Effects of XX-catalysts on quantum annealing spectra with perturbative crossings

Natasha Feinstein* and Louis Fry-Bouriaux
London Centre for Nanotechnology, University College London, WC1H 0AH London, UK

Sougato Bose
Department of Physics and Astronomy, University College London, London WC1E 6BT, UK

P A Warburton
London Centre for Nanotechnology, University College London, WC1H 0AH London, UK and
Department of Electronic & Electrical Engineering,
University College London, WC1E 7JE London, UK

(Dated: October 24, 2024)

In adiabatic quantum computation the required run-time to reach a given ground-state fidelity is dictated by the size of the minimum gap that appears between the ground and first excited state in the annealing spectrum. In general the presence of avoided level crossings demands an exponential increase in the annealing time with the system size which has consequences both for the efficiency of the algorithm and the required qubit coherence times. One promising avenue being explored to produce more favourable gap scaling is the introduction of non-stoquastic XX-couplings in the form of a catalyst - of particular interest are catalysts which utilise accessible information about the optimisation problem in their construction. Here we show extreme sensitivity of the effect of an XX-catalyst to subtle changes in the encoding of the optimisation problem. In particular, we observe that a targeted catalyst containing a single coupling can significantly reduce the gap closing with system size at an avoided level crossing. For slightly different encodings of the same problems however, these same catalysts result in closing gaps in the annealing spectrum. To understand the origin of these closing gaps, we study how the evolution of the ground-state vector is altered by the presence of the catalyst and find that the negative components of the ground-state vector are key to understanding the response of the gap spectrum. We also consider how and when these closing gaps could be utilised in diabatic quantum annealing protocols - a promising alternative to adiabatic quantum annealing in which transitions to higher energy levels are exploited to reduce the run time of the algorithm.

I. INTRODUCTION

Quantum annealing (QA) is a continuous time quantum algorithm proposed as a means of solving NP combinatorial optimisation problems faster than can be achieved classically [1–3]. Being able to solve these kinds of problems has real world applications ranging from portfolio optimisation [4] and resource allocation [5–7] to transport optimisation [8]. In these applications it is of great interest to be able to find the best solution in the shortest possible time.

QA proceeds by initialising a quantum system in the ground-state (GS) of a local transverse field, and evolving to a Hamiltonian whose ground-state encodes the optimal solution of the problem to be solved. If the evolution proceeds adiabatically, the ground-state of the final Hamiltonian is obtained with high probability at the end of the anneal. How fast the anneal can be carried out while keeping the system in its ground-state is determined by the adiabatic theorem [9]. In its simplest form, this states that the total evolution duration must be inversely proportional to the square of the minimum energy gap between the ground and first excited states in the energy spectrum of the total Hamiltonian. The scaling with the problem size of the annealing time required to maintain adiabatic evolution thus depends on scaling of this gap minimum.

Early results that focused on randomly generated instances of the exact-cover problem for small system sizes suggested that the minimum energy gap scaling may be polynomial [10], which would allow QA to find the solution efficiently. However subsequent studies have suggested that we can generally expect exponentially closing gaps in the annealing spectrum resulting from the presence of first order phase transitions [11–15]. A particular problem which has been highlighted is the potential for so called perturbative crossings to form between the low energy eigenstates of the problem Hamiltonian towards the end of the anneal [16, 17]. The corresponding energy gap has been found to be exponentially small with the Hamming distance between the eigenstates, which can generally be expected to grow with the system size [16]. Additionally, it was argued in [17] that the probability of no such crossing forming would decrease exponentially with the number of states that had energy comparable to that of the ground-state, suggesting that the likelihood of not encountering such a transition vanishes with increasing problem size.

Overall, these results confirm the standard assumption that quantum algorithms such as QA are unlikely to be

* ucapnjs@ucl.ac.uk

able to change the complexity class of NP-hard optimisation problems [18]. What is more likely is that they will offer a quantitative improvement to the run-time, even if the time scaling remains exponential with the problem size. Indeed, there have already been some promising results in this regard, with QA showing a quadratic improvement in time scaling over simulated annealing for finding the ground-state of random Ising chains [19] as well as solving the Max Independent Set (MIS) problem [20]. However, even if we are not expecting QA to be a polynomial time algorithm, the exponential scaling in the annealing time still presents a problem in that it demands exponentially long qubit coherence times. Coherent quantum dynamics were recently demonstrated on the D-Wave Advantage hardware [21]. Significant effects of decoherence were however seen for annealing times greater than a few nanoseconds. Other platforms may have longer coherence times than this [20]. Without the error correcting framework that exists for gate based quantum computation, however, exponentially long annealing runs remain intractable.

It is therefore desirable to move the exponential scaling from the anneal time to another part of the algorithm such that, even if the algorithm as a whole has an exponential total run time, the length of each annealing run is sub-exponential and therefore more likely to be within the coherence lifetime of the hardware. The most straightforward way to do this is to carry out an exponential number of anneals with sub-exponential run-time. In this way, one can still achieve a high probability of measuring the GS overall - even if the final ground-state overlap at the end of each individual anneal is exponentially small as a result of being far from the adiabatic limit. One way to improve upon this basic approach would be to find ways to utilise information obtained from each annealing run to alter the annealing path in some targeted way in order to enhance the final GS overlap of the subsequent anneals.

A number of works exist in the literature which examine how different changes to the annealing path affect the gap spectrum, and therefore the final GS overlap. These studies typically focus on the inhomogeneous driving of the transverse field [22–28] or the introduction of an additional non-monotonic term to the Hamiltonian referred to as a catalyst [14, 25, 29–36]. In general, the goal has been to enhance the size of the minimum gap in the annealing spectrum [14, 22–25, 27, 29, 31–33, 36]. However, a higher GS fidelity can also be obtained by manipulating the spectrum such that a diabatic anneal is possible [26, 31, 34]. In this case the goal is to produce a path through the annealing energy spectrum that exploits transitions to higher energy states such that the system subsequently returns to the ground-state. When this can be achieved, the small energy gaps at which these transitions happen no longer present bottlenecks to the algorithm, and if all other gaps in the spectrum close no faster than polynomially with system size, a polynomial QA run will end with a high overlap with the ground-state of the problem Hamiltonian. This exploitation of transitions has been shown to result in exponential speedup for the glued trees problem whose structure results in a symmetric annealing spectrum that is naturally amenable to diabatic QA [37].

For the most part, these studies have not attempted to tailor the path change to the specific problem instance but rather have been studying the effectiveness of different drivers and catalysts in different settings. However there are some examples in the literature of where problem specific information is utilised. It was shown in [23] how knowledge of the local optima could be used to adjust the relative strengths of the local X-fields to remove or weaken perturbative crossings. The same authors worked this idea into a recursive strategy that utilised the fact that sub-exponential annealing runs were likely to return the local optima responsible for the formation of the perturbative crossings in the spectrum. A more recent work [34] showed that by coupling local optima to each other, one can replace a single avoided level crossing with a correlated double avoided level crossing - thus facilitating a diabatic anneal. Further to this, diagonal catalysts that bias towards the GS have been found to result in significant gap enhancement [35]. However, as the authors note, knowledge of the local optima does not always give a good indication of where in the state space the global optimum might be.

There has been significant interest in understanding if the introduction of XX-couplings could result in more favourable gap scaling. In particular, there has been a desire to explore the use of XX-couplings that result in annealing Hamiltonians that are non-stoquastic. Non-stoquastic Hamiltonians are those which contain positive and/or complex off diagonal elements and have been linked to qualitatively different behaviours to what can be observed in the stoquastic setting [34, 38]. For the p-spin [14, 32, 33] and Hopfield models [30], all-to-all non-stoquastic XX-catalysts have been shown, both analytically and numerically, to reduce the first order phase transition to second order. This is in contrast to more recent results [39] which suggested that non-stoquastic Hamiltonians will generally have smaller gaps than stoquastic ones. Additionally, for the weak-strong cluster problem, it was found that whether stoquastic or non-stoquastic XX-couplings removed the phase transition depended on where the couplings were applied [25]. For the frustrated Ising ladder it was found that neither stoquastic nor non-stoquastic XX-interactions were able to remove the first order phase transition for their choice of couplings [40]. The fact that the introduction of XX-interactions can have such strikingly different results, depending on the system they are being applied to and which specific couplings are used, suggests that it may be possible to utilise problem specific information in the construction of such catalysts. Indeed, some light has been shone on how the effect of particular XX-interactions relates to the couplings they produce between states in the problem Hamiltonian [33, 41].

In this work, we aim to further add to the understanding of how different catalyst Hamiltonians affect the annealing spectrum. Specifically, we examine the effect of a targeted XX-catalyst, inspired by ideas introduced in [33, 34, 41], and chosen to enhance the minimum gap size at a perturbative crossing. Using the maximum weighted independent set (MWIS) problem to construct our annealing instances, we show that the effect of this catalyst is strongly dependent on arbitrary small changes to the parameters of the problem Hamiltonian which alter the severity of the exponential scaling associated with the perturbative crossing that forms in the original spectrum. In particular we observe that, when applied to a setting where the exponential closing of the gap minimum is comparatively milder, the catalyst is able to produce significant gap enhancement, *i.e:* an increase in the gap size from its value when no catalyst is applied. When applied to a setting where the gap scaling associated with the perturbative crossing is stronger however, this same catalyst results in closing gaps in the annealing spectrum. More specifically, we observe that as we increase the strength of the catalyst, we first observe a closing of the gap at the perturbative crossing followed by the formation of an additional gap minimum earlier in the anneal.

We structure this paper as follows. We first introduce our annealing problem setting in Section II. The structure of the annealing Hamiltonian and the catalyst term is described in Section II A. In Section II B we describe our process for creating MWIS problems and its relationship with the theory of perturbative crossings. In Section II C we demonstrate how the MWIS problem parameters can be tuned to produce such a perturbative crossing, and also how they affect the scaling of the minimum energy gap against system size. With the problem setting established we then, in Section III, examine the effects of introducing targeted non-stoquastic XX-catalysts. We describe our choice of catalyst and show, in Section III A, that for parameter settings which result in a perturbative crossing with weaker associated gap scaling, the catalyst is able to drastically reduce the gap closing rate with system size. We then, in Section III B, apply the same catalysts to problems for which the perturbative crossing has a stronger associated gap scaling, and in this setting, we show that the effectiveness of the catalyst in enhancing the minimum gap size is drastically reduced. It is in this setting that we observe the aforementioned closing gaps in the annealing spectrum. In order to shed some light on the source of these closing gaps, as well as why the effect of the catalyst differs in these two settings, we examine how the evolution of the ground-state vector is affected by the presence of the catalyst. We also present results for intermediate parameter settings in Section III C. Finally, in Section IV, we discuss the implications of our results with regards to where we expect these catalysts to be useful, how information from short annealing runs could be utilised in their construction, and also how these catalysts may be able to be extended to more complex and general graph structures.

II. PROBLEM SETTING

We are interested in examining the effects of XX-couplings on perturbative crossings - a type of avoided level crossing (AC) that can form towards the end of the anneal between problem states that are close in energy. In particular, we wish to compare the effects of the same XX-coupling on perturbative crossings with different associated gap scaling. We construct a scalable instance of the MWIS problem that allows us to easily adjust key properties of the problem Hamiltonian that are responsible for the formation of such crossings. The MWIS problem is a natural choice due to (a) the free parameters present in the encoding of the problem into an Ising Hamiltonian and (b) the ease with which we can change the energies of the states which are one bit flip apart from each other. (The importance of this will become clear in Section II B.) The MWIS problem is also NP-complete meaning that any other problem in NP can be mapped on to it [42].

A. Hamiltonian

Our annealing Hamiltonian takes the form

$$H(s) = (1-s)H_d + s(1-s)H_c + sH_p, \tag{1}$$

where H_d, H_c and H_p are independent of s and denote the driver, catalyst and problem Hamiltonians respectively. The dimensionless annealing parameter, s, is varied from 0 to 1 over the course of the anneal such that $H(s)$ evolves from $H(0) = H_d$ to $H(1) = H_p$. We consider the anneal in a static setting and so it is not necessary to specify how s is varied with time except to say that it increases monotonically such that the same part of the annealing spectrum is not crossed more than once. We refer to the eigenstates of the total Hamiltonian as the *instantaneous* eigenstates and denote them and their corresponding energies as

$$H(s)\left|E_a(s)\right\rangle = E_a(s)\left|E_a(s)\right\rangle. \tag{2}$$

The states are labelled starting from $a = 0$ in order of increasing energy. Similarly we denote the *problem* eigenstates and their energies as

$$H_p\left|E_a\right\rangle = E_a\left|E_a\right\rangle. \tag{3}$$

Since the problem Hamiltonian is diagonal in the computational basis, the set of problem states $\{|E_a\rangle\}$ is simply the computational basis with the states labelled by energy. Our choice of driver is the conventional homogeneous local X-field,

$$H_d = -\sum_{i=1}^{n}\sigma_i^x, \tag{4}$$

where n is the total number of qubits and σ_x^i denotes the Pauli-X operator on the ith qubit. Its ground-state is the equal superposition over all computational basis states.

