$ unselected vertices, then one selected vertex, then $m_1 \geq 1$ unselected vertices (ensuring the independence constraint), and so on. This clearly defines a one-to-one correspondence between the elements of IS (P_N) and the possible sequences $(l, m_1, \dots, m_{k-1}, r)$. By application of the so-called *stars and bars theorem* from combinatorics, one retrieves the number $[x^k] \text{IP}(P_N)$ of possible choices for $(l, m_1, \dots, m_{k-1}, r)$ with the same result as in Eq. (A5).

3. Weighted paths

Let $w \in (0, 1/2)$ be a weight bias. We endow the path P_{2N} with the weight profile W_w as defined in Eq. (4). By disjunction over which vertices in $\{v_1, v_{2N}\}$ are selected, all ISs can be partitioned in four categories. For instance, if v_1 is selected and v_{2N} is not, then this lets $2N - 3$ remaining vertices that can take any independent configuration from the path P_{2N-3} . Since v_1 contributes to a weight $1/2 + w$, the polynomial $x^{\frac{1}{2}+w} \text{IP}(P_{2N-3})$ characterizes the DoS of this category of ISs. The three other categories can be treated likewise to obtain the independence polynomial for the weighted path P_{2N} :

$$\text{IP}(P_{2N}, W_w) = \text{IP}(P_{2N-2}) + x^{\frac{1}{2}} (x^w + x^{-w}) \text{IP}(P_{2N-3}) + x \text{IP}(P_{2N-4}). \quad (\text{A9})$$

Again, this is another way to write the DoS of the corresponding Hamiltonian $\mathcal{H}_{\text{MWIS}}$. The highest-degree terms count in this regard $\sim N$ non-interpretable configurations located at an energy of only $1/2 - w$. This proves that the manifold of lowly-excited ISs in P_{2N} is filled with non-interpretable configurations.

Appendix B: Crossing lattice

This section is devoted to some formal aspects of the CL that are discussed in our main text. First, we introduce several definitions and enumerate a family of states in the CL, which serves as a lower bound for its DoS. We conclude by a discussion of the numerical parameters used in our study.

1. Notations

Let G be an unweighted graph with $N \gg 1$ vertices. We consider the CL embedding of G for an arbitrary ordering of the vertices in G . Our goal is not to formalize its construction, but to introduce the vocabulary required in our derivations. Given a path $P_{2N} = \{v_1, \dots, v_{2N}\}$ with edges between v_i and v_{i+1} , the vertices $v_1, v_3, \dots, v_{2N-1}$ are said to be *odd vertices* while v_2, v_4, \dots, v_{2N} are called *even vertices*. We use the ordering convention that if P_{2N} is weighted by W_w , then v_1 is always the vertex with weight $1/2 + w$, in accordance with the main text.

Let us consider the path product P_{2N+2}^N , in which each path has the weight profile W_w for some w and is called a *chain*. As shown in Fig. 8 for the case $N = 3$, this graph can be drawn in 2D by arranging the N paths along the lines of a grid lattice, following the construction of Ref. [18]. Instead of introducing UD gadgets to replicate the edge structure from G , one can directly insert copies of the edges in G . We call the resulting graph a *non-planar CL embedding graph* of G and denote it by G_{np} . By identification of P_{2N+2}^N as a subgraph of G_{np} , the N chains are still defined in G_{np} . Clearly, this construction of G_{np} satisfies all the requirements for an embedding with the usual interpretable states defined by having each chain in an antiferromagnetic state with $N + 1$ selected vertices.

Given some configuration $S_{np} \in \text{IS}(G_{np})$, we denote by \mathcal{C}_i the subset of vertices in S_{np} that belong to the i -th chain in G_{np} . Here the ordering of the chains is the same ordering than for the vertices $\{v_i, 1 \leq i \leq N\}$ of G they embed. Given some \mathcal{C}_i in a chain that embeds a vertex $v_i \in G$, we say that this chain is *selected* if \mathcal{C}_i is a subset of the MWIS in G_{np} , namely if the only selected vertices are odd. Because the chain itself is a copy of P_{2N+2} with weight profile




FIG. 8. *a)* The non-planar CL is a weighted path product P_{2N+2}^N . *b)* Upon insertion of new edges, the resulting *non-planar CL embedding graph* satisfies the criteria for an embedding graph except for the planarity. Upon insertion of suitable gadgets, one retrieves a CL embedding graph.

W_w , this MWIS is indeed unique. Conversely, if \mathcal{C}_i only has even selected vertices, then it is said to be *unselected*. Finally, we introduce two independence polynomials

$$A_{N,\pm} = \left(1 + x^{\frac{1}{2} \pm w}\right) (1 + x)^N. \quad (\text{B1})$$

The combinatorial interpretation of these polynomials is that $A_{N,+}$ ($A_{N,-}$) describes all possible subsets of the (un)selected configuration in a weighted chain. Note that except for the empty chain which, by definition, is both selected and unselected, this enumerates all configuration where all vertices have the same parity.

2. Lower bound on the DoS of CL graphs

Due to the dependence of G_{np} in G , and to the hardness of computing the independence polynomials for grid graphs, we are only interested in approximating $\text{IP}(G_{np})$. This can be done by enumerating as many ISs in G_{np} with their respective weight, while ensuring no IS is counted twice and each IS is, indeed, an independent set. Intuitively, the family of ISs that we will build in this section is made by considering interpretable states and allowing in each chain different defects, depending on whether the chain embeds an initial vertex that is selected. This process will ensure that all built configurations remain independent, in spite of the gadgets connecting the chains.

