Targeted catalysts for quantum annealing

by

Natasha Feinstein

Supervisors:

P. A. Warburton

&

Sougato Bose

A thesis submitted in partial fulfilment
of the requirements for the degree of

Doctor of Philosophy

Department of Electronic and Electrical Engineering
University College London

Tuesday 2nd September, 2025

Declaration of Authorship

I, Natasha Feinstein, confirm that the work presented in this thesis, titled "Targeted catalysts for quantum annealing", is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis. Elements of the following parts of this thesis have been published in journals or in preparation:

1. **Chapter 3, 5 & 7:** Natasha Feinstein, Louis Fry-Bouriaux, Sougato Bose, P. A. Warburton, *Effects of XX-catalysts on quantum annealing spectra with perturbative crossings*, Phys. Rev. A, 110 (2024) [1].
2. **Chapter 8:** Natasha Feinstein, Ivan Shalashilin, Sougato Bose, Paul Warburton, *Robustness of diabatic enhancement in quantum annealing*, Quantum Science and Technology, 10 (2025) [2].

Signature: [redacted]

Date: **Tuesday 2nd September, 2025**

Abstract

Quantum annealing is a quantum algorithm for combinatorial optimisation that operates by evolving a system towards a Hamiltonian whose ground state encodes the optimal solution to a problem. The run-time needed to reach a given ground-state fidelity is dictated by the size of the minimum gap that appears between the ground- and first excited-states in the associated energy spectrum. A particular problem for the algorithm is the appearance of so-called perturbative crossings that form as a result of highly competitive local optima and result in gap minima that close exponentially with system size. The associated exponential run-time scaling has consequences for the efficiency of the algorithm and places infeasible demands on qubit coherence times.

We explore the use of a targeted XX-catalyst Hamiltonian that aims to circumvent this bottleneck. We use perturbation theory to elucidate the relationship between the structure of the problem Hamiltonian and the effects of different coupling choices. Through this, we motivate particular choices of XX-couplings that could be identified through knowledge of local optima which the algorithm is likely to return. We explore the effectiveness of these catalysts numerically, utilising a problem graph construction that allows us to controllably produce annealing spectra featuring perturbative crossings with varying properties.

Overall we find that the catalysts perform well, being able to remove perturbative crossings in a range of instances with the number of couplings present in the catalyst scaling linearly with the system size. In some cases we observe additional catalyst effects which complicate the outcome. In these cases we consider when and how they could be exploited in diabatic quantum annealing. Furthermore, we find that the theoretical framework that we have used to understand the effects of XX-catalysts extends well to other coupling choices, potentially proving a useful tool in the design of catalyst Hamiltonians more generally.

Impact Statement

This thesis is a theoretical and numerical investigation into the use of catalyst Hamiltonians in quantum annealing – a quantum algorithm for solving combinatorial optimisation problems. Specifically, this work investigates the use of a targeted catalyst to reduce the time required for an annealing run to return the optimal solution with high fidelity. Reductions to the annealing time are important for two reasons. First, they lower the demand placed on the coherence times of the hardware, facilitating the solving of larger problems on nearer term processors. Shorter annealing times may also offer better time scaling advantages over classical algorithms.

Being able to solve combinatorial optimisation problems has a variety of real word applications. These include portfolio optimisation, resource allocation, transport optimisation and drug design in chemistry. More favourable time scaling will make the solving of larger and more interesting problems more feasible. Furthermore, in many cases it is crucial to be able to solve these problems in the shortest possible time to make the results relevant.

With regards to what our work contributes towards this goal: our investigations highlight a particular form of XX-catalyst for the removal of perturbative crossings in annealing spectra – which are a particular annealing time bottleneck that forms in hard problems. Importantly, the structure of these catalysts relates to potentially accessible information about the problem. This could, in future, facilitate the design of an algorithmic approach to their construction. We make some proposals in this dissertation as to how this could be achieved.

XX-couplings are not yet readily implementable on current hardware and so this currently lies more within the scope of academic exploration. The design of hardware that facilitates such couplings is, however, an active field of research. That the XX-catalysts we have examined perform well offers further motivation for the development of such hardware.

In motivating the catalyst Hamiltonians on which our thesis focuses, we have introduced theoretical tools and a framework which could be applied more generally in catalyst design for quantum annealing. We have also introduced a graph construction which can serve as an aid in related numerical investigations.

Finally, we have highlighted a particular setting, in which the catalyst results in the creation of an additional gap minimum in the annealing spectrum, which could be used as a theoretical test-bed for diabatic quantum annealing. Utilising this setting, we performed an investigation into the robustness of diabatic annealing strategies that offers insight into what considerations would need to be made for such strategies to be practically viable. With the implementation of XX-couplings in hardware, this setting would also serve as a useful *experimental* test-bed – requiring as few as five qubits and one XX-coupler to implement.

Acknowledgements

It is difficult to open an acknowledgments section in a way that doesn't just sound like a cliche, but it really is true that I would not have made it to this point without a great number of wonderful people being there to help me. Anyone who has done a PhD, or I think had even a single conversation with someone doing a PhD, knows that it can be a chaotic and challenging time – both with respect to the work itself and with regards to the angst and self doubt that are inevitable parts of working on a large scale research project for the first time. I am so very grateful to everyone who has been there for me in any way over the last four and a half years and I would like to take a couple of pages to express that.

I'd like to start by thanking my supervisor, Prof. Paul Warburton. Paul, I could not have asked for a better mentor. Thank you for encouraging and supporting me, for being patient with me but also pushing me when I needed it. The time and care, as well as the trust, that you put in me has allowed me to grow into a confident and capable scientist and I am endlessly grateful. I would also like to thank my secondary supervisor, Prof. Sougato Bose, for being the on-call theorist to a baby scientist and an electrical engineer. This project would not have been possible without you.

I am exceptionally grateful to Louis Fry-Bouriaux who was a huge part of the first year of my PhD. Thank you for helping me take my first baby steps into the world of research at a time when I didn't fully believe in myself. Your support made all the difference. I'd also like to give a great big thank you to all the other people who have been a part of

Paul's research group during my time here – Daniel O'Connor, Gioele Consani, Nedeen Al Sharif, Henry Chew, Leon Guerrero, Bobby Banks, Luca Nutricati, Pias Tubsrinuan and, honorary member, Emily Cook. It has been a joy to be surrounded by such a kind and vibrant group of people. Wherever I go next, my colleagues have a lot to live up to. I would like to give a special mention to Bobby for being my partner in crime – a.k.a my fellow theorist. Thank you for lending me your incredible brain when I needed it, for being so encouraging, and for cooking me dinner when I was too stressed to function. I feel very grateful to have made a friend like you.

I'd of course like to thank EPSRC and the centre for doctoral training in Delivering Quantum technologies for making this PhD possible. I'd also like to thank all my fellow students who have been on this journey with me. In particular, I'd like to thank Eleanor Parker for being such a light during my masters year. I don't know that I would have stuck it out if you weren't there with me, sharing the experience and making me laugh until my whole body ached. Dr Abbie Bray – thank you for being such a constant source of support and making every part of PhD life easier, both as a mentor and a friend.

Then of course there are all the wonderful friends outside of UCL who have been there for me over these past four and a half years. There are too many of you to list by name but I do want to make a few special mentions. In particular I would like to give a huge thank you to Willow Weston and Sabrina Parker who have been absolute rocks during this time. I feel so lucky to have such loving and capable humans in my corner. Not everyone gets to be blessed with friends who will be there for them so consistently and with so much to give.

I would also like to give a big thank you to Patrick Scruby. I truly believe that without your help I would not even have made it to the start of this journey. Thank you for carrying me through my undergrad and for being a wonderful friend to me from then until now. Thank you Connie Blott for being such a positive force in my life. You are a

beautiful human and I am endlessly grateful to you for your desire and ability to uplift those around you. Thank you to Joe Klein for keeping me sane by facilitating my regular escape into a fantasy world. Your D&D campaign has genuinely been integral to my wellbeing – I wish I was joking but I am not. Last, but certainly not least, thank you to Adam Searle for getting me through some of the toughest periods of my life – only some of which were PhD related. I couldn't imagine this time without you in it.

Finally, I would like to thank my parents, Marina Solarek and Martin Feinstein. I could truly write pages on how grateful I am to you both for the various ways in which you've helped me through my PhD. Thank you for being a source of strength, while also providing a space to step away from my work during times when it felt all-encompassing. Whatever else was happening, it was always against the backdrop of endless love and unconditional support that you gave me.

Contents

List of Figures

List of Acronyms

AC = **A**voided level **C**rossing

AFM = **A**nti **F**erro**M**agnetic

AQA = **A**diabatic **Q**uantum **A**nnealing

DQA = **D**iabatic **Q**uantum **A**nnealing

FM = **F**erro**M**agnetic

GS = **G**round **S**tate

KES = **K**th **E**xcited **S**tate

K-SAT = **K**-variable Boolean **SAT**isfiability problem

KZ = **K**ibble **Z**urek

LZ = **L**andau **Z**ener

MIS = **M**ax **I**ndependent **S**et

MWIS = **M**ax **W**eighted **I**ndependent **S**et

NP = **N**ondeterministic **P**olynomial

P = **P**olynomial

PM = **P**ara**M**agnetic

QA = **Q**uantum **A**nnealing

QAOA = **Q**uantum **A**pproximate **O**ptimisation **A**lgorithm

QMC = **Q**uantum **M**onte **C**arlo

QUBO = **Q**uadratic **U**nconstrained **B**inary **O**ptimisation

SG = **S**pin **G**lass

TTS = **T**ime **T**o **S**olution

Chapter 1

Introduction

Quantum annealing (QA) is a continuous-time quantum algorithm proposed as a means of solving combinatorial optimisation problems faster than can be achieved classically. The algorithm entered the literature in the 1990s [4, 5, 6], motivated by the idea that quantum tunnelling could be a route to escaping local optima – which pose one of the main obstacles for classical algorithms. In combinatorial optimisation, one is looking to either maximise or minimise some objective function over a discrete, but often exponentially large, configuration space. Being able to solve these kinds of problems has real world applications ranging from portfolio optimisation [7] and resource allocation [8, 9, 10] to transport optimisation [11]. In these applications it is of great interest to be able to find the best solution in the shortest possible time.

From a complexity theory perspective, many combinatorial optimisation problems of interest are considered **NP**-hard. That is, they are at least as difficult to solve as every problem in **NP**. Taking a brief moment to explain what this means: in complexity theory, **NP** is the set of decision problems for which a correct answer can be *verified* in polynomial time. A subset of this complexity class is **P**, which is the set of decision problems that can be *solved* in polynomial time. If **P**$\neq$**NP**, as is widely believed to be the case, we are left with a number of problems that can be *verified* in polynomial time but take a *super*-polynomial time to solve. The hardest problems in this class are the **NP**-complete problems. The time required to obtain the solution to *these* problems, along with the

NP-hard problems of which they are a sub-set, is expected to be at least exponential in the problem size.

This can, in part, be understood through the fact that an exponentially large search space renders anything even approaching an exhaustive search for an optimal solution completely intractable. Another issue is the prevalence of local optima in which heuristic search algorithms can become trapped as they attempt to navigate the complex optimisation landscapes of these problems. To what extent *quantum* algorithms, in general, can offer some kind of speed-up over classical approaches, and what kind of speed-up we can expect this to be (*e.g.* polynomial, exponential...), is a large and complex subject of discussion [12, 13, 14, 15, 16]. Some key examples of provable speed-up are Shor's algorithm [17], offering a super-polynomial speed-up over the fastest known classical algorithms for integer factorization, and Grover's algorithm [18], which offers a quadratic speed-up for unstructured search.

When it comes to QA in particular, there is as of yet, little evidence of definitive speed-up over classical algorithms for the purpose of combinational optimisation – though there do exist some provable speed-ups in other contexts [19, 20, 21]. Outside of current hardware limitations such as coherence times and connectivity, a particular problem that plagues QA is an exponential, and sometimes super-exponential, suppression to the tunnelling rate needed to escape the local optima of a problem and reach to the global optimum. As we shall go on to discuss later in this chapter, the prevalence of this exponential suppression appears to be linked to the hardness of the problem in the classical setting, casting doubt on the idea of quantum tunnelling as a route to speed-up for such problems – as per the initial motivation for the algorithm.

Reflective of the general uncertainty around what kind of advantage we should expect to see from quantum algorithms, the context in which the aforementioned tunnelling suppression is discussed differs between papers within the literature. Whether or not QA, or indeed quantum computing more generally, will prove able to efficiently solve **NP**-hard optimisation problems remains an open question that lies well outside the scope of this work [12, 13, 14, 15, 16, 22, 23, 24, 25]. Certainly, there is still an interest in finding exponential speed-up through QA, either through identifying **NP**-hard problems to which the algorithm is especially suited, or novel variations on the algorithm aimed at circumventing the exponentially suppressed tunnelling rates. More widely though, there is an assumption that quantum algorithms such as QA will *not* have the capacity to change the complexity class of a problem and will therefore not be able to solve **NP**-hard problems efficiently.

This, however, does not rule out the capacity for other types of speed-up in QA that could still offer a significant time advantage over classical algorithms. As such there is an interest in understanding whether QA can offer a *sub*-exponential scaling advantage over classical algorithms and developing strategies that reduce the *severity* of the exponential scaling. There is also an interest in reducing the demand on the coherence times of quantum hardware by shifting the exponential scaling to a different part of the algorithm – we will explain more thoroughly what we mean by this in the following section.

It is with these motivations in mind that we approach our work. In this thesis we focus on a particular bottleneck within the quantum annealing algorithm known as a *perturbative crossing*. These perturbative crossings arise as a result of competing local optima in which the algorithm can get stuck – requiring exponentially long run-times to be able to tunnel out of. In particular, we will be exploring how accessible information of the problem structure could be utilised in a targetted approach to effectively guide the anneal towards the global optimum of the problem.

The remainder of this chapter is organised as follows. We will begin by outlining the quantum annealing algorithm in its most standard form in section 1.1. We then, in section 1.2, give a brief introduction to the main hardware platforms that exist for implementing QA. Section 1.3 then provides an overview of the literature concerning the time scaling of the algorithm in its standard form. In particular, we will highlight the specific bottleneck which this work aims to address. Finally, having introduced the background required to better contextualise our discussion, we end the chapter by giving a structural overview of the rest of the thesis.

1.1 Quantum Annealing

The most standard form of QA operates by interpolating linearly between two non-commuting Hamiltonians. The first is a simple driver Hamiltonian, $\hat{H}_d$, and the second is the problem Hamiltonian, $\hat{H}_p$, which encodes the optimisation problem to be solved. The total Hamiltonian can be written as

$$\hat{H}(s) = (1-s)\hat{H}_d + s\hat{H}_p, \tag{1.1}$$

where s is a dimensionless annealing parameter that is varied monotonically from 0 to 1 over the course of the anneal such that $\hat{H}(0) = \hat{H}_d$ and $\hat{H}(1) = \hat{H}_p$. A typical choice in the literature is to give s the time dependence $s(t) = t/t_a$ where t_a is the total annealing time. The driver Hamiltonian in which the anneal starts is conventionally chosen to be a homogeneous local X-field,

$$\hat{H}_d = -\sum_{i=1}^{n} \hat{\sigma}_i^x, \tag{1.2}$$

where n is the total number of qubits and $\hat{\sigma}_i^x$ denotes the Pauli-X operator on the ith qubit. Its ground state is the symmetric equal superposition over all computational basis states. The problem Hamiltonian is then constructed out of Pauli-Z operators and encodes the optimisation problem to be solved such that the ground state of $\hat{H}_p$ encodes the optimal solution.

Throughout this dissertation we refer to the eigenstates of this total Hamiltonian as the *instantaneous* eigenstates and denote them and their corresponding energies as

$$\hat{H}(s)\,|E_a(s)\rangle = E_a(s)\,|E_a(s)\rangle\,. \tag{1.3}$$

The states are labelled starting from $a = 0$ in order of increasing energy. We refer to the smoothly evolving spectrum of energies, $\{E_a(s)\}$, as the annealing spectrum. Furthermore, since the annealing parameter, s, is varied monotonically from zero to unity, we will often refer to $s = 0$ and $s = 1$ as the start and end of the anneal respectively.

We will denote the *problem* eigenstates and their energies as

$$\hat{H}_p\,|E_a\rangle = E_a\,|E_a\rangle\,. \tag{1.4}$$

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.

The anneal starts with the quantum system initialised in the ground state of the driver Hamiltonian, $|\psi(t = 0)\rangle = |E_0(0)\rangle$. The aim is then for the evolution to proceed adiabatically such that $|\psi(t)\rangle \approx |E_0(s(t))\rangle$ for all t. If this is achieved, the system will end with $|\psi(t_\mathrm{a})\rangle \approx |E_0\rangle$ so that, upon measurement, the ground state of $\hat{H}_p$, and thus the optimal solution to the problem, is obtained with high fidelity. The limit on the rate at which the interpolation can progress while maintaining adiabatic evolution is given by the adiabatic theorem. This theorem has a number of formulations (see [26] and references therein) but in its most basic form it states that the total annealing time, t_a, must scale inversely with the square of the instantaneous gap minimum between the ground and first excited state. That is,

$$t_\mathrm{a} \propto \frac{1}{\min\limits_{0\leq s\leq 1}\,[E_1(s) - E_0(s)]^2}\,. \tag{1.5}$$

If t_a is selected such that equation 1.5 is not satisfied, then ground state amplitude may be lost to higher excited states. Thus, in this formulation of QA, the efficiency of the algorithm depends on how the gap minimum can be expected to scale with increasing problem size. Much work has been devoted to understanding the gap behaviour in annealing spectra and it is now generally agreed that the gap minimum can be expected to close exponentially for hard problems. We discuss these results in section 1.3.

The preceding discussion has equated the total run-time of the algorithm with the time required to approach a unity ground-state (GS) overlap with an anneal. In practice, however, QA as an algorithm consists of more steps than a single annealing run – e.g: pre-processing steps such as mapping the problem to a suitable Hamiltonian [27, 28, 29**?**] and embedding this Hamiltonian into the native hardware graph [30, 31, 32]. More relevant to the discussion in our work is the fact that QA will typically involve multiple annealing runs to increase the likelihood of measuring the ground state at the end of the anneal. A helpful metric to consider is the time to solution (TTS) which gives the total time needed to reach a certain probability of measuring the ground state given the final GS overlap obtained by the individual annealing runs. The expression for this is

$$\mathrm{TTS}(P_{\mathrm{GS}}) = \underbrace{\max\left(\frac{\log(1-P_{\mathrm{GS}})}{\log(1-|\langle\psi(t_{\mathrm{a}})|E_0\rangle|^2)}\,,1\right)}_{\text{required number of runs}} \times t_{\mathrm{a}} \tag{1.6}$$

where P_{GS} is the desired probability of measuring the GS. In general, an increase in t_{a} is expected to result in an increase in $|\langle\psi(t_{\mathrm{a}})|E_0\rangle|$ and so a decrease in the number of runs needed to reach P_{GS}.

In practice, restrictions set by the coherence time of the hardware mean that it may be preferable to select t_{a} such that P_{GS} is reached by performing an exponential number of annealing runs that all achieve an exponentially small overlap with the GS – rather than to attempting to run the anneal slowly enough for this overlap approach unity. We can

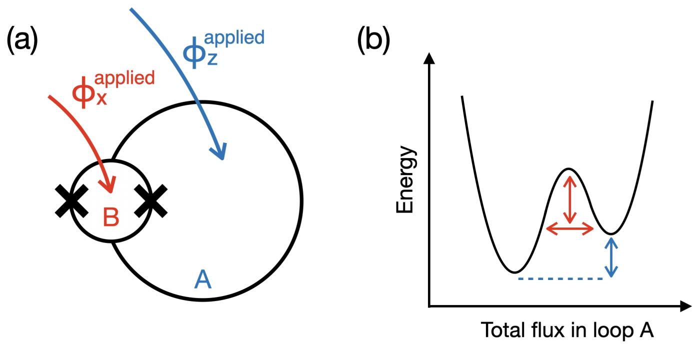

Figure 1.1: A schematic illustrating the key features of a superconducting flux qubit is shown in (a). The crosses denote the locations of the Josephson junctions. In (b) we present a cartoon energy spectrum showing the two lowest lying potential wells associated with the two states that define the qubit. The coloured arrows indicate the features of the energy landscape that can be tuned by varying the applied flux depicted in (a) in the corresponding colours.

think of this as the simplest case of moving the exponential time scaling out of the annealing time to facilitate practical implementation.

This dissertation focuses on the algorithmic side of QA. No explicit assumptions will be made regarding specific hardware limitations such as the coherence time and connectivity that would be associated with a particular platform. However, we will occasionally reference such limitations when discussing our results to help place our findings within the wider context of implementing the algorithm. To better contextualise these discussions, we will now briefly outline the main platforms on which QA has been implemented.

1.2 Hardware platforms

Unlike gate-based quantum computation, QA demands the capacity for smoothly varying time-dependent fields and coupling strengths in its implementation. As such, specialised hardware is required. The state of the art in quantum annealing devices is the D-Wave Advantage2 [33] which consists of around 4400 superconducting flux qubits [34] arranged in the so-called Zephyr topology, in which each qubit is connected to 20 others.

The flux qubits themselves are artificial spin-1/2 systems, where the two energy levels that define the qubit are the two directions of a circulating current around a ring of superconducting material. They consist, in their most basic form, of loops of superconducting wires and Josephson junctions [35] which introduce non-linearity to the response of the system to flux. This non-linearity introduces the anharmonicity required to separate off the two lowest energy levels to be used as a qubit.

A schematic of a superconducting flux a qubit is shown in figure 1.1(a) and a cartoon energy landscape showing the two states defining the qubit is shown in figure 1.1(b). By varying the flux threaded through loop A, the energy difference between the two qubit states can be altered, providing the mechanism for a tunable local Z-field. Meanwhile, the flux threaded through loop B controls the width and height of the potential barrier between the two qubit states, providing a tunable local X-field. ZZ-interactions are implemented through inductive coupling via intermediate coupling qubits. As such, the D-Wave architecture allows for the implementation of local X- and Z- fields as well as ZZ-couplings between qubits – which is sufficient to implement the standard annealing protocol outlined in the preceding section.

As we shall see in our discussion of the literature in chapter 2, there is reason to believe that the introduction of other 2-local couplings may be beneficial to QA. Particularly promising, and the focus of this thesis, is the inclusion of XX-couplings. Such couplings do not currently exist in superconducting annealing hardware. Some proposals have, however, been put forward as to their implementation [36, 37]. Furthermore, rather that implementing them directly, it has been suggested that XX-couplings could be emulated using currently available fields and couplings [38].

Coherent quantum dynamics have been demonstrated on the D-Wave Advantage [39]. However, significant effects from decoherence were seen for annealing times greater than a few tens of nanoseconds. That the coherence times observed in superconducting annealing hardware is significantly shorter than some of the reported coherence times for superconducting qubits in other settings (which can be on the order of $1ms$ [40, 41, 42]) is due to the increased demand for qubit control in QA necessarily introducing more noise channels which adversely affect the qubit lifetimes.

While not as well developed as the D-Wave systems, with regards to their use in quantum annealing, another promising platform for implementing QA is the Rydberg atom array [43, 44, 45]. Here, the qubits are encoded as the ground state and a highly excited Rydberg state of a neutral atom. The atoms are arranged using optical tweezers into the desired 2-dimensional graph structure for which nearest neighbour interactions are facilitated by the Ryderg blockade mechanism, which prohibits the excitation of multiple atoms to the Rydberg state within a certain radius.

This Rydberg blockade mechanism allows for the native implementation of the maximum-independent-set problem [46], This is an NP-complete problem in which one is looking for the largest set of vertices in a graph for which no two vertices are connected by an edge. This is achieved by arranging the atoms, which encode the vertices of the problem graph, such that any two that are connected by an edge in the problem are within each others Rydberg radius. Off-resonant homogeneous laser pulses are then used to drive transitions between the two qubit states with the Rabi frequency and detuning effectively giving us a way to set a global strength for the X- and Z- fields respectively such that an annealing protocol can be carried out. [1] More recently, the capacity for locally tuned Z-fields though local light shifts was demonstrated [47], facilitating the implementation of a weighted variant of the maximum-independent-set problem which is known to be harder to solve.

---

[1] In practice, the schedule used in [46] is a little different to the more traditional annealing schedule given by equation 1.1. However, the overall concept is very much the same.

Because of the planar nature of the arrays, only unit disk graphs can be implemented directly. However, a recent proposal for quantum wires – *i.e.*, chains of atoms in the blockade regime that can link different parts of the graph – may be able to mitigate this problem [48]. Another proposal for using Rydberg atoms to implement the quantum annealing algorithm, which inherently circumvents the connectivity issue (albeit with an additional qubit overhead), is to utilise the Lechner-Hauke-Zoller (LHZ) scheme [45, 49] in which the qubits are arranged in a grid lattice with problem-independent couplings. The problem parameters are then encoded into the local fields.

One advantage that Rydberg atom arrays have over superconducting devices is that they are inherently identical, avoiding potential fabrication errors. The associated downside of this, however, is that they do not possess the same tunability as a manufactured qubit. Another benefit of using Rydberg atom qubits is that they do not require the millikelvin temperatures that the D-Wave devices do. They also have significantly longer coherence times – on the order of $100\mu s$ [46].

It is worth noting that, even with improvements to coherence times in quantum annealing hardware, exponentially long annealing runs remain intractable without the error correcting framework that exists for gate-based quantum computation. As we will see in the following section, hard optimisation problems typically demand an exponentially long run-time for a significant ground state overlap to be achieved at the end of the anneal – hence, a desire for alternative methods in quantum annealing.

1.3 Annealing-time bottlenecks in QA

This section concerns gap- and time-scaling results for QA as defined by the total Hamiltonian in equation 1.1 and driver in equation 1.2. Our discussion focuses on the scaling associated with a single anneal (rather than the algorithm as a whole) and assumes the goal of adiabaticity. We will start with some general results before discussing, in more detail, the specific annealing-time bottleneck on which our work focuses.

Early studies that applied QA to small instances of various **NP**-hard problems provided a mixed outlook as to the expected time scaling for an anneal to reach a constant final GS overlap. In 2001, Farhi *et al.* published results that examined QA applied to random instances of the Exact-Cover problem [50]. Performing numerical simulations of closed system dynamics under the Schrödinger equation, they determined that the median run time required to obtain a constant overlap with the ground state at the end of the anneal scaled polynomially with the system size. Similar results were obtained by Hogg for the k-SAT problem [51]. However, work by Smelyanskiy *et al.* that looked at QA applied to the Set-Partition problem found that the median scaling appeared to be exponential [52]. Looking at the annealing spectra, they associated this exponential scaling with the closing gap at an avoided level crossing (AC) that appeared between the instantaneous ground and first excited states. It was noted that the scaling only became exponential above $n = 10$, indicating that studies on small system sizes may not in general reveal the asymptotic scaling behaviour.

Quantum Monte Carlo (QMC) simulations allowed for investigations up to much larger system sizes. As with the initial small system Schrödinger studies, the early results were promising. In 2008, Young *et al.* published a paper that looked at random instances of the Exact-Cover problem up to system sizes of $n = 128$ and found that the median minimum gap appeared to close only polynomially with system size [53]. However, a subsequent study by the same authors that focused specifically on instances with only one solution (which are hard to solve classically) gave less favourable results [54]. Looking at system sizes up to $n = 256$, they found that the fraction of instances with first order phase transitions tended to one with increasing system size. In 2011, Hen and Young published a paper that studied the gap scaling behaviour for various SAT problems and found the scaling to be exponential [55]. They also found that the the instances that were harder classically also tended to be harder for QA.

Other studies have attempted to understand the gap scaling behaviour through analytical methods. A paper by Znidaric and Horvat published in 2006 derived an expression for the gap minimum for 3-SAT problems [56]. This expression predicted an exponentially closing gap minimum and gave good agreement with numerical results up to $n = 30$. However, the driver Hamiltonian used in this study was not the standard homogeneous local X-field and it is not clear how these findings translate to the more standard setting. Taking a different route, some insight was obtained by Amin as to what properties in a problem Hamiltonian might result in small gap minima [57]. Using perturbation theory it was determined that problems with exponentially-many local optima that are close in energy to but far apart in Hamming distance from the global optimum would have exponentially closing gaps in their annealing spectra. Again, this suggests that problems that are harder to solve classically will also be harder for QA.

Rather than considering optimisation problems, some analytical work has focused on more analytically tractable physical models. The presence of a first order phase transition between a quantum-paramagnetic (PM) and a spin-glass (SG) phase has been analytically confirmed for both the random energy and the p-spin model [58, 59, 60]. The transition that occurs between the quantum PM and the SG phase in the Hopfield model has been shown to only be of second order [61] such that the associated gap minimum can be expected to close only polynomially with the system size. However in this case it was shown that, depending on how the patterns are embedded, the SG phase could contain multiple tunnelling bottlenecks due to competing local optima. These tunnelling occurrences were found to result in exponentially closing gaps.

