IOP Publishing Quantum Sci. Technol. 10 (2025) 045010 https://doi.org/10.1088/2058-9565/adf2d5

Quantum Science and Technology

OPEN ACCESS

RECEIVED
6 December 2024

REVISED
6 April 2025

ACCEPTED FOR PUBLICATION
22 July 2025

PUBLISHED
7 August 2025

PAPER

Enhancing the energy gap of random graph problems via XX-catalysts in quantum annealing

Luca A Nutricati[1,2,*], **Roopayan Ghosh**[3], **Natasha Feinstein**[1], **Sougato Bose**[3] **and P A Warburton**[4]

[1] London Centre for Nanotechnology, University College London, WC1H 0AH London, United Kingdom
[2] Rudolf Peierls Centre for Theoretical Physics, University of Oxford, Parks Road, OX1 3PU Oxford, United Kingdom
[3] Department of Physics and Astronomy, University College London, WC1E 6BT London, United Kingdom
[4] London Centre for Nanotechnology and Department of Electronic & Electrical Engineering, University College London, WC1H 0AH London, United Kingdom
[*] Author to whom any correspondence should be addressed.

E-mail: luca.nutricati@physics.ox.ac.uk

Keywords: quantum annealing, catalysts, enhancing energy gap, removing first-order phase transitions, minimum energy gap

Abstract

One of the main challenges in solving combinatorial optimisation problems with quantum annealers is the emergence of extremely small energy gaps between the ground state and the first excited state of the annealing Hamiltonian. These small gaps may be symptoms of an underlying first-order phase transition, which, according to the adiabatic theorem, can significantly extend the required anneal time, making practical implementation effectively infeasible. In this paper we demonstrate that attaching an XX-catalyst on all the edges of a graph upon which a MWIS (Maximum Weighted Independent Set) problem is defined, significantly enhances the minimum energy gap. Remarkably, our analysis shows that the smaller the energy gap, the more effective the catalyst is in opening it. This result is based on a detailed statistical analysis performed on a large number of randomly generated MWIS problem instances on both Erdős–Rényi and Barabáasi–Albert graphs. We perform the analysis using both stoquastic and non-stoquastic catalysts.

1. Introduction

Optimisation problems are of critical importance in numerous research fields, from protein folding in biology [1], to drug design in chemistry [2–4] and portfolio optimisation in finance [5–7] as well as applications in high energy physics, see for example [8–18]. In particular, combinatorial optimisation problems in graph theory, such as maximum cut, maximum (weighted) independent set, maximum vertex cover and set cover (and generalisations thereof) have ubiquitous applications, from circuit layout design [19] to network routing, where data packets are arranged such that the structure can be modeled as an maximum weighted independent set (MWIS) problem [20].

Given their application in many areas, it is crucial to develop techniques and possibly technologies whereby one can efficiently solve these problems, overcoming their intrinsic complexity from the classical perspective. Many of them are NP-hard, meaning they cannot be classically solved in polynomial time if P $\neq$ NP. With the emergence of powerful near-term quantum devices and algorithms, it is natural to consider if these technologies could offer a qualitatively different solution for NP-hard problems. Specifically, could they provide a novel method to overcome classical techniques leading to a less harsh exponential scaling exponent? Although various quantum computing paradigms have been proposed, one type of quantum computer, known as a quantum annealer, has been designed specifically for the task of solving optimisation problems.

The idea of quantum annealing traces its origins from the pioneering works in [21, 22], where quantum fluctuations have been introduced into the simulated annealing process with the aim to converge to the global minimum quicker than classical simulated annealing (see [23, 24] for comprehensive reviews on quantum annealing).

In current quantum annealers each qubit physically corresponds to a superconducting loop with an associated magnetic flux. Qubits are coupled together with various topologies using inductive couplers. Quantum annealers are machines able to evolve a quantum system initialising it using a local transverse field inducing all the qubits to be in a superposition of states, which corresponds to the ground state of the system. Gradually, using flux biases, the transverse field is shut down and the time-dependent Hamiltonian of the system evolves towards a new Hamiltonian whose ground-state encodes the optimal solution of the problem to be solved. If this evolution proceeds adiabatically, the ground-state of the final Hamiltonian is obtained with high probability at the end of the anneal.

Let us now emphasise some of the crucial aspects of quantum annealing which will be crucial in the following sections. The Hamiltonian that is implemented in available quantum annealers is usually written as

$$H_A \;=\; (1-s)\, H_D + s H_P\,, \tag{1}$$

where $s = s(t)$ is called annealing parameter, which depends upon the physical time t and defines the so-called *anneal schedule*. In the case of *forward annealing*, this is typically an arbitrary function which interpolates from $s = 0$ at $t = 0$ to $s = 1$ at $t = T$, where T is the total anneal time. This function has to be carefully chosen as in some cases its optimisation leads to dramatic performance improvements. However, this is beyond the scope of this paper and from now on we shall focus on the simplest choice of this schedule, i.e. a linear function defined by

$$s(t) \;=\; \frac{t}{T}\,. \tag{2}$$

Finally, H_D and H_P denote the driver and the problem Hamiltonian, respectively. The driver Hamiltonian is a sum of local X-operators

$$H_D \;=\; -\sum_{i=1}^{N} \sigma_x^{(i)}\,, \tag{3}$$

whereas the problem Hamiltonian encodes the problem we aim to solve in the ground of an Ising Hamiltonian

$$H_P \;=\; \sum_{i,j\in E(G)} J_{ij}\,\sigma_z^{(i)}\sigma_z^{(j)} \;+\; \sum_{i\in V(G)} h_i\,\sigma_z^{(i)}\,, \tag{4}$$

where G is the particular graph that constitutes the Ising model, i.e. a set of N nodes (or spins), indicated with $V(G)$, and links between them, indicated by $E(G)$ (couplings) with strength J_{ij}. Finally, h_i are the local biases and $\sigma_z^{(i)}$ the Pauli z matrix, which is related to the longitudinal spin state of the i-th spin.

As previously mentioned, quantum annealers make use of the adiabatic theorem to ensure that the entire evolution of the system from $s = 0$ to $s = 1$ is confined to the ground state of the annealing Hamiltonian in equation (1), without jumps to any excited state. This means that if the adiabatic condition is satisfied, the solution of our optimisation problem is encoded in the final classical values of the spins. However, the condition under which the adiabatic theorem holds dictates that the total anneal time must scale with the inverse of the minimum energy gap between the ground state and the first excited state, i.e.

$$T \;\sim\; \frac{1}{\left[\min_{0\leqslant s\leqslant 1} \Delta E_{01}\,(s)\right]^2}\,, \tag{5}$$

where we recall that T is the total anneal time and $\Delta E_{01}(s) = E_1(s) - E_0(s)$ is the energy gap between the ground and the first excited state of the annealing Hamiltonian H_A for that specific value of the annealing parameter s.

It is well known that some problems exhibit exponentially closing gaps in the system size, resulting in the following scaling with the total number of spins N:

$$\min_{0\leqslant s\leqslant 1} \Delta E_{01}\,(s) \;\sim\; e^{-N}\,. \tag{6}$$

This is the hallmark that a *first-order phase transition* is taking place, which may ultimately lead to an exponentially large anneal time, preventing any possibility to solve the problem with quantum annealers. Indeed, exponentially closing gaps are a common bottleneck in adiabatic quantum computation [25–27] and a great deal of effort has been invested in finding ways to circumvent this problem, by adding, for example, additional drivers and/or catalysts, see [28–35]. Additionally, the works mentioned provide an extensive

description of the use of stoquastic catalysts in contrast with non-stoquastic ones, in particular whether or not the introduction of non-stoquastic interactions enhances the performance compared to the case of the traditional stoquastic formulation. We emphasise that in the rest of the paper we use the term 'stoquastic' to refer to an operator with non-positive off-diagonal matrix elements in the computational basis; whereas the term 'non-stoquastic' includes all the other cases.