Let $S \in \text{IS}(G)$, and assume $v_i \neq v_j$ are two vertices in S . If v_i and v_j are not adjacent in G , then there is in G_{np} no edge between the chain that embeds v_i and the chain that embeds v_j . Thus both chains can be in any of the ISs of P_{2N+2} , without violating the independence criterion. Conversely, let us assume that v_i and v_j are adjacent in G . Then one of the odd vertices in \mathcal{C}_i is adjacent in G_{np} to one of the odd vertices in \mathcal{C}_j . In that case if \mathcal{C}_i is in a configuration where only even vertices are selected, and \mathcal{C}_j itself is in an IS of P_{2N+2} , then there is no independence violation. This can be generalized to the N vertices of G as follows. Let $k = |S|$ be the number of selected vertices, and i_1, \dots, i_k the indices of the vertices in G that are in S . By our previous argument, if each chain $\mathcal{C}_{i_1}, \dots, \mathcal{C}_{i_k}$ is in an IS of P_{2N+2} , and every other chain is unselected, then this defines an IS S_{np} in G_{np} . For a given S , all the ISs that are built by this process define a DoS whose corresponding independence polynomial is $\text{IP}(P_{2N+2}, W_w)^k A_{N,-}^{N-k}$.

As stated above, this DoS depends on one specific $S \in \text{IS}(G)$. We are now interested in obtaining a DoS which corresponds to all the possible $S \in \text{IS}(G)$, and computing its independence polynomial. This is done by adding all the contributions from each possible S , with the only precaution that each configuration can be enumerated at most once. To enforce this condition, let us consider that if $v_i \in S$, then \mathcal{C}_i can take any configuration but cannot be unselected. Intuitively, this allows to invert our construction process by retrieving $S \in \text{IS}(G)$ from the chains that have at least one odd selected vertex. By summation over all possible $S \in \text{IS}(G)$, our enumeration describes a family of ISs in G_{np} whose DoS is represented by the polynomial

$$\text{IP}(G_{np})_{lb} = \sum_{S \in \text{IS}(G)} (\text{IP}(P_{2N+2}, W_w) - A_{N,-})^{|S|} A_{N,-}^{N-|S|} = A_{N,-}^N \text{IP}\left(G; x = \frac{\text{IP}(P_{2N+2}, W_w)}{A_{N,-}} - 1\right), \quad (\text{B2})$$

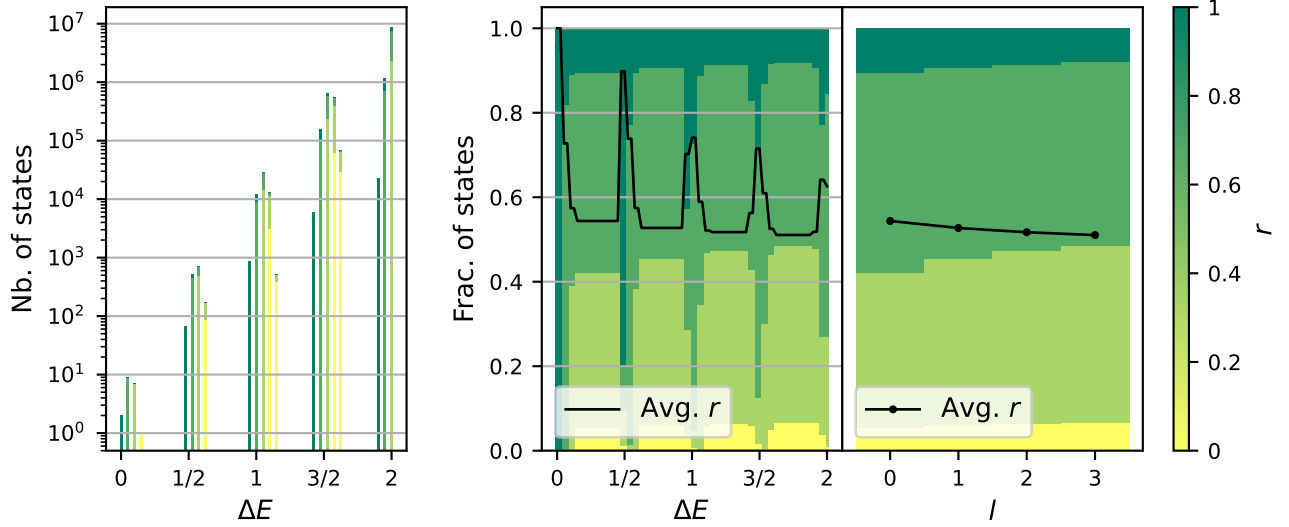


FIG. 9. approximation ratio r after interpretation in a random CL with CL graph with initial size $N = 7$ and weight bias $w = 1/20$. Here the approximation ratio is obtained with the distance strategy. (Left) Low-energy DoS. (Middle) Repartition of r in increasing energy windows $[E_{\text{MWIS}}, E]$ and for the same parameters. (Right) Repartition of r in energy windows defined by $|\Delta E - l/2| \leq 1/4$ for $l = 0, \dots, 3$.

where the subscript lb stands for *lower bound*, in the following sense:

$$\forall w \in \mathbb{R}_+, \quad [x^w] \text{IP}(G_{np})_{lb} \leq [x^w] \text{IP}(G_{np}). \quad (\text{B3})$$

The previous lower bound can be generalized to the CL. As pictured in Fig. 8, the CL is equivalent to the non-planar CL upon insertion of two types of gadgets, so it remains to lower bound the DoS of possible configurations for each gadget given the configuration of the two chains around it. In both types of gadgets, this leaves the possibility to select at most one vertex in the central complete subgraph in the case of interpretable states, with weight 1 (crossing-with-edge) or 2 (crossing). If $|E|$ is the number of edges in the initial graph G , then this yields a lower bound $\text{IP}(G_{\text{CL}})_{lb}$ for the DoS in the CL as follows:

$$\text{IP}(G_{\text{CL}})_{lb} = (1+x)^{|E|} \times (1+x^2)^{N(N-1)/2-|E|} \times \text{IP}(G_{np})_{lb}, \quad (\text{B4})$$

with as before

$$\forall w \in \mathbb{R}_+, \quad [x^w] \text{IP}(G_{\text{CL}})_{lb} \leq [x^w] \text{IP}(G_{\text{CL}}). \quad (\text{B5})$$

More gadget-induced defects can be taken into account by considering separately the cases where zero, one or two boundary vertices of a gadget are selected. The opposite stance can also be taken by first studying the product of the $N(N-1)/2$ gadgets, and then retrieving the CL by inserting small path subgraphs. However these approaches are less tractable and typically require more knowledge of G like its adjacency matrix.