Overall, these results suggest that we can generally expect exponentially closing gaps to appear in annealing gap spectra – sometimes as a result of the transition out of the quantum paramagnetic phase in which it

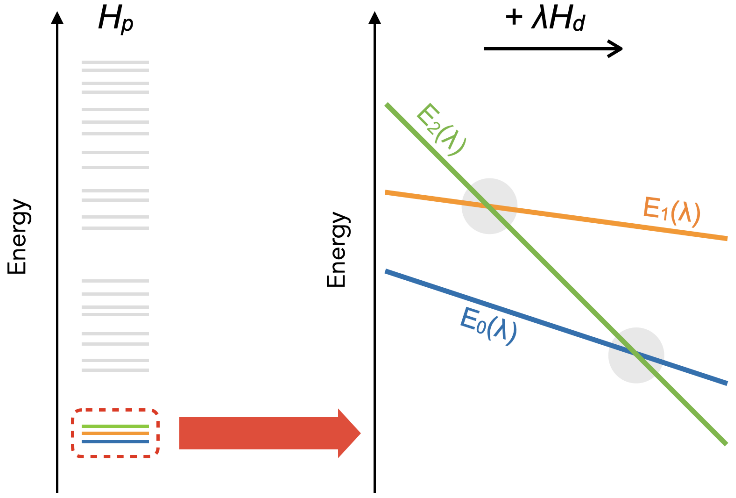

Figure 1.2: Cartoon illustrating the formation of two perturbative crossings. On the left we see a cartoon problem spectrum for which the 1ES and 2ES are close in energy to the GS. On the right, we show example driver perturbations that result in the formation of two perturbative crossings as the perturbed problem 2ES crosses the perturbed problem 1ES and GS.

starts but also as a result of structure in the problem Hamiltonian. In particular, problems that are hard in a classical setting seem to generally also be hard for QA due to tunnelling into the ground state, which was posited as a potential mechanism for speedup in QA, taking exponentially long.

1.3.1 Perturbative crossings

In this section we highlight a particular source of small gaps in annealing spectra known as perturbative crossings. These are localised to localised transitions that can occur towards the end of the annealing spectrum as a result of highly competitive local optima. Here, localisation refers to a state being restricted to a small portion of Hilbert space consisting of states that are close in Hamming distance. Perturbative crossings have also been linked to the idea of Anderson localisation – a phenomenon in which disorder causes a system to become localised [62].

The formation of perturbative crossings can be understood by introducing $\hat{H}_d$ as a perturbation to the problem Hamiltonian and seeing if the energies of any of the perturbed excited states cross with that of the perturbed problem ground state [62, 63]. If such a crossing is predicted by perturbation theory then we can expect the annealing spectrum to exhibit an AC at the corresponding point in the anneal. It is this feature that we refer to as a perturbative crossing. A cartoon illustrating the formation of two such perturbative crossings is presented in figure 1.2. On the left, a cartoon problem spectrum depicts a setting where the first and second excited states are close in energy to the ground state. On the right, example driver perturbations are shown that result in the perturbed 2ES energy crossing both the perturbed 1ES and GS energy resulting in two perturbative crossings.

It is clear that the lower in energy an excited state is, the more likely it is to be involved in a perturbative crossing with the ground state since a smaller perturbation will be required for a crossing to occur. Further to this, insights have been obtained by considering the general form of the perturbative corrections as to (a) precisely under what conditions we can expect a perturbative crossing to occur and (b) what the size of the gap minimum will be at the AC.

By examining the perturbative corrections to second order it is found that, for the states with energy comparable to that of the ground state, those which are coupled by the driver to lower energy states will receive greater perturbations [63, 64]. Thus the perturbed energy of an excited state is likely to cross that of the ground state if (a) it is close in energy to the ground state and (b) the states that are one spin flip apart from it are significantly lower in energy that those that are one spin flip apart from the problem ground state. Following the terminology in [65], we can refer to the set of states one spin flip apart from a particular state as its neighbourhood. That these conditions produce a bottleneck in the algorithm can be understood intuitively as QA having to avoid a wide local optimum [63]. Further, degenerate states that are precisely two spin flips apart have also been shown to produce perturbative crossings with the ground state [63, 64]. This can be understood through level repulsion.

The finite gap at the point of the crossing predicted by perturbation theory is associated with the tunnelling amplitude between the two perturbed problem states and will thus depend on their overlap. It has been noted that a perturbative crossing that forms closer to the end of the anneal can be expected to be associated with a smaller gap [63] due to the fact that the problem states will have undergone less perturbation at the point of the crossing and will therefore be less mixed and have a smaller overlap with each other. The gap size has also been found to be exponential in the Hamming distance between the problem states involved in the perturbative crossing [63]. Since this Hamming distance can generally be expected to grow linearly with system size, this implies an exponentially closing gap.

These analytical findings have been confirmed numerically for a variety of systems. In [65] a tunable maximum-weighted-independent-set problem on a spin chain is used to demonstrate that if a non-degenerate $|E_1\rangle$ has a lower energy neighbourhood than $|E_0\rangle$, a perturbative crossing will form. It is demonstrated in [66] how a perturbative crossing can be predictably created in 3-SAT problems, and [63, 67, 68] contain examples of crossings forming between the ground state and a degenerate excited state. In addition to these results, which show that such crossings can form for the reasons theoretically predicted, it has also been suggested that the likelihood of not observing a perturbative crossing with the ground state will vanish with increasing problem size [62]. This suggestion stems from the fact that the chance of a perturbative crossing forming grows with the number of states with energy comparable to that of the ground state and this is generally expected to grow exponentially with the system size.

1.4 Thesis Outline

In this introductory chapter, we have outlined how quantum annealing operates in its most standard form and discussed the exponential gap (and therefore annealing time) scaling that is expected to be present for the algorithm in this form. We have also drawn attention to a particular source of exponentially closing gap known as a perturbative crossing – which is the annealing-time bottleneck on which this work focuses.

This rest of this dissertation is organised as follows. In chapter 2, we give an overview of some of the variations on QA that have been proposed in order to address the problem of exponentially closing gaps. There are a great many avenues that people have explored in this respect including counter-diabatic diving, in which an additional Hamiltonian is introduced to suppress diabatic transitions [69, 70, 71, 72, 73, 74, 75], and reverse annealing, in which the system is initialised in with diagonal Hamiltonian rather than a transverse field [70, 76, 77, 78, 79, 80]. We will be focusing our attention on methods that involve changing the path that the anneal takes to $\hat{H}_p$, through either inhomogeneous driving or the introduction of what is known as a catalyst Hamiltonian, in order to favourably alter the annealing spectrum.

We will begin by outlining some general strategies that have been considered before moving on to targeted approaches that utilise information about the problem in some way. We also give an overview of the literature relating to *diabatic* quantum annealing, in which transitions into higher excited states are exploited in order to facilitate a faster route to the final ground-state. The possibility of diabatic approaches being utilised is something that we will be returning to at various points when discussing our results.

Chapter 3 concerns the problem instances we will be utilising throughout this work in order to facilitate our numerical investigations. We begin this chapter by outlining the theory behind perturbative crossings in some more detail. We then describe how we construct our problem instances and how their structure relates to the formation of perturbative crossings. Finally, we present numerical results showing that we are able to use these problem instances to produce annealing spectra with the desired properties. Our numerical methods are then outlined in chapter 4. These are included after chapter 3 since they include a Hilbert space reduction that makes use of symmetries present in our problem instances.

Chapters 5-9 contain our results. Chapter 5 introduces the idea of a targeted XX-catalyst for the removal of perturbative crossings. We begin by examining the effect of different XX-couplings using two different perturbative approaches. This analysis allows us to build a picture of how XX-couplings could be chosen in order to mitigate the appearance of a perturbative crossing. We then test these ideas numerically with a few example systems. Overall, we find the effects of the catalyst to be in line with the theoretical predictions and that we are able to extend our analysis to account for further details in our observations. We also discuss some potential ideas for how and when such catalysts could be utilised. In particular, we note that it may be possible to identify suitable XX-couplings using information obtained from catalyst free annealing runs.

The main focus of this dissertation is on the use of XX-couplings. However, in chapter 6, we take a brief detour to examine the effect of other 2-local catalyst choices. We begin with a numerical investigation before using the same theoretical framework as we did in the XX case to shed some light on our results. For the majority of cases, we find that the results from the perturbative analysis match up well with what was observed numerically. Notably, however, we were not able to account for the success of the catalysts involving just one σ^y operator – *i.e.*, the catalyst Hamiltonians with complex components. We end by discussing how these findings could be utilised in the design of effective catalysts.

We then turn our attention back to the targeted XX-catalysts introduced in chapter 5. In chapter 7 we carry out a numerical investigation examining how the capacity of the proposed catalyst to remove a single perturbative crossing scales with system size – as well as how this capacity is affected by the number of couplings included in the catalyst and the severity of the gap scaling in the original annealing spectrum. Our results suggest that a catalyst containing a number of couplings that scales linearly with system size may be sufficient to remove a perturbative crossing, regardless of the severity of the initial gap scaling. As part of this investigation, we explore the effect of introducing a single targeted XX-coupling. Here we observe additional effects from the catalyst which differ depending on the severity of the original gap scaling and include the manifestation of new gap minima in the annealing spectrum. We discuss potential causes for the appearance of these new gap minima as well as the possibility of utilising them as part of a diabatic anneal.

In chapter 8, we use one of these scalable instances as a toy model to investigate the robustness of diabatic annealing in settings where the spectrum is being manipulated in order to facilitate the diabatic anneal. As we will be discussing in section 2.3, the success of diabatic annealing is likely to be far more sensitive to the choice of annealing time than an adiabatic anneal. In settings where the annealing spectrum is being altered to create a diabatic path for the system to follow, any enhancement to the final ground-state overlap will also be sensitive to any parameters involved in this alteration. We begin with a numerical study in which we first confirm a possible enhancement to the final GS fidelity. Then, varying the parameters associated with the anneal, we observe a trade-off between the precision needed in the annealing-time chosen and how the catalyst is introduced. Finally, we introduce a Landau-Zener model to understand the physics behind our observations, allowing us to consider how these findings would extend to more general systems and methods for diabatic annealing.

In our final results chapter (chapter 9) we extend the targeted XX-catalysts discussed so far to settings where multiple perturbative crossings are present. From the discussion in [62], we can surmise that a realistic problem setting is likely to contain a multitude of such crossings. As such, for a catalyst to be an effective solution to this particular bottleneck, it must have the capacity to address multiple perturbative crossings at once. In the first half of this chapter, we consider the capacity of a catalyst consisting of multiple components to remove several perturbative crossings, thereby facilitating an adiabatic anneal for shorter run-times. We then, in the second half of the chapter, examine the possibility of using these catalysts to construct more complex diabatic paths through the spectrum. Our initial results appear promising however we feel that further investigations will be needed before any concrete conclusions can be made.

We conclude this dissertation, in chapter 10, by summarising and discussing our findings, highlighting some key results, and proposing avenues for future work.

Chapter 2

Alternative approaches in quantum annealing

In this chapter we discuss some of the many variations upon the most straightforward formulation of QA, as defined in section 1.1, that have been proposed in order to circumvent the problem of exponentially closing gaps. We will predominantly be focusing on strategies which aim to prevent the appearance of small spectral gaps by altering the path that is taken to the problem Hamiltonian, $\hat{H}_p$. Two ways by which this path-change can be achieved are by introducing inhomogeneity to the driver or by introducing a third term into the annealing Hamiltonian, usually referred to as a catalyst. The potential for success of these approaches will be discussed in sections 2.1.1 and 2.1.2 respectively. There has also been an interest in developing strategies that make use of pre-existing or accessible knowledge of the problem in order to alter the annealing path in some targeted way and this will be discussed in section 2.2 – with targeted inhomogeneous driving being discussed in section 2.2.1 and targeted catalysts being discussed in section 2.2.2.

While we may not expect these strategies to result in a sub-exponential time scaling for the algorithm, they are still appealing for a couple of reasons. Firstly, even if the closing of the spectral gap minimum remains exponential, there may be a quantitative suppression of the scaling exponent allowing for a faster time to solution. Furthermore, a strategy that takes exponentially long to find a path with a polynomially closing gap minimum in the associated spectrum would allow the exponential scaling to be moved out of the annealing run-time. This potentially allows for a single anneal to reach a near unity overlap with the GS without exceeding the coherence time of the hardware. The run-time of the algorithm is then a question of the severity of the exponential scaling associated with finding this annealing path.

For the most part, the discussion in sections 2.1 and 2.2 assumes the goal of adiabaticity. However, it has also been suggested that a faster time to solution may be obtained by utilising the diabatic transitions that occur when the adiabatic condition is not met. In this paradigm, known as diabatic quantum annealing (DQA), the idea is to exploit a path through the instantaneous state spectrum which will result in a high overlap with the ground state at the end of the anneal but which can be followed for shorter run times. Where and how this may be achieved will be discussed in section 2.3.

2.1 Non-targeted approaches

We begin by discussing the non-targeted approaches that exist in the literature to changing the annealing path – *i.e.*, approaches which are applied without any prior knowledge of the problem. We will start with inhomogeneus driving strategies, in section 2.1.1, before going on to the use of catalyst Hamiltonians, in section 2.1.2.

2.1.1 Inhomogeneous driving

One straightforward way of introducing inhomogeneity into the driver is simply to randomise the local field strengths in some way. The possibility of such a strategy being able to remove perturbative crossings was examined for 3-SAT problems [66] with the results indicating that the number of drivers one would need to try before an appropriate path was found would scale polynomially with the number of low energy local optima. Since this number is expected to scale exponentially with

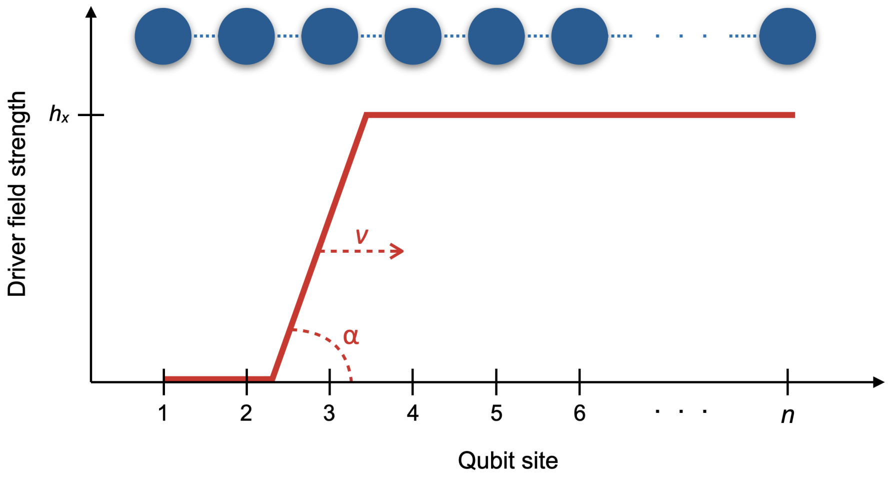

Figure 2.1: Plot illustrating how the local driver fields are varied across an Ising chain in [3]. Here, h_x denotes the strength with which the driver is initialised on each qubit. The driver schedule is characterised by a slope with angle α moving across the chain with a speed of v.

the problem size, the overall scaling of the algorithm would remain exponential. However the annealing runs need only be polynomial in the problem size.

More recently, the effect of highly *non-local* drivers with random field-strengths were investigated [**?**]. These drivers, based on a bosonic spin version of the Sachdev-Ye-Kitaev model, were found to reduce the annealing time by several orders of magnitude in the best cases. Limits to the system sizes that could be investigated meant that meaningful conclusions regarding the scaling could not be reached. However, it seems likely that these non-local drivers would perform at least as well as the randomised *local* drivers since more connectivity betweens states should allow the system to more easily explore the Hilbert space.

The idea of better connectivity within the Hilbert space being beneficial to QA in particular was considered more rigorously in an analytical and numerical investigation led by Roopayan Ghosh [81] with which we were involved. In this work, it was found that including more couplings and of higher order will tend to increase the success probability of the algorithm. In the extreme case, where all couplings, are included we are guaranteed to be connecting the global optimum with any local optimum the algorithm might otherwise become trapped in. Sets of couplings that fall short of this case can still perturbatively couple the global and local optima – with more and higher order couplings offering routes between the optima that require fewer orders of perturbation theory. (This work considers the introduction of additional couplings in the form of a catalyst hamiltonian however the same arguments would apply to a driver.)

Of course, higher order couplings are generally harder to implement in practice. It is also worth noting that, while the use of more chaotic drivers may improve the final GS fidelity, they would be significantly harder to initialise in than the standard homogeneous transverse field due to complex and potentially unknown ground states. Indeed, if our new choice of driver has a GS which is **NP**-hard to find, we have simply replaced one hard problem with another. As we will go on to explain at the start of section 2.1.2, this is an advantage that the introduction of a catalyst Hamiltonian has over changing the make-up of the driver Hamiltonian. Indeed, the authors of [**?**] note this and also apply their drivers in the form of a catalyst in the same work.

Another possibility is to alter the driving schedule in some way. A particular way this can be achieved is to stagger the times at which the local X-fields are reduced [3, 82]. This can be understood through figure 2.1 which illustrates this idea for a spin chain. It can be seen that as the slope moves across the chain at a speed v, only some of the qubits will have their transverse field strengths decreasing at a given time. The rate of this decrease will depend on v as well as the gradient of the slope, α. This driver is motivated by the idea that, if different parts of the system move through a phase transition at different times, excitations may be suppressed as a result of the parts of the system that have already crossed the phase transition communicating their state back to the parts that have not.

This idea has been explored with a ferromagnetic (FM) Ising chain in terms of the Kibble-Zurek (KZ) Mechanism. As the transverse field is decreased, the system must pass from a PM to a FM phase however if the driver is reduced too rapidly, the final state may contain excitations in the form of "kinks" between spin domains. The faster the phase transition is crossed, the higher the kink density and thus the residual energy at the end of the anneal. In [82] an expression is derived which suggests that there is a critical velocity that depends on the KZ correlation length and the energy gap at freeze-out. If the rate at which the critical point propagates through the system is below this, the excitations could be suppressed. These analytical results are supported with numerical simulations of the dynamics. Later work [3] further explored these ideas more in the context of QA. The authors considered a weakly disordered spin-chain and examined the dependence of the residual energy on the angle α and the velocity v in figure 2.1. They found that for very short anneals homogeneous driving was optimal, but for longer anneals, inhomogeneous driving was able to reduce the residual energy of the solutions by several orders of magnitude. Their results also suggested that this strategy was relatively insensitive to the hardness of problem.

A related protocol was proposed in [83]. In this case the driver field strengths are homogeneously reduced for all qubits as in the standard case but with the fields being totally turned off for each qubit in turn as the anneal progresses. It was shown analytically that this could remove the 1st order phase transition present in the p-spin model. The authors support this finding by calculating the gap sizes with and without this inhomogeneity in the driver both by using a semi-classical analysis and numerically. They also demonstrate that this strategy works when random local Z-fields are introduced into the model which is understood to be a harder problem. A later paper [84] by the same authors determined that if non-deal conditions were considered (*e.g.* finite temperatures), a first order phase transition may not be avoidable. However, a quantitative improvement to the tunnelling amplitude was still possible. Similar results to this were found for the Weak-Strong-cluster problem [85]. In this work dynamic simulations were also performed and it was found that not only was the ground state overlap significantly improved but also that the improvement ratio grew with system size. While these results suggest that some form of staggered driver schedule may be beneficial, it is still unclear whether or such a strategy would provide the same advantage in a more realistic problem setting.

2.1.2 Catalysts

Rather than altering the driver terms or schedule, another way to change the annealing path is to introduce an additional term into the annealing Hamiltonian. This Hamiltonian is typically introduced with a pre-factor of $s(1-s)$ into equation 1.1 such that its contribution is zero at the start and end of the anneal. An inherent benefit of this is that the initialisation of the system is unaffected. Before we begin this section, we introduce the concept of "non-stoquasticity" which will be relevant to our discussion. A Hamiltonian is stoquastic if all the off diagonal elements are real and non-positive. Conversely, a Hamiltonian is said to be non-stoquastic if some of the off diagonal elements are positive and/or complex[1]. There is a great deal of discussion in the literature regarding the importance of non-stoquastic Hamiltonians and how they differ to stoquastic Hamiltonians – we will discuss these points when and where they become relevant.

The idea of introducing a random additional Hamiltonian was explored analytically [86] for a highly symmetric 3-SAT problem for which QA exhibited an exponentially closing gap as a result of getting stuck in a local optimum. By examining the effective potentials it was determined that 351/1000 of the randomly generated catalysts were effec-

---

[1] In some cases, the term "non-stoquastic" is 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. In this thesis, we use the term non-stoquastic simply to describe the off-diagonal elements of a Hamiltonian in a particular basis.

tively able to guide QA to the global optimum thereby removing the exponentially closing gap. More evidence that randomly generated catalysts may be able to enhance the success of QA came from a numerical study applying catalysts with the same connectivity as the problem graph to 20-qubit max-2-SAT problems that were hard for QA when a standard linear interpolation was used [87]. Specifically, it was found that stoquastic catalysts almost always enhanced the final ground state overlap – albeit not by very much. Meanwhile, catalysts with complex off-diagonal elements didn't help as often but produced a much greater GS overlap enhancement when they did. Similar results were found for catalysts that contained only local fields [88]. It was found in [87] that the improvement to the final GS overlap generally coincided with an enhancement of the gap minimum. Interestingly however, some cases where observed where complex, *i.e.*, non-stoquastic, catalysts were able to greatly enhance the final GS overlap while the size of the gap minimum remained very small. We will return to this point in section 2.3.

There has also been a great amount of interest in whether catalysts of certain forms could offer a general advantage. Specifically there has been much investigation into the effect of stoquastic and non-stoquastic XX-couplings. A study examining random instances of Ising spin glasses reflected the random catalyst results in that introducing a stoquastic homogeneous XX-field was found to enhance the success probability more often that introducing a non-stoquastic XX-field did [89]. This was again reflected in the effect on the gap minimum which was generally enhanced using the stoquastic catalyst and reduced in may instances using the non-stoquastic catalyst. The stoquastic and non-stoquastic catalysts were also generally found to reduce and increase the frequency of ACs respectively. The same kind of results were found when applying these catalysts to 18-variable 2-SAT problems [90]. However in this case, it was determined that the effect of the non-stoquastic catalyst was highly dependent on the strength with which the catalyst was introduced – with the fraction of instances that showed gap enticement increasing with the catalyst strength. They also found that the non-stoquastic catalyst would in some cases cause the AC present in the original spectrum to become stretched and this was also found to enhance the final GS probability. As in [87], both [89] and [90] found cases where the the success of QA was improved without the gap minimum being enhanced (or the AC becoming stretched).

The various findings that stoquastic catalysts are generally more successful at providing gap enhancement are supported by analytical results comparing the energy spectra of non-stoquastic Hamiltonians to their stoquastic counterparts [91]. For local Hamiltonians that are diagonal in the X-basis, it was shown that the gap between the ground and first excited state could always be increased by making the Hamiltonian stoquastic in the Z-basis. This work also utilised ideas from spectral graph theory to study the properties of the low lying energy spectra of general Hamiltonians and were able to offer some intuition as to why stoquastic Hamiltonians can generally be expected to have larger gaps between their ground and first excited states.

However, there are also examples where non-stoquastic XX-catalysts have been shown to result in significant gap enhancement. A numerical study examining the size of the gap minimum in the annealing of the weak-strong cluster problem found that if the catalyst strength was optimised for each system size, a non-stoquastic XX-catalyst could replace the exponentially closing gap present in the linear anneal with a gap minimum that tended to a constant [67]. The gap scaling remained exponential when a stoquastic catalyst was used. The study also examined a geometrically local Ising example and found that a non-stoquastic XX-catalyst removed the perturbative crossing present in the original annealing spectrum. This effect was linked to the fact that the catalyst coupled the local optima responsible for the perturbative crossing together.

The effects of XX-catalysts have also been studied analytically for the p-spin and Hopfield models – two settings where a first order phase transition is present for the standard linear anneal. It was shown that introducing non-stoquastic XX-couplings to the p-spin model allowed for a path through phase space that avoided all first order phase transitions so long as p was finite and greater than 3 [92]. (In the case of $p \to \infty$ the first order phase transition persists.) This finding was confirmed by numerical results which showed that for a large enough catalyst strength, a non-stoquastic XX-catalyst could reduce the exponentially closing of the gap minimum to a polynomially closing gap [67]. This is in contrast to the stoquastic XX-catalyst which was not able to achieve the same results. Similar results were found for the Hopfield model [93].

The success of a non-stoquastic XX-catalyst in removing the first order phase transition in the p-spin model was initially attributed to the fact that the ground state of the p-spin model had a greater overlap with the GS of this catalyst than with that of the driver (and indeed of the stoquastic XX-catalyst). However, it was later found that higher order non-stoquastic XX-catalysts of the form $(\sum_i \hat{\sigma}_i^x)^k$ had similar effects for both odd and even k [92] – with the overlap between the p-spin ground state and that of the odd k catalyst being exponentially suppressed with increasing system size compared to the GS of the driver Hamiltonian. Another suggestion for why the non-stoquastic XX-catalysts were able to remove the first order phase transition came from a study that examined how its introduction affected the trace-norm distance between the semi-classical spin coherent states and the true ground state, as well as its effect on the level of entanglement present in the system [94]. It was found that both of these quantities were increased by the introduction of the catalyst at the location of the phase transition suggesting that additional quantum effects introduced by the non-stoquasticity may be responsible for the removal of the first order phase transition.

However, it was determined that neither stoquastic nor non-stoquastic XX-couplings were able to remove the 1st order phase transition present in the anneal for the frustrated Ising ladder [95]. It was suggested this may be due to the fact that the two phases associated with the phase transition are separated by different topological structures and the phase transition may therefore be more stable to the introduction of additional terms. Further, it was found that for the weak-strong cluster problem studied in [67], whether a stoquastic or a non-stoquastic XX-catalyst was able to remove the first order phase transition depended on where the couplings were applied [96].

Overall, the results suggest that the effect a given catalyst might have is highly dependent on the system that it is applied to.

2.2 Targeted approaches

That there may not be one single approach that can produce annealing spectra with a polynomially closing gap minimum is perhaps unsurprising since this would imply a polynomial algorithmic run-time. The results presented in the previous section suggest that it may be beneficial to use a different random driver or catalyst for each polynomially long annealing run performed. The number of runs may still need to be exponential in the problem size however at least some of the annealing runs would be likely to achieve a high overlap with the final GS – thereby allowing a quantitative improvement to the scaling over the case where where the standard interpolation is used.

Another route would be to develop strategies that utilise pre-existing or accessible information about a given problem to engineer some targeted path change. Again, we may still expect the scaling associated with finding this path to be exponential – but by guiding the search in some way we may expect a more favourable exponent. In this section we outline a few such ideas from the literature both with regard to inhomogeneous driving and the introduction of catalysts.

2.2.1 Inhomogeneous driving

One avenue for incorporating problem specific information into the driver is to utilise information about the couplings within the system which, unlike the eigenspectrum of the problem Hamiltonian, are known a priori. It has been proposed that a pre-processing step could be used to construct a targeted version of the staggered catalysts discussed in section 2.1.1 where multiple critical fronts are introduced based on the problem structure [97]. This strategy was found to be able to reduce the scaling of previously exponentially closing gaps to sub-polynomial for a variety of Ising Hamiltonians. This kind of approach has also been explored to mitigate the problem of freeze-out – a phenomenon in which certain qubits become frozen long before the computation is completed. In this case, the hypothesis that more strongly coupled qubits would freeze out earlier than those that are only weakly coupled to the rest of the system was utilised to identify the qubits that were more prone to freeze out. Their annealing schedule was then delayed [85]. Testing this approach using the D-Wave 2000Q, it was determined that the TTS could generally be improved with the hardest instances showing the greatest time reduction – often by an order of magnitude.

Rather than altering the schedules, one can of course introduce inhomogeneity into the fields themselves. An analytical and numerical study looking at the 1D Ising model showed that it may be possible to mitigate disorder in the problem Hamiltonian, responsible for the presence of small gap minima, through the introduction of appropriately chosen driver fields [98]. The strengths of the driver fields are chosen based off the adjacent problem coupling strengths which can simply be read off the Hamiltonian.

Another route would be to utilise information about the *eigenspectrum* of $\hat{H}_p$ which is not known a priori, but which it may be possible to obtain in some way. For instance, if the only exponentially closing gaps in the annealing spectra are associated with perturbative crossings then we can expect QA run in polynomial time to return precisely the local optima responsible for perturbative crossings that take the system out of the ground state. Given that it is the neighbourhoods around the problem states (as defined by the driver) that dictate whether or not the perturbative crossings forms, an intuitive approach to removing them is to tune the relative strengths of the local fields – effectively adjusting the relative contributions from the different states in the neighbourhoods. That such an approach can work was demonstrated for the maximum independent set (MIS) problem with an assignment of local fields that only required knowledge of the local optima [64]. This approach was later worked into a recursive strategy in which the driver fields are adjusted after each set of annealing runs, based on what solution states are returned [99]. Simulating this procedure for fifty 64-vertex MIS problems, for which perturbative crossings existed in the original spectrum, it was found that 13 iterations were sufficient to produce an annealing spectrum that allowed QA to reach the ground state in polynomial time.

2.2.2 Catalysts

We now consider how problem-specific information could be utilised to construct a catalyst Hamiltonian that changes the annealing path in some targeted way. One straightforward approach that has been explored is to select a particular form for the catalyst that has free parameters and then optimise these parameters for a particular problem instance. A numerical study applying this approach to randomly generated quadratic unconstrained binary optimisation (QUBO) problems suggested that a catalyst containing only single qubit operators could be optimised to produce an annealing spectrum where the smallest spectral gap between the ground state and the first excited state was at $s = 1$ [88]. Since this gap is only dependent on the problem and not the path taken, this suggests that the optimised catalyst has successfully provided an optimal route from the driver to the problem Hamiltonian. In the procedure, the parameter that was being optimised was the minimal spectral gap. In practice however, one does not have access to this without having already solved the problem and so optimising the catalyst parameters in this way is not a tractable solution. In order to implement such a strategy more practically, one could imagine a similar procedure where the energy of the lowest energy state obtained by a set of annealing runs was used in place of the gap minimum. More generally however, these results suggest that a catalyst choice may always exist that removes the bottlenecks in an anneal.