In this paper we report a statistical analysis using a particular catalyst with multiple XX-interactions on randomly generated MWIS problem instances, with the aim to compare its effectiveness with conventional quantum annealing. We defer the detailed discussion of the catalyst to the following section and, for now, provide an overview of the key points of our analysis. Specifically, by using Arnoldi iterative method to numerically diagonalise the Hamiltonian at each discretised anneal step, we observe statistical evidence of a gap enhancement in most of the instances when a stoquastic catalyst is employed. This reinforces the effectiveness of this class of catalysts, as demonstrated by the general analysis in [36] and the results in [37] regarding the weak-strong cluster problem. We shall repeat the same analysis with the non-stoquastic version of the same catalyst, comparing the two approaches.

The paper is organised as follows: after a brief introduction to the MWIS problem in section 2, we move on to a comparative analysis in section 3 where we discuss the results and effectiveness of our proposed catalyst in randomly constructed MWIS problem instances on two different classes of networks. We finally conclude in section 4.

2. The MWIS problem on quantum annealers

The MWIS problem is a fundamental problem in graph theory and combinatorial optimisation which has a number of real-world applications, from network design [38], scheduling problems [39] and combinatorial auctions [40], to molecular biology [41] and clustering aggregation [42]. In its simplest formulation, it consists in finding a subset of non-connected vertices $\widetilde{V}(G) \subset V(G)$ of a weighted graph G such that the sum of the weights of the nodes in $\widetilde{V}(G)$ is maximised. It is well known that this problem is NP-hard, which means that, if P $\neq$ NP, there is no known polynomial-time algorithm to solve it for general graphs. This makes it computationally challenging, especially for large graphs.

It can be shown that solving the MWIS problem on a weighted graph G with N vertices $V(G) = \{w_i \,|\, i = 1, \ldots, N\}$—where w_i is the weight associated to the node i—and edges $E(G)$, is equivalent to finding the ground state of the Hamiltonian in equation (4) with the following linear and quadratic couplings [43, 44]

$$\begin{aligned} h_i &= \sum_{j \in \mathrm{nbr}(j)} J_{ij} - 2\, w_i \\ J_{ij} &\geqslant \min\left(w_i, w_j\right), \end{aligned} \tag{7}$$

where $\mathrm{nbr}(i) = \{j : (i,j) \in E(G)\}$ is the subset of vertices directly connected with node i. The definition of the couplings in equation (7) guarantees that the solution of the MWIS problem is encoded in the ground state of the Hamiltonian in equation (4) and in our convention will correspond to the subset constituted by the spins pointing upwards.

Being notoriously hard to solve classically, one may wonder if the MWIS problem may be efficiently tackled using quantum annealers, ultimately circumventing all the classical obstructions which prevent its solution in polynomial time. These obstructions are mainly due the exponential growth in the system size of the subsets of independent vertices which have to be checked and sorted by their total weight in order to solve the problem. One might expect significant advantages from using quantum annealing for this type of problem. Certain parameter regimes of the MWIS problem, however, exhibit first-order phase transitions during the adiabatic evolution of the system. This happens when the instantaneous ground and first excited states of $H_A(s)$ are separated by a large Hamming distance but are very close in energy (i.e. small energy gap) for some critical value of s. As outlined in the introduction, this may extend the anneal time to impractically long durations in order to ensure adiabatic evolution. In figure 1 we report a typical example of anneal in a specific instance of the MWIS problem, generated on a random graph, where a small gap between the ground state and the first excited state appears around $s \approx 0.65$. This figure also shows the evolution of a specific quantity we have denominated as 'imbalance', which we quantitatively define as follows

$$I = \sum_{i \,\in\, \mathrm{GS}_\uparrow} \sigma_z^{(i)} - \sum_{i \,\in\, \mathrm{GS}_\downarrow} \sigma_z^{(i)}, \tag{8}$$

where $\mathrm{GS}_\uparrow$ ($\mathrm{GS}_\downarrow$) is the set of spins which are pointing upwards (downwards) when the system is in the ground state of H_p. Computing the evolution of the expectation value of this quantity during the annealing,

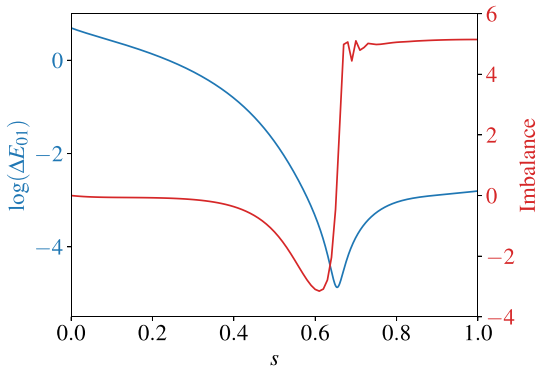

Figure 1. The imbalance I and the log of the energy difference between the ground and the first excited state versus the annealing parameter s for a random generated MWIS problem instance with *10* nodes. The energy gap reaches the minimum around $s \approx 0.65$ where the imbalance exhibits a sudden change.

effectively gives a measure of the Hamming distance between a given instantaneous eigenstate of H_A and the ground state of the problem Hamiltonian. Therefore, this quantity serves as an order parameter for the phase transition. Indeed, as expected, we can see in figure 1 that I exhibits a sharp change when the system reaches the minimum gap between the ground and the first excited states. Of course even first-order phase transitions are rounded out for a finite size system, which explains why we are not observing true discontinuity in the figure above, indicating that this particular system undergoes a crossover rather than a true phase transition. Nevertheless, the term 'first-order phase transition' still is a reasonable way to describe these cases. Crucially, the transitions we analyse do not arise from competition between the local driver and problem Hamiltonian strengths, which typically lead to an Ising-like transition. Instead, they occur in a regime where the problem Hamiltonian dominates, aligning the ground state with specific computational states. This mechanism of state selection transitions, has been repeatedly demonstrated to be a first-order transition in the thermodynamic limit [45–48].

It is crucial now to stress that any attempt to consistently study the scaling of the gap with the system size would be meaningless in this context, since our entire analysis is based on generating individual random instances of the MWIS problem without any possible conception of size scaling. However, perturbative analysis shows that the minimum gap typically scales exponentially with the number of spin flips [48], i.e. the number of spins that need to be flipped to reach the ground state from the first excited state. It is therefore reasonable to expect that instances with very small energy gaps will be generated more frequently when the system size is larger. We shall see later that this expectation matches what we find. However, it is well known that there are systems in which the minimum gap closes algebraically with the system size (see for example [49]), which would imply that the minimum gap we observe in larger systems would be proportional to $k^{-\gamma}$, where k is the number of spin flips and γ a constant that sets the scaling. Although a meaningful way to estimate this scaling remains elusive in our setup, we note that algebraically closing gaps are typically observed when the ground state is highly degenerate, which is in general not the case of the MWIS problem.

In the following, we will use the minimum gap as the sole criterion to identify potentially difficult MWIS problem instances, where a first-order phase transition might take place. Specifically, we shall start from generating a large number of random instances and proceed to classifying their hardness based on how small the energy gap is. Among the various strategies to improve quantum annealing performance, we focus on one that uses catalysts. Catalysts can be designed to reshape the energy landscape, increasing the minimum energy gap and enabling more efficient adiabatic evolution. With careful selection and tuning, they offer a powerful means to enhance gap sizes and improve the annealing process. In the following we shall investigate on how a particular class of catalysts leads to a general gap improvement on randomly generated MWIS problem instances.

3. Adding a catalyst term in the Hamiltonian

Before going through the details of our analysis, let us provide a general overview on how catalysts can be used in quantum annealing. A catalyst corresponds to a set of new interactions appearing in the annealing Hamiltonian H_A in equation (1), which now reads

$$H_A = (1-s)\,H_D + sH_P + s(1-s)\,H_C, \tag{9}$$

where the catalyst term H_C is in general a combination of the X,Y and Z-operators. The catalyst term does not have any effect at the beginning and at the end of the evolution, i.e. at $s=0$ and $s=1$, due to the prefactor $s(1-s)$. Instead, it affects the energy levels during the adiabatic evolution of the system with an impact that increases and decreases quadratically in the anneal parameter s.