3. Approximation ratio

In the scope of the CL embedding, the relevant approximation ratio is the one obtained after interpretation, r . For an interpretable configuration S_{CL} , it is always $r(f^{-1}(S_{\text{CL}}))$. For non-interpretable configurations, it depends on the interpretation strategy. Using the distance strategy, it would be the approximation ratio of the nearest interpretable configurations. In the (rare) case of several equidistant configurations, we chose the approximation ratio of the last one with respect to the order returned by the solver. Note that Sec. IV uses a different convention (choosing the smallest approximation ratio), and justifies that the exact convention should not change the results.

In Fig. 9 (left panel), we plot the low-energy DoS obtained with the same CL graph than in Fig. 6 (a, left column), with $w = 1/20$. Here the color indicates the approximation ratio r produced after applying the distance strategy. In

this regime where $w \ll 1$, we observe that the approximation ratio quickly decreases towards a value which is the average r over all ISs of the initial graph, with sharp increases at energies allowing new defects. This is explained by the fact that the CL DoS embeds a copy of the initial DoS at energies arbitrarily close to the ground-state energy when $w \rightarrow 0^+$. The general trend observed as a function of the typical number l of domain walls is shown in Fig. 9 (*middle and right panel*). It indicates a slight decrease of the approximation ratio when l increases, and this behavior is consistent across all other random graphs (not shown). This raises the question of whether different gadgets would lead to an increase of the variation of r with respect to l , a question we leave for further work.

In addition to studying the approximation ratio for the distance strategy, we compare in Fig. 10 the evolution with w of the approximation ratio distribution, for both strategies. This data is for the CL of the same random graph used in Sec. IV with the same parameters also, except for the values of w , with $1/w \in \{1, 3, 5, 14\}$. The black lines correspond to the average approximation ratio obtained when sampling at random a configuration in an energy window $[E_{\text{MWIS}}, E_{\text{MWIS}} + \Delta E]$.

Here, the value $w = 1$ is particularly interesting, since in the literature w is typically restricted to $w < 1/2$ [18, 23]. This constraint ensures that the weights defined in Eq. (4) are all positive, and thus that any defect in the antiferromagnetic ordering has a positive energy cost. With $w \geq 1/2$, some defective configurations become energetically favored compared to interpretable configurations, so that the CL ground state (and the low energy states) is in general not interpretable. Nevertheless, an interpretation strategy can still correct them and provide an approximation ratio $r = 1$. This is what we observe in Fig. 10, where states with $\Delta E \lesssim 1/2 - w$ are still interpreted correctly by both strategies.

For the deselection strategy, this can be explained as follows. Suppose one starts from an interpretable CL configuration and add defects at random. If only domain walls are considered, adding a defect to a selected chain (in the sense of App. E, that is when the chain is antiferromagnetically ordered and its endpoint with weight $1/2 + w$ is selected) costs an energy $1/2 + w$, while adding a defect to an unselected chain costs an energy $1/2 - w$. By construction, for $0 \leq w \leq 1/2$, both cost energy, so that interpretable states are favored. But for $w > 1/2$, adding a defect to an unselected chain is energetically favorable, while adding a defect to a selected chain is energetically prohibitive. In this w regime, low energy states therefore contain defects in chains that would otherwise be unselected. The strategy of deselecting defective chains then allows to recover good solutions to the original MIS problem despite the presence of potentially many defects. Finally, we point out that this argument also holds for gadget-induced defects with a higher energy cost (like $3/2 - w$).

4. Effect of the parameter w

Our computations use random unweighted graphs G generated with $N \in \{3, \dots, 7\}$ vertices and with an edge probability $p = 1/2$. The ordering of the vertices, on which their CL embedding graph depends, is thus chosen at random. These graphs have in average 5 (for $N = 3$) to 22 ($N = 7$) ISs, and a MIS size $|\text{MIS}|$ going from 1 to 4. In each resulting CL graph, the embedded DoS thus spans a quite restricted energy window with $|\text{MIS}| + 1$ peaks that are evenly spread across $\Delta E \in [0, 2w|\text{MIS}|]$. An interesting consequence is that some visual features of the DoS that depend on $w \cdot |\text{MIS}|$ can be deduced for large graphs (high $|\text{MIS}|$) from small graphs, simply by tuning w . In the main text, the values of $1/w$ are taken in [3, 20]. For several values this implies that the upper bound of this reduced energy window is always less than 1 and thus that all the interpretable states are energetically separated from states with domain walls, like in Fig. 6.

The regime of particularly small weight bias is illustrated in Fig. 11 and is the basis of our analysis with respect to the domain wall number l in the main text. Recall that in the case of path graphs, under our weighting convention, a domain wall always induces a weight reduction by $1/2 \pm w$. This remains true in the CL graph, except for some gadget induced-defects. Indeed, due to the presence of vertices with weight 2 in the non-crossing gadget shown on Fig. 8 b), some defects are associated with a weight decrease $3/2 \pm w$. Since $3/2$ is itself a multiple of $1/2$, any configuration has an energy which is an integral multiple of $1/2$ in the limit $w \rightarrow 0^+$.

To sum up, at small values of w , the interpretable states are located in a narrow energy window above the MWIS. The width of this embedded DoS scales linearly with N for a broad class of random graphs [38], but finding its exact width (or equivalently the MWIS weight in the initial problem) is itself an NP-complete problem [7]. The regime of small w has the advantage of increasing the fraction $\tau_0(E)$ for small enough E , meaning that annealing outputs are more likely to be interpretable. The downfall is that interpretable states with a poor approximation ratio are also more likely to be found, decreasing the quality of the output after interpretation.