Another idea that has been introduced is to use a diagonal catalyst that biases the Hamiltonian towards the global optimum [68]. For the p-spin model it was shown that so long as sufficiently many spins are biased in the correct direction the 1st order phase transition could be removed for an optimised catalyst strength. That the phase transition re-appears for catalyst strengths greater than this can be understood as a result of the incorrectly biased spins. Similar results were obtained for the weak-strong cluster problem. The feasibility of this as a strategy however, depends on how readily one can identify a state that is close in Hamming distance to the ground state without knowledge of the ground state itself. One might consider that knowledge of the local optima, which could be identified through fast annealing runs, classical, or hybrid methods, could be used in some way. However, as the authors note, in many cases the local optima are very far in Hamming distance from the global optimum. In fact, such cases are likely to correspond to the hardest QA instances [63]. Nevertheless, it remains a possibility that such catalysts may be of use provided some mechanism for selecting a state to bias towards is determined.

As discussed in section 2.1.2, there has been a particular interest in the use of XX-couplings in catalysts. More recently, there has been a focus on understanding how the effect a particular XX-catalyst has on the spectrum relates to the couplings it introduces between problem states. We noted earlier that the success of a non-stoquastic XX-catalyst in removing the perturbative crossing for a particular geometrically-local Ising example was attributed to the fact that this catalyst coupled the local optima responsible for the perturbative crossing to each other [67]. A similar phenomenon was observed for a maximum-weighted independent set (MWIS) problem [100], however in this case the catalyst was specifically crafted to introduce the desired couplings between the local optima. As in [67], it was demonstrated both analytically and numerically that introducing the catalyst with the same sign as the driver, such that the total Hamiltonian was stoquastic, had the opposite effect of reducing the size of the gap minimum. (We note that in this setting, additional structure in the problem means that the effect of the non-stoquastic catalyst becomes more complicated for higher catalyst strengths – this will be discussed further in section 2.3.) In addition to demonstrating that such a targeted catalyst has the desired effect, this work also proposed a method for determining the appropriate couplings in polynomial time using information obtained from fast annealing runs with the standard schedule.

That there may be a way to obtain a gap spectrum with a polynomially closing gap in polynomial time is a more promising result than is generally expected. While it is not impossible that poly-time quantum algorithms could be found for solving **NP**-hard problems, it is worth noting that the strategy proposed in [100] makes some assumptions regarding the structure of the problem which might account for the efficiency of the proposed algorithm. That being said, the required graph structure is precisely that which causes the perturbative crossing in the original annealing spectrum, and so the method may be applicable to some appreciable fraction of hard MWIS instances.

Another example of targeted XX-couplings being introduced to remove a perturbative crossing can be found in [65]. (Here the couplings are simply added to the driver rather than in a catalyst. The analytical reasoning for its success, however, would also apply had they been introduced in a separate Hamiltonian.) In this case the perturbative crossing in the original annealing spectrum is the result of a non-degenerate first excited state with a significantly lower energy neighbourhood than the ground state – *i.e.*, the 1ES is coupled to more low energy states by the driver than the GS is. Here the couplings are introduced with the same sign as the driver (such that the total Hamiltonian remains stoquastic) and couples the ground state low energy states, while coupling the first excited state to significantly higher energy neighbours. The catalyst effectively extends the neighbourhoods of the two states in such a way that the ground state becomes the state with the lower energy neighbourhood. The approach is not dissimilar to the idea of adjusting the local field strengths [64] described in 2.2.1. However in this case, rather than tuning the weight of the contributions from different neighbours, new neighbours are being added.

We note that the author of [65] does not propose this as a strategy for removing perturbative crossings but rather the results are used to highlight the importance of the neighbourhoods around the low energy states. Indeed, choosing a catalyst based on to which states it will couple the problem ground state is, of course, not a practical strategy as it assumes knowledge of the solution to the problem.

2.3 Diabatic quantum annealing

Rather than attempting to enhance the gap minimum in order to facilitate adiabaticity, another route being explored to reduce the run-time of QA is to exploit transitions to higher energy levels in some way. Annealing protocols that do this are known as Diabatic Quantum Annealing (DQA) protocols. Crucially, not having to remain in the GS means that exponentially closing gaps do not necessarily mean the anneal must be run exponentially slowly.

We have already discussed how hardware restrictions may make it necessary to run exponentially many short anneals for which adiabaticity is not preserved. In this situation, the diabaticity is not helping the anneal reach the GS but rather is a unwanted side effect of running the anneal within the coherence time of the hardware that we must compensate for.

However, there are also examples where allowing the system to evolve diabatically is useful in helping the system reach a higher overlap with the GS.

The first, and arguably still the most decisive, example of this is for the oracular glued trees problem [19]. As a result of its symmetric nature, two exponentially closing gap minima form in the annealing spectrum which can be exploited. Specifically, rather than running the anneal slowly enough that the evolution remains adiabatic, an annealing time can be selected for which the system transitions into the first excited state at the first gap minimum – since it can transition back into the ground state at the second. By carefully studying the gap spectrum and the dynamics both at and away from the exponentially closing gaps, the authors are able to show that QA offers an exponential speedup over classical algorithms for this problem.

While not offering as decisive a speedup, there have since been other examples of utilising transitions out of the ground state to enhance the success of the QA algorithm. It has been noted in multiple settings [87, 101, 102, 103] that when restricted to modest annealing times, faster run times actually achieved a greater overlap with the ground state at the end of the anneal. This is accounted for by the fact that if there is a small gap towards the end of the anneal then it may be beneficial for some of the amplitude to leak out of the ground state prior to this point. That way some of the amplitude may be returned to the ground at the location of the gap minimum – rather than the evolution being totally adiabatic until the small gap at which point all the GS amplitude may be lost.

A similar idea was discussed in [104] which looked at permutation symmetric problems. Here it was found that the system was able to return to the ground state through a cascade of diabatic transitions after spending the majority of the anneal in higher energy states. Utilising this, DQA was found to be optimal over Adiabatic QA (AQA) in this setting. It was later suggested however that this diabatic speedup is lost when the symmetry of the problem is broken [105]. In addition to this, there is also the issue of how precise the annealing time must be in order to achieve this speedup [104, 105].

So far we have discussed cases where the annealing spectrum already has properties that are in some way exploitable by DQA. However, there are also examples in the literature where a diabatic anneal is facilitated by a change to the annealing path. As discussed in section 2.1.2, non-stoquastic catalysts have been linked to an increase in the number of ACs involving the ground state [89, 90] which naturally suggests a greater potential for diabatic evolution. Further, it was noted that the introduction of a non-stoquastic catalyst was sometimes able to drastically increase the final GS overlap while not actually enhancing the size of the gap minimum [87, 89, 90]. This enhancement to the success probability was explicitly linked to the occurrence of diabatic transitions in [90] through numerical simulation of the dynamics, with the increase in ACs effectively providing potential routes back to the ground state.

That a non-stoquastic catalyst generally increases the number of ACs between the ground and first excited state is in line with the findings that non-stoquastic Hamiltonians tend to have smaller gaps in their low energy spectra [91]. It is also in line with ideas introduced in [100] where it is demonstrated that there are more ways for an AC to form with the ground state if it has negative vector components – which it can only have if the Hamiltonian is non-stoquastic.

While the introduction of additional ACs in the spectra have a chance of offering a diabatic path to the ground state that can be followed efficiently, the complicated spectra that are produced in this way may not always aid the computation. Recently however, an example of a more targeted approach has been put forward [100]. Considering a MWIS problem where the graph structure results in a perturbative crossing between the GS and a highly degenerate 1st excited state (1ES), the author demonstrates how a catalyst can be constructed that replaces the single AC at the end of the anneal with a correlated double AC which can be exploited to reach the ground state diabatically in polynomial time. The catalyst that achieves this is the same catalyst that was found to result in gap enhancement when introduced with a different strength (discussed in section 2.2.2) for which a polynomial time method is proposed for acquiring the information required to construct it. Another potential strategy for controllably altering the annealing spectrum has been proposed in which a pre-anneal is introduced to create an additional small gap at the start of the annealing spectrum. This has the to potential to circumvent the existence of a small gap later in the spectrum [106].

It is also possible that, rather than altering the annealing spectrum, more successful anneals that utilise diabatic dynamics could be achieved through clever optimisation of the annealing schedule. The main evidence for this comes from a study on the quantum approximate optimisation algorithm (QAOA) – which can be thought of as a Trotterised QA algorithm where the lengths of time that the driver and the problem Hamiltonians are applied at each step are parameters to be optimised. As such, annealing schedules and QAOA parameters can be mapped on to each other. A study comparing the success of QAOA and QA for Max-Cut problems [102] determined that for QAOA instances which corresponded to annealing spectra with larger gaps, the TTS with the optimal QAOA parameters followed the Landau-Zener (LZ) prediction – suggesting that the optimal QAOA parameters corresponded to an adiabatic annealing trajectory. However, for instances corresponding to annealing spectra with smaller gaps, the QAOA TTS fell away from the LZ prediction and plateaued. This suggests that the optimal QAOA schedule corresponds to an annealing schedule which utilises some kind of diabatic dynamics.

Comparing the QAOA TTS for these instances to the TTS for a standard linear anneal with the same run-time for the individual runs, it was found that QAOA significantly outperformed QA. This suggests that the optimised QAOA parameters correspond to some cleverly optimised annealing schedule which makes more efficient use of diabatic dynamics to end up with a higher ground state overlap – rather than just running the anneal more quickly. Mapping the QAOA parameters to a QA schedule for a representative instance and looking at the dynamics, it was confirmed that this was indeed the case. Further evidence that the optimal annealing schedule may be one that cleverly makes use of diabaticity came a few years later in a study that looked at optimising annealing schedules directly [107]. Here, optimising the annealing schedule was found to reduce the exponential run-time needed to find the GS of the frustrated ring model with QA to a polynomial run-time. Examining the dynamics associated with the optimised schedules, it was found that this improvement was being facilitated by a complex diabatic evolution.

There is of course the question of what we an expect of the time scaling associated with optimising the annealing schedule. In [107], the time requirements for the systems examined, which had up to 39 spins, suggested that the scaling was sub-exponential. It is worth noting however that, while hard to solve with QA, the frustrated ring model is easy to solve classically. It is possible that the non-exponential time scaling to find an optimal path is related to this.

Aside from the examples discussed here in which diabatic dynamics are demonstrated to be able to help QA reach the ground state faster, there are additional reasons why diabatic QA may be considered superior to its adiabatic counterpart. Firstly, DQA allows for universal computation which adiabatic QA does not without the introduction of non-stoquasticity [22, 108]. Currently, the overheads involved in any of the proposed methods mean that the use of QA for universal computation is intractable – however the fact that DQA allows for universality suggests that it is inherently more powerful that adiabatic QA.

However, DQA comes with its own problems. It has already been noted that the run-time may need to be chosen very precisely in order for DQA to outperform adiabatic QA [104, 105]. Further to this, an annealing protocol that attempts to utilise excited states as a path to the ground state is likely to be more adversely affected by certain open system effects such as relaxation.

2.4 Discussion

In this chapter we have discussed some of the variants on quantum annealing that exist in the literature. These variants included the use of inhomogeneous drivers, the introduction of catalyst Hamiltonians and the exploitation of diabatic transitions in order to circumvent the problem of exponentially closing gaps in the annealing spectrum. Let us highlight some of the key takeaways.

With regards to driver and catalyst selection, inhomogeneous drivers and catalyst hamiltonians with randomly selected components seemed to have some capacity to help QA achieve a high GS fidelity for shorter annealing times [66, 86, 87, 88**?**]. Considering the use of catalysts specifically, numerical studies found that *stoquastic* catalysts resulted in gap enhancement more frequently than *non*-stoquastic ones [87, 88, 89, 90] – a finding that was corroborated by theoretical work [91]. Meanwhile, non-stoquastic catalyst choices were found to have the capacity to facilitate diabatic annealing through the introduction of additional gap minima [87, 89, 90, 100]. The majority of the literature surrounding catalyst Hamiltonians has focused on the use of XX-couplings. However, the introduction of higher order X-terms were also found to be beneficial [81, 92**?**] with respect to enhancing the size of the spectral gap minimum. That higher order couplings tend to perform well can be understood as the system more easily being able to traverse between local and global optima.

Outside of introducing new couplings and/or altering the local-field strengths, another promising approach was to stagger the driver schedule such that different parts of the system would be pushed through a phase transition at different times [3, 82, 83, 85]. It is, however, hard to say exactly how any of these approaches will perform in general problem settings. Where concrete results demonstrating a quantifiable speed-up over the standard formulation of QA do exist [82, 83, 92, 93], they relate to analytically tractable models which do not reflect the typical structure one would expect to see in a Hamiltonian that encodes a difficult to solve optimisation problem. Indeed, by virtue of *being* analytically tractable, the models in question are only difficult for QA and do not actually correspond to hard problems. This does not necessarily invalidate the conclusions from these works – in which polynomial annealing times were facilitated by the removal of a first order phase transition. A strategy that is able to remove a first order phase transition in some simple model may well have a similar effect on a first-order phase transition present in a more complex setting.

That being said, if we assume that QA will *not* be able to solve **NP**-hard problems efficiently, that necessarily implies that there will be no "one size fits all" approach that facilitates polynomial time annealing runs that can reach the final GS with high fidelity. (Although, in theory, there could exist generic approaches that offer sub-exponential improvements to the annealing time.) And indeed this is reflected in the literature as a whole – motivating approaches which attempt to facilitate polynomial annealing run-times even if the overall scaling of the algorithm remains exponential in the problem size. For instance, while randomly generated drivers and catalysts may offer faster routes to the GS, the results in [66] suggest that an exponential number would likely need to be tried before finding a suitable choice. There is also the fact that different approaches perform differently in different settings. While stoquastic catalysts tend to result in more gap enhancement than non-stoquastic ones, there are cases where *non*-stoquastic catalysts are able to remove first-order phase transitions [67, 92, 93] where *stoquastic* ones cannot. Other studies have highlighted the importance of where the catalyst couplings are placed within the problem graph structure [81, 96, 100] as well as the sensitivity of their effects to the strength with which the catalyst is introduced [90, 100]. This naturally leads to the idea that it may be beneficial to actively incorporate problem-specific information into the way in which the algorithm is run.

Of course not all the information we might wish to utilise is available to us. For instance, optimising the catalyst parameters based on information on the instantaneous gap spectrum [88], or biasing towards the problem GS [68] are not likely to be viable strategies in practice – though there may be ways to utilise the ideas introduced in these works in other approaches. Information to which we *would* have access from the specification of the problem alone includes surface level information on the graph structure – *e.g.* what couplings are present. Ideas have been put forward that incorporate this kind of information into the selection of local-driver field-strengths [98], the design of inhomogeneous driving schedules [85, 97] and the construction of n-local XX-catalysts [81]. There have also been proposals to use knowledge of the local optima of a problem, which could be obtained through classical methods or "failed" annealing runs, in driver [64, 99] or catalyst [65, 67, 100] design.

Overall, a number of potential avenues for enhancing QA have been proposed – with different approaches likely to have different levels of viability on different hardware setups. In this dissertation, we aim to add to the arsenal of these. In particular, we will be focusing our attention on the inclusion of XX-couplings in the form of a catalyst Hamiltonian for the removal of perturbative crossings. Specifically, we will be examining the relationship between the structure of the problem spectrum and which XX-terms are most effective, thereby motivating a from of targeted catalyst Hamiltonian.

Chapter 3

Problem Instances

We are interested in examining the effects of different XX-catalysts on perturbative crossings. Specifically, we will be considering perturbative crossings that form as a result of non-degenerate excited states with lower energy neighbourhoods than the ground state. Not only do we wish to investigate how the couplings present in the catalyst impact the annealing spectrum but also how the effect of the catalyst depends on the properties of the perturbative crossings. To facilitate this, 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 adjust the neighbourhoods around the local optima.

We begin this chapter by going through the theory behind the formation of perturbative crossings. We then give a description of the MWIS problem and its encoding into an Ising Hamiltonian in section 3.2. Section 3.3 then describes the problem instances used in this work, how their structure relates to the theory of perturbative crossings and how this enables us to produce annealing spectra with the desired properties. Section 3.4 covers how we select and normalise the parameters that go into the MWIS Ising Hamiltonian. Finally, in section 3.5, we present numerical results showing that we can indeed controllably manipulate the annealing spectrum in the ways described.

3.1 Perturbative crossings

As described in section 1.3.1, we can understand perturbative crossings by introducing $\hat{H}_d$ as a perturbation to the problem Hamiltonian such that the total Hamiltonian is written

$$\hat{H}(\lambda) = \hat{H}_p + \lambda \hat{H}_d. \tag{3.1}$$

The perturbed energies up to second order corrections are then given by

$$\begin{aligned} E_a(\lambda) &= E_a + \lambda \langle E_a|\hat{H}_d|E_a\rangle + \lambda^2 \sum_{c\neq a} \frac{|\langle E_c|\hat{H}_d|E_a\rangle|^2}{E_a - E_c} \\ &= E_a + \lambda^2 \sum_{c\in N_x(a)} \frac{1}{E_a - E_c}, \end{aligned} \tag{3.2}$$

where we have used $N_x(a)$ to denote the neighbourhood of a – *i.e.*, the set of states to which a is coupled by the driver. To get from the first to the second line of equation 3.2, we have used the fact that $|\langle E_b|\hat{H}_d|E_a\rangle| = 1$ for $b \in N_x(a)$ and 0 otherwise. The first order term vanishes since $\hat{H}_d$ cannot couple a problem state to itself. Note that we can transform between the perturbed Hamiltonian in equation 3.1 and the annealing Hamiltonian in equation 1.1 with $s = 1/(1+\lambda)$.

We can say that a crossing occurs between two problem states a and b ($b > a$) if there exists some λ_x such that $E_b(\lambda) < E_a(\lambda)$ for $\lambda \geq \lambda_\mathrm{x}$. If this λ_x is within the convergence radius of the expansion, this indicates the formation of an AC in the annealing spectrum at the corresponding $s_\mathrm{x} = 1/(1+\lambda_\mathrm{x})$. In general, the crossings of interest are those between the lowest energy instantaneous states and, in particular, those involving the instantaneous ground state. We therefore restrict ourselves to considering the problem ground state and problem states with $E_0 < E_a \leq E_0 + \delta$ where δ is small compared to the energy gap between the problem ground state and the rest of the spectrum. We refer to this set of problem states as L_δ. If all states in L_δ are more than one spin flip apart from each other, as will be the case if these states correspond to local optima, then each term in the sum from equation 3.2 will be negative for all states $a \in L_\delta$. For clarity we can write, for the states $a \in L_\delta$,

$$E_a(\lambda) = E_a - \lambda^2 \sum_{c \in N_x(a)} \frac{1}{|E_c - E_a|}. \tag{3.3}$$

Taking of the absolute value in the denominator is redundant since $E_c - E_a$ will always be positive in this context however we include it to make the sign of these terms explicit. Since the contributions to the sum are all cumulative, it is clear that a state $a \in L_\delta$ will receive a greater negative perturbation if it has more states in its neighbourhood that are close in energy to it. And since L_δ contains the very lowest energy states, the neighbours that provide the greatest contributions to the sum in equation 3.3 will be the lowest energy neighbours.

Thus, for a crossing to occur between two states a and $b \in L_\delta$ ($b > a$), $N_x(b)$ must contain more low energy states than $N_x(a)$. Explicitly, there must be a positive λ_x for which

$$\begin{aligned}
E_a(\lambda_\mathrm{x}) &= E_b(\lambda_\mathrm{x}), \\
E_a - \lambda_\mathrm{x}^2 \sum_{c \in N_x(a)} \frac{1}{|E_c - E_a|} &= E_b - \lambda_\mathrm{x}^2 \sum_{c \in N_x(b)} \frac{1}{|E_c - E_b|}, \\
\Delta E_{ab} &= \lambda_\mathrm{x}^2 \left(\sum_{c \in N_x(b)} \frac{1}{|E_c - E_b|} - \sum_{c \in N_x(a)} \frac{1}{|E_c - E_a|} \right), \qquad (3.4)
\end{aligned}$$

where we have used ΔE_{ab} to denote the energy difference $E_b - E_a$. In order for the crossing to occur at small λ_x, such that perturbation theory remains valid, the difference in the energies of the neighbourhoods around a and b must be large enough compared to the unperturbed energy gap $\Delta E_{ab} = \mathcal{O}(\delta)$.

Before moving on, we note briefly that if we were to instead use an inhomogeneous driver of the from

$$\hat{H}_d = -\sum_{i=1}^{n} \gamma_i \hat{\sigma}_i^x, \tag{3.5}$$

then the numerators in equation 3.3 will be replaced by the $|\gamma_i|^2$ for which $\langle E_c | \hat{\sigma}_i^x | E_a \rangle = 1$. As such, the local field strengths γ_i can be used to tune the relative contributions of the neighbours as discussed in section 2.2.1. We do not utilise inhomogeneous driving in this work however the ideas behind the catalysts that we use are not unrelated.

Now that we have established the conditions under which a perturbative crossing occurs, we can move on to considering the size of the gap minimum that occurs as a result. We noted in section 1.3.1 that the degeneracy at the crossing point is removed by tunnelling between the two perturbed states which will depend on their overlap. We can write the perturbed states as

$$|E_a(\lambda)\rangle = \sum_{c=0}^{2^n-1} c_{ac}(\lambda) |E_c\rangle, \tag{3.6}$$

where the $c_{ac}(\lambda)$ values are the components of the perturbed vector $|E_a(\lambda)\rangle$ in the computational basis – or, equivalently, the overlaps of $|E_a(\lambda)\rangle$ with the problem eigenstates. Thus, for the gap size at the perturbative crossing, we can write

$$\Delta E_{ab}(\lambda_{\text{x}}) \propto \langle E_a(\lambda) | E_b(\lambda) \rangle = \sum_{c=0}^{2^n-1} c_{ac}(\lambda_{\text{x}}) c_{bc}(\lambda_{\text{x}}). \tag{3.7}$$

Perturbation theory tells us that, for any perturbed problem state a, the magnitude of $c_{ac}(\lambda_{\text{x}})$ will be exponentially small in the Hamming distance between $|E_a\rangle$ and $|E_c\rangle$ – since $|E_c\rangle$ will only enter the perturbative expansion on the order of this Hamming distance. Because each term in equation 3.7 will depend on the Hamming distance between each state, a and b, and the state c, we can expect $\Delta E_{ab}(\lambda_{\text{x}})$ to decrease exponentially with the Hamming distance with between $|E_a\rangle$ and $|E_b\rangle$. This Hamming distance will, to some extent, be instance specific – in that it depends on how well separated the problem's local optima are. However, we can generally expect this Hamming distance to grow linearly with the number of spins in the system resulting in a gap minimum that closes exponentially with the problem size.

3.2 The MWIS problem

Now that we have outlined the theory behind the annealing bottleneck we wish to study, let us move on to describing our problem setting. The MWIS problem 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. 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. More specifically, the problem Hamiltonian is given by

$$\hat{H}_p = \sum_{i \in \{\text{vertices}\}} (d_i J_{zz} - 2w_i)\hat{\sigma}_i^z + \sum_{(i,j) \in \{\text{edges}\}} J_{zz}\hat{\sigma}_i^z\hat{\sigma}_j^z, \tag{3.8}$$

where d_i is the degree and w_i the weight of vertex i. The appearance of the edge penalty strength, J_{zz}, in the local field terms is to account for the fact that two down spins coupled by an edge should not be penalised. *i.e.*, it is acceptable for two adjacent spins to not be included in the set.

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

We now describe how we construct our family of MWIS instances.

3.3 Problem graph structure

In order to produce annealing spectra with a desired perturbative crossing structure we need to be able to (a) select the number of excited states that have an energy close to that of the ground state, (b) tune how close in energy to the ground state these excited states are, and (c) adjust the energy spectra of the neighbourhoods around these states. The problems we construct, as well as the ways we utilise the free parameters in the corresponding Ising Hamiltonian encodings, are inspired by the MWIS instances used in [65].

As depicted in figure 3.1(a), the graphs used in this work are complete k-partite graphs with n_a vertices in each sub-graph, G_a. As such, the graphs have k maximally independent sets, giving the MWIS problem on these graphs k local optima – with each local optimum corresponding to the selection of all the vertices within one sub-graph. We allocate a total weight, W_a, to each sub-graph which we split evenly between the vertices such that each vertex in sub-graph G_a has a weight W_a/n_a. Given this problem structure, the problem Hamiltonian given in equa-

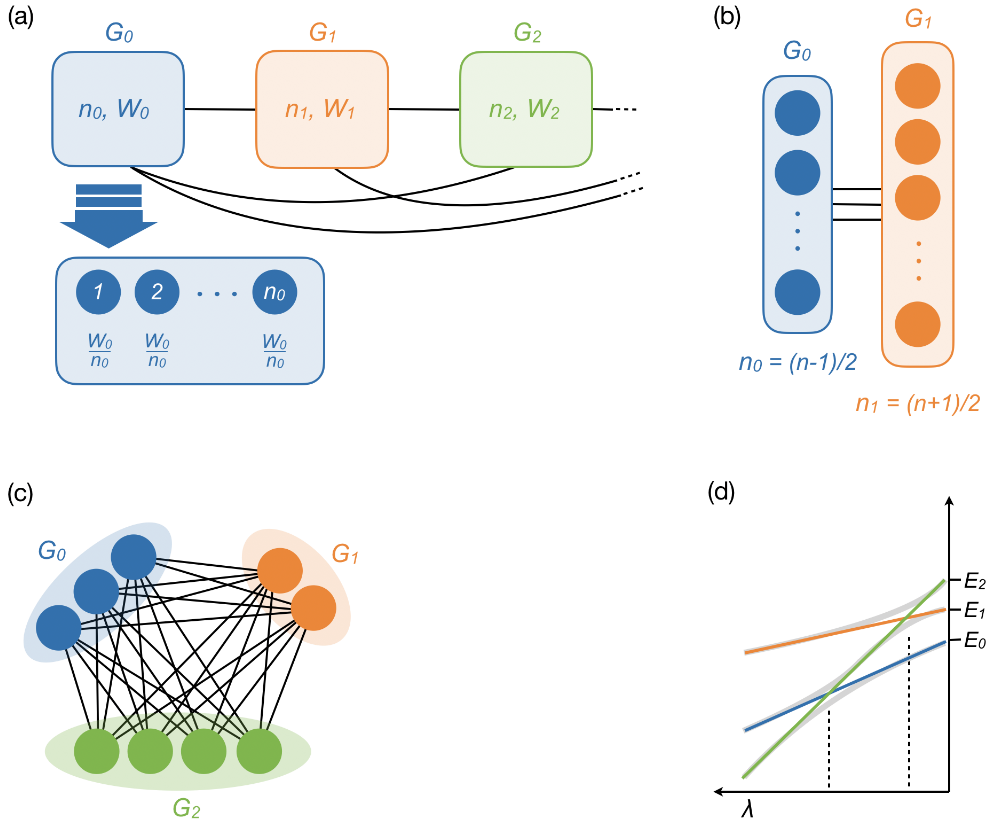

Figure 3.1: An illustration of our graph structure, as described in section 3.3, is shown in (a). (b) shows the bipartite graph scaling used to obtain the results in figure 3.4(d). An example graph with $n_0 = 3$, $n_1 = 2$ and $n_2 = 4$ is shown in (c) and a cartoon showing the perturbations to the three lowest energy states of the corresponding problem Hamiltonian is shown in (d). Here the coloured lines represent the perturbed problem energies up to second order and the grey curves illustrate the resultant instantaneous energies.

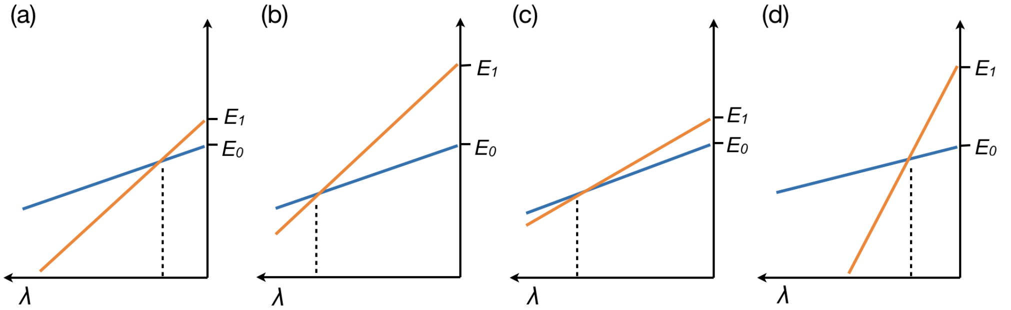

Figure 3.2: Cartoons illustrating the perturbed ground and first excited state energies of a problem Hamiltonian corresponding to a bipartite graph where $n_1 > n_0$. Taking the setting depicted in (a) as a base, (b) demonstrates the effect of increasing the difference between W_0 and W_1 and (c) demonstrates the effect of decreasing J_{zz} – both of which increase the value of λ for which the crossing occurs. (d) then demonstrates the effect of increasing both the weight difference and J_{zz}, keeping the location of the crossing the same.

tion 3.8 can be written as follows:

$$\hat{H}_p = \sum_{a=0}^{k-1} \left[\left(\underbrace{J_{zz} \sum_{b \neq a} n_b}_{d_i \forall i \in G_a} \underbrace{-2 \frac{W_a}{n_a}}_{w_i \forall i \in G_a} \right) \sum_{i \in G_a} \hat{\sigma}_i^z + J_{zz} \sum_{i \in G_a} \sum_{b > a} \sum_{j \in G_b} \hat{\sigma}_i^z \hat{\sigma}_j^z \right]. \tag{3.9}$$

The outer sum goes over all the sub-graphs in the problem. The first term introduces the local fields on the qubits associated with the vertices in a given sub-graph. For clarity, we highlight the parts corresponding to the degree, d_i, and weight, w_i, of the vertices. The second term in equation 3.9 introduces the necessary ZZ-couplings – each with a strength of J_{zz}. $2k + 1$ parameters are required to define a problem instance and its encoding into an Ising Hamiltonian. These are the k weights, W_a, the k sub-graph sizes, n_a, and the edge penalty, J_{zz}.