In certain MWIS problem instances with a bi-partite structure one can observe that a single XX-catalyst defined as the product of two X-operators acting on two nodes of the graph is effective in enhancing the minimum energy gap [34]. In addition, in [50] we show on how specific catalysts can completely remove first-order phase transitions in the same class of problems with specific characteristics. However, these approaches would at least require partial information on the specific instance, for example the partition structure, symmetries, general properties of the graph, *etc*. In contrast, here we will adopt a complementary approach by randomly generating thousands of problem instances and operating without any *a-priori* knowledge of the structural properties of the graphs. This method allows us to comprehensively assess the average efficacy of the proposed catalyst on the MWIS problem in its full generality.

From now on we shall focus on a particular catalyst which will be at the centre of our entire analysis. We define

$$H_C = J_c \sum_{i,j\in E(G)} \sigma_x^{(i)}\sigma_x^{(j)}, \tag{10}$$

where we note that the summation is limited to i's and j's which are directly connected by an edge in the problem graph and where J_c is a parameter that tunes the strength of the catalyst. We note that this catalyst has already been adopted in [31, 32] though in different contexts from the one we shall analyse. We will we set $J_c=-1$ for the rest of this paper, meaning that we are introducing a *stoquastic* catalyst. Many works have been directed to the study of non-stoquastic catalysts, see for example [28–31, 51–53], demonstrating their efficacy in improving the minimum energy gap in various scenarios. However, parallel studies have proved that stoquastic catalysts can lead to similar improvements under certain conditions for specific problem settings [36, 37, 50]. In this section, we shall later discuss the features of these catalysts in terms of our findings.

In particular, the catalyst in equation (10) consists of a set of XX (transverse) interactions between pairs of spins whose corresponding vertices in the graph are connected by an edge. In the following we shall focus on two classes of graphs: random graphs and scale-free networks—picking Erdõs–Rényi graphs and Barabáasi–Albert networks as archetypal examples of the two categories, respectively. In both cases we will generate MWIS problem instances on randomly generated graphs with random parameters uniformly distributed in the following ranges

$$\begin{aligned} J_{ij} \equiv J &\in [1,2] \\ w_i &\in [0,1], \end{aligned} \tag{11}$$

where we choose the couplings J_{ij} to be constants, i.e. independent of the sites i and j, such that the second line of equation (7) is always satisfied. Note that this can be done without any loss of generality, i.e. without restricting the number of possible realisations of the MWIS problem.

Let us now describe in more detail the kind of networks which will serve as a testing ground to benchmark the effectiveness of the catalyst in equation (10).

3.1. Erdős–Rényi graphs

Erdős—Rényi (ER) graphs are a fundamental model in the study of random graphs, widely applied across fields such as computer science, physics, biology, and social sciences. An Erdõs—Rényi graph $G(N,p)$ is defined by two parameters: the number of vertices N and the probability p with which each pair of vertices is connected by an edge. The parameter p in this model can be thought of as a weighting function; as p increases the model becomes more and more likely to include graphs with more edges.

To conduct a statistical analysis, we have generated random Erdõs–Rényi graphs with $N=10$ nodes and constructed random MWIS problems on them, using parameters uniformly chosen from the ranges specified in equation (11). In particular we have analysed more than 10 000 problem instances each for four different values of p and compared the minimum energy gap with and without the catalyst. In figure 2 we can see the effect of using this catalyst for $p=0.25, 0.4, 0.6, 0.8$, showing that this method enhances the minimum energy gap for most of the instances. Of particular interest is the central region of these plots, where the approach without the catalyst leads to a minimum energy gap between 10^{-5} and 10^{-3}. The instances in this region receive a gap enhancement when the catalyst is used, increasing the minimum gap by up to three orders of magnitude. To better assess the effectiveness of this catalyst, we have added an inset to each figure showing a box plot representing the distribution of the gap improvement, measured as Δ_c/Δ, for various

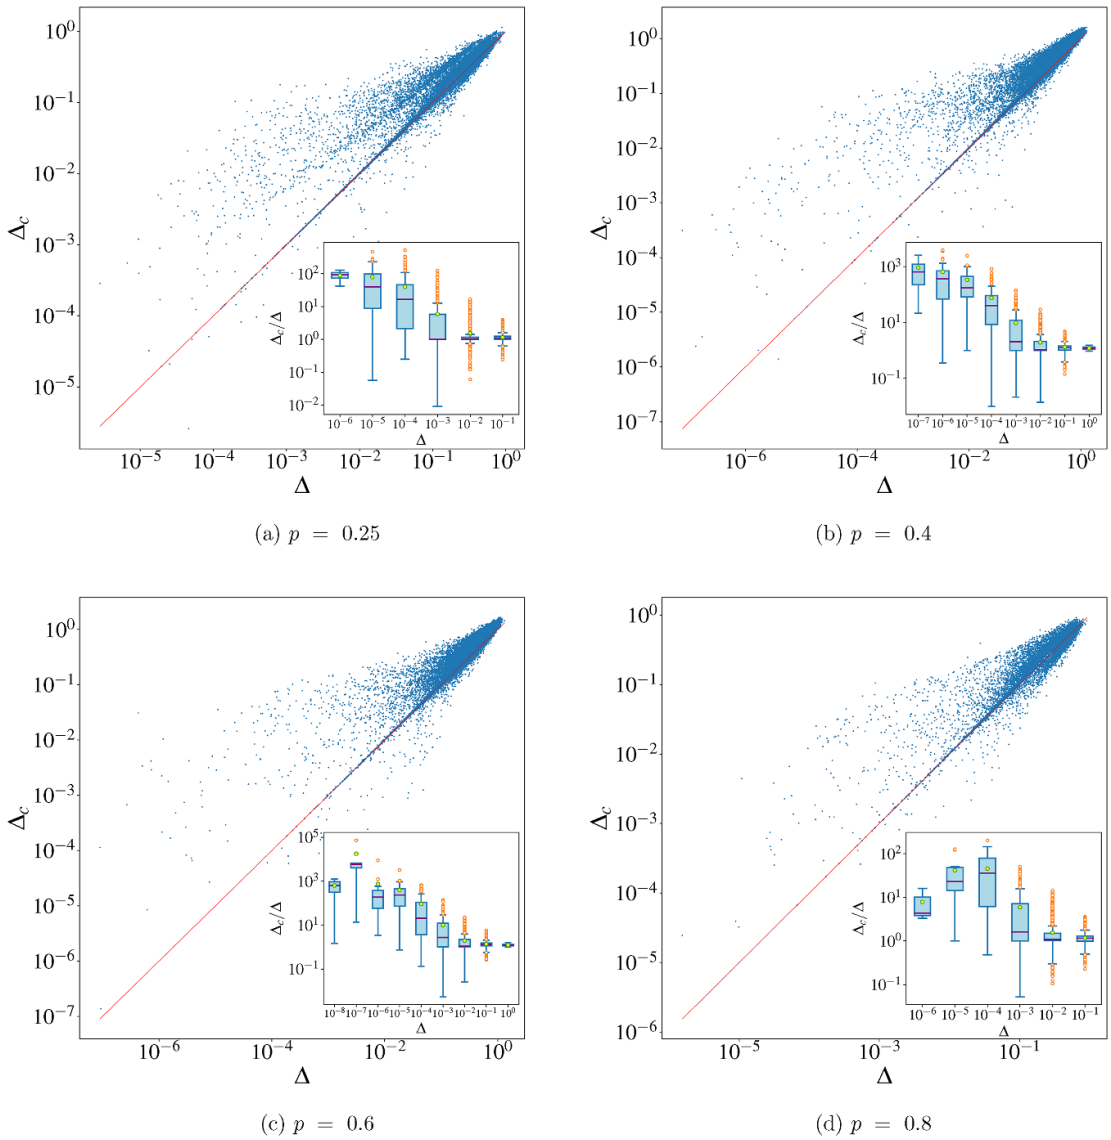

Figure 2. Randomly generated MWIS problem instances on Erdős–Rényi networks for different values of edge probability p. For each of these values we generated more than 10 000 instances and plotted Δ_c, the minimum gap using the catalyst versus the minimum gap, Δ, in the catalyst-free case. In all these cases we observe a general gap enhancement, meaning that the majority (74% on average) of the instances lay above the critical red line $\Delta_c = \Delta$ which divides the region in which the catalyst enhances the minimum gap from the region in which its effect is opposite. The box plots show the distribution of the ratio Δ_c/Δ for various orders of magnitude of the original minimum energy gap Δ. The purple line and the yellow circle indicate the median and the mean respectively. The top and bottom edges of the box show the upper and lower inter-quartile range. Outliers, depicted as orange circles, are determined if they are more than 1.5 times more than the inter-quartile range away from the median. The two caps attached to each box show the maximum and minimum data points, excluding outliers. The labels 10^{-n} on the horizontal axis must be interpreted as a shorter notation to indicate the range from 10^{-n} to $10^{-(n+1)}$.