On contrary, large values of w penalize such sub-optimal interpretable solutions by endowing them with a high

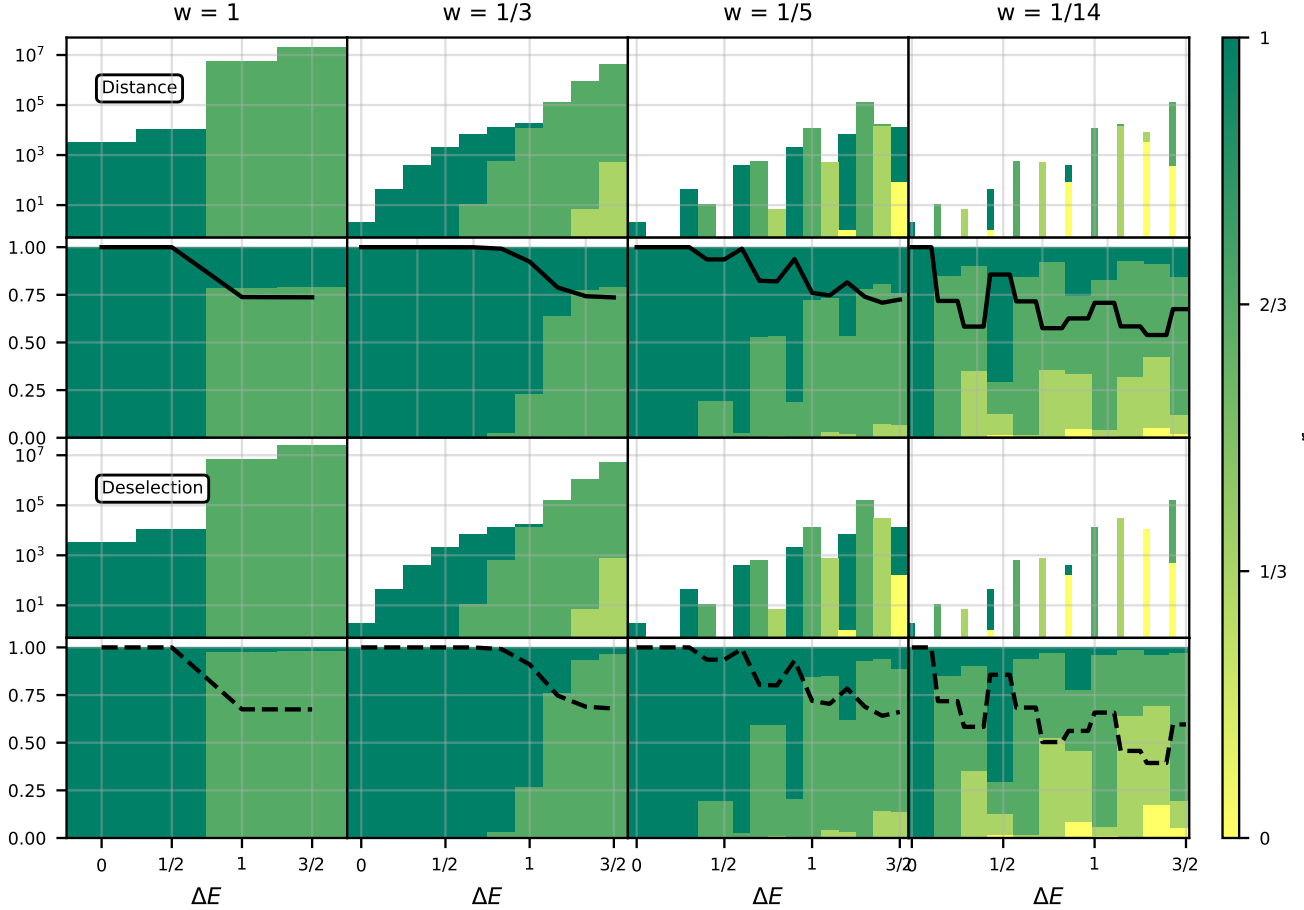


FIG. 10. Distribution of the approximation ratio obtained with the distance strategy (*first two rows*) and with the deselection strategy (*last two rows*) as described in Sec. IV and with the same graph used for Fig. 7. As before, the distribution for each value of $1/w \in \{1, 3, 5, 14\}$ is shown with the actual DoS (first and third rows) and after cumulating the DoS in increasing energy windows $[E_{\text{MWIS}}, E]$ (second and last rows). Black lines: average r obtained when sampling at random a configuration in the increasing energy window.

energy. This different regime, however, is plagued by the small energy $1/2 - w$ of defects, making non-interpretable configurations more likely to be found even at about $\simeq 1$ above the groundstate energy.

5. Example of high- d configuration

In Fig. 6, the highest distance d that is achieved (in the probed energy range) is $d = 70$. We illustrate in Fig. 12 a configuration that achieves this distance (*left pannel*) and highlight a possible choice of d bitflips to interpret it (*right panel*). Despite the hardness to interpret it, this configuration is located at a relatively low energy E such that $\Delta E \simeq 2$.

Appendix C: Subdivision embedding

1. 2D and 3D graph subdivisions

We discuss in this section a different kind of embedding than the CL scheme. Given a graph, an *even-edge subdivision* is the operation of inserting an even number of vertices v_1, \dots, v_{2M} in a given edge. The new edges