It serves us at this point to introduce the notion of stoquasticity. A stoquastic Hamiltonian, by its simplest definition, is one whose off-diagonal elements are all real and non-positive. The annealing Hamiltonian, as we have introduced it so far, is an example of a stoquastic Hamiltonian. Any Hamiltonian containing a non-zero number of positive and/or complex off-diagonal elements is considered non-stoquastic. There has been much discussion in the literature regarding the importance of non-stoquasticity in QA [34, 38, 43]. This discussion includes its relationship to complexity theory, universal quantum computation, the Monte-Carlo sign problem and the introduction of negative vector components to the instantaneous GS. In some cases, the term "non-stoquastic" has been reserved for Hamiltonians for which there is no possible rotation into a stoquastic form. However, determining whether or not such a rotation exists is, in general, NP-hard [44]. In this work, we use the term "non-stoquastic" simply to describe the off-diagonal elements of a Hamiltonian in the basis we are working in.

Let us now introduce our catalyst Hamiltonian. This Hamiltonian contains a single XX-coupling between two qubits, i and j, and can thus be written,

$$H_c = J_{xx}\sigma_i^x\sigma_j^x. \tag{5}$$

The catalyst strength, J_{xx}, is always chosen to be positive such that the total Hamiltonian is non-stoquastic for $s \neq 0, 1$. Its magnitude can be tuned to adjust the strength of the catalyst relative to H_d and H_p. We discuss the selection of a suitable i and j at the start of section III.

We now go on to describe the motivation for and construction of our problem Hamiltonians.

B. Problem Graph

Our problem instances are constructed to produce annealing spectra with perturbative crossings - a well understood bottleneck in QA that can arise when excited states of the problem Hamiltonian have an energy comparable to that of its ground-state [16, 17]. Their presence can be understood by treating the driver H_d as a perturbation to the problem H_p, and looking at the energy corrections to the problem eigenstates. Rather than performing this perturbative analysis for specific problem instances, this understanding can be reached through looking at the general form of the perturbations to the low energy eigenstates of H_p. We briefly outline this technique before going on to describe our problem instances - more precise descriptions of how these ACs form can be found in [16, 17].

Writing the perturbed energies as $E_a(\lambda)$, we can say that a crossing occurs between two problem states a and b if $E_b(\lambda) < E_a(\lambda)$ $(b > a)$ for some λ. If this λ is small enough that perturbation theory remains valid, this indicates the formation of an AC towards the end of the instantaneous gap spectrum. Because the driver Hamiltonian couples problem states that are one spin-flip apart, the perturbation to the state $|E_a\rangle$ is dependent on the set of states that are within unit Hamming distance of $|E_a\rangle$ - we refer to this set as the *neighbourhood* of $|E_a\rangle$. It can be shown that the lowest energy problem states will experience a reduction in energy as a result of the perturbation as long as the Hamming distance between any pair of these states is greater than 1. The rate of this reduction with respect to λ for each state depends on the energy of the states in its neighbourhood, with lower energies resulting in a greater reduction. Thus a crossing can occur between the ground-state and some other problem state $|E_a\rangle$ that is close in energy if the neighbourhood around the state $|E_a\rangle$ contains lower energy states than that around the ground-state. The value of λ for which the crossing occurs, and thus whether or not an AC forms, will depend on the difference in energy between the two unperturbed problem states, $\Delta E_{0a} = E_a - E_0$, and how different in energy the neighbourhoods around them are.

The top row of Figure 1 shows cartoon problem energy spectra with E_0 and its neighbourhood highlighted in blue and E_1 and its neighbourhood in orange. (For clarity we include arrows showing the driver couplings from $|E_0\rangle$ and $|E_1\rangle$ to their respective neighbourhoods.) Corresponding cartoons illustrating the perturbed energies are shown below with the magnitude of the driver increasing from right to left. Figure 1(a) depicts a setting where $|E_0\rangle$ has a lower energy neighbourhood than $|E_1\rangle$ such that it receives a greater negative perturbation and no crossing forms. Figure 1(b) depicts the case where the neighbourhoods are reversed such that $|E_1\rangle$ receives a greater negative perturbation and the perturbed energies cross. Figure 1(c) shows a similar scenario but with a greater ΔE_{01} such that a higher λ is needed before the energies cross. Finally, in Figure 1(d) ΔE_{01} is the same as in Figure 1(a) but the difference between the energies of their neighbourhoods is decreased such that the difference between the gradients of the perturbed energies are smaller. This also increases the value of λ for which the crossing occurs.

In addition to changing the value of λ for which the crossing occurs, and thus the value of s at which we observe an AC in the annealing spectrum, ΔE_{01} and the neighbourhoods around $|E_0\rangle$ and $|E_1\rangle$ also affect the extent to which the perturbed state vectors $|E_a(\lambda)\rangle$ have evolved away from the problem state vectors $|E_a\rangle$ at the point of the crossing. If the introduction of problem states, $|E_b\rangle \neq |E_a\rangle$, to the perturbed state, $|E_a(\lambda)\rangle$, is minimal, the overlap between the instantaneous ground-state before and after the AC will be very small. However if the perturbed problem states have become more mixed, then the overlap may be larger. The extent of this overlap affects the size of the gap minimum associated with the AC, with a larger overlap leading to a bigger gap [45].

In order to produce an annealing spectrum with a perturbative crossing we thus require a problem setting where E_0 and E_1 are close in energy and where the neigh-

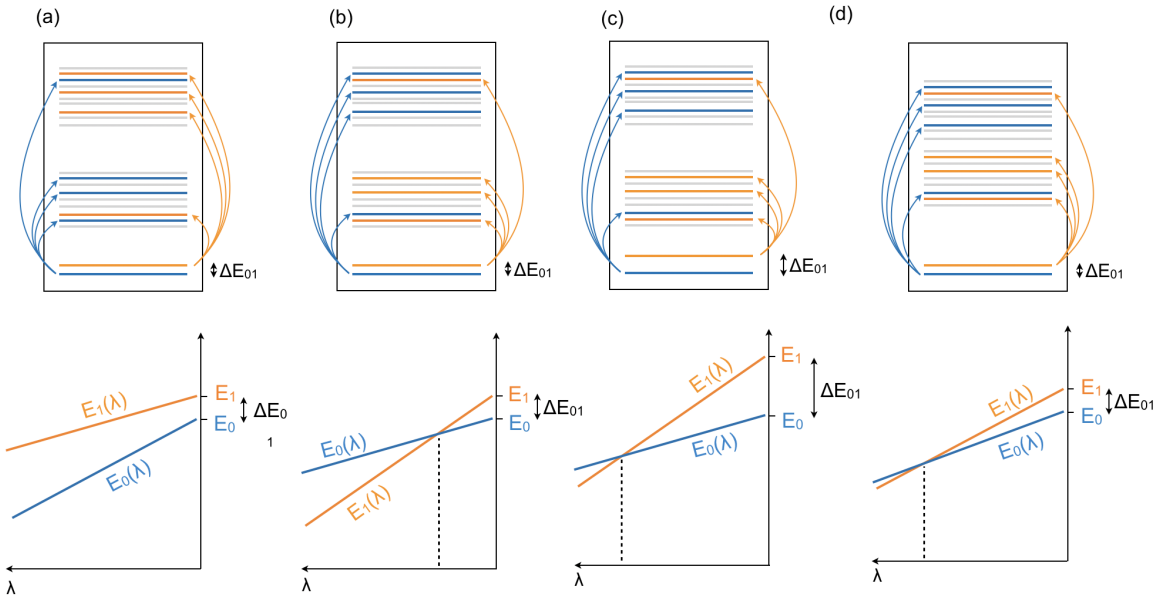

FIG. 1: Top: Cartoon problem energy spectra with the ground and first excited state, as well as their neighbourhoods, highlighted in blue and orange respectively. Bottom: Corresponding cartoons of the perturbed problem energies where the perturbation is the driver Hamiltonian and the perturbative parameter is λ. (a) $|E_0\rangle$ has a lower energy neighbourhood than $|E_1\rangle$ such that no perturbative crossing forms. (b)$|E_1\rangle$ has a lower energy neighbourhood than $|E_0\rangle$ such that a perturbative crossing does form. (c) Similar setting to (b) but with a larger ΔE_{01} such that the crossing happens at larger λ. (d) Similar setting to (b) but where the neighbourhoods are closer in energy, reducing the difference between the gradients of the perturbations, such that the crossing happens at larger λ.

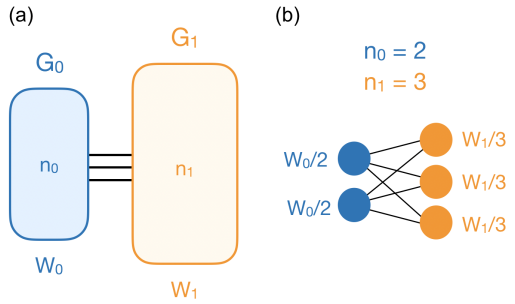

FIG. 2: An illustration of our graph structure, as described in Section II B, is shown in (a). An example with $n_0 = 2$ and $n_1 = 3$ is shown in (b).

bourhood around $|E_1\rangle$ contains lower energy states than that around $|E_0\rangle$. Additionally, if we wish to be able to adjust the gap size associated with the crossing, we need to be able to easily tune ΔE_{01} and the energy differences of the neighbourhoods. To achieve this, we utilise the MWIS problem, which takes as its input an undirected, weighted graph and aims to find the set of vertices with the largest weight for which no two vertices are connected by an edge. The problems we construct, as well as the ways we utilise the free parameters in its Ising Hamiltonian encoding, are inspired by the MWIS instances used in [41]. We now describe how the MWIS problem Hamiltonian is constructed before going on to describe our specific problem setting.

In its Ising formulation, each vertex of the problem graph is represented by a spin. Each basis state then represents a set of vertices, with spin up denoting a vertex that is in the set and spin down denoting a vertex that is not. For instance, the problem state $|E_a\rangle = |\downarrow\uparrow\uparrow\downarrow\uparrow\rangle$ corresponds to the set of vertices $\{2, 3, 5\}$. Note that for this encoding flipping a spin corresponds to either adding or removing a vertex from the set. The vertex weights are implemented with local Z-fields and the independent set condition by introducing an edge penalty. This penalty is achieved by adding an anti-ferromagnetic ZZ-coupling, with a strength J_{zz}, between any two qubits corresponding to vertices connected by an edge. Overall, the problem Hamiltonian is given by

$$H_p = \sum_{i\in\{\text{vertices}\}} (c_i J_{zz} - 2w_i)\sigma_i^z + \sum_{(i,j)\in\{\text{edges}\}} J_{zz}\sigma_i^z\sigma_j^z \tag{6}$$

where c_i is the number of edges connected to vertex i and w_i is the weight on vertex i. The fact that the edge penalty appears in the local field terms is to account for the fact that two neighbouring down spins should not be penalised.

Assuming a high enough J_{zz} is chosen, the energies of the problem states will form clusters based on how many edges are contained in the corresponding set of vertices. The higher the magnitude of J_{zz}, the greater the sepa-

ration between these clusters. Within the clusters, the energies of the states will be ordered by the total weight of the corresponding set, with the highest weighted set having the lowest energy.

A diagram of the scalable MWIS problem used in this work is shown in Figure 2(a). It is a complete bipartite graph with n_0 vertices in sub-graph G_0 and n_1 vertices in sub-graph G_1. As a result, the graph has two maximally independent sets giving the MWIS problem on this graph one global and one local optimum. We allocate a weight to each sub-graph, W_0 and W_1, which we split equally between the vertices. We give each sub-graph a base weight of W and then increase the weight on sub-graph G_0 by δW such that $W_1 = W$ and $W_0 = W_1 + \delta W = W + \delta W$. So long as δW is chosen to be small enough, this results in the set containing all the vertices in G_0 being the highest weighted independent set and the set containing all the vertices in G_1 being the independent set with the second highest weight. An example graph with $n_0 = 2$ and $n_1 = 3$ is shown in Figure 2(b).

The relative sizes of the sub-graphs dictate the energy spectra of the neighbourhoods around $|E_0\rangle$ and $|E_1\rangle$. Recall that flipping one of the spins in $|E_a\rangle$ from down to up corresponds to adding a vertex to the corresponding set and that flipping a spin from up to down corresponds to removing a vertex. This means that the neighbourhood around $|E_0\rangle$ has n_0 neighbours corresponding to independent sets and n_1 neighbours corresponding to dependent sets - and vice versa for $|E_1\rangle$. Recalling further that the states corresponding to independent sets have lower energies than those corresponding to dependent sets, we can say that if $n_1 > n_0$ then $|E_1\rangle$ has more low energy neighbours than $|E_0\rangle$.

Given a particular problem graph, we can use δW and J_{zz} to tune ΔE_{01} and the differences between the energies of the neighbourhoods around $|E_0\rangle$ and $|E_1\rangle$. Roughly speaking, we use δW to adjust ΔE_{01} and the edge penalty, J_{zz}, to change the separation between the clusters of problem energies associated with different numbers of edge violations. In practice however, the way in which we normalise H_p to keep the energy scale consistent means that both parameters have some effect on all the energy gaps in the spectrum. We are interested in how the effect of an XX-catalyst changes depending on the strength of the AC and so we choose our values of δW and J_{zz} to adjust the extent to which the problem state vectors are perturbed before the crossing while keeping the value of s for which the AC occurs for the 5-vertex instance (on which we focus our discussion) unchanged. Further details on our Hamiltonian and how we select our parameters can be found in Appendix A.