orders of magnitude of the minimum energy gap Δ. The purple line and the yellow circle represent the median and the mean, respectively. From the box plots we note that harder instances, populating the left hand region of each scatter plot, are distributed around larger values of Δ_c/Δ meaning that these instances get a larger gap improvement with respect to easier problems with bigger minimum gaps, confirming what one can qualitatively deduce from the corresponding scatter plot.

We would also like to stress that using the non-stoquastic version of the same catalyst, i.e. setting $J_c > 0$ in equation (10), leads to different results. The scenario that emerges using the non-stoquastic version of the same catalyst is only comparable to the one depicted in figure 2 for harder instances (meaning instances with $\Delta < 10^{-3}$). Whereas when Δ increases to approach 10^{-2}, the non-stoquastic catalyst on average, leads to a further shrinkage of the minimum gap, becoming disadvantageous in this regime. Overall, most of the instances in this case lie below the critical line $\Delta_c = \Delta$, meaning $\Delta_c < \Delta$ in the majority of the cases. Note that, we have analysed the non-stoquastic case for different positive values of J_c, all leading to similar results as the case discussed in appendix, where we show and discuss the results for $J_c = 1$.

3.2. Barabáasi–Albert graphs

Barabáasi–Albert (BA) graphs constitute a different class of graphs than the previous scenario as they are 'scale-free'. This arises from the fact that the distribution of degrees across different vertices (nodes) of this graph follows a power law decay, distinct from, for example, the normal distribution decay in the previous case. This in turn implies that several of the nodes in such a graph have a preferential attachment to many other nodes, and are called hubs. This introduces a new kind of complexity in the problem compared to the Erdõs–Rényi networks.

Apart from the theoretical motivation, studying MWIS problems in these graphs might have a strong impact on real-world problems. All large scale networks such as citation networks, the internet or social networks have a BA graph structure [54–56]. In principle, we can weigh each node by its degree, quantifying the 'importance' of an individual in the network and then the MWIS problem will help identify the most important individuals in the network who have no connection with each other. This would be greatly useful to model and understand social interactions.

A Barabási-Albert network is defined through a parameter m, which determines the number of edges that each new node introduces when it is added to the network. Essentially, m represents how many connections a new node will form with existing nodes. When a new node is added to the network, it does not connect to existing nodes randomly. Instead, it follows the principle of preferential attachment, meaning that nodes with higher degrees (more connections) are more likely to receive new links. This mechanism ensures that well-connected nodes continue to attract more connections, leading to the formation of hubs. The choice of m significantly impacts the structure of the network. A larger m results in a network with higher average connectivity, as each new node introduces more connections.

In what follows, we take the same approach as before by putting random weights and random couplings as specified in equation (7). We then asses the efficiency of the proposed catalyst on this new set of graphs in the same way we did for the case of Erdõs–Rényi networks, i.e. generating more than ten thousand problem instances for four different values of m and comparing the minimum energy gap with and without the catalyst. In figure 3 we show the result for $m = 3, 4, 5, 6$. We note that this catalyst is again statistically effective in enhancing the minimum energy gap. The 'hard instances' region, $\Delta \lesssim 10^{-2}$, is populated by points which are mostly distributed above the diagonal, meaning that the catalyst is very efficient in opening the gap in these cases. As before, the insets show the distribution of the gap improvement for different orders of magnitude of Δ, leading to the same conclusion as in the Erdõs–Rényi case.

We finally notice that the scenario that emerges using the non-stoquastic counterpart of this catalyst is qualitatively the same as the one we observe in the Erdõs–Rényi case. See appendix for a more detailed discussion.

3.3. Ground state fidelity improvement

Having demonstrated statistical evidence that this catalyst achieves a general gap enhancement in randomly generated MWIS problems with $N = 10$ nodes, we now examine whether this results in a consistent reduction of the total anneal time T required to reach the ground state. To assess this, we use ground state fidelity, defined as

$$F \equiv \langle \psi(T) | \mathrm{GS}(H_P) \rangle^2 , \tag{12}$$

where we recall that T is the total anneal time and $\mathrm{GS}(H_P)$ is the ground state of the problem Hamiltonian. The fidelity is then computed by numerically solving the time-dependent Schrödinger equation starting from the initial ground state at $t = 0$, to $t = T$, and computing its overlap with exact problem ground state. In particular, the state $|\psi\rangle$ is calculated such that $|\psi(0)\rangle$ corresponds to the ground state of H_D. We shall compute the fidelity F by evolving $|\psi\rangle$ using the Hamiltonian with and without the catalyst, denoting the associated fidelities as F_c and F, respectively. In figure 4 we can see an example in a specific instance where the catalyst improves the minimum energy gap by an order of magnitude, precisely from 10^{-3} to 10^{-2}.

We now broaden this analysis to several problem instances, in order to check if, on average, the energy gap improvement consistently results in a reduction of the total anneal time required to reach the ground state of the problem Hamiltonian. In order to do so, we have randomly selected around one thousand instances from the ones shown in figures 2 and 3 and computed the final ground state fidelity with and without the catalyst, fixing the total anneal time to $T = 5\,\mu s$. We have collected these results in figure 5. We specify that the selected instances are those with $\Delta < 10^{-2} < \Delta_c$, as we are interested in cases where the catalyst is effective in opening the gap. In principle, we could have chosen all the cases with $\Delta_c > \Delta$, but this would have included instances with very small energy gap, even if the catalyst is effective in increasing it by several orders of magnitude. This implies that that the total anneal time required to reach the ground state with acceptable accuracy would have been prohibitive in terms of our computational capability.