The sub-graph weights, W_a, will decide which of the k local optima is the global optimum and the ordering of the other local optima. If the weights are selected to be sufficiently close together, the k local optima will be the highest weighted independent sets and will thus correspond to the k lowest energy states in the associated problem spectrum[1]. How close the sub-graph weights are to each other will dictate how close in energy these states are. We have thus established how this graph structure can be used to select both the number of states with energy comparable to the ground state and also how close in energy these states are, as desired.

Next, we describe how we can control the spectra of the neighbourhoods around the states corresponding to the local optima. This is achieved through the allocation of sub-graph sizes. 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 a state $|E_a\rangle$, associated with the local optimum corresponding to sub-graph G_a, will have n_a neighbours corresponding to independent sets and $\sum_{b\neq a} n_b$ neighbours corresponding to dependent sets. By our selection of J_{zz}, we can ensure that all the states corresponding to independent sets have lower energies than those corresponding to dependent sets. As such, we can say that states corresponding to local optima associated with larger sub-graphs will have lower energy neighbourhoods. This means that the parameters n_a can be used to set the strengths of the perturbations. Naturally, increasing the difference between the sub-graph sizes will enhance the difference between the energies of the neighbourhoods and thus the difference in the level of the perturbation from the driver. However, we can also enhance this difference by increasing J_{zz} since this will increase the energy difference between the states corresponding to independent and dependent sets. Note that increasing J_{zz} does not alter the problem that is encoded.

To summarise the above, we can use the number of sub-graphs, k, to

---

[1] For a sufficiently large weight differences, sub-sets of one of the higher weighted sub-graphs may have a larger weight than one of the other maximally independent sets. When this is the case, the local optima will not correspond to the very lowest excited states. This does not affect the way in which the desired perturbative crossings form but will of course impact the labeling of the states.

set the number of problem states with energy comparable to that of the ground state and we can use the sub-graph weights, W_a, to decide how close in energy these states are. We can then decide on the level of perturbation that each problem state will receive from the driver with the sub-graph sizes, n_a, as well as enhance or reduce the difference between these perturbations by tuning J_{zz}. For ease of notation we will always set the weights such that $W_0 > W_1 > \ldots$. This is so that, for $a < k$, we can equate the problem state $|E_a\rangle$ with the local optimum associated with selecting all the vertices in sub-graph G_a.

To better understand how these parameters will result in a particular perturbative crossing structure we consider the tripartite example graph in figure 3.1(c) with $n_0 = 3$, $n_1 = 2$ and $n_2 = 4$. These parameter choices mean that out of the three problem states corresponding to the three local optima, the 2ES will have the lowest energy neighbourhood, followed by the GS and then the 1ES. We therefore expect the 2ES to receive the greatest negative perturbation from the driver and the 1ES to receive the smallest. A cartoon of the resultant crossing structure is shown in figure 3.1(d). Effectively, by choosing the sub-graph sizes we can decide how the energies of the problem states re-order themselves.

It is straightforward to understand how changing the sub-graph weights and J_{zz} will effect the nature of the perturbative crossings using the ideas introduced in section 3.1. Consider the cartoon in figure 3.2(a) which depicts the perturbed problem ground and first excited state energies in a setting where we have chosen $n_1 > n_0$ such that a perturbative crossing forms. By increasing the weight difference, $\delta W = W_0 - W_1$, we can increase the energy difference, $\Delta E_{01} = E_1 - E_0$, between the unperturbed problem states. As shown in figure 3.2(b), this will result in the crossing happening for a larger value of λ and thus a smaller value of s. Lowering the strength of the edge penalty J_{zz} will have a similar effect since the difference between the perturbation strengths will be reduced – this is illustrated in figure 3.2(c). Finally, figure 3.2(d) depicts a setting where we have increased both δW and J_{zz} such that the location of the crossing remains the same but the 1ES has received a significantly greater perturbation at the crossing point compared to the setting depicted in (a). Recall from the discussion in section 3.1 that we expect a larger perturbation of the states at the point of crossing to result in a greater overlap between the two states at this point and so a larger gap minimum. We therefore expect setting (d) to have a larger gap minimum than setting (a) in its associated annealing spectrum but for the location of this gap minimum to be the same.

In section 3.5 we will present numerical results confirming that we are able to controllably produce annealing spectra as outlined in this section. First, however, we give more detail on the parameters that enter $\hat{H}_p$.

3.4 Problem parameters

The preceding section described how the parameters, $\{n_a\}$, $\{W_a\}$ and J_{zz}, will be used to to alter the gap spectrum of $\hat{H}_p$ in order to tune the properties of the perturbative crossings that form in the annealing spectrum. In this section we discuss how these values are actually chosen and how we normalise the parameters that enter equation 3.9.

First, it is important to differentiate between the MWIS problem itself and the Hamiltonian that is constructed to encode it. Note that the problem is defined solely by the parameters $\{n_a\}$ and $\{W_a\}$ while J_{zz} is a parameter associated exclusively with the problem's encoding into an Ising Hamiltonian. With regards to the MWIS problem, one need not worry about selecting an appropriate scale with regards to the magnitudes of $\{W_a\}$. However, once these values are substituted into 3.9, we must consider what this means for the energy scale of $\hat{H}_p$ in relation to $\hat{H}_d$ and, if a catalyst is present, $\hat{H}_c$. As such, a consistent means of re-scaling the parameters is required to ensure that $\hat{H}_p$ has an appropriate energy scale. (The parameters, $\{W_a\}$, will also affect the minimum value of J_{zz} that ensures that all the states corresponding to dependent sets are higher in energy than those corresponding to independent sets.

Selecting J_{zz} values will be discussed towards the end of this section.)

We want to be able to adjust the values of $\{W_a\}$ and J_{zz} to alter the *relative* gaps between states in the problem spectrum without changing the energy scale of $\hat{H}_p$. With all other factors being equal, increasing this energy scale will result in gap enhancement and so we wish to keep it constant for a given system size, $n = \sum_{a=0}^{k-1} n_a$. Furthermore, we wish to maintain a consistent relationship between the energy scales of $\hat{H}_p$ and $\hat{H}_d$ with increasing n.

With the preceding discussion in mind, we introduce the normalisation factor

$$F = \frac{E_{2^n-1}(0) - E_0(0)}{E_{2^n-1} - E_0}, \tag{3.10}$$

where the numerator and denominator are the spectral range of $\hat{H}_d$ and $\hat{H}_p$ respectively. The numerator is trivially

$$E_{2^n-1}(0) - E_0(0) = 2 \times n = 2 \sum_{a=0}^{k-1} n_a. \tag{3.11}$$

We can use equation 3.9 to obtain expressions for E_0 and E_{2^n-1} in terms of the parameters $\{n_a\}$, $\{W_a\}$ and J_{zz}. Assuming that a sufficiently high J_{zz} has been chosen (and recalling that we allocate the largest weight to G_0) we know that the ground state will have all the spins in G_0 pointed up and the rest pointed down. Similarly, we know that the highest excited state, E_{2^n-1}, will be the state with all spins pointed up since this corresponds to the set with the maximum number of edge violations. From this we obtain the following expressions:

$$E_0 = -2W_0 + \sum_{a=1}^{k-1} \left[2W_a - J_{zz} n_a \sum_{b=0}^{a-1} n_b \right] , \tag{3.12}$$

$$E_{2^n-1} = \sum_{a=0}^{k-1} \left[J_{zz} n_a \left(\sum_{b=0}^{a-1} n_b + 2 \sum_{b=a+1}^{k-1} n_b \right) - 2W_a \right] . \tag{3.13}$$

Substituting equations 3.11, 3.12 and 3.13 into equation 3.10 we get:

$$F = \frac{\sum_{a} n_a}{\sum_{a=1}^{k-1}\left[n_0 n_a J_{zz} + n_a J_{zz} \sum_{b \neq a} n_b - 2W_a\right]} \, . \tag{3.14}$$

For an un-normalised set of problem weights, $\{W_a\}$, and an edge penalty, J_{zz}, we can now obtain the normalised parameters:

$$W_a' = E_{\text{scale}} \times F \times W_a, \; J_{zz}' = E_{\text{scale}} \times F \times J_{zz} \, . \tag{3.15}$$

Here, E_{scale} sets the energy scale of $\hat{H}_p$ in relation to $\hat{H}_d$. $E_{\text{scale}} = 7.5$ is used throughout this work. This value has been chosen because it allows key features of the annealing spectra to be easily distinguished for the smaller system sizes – such that we can use the full instantaneous spectra of these instances as illustrative examples. Selecting a much smaller E_{scale} results in the perturbative crossings being squashed at the very end of the anneal. Significantly larger values of E_{scale} lead to less pronounced gap minima.

In selecting the un-normalised weights, the lowest weighted sub-graph is always given a weight of unity (with the exception of one instance in chapter 9). The other sub-graphs are then given weights $1 < W_a \leq 1.7$ depending on how close in energy we want to make the states corresponding to the local optima. J_{zz} is then chosen in response to these values in order to produce an annealing spectrum with the desired properties. Note that the values of $\{W_a\}$ and J_{zz} quoted throughout this text are the un-normalised parameters. The corresponding normalised parameters can be found in appendix B.

Throughout this section we have referred to choosing a *sufficiently high* J_{zz} to ensure that the states corresponding to independent sets all have lower energies than those corresponding to dependent sets. It is clear that the highest energy state corresponding to an *independent* set will be the state with all spins pointed down – representing the empty set. While it is easy enough to obtain an expression for this energy from equation 3.9, determining an expression for the lowest energy state corresponding to a *dependent* set is not so straightforward. This is because the set of vertices that this state corresponds to is instance dependent. We know that the state in question should contain only one edge violation and so should consist of one vertex from one sub-graph and one vertex from another. However, the selection of sub-graphs that results in the largest weighted set of two vertices will depend on both the total sub-graph weights and their sizes since it is the weights on the individual vertices that is relevant. In practice, therefore, we simply check that the problem spectrum contains no states corresponding to dependent sets that are lower in energy than the all spin down state.

Finally, the values of n_a that have been used vary between 2 and 37 and the total system sizes we examine range from $n = 5$ to $n = 73$. These larger system sizes are tractable as a result of symmetries in the problem. How these symmetries were used to reduce the size of the Hilbert space is discussed in section 4.2.

3.5 Spectral Properties

We now present numerical results corresponding to different parameter settings to confirm that we are able to produce the desired perturbative crossings and that we are able to alter their properties in the way described. We begin by demonstrating that we can controllably produce different AC structures between the k lowest instantaneous states using our k-partite graphs. We then move to a bipartite graph for simplicity to show how tuning the edge penalty and weight difference, $\delta W = W_0 - W_1$, alters the properties of the perturbative crossing that forms.

In figure 3.3 we present annealing spectra corresponding to the three example graphs shown in the top row. These results are obtained through

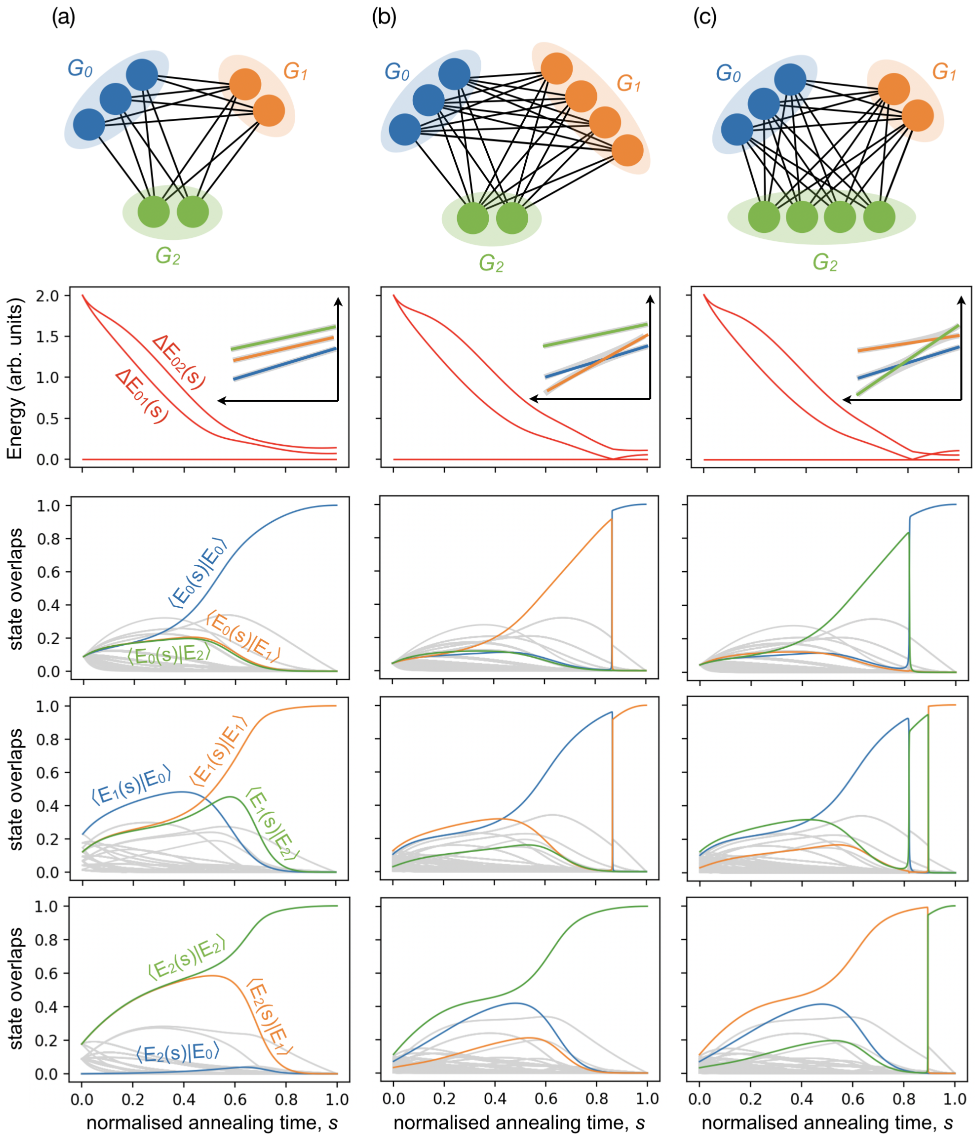

Figure 3.3: Numerical results for the annealing spectra associated with three different example graphs. Each column corresponds to a different graph which is depicted on the top row. All three instances share the same un-normalised parameters: $W_0 = 1.04$, $W_1 = 1.02$, $W_2 = 1.00$ and $J_{zz} = 2$. The second row shows the gap spectra as well as an inset cartoon illustrating the perturbations to the problem state energies. Rows three, four and five show the evolution of the ground, first and second excited states respectively in terms of their overlaps with the problem states – the problem ground first and second excited state overlaps are highlighted in blue, orange and green respectively.

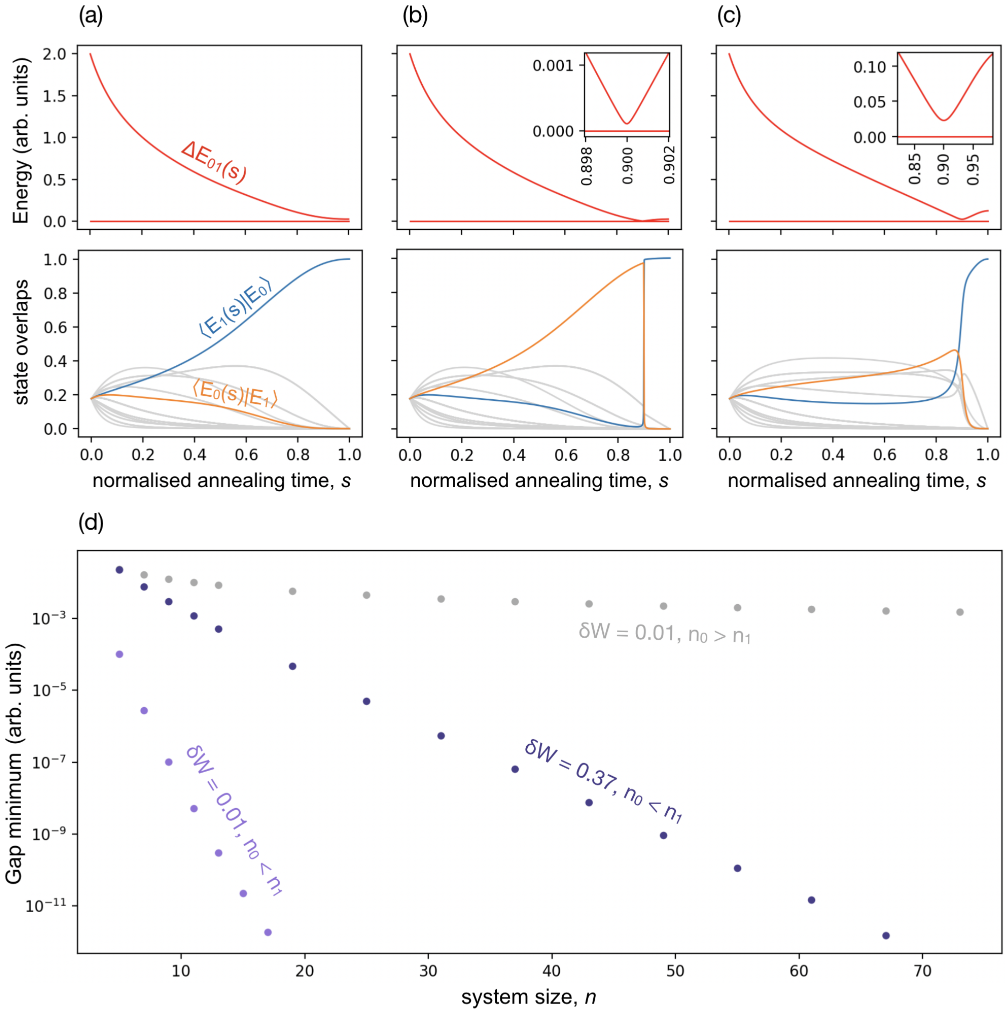

Figure 3.4: (a-c) show numerical results for annealing spectra corresponding to different bipartite MWIS problems. Gap spectra are shown on the top and the evolution of the instantaneous ground state vectors are shown on the bottom. The gap spectra insets show a close-up view of the gap minima. The problem parameters are: (a) $n_0 = 3$, $n_1 = 2$, $W_0 = 1.01$, $W_1 = 1.00$, $J_{zz} = 5.33$. (b) $n_0 = 2$, $n_1 = 3$, $W_0 = 1.01$, $W_1 = 1.00$, $J_{zz} = 5.33$ and (c) $n_0 = 2$, $n_1 = 3$, $W_0 = 1.37$, $W_1 = 1.00$, $J_{zz} = 37.5$. (d) shows the minimum gap with increasing system size for three different parameter settings. The grey line corresponds to the parameter settings in (a) and the two purple lines correspond to the parameter settings in (b) and (c). 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.

numerical diagonalisation of the total Hamiltonian at different values of s. Details on the tools we use to perform the diagonalisation will be given in the following chapter. That our numerical methods are not presented prior to this point is because they include a Hilbert space reduction that makes use of the graph structure introduced in this chapter.

In the second row of figure 3.3, we plot the energy gaps ΔE_{0a} as a function of s and also include an inset cartoon depicting the perturbations to the problem state energies. In the third, fourth and fifth rows of figure 3.3 we show the corresponding evolution of $|E_0(s)\rangle$, $|E_1(s)\rangle$ and $|E_2(s)\rangle$ respectively.

We have not attached a unit to the energies we present in figure 3.4 as it is not currently relevant to the discussion. This is because we are currently only interested in particular features of the energy landscape (such as the formation of gap minima) and the relative energy gaps associated with different parameter settings. The values can simply be understood as being in relation to the local driver fields being introduced with a magnitude of unity. Since the parameters of H_p are normalised with respect to the driver (as discussed in the previous section) and we are currently only working with the normalised annealing time ($s = t/t_a$), this is sufficient to define the setting. More formally, for a driver introduced with energy h_x, the plotted values are the resultant energy gaps divided by h_x. [2]

The sub-graph weights associated with each of the example graphs are $W_0 = 1.04$, $W_1 = 1.02$ and $W_2 = 1.00$ and the edge penalty chosen is $J_{zz} = 2$. Figure 3.3(a) depicts a setting where the ground state has the lowest energy neighbourhood by virtue of the corresponding local optimum containing the most vertices. The first and second excited states corresponding to the other two local optima both have neighbourhoods with comparable energy spectra. This is reflected in the corresponding

---

[2] For a realistic annealing setting, one could consider these energies to be in GHz – which is reflective of the energy scales used by the D-Wave quantum annealers [109]. In later sections of this dissertation, where it does become relevant, this is indeed the energy scale that we use.

cartoon where we show the problem ground state receiving the greatest negative perturbation from the driver and the other two states receiving the same perturbations as each other such that no perturbative crossings are expected to form. Looking at the associated gap spectrum we see that this is indeed the case in that no small gap minima form between the instantaneous energy levels. Looking at the instantaneous vectors, we see that the evolution proceeds monotonically towards the corresponding problem state – *i.e.*, $|E_a(s)\rangle$ evolves monotonically towards $|E_a\rangle$. Considering the large s values, which is where we would expect to see perturbative crossings were there any present, we see that each state $|E_a(s)\rangle$, remains dominated by the corresponding problem state, $|E_a\rangle$, rather than there being a sharp amplitude exchange into a different problem state which we would expect to see at a perturbative crossing.

Next, consider the problem graph shown in figure 3.3(b) where we have now given the problem 1ES a lower energy neighbourhood than the ground state. As depicted in the corresponding cartoon, we now expect a perturbative crossing to form between these two states resulting in an AC between the instantaneous ground and first excited states. Looking at the associated gap spectrum we do indeed observe the appearance of a small gap minimum forming between the two states. Considering the evolution of the instantaneous ground state, moving backwards from $s = 1$, $|E_0(s)\rangle$ can be seen to be dominated by $|E_0\rangle$ until there is a sharp exchange into $|E_1\rangle$ at the location of the gap minimum – reflecting the perturbed problem 1ES becoming lower in energy than the perturbed problem GS. The reverse behaviour is observed for the instantaneous 1ES. Conversely, the instantaneous 2ES evolves smoothly away from $|E_2\rangle$ as s decreases away from $s = 1$, reflecting the fact that it does not partake in a perturbative crossing.

Finally, consider the setting depicted in figure 3.3(c). This is the same example as that discussed in section 3.3 and we expect the problem energies to re-order themselves such that two perturbative crossings form as illustrated in the corresponding cartoon. Once again, this is reflected in the gap spectrum where we observe a gap minimum forming between the ground and first excited state at around $s = 0.81$, followed by a gap minimum between the first and second excited state at around $s = 0.88$. The expected abrupt changes in the instantaneous states are observed at the locations of the gap minima. In particular, note how the $|E_2\rangle$ amplitude can be traced through from the instantaneous ground state, through the first excited state and into the second excited state as the perturbed problem 2ES goes from having the lowest energy to the highest with increasing s (or decreasing λ).

Before moving on, we briefly consider the path a system might take through these annealing spectra. Let us assume that an annealing time is chosen that allows the evolution to proceed adiabatically at all points other than at the perturbative crossings at which point all its amplitude is transferred. Under these conditions, the system will end with a high overlap with the state that has the lowest energy neighbourhood – the ground, first and second excited states for the three examples respectively. Effectively quantum annealing with its standard driver and linear interpolation searches for the low energy state with the lowest energy neighbourhood – rather than the very lowest energy state.

We now demonstrate the effects of altering the sub-graph weights and edge penalty. For simplicity we move to a bipartite graph with $n_1 > n_0$ such that we produce one perturbative crossing whose properties we can tune. Numerical results for a graph with $n_0 = 2$ and $n_1 = 3$, and two different settings for the weights and edge penalty, are presented in figures 3.4(b) and (c). The gap spectra are presented in the top row and the evolution of the instantaneous ground state on the bottom.

The results in figure 3.4(b) correspond to a setting with $W_0 = 1.01$, $W_1 = 1.00$ and $J_{zz} = 5.33$. The result is that the perturbed problem states have undergone little change before the crossing point at $s = 0.9$. This is reflected in the fact the instantaneous ground state goes from $|E_0(s)\rangle \approx |E_1\rangle$ to $|E_0(s)\rangle \approx |E_0\rangle$ at the perturbative crossing. For comparison, figure 3.4(a) depicts a setting with the same parameters but with the sub-graph sizes swapped so that no perturbative crossing forms.

The parameters used for the anneal in figure 3.4(c) are $W_0 = 1.37$, $W_1 = 1.00$ and $J_{zz} = 37.5$ – that is, we have increased both the weight difference and the edge penalty. (The specificity of this edge penalty and that from the preceding parameter set, $J_{zz} = 5.33$, are so that both perturbative crossings occur at $s = 0.9$.) Looking at the evolution of the instantaneous GS in this setting, we see that $|E_0(s)\rangle$ is significantly more mixed just before the perturbative crossing reflecting the fact that the problem 1ES has undergone a greater perturbation at the point of the crossing. These two settings correspond to the cartoons in figures 3.2(a) and (d) respectively. As expected, we see that the gap minimum in the second setting is significantly larger than that in the first. We also note that the exchange that happens in the instantaneous ground state at the crossing point is more abrupt in the setting with the smaller gap. 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 [110].

Finally, we consider the scaling of the gap minimum between the ground and first excited state with increasing system size – these results are presented in figure 3.4(d). We plot the results corresponding to the parameter settings from figures 3.4(b) and (c) in light and dark purple respectively. The sub-graph sizes are scaled as shown in figure 3.1(b) such that G_1 always has one more vertex than G_0. We see that while both gap minima appear to close exponentially, the scaling exponent is more severe for the parameter setting associated with figure 3.4(b). We will be utilising these two parameter settings throughout this work and refer to them as the weak gap scaling (WGS) and strong gap scaling (SGS) settings. For comparison we also include the scaling of the gap minimum when there is no perturbative crossing. This is plotted in grey and we see that we no longer observe the same closing of the gap.

This scaling behaviour further confirms that an AC can be controllably introduced into the annealing spectrum. In addition, we have demonstrated that we are able to tune the problem parameters to quantitatively change the nature of the AC without changing the optimal solution that $\hat{H}_p$ encodes or the structure of the local optima.

3.6 Summary

In this chapter, we have introduced the MWIS problem instances that we will be using throughout this work. We began by describing the theory behind perturbative crossings and, in particular, what spectral features of the problem Hamiltonian would cause one to form. We then introduced our problem graphs and linked their structure and defining parameters to the these key spectral features. We outlined and demonstrated how these graphs could be used to controllably produce annealing spectra with the desired perturbative crossing structures.

These graphs will provide a numerical test-bed for different catalyst choices in chapters 5 and 6 which, alongside our theoretical analysis, will allow us to motivate the inclusion of particular XX-coupling choices. Furthermore, we will be utilising the the straightforward scalability of these graphs to study the scaling behaviour of the proposed catalysts with both the size of the problem graph and also the number of local optima in chapters 7 and 9 respectively.

Chapter 4

Numerical Methods

This chapter outlines the tools we have used to obtain our numerical results. These include standard pre-written packages as well as tools that we have implemented. Section 4.1 concerns the diagonalisation of our Hamiltonains. Next, in section 4.2, we describe how we make use of symmetries in the system in order to reduce the size of the Hilbert space. Finally, in section 4.3, we discuss the closed system dynamics.

All simulations were performed with `python` on a 2.4 GHz Quad-Core Intel processor with 16GB of RAM. The two main packages we have made use of are `numpy` and `qutip`.

4.1 Diagonalisation

The main framework by which we consider the effects of different catalyst Hamiltonians is through their impact on the instantaneous annealing gap spectra and the evolution of the instantaneous states – as obtained by numerical diagonalisation of the total Hamiltonian at different time-steps. We use the `numpy` function `numpy.linalg.eigh` [111] which takes, as its input, a complex hermitian matrix or a real symmetric matrix. The function makes use of `LAPAK` routines [112] – specifically, `syevd` for real symmetric matrices and `heevd` for complex hermitian matrices. These routines compute all the eigenvalues and, optionally, the eigenvectors of the matrix using divide and conquer algorithms.

4.1.1 Global Sign Selection

One additional complexity is the selection of the global signs associated with the instantaneous eigenvectors at each time step. In terms of the physics, the global phase has no physical significance. And indeed, as far as the diagonalisation is concerned, the global sign of the vector does not have a definite value. However, as we shall now go on to explain, interpreting our data will require us to establish some kind of consistent global sign.

Figure 4.1(a) shows the the instantaneous GS vectors returned by the diagonaliser for an example 3-spin system with the different vector components plotted in different colours. The results at each time step that we have sampled are technically completely correct. However, in order for us to understand the *evolution* of this state, some post-processing is still required. If we simply connect up the raw data, as shown in figure 4.1(b), we end up with a picture that suggests rapid changes in the GS vector at various points during the anneal. For the data presented, it is relatively easy to deduce intuitively that the rapid changes indicated by following the evolution in this way are not actually happening. However, there are also some more rigorous ways to understand that the deduced evolution presented in figure 4.1(b) is incorrect. Firstly, the instantaneous vectors must be normalised to one, meaning that we cannot have as many vector components crossing zero as figure 4.1(b) would suggest. We can also check the interpolated evolution by sampling points in between two time steps and seeing if they lie along, or close to, the previously interpolated trajectory.