C. Spectral Properties

We now present numerical results corresponding to different parameter settings to confirm that we produce the desired AC and that we are able to alter its properties in the way described above. The results that we plot are the energy gaps between the instantaneous ground and first excited state, $\Delta E_{01}(s)$, as well as the evolution of the instantaneous ground-state, $|E_0(s)\rangle$. The evolution of $|E_0(s)\rangle$ we consider in terms of its overlaps with the problem states - in particular with $|E_0\rangle$ and $|E_1\rangle$ since it is between these two states that we expect a crossing. Our results are obtained by numerical diagonalisation of $H(s)$ at different values of s with no catalyst present.

We plot on the top row of figures 3(a)-(c) the energy difference ΔE_{01} as a function of s. On the bottom row of figures 3(a)-(c) we show the corresponding evolution of $\langle E_0(s)|\, E_{0,1}\rangle$. In Figure 3(a) we show the case where $n_0 > n_1$ for a given set of problem parameters, which results in the absence of an AC since $|E_0\rangle$ has a lower energy neighbourhood than $|E_1\rangle$. Figure 3(b) shows the case where $n_1 > n_0$ but with the other problem parameters (δW and J_{zz}) kept the same as in (a). This now results in an AC occurring at $s = 0.9$ since $|E_1\rangle$ now has a lower energy neighbourhood than $|E_0\rangle$. In both (a) and (b), the MWIS problem parameters are $\delta W = 0.01$ and $J_{zz} = 5.33$. Figure 3(c) illustrates the case where $n_1 > n_0$, but the MWIS problem parameters have been increased to $\delta W = 0.37$ and $J_{zz} = 37.5$. This results in an increase of ΔE_{01} and enhances the differences between the energies of the neighbourhoods of $|E_0\rangle$ and $|E_1\rangle$. As in Figure 3(c), we observe a gap minimum at $s = 0.9$ indicative of an AC.

In the bottom plots of Figures 3(b) and (c) we see that the instantaneous ground-state evolves towards $|E_1\rangle$ before there is a sharp exchange at $s = 0.9$, after which $|E_0(s)\rangle$ becomes dominated by $|E_0\rangle$. This exchange, coinciding with the gap minimum, is what we expect to see at a perturbative crossing [46]. This is in contrast to the bottom plot in Figure 3(a) where the instantaneous ground-state is seen to evolve smoothly towards $|E_0\rangle$. Finally, in Figure 3(d), we show the scaling of the gap minimum with system size. For the two cases where an AC is produced we scale the sub-graph sizes as $n_0 = (n-1)/2$ and $n_1 = (n+1)/2$ and for the case where we do not produce an AC these sizes are reversed. We use, for each system size, the same values of δW and J_{zz} as for our 5 vertex examples. We see that the minimum gap size does indeed appear to close exponentially in the cases where $n_0 < n_1$ - further indicating the presence of an AC.

Figures 3(a)-(d) show that the AC can be manipulated in the desired way by changing our problem parameters. First, we are able to control whether or not an AC occurs in the spectrum by changing the relative sizes of n_0 and n_1 as intended. Secondly, we are able to adjust the extent to which the instantaneous states have evolved away from the problem states at the point of the crossing by changing δW and J_{zz}. This alters the scaling of the minimum gap against system size. We observe that as δW and J_{zz} are increased for $n_1 > n_0$, the gap does still appear to close exponentially, but the scaling has been 'weakened', *i.e:* the scaling exponent is reduced. Therefore, in the following, we will refer to our two problem

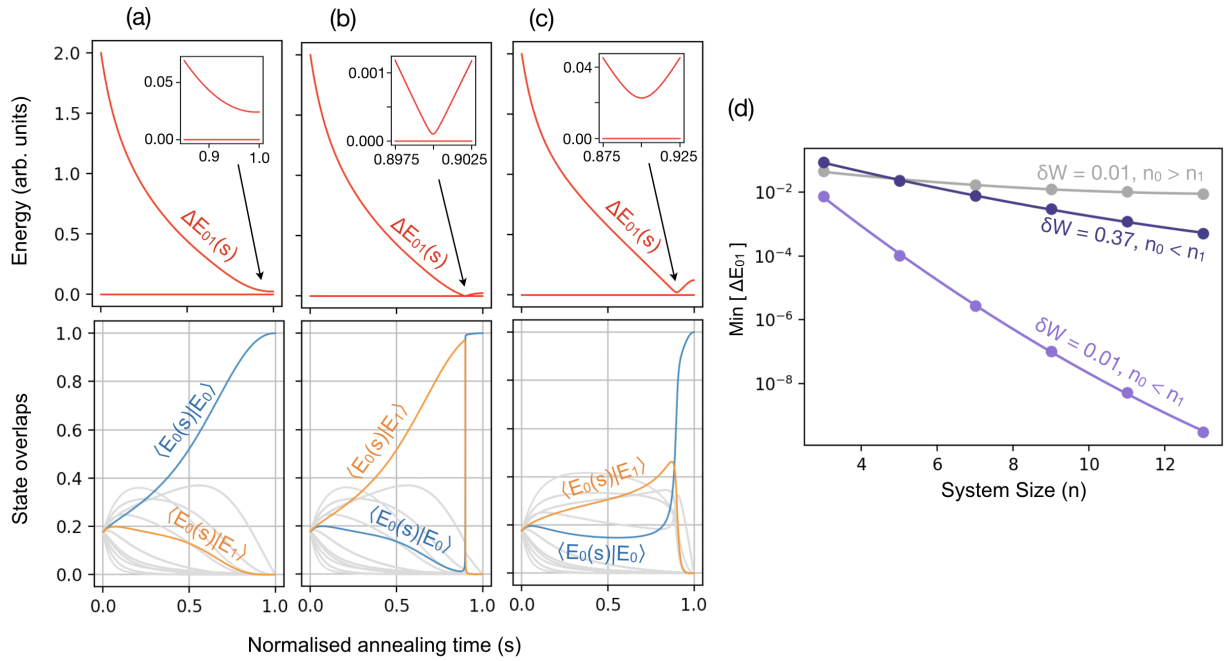

FIG. 3: (a-c) show numerical results for different anneals without the presence of a catalyst. Gap spectra are shown on the top and the evolution of the instantaneous ground-state vectors are shown on the bottom in terms of their overlaps with the problem states - the problem ground and first excited state overlaps are highlighted in blue and orange respectively. The problem parameters are: (a) $n_0 = 3$, $n_1 = 2$, $\delta W = 0.01$, $J_{zz} = 5.33$. (b) $n_0 = 2$, $n_1 = 3$, $\delta W = 0.01$, $J_{zz} = 5.33$ and (c) $n_0 = 2$, $n_1 = 3$, $\delta W = 0.37$, $J_{zz} = 37.5$. (d) shows the minimum gap with increasing system size for three different parameter settings. The grey line corresponds to the parameter settings in (a). The two purple lines correspond to the parameter settings in (b) and (c) - which are associated with the strong- and weak-gap-scaling settings respectively. For the cases where an AC is produced the sub-graph sizes are scaled as $n_0 = (n-1)/2$ and $n_1 = (n+1)/2$. For the case where we do not produce an AC the sub-graph sizes are reversed. Lines in (d) connecting the data points are a guide to the eye.

parameter settings (which produce an AC) as the weak-gap-scaling (WGS) and strong-gap-scaling (SGS) cases. The parameters, δW and J_{zz}, can be varied smoothly between, and indeed beyond, the two settings depicted in Figure 3. These are simply the two settings we use as illustrative examples.

III. XX-CATALYSTS

With our problem setting established, we now turn to the effects of introducing an XX-catalyst into the annealing Hamiltonian. As stated in Section II A, the catalysts used in this work contain a single XX-coupling which is introduced into $H(s)$ with the opposite sign to the driver and a strength of J_{xx} (equation 5). This results in a non-stoquastic $H(s)$ for $s \neq 0, 1$ and as such it is possible for components of the instantaneous GS vector to become negative during the anneal. That the relative signs of these vector components can be key to understanding the structure of quantum annealing spectra was highlighted in [34] - and we will see that they also play a crucial role in our results.

The XX-coupling is applied between a single pair of vertices in G_1. (Permutation symmetry within the sub-graphs means that there is no need to specify which two vertices are selected.) The result of this XX term is to couple $|E_1\rangle$ to a state corresponding to an independent set and $|E_0\rangle$ to a state corresponding to a dependent set - that is, the catalyst couples $|E_1\rangle$ to a state with significantly lower energy than that to which it couples $|E_0\rangle$.

This approach can intuitively be thought of as a counter-approach to giving $|E_0\rangle$ additional low energy neighbours with a *stoquastic* catalyst - *i.e:* one that enters with the same sign as the driver. This stoquastic approach can be understood through the same arguments as those made in Section II B regarding the formation of perturbative crossings and was shown in [41] to enhance the minimum gap size. There is some numerical evidence that taking the reverse approach with a non-stoquastic catalyst (*i.e:* one that has been introduced with the opposite sign to the driver) is able to weaken or remove a perturbative crossing [33, 34] - although in these examples the catalysts are specifically coupling together local optima and so are not a perfect analogy to the work in [41]). [34] also introduces a subtly different perturbative argument to explain why such non-stoquastic catalysts are able to enhance the gap size at an AC. We apply the same perturbative approach to our setting in Appendix B in order to further support our choice of XX-coupling.

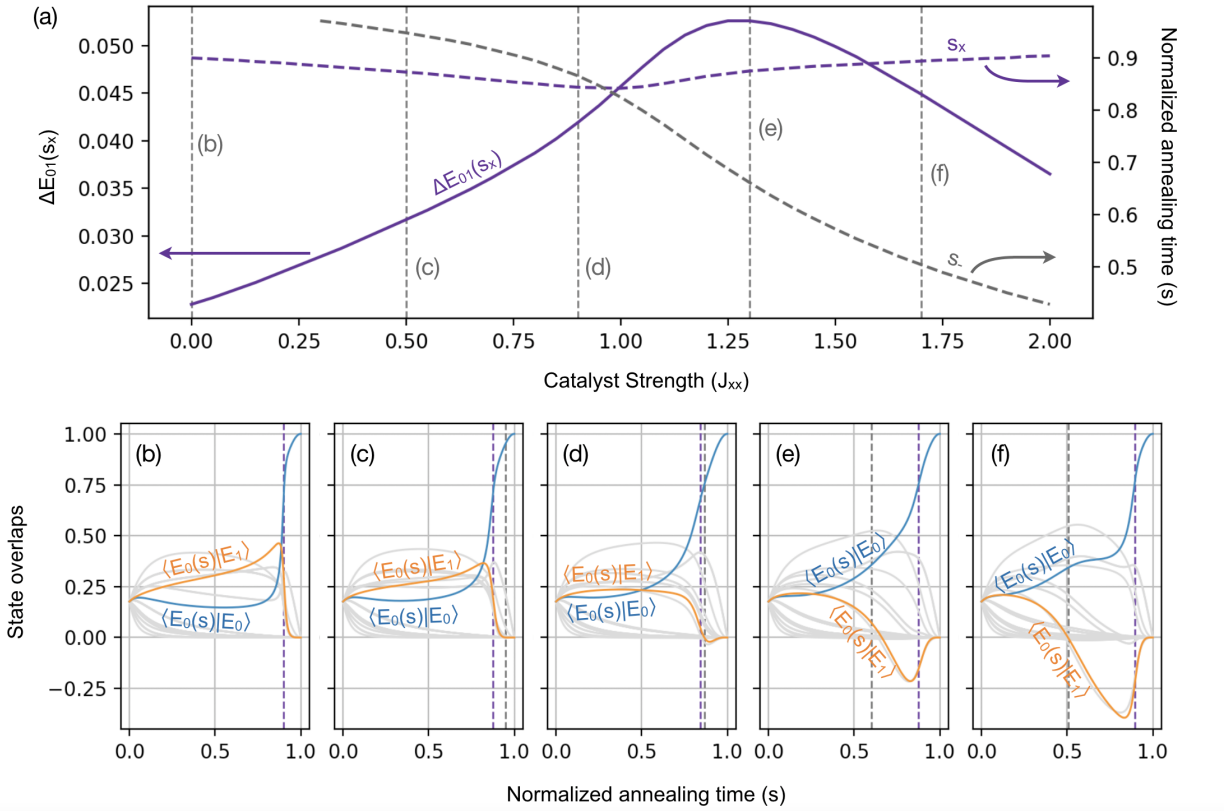

FIG. 4: Numerical results for a problem instance with $n_0 = 2$, $n_1 = 3$, $\delta W = 0.37$ and $J_{zz} = 37.5$ (*i.e:* the parameters associated with the WGS). A catalyst is applied between a single pair of vertices in sub-graph G_1. (a) shows the dependence on catalyst strength of the gap size at the AC, $\Delta E_{01}(s_{\mathrm{x}})$ (solid purple), the location of the minimum gap, s_{x} (dashed purple), and the value of s for which $\langle E_0(s)|E_1\rangle$ becomes negative, s_- (dashed grey). The evolution of the instantaneous ground-state for different catalyst strengths is shown in (b-f). These plots have s_{x} and s_- marked with purple and grey dashed lines respectively. The catalyst strengths for which we show the evolution are marked on (a) with vertical grey dashed lines.