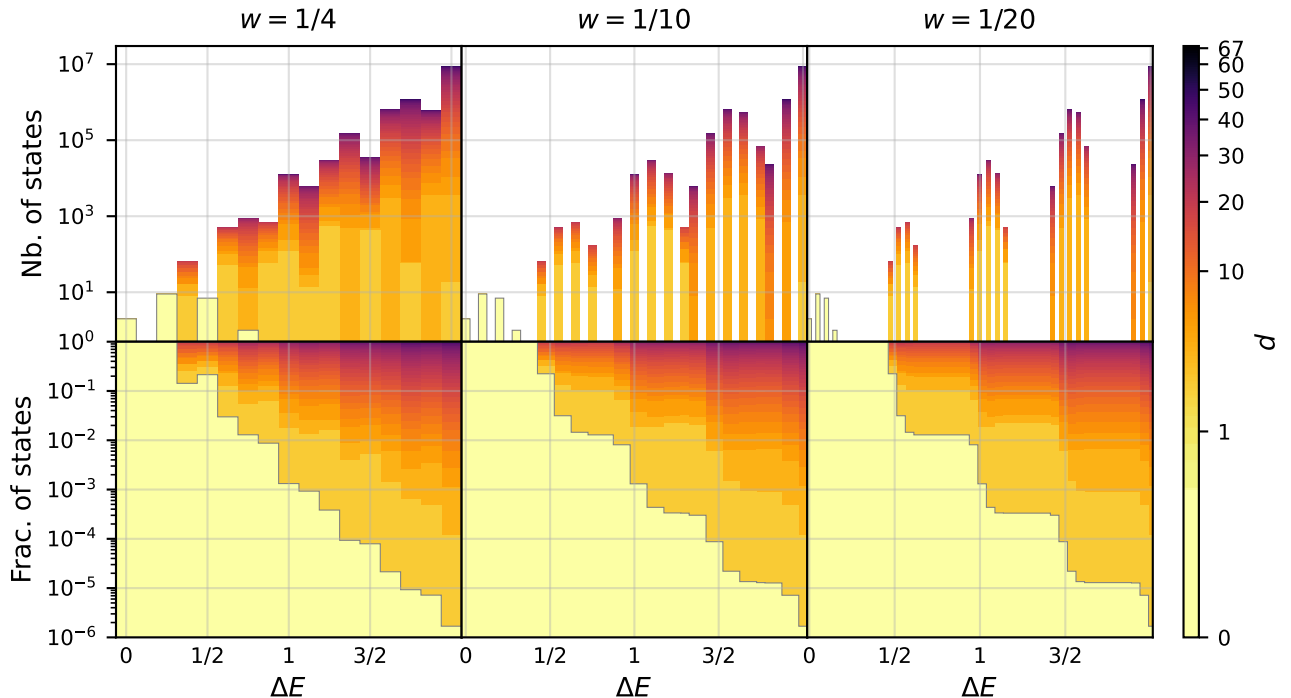


FIG. 11. Effect of the weight bias w for three values of $1/w \in [4, 10, 20]$ on a random CL graph with initial size $N = 6$. (*Top*) DoS of the CL graph with the usual weight profile W_w . When $w \rightarrow 0$, interpretable states corresponding to sub-optimal solutions in the initial graph are not energetically penalized. As a result, the energy of any configuration is given by the energy of its defects, which is approximately an integral multiple of $1/2$. (*Bottom*) Corresponding distribution of the fraction of d -interpretable states after cumulating with respect to E , like in Fig. 6.

are $(v_l, v_1), (v_1, v_2), \dots, (v_{2M}, v_r)$, so that the initial edge has been replaced by an even-length path subgraph.

Even-edge subdivisions have been used to embed non-planar or non-UD graphs into UD graphs [1–3]. This techniques can be generalized in 3D where the analogous to UD graphs is the class of *unit-ball* graphs [4, 17] (for which efficient approximation schemes also exist [13]).

Because even-length path subgraphs are used here to embed the state of a pair (v_l, v_r) of adjacent vertices in G , the two antiferromagnetic configurations are not in one-to-one correspondence with the three possible configurations for (v_l, v_r) . The generally used convention is thus that the antiferromagnetic states of the inserted path only encode the case when one vertex at the endpoints is selected. When the path has one domain wall, the configuration can be considered as frustrated and discarded [4], or both endpoints can be considered to be unselected [1]. In the following we adopt an intermediate stance where interpretation is done by contracting the subdivided edges to retrieve the embedded graph. Independence violations are discarded, typically by deselecting the vertices in violation. Note that in any case, the formalization of this embedding relies on a multi-valued embedding map f with respect to the criteria discussed in the main text, whence a larger domain of definition for the interpretation map f^{-1} .

2. Limitations of interpretability and weight profile

Here, we point out that subdivision embedding schemes are prone to independence violation upon interpretation in unweighted subdivided graphs. An example of this issue is given in Fig. 13. The instance of initial graph (*top*) has the particularity of having multiple odd-length cycle subgraphs. (The fact that it is UD and connected bears no importance here.) By performing arbitrary even-edge subdivisions, one obtains an embedding graph where the number of odd-length cycle subgraphs remains identical (*bottom*). However, odd-length cycles have a highly degenerate MIS akin to the one in unweighted path graphs P_{2N} , with at most one domain wall. Our example builds upon this property so that upon interpretation, subdivided edges can be interpreted into initial edges where both endpoints are selected, namely they lead to an independence violation that must be discarded. Furthermore, if an MIS of this embedding




FIG. 12. (Left) Non-interpretable configuration taken from the configurations shown in Fig. 6, here with the distance $d = 70$. Four domain walls are highlighted in light red, in good agreement with the energy $\Delta E \simeq 2$ of this configuration. (Right) Example of 70 vertices that should be flipped to reach the closest interpretable configuration.




FIG. 13. Initial graph (b) and embedding graph (a) in a pathological example of even-edge subdivision. If left unweighted, the embedding graph has a highly degenerate MIS. Upon interpretation by edge contraction (b), the average number of independence violations (red edges) scales linearly with the vertex count of the initial size.

graph is taken at random, then each cycle has a non-zero probability (here $1/2$) to have a domain wall and the average number of violations scales linearly with the graph size. Note that the initial graph is connected, showing that odd-length cycles can give raise to independence violation even when embedded in more complex graph. The initial graph is also UD and with small maximum degree, for the clarity of representation.

This example generally proves that graphs with odd-length cycle subgraphs are not good candidates for unweighted even-edge subdivisions. In theory, this issue can be avoided by adopting a weight profile like the W_w from path graph that energetically penalizes domain walls. This solution faces two obstacles in practice. First, as shown in the main text, imperfect UD-MWIS optimization still raises interpretability issues. In the case of subdivided graphs, this is

expected to yield states with multiple independence violations akin to the example of Fig. 13.