In most cases, it is straightforward but time consuming to “fix” these global sign flips manually. Given the prevalence of this form of data in our research, we chose to implement tools to automate this process. Again, as far as the sampled data points are concerned, there is no “incorrect” choice of global phase that needs to be “fixed” – it is only the suggested evolution from connecting these points that is incorrect. Specifically then, what we want is a means of selecting the global sign

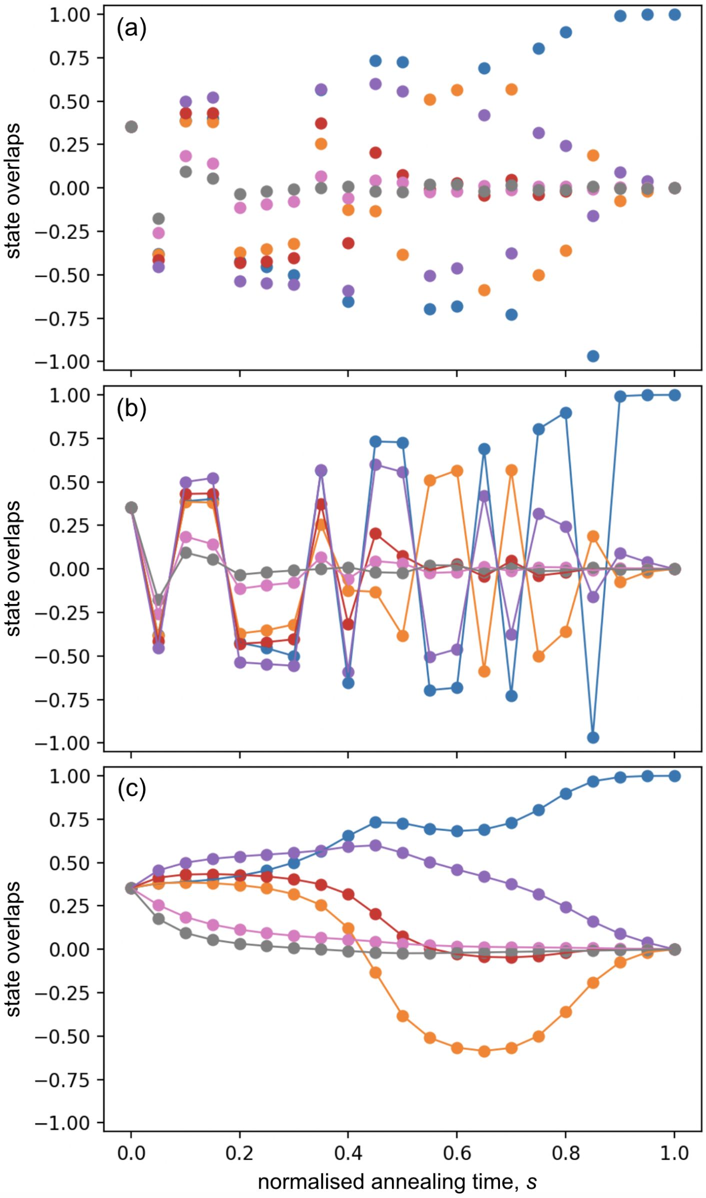

Figure 4.1: Plots showing the diagonalisation results for the instantaneous GS of a 3-spin example instance at 21 points in the anneal. The different problem state overlaps, $\langle E_0(s)|E_i\rangle$, are plotted in different colours. The problem parameters associated with this example are $n_0 = 1$, $n_1 = 2$, $W_0 = 1.2$, $W_2 = 1.0$ and $J_{zz} = 20$. A catalyst Hamiltonian that consists of a single XX-coupling between the two vertices in G_1 is introduced with a strength of 2. Plot (a) shows the raw results as returned by `numpy.linalg.eigh`. Plot (b) shows the same results but where we have connected the data points. Finally, plot (c) shows the data once processed as discussed in section 4.1.1.

of each sampled vector such that the implied evolution between these points reflects the actual evolution that would be obtained if more points were sampled in between. We do this by combing through the sampled data points and deciding whether or not a non-physical global sign flip has occurred.

In order to automatically search for these global sign flips we implemented two heuristic methods. (Simply checking if all the vector components have changed sign between two points is unfortunately not sufficient since a vector component that has changed its *relative* sign between these points will have kept its sign if a global sign change has occurred.) The two methods work as follows:

1. For a given threshold, M_{thresh}, assume a global sign change has occurred if all vector components with a magnitude exceeding M_{thresh} change sign.
2. For a given threshold, N_{thresh}, assume a global sign change has occurred if more than N_{thresh} vector components change sign.

Overall we found that the first method worked better and that selecting $M_{\text{thresh}} = 0.05$ tended to perform well. In practice it is straightforward to confirm that the global signs have been chosen appropriately and in some cases we switched between methods, adjusting M_{thresh} and N_{thresh} as needed. Figure 4.1(c) shows the evolution of our example instance after this post-processing.

4.2 Hilbert space reductions

In order to make our simulations more tractable, we utilise symmetries present in the Hamiltonians to reduce the size of the Hilbert space – taking a similar approach to that used in [67]. We would like to thank Tameem Albash for taking the time to discuss this with us.

As described in chapter 3, the problem instances we will be working with are defined on complete k-partite graphs constructed out of k disconnected sub-graphs, G_a (where $a = 0, 1 \ldots k-1$). Each sub-graph, G_a, is allocated a total weight, W_a, which is split evenly between the n_a vertices in the sub-graph. As a result, $\hat{H}_p$ is invariant under permutations of qubits within the sub-graphs. Since the driver, $\hat{H}_d$, is simply a homogeneous local X-field across all the qubits, the total Hamiltonian, $\hat{H}(s) = (1-s)\hat{H}_d + s\hat{H}_p$, is also invariant under such transformations. Let us first describe how we utilise these symmetries in the catalyst free setting before going on to describe how we augment this approach to account for the introduction of the catalyst.

The crux of the approach is that the system is initialised in the equal superposition state such that the evolution starts in the subspace spanned by the states that are symmetric under the aforementioned permutations. In the closed system setting, the evolution is confined to this subspace which means that we can drastically reduce the size of the Hilbert space we need to consider. A natural basis for this subspace can be constructed from Dicke states associated with each of the sub-graphs G_a.

For each sub-graph, G_a, we can write the total Z-spin operator,

$$\hat{S}^z_a = \frac{1}{2} \sum_{i \in G_a} \hat{\sigma}^z_i, \tag{4.1}$$

of which the Dicke states, $|s_a, m_a\rangle$, are eigenstates – s_a being the total spin associated with the state. Their eigenvalues are m_a which run from $-s_a$ to s_a in integer steps. Equivalently, we can write out the corresponding X and Y operators as

$$\hat{S}^x_a = \frac{1}{2} \sum_{i \in G_a} \hat{\sigma}^x_i \tag{4.2}$$

and

$$\hat{S}^y_a = \frac{1}{2} \sum_{i \in G_a} \hat{\sigma}^y_i. \tag{4.3}$$

In the catalyst free case, and indeed for most of the catalysts we examine, we will only need $\hat{S}^z_a$ and $\hat{S}^x_a$ to construct our Hamiltonians. The only point at which we will need $\hat{S}_a^y$ is in chapter 6.

In order to construct $\hat{H}(s)$ out of these total spin operators, we explicitly compute the matrix elements

$$\langle s_a, m_a | \hat{S}_a^z | s_a, m_a' \rangle = \delta_{m_a, m_a'} \times m_a, \tag{4.4}$$

and

$$\begin{aligned} &\langle s_a, m_a | \hat{S}_a^x | s_a, m_a' \rangle = \\ &\qquad \delta_{m_a, m_a'-1} \times \frac{1}{2}\sqrt{s_a(s_a+1) - m_a(m_a+1)} \\ &\qquad\qquad + \delta_{m_a, m_a'+1} \times \frac{1}{2}\sqrt{s_a(s_a+1) - m_a(m_a-1)}. \end{aligned} \tag{4.5}$$

The matrix elements of S_a^y, when we need them, are calculated as

$$\begin{aligned} &\langle s_a, m_a | \hat{S}_a^y | s_a, m_a' \rangle = \\ &\qquad \delta_{m_a, m_a'-1} \times \frac{i}{2}\sqrt{s_a(s_a+1) - m_a(m_a+1)} \\ &\qquad\qquad - \delta_{m_a, m_a'+1} \times \frac{i}{2}\sqrt{s_a(s_a+1) - m_a(m_a-1)}. \end{aligned} \tag{4.6}$$

The driver Hamiltonian can then be written in much the same way as it was using the Pauli-X operators, as

$$\hat{H}_d = 2\sum_a \hat{S}_a^x. \tag{4.7}$$

Similarly, we can write the problem Hamiltonian as

$$\hat{H}_p = 2\sum_a h_a \hat{S}_a^z + 4J_{zz} \sum_{a,\, b>a} \hat{S}_a^z \hat{S}_b^z, \tag{4.8}$$

where h_a and J_{zz} are, respectively, the local field, and coupling strengths as calculated in section 3.4. We have dropped some identity terms that would make this problem Hamiltonian fully equivalent to the Hamiltonian in equation 3.8. However, the effect of this is only to introduce a shift in the absolute values of the energies. The spectral gaps are unchanged.

The eigenstates of equation 4.8, which span the symmetric subspace to which the evolution is restricted, are products of the Dicke states associated with the different permutation symmetries of the sub-graphs:

$$
\begin{aligned}
&|m_0, m_1, \ldots, m_{k-1}\rangle \\
&\qquad \equiv |\frac{n_0}{2}, m_0\rangle \otimes |\frac{n_1}{2}, m_1\rangle \otimes \cdots \otimes |\frac{n_{k-1}}{2}, m_{k-1}\rangle \\
&\qquad\qquad = \prod_{0 \leq a < k-1} |\frac{n_a}{2}, m_a\rangle \,. \qquad (4.9)
\end{aligned}
$$

The total number of states is

$$\prod_{0 \leq a < k-1} (n_a + 1), \qquad (4.10)$$

such that the Hilbert space now scales linearly with the increase of vertices in any one of the sub-graphs, G_a. In practice, we will usually be increasing the size of all sub-graphs when scaling our graph structures such that the size of the Hilbert space scales as a polynomial of order k. The Hilbert space scales exponentially with the *number* of sub-graphs, k.

The situation is slightly complicated by the introduction of different catalyst Hamiltonians, $\hat{H}_c$, which we will be introducing to our total Hamiltonian with the pre-factor $s(1-s)$. Depending on our choice of $\hat{H}_c$, this catalyst may disrupt the symmetry shared by $\hat{H}_d$ and $\hat{H}_p$. We will be exploring a variety of different catalysts over the course of this dissertation – and the exact nature of these different catalysts will be discussed within the chapters in which they are introduced. However, they all share the property that the resultant symmetries can be captured by splitting off some of the vertices from one or more sub-graphs. That is, we can rewrite the necessary sub-graphs as $G_a = G_{a'} \cup G_{a''}$ such that $n_a = n_{a'} + n_{a''}$.

The effect on the construction of $\hat{H}_d$ and $\hat{H}_p$ in equations 4.7 and 4.8 is that we must simply replace the relevant total spin operators with $\hat{S}_a^\alpha = \hat{S}_{a'}^\alpha + \hat{S}_{a''}^\alpha$ (where $\alpha = x$, y or z). However, we now have the capacity to construct various 2-local catalyst Hamiltonians. To give some examples, all-to-all couplings can be introduced within one of the sub-graphs or a part of one of the sub-graphs with $2(\hat{S}_a^\alpha \otimes \hat{S}_a^\alpha)$ or $2(\hat{S}_{a''}^\alpha \otimes \hat{S}_{a''}^\alpha)$ respectively. As a sub-case of this, a single coupling can be introduced if n'' is chosen to be 2. We can introduce all-to-all couplings between one part of a sub-graph another another with $4(\hat{S}_{a'}^\alpha \otimes \hat{S}_{a''}^\alpha)$ or all-to-all couplings between two different sub-graphs with $4(\hat{S}_a^\alpha \otimes \hat{S}_b^\alpha)$.

4.3 Closed system dynamics

Simulations of the closed system dynamics were done using the `qutip` function, `qutip.sesolve` [113] which is an ordinary differential equation (ODE) solver. (Strictly speaking, we use the `qutip.mesolve` [114] function which is a Lindblad master equation solver. This was so that, if needed, we could easily extend to open systems. `qutip.mesolve` defaults to `qutip.sesolve` in the closed system setting.) The function is handed the Hamiltonian, the initial state of the system and a list of times for which to obtain the state of the system. The function can also be handed a set of solver options, the default values for which can be found at [115]. The only parameter that we change from the default value is `nsteps` which is the maximum number of steps the ODE solver takes at each specified time. The default value for this is 1000. In order for the integration to converge, however, this value needed to be raised as high as 10^8 in the most extreme case.

Outside of our dynamic simulations, we represent the various Hamiltonians and state vectors using `numpy` arrays. In order to use `qutip.sesolve`, we must first create corresponding `qutip` quantum objects for each of the relevant Hamiltonians and states. This can be done straightforwardly in each case by handing the constructor, `qutip.Qobj` the associated `numpy` array. The time-dependent Hamiltonians are represented with `qutip.QobjEvo` objects. These are constructed out of the time-*independent* Hamiltonians ($\hat{H}_d$, $\hat{H}_p$ and $\hat{H}_c$ in our case) and their schedules in the form of functions.

`qutip.sesolve` returns a results object from which we can extract `numpy` arrays representing the state vectors to use in the rest of our code. Specifically, we obtain the overlaps of these state vectors with the instantaneous eigenstates, $|E_a(s)\rangle$, and with the problem eigenstates, $|E_a\rangle$.

4.4 Summary

In this chapter, we have outlined how we perform the numerical diagonalisation and closed system dynamics that make up a large portion of this thesis. We have also described how we perform a Hilbert space reduction which, alongside the scalable graphs introduced in chapter 3, allows us to study the scaling behaviour of the catalysts we examine in this dissertation.

Numerical diagonalisation of the annealing Hamiltonian is used throughout all of our results chapters to understand how the catalysts affect both the gap spectra and the evolution of the instantaneous eigenstates associated with the anneal. Simulation of the closed system dynamics is only used in chapters 8 and 9 where we explore diabatic annealing processes that are difficult to analyse using the annealing spectrum alone.

Chapter 5

Targeted XX-Couplings

This chapter lays the groundwork for the theory behind the targeted XX-couplings that we explore in this work. It introduces key ideas that will be returned to throughout this dissertation, and illustrates these ideas with numerical examples. The findings in this chapter give the first indications that the specific catalysts that we examine in this work have the desired effect of enhancing the gap minimum at a perturbative crossing. This chapter also establishes how the theory behind the targeted couplings relates to our graph structure.

The catalyst Hamiltonian, $\hat{H}_c$, with which we introduce the additional couplings enters $\hat{H}(s)$ as follows:

$$\hat{H}(s) = (1-s)\hat{H}_d + s(1-s)\hat{H}_c + s\hat{H}_p. \tag{5.1}$$

The pre-factor $s(1-s)$ ensures that these additional couplings leave the initial and final ground states unchanged such that the ground state we initialise in remains the equal superposition state and we can read off the solution from the final ground state. Throughout this chapter, and indeed this dissertation, we associate the terms *stoquastic* and *non-stoquastic* with the catalysts that we are using. It should be noted that the catalyst Hamiltonians are not in fact non-stoquastic in and of themselves but rather their addition results in a total Hamiltonian that is non-stoquastic for $s \neq 0, 1$. There has been much discussion in the literature regarding the importance of non-stoquasticity in QA [100, 108, 116]. In our use of the term we are simply describing a Hamiltonian with positive as well as negative off-diagonal elements and make no assumptions regarding the relationship between non-stoquasticity and computational complexity.

The catalysts that we consider in this section have the form

$$\hat{H}_c = J_{xx} \sum_{(i,j)\in E_{\mathrm{cat}}} \hat{\sigma}_i^x \hat{\sigma}_j^x. \tag{5.2}$$

That is, the catalysts consist of some sub-set, E_{cat}, of possible XX-couplings between vertices that are all introduced with the same sign and magnitude. We allow for the inclusion of vertex pairs in E_{cat} that do not correspond to edges in the problem graph. Note that a negative J_{xx} will result in the XX-couplings entering $\hat{H}(s)$ with the same sign as the driver couplings, resulting in a stoquastic Hamiltonian for all s. If we choose J_{xx} to be positive however, the total Hamiltonian will be non-stoquastic for $s \neq 0, 1$.

This chapter is structured as follows. Section 5.1.1 and 5.1.2 describe two theoretical approaches to understanding how the presence of a perturbative crossing will be affected by the introduction of different XX-couplings. Both approaches make use of perturbation theory. Numerical results that confirm the predictions made in sections 5.1.1 and 5.1.2 are presented in section 5.2. Section 5.3 then draws attention to the importance of the relative signs between the instantaneous vector components, how these are affected by the introduction of a non-stoquastic catalyst and what impact this has on the effectiveness of the catalysts. Finally, we take some time to discuss the implications of our results and how they could be used in section 5.4 before summarising our findings in section 5.5.

before summarising out findings

we summarise our findings in section 5.4 and discuss how they may be used in practice.

5.1 Theoretical motivation

5.1.1 Perturbative introduction of $(\hat{H}_d + \hat{H}_c)$ to $\hat{H}_p$

One can understand the effects of specific *stoquastic* XX-couplings on the presence of perturbative crossings in the annealing spectrum in much the same way that we understood the formation of these crossings in the first place (section 3.1). That is, we introduce $\hat{H}_c$ as an additional perturbation to $\hat{H}_p$, writing the perturbed Hamiltonian as

$$\hat{H}(\lambda) = \hat{H}_p + \lambda(\hat{H}_d + \hat{H}_c). \tag{5.3}$$

In equation 5.1 we see that $\hat{H}_d$ and $\hat{H}_c$ do not enter with the same prefactor. Assuming, however, that $\hat{H}_c$ enters the Hamiltonian linearly is an acceptable approximation for small values of λ. Such an approach is taken in [67] to understand why the introduction of a non-stoquastic all-to-all XX-field is able to elevate the exponentially closing gap in one of the example problems studied. Their setting is however somewhat different to ours in that the results hinge on the degeneracy of the states. We also note that [65] outlines some ideas similar to those we introduce here. However, the inclusion of XX-couplings is only a brief aside to illustrate their larger point which relates more closely to the formation of perturbative crossings in a standard QA protocol. Our aim here is to discuss these ideas more fully with the aim of drawing some general conclusions regarding the expected effects of different XX-couplings.

Carrying out the perturbation in equation 5.3, the perturbed problem states are given to second order by

$$\begin{aligned} E_a(\lambda) = E_a &+ \lambda(\langle E_a|\hat{H}_d|E_a\rangle + \langle E_a|\hat{H}_c|E_a\rangle) \\ &+ \lambda^2 \sum_{c\neq a} \left(\frac{|\langle E_c|\hat{H}_d|E_a\rangle|^2}{E_a - E_c} + \frac{|\langle E_c|\hat{H}_c|E_a\rangle|^2}{E_a - E_c} \right). \end{aligned} \tag{5.4}$$

As in the case without the catalyst, the first order corrections are zero since the problem states cannot be coupled to themselves by either $\hat{H}_d$

or $\hat{H}_c$. Substituting in $\hat{H}_d$ as before, and now also the catalyst from equation 5.2, we arrive at

$$E_a(\lambda) \quad = \quad E_a \, + \, \lambda^2 \left(\sum_{c \in N_x(a)} \frac{1}{E_a - E_c} + J_{xx}^2 \sum_{c \in N_{xx}(a)} \frac{1}{E_a - E_c} \right). \tag{5.5}$$

We have used $N_{xx}(a)$ to denote the set of states that a is coupled to by the catalyst.

As in our original discussion of the formation of perturbative crossings, let turn our focus to the states $a \in L_\delta$. From the arguments in section 3.1, we make the same assumption that $E_c > E_a$ for $c \in N_{xx}(a)$ and $a \in L_\delta$, meaning that all the contributions to the two sums will be negative. As in section 3.1, we re-write equation 5.5 to make this explicit:

$$E_a(\lambda) \quad = \quad E_a \, - \, \lambda^2 \left(\sum_{c \in N_x(a)} \frac{1}{|E_c - E_a|} + J_{xx}^2 \sum_{c \in N_{xx}(a)} \frac{1}{|E_c - E_a|} \right). \tag{5.6}$$

We to see that the catalyst couplings contribute to the perturbation in much the same way as the driver couplings. That is, a state in L_δ that is coupled to more low energy neighbours will receive a greater negative perturbation. As for how this can remove a perturbative crossing, consider a setting where

$$\sum_{c \in N_x(b)} \frac{1}{|E_c - E_b|} > \sum_{c \in N_x(a)} \frac{1}{|E_c - E_a|} \tag{5.7}$$

for two states ($b > a$) in L_δ such that a perturbative crossing forms between states $|E_a\rangle$ and $|E_b\rangle$. There may be a choice of XX-couplings between qubits which results in $N_{xx}(a)$ containing generally lower energy states than $N_{xx}(b)$, such that:

$$\sum_{c \in N_{xx}(b)} \frac{1}{|E_c - E_b|} < \sum_{c \in N_{xx}(a)} \frac{1}{|E_c - E_a|}. \tag{5.8}$$

Looking at equation 5.6, it is clear that increasing the magnitude of J_{xx} enhances the negative perturbation to both $|E_a\rangle$ and $|E_b\rangle$. However, the increase will be greater for $|E_a\rangle$ such that, for some J_{xx}, the magnitude of the second order correction to $|E_a\rangle$ will be larger than that to $|E_b\rangle$:

$$\sum_{c \in N_x(b)} \frac{1}{|E_c - E_b|} + J_{xx}^2 \sum_{c \in N_{xx}(b)} \frac{1}{|E_c - E_b|} < \sum_{c \in N_x(a)} \frac{1}{|E_c - E_a|} + J_{xx}^2 \sum_{c \in N_{xx}(a)} \frac{1}{|E_c - E_a|}. \quad (5.9)$$

In fact, equation 5.9 does not need to hold for us to expect the perturbative crossing to be lifted. It is sufficient that the difference in magnitude between the two sides of the equation is small enough (compared to the energy gap ΔE_{ab} in the problem spectrum) that the value of λ at which the crossing would occur sits outside the range for which perturbation theory remains valid. We would also expect any reduction in difference between the two sides of the equation to result in some gap enhancement since the crossing states will be more mixed and so will have a larger overlap.

To summarise, if a perturbative crossing exists between $|E_a\rangle$ and $|E_b\rangle$ ($b > a$), this crossing can be removed by introducing XX-terms that couple $|E_a\rangle$ to lower energy states than they couple $|E_b\rangle$ to. That this works is confirmed numerically in section 5.2. This is not dissimilar to the ideas introduced in [64]. However, in their case, this is achieved through inhomogeneous driving to re-scale the contributions from different coupled states rather than introducing new ones.

The fact that J_{xx} only appears in equation 5.6 squared suggests that this effect is insensitive to the sign with which these couplings enter $\hat{H}(s)$. From our numerical results, however, we find that the above conclusions only apply when the J_{xx} is introduced with the same sign as the driver. Introducing the XX-couplings in this way results in a stoquastic $\hat{H}(s)$ $\forall s$. It is not entirely clear how to reconcile the numerical results (as well as the theoretical results in the following section) with the perturbative argument we have made here. The best we can suggest is that this perturbative argument does not account for some of the additional complexities that can arise when $\hat{H}(s)$ is non-stoquastic – at least not without including higher order corrections.

Let us now consider how one might want to use the ideas introduced here to design a targeted catalyst. Typically, we are interested in removing perturbative crossings which involve the ground state since these are what will result in exponentially closing gaps between the instantaneous ground and first excited states. In order to remove these perturbative crossings in the way described in this section, one would need to be able to identify coupling(s) that result in $N_{xx}(0)$ containing states with very low energy. However, the capacity to identify such couplings assumes access to knowledge of the problem ground state which one may not have. Certainly, if one had full knowledge of the problem ground state this would mean that the optimal solution to the problem was already known in advance. In practice, therefore, the ideas introduced here may not be the most useful with regards to coupling selection – though it is not inconceivable that a setting may exist in which one has sufficient knowledge of the global optimum to select suitable couplings while still not knowing the optimal solution in full.

Far more useful, however, would be the capacity to design a catalyst Hamiltonian based of knowledge of the *local* optima. Indeed, understanding how this can be done and how effective the catalysts that have been designed in this way are is the focus of this thesis. In the following two sections we will see how couplings introduced with the *opposite* sign to the driver ($J_{xx} > 0$) have the potential to remove a perturbative crossing between two problem states, $|E_a\rangle$ and $|E_b\rangle$ ($b > a$), if chosen to give the *higher* energy state, $|E_b\rangle$, a low energy neighbourhood, $N_{xx}(b)$. This can be intuitively thought of as a reversal of the stoquastic approach introduced in this section. However, as we have already noted, the perturbative argument made here is not sufficient to differentiate between the two settings. We now go through a different perturbative argument which does allow us to understand the behaviour of the case where $J_{xx} > 0$.

5.1.2 Perturbative introduction of $\hat{H}_c$ to $\hat{H}(\lambda)$

In the previous section we introduced $(\hat{H}_c + \hat{H}_d)$ as a perturbation to $\hat{H}_p$ in order to understand how the introduction of $\hat{H}_c$ affected the presence of perturbative crossings in the spectrum. The approach we take in *this* section is to introduce $\hat{H}_c$ to the perturbed Hamiltonian, $\hat{H}(\lambda) = \hat{H}_p + \lambda \hat{H}_d$, from section 3.1. The new total Hamiltonian reads

$$\hat{H}(\lambda, \mu) = \hat{H}(\lambda) + \mu \hat{H}_c, \tag{5.10}$$

and the perturbed energies, to first order, are given by

$$\begin{aligned} E_a(\lambda, \mu) &= E_a(\lambda) + \mu \langle E_a(\lambda)| \hat{H}_c |E_a(\lambda)\rangle \\ &= E_a(\lambda) + \mu \left(\sum_{b=0}^{2^n-1} \langle E_a(\lambda)|E_b\rangle \langle E_b| \right) \hat{H}_c \left(\sum_{c=0}^{2^n-1} |E_c\rangle \langle E_c|E_a(\lambda)\rangle \right). \end{aligned} \tag{5.11}$$

Substituting in equation 5.2 for $\hat{H}_c$, we can write

$$E_a(\lambda, \mu) = E_a(\lambda) + 2\mu J_{xx} \sum_{(b,c)\in C_{xx}} \langle E_a(\lambda)|E_b\rangle \langle E_c|E_a(\lambda)\rangle, \tag{5.12}$$

where we have used C_{xx} to denote the set of problem state pairs that are coupled by $\hat{H}_c$. Note that every coupling included in E_{cat} will pair up the eigenstates of $\hat{H}_p$ – introducing $2^n/2$ pairs to C_{xx}. As such, C_{xx} will contain $n_{\text{edges}} \times 2^{(n-1)}$ pairs. Something similar to this is carried out in [100] for the specific Hamiltonian considered there.

We are interested in the effect of $\hat{H}_c$ on two perturbed states whose energies cross, resulting in a perturbative crossing in the annealing spectrum. Specifically, we are focusing our attention on crossings between states in L_δ. From the discussion in section 3.1 we know that $|E_a(\lambda)\rangle$ will be dominated by $|E_a\rangle$ and states that are close in energy and Hamming distance to $|E_a\rangle$. As such, we can expect larger terms to enter the sum in equation 5.12 if we choose XX-terms that couple $|E_a\rangle$ to states that are close in energy to it – as well as coupling the states that are close in energy and Hamming distance to $|E_a\rangle$ to each other. (Recall that, since L_δ are the lowest energy states, the problem states that will be close in energy to those in L_δ will be other low energy states.)

To understand what this means for the energies, $E_a(\lambda, \mu)$, we must gain some insight into the signs of the overlaps, $\langle E_a(\lambda)|E_b\rangle$, that appear in the sum, *i.e.*, the signs with which the different problem states enter $|E_a(\lambda)\rangle$. In the stoquastic setting, the ground state vector of the total Hamiltonian is guaranteed to contain no negative components. However, recall that $|E_a(\lambda)\rangle$ are the perturbed problem states and not the instantaneous eigenstates of the total Hamiltonian such that the perturbed GS, $|E_0(\lambda)\rangle$, may not always be the lowest energy state. In fact, in the settings we are most interested in (*i.e.*, settings where the GS undergoes a perturbative crossing), $|E_0(\lambda)\rangle$ will certainly not remain the lowest energy state as λ is increased. Furthermore, we also need to understand the signs of the higher energy states.

Fortunately, perturbation theory tells us that we can expect the prefactors of the problem states with the largest contributions to $|E_a(\lambda)\rangle$ to be positive. As such, we expect that introducing the catalyst with a positive J_{xx} will result in $E_a(\lambda, \mu > 0) > E_a(\lambda, 0)$ and introducing the catalyst with a negative J_{xx} will result in $E_a(\lambda, \mu > 0) < E_a(\lambda, 0)$. We will come back to the implications this has for the catalyst's effect on a perturbative crossing at the end of this section. First, we will go through the argument for why we expect the dominant problem states to enter $|E_a(\lambda)\rangle$ with positive coefficients.