We briefly note that taking the stoquastic approach in our problem setting (*i.e:* introducing a negative XX-coupling into G_0 to give the GS an additional low energy neighbour) results in the same kind of gap enhancement seen in [41]. In practice however, identifying such a coupling assumes knowledge of the global optimum. The catalyst that we examine in this work instead effectively targets a local optimum in which the algorithm may become trapped - such that partial information about the problem may allow for its construction. We will discuss this point further in Section IV.

A. Weak-Gap-Scaling Case

In Figure 4 we present numerical results of varying the XX catalyst strength J_{xx} on the annealing problem described in Section II C, with $n_0 = 2$, $n_1 = 3$, $\delta W = 0.37$ and $J_{zz} = 37.5$. Recall that, in the catalyst-free case, these parameter choices result in the comparatively weak AC shown in Fig 3(c). Figure 4(a) shows the variation of the location, s_{x}, and size, $\Delta E_{01}(s_{\mathrm{x}})$, of the gap minimum with the catalyst strength J_{xx}. Also plotted is the value of s at which $\langle E_0(s)|\, E_1\rangle$ becomes negative, s_-. Figures 4(b)-(f) illustrate the behavior of the coefficients $\langle E_0(s)|\, E_{0,1}\rangle$ at selected values of J_{xx}. It can be seen that $\langle E_0(s)|E_1\rangle$ can become negative at sufficiently large values of J_{xx}. We observe that the introduction of the catalyst results in a clear maximum value of $\Delta E_{01}(s_{\mathrm{x}})$ at $J_{xx} = 1.30$. The existence of an optimal J_{xx}, rather than a monotonic improvement, was also observed in other work [33, 34].

We obtain similar results when increasing the system size. Figure 5(a) compares the gap scaling without the catalyst (black) and the gap scaling when we use a catalyst with J_{xx} optimized for each system size (purple) - the inset shows the optimal values of J_{xx} for each system size. We see that by using the optimal values we are able to greatly reduce the severity of the gap scaling with the optimized scaling appearing to be sub-exponential - although for these system sizes it is difficult to discern what the true scaling behaviour is. We also note that the optimal J_{xx} values appear to quickly tend to a constant value as the system size is increased. As we will go on to discuss in Section IV, however, it is unclear whether or not this will extend to larger system sizes.

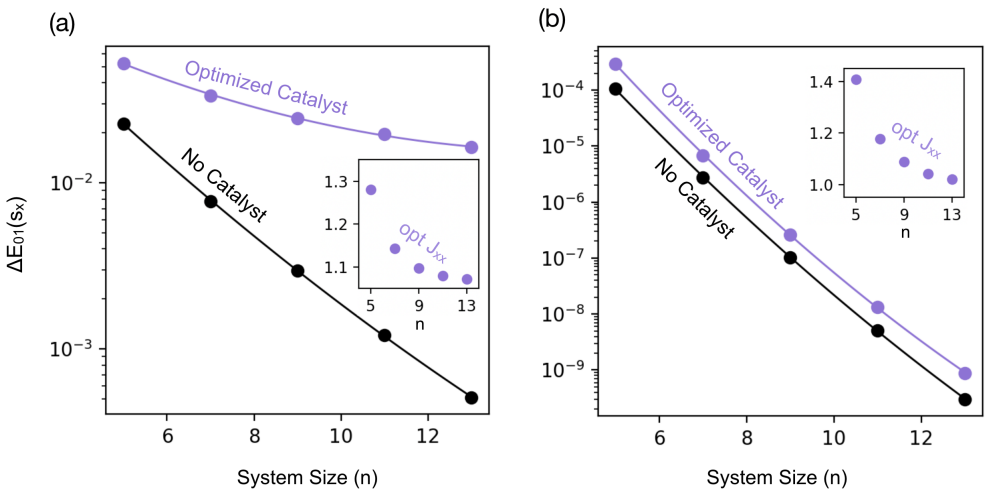

FIG. 5: Results for the size of the gap minimum at the AC with increasing system size. The sub-graph sizes are scaled as $n_0 = (n-1)/2$ and $n_1 = (n+1)/2$. (a) shows results for the WGS setting ($\delta W = 0.37$ and $J_{zz} = 37.5$) and (b) shows results for the SGS setting ($\delta W = 0.01$ and $J_{zz} = 5.33$). In each case we show the gap size without a catalyst in black and with J_{xx} chosen to maximise the gap size in purple. Lines connecting the data points are a guide to the eye. The values of J_{xx} that maximise the gap for each n are plotted in the insets.

Turning our attention to the evolution of the problem state overlaps $\langle E_0(s)|\, E_{0,1}\rangle$ shown in Figures 4(b-f), we observe that changes in $\langle E_0(s)|\, E_{0,1}\rangle$ are overall less significant when $\Delta E_{01}(s_{\rm x})$ is largest. The relation between the rate of change of the instantaneous ground-state and the gap separating it from $|E_1(s)\rangle$ is well understood and expressions explicitly linking them can be found in [47].

B. Strong-Gap-Scaling Case

We now examine the effect of this same catalyst in the SGS setting, *i.e:* for the MWIS problems with $\delta W = 0.01$ and $J_{zz} = 5.33$. Figure 5(b) shows the scaling of the minimum gap in the SGS case as a function of system size with and without the catalyst in purple and black respectively. As in the WGS case, the value of J_{xx} used for each n is chosen to maximise the gap enhancement and is shown in the inset. An important difference when compared to the WGS case is that while we do still observe some gap enhancement at the AC, the catalyst does not significantly alter the scaling with system size. This suggests that, unlike in the WGS case, the use of such an XX catalyst in this situation is of little use for alleviating the exponential drop-off in fidelity due to the exponentially closing gap.

Numerical results for the size of the gap minimum at the AC for the 5-vertex instance at different catalyst strengths are presented in Figure 6(a). The location of this gap minimum is denoted as $s_{\rm x}$ as before. We see that, as for the weaker AC, the gap size does reach a maximum for an optimum catalyst strength. However, we now observe a closing of the gap for $J_{xx} \approx 0.3$. In addition to this we also observe the formation of a new, second local gap minimum earlier in the anneal, at a location we denote as $s_{\rm n}$. The dependence of the size of this new local gap minimum on J_{xx} is shown in Figure 6(b). Note the difference in energy scales between Figures 6(a) and (b). To illustrate the manifestation of this second local gap minimum during an anneal, we plot the evolution of $\Delta E_{01}(s)$ for different catalyst strengths in Figure 7, which clearly shows the occurrence of a double local energy minima in the annealing spectrum. As with the gap minimum at $s_{\rm x}$, we appear to be able to bring this gap at $s_{\rm n}$ arbitrarily close to zero by increasingly fine tuning of J_{xx}, albeit at a different catalyst amplitude. Note that the plots in Figures 6(a)-(b) correspond to exactly the same setting and that our separating the results for the two gap minima into different plots is only for readability of the data.

As in Figure 4(a), Figure 6(a) also shows the location of the gap minimum, $s_{\rm x}$ with a dashed purple line, and the point at which the sign change in ground-state vector components occurs, s_-, with a dashed grey line. As J_{xx} increases, the value of s_- decreases which can be understood by the fact that this change in signs is associated with the non-stoquasticity introduced by the catalyst. This argument can be made more explicit using ideas introduced in [34]. We see that the closing gap is observed for the J_{xx} value at which $s_- = s_{\rm x}$. Figure 6(b) shows the same data but for the additional gap minimum that forms in the strong-AC setting. As with the previously discussed gap minimum, this new gap minimum approaches zero for the J_{xx} value at which $s_- = s_n$. Associating the closing gaps with values of J_{xx} at which the minimum gap locations coincide with s_- is however not sufficient to account for the differing behaviours in the WGS and SGS settings. In Figure 4(a) from the preced-