A second obstacle is that the weight profile is itself constrained by experimental needs. For Rydberg atoms, the van der Waals interaction has long-range effects that are not taken into account in the model of UD graphs. In particular, higher detunings have to be applied to the atoms at the corner of the lattice or with a high number of neighbours [1, 5, 17]. This modification aims at counteracting the long-range repulsion between the second-nearest neighbors. However, according to the previous discussions, this causes an adversarial effect with respect to the weight profile W_w . This weight profile penalizes defects by having lower weights on the endpoints of the antiferromagnetic chains than on the inner vertices. This suggests that subdivision embeddings face the issue of degenerate domain walls, with constraints on the weight profile that would otherwise be used to alleviate these defects.

Appendix D: Mitigation

In this section, we study a weight profile on the CL embedding that is intended to penalize the localization of defects in the bulk of the CL graph. The motivation of this approach is to extend the defect penalization from the weight profile W_w introduced in the main text, which worked by lowering the weights on the endpoints of path graphs. However, this modification has to satisfy an energy rescaling condition to ensure it describes the realistic constraint that the sum of all weights (or detunings in the case of Rydberg atoms) is a constant. We will see the modified weight profile has unwanted effects because of this condition, but that it is beneficial to the interpretation if the energy budget is increased.

1. Modified weight profiles

a. Case of paths graphs

Let P_{2N} be some path graph with vertices v_1, \dots, v_{2N} and edges between v_i, v_{i+1} . For each vertex v_i , we introduce its *distance to the endpoints* $d_e(v_i) = \min(i, 2N - i) \in \mathbb{N}$. We seek to modify the weight profile W_w from Eq. (4) in order to penalize the presence of domain walls with a high distance to the endpoints. To do so, let us introduce the penalty function π defined by

$$\forall \mu \geq 0, \quad \forall \nu \geq 1, \quad \forall n \in \mathbb{N}, \quad \pi(n; \mu, \nu) = (n + 1)^\mu \cdot \nu^n. \quad (\text{D1})$$

A generalized weight profile $\tilde{W}_{w, \mu, \nu}$ to Eq. (4) is then

$$\tilde{W}_{w, \mu, \nu}(v_i) = W_w(v_i) \cdot \pi(d_e(v_i); \mu, \nu). \quad (\text{D2})$$

This weight profile on P_{2N} inherits from d_e and W_w the properties required for an embedding such as the weight ordering condition from Eq. (2). In addition, the $\pi(n; \mu, \nu)$ factor penalizes the localization of domain walls far from both endpoints of P_{2N} .

b. Case of the CL graph

This scheme can be generalized to the CL embedding at the cost of a more refined weight profile. We still use the gadgets from Ref. [18] which are delimited as blocks of 4×4 sites on a grid lattice. The CL graph can in this case be partitioned in 4×4 blocks where some blocks do not correspond to gadgets but rather to pairs of vertices from the diagonal, horizontal and vertical boundaries. For each block corresponding to a gadget or placed on the diagonal, we define its *block distance to the endpoints* $b_e \in \mathbb{N}$ as proportional to the distance of this block to the two blocks located at the endpoints of the diagonal in the CL graph. A block distance $b_e = 0$ is assigned to all remaining blocks, which are either on the vertical or on the horizontal boundary of the CL. Each vertex v_i being located in a block with a well-defined block distance b_e , we use the notation $b_e(v_i)$ to denote this mapping. The weight profile W_w , this time over the CL graph, is then modulated to define a new profile $\tilde{W}_{w, \mu, \nu}$ using the block distance:

$$\tilde{W}_{w, \mu, \nu}(v_i) = W_w(v_i) \cdot \pi(b_e(v_i); \mu, \nu). \quad (\text{D3})$$