Note that at this point we are considering perturbations to the *problem states* from the introduction of the *driver* (such that we can understand the corrections to the energies of these perturbed states as a result of a separate perturbation from the *catalyst*). Perturbative corrections of different orders introduce problem states to $|E_a(\lambda)\rangle$ with different Hamming distances from $|E_a\rangle$. Specifically, the k^{th} order corrections will introduce states that are a Hamming distance k from $|E_i\rangle$ as well as some states with a Hamming distance $< k$. (The first, second and third order corrections can be found in Appendix C). Let us consider some state, $|E_b\rangle$, that has a Hamming distance, d, from $|E_a\rangle$. The contribution of $|E_b\rangle$ to the perturbed state, $|E_a(\lambda)\rangle$, can come from multiple corrections of order $k \geq d$ however the largest contribution will come from the d^{th} order correction. And this we find will always be positive for $a \in L_\delta$.

The equation associated with the k^{th} order correction is

$$(E_a - \hat{H}_p)\,|E_a^{(k)}\rangle = \hat{H}_d\,|E_a^{(k-1)}\rangle - \sum_{i=2}^{k} E_a^{(i)}\,|E_a^{(k-i)}\rangle\,. \tag{5.13}$$

The sum starts at $i = 2$ and not one because in our setting the first order correction to the energies is zero. Pre-multiplying by $|E_b\rangle$ and rearranging we obtain

$$\langle E_b|E_a^{(k)}\rangle = \frac{\langle E_b|\hat{H}_d|E_a^{(k-1)}\rangle}{(E_a - E_b)} - \sum_{i=2}^{k} E_a^{(i)}\,\frac{\langle E_b|E_a^{(k-i)}\rangle}{(E_a - E_b)}. \tag{5.14}$$

The first term in equation 5.14 is responsible for the lowest order (and therefore presumably largest) contribution to the states that have a Hamming distance $d = k$ from $|E_a\rangle$. Starting with

$$|E_a^{(1)}\rangle = \sum_{b \neq a} \frac{\langle E_b|\hat{H}_d|E_a\rangle}{(E_a - E_b)}\,|E_b\rangle\,,$$

we can recursively determine that the first term in equation 5.14 will be a sum of terms of the following form: The numerators are a product of k matrix elements, $\langle E_b|\hat{H}_d|E_c\rangle$ (where the states, $|E_b\rangle$ and $|E_c\rangle$, can be the problem state being perturbed, $|E_a\rangle$, or any other problem state), while the denominators will be products of k energy differences, $(E_a - E_b)$. The matrix elements will all be either -1 or zero depending on whether or not the two problem states are one spin flip apart. Furthermore, because we are considering perturbations to states in L_δ, we can continue to assume that the energy differences, $(E_a - E_b)$, will be negative. As, a result, each term in the sum will be positive such that the lowest order term contributing to each problem state that enters the perturbative correction is also positive. It is of course possible that, as a result of higher order corrections, some problem states may end up with a negative pre-factor however it is reasonable to assume that these are likely to be negligible.

We have considered the effect of an XX-catalyst, of the form given in equation 5.2, on the perturbed energies of the low energy local optima. So far we have established the following:

- The catalyst will have the greatest effect on the perturbed energy, $E_a(\lambda)$, if the XX-terms are chosen to couple $|E_a\rangle$ to other low energy states and/or couple together low energy states that are close in Hamming distance to $|E_a\rangle$.
- The energy will *increase* if the catalyst is introduced with $J_{xx} > 0$ - *i.e.*, the *opposite* sign to the driver (*non-stoquastic* total $\hat{H}(s)$).
- The energy will *decrease* if the catalyst is introduced with $J_{xx} < 0$ - *i.e.*, the *same* sign as the driver (*stoquastic* total $\hat{H}(s)$).

Note that the final point is in agreement with the findings presented in the previous section. Thus, this approach also leads us to conclude that introducing negative XX-terms that couple the problem GS to other low energy states will help to remove any perturbative crossings it is participating in and hence enhance the minimum spectral gap between the instantaneous ground and first excited state. In addition to this however, the conclusions from this section also suggest that coupling together other low energy states that are simply close in Hamming distance to the GS could result in a similar effect.

In *contrast* to the previous section, the conclusions in this section suggest that catalysts with $J_{xx} > 0$ will have the *opposite* effect to those introduced with $J_{xx} < 0$ – *i.e.*, they will *increase* the energy of the perturbed state rather than reduce it. As such, these results suggest that a perturbative crossing between the GS and the a^{th}ES could be removed by instead introducing *positive* XX-terms that couple $|E_a\rangle$ to low energy states – or by coupling low energy states that are close in Hamming distance to $|E_a\rangle$ to each other. It is catalysts motivated by this observation on which the majority of this thesis is concerned. However, in the following, we present numerical results confirming the predictions regarding the stoquastic catalysts as well.

5.2 Numerical results

We now present numerical data confirming the validity of the conclusions reached in the preceding section. To recap: we expect an XX-catalyst to have the greatest impact on the energy of a perturbed state $|E_a(\lambda)\rangle$ ($a \in L_\delta$) when the XX-terms couple the problem state $|E_a\rangle$ to other low energy states and/or low energy states that are close in Hamming distance to $|E_a\rangle$ to each other. Throughout this work we will refer to this as selecting couplings that "target" the state $|E_a\rangle$. We expect the effect of such a catalyst to be an increase in the energy of $|E_a(\lambda)\rangle$ if the catalyst enters with the opposite sign to the driver ($J_{xx} > 0$) and a decrease in the energy if the catalyst enters with the same sign as the driver ($J_{xx} < 0$). As such, a perturbative crossing between two states, $|E_a\rangle$ and $|E_b\rangle$ ($a, b \in L_\delta$ and $b > a$), can potentially be removed by a selecting couplings that target $|E_a\rangle$ and choosing $J_{xx} > 0$ or by selecting couplings that target $|E_b\rangle$ and choosing $J_{xx} < 0$. Correspondingly, we expect that choosing the same couplings but flipping the signs of J_{xx} will result in the catalyst increasing the severity of the perturbative crossing.

To demonstrate these effects numerically we must establish what these coupling selections look like with respect to the problem instances we make use of in this work. The MWIS problem instances we construct result in perturbative crossings forming between states corresponding to the local optima associated with selecting all the vertices in one of

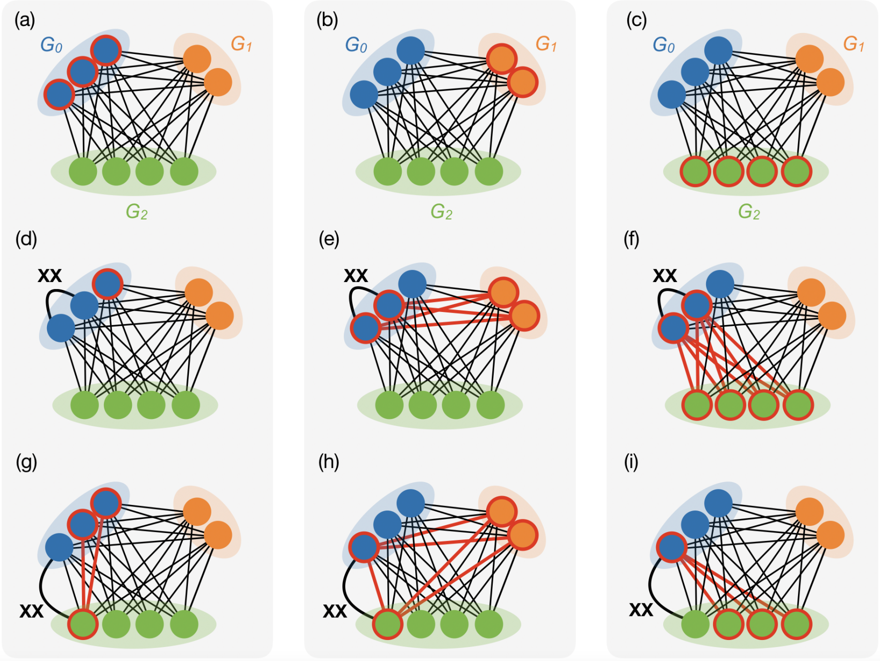

Figure 5.1: This figure demonstrates how introducing different XX-terms results in different couplings between the problem states. Each sub-figure represents a different set of vertices or, equivalently, an eigenstate of the problem Hamiltonian. The selected vertices are circled in red and any edge dependencies are highlighted. For instance, figure (g) shows the set containing two vertices in G_0 and one in G_2 which is encoded by the problem state with all the corresponding spins pointed up. The figures in the top row (a-c) show the local optima of the problem. Below these (d-f), we show the neighbours introduced to these states by an XX-term between two vertices in G_0. Finally, on the bottom row (g-i), we show the neighbours introduced to the local optima by an XX-term between G_0 and G_2.

the disconnected sub-graphs. To couple one of these states, $|E_a\rangle$, associated with the sub-graph, G_a, to a low energy state, we can introduce an XX-coupling between two vertices in G_a. This will couple $|E_a\rangle$ to a state corresponding to an independent set containing all but two of the vertices in G_a. It will also couple states corresponding to sub-sets of G_a (which are low energy states close in Hamming distance to $|E_a\rangle$) to each other. Meanwhile, the states corresponding to the other local optima will be coupled to dependent sets that contain all the vertices in their associated sub-graphs plus two from G_a.

This is demonstrated in figure 5.1 for a problem graph with $n_0 = 3$,

$n_1 = 2$ and $n_2 = 4$. Figures 5.1(a-c) illustrate the states corresponding to the three local optima of the problem with the selected vertices highlighted in red. Below these, in figures 5.1(d-f), we show the states that these local optima become coupled to through the introduction of an XX-term in G_0. Edge dependencies are highlighted in red. We see that the local optimum consisting of all the vertices in G_0 becomes coupled in an independent set while the other two become coupled to dependent sets. Because of how the energy cost penalising dependent sets is introduced (see chapter 3) the independent set will be one of the lowest energy states in the spectrum despite having only one vertex in the corresponding set. The dependent sets will be significantly higher in energy.

We also demonstrate in this figure the result of introducing an XX-coupling between two vertices in different sub-graphs – specifically, between one vertex in G_0 and another in G_2. The states that the independent set states (a-c) become coupled to are shown in the bottom row (d-f). As indicated by the highlighted red edges, all of these states correspond to independent sets. Again, the reason for using the graph structure that we do is not because of specific interest in the graph structure itself because but because of how we can relate it to the properties of the problem spectrum. As such these couplings are not of particular interest for our investigation. For completeness, we have included some results where couplings have been placed between sub-graphs in Chapter 6. However, we do not touch on these for the remainder of this Chapter.

5.2.1 Gap minima

We now present, in figure 5.2, numerical data for the size of the gap minimum at the perturbative crossing for a problem graph with the sub-graph sizes $n_0 = 3$, $n_1 = 4$ and $n_2 = 2$ – see figure 3.3(b) for an illustration. The result of allocating these sizes to the sub-graphs is that a perturbative crossing forms between the ground and first excited state

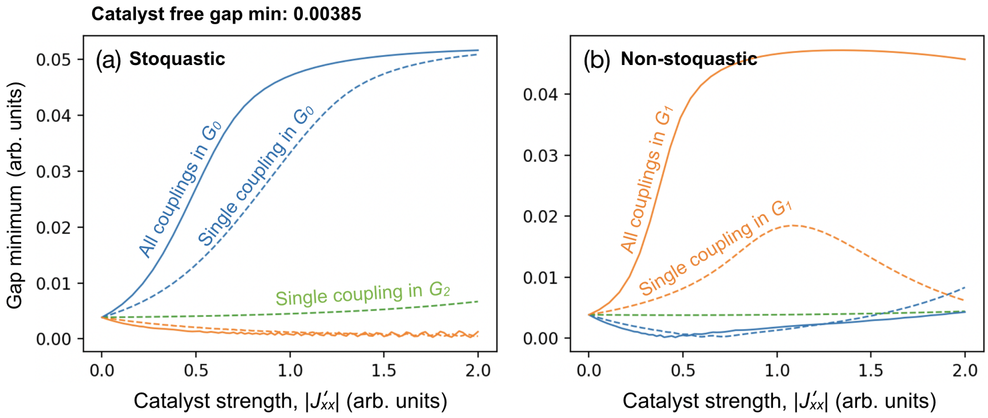

Figure 5.2: Size of the gap minimum as a function of catalyst strength, $|J'_{xx}|$, for a problem with parameters $n_0 = 3$, $n_1 = 4$, $n_2 = 2$, $W_0 = 1.70$, $W_1 = 1.35$, $W_2 = 1.00$ and $J_{zz} = 35$. Plots (a) and (b) show the results when the catalyst is introduced with the same and the opposite sign to the driver respectively (*i.e.*, the stoquastic and non-stoquastic case). In blue, orange and green we show the results when the catalyst consists of couplings within G_0, G_1 and G_2 respectively. Results using all the possible couplings within the sub-graph are plotted with solid lines and the results when only one coupling is used are plotted with dashed lines.

while the second excited state does not participate in a perturbative crossing. This effectively allows for a "dud" coupling location, the use of which is expected to have minimal effect on the size of the gap minimum. The weights allocated to the sub-graphs are $W_0 = 1.70$, $W_1 = 1.35$ and $W_2 = 1.00$ and the edge penalty is $J_{zz} = 35$. These parameters result in perturbative crossings closer in nature to those produced by the WGS parameters from section 3.5. As we will be exploring, in chapter 7, the effects of the catalyst can be a little more involved in settings associated with a more severe gap scaling.

We choose to scale the catalyst strength by the number of couplings present in the catalyst as follows:

$$\hat{H}_c = \frac{J'_{xx}}{\sqrt{n_{\text{edges}}}} \sum_{(i,j)\in E_{\text{cat}}} \hat{\sigma}_i^x \hat{\sigma}_j^x, \tag{5.15}$$

where n_{edges} is the number of terms in E_{cat}. The scaling compensates for the fact that introducing more couplings will result in more terms entering the sums in equations 5.6 and 5.12 and so larger perturbations. This allows us to make a more direct comparison between the effect of using different numbers of couplings in the catalyst. Taking the square root of n_{edges} in equation 5.15 was motivated specifically by equation 5.6 where J_{xx} enters at second order. However, the catalyst strength, J_{xx}, enters equation 5.12 at first order. In this sense there is no clear "correct" way to scale the catalyst strength and so when comparing the effects from different numbers of couplings one should not focus on the precise values of J'_{xx} at which the effects of the catalyst are observed.

Figure 5.2(a) shows how the size of the gap minimum varies with $|J'_{xx}|$ for a stoquastic catalyst ($J'_{xx} < 0$). As in section 3.5, we present the data with arbitrary units – *i.e.*, we have scaled the catalyst strength and energy gaps to the driver Hamiltonian. (Indeed, since we are often only interested in a comparison between the gap sizes with and without a catalyst, we present our data this way throughout this chapter as well as in chapters 6, 7 and parts of chapter 9.) The results are colour coded based on which sub-graph the XX-couplings are introduced in: blue for G_0, orange for G_1 and green for G_2. Results when a single coupling within the sub-graphs are used are plotted in dashed lines and results when all possible couplings within the sub-graphs are used are plotted in solid lines. There is no solid line corresponding to G_2 since there is only one possible coupling.

We see that introducing XX-couplings into G_0 (*i.e.*, selecting the couplings to target $|E_0\rangle$) results in an enhancement to the gap minimum. Meanwhile, introducing XX-couplings into G_1 (*i.e.*, selecting the couplings to target $|E_1\rangle$) results in a *suppression* of the gap minimum. This is in line with the theoretical predictions from the previous two sections. As expected, minimal effect is seen from introducing a coupling within G_2. Some slight gap enhancement is seen which may be a result of some contribution to the perturbative crossing from the perturbed problem state $|E_2(\lambda)\rangle$. (Recall that greater mixing of states will result in a larger gap minimum due to an increased tunnelling rate). However, since this enhancement is very minor, we will not dwell on understanding its origin.

Comparing the single-coupling results to those where all-to-all couplings within a given sub-graph are used, we see qualitatively similar effects. In the case of coupling(s) being included in G_0, the size of the gap minimum approaches the gap size at the end of the anneal, $\Delta E_{01} (= 0.052)$, which is the maximum enhancement possible since $\hat{H}(s = 1)$ is unchanged. The size of the gap minimum continues to tend to ΔE_{01} as $|J'_{xx}|$ is increased above 2. The behaviour of the stoquastic catalysts containing couplings in G_1 is for the size of the gap minimum to tend to 0. In addition to this however, oscillations in the data can be seen from around $|J'_{xx}| = 1$. Note that the "jagged" look to these oscillations is not an artifact of the sampling grain of either s for the individual spectra or $|J'_{xx}|$. These oscillations do not impact the main results presented here or in the rest of this dissertation and so, once again, we have chosen not to dwell on their origin.

Let us now turn to figure 5.2(b) which shows the same data for a *non*-stoquastic catalyst ($J'_{xx} > 0$). Here we can see that it is coupling in G_1 that enhance the size of the gap minimum and couplings in G_0 that suppress it – in line with the theoretical predictions from the previous section. As before, introducing a coupling within G_2 has minimal effect. In this case, it has even less effect which is to be expected since increasing the energy of $|E_2(\lambda)\rangle$ will not increase its involvement in the perturbative crossing.

There are some notable differences between the stoquastic and non-stoquastic case – beyond the results for G_0 and G_1 being flipped. In the stoquastic setting, monotonic enhancement of the gap minimum towards ΔE_{01} is seen when couplings in G_0 are used. Conversely, in the non-stoquastic setting (using couplings in G_1) the gap minimum reaches a maximum value for some $|J'_{xx}|$ before decreasing again. Similarly, rather than tending to zero, introducing non-stoquastic couplings into G_0 results in the size of the gap minimum approaching zero before increasing again. We can relate this behaviour to the appearance of negative problem state overlaps which arise for sufficiently large J'_{xx}. This will be discussed in section 5.3.

Returning to the use of non-stoquastic couplings in G_1, there is also a notable difference in the performance between the single and the all-to-all coupling cases. When all-to-all couplings are used, the maximum gap enhancement reached by the catalyst is comparable to that reached by the stoquastic G_0 catalyst. However, the enhancement achieved when only a single coupling is used is significantly lower. This can also be accounted for by considering the signs of the problem state overlaps.

So far we have confirmed that we are able to affect the gap minima at the perturbative crossings using XX-couplings in a way that is consistent with the theory introduced in the previous sections. We have also highlighted some key differences between the stoquastic and non-stoquastic settings and suggested that their origin may relate to appearance of negative components to the instantaneous state vectors. This discussion will be expanded in section 5.3. We will now examine how the introduction of the different XX-couplings affect the evolution of the instantaneous ground state vector – specifically the presence of the different problem states involved in the perturbative crossings. This will help confirm that the gap enhancement and suppression we have seen in this section are indeed the results of the predicted energy shifts to the perturbed problem states.

5.2.2 Evolution of the instantaneous GS

In this section we present numerical results for the evolution of the instantaneous GS vector for an example instance with $n_0 = 2$, $n_1 = 3$ and $n_2 = 4$. This selection of sub-graph sizes results in all three of the problem states participating in perturbative crossings. An illustration of the graph structure can be seen in figure 5.3(a) and a cartoon showing the expected perturbations is shown in figure 5.3(b). We can see that a perturbative crossing forms between $|E_0(\lambda)\rangle$ and $|E_1(\lambda)\rangle$ which takes place

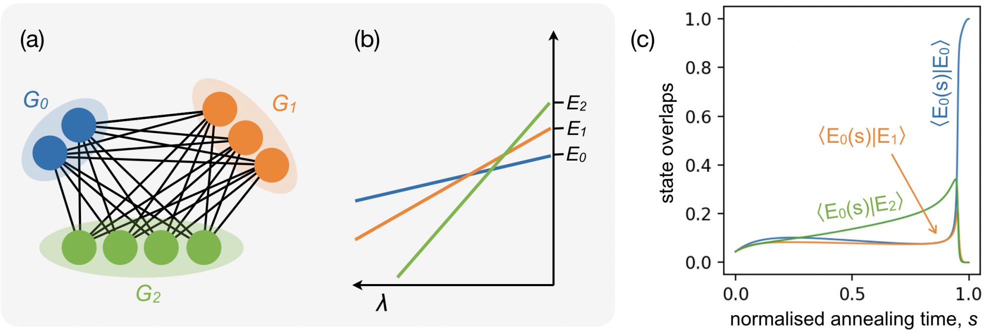

Figure 5.3: (a) shows a problem graph with $n_0 = 2$, $n_1 = 3$ and $n_2 = 4$. A cartoon illustrating the expected perturbations to the three lowest energy states of the corresponding problem Hamiltonian is shown in (b). (c) shows numerical results for the overlap of the instantaneous ground state with the three lowest energy problem states when $W_0 = 1.70$, $W_1 = 1.35$, $W_2 = 1.00$ and $J_{zz} = 35$.

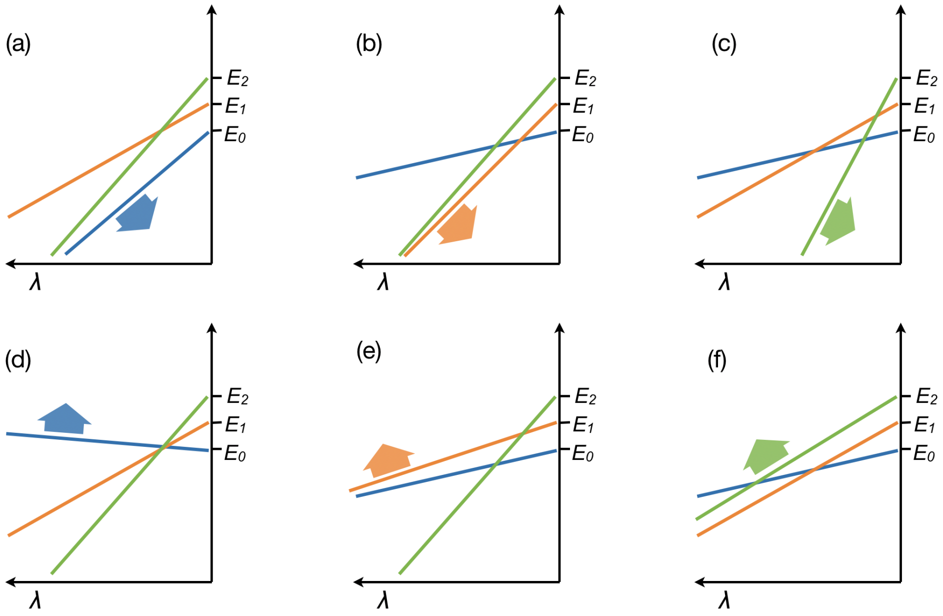

Figure 5.4: Cartoons depicting the action of different catalyst Hamiltonians on the perturbed problem states depicted in figure 5.3(b). The catalysts are as follows: (a) stoquastic in G_0, (b) stoquastic in G_1, (c) stoquastic in G_2, (d) non-stoquastic in G_0, (e) non-stoquastic in G_1, (f) non-stoquastic in G_2.

between the instantaneous first and second excited state. There are also two perturbative crossings involving $|E_2(\lambda)\rangle$ – one with $|E_1(\lambda)\rangle$ and one with $|E_0(\lambda)\rangle$. These form between the instantaneous first and second excited states and the ground and first excited states respectively. We expect the instantaneous ground state vector to have a sizable contribution from $|E_2\rangle$ prior the AC in which it is involved before transitioning to being dominated by $|E_0\rangle$. Figure 5.3(c) shows numerical results for the evolution of the instantaneous GS vector in terms of its overlap with the three lowest energy problem states. We see that at $s \approx 0.9$, $|E_0\rangle$ and $|E_1\rangle$ have similar contributions while $|E_2\rangle$ has a much greater presence. The AC occurs at $s \approx 0.95$ at which point the instantaneous GS transitions sharply into being dominated by $|E_0\rangle$.

Now let us consider the impact of different catalyst Hamiltonians on the evolution of the instantaneous ground state. Figure 5.4 shows six cartoons illustrating the impact of different choices of XX-couplings on the perturbed problem states associated with the example in figure 5.3. The top (a-c) and bottom (d-f) rows illustrate the expected effect of the stoquastic and non-stoquastic catalysts respectively. The three columns, going from left to right, correspond to the selection of XX-couplings within G_0, G_1 and G_2 respectively. These cartoons will help aid our discussion of the numerical results.

The plots in 5.5 show the overlap of the instantaneous GS with $|E_0\rangle$, $|E_1\rangle$ and $|E_2\rangle$ in blue, orange and green respectively. The dashed lines (which are the same in each plot) show the results without a catalyst and the solid lines show the results using different catalysts in each plot. Where it is hard to see the dashed lines we advise the reader to refer to figure 5.3(c). Each of the catalysts contains a single coupling introduced with strength $|J'_{xx}| = |J_{xx}| = 0.5$. This value has been chosen to be low enough for there to be minimal negative components to the GS vector in the non-stoquastic setting but high enough that the changes to its evolution are easily observable. (We will not always restrict ourselves to the regime where $|E_0(s)\rangle$ has negligible negative components however

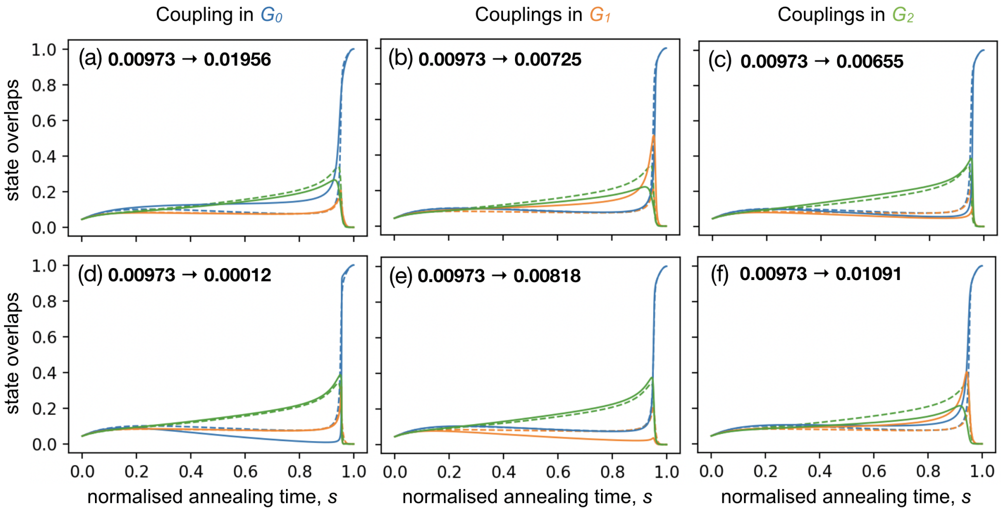

Figure 5.5: The plots in this figure show the absolute value of the overlap of the instantaneous ground state with the three lowest energy problem states for an instance with the parameters: $n_0 = 2$, $n_1 = 3$, $n_2 = 4$, $W_0 = 1.70$, $W_1 = 1.35$, $W_2 = 1.00$ and $J_{zz} = 35$. The overlaps with $|E_0\rangle$, $|E_1\rangle$ and $|E_2\rangle$ are plotted in blue, orange and green respectively. Each plot shows the effect on the spectrum of a different catalyst Hamiltonian. In each case the overlaps without the presence of the catalyst are plotted with dashed lines (these are the same in each plot) and the solid lines show the overlaps with the catalyst introduced. The top and bottom rows show the results when J'_{xx} is negative and positive respectively and the first, second and third columns show the results when the catalyst is introduced between vertices in G_0, G_1 and G_2 respectively. In each case only one coupling in the sub-graph is used. The catalyst strength is $|J'_{xx}| = 0.5$. The resultant change to the size of the gap minimum is given on each plot.

they would only serve to complicate the discussion at this point.) The six plots correspond to the cartoons in figure 5.4. That is, the top row (a-c) shows the results when a stoquastic catalyst is used and the bottom row (d-f) shows the results for a non-stoquastic catalyst. Results when the coupling is placed between two vertices in G_0, G_1 and G_2 are shown in the first, second and third columns respectively.

Looking at the effect of introducing a stoquastic catalyst into G_0 we see an increase in the presence of $|E_0\rangle$ and, correspondingly, a slight suppression to the presence of $|E_2\rangle$. The presence of $|E_1\rangle$ is largely unchanged. Looking at the associated cartoon in figure 5.4(a), we can understand these changes to be a result of the reduction in energy to the perturbed problem GS, bringing a greater contribution from $|E_0\rangle$ to the

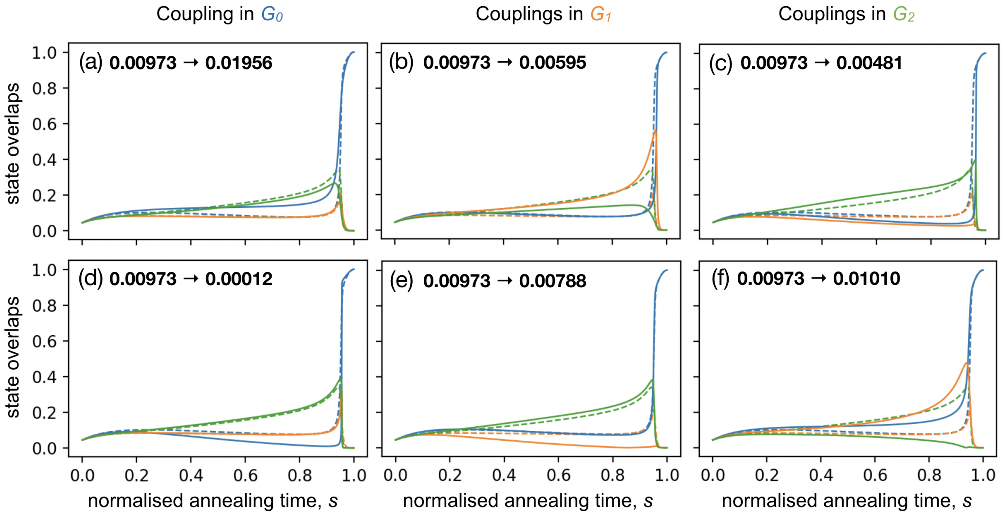

Figure 5.6: The plots in this figure show the absolute value of the overlap of the instantaneous ground state with the three lowest energy problem states for an instance with the parameters: $n_0 = 2$, $n_1 = 3$, $n_2 = 4$, $W_0 = 1.70$, $W_1 = 1.35$, $W_2 = 1.00$ and $J_{zz} = 35$. The overlaps with $|E_0\rangle$, $|E_1\rangle$ and $|E_2\rangle$ are plotted in blue, orange and green respectively. Each plot shows the effect on the spectrum of a different catalyst Hamiltonian. In each case the overlaps without the presence of the catalyst are plotted with dashed lines (these are the same in each plot) and the solid lines show the overlaps with the catalyst introduced. The top and bottom rows show the results when J'_{xx} is negative and positive respectively and the first, second and third columns show the results when the catalyst is introduced between vertices in G_0, G_1 and G_2 respectively. In each case, all potential couplings within the sub-graph are used. The catalyst strength is $|J'_{xx}| = 0.5$. This strength is scaled as shown in equation 5.15. The resultant change to the size of the gap minimum is given on each plot.

instantaneous GS. Conversely, a non-stoquastic XX-coupling between vertices in G_0 suppresses the presence of $|E_0\rangle$ as the perturbed problem GS is increased – as illustrated in figure 5.4(d). In this case, $|E_0\rangle$ has a near zero presence immediately prior to the perturbative crossing such that there is a very small overlap between the instantaneous GS before and after the AC – resulting in a highly suppressed tunnelling rate. As expected, the stoquastic catalyst enhances the size of the gap minimum while the non-stoquastic catalyst suppresses it. Accompanying this, we see that the exchange at the AC is visibly softened in figure 5.5(a) while in figure 5.5(d) it becomes more abrupt.