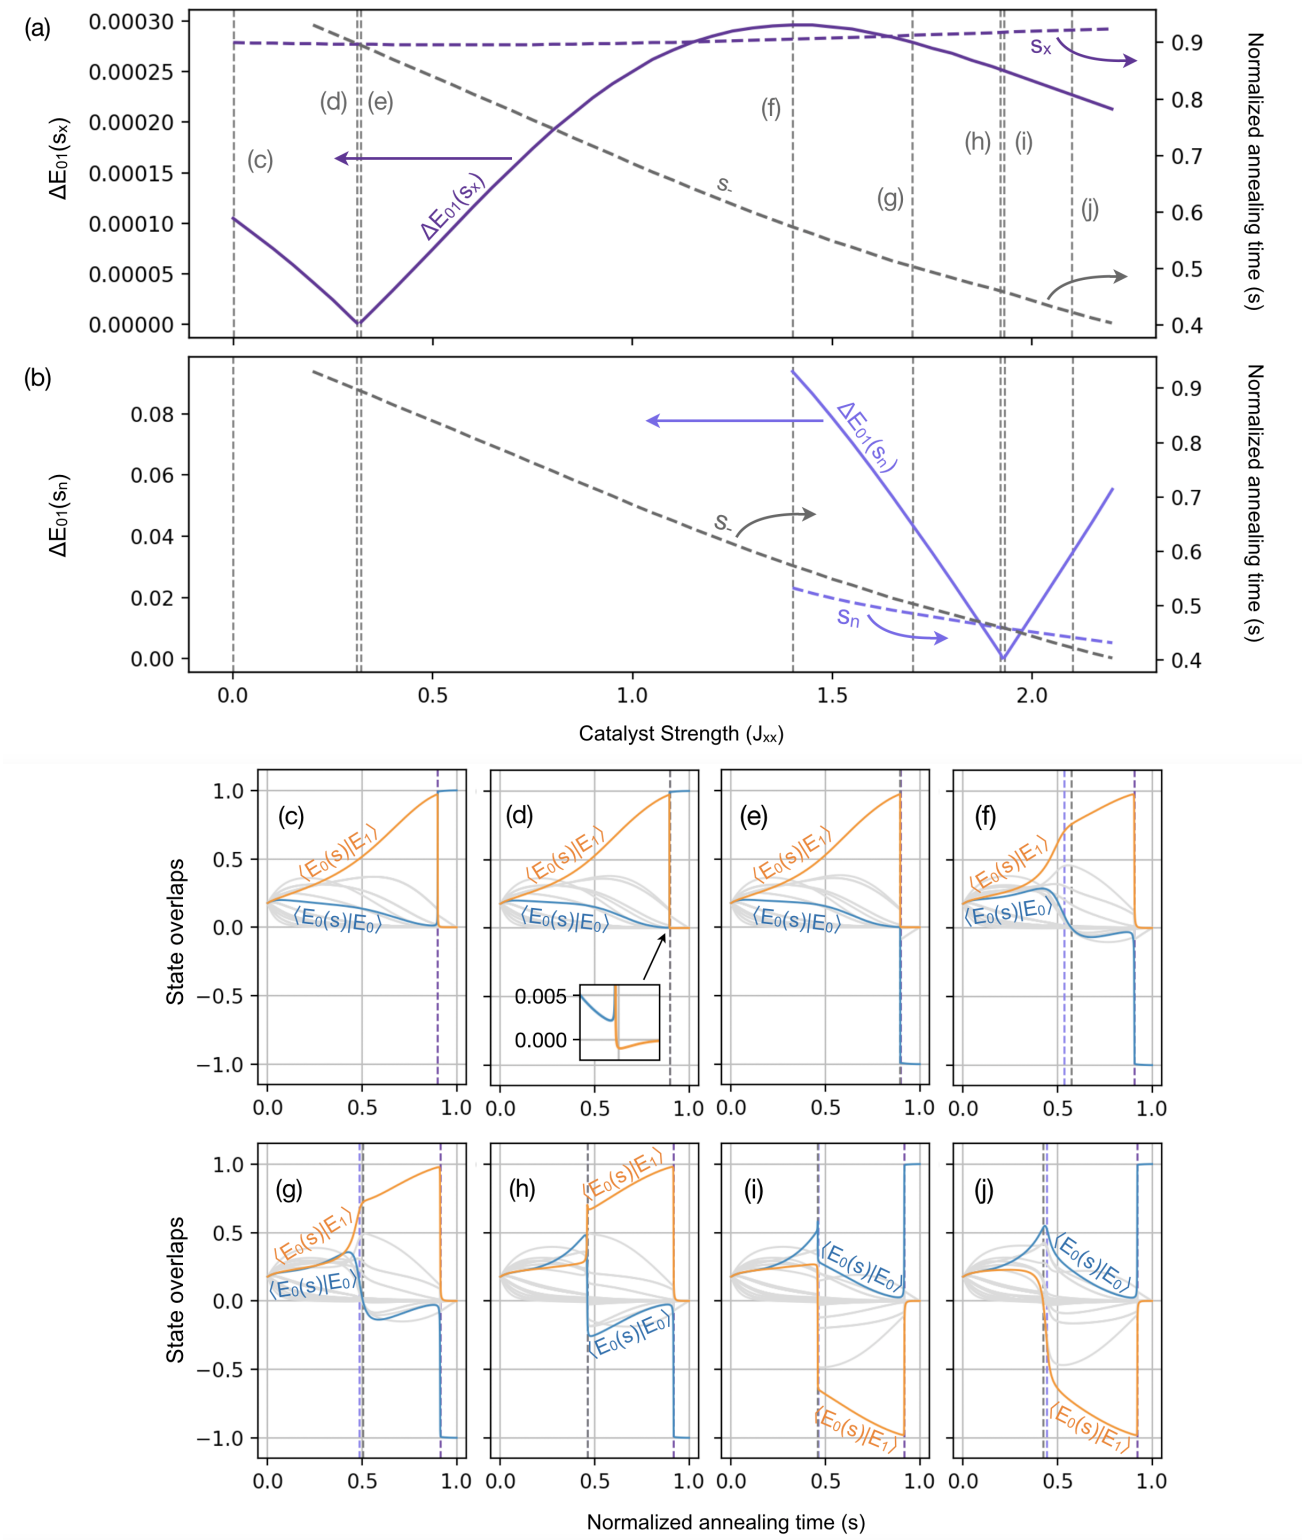

FIG. 6: Numerical results for a problem instance with $n_0 = 2$, $n_1 = 3$, $\delta W = 0.01$ and $J_{zz} = 5.33$ (*i.e:* the parameters associated with the SGS). A catalyst is applied between a single pair of vertices in sub-graph G_1. (a) shows, for increasing catalyst strength, results for the gap size at the AC, $\Delta E_{01}(s_{\rm x})$ (solid purple), the location of the minimum gap, $s_{\rm x}$ (dashed purple), and the value of s for which either $\langle E_0(s)|E_0\rangle$ or $\langle E_0(s)|E_1\rangle$ becomes negative, s_- (dashed grey). The lower plot shows, for the same J_{xx} values, the gap size of the additional minimum gap that forms, $\Delta E_{01}(s_{\rm n})$ (solid lighter purple), the location of this additional minimum, $s_{\rm n}$ (dashed lighter purple) and s_- (dashed grey). The evolution of the instantaneous ground-state for different catalyst strengths is shown in (c-j). These plots have $s_{\rm x}$, $s_{\rm n}$ and s_- marked with the deeper purple, the lighter purple and grey respectively. The catalyst strengths for which we show the evolution are marked on (a) and (b) with vertical grey dashed lines.

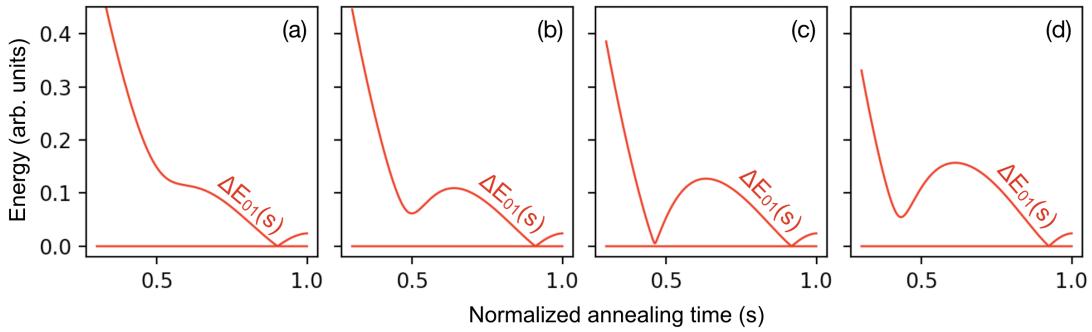

FIG. 7: Gap spectra for four different catalyst strengths showing the formation of the additional minimum gap between the instantaneous ground and first excited state in the SGS setting. The problem parameters are $n_0 = 2$, $n_1 = 3$, $\delta W = 0.01$ and $J_{zz} = 5.33$ and the catalyst is applied between two vertices in sub-graph G_1. The catalyst strengths are J_{xx} = (a) 1.3 (b) 1.6 (c) 1.9 (d) 2.2

ing section, we see that there is also a value of J_{xx} for which $s_- = s_{\mathrm{x}}$ however no closing of the gap at the AC is observed.

To better understand the differences in the behavior between these two regimes we compare how the evolution of the ground-state vector changes with J_{xx}. Similarly to the WGS setting, we show the behavior of the overlaps $\langle E_0(s)|\, E_{0,1}\rangle$ at specific values of J_{xx} in Figures 6(c)-(j). In both the WGS (Figures 4(b)-(f)) and SGS (Figures 6(c)-(j)) cases, we clearly see that the components corresponding to the problem ground and first excited states take opposite signs at some s given sufficiently large J_{xx}. We note however that, in contrast to the WGS case where it is always $\langle E_0(s)|E_1\rangle$ that becomes negative, whether $\langle E_0(s)|E_0\rangle$ or $\langle E_0(s)|E_1\rangle$ takes a negative value in the SGS case depends on the value of J_{xx}.

Given that the relative signs that the vector components take are consistent with varying J_{xx}, it may be tempting to think that there is no physical change occurring. However, as one of the vector components changes sign, the overlap of the instantaneous GS with that problem state will first go to zero before increasing again. As such, a change in which vector component flips its sign will result in a measurable difference in the evolution of a system following the ground state. Note that in plotting the instantaneous vector components, we choose the global sign at each time-step such that any apparent change of sign is indicative of that vector component going to zero in between those two points. With the preceding discussion in mind we make some further notes on the sign behaviour of the GS vector components.

In each plot in Figures 6(c)-(j), s_{x}, s_n and s_- are marked in dark purple, light purple and grey respectively and the J_{xx} values to which these plots correspond are marked in Figure 6(a) with vertical dashed grey lines. Looking at the plots corresponding to $J_{xx} = 0$ and 0.31, Figures 6(c) and (d), we see that as we increase J_{xx} towards the value at which $\Delta E_{01}(s_{\mathrm{x}})$ vanishes, there is a sharpening of the change in the vector components around s_{x}. Then, as s_- passes s_{x}, which happens between Figures 6(d) and (e), there appears to be a discontinuous change in which overlap crosses zero. After this point, the change in $|E_0(s)\rangle$ around s_{x} begins to soften again. Looking at the Figures 6(f)-(j), we see the same sharpening (up to $J_{xx} = 1.92$ in Figure 6(h)) and then weakening of the rate of change in $|E_0(s)\rangle$ around the location of the new gap minimum, s_n. We also observe another shift in which vector component crosses zero as s_- passes s_n; this occurs between Figures 6(h)-(i).

This behaviour is in stark contrast to the WGS case where no distinct changes are observed around the J_{xx} value where $s_- = s_{\mathrm{x}}$. It is unclear why exactly this should be the case. However we can comment on some differences we observe between the two settings. In the WGS case, the AC has already become significantly smoothed out for the value of J_{xx} at which $s_- = s_{\mathrm{x}}$. This is in contrast to the SGS setting where the magnitudes of $|\langle E_0(s)|E_0\rangle|$ and $|\langle E_0(s)|E_1\rangle|$ that are exchanged at the AC are largely unchanged for the J_{xx} value at which $s_- = s_{\mathrm{x}}$. This is at least partially explained by the fact that if the initial AC is weaker, we can expect that a smaller change to the Hamiltonian is required to lift it. We also note that the catalyst strength for which s_- passes s_{x} is over three times as high as it is for the SGS case. It may therefore be possible to lift the AC to a greater extent before reaching the value of J_{xx} for which $s_- = s_x$.

C. Intermediate Regime

We now present data examining intermediate parameter settings to further understand the changing behaviour as we move from the SGS to the WGS case. In Figure 8 we plot the J_{xx} values associated with the two closing gaps for the 5-vertex instance with seven different parameter settings. The x-axis gives the δW used for each parameter setting. Note however that the edge penalty, J_{zz}, is also changing between each data point, with its value in each case chosen such that $s_{\mathrm{x}} = 0.9$ without

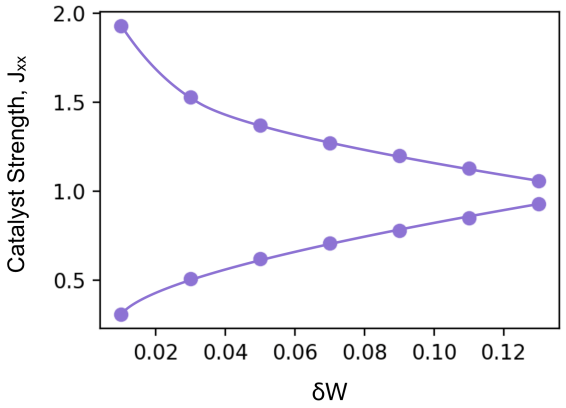

FIG. 8: J_{xx} at which the two gaps go to zero for different parameter settings. The x-axis shows the δW values for each parameter setting. The corresponding J_{zz} value for each case is chosen such that $s_{\mathrm{x}} = 0.9$. Lines connecting the data points are a guide to the eye.

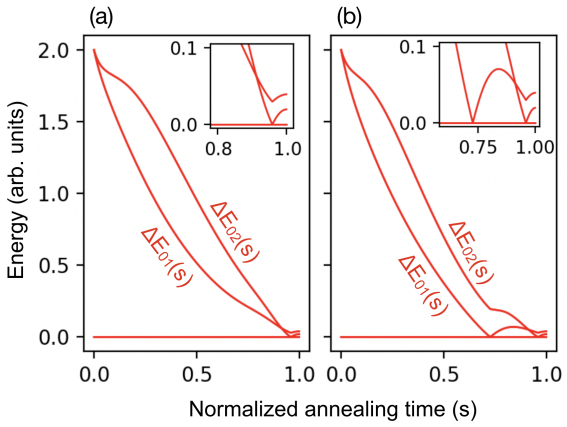

FIG. 9: Gap spectra for an anneal corresponding to a tri-partite MWIS problem with two local optima in addition to the global optimum. The results with and without a catalyst are shown in (b) and (a) respectively. The sub-graph sizes associated with the problem instance are $n_0 = 2$, $n_1 = 4$ and $n_2 = 3$. The weight spacing between the sub-graphs and edge penalty are $\delta W = 0.02$ and $J_{zz} = 5.33$ respectively.

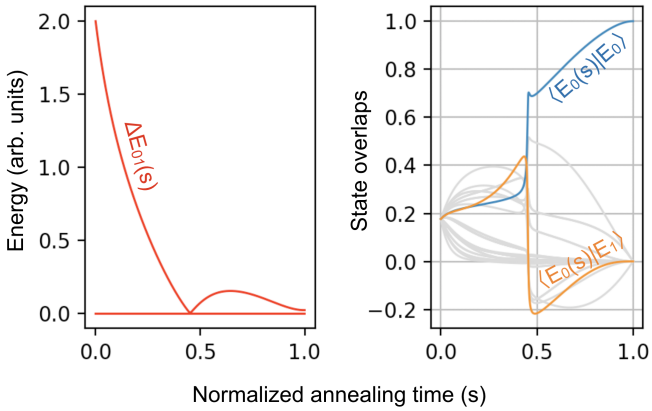

FIG. 10: Gap spectrum and ground-state evolution for a problem instance with $n_0 = 3$, $n_1 = 2$, $\delta W = 0.01$ and $J_{zz} = 5.33$ and a catalyst containing one coupling within G_0. Results for this problem instance without a catalyst can be found in Figure 3(a).

the presence of a catalyst. The size of the gap minimum associated with these parameter settings in the catalyst free case increases from left to right - varying by around one order of magnitude. (See the 5-spin results in Figure 3(d) for an indication of the energy scale.) We observe that the two J_{xx} values approach each other as we adjust the parameters to weaken the AC present in the original annealing spectrum until both values disappear at $\delta W \approx 0.13$. Given that the closing gaps are associated with a change in which vector component changes sign, one interpretation of these results could be that the lack of closing gaps in the weak-AC setting is due to the vanishing of the J_{xx} range for which the vector component corresponding to the problem ground-state crosses zero.

IV. DISCUSSION

An important effect we have observed is that, while the targeted catalyst was able to suppress the severity of the gap scaling in the WGS case, it did not result in the intended suppression to the scaling in the SGS setting. The resulting spectrum in the latter case, however, appears amenable to diabatic quantum annealing (DQA)[26, 34, 48]. Rather than running the algorithm slowly enough that the system remains in the ground-state, the system could be allowed to transition into the first excited state at the first small gap (the new gap minimum produced by the catalyst) and then back into the ground-state at the second (the gap minimum corresponding to the AC present in the original annealing spectrum). Simulations of the dynamics of the 5-qubit problem in a closed-system setting suggest that this is indeed possible and results in a significant reduction of the annealing time required to reach the problem ground-state with high fidelity. These results are presented in Appendix C.

So far we have only considered the effect of the catalyst on annealing spectra corresponding to the very simple MWIS graph shown in Figure 2. This was so we could examine the differing effects from the catalyst in the most straightforward setting. However, note that we observe these closing gaps for more general MWIS instances with additional local optima such that the problem Hamiltonian has more excited states with energy comparable to that of the ground-state. These instances can be constructed in much the same way as the problem graph shown in Figure 2(a), but with additional sub-graphs such that we have a complete k-partite graph. The catalysts were applied to instances with up to four sub-graphs and in each case the coupling was chosen to be within

the sub-graph corresponding to the excited problem state associated with the AC involving the ground-state. Further details on the construction of these more general instances can be found in Appendix D. For all the examples examined (where the problem parameters were chosen to produce a perturbative crossing with strong-gap-scaling) we observed a closing of the gap at the AC as well as the formation of an additional gap minimum between the instantaneous ground and first excited states - suggesting that the effects from the catalyst described in Section III B may be a more general phenomenon. An example spectrum with and without a catalyst is shown in Figure 9. Details on the construction of this problem instance are given in Appendix D.