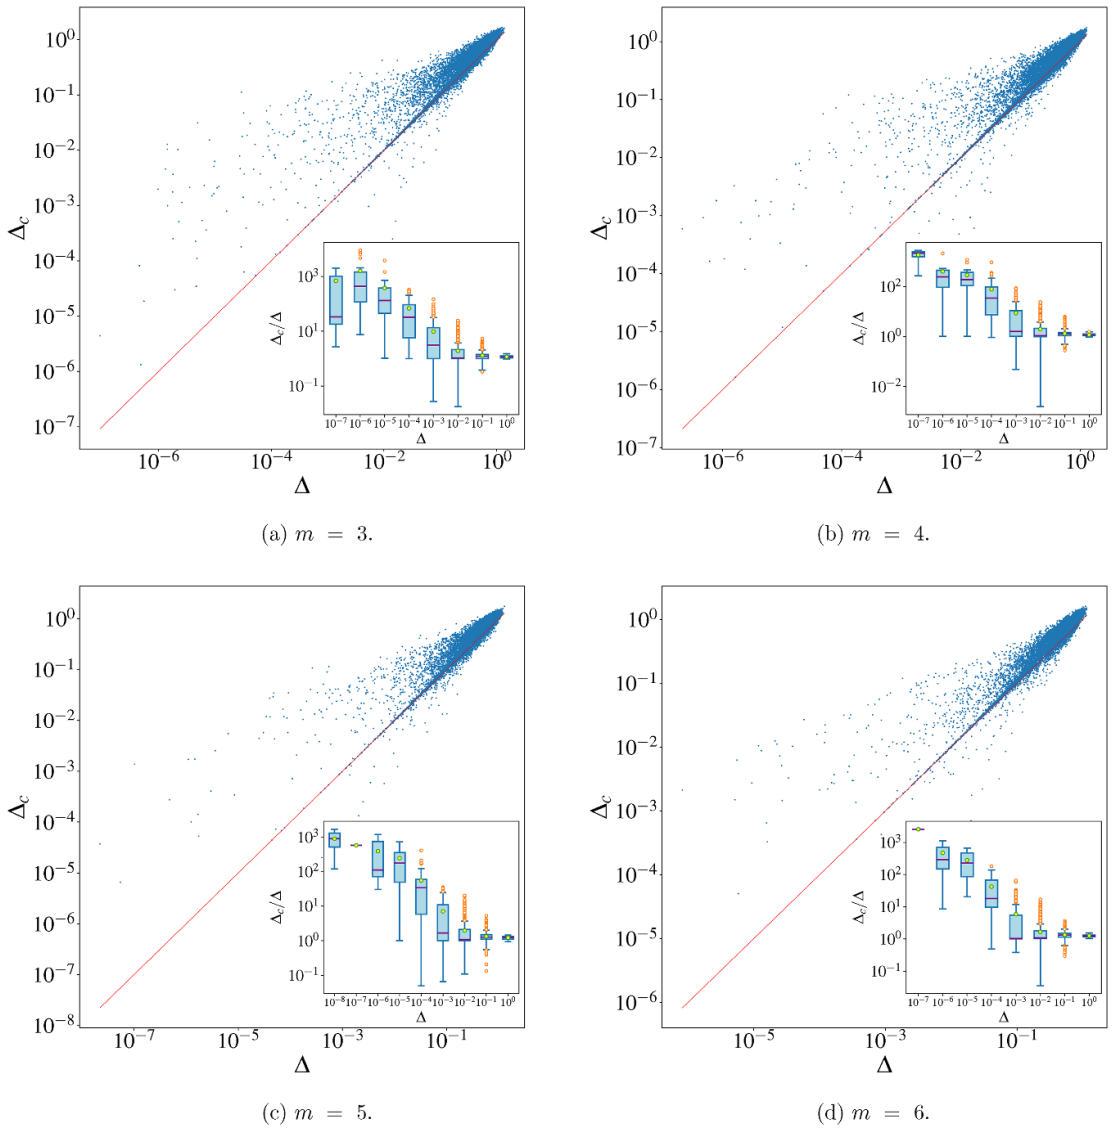

Figure 3. Randomly generated MWIS problem instances on Barabáasi–Albert networks for different values of m. For each of these values we generated more than 10 000 instances and plotted Δ_c, the minimum gap using the catalyst versus the minimum gap, Δ, in the catalyst-free case. In all these cases we observe a general gap enhancement, meaning that the majority of the instances (81% on average) lay above the critical red line $\Delta_c = \Delta$ which divides the region in which the catalyst enhances the minimum gap from the region in which its effect is opposite. The box plots show the distribution of the ratio Δ_c/Δ for various orders of magnitude of the original minimum energy gap Δ. The purple line and the yellow circle indicate the median and the mean respectively. The top and bottom edges of the box show the upper and lower inter-quartile range. Outliers, depicted as orange circles, are determined if they are more than 1.5 times more than the inter-quartile range away from the median. The two caps attached to each box show the maximum and minimum data points, excluding outliers. The labels 10^{-n} on the horizontal axis must be interpreted as a shorter notation to indicate the range from 10^{-n} to $10^{-(n+1)}$.

Using the catalyst, we were able to achieve improvements of several orders of magnitude in the final fidelity. In fact, less than 1% of the instances present fidelity F_c smaller than 0.6 after the anneal evolution. This also provides an indirect estimate of the matrix elements of $\mathrm{d}H/\mathrm{d}t$ between the instantaneous ground state and the first excited state of the annealing Hamiltonian, showing that they do not act as a bottleneck in our setup.

3.4. Scaling properties

Having analysed the effectiveness of the catalyst in improving the energy gap for fixed size problem instances, let us now extend the treatment to assess how this improvement scales with the system size. In figure 6 we have plotted the fraction of potential harder instances (meaning those ones with $\Delta < 10^{-1}$) which receive a gap improvement from the catalyst (i.e. the fraction of instances with $\Delta_c > \Delta$) for various numbers of nodes, where Δ_c and Δ are the minimum gaps with and without the catalyst, respectively. Additionally, in figure 7 we show the same scaling plot but with a different metric, denoted as the average gap improvement and quantified as the geometric mean of the ratio Δ_c/Δ. Since this ratio represents an improvement rate, the

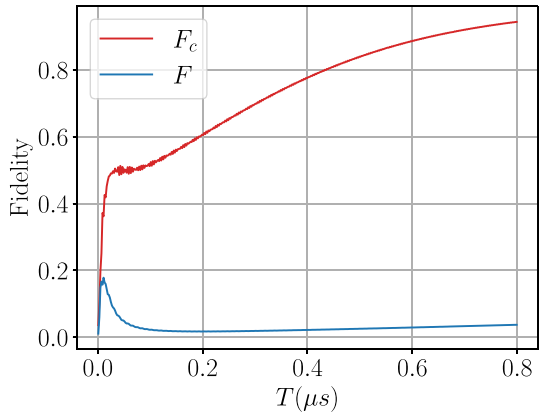

Figure 4. Final fidelity for different values of total anneal time T. The red curve refers to the case with the catalyst, the blue to the catalyst-free case.

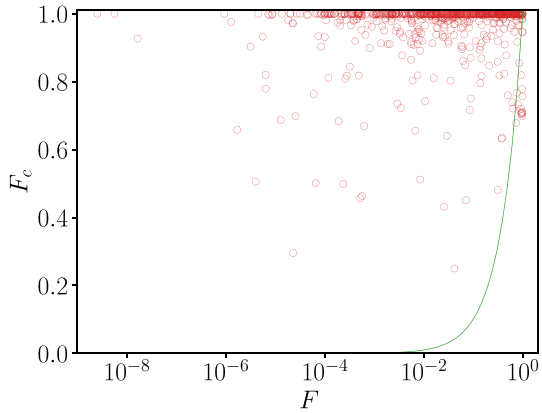

Figure 5. Scatter plot of the final fidelity (total anneal time $T = 5\,\mu s$) using the catalyst versus the same quantity in the catalyst-free case. This comparison includes around one thousand of problem instances with $\Delta < 10^{-2} < \Delta_c$, randomly chosen from the instances presented in the previous analysis. The plot includes the line $F_c = F$ in green.

geometric mean is a more meaningful estimator than the arithmetic mean. For the Erdõs–Rényi case the data are collected by generating random graphs with $p = 0.5$, whereas for the Barabáasi–Albert we analysed two cases: one with $m = 3$ independently of the number of the spins N and one maintaining a constant ratio $\mathrm{N\,m^{-1}} = 2$. Figure 6 shows that the majority of the randomly generated MWIS problem instances present a larger minimum energy gap once the catalyst is used. We note a monotonic decreasing behaviour only in the BA case with constant m, while the other two cases present oscillations of the order of a few percent. Given the limited number of data points it is not possible to extrapolate the behaviour for large N in the three cases analysed, which certainly needs to be addressed in a future study.

On the other hand, figure 7 displays an overall increasing behaviour in terms of the average gap improvement for both scale-free and random networks as the number of total nodes in the graph increases. Even more notably, the scaling persists although the fraction of harder instances —i.e. those ones with $\Delta < 10^{-2}$, increases with N as well, as one can see in figure 8.

Specifically, the fraction of these harder instances almost triples from $N = 6$ to $N = 12$ for instances with $\Delta < 10^{-2}$ and nearly quadruples if we restrict to cases with $\Delta < 10^{-3}$. Although the relative number of harder instances increases, the average gap improvement continues to expand.