Let us now turn to the case where we introduce the XX-coupling between two vertices in G_1 – figures 5.5(b) and (e). Similar to the G_0 coupling, we see that in the stoquastic case, the presence of $|E_1\rangle$ is enhanced while, in the non-stoquastic case, its presence is suppressed. The changes in size of the gap minimum are less pronounced for these two catalysts however we do see some suppression to the gap minimum for both. For the stoquastic catalyst, we see that the presence of $|E_1\rangle$ is now greater than that of $|E_2\rangle$ in the catalyst free setting. Meanwhile, the presence of $|E_0\rangle$ is largely unaffected. Effectively, the reduction to the energy of the perturbed problem 1ES, as shown in 5.4(b), has replaced the original perturbative crossing with a more severe one. The gap suppression in figure 5.5(e) may be a result of the instantaneous GS being less mixed prior to the crossing.

Finally, we turn our attention to the case where the XX-coupling is in G_2. As before, we see the expected enhancement and suppression to the presence of $|E_2\rangle$ in the stoquastic and non-stoquastic settings respectively. In the stoquastic setting, this enhancement is accompanied by a suppression to the presence of $|E_0\rangle$ and $|E_1\rangle$ and a sharpening of the exchange at the AC. As expected, this results in a reduction to the size of the gap minimum. In the non-stoquastic setting, the suppression to $|E_2\rangle$ is accompanied by an enhancement to the presence of both $|E_0\rangle$ and $|E_1\rangle$ and the exchange at the AC becomes dominated by an exchange between $|E_0\rangle$ and $|E_1\rangle$. The gap minimum is very slightly enhanced. However, due to the 2ES-GS perturbative crossing being replaced by a 1ES-GS perturbative crossing, as seen in 5.4(f), this enhancement is minimal.

Figure 5.6 shows equivalent data to that in figure 5.5 but with catalysts containing all the couplings in the specified sub-graph. The findings are very much the same but with small quantitative differences. (Note that we plot the absolute value of the state overlaps. In figure 5.5 this was irrelevant because all the overlaps were positive. However, in figure 5.6, two state overlaps dip very slightly into the negative. This can be seen in figure 5.6(e) at $s \approx 0.9$ and in figure 5.6(f) at $s \approx 0.95$.) That the results using a single coupling and those using all-to-all couplings with a scaled $|J'_{xx}|$ are qualitatively the same gives further weight to the theoretical arguments put forward in the previous section. *i.e.*, the effect of the catalyst on the perturbative crossings is through cumulative perturbations to the energies, $|E_a(\lambda)\rangle$, in line with equations 5.6 and 5.12.

Effectively, we can think of introducing a stoquastic (non-stoquastic) catalyst that targets a particular problem state as guiding the anneal towards (away) from that state. As such, we would expect a stoquastic catalyst targeting the problem GS, to always offer significant enhancement to the gap minimum between the ground the first excites state. Furthermore, if perturbative crossings are the only source of gap minima in the annealing spectrum, we might expect such a catalyst to always result in $E_{01}^{\text{min}}(s)] \approx E_{01}(s=1)$ for sufficiently large $|J'_{xx}|$.

The capacity for the non-stoquastic catalysts to enhance the gap minimum however will be more dependent on the precise structure of the perturbative crossings in the annealing spectrum. In the previous section, we saw that the all-to-all non-stoquastic catalyst targeting $|E_1\rangle$ was able to achieve similar gap enhancement to the stoquastic catalysts targeting $|E_0\rangle$ (figure 5.2). There, $|E_1\rangle$ was the only state engaged in a perturbative crossing with $|E_0\rangle$ such that guiding the anneal towards $|E_0\rangle$ or away from $|E_1\rangle$ have similar effects. For the problem instance we have used as our example in this section, this is not the case. Looking at figure 5.3(b) we see that, even if we are able to remove the perturbative crossing between $|E_0\rangle$ and $|E_2\rangle$ by guiding the anneal away from $|E_2\rangle$, we will be left with a perturbative crossing between $|E_0\rangle$ and $|E_1\rangle$ – as seen in figures 5.5(f) and 5.6(f). Chapter 9 of this dissertation explores the possibility of removing multiple perturbative crossings with a non-stoquastic catalyst that contains couplings targeting different problem states.

We now turn to turn our attention to the signs of the instantaneous GS vector components and how these relate to the effect of the catalyst.

5.3 Signs of the ground-state vector components

In the preceding sections we have seen that the stoquastic and non-stoquastic catalysts we examine have, in many ways, inverted effects. There are however some qualitative differences between the two cases which we have suggested may relate to how the signs of the instantaneous vector components change in response to the introduction of the catalyst. That the relative signs between vector components corresponding to different problem states could play an important role in understanding the effect of a catalyst Hamiltonian is highlighted in [100]. With regards to our work, recall that the theoretical arguments we presented in section 5.1 hinged on the non-negligible vector components of the relevant states being non-negative. And while we were able to show that this was indeed the case prior to the introduction of the catalyst, it is reasonable to expect this condition to no longer be met for sufficiently large J_{xx}.

In this section we explore how the introduction of non-stoquastic XX-couplings affect the signs that the different vector components take over the course of the anneal. We will focus our attention on the instantaneous GS where we are able to make some analytical observations. It is also perturbative crossings with the GS that are of the greatest interest since these are what will bottleneck the algorithm. In particular, we will examine how the coupling choices affect the relative signs between vector components corresponding to different problem states and at what point in the anneal we expect negative GS vector components to emerge. We relate these findings back to how we can expect the effect of a non-stoquastic catalyst to diverge from the predictions in section 5.1 as J_{xx} increases. The conclusions from this section will also be particularly helpful to our discussion in chapter 7.

We begin by recounting an argument introduced in [100]. In this work, the author uses the expectation value of the total Hamiltonian to motivate the introduction of specific XX-couplings in order to achieve particular relative signs between certain GS vector components. Here we effectively do the reverse in that we consider the introduction of XX-terms that couple problem states together as motivated by the preceding sections and seek to understand what this means for the signs with which these problem states contribute to the instantaneous GS. We also use this energetic argument to derive an approximate expression for the value of s at which the first negative vector components appear in the GS. We supplement this discussion with numerical results using our example graphs. These numerical results show good agreement with the theoretical predictions.

The expectation value of $\hat{H}(s)$ for the instantaneous ground state is given by

$$\begin{aligned}\langle \hat{H}(s) \rangle = \langle E_0(s) | \, \hat{H}(s) \, | E_0(s) \rangle = \\ &- (1-s) \sum_{(a,b) \in C_x} \langle E_0(s) | E_a \rangle \langle E_b | E_0(s) \rangle \\ &+ s(1-s) J_{xx} \sum_{(a,b) \in C_{xx}} \langle E_0(s) | E_a \rangle \langle E_b | E_0(s) \rangle \\ &+ s \sum_{a=0}^{2^n - 1} E_a | \langle E_0(s) | E_a \rangle |^2, \end{aligned} \tag{5.16}$$

where C_x and C_{xx} denote the set of pairs of states that are coupled by the driver and the catalyst respectively. The final term in this expression, which arises as a result of the problem Hamiltonian, is irrelevant to our discussion since it is unaffected by the signs of the vector components. With respect to the driver term, which enters the expression with a negative sign, it is energetically favourable for components corresponding to pairs of states in C_x to have the same sign. This can only be true for every pair if all the components have the same sign. In the stoquastic setting, where J_{xx} is negative, it will also be energetically favourable for pairs of states in C_{xx} to have the same sign. That the lowest energy is achieved by all the components of the ground state having the same sign is consistent with what we already know from the Perron-Frobeneus theorem – that the GS of a stoquastic Hamiltonian will have only positive vector components. When J_{xx} is positive however, pairs of states appearing in C_{xx} will reduce the energy when the corresponding components have *opposite* signs to each other.

Let us now consider how this may affect the signs of the GS vector components over the course of the anneal. At $s = 0$, when the catalyst contribution is zero, it is energetically favourable for all the components to have the same sign. However, for high enough J_{xx}, we may expect one or more of the terms in the catalyst sum to become dominant enough at some s that the energy would be lowered by certain components becoming negative. Determining what combination of components being negative will minimise the energy at any given s is in general likely to be very complicated. Not only is there the balance between the driver and the catalyst terms in equation 5.16 to consider but, if there are multiple couplers in the catalyst, there may be conflicting pairs of states whose components would lower the energy by having the opposite sign.

The situation becomes a little easier to understand if we limit ourselves to a single XX-coupling such that there are no conflicting pairs of coupled states with respect to the catalyst. In the limit where $1/J_{xx} \to 0$ (*i.e.*, the catalyst dominates significantly over the driver), it would be energetically favourable for the components of the instantaneous GS vector to split into two oppositely signed sets with one problem state from each pair in C_{xx} appearing in each set. In practice we might expect some of the pairs entering the sum to have a larger product such that they take opposite signs first while it remains energetically favourable for others to keep the same sign as each other. The non-stoquastic catalysts we consider here are those containing XX-terms chosen to couple some low energy local optimum, $|E_a\rangle$, participating in a perturbative crossing with the ground state, to another low energy state, $|E_{N_{xx}(a)}\rangle$. By the fact that there is a crossing between $|E_a(\lambda)\rangle$ and $|E_0(\lambda)\rangle$ we can expect both $|E_a\rangle$ and $|E_{N_{xx}(a)}\rangle$ to have a large contribution to the instantaneous GS.

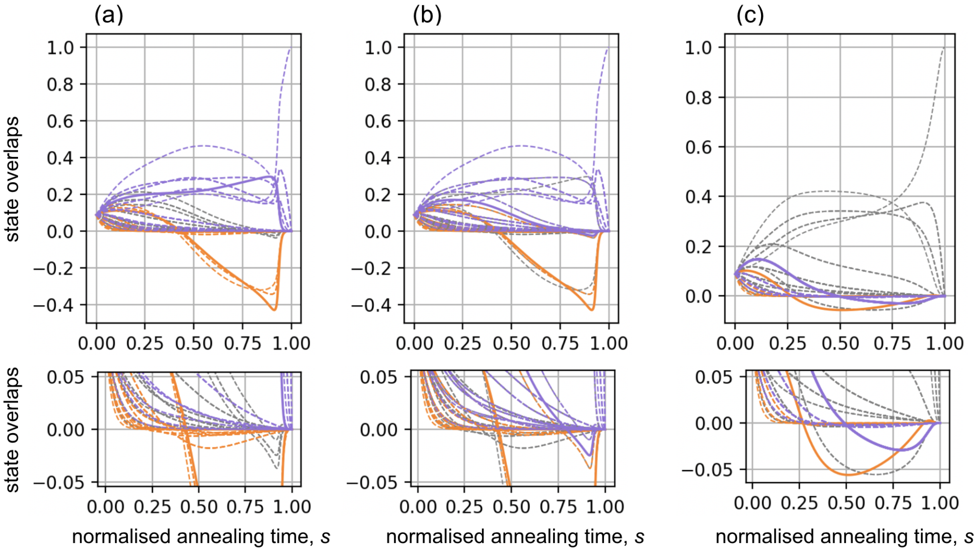

Figure 5.7: Plots showing the evolution of the instantaneous GS vector in terms of its overlap with the eigenstates of $\hat{H}_p$. The parameters associated the problem instance are $n_0 = 3$, $n_1 = 4$, $W_0 = 1.37$, $W_1 = 1.00$ and $J_{zz} = 37.5$. Plots (a) and (b) show the results when a catalyst Hamiltonian containing an XX-coupling with $J_{xx} = 2$ is introduced between the first two vertices in G_1 and plot (c) shows the results when all XX-couplings in G_1 are included. In all the plots, the overlap with the problem 1ES, $|\downarrow\downarrow\downarrow\uparrow\uparrow\uparrow\uparrow\rangle$, is plotted with a solid orange line. In (a), the solid purple line shows the overlap with $|\downarrow\downarrow\downarrow\downarrow\downarrow\uparrow\uparrow\rangle$ – *i.e.*, the state that the problem 1ES becomes coupled to by the catalyst. In (b), the solid purple line shows the overlap with $|\downarrow\downarrow\downarrow\downarrow\uparrow\downarrow\uparrow\rangle$. In (c) the overlaps with all the states that have two up and two down spins in G_1 are plotted with a solid purple line. In this case, the catalyst does not break the symmetry of the problem and so their evolutions are all the same. The rest of the states are colour coded according to which of these states they are closest to in Hamming distance. Equidistant states are plotted in grey. The plots on the bottom row show the same data as the top row plots but for a smaller range around zero.

As such, we may expect this to be the first term in the catalyst sum to affect the signs of the problem state overlaps of the GS and so for the corresponding GS components, $\langle E_0(s)|E_a\rangle$ and $\langle E_0(s)|E_{N_{xx}(a)}\rangle$, to take opposite signs. If the product of the other pairs of overlaps that enter the sum are significantly smaller we might expect the driver terms to dictate the rest of the signs such the rest of the overlaps share the sign of whichever state, $\langle E_0(s)|E_a\rangle$ or $\langle E_0(s)|E_{N_{xx}(a)}\rangle$, they are closer to.

Looking at the signs of the GS coefficients for various example graphs we find that this is indeed the case. We present, as an example, results for a bipartite instance in which an XX-term is introduced that couples $|E_1\rangle$ to another low energy state, $|E_{N_{xx}(1)}\rangle$ – see figure 5.7(a). The overlaps of the instantaneous GS with $|E_1\rangle$ and $|E_{N_{xx}(1)}\rangle$ are plotted in solid orange and purple respectively. The overlaps with the states that are closer in Hamming distance to each are plotted in the corresponding colours with dashed lines. States that are equidistant from the two states are shown in grey. We see that, at $s \approx 0.4$, the GS overlaps with $|E_1\rangle$ and $|E_{N_{xx}(1)}\rangle$ take different signs – with $\langle E_0(s)|E_1\rangle$ becoming negative. Furthermore, we see that the overlaps with the states that are closer in Hamming distance to $|E_1\rangle$ also become negative.

In figure 5.7(b) we show the same data but with the solid purple line highlighting the overlap with a state that is symmetrically equivalent to $|E_{N_{xx}(1)}\rangle$ in the catalyst free setting – both represent sets containing all but two vertices in G_1. The dashed purple lines show the overlaps with states that are closer in Hamming distance to this state than $|E_1\rangle$. We see that the states sorted in this way do not consistently share the same signs – lending weight to the arguments we have presented regarding the key elements that contribute to the signs taken by the problem state overlaps.

We do not necessarily expect this exact pattern to hold for all settings where a single XX-term is introduced – *i.e.*, the largest pair that enters the catalyst sum in equation 5.16 dictates two overlaps that take the opposite sign to each other while the rest of the states become sorted by which state in this pair they have the shortest Hamming distance from. It may be that more than one term in this sum becomes sufficiently large, for comparable J_{xx} and s, such that the states having opposite signs is more energetically favourable than the states sorting themselves by Hamming distance from one pair of states coupled by the XX-term.

The crucial observation however is this: We have motivated the introduction of XX-terms that couple problem states together that have a large presence in the perturbed problem state, $|E_a(\lambda)\rangle$, responsible for a perturbative crossing with the GS. The purpose of these couplings is to introduce large terms into the sum in equation 5.12 so that $|E_a(\lambda)\rangle$ receives large positive perturbations. Such couplings will only result in a *positive* perturbation if the overlaps with the pairs of coupled problem states share the same sign. The preceding discussion however suggests that these overlaps will take opposite signs for sufficiently large J_{xx} precisely *because* of their large contribution to $|E_a(\lambda)\rangle$. At this point equation 5.12 implies that continuing to increase J_{xx} will result in a *negative* perturbation to $|E_a(\lambda)\rangle$. In contrast to the stoquastic case, this suggests a limit to the gap enhancement that can be achieved in this way before we start to see it begin to close again – which is in line with the numerical results presented in figure 5.2.

That the gap enhancement from the non-stoquastic catalyst is expected to peak for some J_{xx} rather than continue to increase monotonically does not necessarily impose a problem for the use of such a catalyst. The crucial consideration will be how this gap enhancement scales with system size – which we investigate in the following chapters. We have also so far only focused our attention on the case where the catalyst contains a single XX-term. As we will now see, the situation is somewhat helped by the introduction of additional couplings.

In figure 5.7(c) we show the evolution of the instantaneous GS when the catalyst contains all possible XX-couplings within G_1. The overlaps with the different problem states are coloured as in figure 5.7(a). In this case however all the states corresponding to sets consisting of all but two of the vertices in G_1 are symmetrically equivalent and so their overlaps with the GS have the same evolution (shown by the solid orange curve). Here we see that, for values of s at which we see negative overlaps, the overlap of the ground state with $|E_1\rangle$ does *not* consistently have the opposite sign to the states it becomes coupled to by the catalyst. We also see that the magnitude of the negative overlaps are suppressed compared to the case where a single coupling is used.

We suggest that these findings relate to the fact that the catalyst now contains conflicting pairs of states in C_{xx} with non-negligible contributions to the GS. That is, $|E_1\rangle$ is coupled to each state corresponding to a set that consists of all but two vertices in G_1 and these states also become couples to each other. It is clear that not all of the pairs of states can have overlaps with the opposite sign to each other. While this particular observation relates specifically to our graph structure and choice of catalyst, we suggest that introducing more XX-couplings will tend to result in more conflicting pairs of problem states such that these findings are likely to extend to other settings as well. This suggests that using multiple couplings will tend to suppress the problem of key pairs of overlaps taking the opposite sign and so result in the capacity for greater gap enhancement. This is again in agreement with the numerical results in 5.2(b).

5.3.1 Sign change location

In order to provide some more rigour to the preceding discussion, we use equation 5.16 to derive an expression for the value of s at which we expect the appearance of negative vector components and compare this to numerical results. In order to do this we must return to the case where the catalyst contains a single coupling. As part of this derivation we assume that the problem states can be split into two sets such that the ground state overlaps with all the states in one set change sign at some value of s. Rewriting equation 5.16 in terms of these two sets, which we denote A and B, we get

$$
\begin{aligned}
\langle \hat{H}(s) \rangle = -(1-s)\Bigg(\sum_{(a,b)\in C_x^A} \langle E_0(s)|E_a\rangle \langle E_0(s)|E_b\rangle \\
+ \sum_{(a,b)\in C_x^B} \langle E_0(s)|E_a\rangle \langle E_0(s)|E_b\rangle + \sum_{(a,b)\in C_x^I} \langle E_0(s)|E_a\rangle \langle E_0(s)|E_b\rangle \Bigg) \\
+ s(1-s)J_{xx}\Bigg(\sum_{(a,b)\in C_{xx}^A} \langle E_0(s)|E_a\rangle \langle E_0(s)|E_b\rangle \\
+ \sum_{(a,b)\in C_{xx}^B} \langle E_0(s)|E_a\rangle \langle E_0(s)|E_b\rangle + \sum_{(a,b)\in C_{xx}^I} \langle E_0(s)|E_a\rangle \langle E_0(s)|E_b\rangle \Bigg) \\
+ s \sum_{a\in A\cup B} E_a |\langle E_0(s)|E_a\rangle|^2, \quad (5.17)
\end{aligned}
$$

where we have used used the superscripts A and B to denote the sets of pairs of states where both states are in A or B respectively and the superscript I to denote pairs of states for which one state is in A and the other is in B.

We then write the expression in terms of the magnitudes of the coefficients in the case where all the coefficients have the same sign and in the case where the coefficients in different sets have opposite signs. By setting the expression for the latter case as smaller than the expression for the first and re-arranging, we obtain,

$$
sJ_{\mathrm{xx}} > \frac{\sum_{(a,b)\in C_x^I} |\langle E_0(s)|E_a\rangle \langle E_0(s)|E_b\rangle|}{\sum_{(a,b)\in C_{xx}^I} |\langle E_0(s)|E_a\rangle \langle E_0(s)|E_b\rangle|}. \quad (5.18)
$$

This expression tells us, for a given s, whether or not the two sets having the opposite signs will reduce the energy of the ground state. Assuming that the sign change happens as soon as it would lower the energy, we can use this expression to determine an approximation for the value of s at which the sign change occurs. We refer to this value as s_n. To use this

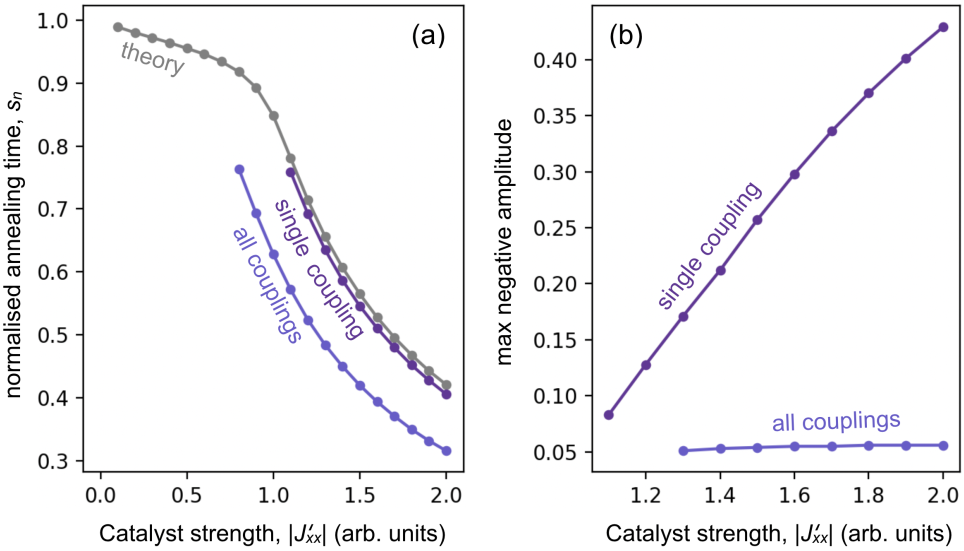

Figure 5.8: Plots showing the properties of the negative GS vector components for an anneal to a problem instance with the parameters $n_0 = 4$, $n_1 = 5$, $W_0 = 1.37$, $W_1 = 1.00$ and $J_{zz} = 37.5$. (a) shows the value of s for at which the first vector components become negative with increasing $|J'_{xx}|$. (b) shows the largest magnitude reached by the negative vector components. The numerical results when using a single XX-coupling and all-to-all XX-couplings in G_1 are plotted in the darker and lighter purple respectively. We dismiss the vector components which never drop below -0.05. The theoretical predictions for the sign change location when a single coupling is used are plotted in grey.

expression we sort the problem states by Hamming distance from some key pair of coupled states that we assume will take opposite signs. The states that are equidistant we add to the set associated with the lower energy problem state. Similar results are however obtained if they are added to the other set. We numerically sample the magnitudes of the GS vector components at different points in the anneal for a given J_{xx}. These values can then be plugged into equation 5.18 for each s to determine whether we expect the overlaps in the different sets to have the opposite sign at this point in the anneal.

Figure 5.8(a) shows how s_n varies with increasing catalyst strength for an example graph with $n_0 = 4$, $n_1 = 5$, $W_0 = 1.37$, $W_1 = 1.00$, $J_{zz} = 37.5$ and a catalyst introduced into G_1. The theoretical predictions are shown in grey and the numerical results when a single coupling is used are shown in the darker purple. In fact there are a number of negligible problem state overlaps that dip very slightly into the negative. As such, we discount any overlaps that do not drop below 0.05 when obtaining s_n from the numerical data. (This is why the numerical data does not extend to $J'_{xx} = 0$.) We see that there is good agreement between the numerical data and the theoretical prediction. We also include numerical results for the case when the catalyst contains all couplings within G_1 – plotted in the lighter purple. We see the results are qualitatively similar to the single coupling case. Once again, our way of scaling J_{xx} by the number of couplings included in $\hat{H}_c$ means that one should not read too much into lower values of s_n associated with the all-to-all catalyst.

Finally, we return very briefly to the magnitudes of the negative vector components. Figure 5.8(b) shows the largest magnitude reached by a negative problem state overlap with increasing catalyst strength. The results when $\hat{H}_c$ contains a single coupling and all-to-all couplings in G_1 are plotted in the darker and lighter purple respectively. In the single coupling case, this magnitude is clearly seen to increase with J_{xx} – as one might expect from the preceding discussion. There is also a slight increase in the all-to-all coupling case. The magnitude is however significantly suppressed compared to the single coupling case. Again, this suggests that introducing more couplings will allow the catalyst to achieve greater gap enhancement by alleviating the issues associated with the negative vector components.

5.4 Discussion

Using perturbation theory, we determined that a perturbative crossing between two states, $|E_a\rangle$ and $|E_b\rangle$ ($b > a$), could be removed by either:

- selecting XX-terms that couple $|E_a\rangle$ to other low energy states and/or low energy states that are close in Hamming distance to $|E_a\rangle$ to each other and introducing these terms with the *same* sign as the driver. (This results in a total Hamiltonian that is stoquastic.)

Or

- selecting XX-terms that couple $|E_b\rangle$ to other low energy states and/or low energy states that are close in Hamming distance to $|E_b\rangle$ to each other and introducing these terms with the *opposite* sign to the driver. (This results in a total Hamiltonian that is non-stoquastic.)

While both the stoquastic and non-stoquastic catalysts resulted in gap enhancement, our numerical investigation revealed differences between the effect of the two catalysts that our initial theoretical investigations did not account for. We found that, while the stoquastic catalysts resulted in monotonic enhancement to the spectral gap with increasing J_{xx}, the gap enhancement from the non-stoquastic catalysts reached a maximum for some J_{xx}. Increasing J_{xx} beyond this point caused the gap minimum to begin decreasing once more. Furthermore, the stoquastic catalyst produced similar gap enhancement whether just one or all possible couplings were used. Conversely, the maximum gap enhancement reached by the non-stoquastic catalyst was greater if more couplings were used.

We attributed the behaviour in the non-stoquastic setting to the effect of the catalyst on the relative signs of the instantaneous vector components. By considering the expectation value of the ground state we were able to shed some light on the relationship between the XX-terms in the catalyst and the relative signs between key vector components. In addition to explaining the aforementioned differences between the stoquastic and non-stoquastic settings, this also allowed us to deduce that these results are likely not specific to our problem graphs. Rather, any non-stoquastic XX-terms chosen as motivated in section 5.1 will result in the gap enhancement reaching some maximum value rather than increasing monotonically with J_{xx} – and this maximum can be increased by introducing additional XX-terms.

Both our analytical discussion and numerical investigations have considered perturbative crossings that form as a result of highly competitive *non*-degenerate local optima. And indeed it is perturbative crossings of this form on which our dissertation is focused. However, as noted in sections 1.3.1 and 3.1, it is also possible for perturbative crossings to form as a result of *degenerate* local optima separated by a Hamming distance of two. In this case, if we introduce XX-couplings that target the local optima as we have motivated here, there are two scenarios that could arise. The first is a similar scenario to that which occurs for the non-degenerate case – *i.e.*, one or more of the degenerate local optima could become coupled to one or more other low energy states. However, there is now also the possibility of coupling the degenerate local optima to each other. For latter case, it has already been shown that such coupling choices have the capacity to remove the perturbative crossing when introduced with the opposite sign to the driver [67, 100]. With regards to the first of these scenarios: whether or not the arguments made in section 5.1.2 apply to the degenerate case will depend on whether the dominant vector components of the relevant perturbed problem states are positive – as we showed for the non-degenerate setting. This, however, we leave as a question for future work.

With regards to understanding how useful the catalysts we have introduced will be in helping quantum annealing to find the global optimum of a problem, four important questions are: (1) How do the effects we have observed in this chapter scale to larger system sizes? (2) How sensitive are the effects of the catalyst to the specifics of the setting into which they are introduced? (3) Are XX-terms that introduce the desired couplings between problem states always guaranteed to exist in more general settings? (4) Assuming such XX-terms do exist, how easily could appropriate couplings be identified using accessible information?