We note however that the production of this additional gap minimum cannot generally be expected to result in a setting suitable to DQA. For instance, consider the example energy spectra presented in Figure 9 where we have introduced an additional local optimum. Here, the additional gap minimum results in a setting where the system is likely to evolve to the second excited state if the anneal is run diabatically. We note that if the anneal is run without a catalyst, the energy spectrum suggests that the system is likely to evolve to the first excited state - meaning that, in this case, the additional gap minimum not only fails to help the system reach the ground-state, but it actually reduces the quality of the solution found. We therefore suggest that, while the mechanism responsible for this additional gap minimum may be of interest for manipulating the gap spectrum, the changes to the gap spectrum we observe in this work are not, by themselves, useful for creating a gap spectrum suitable to DQA. However, as we note towards the end of this section, a simple extension to this catalyst may be able to produce useful annealing spectra in a wider variety of settings.

We note that the formation of the additional gap minimum that we observe in the SGS case is not intrinsically linked to the presence of a perturbative crossing in the original annealing spectrum. We demonstrate this in Figure 10 which shows results for the gap spectrum and ground-state evolution when a catalyst is applied to the setting without an AC shown in Figure 3(a). The catalyst contains an XX-coupling within G_0 such that it couples $|E_0\rangle$ to a lower energy state than it couples $|E_1\rangle$ to. We suggest that the formation of the new gap minimum may relate to the existence of a J_{xx} value for which both the problem ground and first excited state components of the instantaneous ground-state vector have appreciable magnitudes in the vicinity of s_- such that either component crossing zero results in a sharp change in the ground-state vector.

We argue that the reason such a value is observed for the setting without an AC as well as the strong-AC setting (but not for the weak AC) relates to the magnitudes of key ground-state vector components in the original annealing spectrum. Effectively, what we are doing with the introduction of our catalyst is to guide the anneal away from a particular problem state, which has the effect of suppressing the magnitude of the corresponding vector component in the instantaneous ground-state. For the two settings with perturbative crossings that we examined this is the component corresponding to the problem first excited state and for the setting without a perturbative crossing this is the component corresponding to the problem ground-state. What the setting without an AC and the strong-AC setting have in common is that the magnitude of the component being suppressed by the catalyst ($\langle E_0(s)|\, E_0\rangle$ for the no AC case and $\langle E_0(s)|\, E_1\rangle$ for the SGS AC case) approaches 1 in the original annealing spectrum without a catalyst - as can be seen in Figures 3(a) and (b). Looking at Figure 3(c), we see that in the weak-AC setting, the maximum magnitude reached by the component being suppressed ($\langle E_0(s)|\, E_1\rangle$) is only around 0.45. Thus, the problem state that our catalyst is guiding the anneal away from reaches greater levels of suppression for the same J_{xx} values in this setting than it does in the other two. This goes hand in hand with our earlier observation that, unlike in the SGS setting, the AC is already significantly lifted in the WGS setting for the J_{xx} value at which $s_{\mathrm{x}} = s_-$.

It is not yet clear whether the behaviour we observe in these two settings will scale to larger systems. Regarding the gap enhancement in the weak-AC setting, the optimal J_{xx} appears to be quickly approaching a constant - see the inset in Figure 5(a). That being said, the results we have presented are for very small system sizes and may not be sufficient to determine the true scaling behaviour. Our initial investigations into larger systems has revealed that the response of the spectrum becomes more complex as we increase the number of spins. For $n \gtrsim 25$ we start to see the appearance of two further gap minima in the spectrum (in addition to the effect of the original gap minimum being enhanced) with some transitional behaviour starting at $n \approx 17$. While these additional gap minima may not necessarily bottleneck the algorithm, since they result in a similar diabatic path to that seen for the SGS setting, it means the gap minimum is no longer a useful metric to track the enhancement to the algorithm. Determining whether or not the catalyst examined here continues to be of use in larger system sizes, and how the optimal J_{xx} scales, will require an examination of the dynamics.

Potentially a more fruitful avenue of further work however may be to apply these catalysts to problem settings with more general and interesting structures. The way in which we scale our problem graph in this work is to increase the number of vertices while keeping the number of local optima constant. However, in a realistic setting, one would typically expect the number of local optima to scale with the problem size [34, 49, 50]. The results presented here suggest that a single non-stoquastic XX-coupling may be sufficient to remove one perturbative crossing if it is chosen to couple the local optimum responsible for the perturbative crossing to another low energy state. In a setting with multiple local optima, several perturbative crossings may have to be removed

before the minimum gap scaling between the ground and first excited state becomes polynomial. We have begun investigating whether these catalysts could be extended to include multiple couplings to target the different perturbative crossings in the spectrum - with our initial results suggesting that a similar reduction to the gap scaling can be achieved in this way. We suggest that an algorithmic approach similar to that in [49] could be developed, where couplings are introduced and their strengths adjusted according to the local optima returned by sub-exponential QA runs. We noted in the introduction to this paper that quantum algorithms like QA are not expected to alter the complexity class of a problem. It is plausible that, if such a process could produce a spectrum that allowed a sub-exponential annealing run to achieve a high GS fidelity, it may require exponentially many calls to the quantum annealer.

Returning to the SGS setting, it also remains to be seen whether or not there continues to exist a regime where we observe the closing gaps and where the catalyst does not reduce the gap scaling as we move to larger problems. By exploiting permutation symmetries in the problem [33] we have been able to begin preliminary investigations into graphs with up to 70 vertices and have confirmed the existence of the additional closing gaps up to these system sizes. Something it may be beneficial to investigate is how the boundary between the two regimes changes with respect to different problem parameters with increasing system size. With regards to the existence of additional local optima, we have so far confirmed the behaviour we associated with the strong-gap-scaling for graphs with up to five local optima. As with the WGS setting, our preliminary results have also suggested that a recursive strategy may be of use in which an XX-coupling in introduced to target each perturbative crossing in the spectrum. The result in this case is a spectrum with a diabatic route to the GS - rather than an enhanced gap minimum.

Finally, we note that while our investigation has made use of the MWIS problem, there is no reason to think these findings should not extend to other problem settings. Crucially, the motivation behind the catalyst, and our further discussion of its effects, have hinged on the structure of the problem spectrum - rather than specific features of the MWIS problem. We would therefore expect to see similar results for any setting where an XX-term can be introduced to couple a local optimum responsible for a perturbative crossing to another low energy state. Further to this, we suggest that catalysts containing higher order X couplings, or couplings of a different form (e.g: YY, XY etc...), could also be understood in this way - with the effect that they have on the annealing spectrum being tied to the couplings which they create between states in the problem spectrum.

V. CONCLUSION

We have examined the effects of specific XX-catalysts on annealing spectra corresponding to small instances of the MWIS problem. In particular, we have examined how the effect of a targeted non-stoquastic XX-coupling, chosen to enhance the gap size at a perturbative crossing, differs depending on the nature of the AC that is present in the spectrum. We found that the response of the gap spectra to the introduction of the catalyst was highly sensitive to subtle changes in the encoding of the problem which affected the scaling of the gap at the AC with system size

More specifically, we found that for parameter settings that resulted in a AC with milder (but still exponential) gap scaling, the catalyst resulted in an enhancement of the gap minimum, and that if an optimized catalyst strength was used for each system size, the gap scaling could be significantly improved. We then applied the same catalysts to the case where we chose our problem parameters to produce an AC with a stronger gap scaling and found that, while some gap enhancement was possible, the catalyst was not able to provide the same improvement in scaling. In addition, we found that in this setting the catalyst could result in the closing of the gap at the AC, as well as the production of an additional gap minimum in the spectrum, if introduced with particular magnitudes. We determined that these magnitudes were the ones for which the s value at which the gap minimum occurred was the same as that for which the sign of key ground-state vector components changed.

These results suggest that the catalysts we examine here may be less successful in suppressing the gap scaling at an AC when the original gap scaling is less favourable. More generally, they show that small changes to the problem parameters can result in the same catalyst having strikingly different effects indicating that great care is required when designing a catalyst.

With regards to the creation of the new gap minimum, we found that the resultant spectrum allowed a higher GS fidelity to be reached for shorter run times as a result of diabatic transitions. We noted however that the introduction of the additional gap minimum could not generally be expected to produce such a spectrum and so was not, by itself, a strategy for creating spectra suitable for DQA.

As well as helping the guide the development of QA algorithms, these findings may also have relevance for other NISQ-era quantum algorithms such as the Quantum Approximate Optimisation Algorithm (QAOA) [51] which can be thought of as a trotterised QA. Coherence times limit QAOA to shallow circuit depths in much the same way that individual annealing runs are limited to short duration. Hamiltonians that reduce the required runtime in QA may also be a useful tool in helping QAOA reach higher GS fidelities for smaller circuit depths.

ACKNOWLEDGMENTS

We gratefully acknowledge Vicky Choi, Tameem Albash, Daniel O'Connor and Robert Banks for inspiring discussion and helpful comments. This work is supported by EPSRC, grant references EP/S021582/1 and EP/T001062/1.

[1] B. Apolloni and D. D. Falco, Quantum Stochastic Optimization, Stochastic Processes and their Applications **33** (1989).

[2] A. B. Finnila, M. A. Gomez, C. Sebenik, C. Stenson, and J. D. Doll, Quantum annealing: A new method for minimizing multidimensional functions, Chemical Physics Letters **219**, 10.1016/0009-2614(94)00117-0 (1994), arXiv:9404003 [chem-ph].

[3] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, Physical Review E **58**, 10.1103/PhysRevE.58.5355 (1998), arXiv:9804280 [cond-mat.stat-mech].

[4] Álvaro Rubio-García, J. J. García-Ripoll, and D. Porras, Portfolio optimization with discrete simulated annealing, (2022), arXiv:2210.00807.

[5] E. Gabbassov, Transit facility allocation: Hybrid quantum-classical optimization, PLOS ONE **17**, 10.1371/journal.pone.0274632 (2023).

[6] A. Witt, C. Körber, A. Kirstädter, and T. Luu, Tactile network resource allocation enabled by quantum annealing based on ilp modeling, (2022), arXiv:2212.07854.

[7] T. Otsuka, A. Li, H. Takesue, K. Inaba, K. Aihara, and M. Hasegawa, High-speed resource allocation algorithm using a coherent ising machine for noma systems, (2022), arXiv:2212.01578 [cs.IT].

[8] J. F. A. Sales and R. A. P. Araos, Adiabatic quantum computing for logistic transport optimization, Frontiers in Computer Science **5** (2023), arXiv:2301.07691.

[9] Tosio Kato, On the Adiabatic Theorem of Quantum Mechanics, Journal of the Physical Society of Japan **5** (1950).

[10] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, A quantum adiabatic evolution algorithm applied to random instances of an np-complete problem, Science **292**, 10.1126/science.1057726 (2001).

[11] A. P. Young, S. Knysh, and V. N. Smelyanskiy, First-order phase transition in the quantum adiabatic algorithm, Physical Review Letters **104**, 10.1103/physrevlett.104.020502 (2010), arXiv:0910.1378 [cond-mat.stat-mech].

[12] T. Jörg, F. Krzakala, J. Kurchan, and A. C. Maggs, Quantum annealing of hard problems, Progress of Theoretical Physics **184**, 10.1143/ptps.184.290 (2010), arXiv:0910.5644 [quant-ph].

[13] T. Jörg, F. Krzakala, J. Kurchan, A. C. Maggs, and J. Pujos, Energy gaps in quantum first-order mean-field–like transitions: The problems that quantum annealing cannot solve, EPL (Europhysics Letters) **89**, 10.1209/0295-5075/89/40004 (2010).

[14] B. Seoane and H. Nishimori, Many-body transverse interactions in the quantum annealing of the p-spin ferromagnet, Journal of Physics A: Mathematical and Theoretical **45**, 10.1088/1751-8113/45/43/435301 (2012), arXiv:1207.2909 [quant-ph].

[15] S. Knysh, Zero-temperature quantum annealing bottlenecks in the spin-glass phase, Nature Communications **7**, 10.1038/ncomms12370 (2016).

[16] M. H. Amin and V. Choi, First-order quantum phase transition in adiabatic quantum computation, Physical Review A **80**, 10.1103/PhysRevA.80.062326 (2009), arXiv:0904.1387 [quant-ph].

[17] B. Altshuler, H. Krovi, and J. Roland, Anderson localization makes adiabatic quantum optimization fail, Proceedings of the National Academy of Sciences **107**, 10.1073/pnas.1002116107 (2010).

[18] J. Preskill, Quantum computing in the nisq era and beyond, Quantum **2** (2018).

[19] T. Zanca and G. E. Santoro, Quantum annealing speedup over simulated annealing on random ising chains, Physical Review B **93**, 10.1103/PhysRevB.93.224431 (2016).

[20] 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, Quantum optimization of maximum independent set using rydberg atom arrays, Science **376**, 10.1126/science.abo6587 (2022).