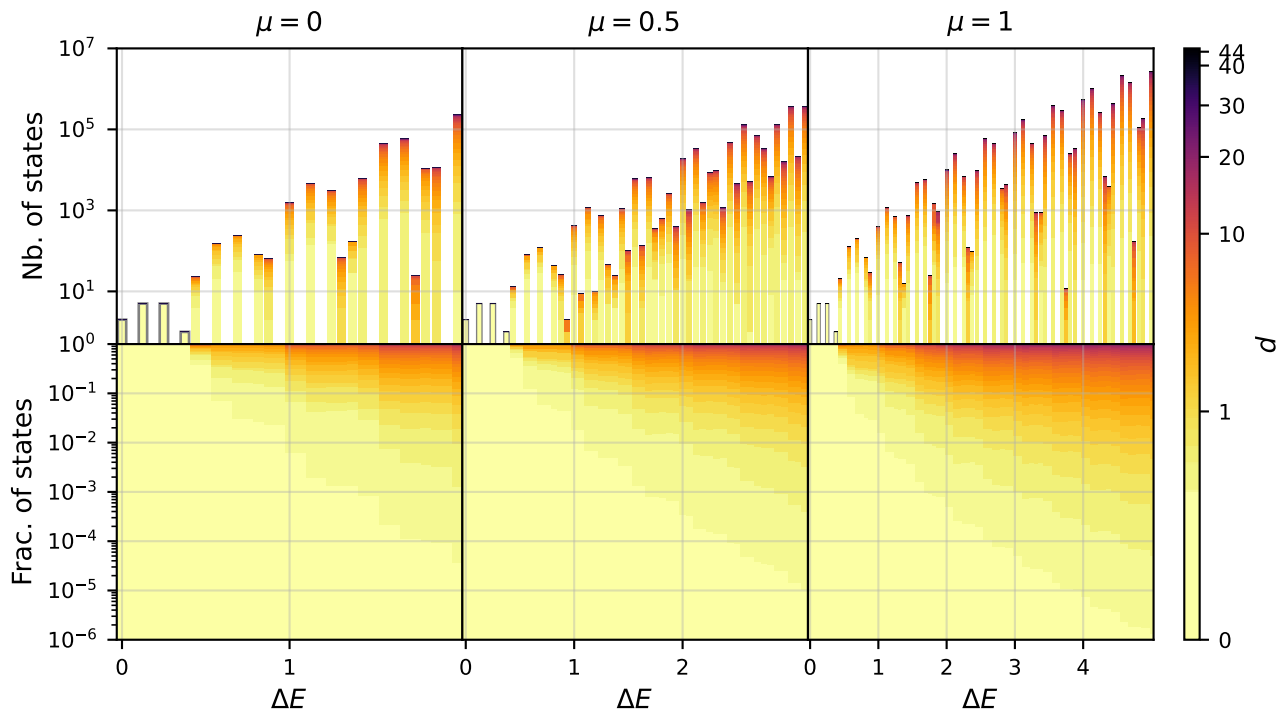


FIG. 14. (Top) Low-energy DoS for three set of parameters of the modified weight profile $\tilde{W}_{w,\mu,\nu}$ for a random CL graph with initial size $N = 5$. The indicated energies correspond to the energies in the CL before renormalization by the total energy (sum of all detunings), which depends on the weight profile. However the probed window is chosen so that the same fraction of the total energy is shown in all plots. (Bottom) Corresponding distribution of the fraction of d -interpretable states after cumulating with respect to the energy.

The weight ordering property and the optimality of the embedded MWISs remain guaranteed. Indeed, since all vertex in each gadget are modulated by the same factor, the 3 or 4 interpretable sub-configurations in each gadget have the same weight irrespective of (μ, ν) . The fact that $\pi(b_e(v); \mu, \nu) = \pi(0; \mu, \nu) = 1$ for any vertex v with $W_w(v) = 1/2 \pm w$ ensures that all the weight biases undergo the same modulation. (Note that other weight profiles with different weights *inside* the gadgets could also be engineered.) Eventually, this weight profile can be seen as a generalization of the one from Ref. [18], which is recovered by taking $(\mu, \nu) = (0, 1)$.

2. Results

We study the low-energy DoS for $\nu = 1$ and $\mu \in [0, 0.5, 1]$ in one of our random CL graphs with initial size $N = 5$. The sum of the weights $\sum_{v_i \in G_{\text{CL}}} \tilde{W}_{w,\mu,\nu}(v_i)$ in the graph increases with μ . Since this quantity would be proportional to the sum of all detunings used with an array of Rydberg atoms, some rescaling is needed to ensure we model different weight profiles with the same amount of energy supplied through the detunings. We do so by studying the low-energy DoS in a window $[E_{\text{MWIS}}, E_\mu]$ such that the ratio between $E_\mu - E_{\text{MWIS}}$ and $\sum_{v_i \in G_{\text{CL}}} \tilde{W}_{w,\mu,\nu}(v_i)$ is the same for all values of μ , with $E_{\mu=0} = 2$.

In Fig. 14, we plot the three resulting low-energy DoS. We see that the modified weight profile bends the curve of the fraction $\tau_d(E)$, at least for small distances d . After the energy renormalization both the value of $\tau_d(E)$ and the lowest-energy of non-interpretable states decrease when μ increases. These observations suggest the energy rescaling indirectly favors defects with a low b_e in a way that dwarfs the benefits from penalizing high- b_e defects, ultimately making the modified weight profile detrimental to the interpretation of approximate UD-MWIS solutions.

Despite the unwanted variation in the distribution of $\tau_d(E)$, we verify that the weight profile $\tilde{W}_{w,\mu,\nu}$ does penalize defects with a high b_e as follows. Given a configuration S in the CL graph and a vertex v_i , we denote by $\xi(v_i|S) \in \mathbb{N}$ the number of domain walls located at v_i (and at least one of its neighbors). This is done by looking for two successive unselected vertices along each chain. In the crossing (resp. non-crossing) gadgets, this approach readily generalizes




FIG. 15. Distribution of the defects in a CL graph corresponding to the probed configurations from Fig. 14. We observe that increasing μ leads to relatively less defects located at the center of the CL.

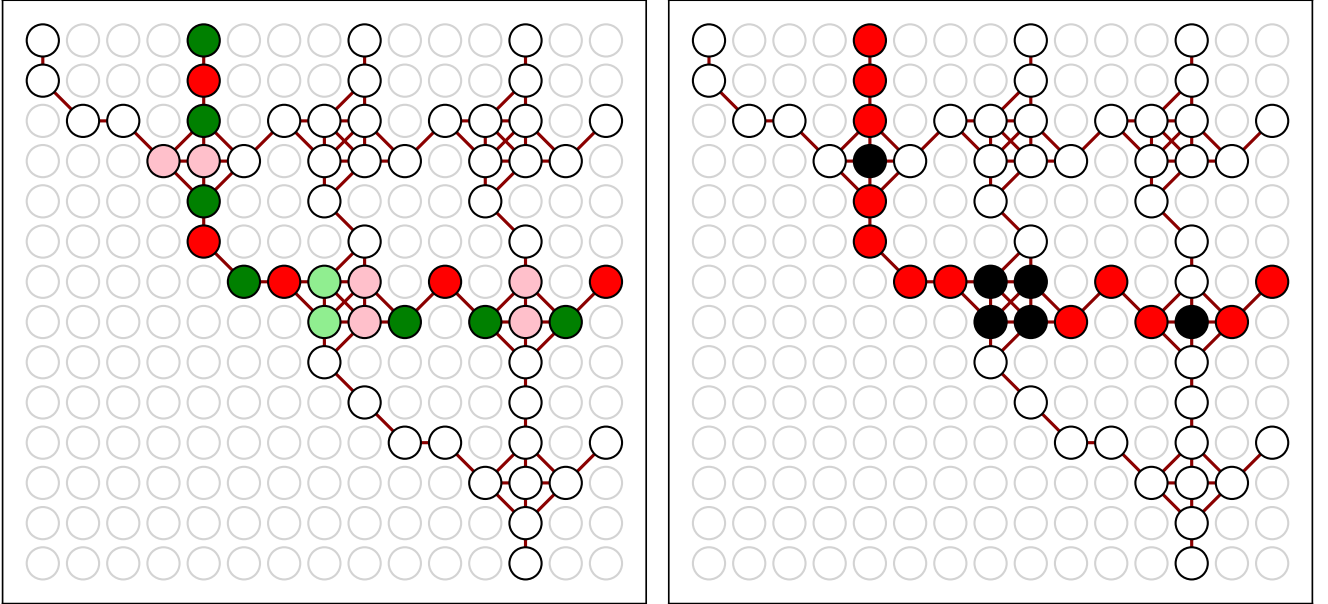


FIG. 16. (Left) Example of a chain as defined in App. D 2, where gadgets act by splitting some vertices (lighter colors). Thanks to this, it is possible to define a parity along the chain (red and green colors) and to identify any defect as the succession of two different colors with no selected vertex. (Right) Example of path nodes (red) and gadget nodes (black) as defined in App E. The CL can be partitioned between the path nodes of all N chains and the gadgets nodes. In an interpretable configuration, the path nodes have an antiferromagnetic ordering specified by whether the corresponding path is selected or not.