Finally, we now give a brief explanation about why, in general, this catalyst is efficient in enhancing the gap between the ground and the first excited state of the problem Hamiltonian in most of the MWIS problem instances analysed. In [50] we discuss how an informed choice of the catalyst may have a huge impact on the minimum energy gap, with the possibility of removing a first-order phase transition in several cases. In particular, it turns out the efficacy of a catalyst strongly depends on how its composition reflects the structure of the ground and first excited states of the problem Hamiltonian in terms of their Hamming distance. Indeed, it turns out that to maximise the effect of the catalyst (which would correspond to a complete removal of the first-order phase transition), one would need to introduce a catalyst which includes all the

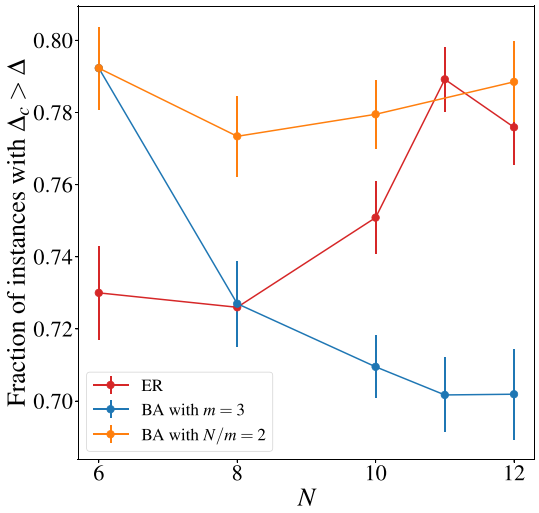

Figure 6. Fraction of potential hard (i.e. with $\Delta < 10^{-1}$) MWIS problem instances on BA and ER graphs which benefit from the use of the catalyst versus N, the total number of nodes. We have performed the analysis for both Erdős–Rényi with $p = 0.5$ (red) and Barabáasi–Albert with $m = 3$ constant (blue) and m variable such that $\mathrm{N\ m^{-1}} = 2$ (orange). For each point in the graph we generated more than ten thousand problem instances. The associated error bars are computed via bootstrapping methods, resampling from the original dataset.

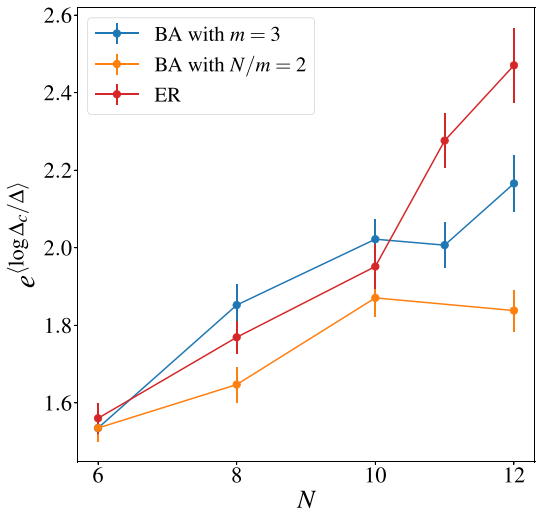

Figure 7. Average catalyst gap improvement computed as the geometric mean of Δ_c/Δ versus total number of vertices in the graph. In computing this average we used over ten thousand MWIS problem realisations for each data point, taking into account instances with $\Delta < 10^{-1}$. The error bars are computed taking into account how the error on $\langle \Delta_c/\Delta \rangle$ propagates.

spin flip operators that connect the ground and the first excited state of H_P. This of course would require knowledge of the structure of the ground state which is as hard as solving the optimisation problem itself. However, one can achieve satisfactory results without any a-priori knowledge of the structure of the global minimum or eventual symmetries of the problems. This is indeed the scenario which is emerging from the analysis conducted in this paper, where the catalyst in equation (10) is able to improve the minimum energy gap in the majority of the MWIS problem realisations, regardless of all the features that the specific instance has. In this context, this means that the catalyst used throughout this whole analysis introduces (on average) a sufficient number of spin flip operators to connect, at least perturbatively, the ground and first excited states in a statistically significant number of MWIS problem cases.

4. Discussion and conclusions

In this paper, we have conducted a study of the enhancement of the minimum energy gap in MWIS problems using a catalysts that attaches an XX-interaction to each edge of the problem graph. Our findings provide robust statistical evidence that the introduction of this catalyst can lead to significant gap improvements in

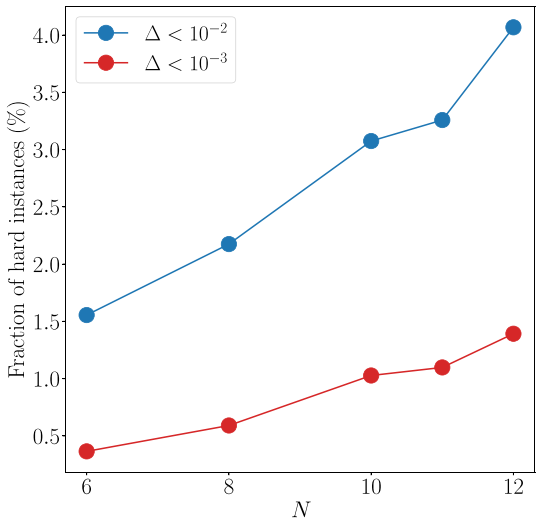

Figure 8. Fraction of hard instances for increasing values of the system size. The blue and red points refer to instances with minimum gap $\Delta < 10^{-2}$ and $\Delta < 10^{-3}$, respectively. Lines connecting points are a guide for the eye.

randomly generated MWIS problem instances. This improvement is notable as it may help reduce the effects of first-order phase transitions, which are a common bottleneck in adiabatic quantum computation.

For our analysis, we randomly generated thousands of MWIS problem instances on two types of graphs: Erdős–Rényi and Barabáasi–Albert, both of which have widespread real-world applications. Firstly, we fixed the total number of nodes in the graphs at $N = 10$. From this analysis, we found that almost 80% of the instances exhibited a noticeable improvement in the gap due to the introduction of the catalyst, regardless of the type of graph—whether scale-free or random—on which the MWIS instance was constructed. Interestingly, this analysis revealed that the catalyst is more effective on specific problem instances that feature smaller minimum gaps, Δ. Specifically, the catalyst demonstrates greater efficacy on instances with potential first-order phase transitions, characterised by gaps $\Delta \lesssim 10^{-2}$. For these harder instances, the average improvement ratio, Δ_c/Δ, shows a significant increase, suggesting that the catalyst specifically enhances performance on these more challenging cases. Conversely, instances with larger minimum gaps do not experience the same degree of improvement. This selective targeting of more complex instances highlights the potential of this catalyst as a powerful tool for addressing hard-to-solve cases, where the presence of vary small gaps plays a critical role in determining performance outcomes.

To cement our results, we have used simulation of real-time annealing aiming to asses the effectiveness of the catalyst in reducing the total anneal time T required to reach the ground state with sufficient accuracy. In particular, we tested the validity of our findings by checking whether a gap enhancement actually corresponded to an improvement in terms of total anneal time T. Remarkably, when the catalyst is used, we observe several orders of magnitude improvement in the final ground state fidelity for fixed total anneal time $T = 5\,\mu$s.

Let us now make some comments on how the stoquastic catalyst compares with its non-stoquastic counterpart. Although most of the analysis has been carried out with the stoquastic version, in appendix we analysed the non-stoquastic catalyst in terms of its efficacy in improving the energy gap. Comparing the two catalyst with respect to this aspect, we conclude that our findings underscore the potential of both of them in enhancing the minimum energy gap for hard instances of the MWIS problem, i.e. instances where the minimum gap is $\Delta < 10^{-3}$. Indeed, if we limit our analysis to those instances, it is clear that both catalysts have similar performances, making it difficult to categorically state which one is the most effective. On the other hand, our analysis on MWIS problems indicates that differences in performance emerge when Δ approaches values on the order of 10^{-2}, i.e. where the problems should become in principle easier to solve. In this regime, the non-stoquastic catalyst seems somehow counterproductive as it leads to a shrinkage of the minimum gap in relatively easy problems. Since we did not extend this comparative study to other aspects (e.g. final ground state fidelity, scaling properties,...), we do not have sufficient evidence to clearly determine which catalyst is the most suitable in this context.