[21] A. D. King, S. Suzuki, J. Raymond, A. Zucca, T. Lanting, F. Altomare, A. J. Berkley, S. Ejtemaee, E. Hoskinson, S. Huang, E. Ladizinsky, A. J. R. MacDonald, G. Marsden, T. Oh, G. Poulin-Lamarre, M. Reis, C. Rich, Y. Sato, J. D. Whittaker, J. Yao, R. Harris, D. A. Lidar, H. Nishimori, and M. H. Amin, Coherent quantum annealing in a programmable 2000-qubit ising chain, Nature Physics **18**, 10.1038/s41567-022-01741-6 (2022).

[22] E. Farhi, J. Goldstone, D. Gosset, S. Gutmann, H. B. Meyer, and P. Shor, Quantum adiabatic algorithms, small gaps, and different paths, Quantum Information and Computation **11**, 10.26421/qic11.3-4-1 (2011), arXiv:0909.4766 [quant-ph].

[23] N. G. Dickson and M. H. Amin, Does adiabatic quantum optimization fail for NP-complete problems?, Physical Review Letters **106**, 10.1103/PhysRevLett.106.050502 (2011), arXiv:1010.0669 [quant-ph].

[24] Y. Susa, Y. Yamashiro, M. Yamamoto, and H. Nishimori, Exponential speedup of quantum annealing by inhomogeneous driving of the transverse field, Journal of the Physical Society of Japan **87**, 10.7566/JPSJ.87.023002 (2018), arXiv:1801.02005 [quant-ph].

[25] K. Takada, Y. Yamashiro, and H. Nishimori, Mean-field solution of the weak-strong cluster problem for quantum annealing with stoquastic and non-stoquastic catalysts, Journal of the Physical Society of Japan

89, 10.7566/JPSJ.89.044001 (2020), arXiv:1912.09189 [quant-ph].
[26] L. Fry-Bouriaux, D. T. O'Connor, N. Feinstein, and P. A. Warburton, Locally suppressed transverse-field protocol for diabatic quantum annealing, Physical Review A **104**, 10.1103/physreva.104.052616 (2021).
[27] J. I. Adame, P. L. McMahon, and P. L. McMahon, Inhomogeneous driving in quantum annealers can result in orders-of-magnitude improvements in performance, Quantum Science and Technology **5**, 10.1088/2058-9565/ab935a (2020), arXiv:1806.11091 [quant-ph].
[28] M. M. Rams, M. Mohseni, and A. D. Campo, Inhomogeneous quasi-adiabatic driving of quantum critical dynamics in weakly disordered spin chains, New Journal of Physics **18**, 10.1088/1367-2630/aa5079 (2016), arXiv:1606.07740 [quant-ph].
[29] E. Farhi, J. Goldstone, and S. Gutmann, Quantum adiabatic evolution algorithms with different paths, (2002), arXiv:0208135 [quant-ph].
[30] Y. Seki and H. Nishimori, Quantum annealing with antiferromagnetic transverse interactions for the Hopfield model, Journal of Physics A: Mathematical and Theoretical **48**, 10.1088/1751-8113/48/33/335301 (2015), arXiv:1410.0450 [cond-mat.stat-mech].
[31] L. Hormozi, E. W. Brown, G. Carleo, and M. Troyer, Nonstoquastic Hamiltonians and quantum annealing of an Ising spin glass, Physical Review B **95**, 10.1103/PhysRevB.95.184416 (2017), arXiv:1609.06558 [quant-ph].
[32] H. Nishimori and K. Takada, Exponential enhancement of the efficiency of quantum annealing by non-stoquastic Hamiltonians, Frontiers in ICT **4**, 10.3389/fict.2017.00002 (2017), arXiv:1609.03785 [quant-ph].
[33] T. Albash, Role of nonstoquastic catalysts in quantum adiabatic optimization, Physical Review A **99**, 10.1103/PhysRevA.99.042334 (2019), arXiv:1811.09980 [quant-ph].
[34] V. Choi, Essentiality of the Non-stoquastic Hamiltonians and Driver Graph Design in Quantum Optimization Annealing, (2021), arXiv:2105.02110 [quant-ph].
[35] T. Albash and M. Kowalsky, Diagonal catalysts in quantum adiabatic optimization, Physical Review A **103**, 10.1103/PhysRevA.103.022608 (2021), arXiv:2009.05726 [quant-ph].
[36] V. Mehta, F. Jin, H. De Raedt, and K. Michielsen, Quantum annealing with trigger hamiltonians: Application to 2-satisfiability and nonstoquastic problems, Physical Review A **104**, 10.1103/physreva.104.032421 (2021).
[37] R. D. Somma, D. Nagaj, and M. Kieferová, Quantum speedup by quantum annealing, Physical Review Letters **109**, 10.1103/physrevlett.109.050501 (2012).
[38] L. Gupta and I. Hen, Elucidating the interplay between non-stoquasticity and the sign problem, Advanced Quantum Technologies **3**, 10.1002/qute.201900108 (2019).
[39] E. Crosson, T. Albash, I. Hen, and A. P. Young, Designing Hamiltonians for quantum adiabatic optimization, Quantum **4**, 10.22331/Q-2020-09-24-334 (2020), arXiv:2004.07681 [quant-ph].
[40] K. Takada, S. Sota, S. Yunoki, B. Pokharel, H. Nishimori, and D. A. Lidar, Phase transitions in the frustrated ising ladder with stoquastic and nonstoquastic catalysts, Physical Review Research **3**, 10.1103/physrevresearch.3.043013 (2021).
[41] V. Choi, The effects of the problem Hamiltonian parameters on the minimum spectral gap in adiabatic quantum optimization, Quantum Information Processing **19**, 10.1007/s11128-020-2582-1 (2020), arXiv:1910.02985 [quant-ph].
[42] R. Karp, Reducibility among combinatorial problems, Complexity of Computer Computations **40**, 10.1007/978-3-540-68279-0_8 (1972).
[43] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev, Adiabatic quantum computation is equivalent to standard quantum computation, SIAM J. Comput. **37**, 10.22331/q-2018-08-06-79 (2005), arXiv:0405098 [quant-ph].
[44] M. Marvian, D. A. Lidar, and I. Hen, On the computational complexity of curing non-stoquastic hamiltonians, Nature Communications **10**, 10.1038/s41467-019-09501-6 (2019).
[45] M. H. Amin, Effect of local minima on adiabatic quantum optimization, Physical Review Letters **100**, 10.1103/PhysRevLett.100.130503 (2008), arXiv:0709.0528 [quant-ph].
[46] V. Choi, Avoid First Order Quantum Phase Transition by Changing Problem Hamiltonians, (2010), arXiv:1010.1220 [quant-ph].
[47] A. Braida and S. Martiel, Anti-crossings and spectral gap during quantum adiabatic evolution, Quantum Information Processing **20**, 10.1007/s11128-021-03198-7 (2021), arXiv:2102.00987 [quant-ph].
[48] E. J. Crosson and D. A. Lidar, Prospects for quantum enhancement with diabatic quantum annealing, Nature Reviews Physics **3**, 10.1038/s42254-021-00313-6 (2021).
[49] N. G. Dickson and M. H. Amin, Algorithmic approach to adiabatic quantum optimization, Physical Review A **85**, 10.1103/PhysRevA.85.032303 (2012), arXiv:1108.3303 [quant-ph].
[50] M. Žnidarič and M. Horvat, Exponential complexity of an adiabatic algorithm for an NP-complete problem, Physical Review A **73**, 10.1103/physreva.73.022329 (2006).
[51] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, (2014), arXiv:1411.4028.

Appendix A: MWIS problem parameters

The MWIS problem Hamiltonian is given by

$$H_p = \sum_{i\in\{\text{vertices}\}} (c_i J_{zz} - 2w_i)\sigma_i^z + \sum_{(i,j)\in\{\text{edges}\}} J_{zz}\sigma_i^z\sigma_j^z \tag{A1}$$

where c_i is the number of edges connected to vertex i, w_i is the weight on vertex i, and J_{zz} is the edge penalty. As discussed in Section II B, our problem instances are on complete bipartite graphs where we label our two disconnected sub-graphs as G_0 and G_1 - as shown in Fig 2. We allocate a total weight of 1 to G_1 and a total weight of $1 + \delta W$ to G_0. This weight is split evenly between the vertices in each sub-graph. To get a clearer picture of what parameters we can change to define our problem instances within this structure, we can re-write equation

A1 as

$$H_p = (n_1 J_{zz} - \frac{2(1+\delta W)}{n_0}) \sum_{i \in G_0} \sigma_i^z + (n_0 J_{zz} - \frac{2}{n_1}) \sum_{i \in G_1} \sigma_i^z + J_{zz} \sum_{i \in G_0} \sum_{j \in G_1} \sigma_i^z \sigma_j^z, \quad \text{(A2)}$$

where n_0 and n_1 are the number of vertices in G_0 and G_1 respectively. For a particular problem graph defined by n_0 and n_1, the free parameters are then δW and J_{zz}.

So that we do not change the energy scale of the problem Hamiltonian when we adjust the values of δW and J_{zz}, but rather just the relative gaps between different energy levels, we normalise the parameters that go into equation A1. For a set of un-normalized parameter values $(\delta W', J'_{zz})$ that define the problem instance, we first calculate the un-normalized vertex weights as $w'_i = (1 + \delta W')/n_0$ for $i \in G_0$ and $w'_i = 1/n_1$ for $i \in G_1$. The normalized parameters that enter equation A1 are then obtained as

$$w_i = E_{\text{scale}} \times K \times w'_i,$$

$$J_{zz} = E_{\text{scale}} \times K \times J'_{zz}$$

where E_{scale} sets the energy scale in relation to the driver and K is a normalisation factor given by

$$K = \frac{n_0 + n_1}{4(n_0 \times n_1 \times J'_{zz} - 1)}.$$

This factor is the reciprocal of the energy gap between the ground and highest excited state of H_p as it would be with the *un*-normalized parameters. The expression has been obtained by using equation A2 to obtain exact expressions for these two energies in terms of the sub-graph sizes and the un-normalized edge penalty and weight difference.

We choose our problem parameters with reference to our five vertex instance ($n_0 = 2$ and $n_1 = 3$) since this is the instance for which we present data for the whole anneal rather than just the minimum gap values. Our different parameter sets $(\delta W', J'_{zz})$ are chosen by first selecting $\delta W'$ and then, through trial and error, adjusting the *un-normalized* J'_{zz} so that in the annealing spectrum to corresponding to the *normalized* problem Hamiltonian, $s_{\text{x}} = 0.9$ for our 5-vertex example. Note that the values for δW and J_{zz} quoted in the body of this work are in fact the un-normalized values $\delta W'$ and J'_{zz}.

Appendix B: Perturbative introduction of the catalyst

This section follows a perturbative argument made in [34] which offers some insight into why the catalysts used in this work result in gap enhancement at the perturbative crossings in our systems. Unlike the perturbative argument outlined in Section II B, in which H_p was taken to be the initial Hamiltonian, here we consider the effect of introducing the catalyst as a perturbation to the total Hamiltonian for a given s. That is, we take the initial Hamiltonian to be $H(s) = (1-s)H_d + sH_p$ and then write our perturbed Hamiltonian as

$$H(s, \mu) = H(s) + \mu H_c.$$

To first order, the perturbed energies are then

$$E_i(s, \mu) = E_i(s) + \mu \langle E_i(s)| H_c |E_i(s)\rangle = E_i(s) + 2\mu J_{xx} \sum_{(j,k) \in C_{xx}} \langle E_i(s)|E_j\rangle \langle E_k|E_i(s)\rangle \quad \text{(B1)}$$

where we have used C_{xx} to denote the set of pairs of problem states that are coupled by the single XX-coupling included in the catalyst.

We are then interested in what happens to the instantaneous ground and first excited states ($i = 0, 1$) around the point of the perturbative crossing, s_{x}. In order to use equation B1 to gain some insight as to the effect of the catalyst we can consider the perturbative argument from Section II B to reason what problem states will dominate the two instantaneous states *prior* to the introduction of the catalyst - *i.e:* for $H(s, \mu = 0)$. We know that for $s > s_{\text{x}}$, $|E_0\rangle$ will make up a large part of $|E_0(s)\rangle$ and that $|E_1\rangle$ will make up a large part of $|E_1(s)\rangle$. As for what other problem states will have a reasonable presence in $|E_{0,1}(s)\rangle$, we know from perturbation theory that the introduction of the driver will introduce states with a weight that depends on their energy difference and Hamming distance from the state being perturbed - *i.e:* from $|E_{0,1}\rangle$. We thus reason that for $s > s_{\text{x}}$, $|E_0(s)\rangle$ is occupied by problem states corresponding to sub-sets of G_0 and that $|E_1(s)\rangle$ is occupied by problem states corresponding to sub-sets of G_1. That is, they are occupied by states corresponding to independent sets (*i.e:* low energy states such that the energy differences between them and $E_{0,1}$ are small) that are closest in Hamming weight to the two respective problem states. We denote these sets of problem states as $P^{\text{ext}}_{0,1}$.