by considering that 1 (resp. 2) vertex along the chain undergoes a vertex splitting, so that the corresponding defect is identified as 3 (resp. 4) unselected vertices. An example is given in Fig. 16. (Note that if two or more defects are adjacent one vertex v_i can have $\xi(v_i|S) > 1$.) The value of $\xi(v_i|S)$ is then summed across all probed configurations, all other parameters being the same as in Fig. 14. For each μ , the result is linearly remapped so that its extremal values across all v_i are 0 and 1. The remapped value is plotted in Fig. 15 and can be understood as a measure of the prevalence of defects across the CL graph. As expected, we observe that for higher values of μ than the default $\mu = 0$, vertices with a high block distance b_e are energetically penalized.

Appendix E: Finding the nearest interpretable state as a QUBO problem

Here we describe the task of finding the nearest interpretable state of a CL configuration S as a QUBO problem closely related to an MWIS problem on a subgraph of the original graph.

We want to minimize the distance d between a fixed configuration S and any interpretable configuration S_{int} . This distance can be decomposed into contributions from path nodes (paths are indexed by p) and gadget nodes (gadgets are indexed by (p, q))

$$d = \sum_p d_p + \sum_{p>q} d_{pq}. \quad (\text{E1})$$

In this appendix, we use the following definitions. Gadget nodes are either i) the central node with 4 edges in the crossing-with-edge gadget, or ii) the central 4-clique in the crossing-without-edge gadget. Path nodes are all the other nodes and they can be grouped by path in a straightforward way. An example is shown in Fig. 16 (*right*), where all vertices along a path are colored depending on whether they are considered as path nodes (*red*) or gadget nodes (*black*). In an interpretable state, path nodes selection is determined only by the status (selected/unselected) of the path in which they belong, any ambiguity has been moved out to gadget nodes. We denote $n_p = 0$ if path p is unselected in S_{int} and $n_p = 1$ if it is selected.

For each path p , denote d_p^0 (d_p^1) the distance generated from S by deselecting (selecting) path nodes along path p (excluding gadget nodes). Then we have

$$d_p = d_p^0 + (d_p^1 - d_p^0)n_p. \quad (\text{E2})$$

In an interpretable configuration, each gadget can be in at most 4 states, corresponding to all possibilities of selecting or deselecting adjacent paths, which we label by 00, 01, 10 and 11. Therefore, for a given S , each gadget (p, q) can contribute to one of only 4 possible values of distance, which we denote d_{pq}^{ab} for $a, b \in \{0, 1\}$. In fact, crossing-with-edge gadgets can only be in two states, since they are a single node in our definition. For these, we will still use the same labels and keep in mind that $d_{pq}^{01} = d_{pq}^{10}$ and set $d_{pq}^{11} = 0$ since the case where both p and q are selected does not exist in an interpretable state. The contribution of an arbitrary gadget (p, q) to the distance is therefore

$$d_{pq} = d_{pq}^{00}[1 - n_p][1 - n_q] + d_{pq}^{01}[1 - n_p]n_q + d_{pq}^{10}n_p[1 - n_q] + d_{pq}^{11}n_pn_q. \quad (\text{E3})$$

For crossing-with-edge gadgets this simplifies to

$$d_{pq} = d_{pq}^{00} + [n_p + n_q][d_{pq}^{01} - d_{pq}^{00}], \quad (\text{E4})$$

using the fact that with such gadget, n_pn_q is always zero.

All together, we see that d is a quadratic form

$$d = C + \sum_p D_p n_p + \sum_{p>q} Q_{pq} n_p n_q \quad (\text{E5})$$

which must be minimized under the constraint that $\{n_p\}$ should form an independent set of the original graph G . This constraint can be formally introduced with a large extra term \bar{Q}_{pq} for each edge (p, q) of G , namely

$$d = C + \sum_p D_p n_p + \sum_{(p,q) \in \bar{E}} Q_{pq} n_p n_q + \sum_{(p,q) \in E} \bar{Q}_{pq} n_p n_q \quad (\text{E6})$$

where E is the set of edges of G , \bar{E} its complementary, and $\bar{Q}_{pq} \gg Q_{p'q'}, D_{p'}$. In this form this is a QUBO problem. We now show the proximity of this problem to an instance of MWIS. Observing that the quadratic term Q only comes from crossing-without-edge gadgets, it can be seen as a perturbation compared to the stronger independence constraint. If we ignore it, the task become equivalent to maximizing

$$f = \sum_p (d_p^0 - d_p^1)n_p \quad (\text{E7})$$

under the independence constraint. Paths p such that $d_p(0) \leq d_p(1)$ can be ignored, since they must be unselected in the optimal state (if it was not the case, then deselecting it would increase f without breaking the independence constraint). We can thus remove these paths from the optimization and we are left with positive weights only. Therefore, this problem is equivalent to finding an MWIS on the subgraph of the initial graph composed of those nodes where $d_p(0) > d_p(1)$, with weights $d_p(0) - d_p(1)$. Informally, $d_p(0) > d_p(1)$ means that the path p is closer to being selected than unselected in S . Interestingly, this tells us that, in the search for the distance, paths that are closer to being unselected can be set to unselected right away, if one is happy to ignore the contribution from crossing-without-edge gadgets.