Using the stoquastic version of the catalyst, we then examined how the its efficacy scales with problem size, ranging from $N = 6$ to $N = 12$ nodes. The results show that the catalyst consistently improves the minimum energy gap, leading to $\Delta_c > \Delta$ from 70% to almost 80% of the cases, depending on the class of the MWIS graph. Due to limited computational capabilities, it has not been possible to investigate the regions

beyond $N = 12$. Furthermore, it is not even feasible to reliably extrapolate the scaling behaviours for larger N because of the presence of substantial statistical fluctuations. These fluctuations can introduce large oscillations in the data, obscuring any clear trends or patterns, and making it difficult to draw meaningful conclusions about the underlying scaling behaviour based solely on the available information.

By contrast, analysing the average gap improvement $\exp\langle \log \Delta_c/\Delta \rangle$, we observe a notable trend: this quantity consistently increases as the problem size grows. What is particularly striking is that this behaviour holds even as the fraction of more challenging problems rises significantly with increasing N, as we have discussed in the previous section. This suggests a robust scaling property that persists despite the increasing difficulty associated with larger problem sizes. Although our current computational capabilities limited the exploration to graphs with up to $N = 12$ nodes, the observed trends suggest that a similar improvement can be achieved for larger graphs.

Finally, let us conclude with a few comments. First of all, we remark that this setup precluded any possible meaningful way of size-tunability while keeping the phase transition fixed during the process. Hence, it was not possible to determine to what extent this catalyst could be potentially effective against exponentially closing gaps. However, perturbative arguments [48] show that, in these setups, the gap closes exponentially with the number of spin flips required to map the first excited state to the ground state (i.e. the Hamming distance between the two states). This observation agrees with what we find in figure 8, where the relative number of harder instances increases with the system size, which reflects the fact that when the number of nodes increases, the system has more spins available to widen the Hamming distance between the ground and first excited state. Despite this, we observe that the catalyst is consistent in enhancing the minimum energy gap showing an overall improvement in performance even when the fraction of harder instances increases. For these reasons, we believe that this catalyst could potentially be useful in mitigating the effects of first-order phase transitions in MWIS setups and possibly beyond this specific problem. This is definitely worth exploring more closely in future studies.

Data availability statement

The data that support the findings of this study will be openly available following an embargo on GitHub. Data will be available from 01 January 2026.

Acknowledgments

LAN, RG, SB and PAW are supported by EPSRC Grant EP/Y004590/1 'MACON-QC'.

Appendix. Results using a non-stoquastic XX-catalyst

In figure A1, we present the analogous results to figures 2 and 3, but for the non-stoquastic version of the catalyst defined in equation (10). For this analysis, we set $J_c = 1$ and selected $p = 0.4, 0.6$ for Erdős–Rényi graphs and $m = 4, 5$ for Barabási–Albert networks, both with $N = 10$, as representative examples. In this scenario, we observed that approximately 90% of the data points fall below the red line, $\Delta_c = \Delta$, indicating that the majority of instances experience a reduction in the minimum energy gap after applying the catalyst. This suggests that, unlike the stoquastic version, the non-stoquastic catalyst tends to reduce the minimum gap for most instances.

Interestingly, choosing other positive values of J_c did not result in substantial changes compared to the case shown here. Despite the overall gap shrinkage, we note that harder problem instances, specifically those with $\Delta \lesssim 10^{-3}$, still exhibit an improvement in their minimum gap, comparable with the behaviour seen in the stoquastic version. On the other hand, a notable feature in all cases is the presence of a tail of instances originating from the top-right corner of each figure and extending into the region where $\Delta_c < \Delta$. This is further confirmed by the box plots, which display a pronounced tail corresponding to data points in the ranges $10^{-2} \div 10^{-1}$ and $10^{-1} \div 10^{0}$. These instances show significant reductions in the gap ratio, with Δ_c/Δ values dropping as low as 10^{-3}.

Although the non-stoquastic catalyst provides comparable improvements for harder instances with $\Delta < 10^{-3}$, it becomes disadvantageous when the minimum gap approaches the order of 10^{-2}. In these cases, it leads to a significant reduction in the gap ratio, Δ_c/Δ. However, this is limited to cases where the gap is not too small, making it reasonable to exclude those cases from the analysis.

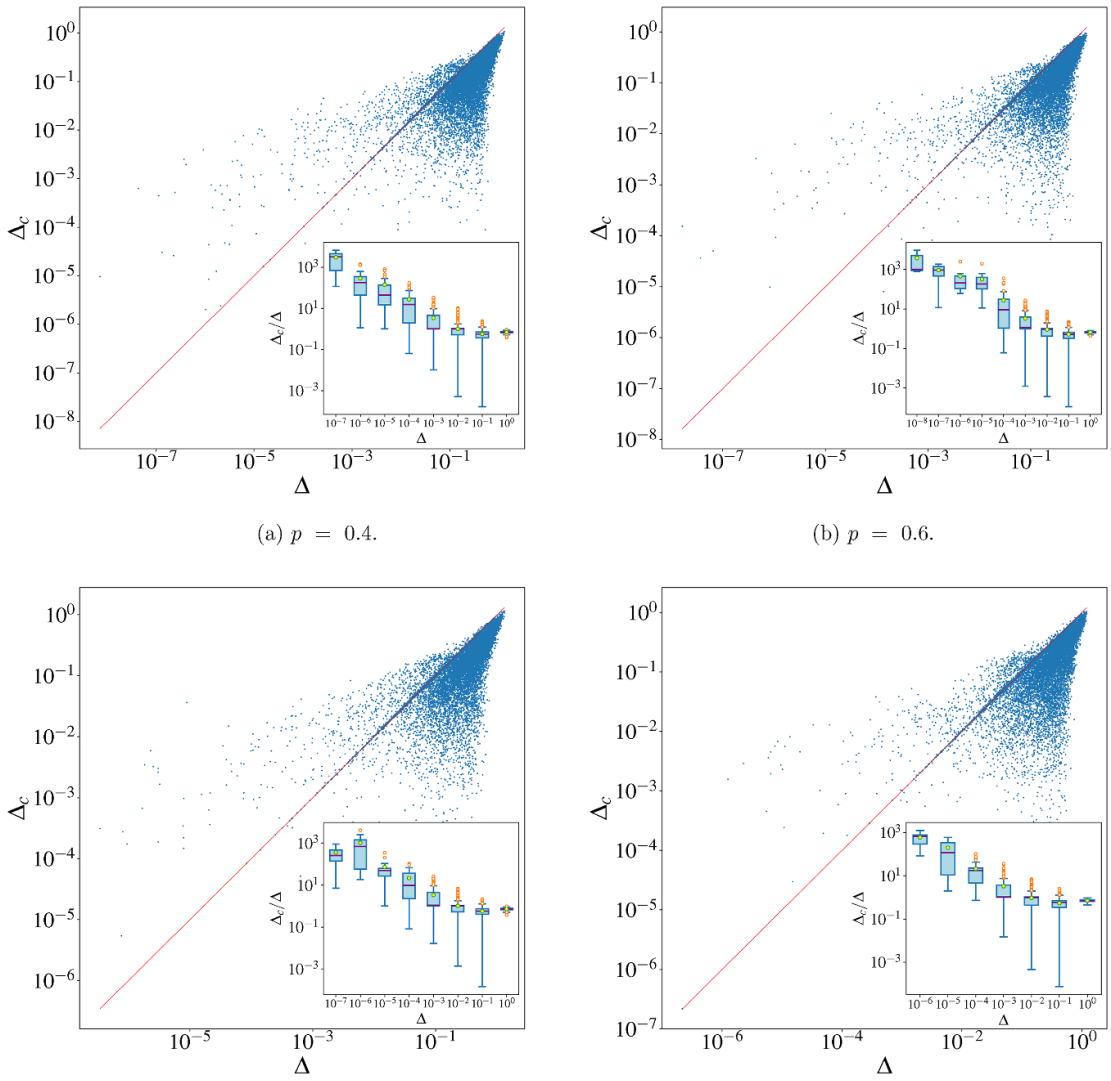

(a) $p = 0.4$. (b) $p = 0.6$.