We wish to examine the effect of the catalyst at points in the anneal shortly before and after the AC. Following the notation in [34], we denote these points as s^-_{x} and s^+_{x} respectively. From the above, we can write

$$|E_0(s^+_{\text{x}})\rangle \approx \sum_{i \in P^{\text{ext}}_0} \langle E_0(s^+_{\text{x}})|E_i\rangle |E_i\rangle \equiv |E_0\rangle^{\text{ext}}_0$$

where we are not interested in the specific values of $\langle E_0(s)|E_i\rangle$. Note that the subscript inside the ket refers to the instantaneous ground-state while the subscript outside the ket refers to the set of states P^{ext}_0. Similarly, we can write

$$|E_1(s^+_{\text{x}})\rangle \approx |E_1\rangle^{\text{ext}}_1 .$$

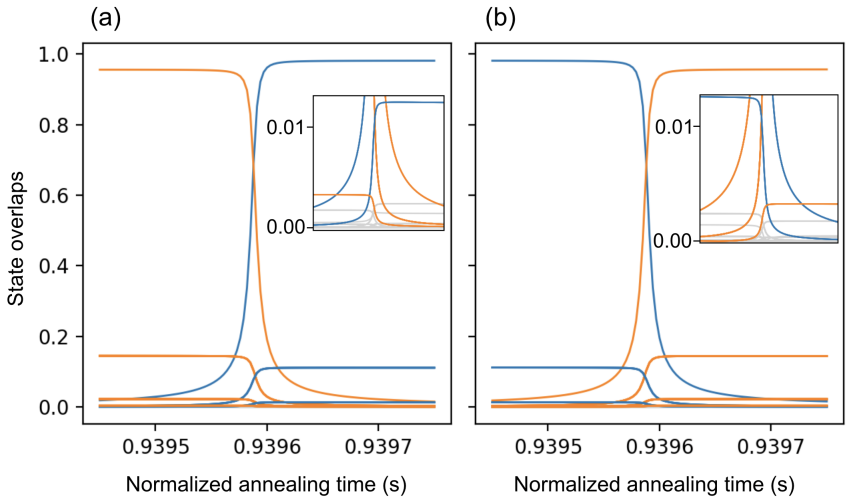

FIG. 11: Evolution of the instantaneous ground-state (a) and first excited state (b) around the location of the perturbative crossing in an anneal corresponding to a problem instance with $n_0 = 3$, $n_1 = 4$, $\delta W = 0.01$ and $J_{zz} = 5.33$. The instantaneous states are represented in terms of their overlaps with the problem states and these are colour coded with respect to which of the sets introduced in Appendix B they are in. Overlaps with problem states in P_0^{ext} and P_1^{ext} and plotted in blue and orange respectively and overlaps with problem states that are in neither are plotted in light grey. The insets show the same data over a smaller range of the state overlaps so that the overlaps with smaller contributions are visible.

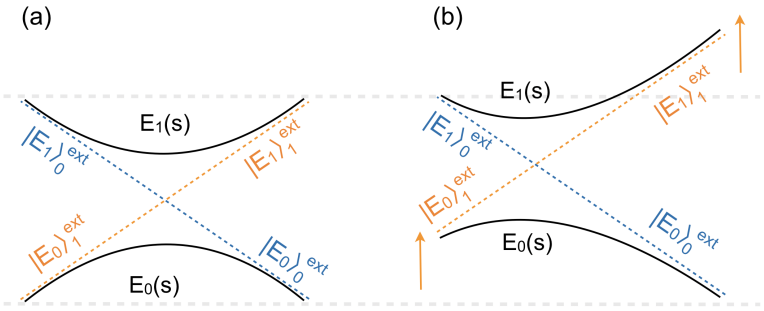

FIG. 12: Cartoons illustrating the arguments made in Appendix B regarding how the introduction of our choice of catalyst would perturb the instantaneous energy levels around the perturbative crossing. (a) depicts an avoided level crossing in which the ground-state goes from being dominated by the set of problem states P_1^{ext} to the set of problem states P_0^{ext} while the first excited state goes from being dominated by P_0^{ext} to P_1^{ext}. The two coloured dotted lines illustrate the exchange in problem states between the instantaneous ground and first excited states and the solid black lines illustrate the energies of the two instantaneous states. (b) then depicts the effect of the catalyst as argued in Appendix B. That is, there is no effect on a state consisting of P_0^{ext} while a state consisting of P_1^{ext} will experience an increase in energy. The cartoon illustrates why we expect this to move the crossing earlier in the anneal.

At the AC, the two instantaneous states exchange their properties and so we can further write

$$|E_0(s_x^-)\rangle \approx |E_0\rangle_1^{\text{ext}} ,$$

$$|E_1(s_x^-)\rangle \approx |E_1\rangle_0^{\text{ext}} .$$

We support our assumptions about the makeup of the instantaneous states around the AC with numerical results shown in Figure 11. These results correspond to a problem instance with $n_0 = 3$, $n_1 = 4$, $\delta W = 0.01$ and $J_{zz} = 5.33$. Figures 11(a) and (b) show the evolution of the instantaneous ground and first excited state respectively around the location of the perturbative crossing. Plotted in blue are the overlaps of the instantaneous states with the problem states in P_0^{ext} and in orange we show the overlaps with the problem states in P_1^{ext}. The rest are plotted in grey. These results are in agreement with our predictions.

We now consider what the perturbation in equation B1 looks like when the state being perturbed is $|E_{0,1}\rangle_0^{\text{ext}}$ or $|E_{0,1}\rangle_1^{\text{ext}}$. Due to the fact that the XX-coupling we introduce is between two spins in G_1, all the states in P_0^{ext} are coupled by the catalyst to states corresponding to dependent sets. That is there are no pairs of states in C_{xx} for which both states have a significant presence in

$|E_{0,1}\rangle_0^{\text{ext}}$. As a result, $\langle E_{0,1}|_0^{\text{ext}} H_c |E_{0,1}\rangle_0^{\text{ext}} = 0$. Thus, $E_0(s_{\text{x}}^+, \mu) \approx E_0(s_{\text{x}}^+)$ and $E_1(s_{\text{x}}^-, \mu) \approx E_1(s_{\text{x}}^-)$.

On the other hand, there are pairs of states in C_{xx} for which both states have a significant presence in $|E_{0,1}\rangle_1^{\text{ext}}$. Let us first consider what this means for $E_0(s_{\text{x}}^-, \mu)$. Because $H(s)$ is stoquastic, the ground-state vector components, $\langle E_0(s)|E_i\rangle$, must all be positive meaning that all the terms in the sum in equation B1 are also positive. Thus, for $J_{xx} > 0$ (as is the case in our setting), $E_0(s_{\text{x}}^-, \mu) > E_0(s_{\text{x}}^-)$. We now consider the perturbed energy $E_1(s_{\text{x}}^-, \mu)$. In this case, the individual vector components cannot be guaranteed to be positive since the first excited state can have both positive and negative vector components. However we can use an observation made in [34], regarding the behaviour of the vector components of two instantaneous states at an AC, to argue that $sign(\langle E_0(s)|E_i\rangle \langle E_0(s)|E_j\rangle) = sign(\langle E_1(s)|E_i\rangle \langle E_1(s)|E_j\rangle)$. Thus, we can also say that all the terms in the sum will be positive and that $E_1(s_{\text{x}}^+, \mu) > E_1(s_{\text{x}}^+)$.

In order to make a rigorous case for this enhancing the minimum gap size at the AC, further ideas from [34] must be introduced. However, some intuition as to why the gap size increases can be had by considering the cartoons shown in Figure 12. Figure 12(a) depicts the kind of AC described above where the instantaneous ground-state goes from $|E_0\rangle_1^{\text{ext}}$ to $|E_0\rangle_0^{\text{ext}}$ while the instantaneous first excited state goes from $|E_1\rangle_0^{\text{ext}}$ to $|E_1\rangle_1^{\text{ext}}$. The dashed lines illustrate the $|E_{0,1}\rangle_{0,1}^{\text{ext}}$ states (which can also be thought of as the perturbed states obtained when introducing the driver perturbatively to the problem Hamiltonian as in Section II B) and the solid black lines show the two instantaneous energy levels. Above we have argued that $|E_{0,1}\rangle_1^{\text{ext}}$ is increased by the introduction of the catalyst while $|E_{0,1}\rangle_0^{\text{ext}}$ is unchanged. Figure 12(b) depicts the effect of introducing the catalyst. We can see that the result is to move the AC to a lower value of s_{x}. Thus the strength of the driver Hamiltonian will be greater at the point of the crossing and the vectors will be more mixed - thus resulting in a greater overlap between the instantaneous ground-state before and after the crossing and so a larger gap size. We note that the decrease in s_{x} is also consistent with our numerical results in Section III.

Appendix C: Diabatic Anneal on the 5-vertex Graph

We present here numerical results for the dynamics of an anneal on the 5-vertex problem instance in the SGS setting. These results are obtained using closed system spin models and the evolution of the system is presented in terms of its overlap with the instantaneous ground and first excited states. Since we have not specified a true energy scale we cannot discuss the evolution with respect to an actual annealing time but only in relation to the magnitude with which we introduce our Hamiltonian.

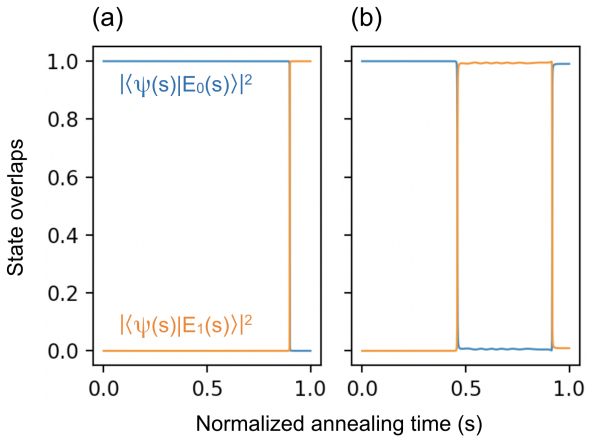

FIG. 13: Results showing the dynamics of an anneal to a problem graph with $n_0 = 2$, $n_1 = 3$, $\delta W = 0.01$ and $J_{zz} = 5.33$ in a closed system setting. This evolution is presented in terms of the system's overlap with the instantaneous ground and first excited states in blue and orange respectively. The total annealing time used is $T_{\text{anneal}} = 1000$ in comparison to the local driver fields being introduced with a magnitude of 1. (a) shows the results without a catalyst and (b) shows the results using a catalyst with $J_{xx} = 1.92$.

The annealing time used in these simulations was 1000, in relation to the local driver fields being introduced with a magnitude of 1. Figure 13(a) shows the evolution without a catalyst and we see that for this annealing time the system has a negligible overlap with the problem ground-state at the end of the anneal after transitioning into the first excited state at the avoided level crossing. Figure 13(b) shows the results when a catalyst is introduced with $J_{xx} = 1.92$, such that another small gap in the spectrum is produced. We observe that in this case the system ends the anneal with a near unity overlap with the problem ground-state as a result of transitioning into the first excited state at the first small gap and then back into the ground-state at the second.

We stress that these results are for a very specific system where only one perturbative crossing is present. The results presented in Figure 13 show that the spectrum in Figure 7(c) does provide a diabatic path that allows the GS to be reached for shorter run times. However, as we discuss in Section IV, we do not expect the creation of the additional gap minimum to always produce such a spectrum.

Appendix D: Introducing additional local optima

We describe here how our problem graph can be generalised to include additional local optima. This can be done straightforwardly by adding further sub-graphs for each additional local optima to produce a complete k-partite graph. That is we have k sub-graphs where each

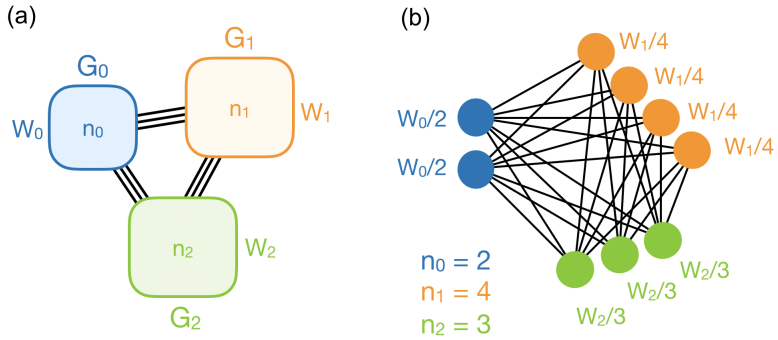

FIG. 14: An illustration of tri-partite graph constructed as described in Appendix D is shown in (a). An example with $n_0 = 2$, $n_1 = 4$ and $n_2 = 3$ is shown in (b). This is the example corresponding to the annealing spectra in Figure 9.

vertex in every sub-graph is connected to every vertex in every other sub-graph and there are no connections within the sub-graphs themselves. Each local optimum then corresponds to picking all the vertices from one of the sub-graphs. As with our bipartite graphs, each sub-graph G_a is given n_a vertices which determines the energy spectrum of the neighbourhood of the corresponding problem state. It is also given a total weight, W_a, which is shared out equally between its vertices such that each vertex in G_a has a weight of W_a/n_a. A general tri-partite graph is shown in Figure 14(a) and a specific example is shown in Figure 14(b). However, more sub-graphs can be added.