(c) $m = 4$. (d) $m = 5$.

Figure A1. Comparison of the minimum energy gap using a non-stoquastic version of the catalyst (setting $J_c = 1$) with the minimum energy gap in the catalyst free case on randomly generated MWIS problem instances constructed on both Erdős–Rényi, in (a) and (b), and Barabási–Albert graphs, in (c) and (d). While harder instances (those with $\Delta < 10^{-3}$) still benefit from a gap improvement comparable to what is achieved with the stoquastic version, it is important to highlight a different trend for instances with moderately larger minimum gaps. These instances, particularly those with gaps $\Delta \gtrsim 10^{-2}$, experience a reduction in the gap due to the catalyst, as is evident from the insets. This suggests that although the catalyst is still effective for more difficult problems, it can negatively impact the performance on relatively easier instances by shrinking their gaps. As a consequence, more than 90% of the MWIS realisations lie in the region below the red line, contrary to what happens in the stoquastic case.

ORCID iD

Luca A Nutricati 0000-0002-5045-5113

References

[1] Levinthal C 1969 *Mossbauer Spectroscopy in Biological Systems: Proceedings of a Meeting Held at Allerton House* (*Monticello, Illinois*)
[2] Aspuru-Guzik A, Dutoi A, Love P and Head-Gordon M 2005 *Science* **309** 1704–7
[3] Lanyon B *et al* 2010 *Nat. Chem.* **2** 106–11
[4] Elfving V E, Broer B W, Webber M, Gavartin J, Halls M D, Lorton K P and Bochevarov A 2020 How will quantum computers provide an industrially relevant computational advantage in quantum chemistry? (arXiv:2009.12472)
[5] Moody J, Wu L, Liao Y and Saffell M 1998 *J. Forecast* **17** 441–70
[6] McNeil A, Frey R and Embrechts P 2015 *Quantitative Risk Management: Concepts, Techniques and Tools: Revised edition* pp 1–699
[7] Tsantekidis A, Passalis N, Tefas A, Kanniainen J, Gabbouj M and Iosifidis A 2017 Using deep learning to detect price change indications in financial markets pp 2511–5
[8] Allanach B C, Grellscheid D and Quevedo F 2004 *J. High Energy Phys.* JHEP07(2004)069
[9] Douglas M R and Taylor W 2007 *J. High Energy Phys.* JHEP01(2007)031
[10] Abel S and Rizos J 2014 *J. High Energy Phys.* JHEP08(2014)010

[11] Halverson J and Ruehle F 2019 *Phys. Rev. D* **99** 046015
[12] Cole A, Schachner A and Shiu G 2019 *J. High Energy Phys.* JHEP11(2019)045
[13] Halverson J, Nelson B and Ruehle F 2019 *J. High Energy Phys.* JHEP06(2019)003
[14] Ruehle F 2020 *Phys. Rep.* **839** 1–117
[15] Larfors M and Schneider R 2020 *Fortschr. Phys.* **68** 2000034
[16] Abel S, Constantin A, Harvey T R and Lukas A 2022 *Fortschr. Phys.* **70** 2200034
[17] Abel S A, Constantin A, Harvey T R, Lukas A and Nutricati L A 2024 *Fortschr. Phys.* **72** 2300260
[18] Abel S A, Nutricati L A and Rizos J 2023 *Fortschr. Phys.* **71** 2300167
[19] Barahona F, Grötschel M, Jünger M and Reinelt G 1988 *Oper. Res.* **36** 493–513
[20] Ramaswami R, Sivarajan K and Sasaki G 2009 *Optical Networks: A Practical Perspective Morgan Kaufmann Series in Networking* (Elsevier)
[21] Kadowaki T and Nishimori H 1998 *Phys. Rev. E* **58** 5355–63
[22] Ray P, Chakrabarti B K and Chakrabarti A 1989 *Phys. Rev. B* **39** 11828–32
[23] Das A and Chakrabarti B K 2008 *Rev. Mod. Phys.* **80** 1061–81
[24] Albash T and Lidar D A 2018 *Rev. Mod. Phys.* **90** 015002
[25] Jörg T, Krzakala F, Kurchan J and Maggs A C 2010 *Prog. Theor. Phys. Suppl.* **184** 290–303
[26] Jörg T, Krzakala F, Kurchan J, Maggs A C and Pujos J 2010 *Europhys. Lett.* **89** 40004
[27] Knysh S 2016 *Nat. Commun.* **7** 1–9
[28] Seki Y and Nishimori H 2012 *Phys. Rev. E* **85** 051112
[29] Seoane B and Nishimori H 2012 *J. Phys. A: Math. Theor.* **45** 435301
[30] Nishimori H and Takada K 2017 *Front. ICT* **4** 2
[31] Albash T 2019 *Phys. Rev. A* **99** 042334
[32] Takada K, Sota S, Yunoki S, Pokharel B, Nishimori H and Lidar D A 2021 *Phys. Rev. Res.* **3** 043013
[33] Albash T and Kowalsky M 2021 *Phys. Rev. A* **103** 022608
[34] Feinstein N, Fry-Bouriaux L, Bose S and Warburton P A 2022 arXiv:2203.06779
[35] Feinstein N, Shalashilin I, Bose S and Warburton P 2024 arXiv:2402.13811
[36] Crosson E, Albash T, Hen I and Young A P 2020 *Quantum* **4** 334
[37] Takada K, Yamashiro Y and Nishimori H 2020 *J. Phys. Soc. Japan* **89** 044001
[38] Zhao Z, Verma G, Rao C, Swami A and Segarra S 2021 Distributed scheduling using graph neural networks *ICASSP 2021—2021 IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)* pp 4720–4
[39] Arkin E M and Silverberg E B 1987 *Discrete Appl. Math.* **18** 1–8
[40] De Vries S and Vohra R V 2003 *Infor. J. Comput.* **15** 284–309
[41] Bafna V, Narayanan B and Ravi R 1996 *Discrete Appl. Math.* **71** 41–53
[42] Li N and Latecki L 2012 Clustering aggregation as maximum-weight independent set *Advances in Neural Information Processing Systems* vol 25, ed F Pereira, C Burges, L Bottou and K Weinberger (Curran Associates, Inc.) (available at: https://proceedings.neurips.cc/paper_files/paper/2012/file/01386bd6d8e091c2ab4c7c7de644d37b-paper.pdf)
[43] Choi V 2008 Minor-embedding in adiabatic quantum computation: I. The parameter setting problem (arXiv:0804.4884)
[44] Choi V 2010 arXiv:1004.2226
[45] Bapst V, Foini L, Krzakala F, Semerjian G and Zamponi F 2013 *Phys. Rep.* **523** 127–205
[46] Altshuler B, Krovi H and Roland J 2010 *Proc. Natl Acad. Sci.* **107** 12446–50
[47] Farhi E, Goldstone J, Gutmann S and Sipser M 2000 arXiv:quant-ph/0001106
[48] Amin M H S and Choi V 2009 *Phys. Rev. A* **80** 062326
[49] Laumann C R, Moessner R, Scardicchio A and Sondhi S L 2012 *Phys. Rev. Lett.* **109** 030502
[50] Ghosh R, Nutricati L A, Feinstein N, Warburton P A and Bose S 2024 arXiv:2409.13029
[51] Crosson E, Farhi E, Lin C Y Y, Lin H H and Shor P 2014 arXiv:1401.7320
[52] Seki Y and Nishimori H 2015 *J. Phys. A: Math. Theor.* **48** 335301
[53] Hormozi L, Brown E W, Carleo G and Troyer M 2017 *Phys. Rev. B* **95** 184416
[54] Redner S 1998 *Eur. Phys. J. B* **4** 131–4
[55] Albert R, Jeong H and Arabasi A L 1999 (available at: Diameter of the World-Wide Web)
[56] Baraba-si A-L and Albert R 1999 *Science* **286** 509–12