The first two of these questions are what much of the rest of this thesis addresses. So far the numerical results that we have presented have been illustrative examples for which we have kept to a relatively consistent, small, system size and parameters that we expect to result in a comparatively *milder* exponential gap scaling – as discussed in Chapter 3. Chapter 7 explores both the scaling behaviour of the catalysts and their sensitivity to parameter changes in $\hat{H}_p$ with a bi-partite graph – *i.e.*, in the setting where there is a single perturbative crossing. Chapter 9 then considers how these catalysts scale with the introduction of additional local optima resulting in additional perturbative crossings.

A rigorous investigation into the second two questions lies somewhat outside of the scope of this work. We will however take a moment to discuss them here – starting with the question of identifying appropriate couplings, assuming they exist. We noted before that the stoquastic catalyst was not likely to be a practical solution to removing perturbative crossings since identifying appropriate XX-terms would likely require more knowledge of the problem GS (encoding the global optimum) than one would have access to. For the non-stoquastic catalysts however, it may be possible to identify suitable couplings using knowledge of the relevant local optima.

Crucially, if the only exponentially closing gaps in the annealing spectrum are the perturbative crossings, a polynomial-time annealing run could be used to identify the relevant local optima. Consider the setting presented in figure 3.3(c) as an example. An exponentially closing gap appears between the instantaneous GS and 1ES as a result of a perturbative crossing between the problem GS and 2ES. Tracing through the gap spectrum, it is clear that an anneal for which the majority of state transfer happens at the perturbative crossings is likely to return the problem 2ES. In theory then, a polynomial time annealing run can be used to find precisely the local optimum responsible for the algorithmic bottleneck to an adiabatic anneal.

The question then is: *can appropriate XX-terms be identified if we know the local optimum and corresponding state which we wish to target?* In the Ising formulation of the MWIS problem it is straightforward to identify low energy states that have a Hamming distance two from a given local optimum by simply selecting two up spins to flip. This will result in a state corresponding to another independent set. It is possible that similar strategies could be identified for other problems. For instance, in the max-clique problem, one may want to select a vertex that is connected to as many vertices in the current clique as possible and introduce an XX-coupling between this vertex and one already in the clique that it is *not* connected to. (This is not as clear cut a route to a low energy state as what we suggest for the MWIS problem; however it is also likely that better approaches exist.) We also note that the MWIS problem is NP-complete. Thus, if it is significantly harder to identify suitable couplings in other problem settings, the problem could always be represented as an MWIS problem to facilitate using the catalysts introduced here.

Recall, however, that it is not sufficient to couple the state corresponding to the relevant local optimum to low energy states. The couplings chosen must also couple the GS to states that are, overall, significantly higher in energy – otherwise both problem states will receive a similar perturbation from the catalyst. In our problem setting, selecting a coupling that avoids this is straightforward since none of the maximally independent sets share vertices. If, however, the local optimum we wish to target shares at least two vertices with the global optimum, then there is a chance that an XX-term selected based only on knowledge of the local optimum will also couple the GS to a low energy state. We suggest that this should not pose too much of a problem since the two maximally independent sets cannot share all their vertices and so, by introducing enough XX-terms, we are guaranteed to end up with some suitable couplings in the catalyst. Furthermore, as we saw in section 3.1, the size of the gap minimum that forms as a result of a perturbative crossing between two problem states depends inversely on the Hamming distance between them. This means that the more vertices the local and global optimum share, the larger the associated spectral gap. Therefore, the settings where the introduction of a catalyst will be most necessary are those in which the selection of an unsuitable coupling is less likely.

In considering how appropriate couplings may be unidentified in practice, we have also partially answered the question of whether or not appropriate XX-couplings are guaranteed to exist. Certainly for the MWIS graphs we are using, and the potential generalisation thereof in which local optima can share vertices, it would seem that suitable (and identifiable) couplings will always exist. We do not however presume to have covered all possible subtleties in the preceding discussion. While we leave a rigorous analysis of this topic for future work, we feel the discussion here gives some indication that catalysts of the form we examine have the potential to be constructed using accessible information.

5.5 Summary

In this chapter we have motivated two potential forms of XX-catalyst for the removal of perturbative crossings between low-energy local optima – which are the perturbative crossings of primary concern for quantum annealing. The first was a catalyst which enters with the same sign as the driver and targets the lower excited state involved in the perturbative crossing and the second was a catalyst that enters with the *opposite* sign to the driver and targets the *higher* excited state. Studying these catalysts numerically, we confirmed that both were indeed able to increase the size of the gap minimum associated with a perturbative crossing. Furthermore, examining the effect that the catalysts had on the evolution of the instantaneous eigenstates, we were able to confirm that this gap enhancement was a result of the expected changes to the energies of the perturbed problem states.

In our numerical investigations we observed some differences between the effects of the two XX-catalysts, which we were able to account for though analysis of the changing signs of the instantaneous GS vector components in the non-stoquastic case – *i.e.*, the case where the catalyst entered with the opposite sign to the driver. Specifically, we found that the appearance of negative GS vector components, which is an inherently non-stoquastic phenomenon, limited the maximum gap enhancement that the catalyst was able to achieve compared to the stoquastic case. Despite this, we believe that the non-stoquastic catalyst is the more promising route due to the fact that it may be possible to identify the necessary XX-couplings for its implementation using accessible information.

Before moving on, let us briefly comment on the extent to which the concept of stoquasticity is relevant to our discussion. When motivating our catalysts, the selection of what sign the couplings should be given in relation to the driver related to the perturbations that we wanted to introduce – rather than a need for the annealing Hamiltonian to be either stoquastic or non-stoquastic. In this sense, the fact that one from of catalyst results in a non-stoquastic $H(s)$ and the other doesn't is incidental. On the other hand, we have already observed that the behaviour of the catalyst introduced with the opposite sign to the driver (*i.e.*, the catalyst on which we will be focusing) is intrinsically tied to the non-stoquasticity that it introduces.

Overall, the concept on non-stoquasticity will not play a major role in our discussion moving forwards, becoming relevant only when we discuss certain effects relating to the signs of the GS vector components.

Chapter 6

Alternative Couplings

The focus of this thesis is on the use of XX-couplings. However, in this chapter, we will briefly examine the effect of the other possible 2-local Pauli interactions – YY, ZZ, XY, XZ and YZ. We will begin, in section 6.1, by examining numerical results for the effect on the gap minimum of the different couplings when introduced at different graph locations. Then, in section 6.2, we employ the same kind of perturbative argument made in section 5.1.2 to shed some light on our observations. We discuss what these findings imply for the use of the different coupling choices in alleviating the algorithmic bottleneck associated with a perturbative crossing in section 9.3 before providing a brief summary of the key results in section 5.5.

6.1 Numerical results

The results presented in this section are for a tri-partite graph with $n_0 = 4$, $n_1 = 5$ and $n_2 = 3$ – such that we can create a setting where a perturbative crossing forms between the problem states associated with G_0 and G_1. The state associated with G_2 does not participate in a perturbative crossing with the former two states, effectively providing three "dud" vertices in the graph. The un-normalised sub-graph weights are $W_0 = 1.70$, $W_1 = 1.35$ and $W_2 = 1.00$ and the un-normalised edge penalty is $J_{zz} = 35$. This results in a comparatively mild avoided level crossing between the instantaneous GS and 1ES – with regards to the parameter tuning discussed in chapter 3.

We present numerical results for the introduction of XX, YY, ZZ, XY, XZ and YZ couplings. By the symmetries of the problem there are six potential coupling locations. We can place a coupling between two vertices within the same sub-graph or we can couple two vertices in different sub-graphs. For each of these cases we examine the effect of introducing all possible couplings of the given type. (*i.e.*, when looking at the effect of introducing YY-couplings between G_0 and G_1, we couple every vertex in G_0 to every vertex in G_1.)

For the case where the couplings are introduced *within* a sub-graph, the catalyst Hamiltonians are given by

$$H_{\text{cat}}^{a,\alpha\beta} = \frac{J_{\alpha\beta}^{\text{cat}}}{2} \sum_{i,j \in G_a} \sigma_i^\alpha \sigma_j^\beta, \tag{6.1}$$

where a denotes the sub-graph and α and β denote the type of coupling – X, Y or Z. The result is that for the cases where both Pauli operators are the same, each coupling has a magnitude of $J_{\alpha\beta}^{\text{cat}}$. In the cases where the Pauli operators are different each edge included in the catalyst will be coupled "both ways round" with a magnitude of $J_{\alpha\beta}^{\text{cat}}/2$. That is, the coupling along each edge is given by $(J_{\alpha\beta}^{\text{cat}}/2)(\sigma_i^\alpha \sigma_j^\beta + \sigma_i^\beta \sigma_j^\alpha)$. The sum in equation 6.1 goes over all the vertices in G_a such that the catalyst contains all-to-all couplings within sub-graph G_a.

Where the couplings are between sub-graphs, the Hamiltonians can be written as

$$H_{\text{cat}}^{ab,\alpha\beta} = J_{\alpha\beta}^{\text{cat}} \sum_{i \in G_a} \sum_{j \in G_b} \sigma_i^\alpha \sigma_j^\beta, \tag{6.2}$$

where we have introduced b to denote the second sub-graph. Here, each coupling $\sigma_i^\alpha \sigma_j^\beta$ is introduced with a strength of $J_{\alpha\beta}^{\text{cat}}$. All σ_i^α operators act on qubits in sub-graph G_a while all σ_j^β operators act on qubits in sub-graph G_b. All available couplings between the sub-graphs are included in the catalyst. Note that we have made a few changes to the notation in this chapter in order to avoid confusion as our discussion progresses.

We have introduced the superscript "cat" to the catalyst strength to differentiate it from the problem ZZ-couplings when $\alpha = \beta = z$. We have also made the exchange $H_c \rightarrow H_{\text{cat}}$ to avoid confusion when we start using c to index problem states.

There are 45 unique cases in total which we can summarise as follows:

- $G_0 \leftrightarrow G_0$: XX, YY, ZZ, XY, XZ, YZ
- $G_1 \leftrightarrow G_1$: XX, YY, ZZ, XY, XZ, YZ
- $G_2 \leftrightarrow G_2$: XX, YY, ZZ, XY, XZ, YZ
- $G_0 \leftrightarrow G_1$: XX, YY, ZZ, XY, YX, XZ, ZX, YZ, ZY
- $G_0 \leftrightarrow G_2$: XX, YY, ZZ, XY, YX, XZ, ZX, YZ, ZY
- $G_1 \leftrightarrow G_2$: XX, YY, ZZ, XY, YX, XZ, ZX, YZ, ZY.

We present results for the size of the gap minimum between the ground and first excited state with varying catalyst strength, $J^{\text{cat}}_{\alpha\beta}$, for these 45 cases in figures 6.1 and 6.2. The results for $G_0 \leftrightarrow G_0$, $G_1 \leftrightarrow G_1$ and $G_2 \leftrightarrow G_2$ couplings are presented in Figures 6.1(a), (b) and (c) respectively and the results for $G_0 \leftrightarrow G_1$, $G_0 \leftrightarrow G_2$ and $G_1 \leftrightarrow G_2$ couplings are presented in Figures 6.2(a), (b) and (c). We will now make a few comments on our numerical results before going into the theory behind these observations in section 6.2.

First, we briefly note that the XX results are in line with the findings from the previous chapter. If included within G_0, the gap is enhanced when $J^{\text{cat}}_{\alpha\beta}$ is negative and, if included within G_1, the gap is enhanced for positive $J^{\text{cat}}_{\alpha\beta}$. The couplings have a limited effect when included within G_2. (Note the difference in axis scales.) We also see that the XX-couplings have a limited effect when introduced between the sub-graphs. This makes sense given the perturbative argument made in section 5.1.2, as we will go on to discuss in the following section.

Turning now to the other couplings, let us start by considering the couplings *within* the different sub-graphs. Looking at figure 6.1(a), we can

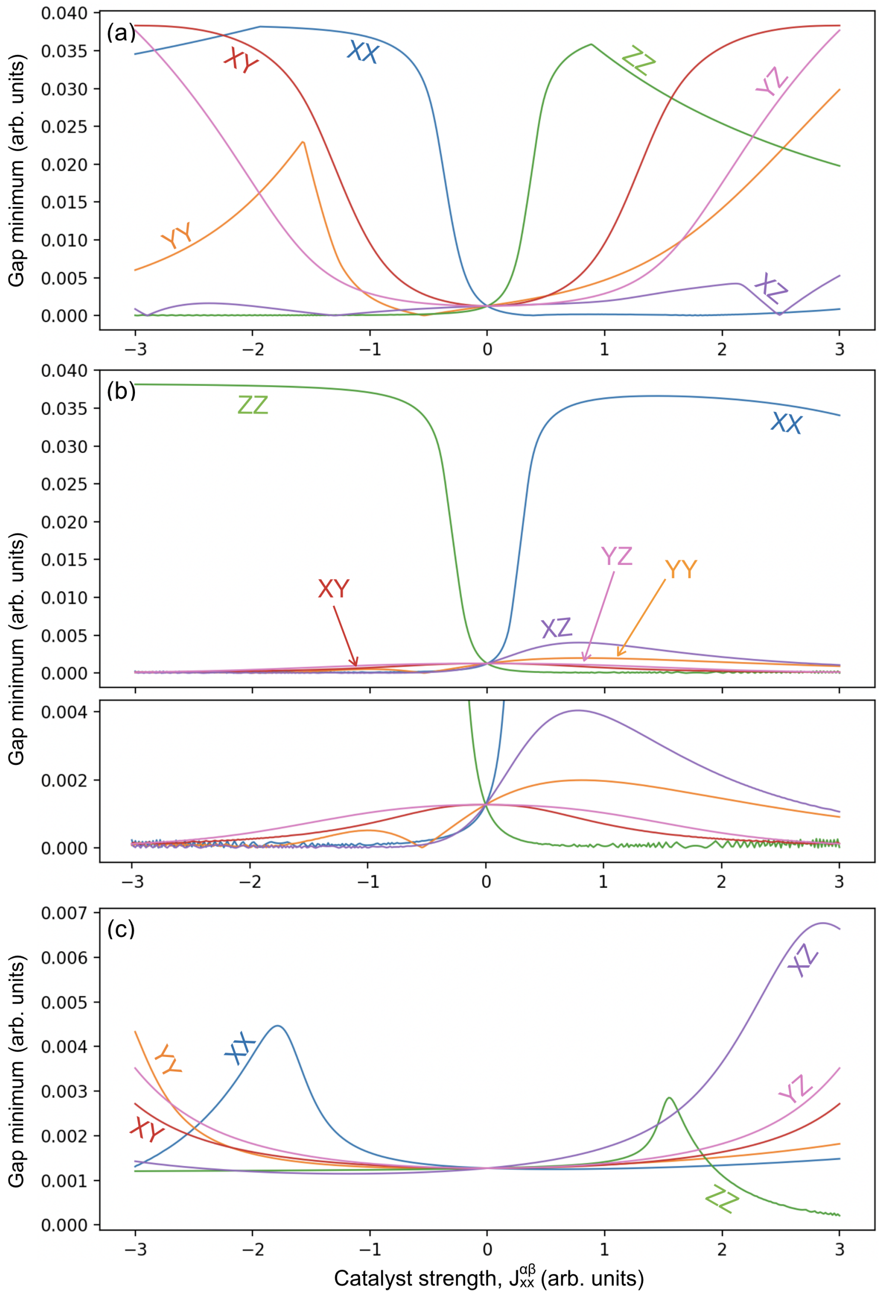

Figure 6.1: Numerical results for the size of the minimum spectral gap, $\Delta E_{01}^{\min}(s)$, with varying catalyst strength and different coupling choices. The results are for a problem graph with with $n_0 = 4$, $n_1 = 5$, $n_2 = 3$, $W_0 = 1.70$,$W_1 = 1.35$, $W_2 = 1.00$ and $J_{zz} = 35$. The plots in (a), (b) and (c) are for catalysts with all-to-all couplings within sub-graph G_0, G_1 and G_2 respectively. Results using different couplings are plotted in different colours. We include an additional panel in (b) showing a zoomed-in view of the results closer to zero gap.

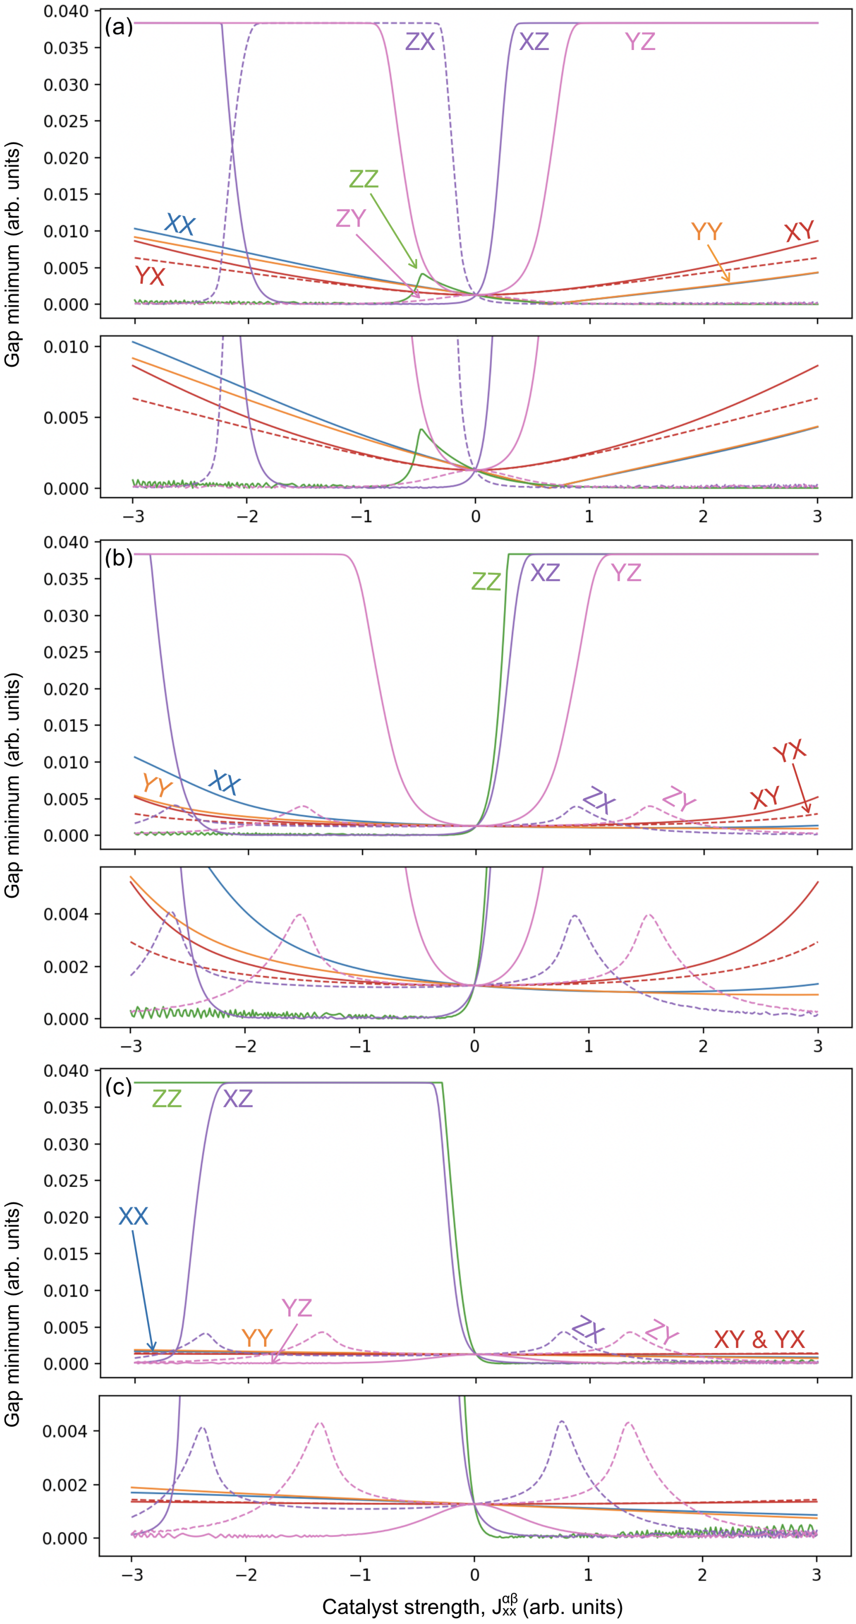

Figure 6.2: Same as in figure 6.1 but in this case the sub-plots show the results for catalysts containing all to all couplings between vertices in G_0 and G_1 (a), G_0 and G_2 (b) and G_1 and G_2 (c).

see that most of the coupling choices perform reasonably well when introduced between vertices in G_0. Intuitively, this makes sense given that we can think of introducing couplings into G_0 as effectively targeting the problem GS. However, as we will see in the second half of this chapter, this intuition is only backed up by our perturbative arguments in a couple of cases. Looking at figure 6.1(c) we see that none of the coupling choices perform well when introduced into G_2. As for the couplings in G_0, this intuitively makes sense given that G_2 is associated with a state that does not participate in a perturbative crossing. In this case, we find that this intuition is backed up by the theory. Finally, in figure 6.1(b), we see that only two of the couplings perform well when introduced into G_1. These are XX, which we have already discussed, and ZZ. We also find these results to be in line with the perturbative analysis in the following section.

Turning our attention to Figure 6.2, which shows the results for the inter-sub-graph couplings, we see that a number of the coupling choices perform well – with many reaching the maximum possible gap enhancement (ΔE_{01}) and seemingly plateauing. For all but the YZ couplings, we find that we are able to make sense of performance of the different coupling choices.

We believe that the erratic oscillations that can be seen for the very smallest values of the gap sizes are a numerical artifact. However, we have not properly investigated their origin since they have minimal impact on the results.

6.2 Perturbative analysis

To understand the observations we have made for the different coupling choices, we utilise the perturbative argument that we made in section 5.1.2. To do this, let us write down a more general version of equation 5.12 which was specifically associated the perturbative introduction of XX-couplings to the driver perturbed problem states, $|E_a(\lambda)\rangle$. We can write the perturbed energies associated with a general catalyst as

$$E_a(\lambda, \mu) = E_a(\lambda) + \mu J^{\text{cat}}_{\alpha\beta} \sum_{(i,j)\in E_{\text{cat}}} \sum_{b,c=0}^{2^n-1} f^{\alpha\beta,ij}_{bc} \langle E_a(\lambda)|E_b\rangle \langle E_c|E_a(\lambda)\rangle \, , \quad (6.3)$$

where $f^{\alpha\beta,ij}_{bc} = \langle E_b|\sigma^\alpha_i \sigma^\beta_j|E_c\rangle$ (with α and $\beta = x, y, z$) and E_{cat} denotes the edges, or couplings, included in the catalyst Hamiltonian. Note that we are now using a, b and c to denote problem states and not sub-graphs. The fact that we now have two sums is because there is now the possibility of two states being connected by more than one coupling. For instance, $|\uparrow\uparrow\downarrow \ldots\rangle$ can become coupled to $|\downarrow\uparrow\downarrow \ldots\rangle$ by an XZ-coupling acting on the first and second spins or the first and third spins.

We are of course interested in the perturbations to the two states associated with a perturbative crossing. Let us assume that one is the perturbed GS, $|E_0(\lambda)\rangle$, associated with the global optimum and the other is some other state $|E_l(\lambda)\rangle$ associated with a local optimum. Let us also refer to the set of problem states that dominate these two states as S_0 and S_l. To remove the perturbative crossing we want to introduce negative perturbations to $|E_0(\lambda)\rangle$ and/or positive perturbations to $|E_l(\lambda)\rangle$. For the XX-catalyst, we determined that we could achieve this by selecting XX-terms that coupled states within $|E_0(\lambda)\rangle$ to each other with $J^{\text{cat}}_{\alpha\beta} < 0$ or selecting XX-terms that coupled states within $|E_l(\lambda)\rangle$ to each other with $J^{\text{cat}}_{\alpha\beta} > 0$. (For our graph structure this amounted to placing XX-terms within the sub-graphs corresponding to the global and local optima respectively.)

This result hinged on two main things. First, the fact that the matrix elements, $f^{xx,ij}_{bc}$, are all either zero or unity. And secondly, the dominant vector components in $|E_0(\lambda)\rangle$ and $|E_l(\lambda)\rangle$ were all positive – which we demonstrated in Section 5.1.2. This second point of course holds regardless of the catalyst we choose since it relates to the perturbations from the driver. However, depending on the catalyst chosen, $f^{\alpha\beta,ij}_{bc}$ can take different values. To aid our discussion, let us write out equation

6.3 for $E_0(\lambda,\mu)$ and $E_l(\lambda,\mu)$ explicitly. Using the fact that, in our setting, $\langle E_a(\lambda)|E_b\rangle\langle E_c|E_a(\lambda)\rangle = \langle E_b(\lambda)|E_b\rangle\langle E_c|E_a(\lambda)\rangle$, we can write these energies as

$$E_0(\lambda,\mu) \approx E_0(\lambda) + \mu J^{\text{cat}}_{\alpha\beta} \sum_{\substack{a,b\in S_0 \\ b\geq a}} \langle E_0(\lambda)|E_a\rangle\langle E_b|E_0(\lambda)\rangle \sum_{(i,j)\in E_{\text{cat}}} (f^{\alpha\beta,ij}_{ab} + f^{\alpha\beta,ij}_{ba}) \tag{6.4}$$

and

$$E_l(\lambda,\mu) \approx E_l(\lambda) + \mu J^{\text{cat}}_{\alpha\beta} \sum_{\substack{a,b\in S_l \\ b\geq a}} \langle E_l(\lambda)|E_a\rangle\langle E_b|E_l(\lambda)\rangle \sum_{(i,j)\in E_{\text{cat}}} (f^{\alpha\beta,ij}_{ab} + f^{\alpha\beta,ij}_{ba}). \tag{6.5}$$

Note the change in indices $b \to a$ and $c \to b$ since we are no longer using a to denote the perturbed state.

To understand the effect of a particular choice of catalyst there are two main components. First, there is the question of whether or not the catalyst includes terms that result in couplings between states within one or both of the two sets, S_0 and S_l. This will dictate whether $E_0(\lambda)$ and/or $E_l(\lambda)$ receive any perturbations at all – at least according to equations 6.4 and 6.5. (Recall, from our discussion in section 3.1, that $S_{0/l}$ will contain $E_{0/l}$ as well as the problem states that are close in both energy and Hamming distance to $E_{0/l}$.) The second element is the values that the non-zero matrix elements, $f^{\alpha\beta,ij}_{ab}$ and $f^{\alpha\beta,ij}_{ba}$, associated with these pairs of coupled states take.

To help with this discussion, we include two tables that cover the key elements that we have highlighted. Table 6.1 outlines what coupling choices result in internal couplings within sets S_0 and S_l. To understand these results, recall that flipping any number of spins in one of the sub-graphs, G_a will result in states corresponding to any sub-set (including the complete set) of vertices in G_a becoming coupled to states corresponding to other sub-sets of G_a. Meanwhile states corresponding to subsets of other sub-graphs, G_b will become coupled to states corresponding to sets that include vertices from both G_a and G_b. From our understanding of the make-up of the sets, S_0 and S_l, flipping spins in G_0 will result in internal couplings between states within S_0 but not in S_l while the reverse will be true if we flip spins in G_1. Flipping spins in G_2 won't result in internal couplings between states in either S_0 or S_l.

	XX	YY	ZZ	XY	YX	XZ	ZX	YZ	ZY
$G_0 \leftrightarrow G_0$	S_0	S_0	S_0 & S_l	S_0	S_0	S_0	S_0	S_0	S_0
$G_1 \leftrightarrow G_1$	S_l	S_l	S_0 & S_l	S_l	S_l	S_l	S_l	S_l	S_l
$G_2 \leftrightarrow G_2$	-	-	S_0 & S_l	-	-	-	-	-	-
$G_0 \leftrightarrow G_1$	-	-	S_0 & S_l	-	-	S_0	S_l	S_0	S_l
$G_0 \leftrightarrow G_2$	-	-	S_0 & S_l	-	-	S_0	-	S_0	-
$G_1 \leftrightarrow G_2$	-	-	S_0 & S_l	-	-	S_l	-	S_l	-

Table 6.1: Table showing which graph coupling choices result in couplings between states within S_0 and/or S_l – as defined in section 6.2.

$b_i\ b_j$	$a_i\ a_j$	XX	YY	ZZ	XY	YX	XZ	ZX	YZ	ZY
↑↑	↑↑	0	0	1	0	0	0	0	0	0
↑↑	↑↓	0	0	0	0		0	1	0	i
↑↑	↓↑	0	0	0	0	0	1	0	i	0
↑↑	↓↓	1	-1	0	i	$-i$	0	0	0	0
↑↓	↑↑	0	0	0	0	0	0	1	0	$-i$
↑↓	↑↓	0	0	-1	0	0	0	0	0	0
↑↓	↓↑	1	1	0	$-i$	i	0	0	0	0
↑↓	↓↓	0	0	0	0	0	-1	0	$-i$	0
↓↑	↑↑	0	0	0	0	0	1	0	$-i$	0
↓↑	↑↓	1	1	0	i	$-i$	0	0	0	0
↓↑	↓↑	0	0	-1	0	0	0	0	0	0
↓↑	↓↓	0	0	0	0	0	0	-1	0	$-i$
↓↓	↑↑	1	-1	0	$-i$	i	0	0	0	0
↓↓	↑↓	0	0	0	0	0	-1	0	i	0
↓↓	↓↑	0	0	0	0	0	0	-1	0	i
↓↓	↓↓	0	0	1	0	0	0	0	0	0

Table 6.2: Table showing the values of different matrix elements $f_{ab}^{\alpha\beta,ij} = \langle E_a|\sigma_i^\alpha \sigma_j^\beta|E_b\rangle$ (α and $\beta = x,\ y,\ z$). The first and second columns show the orientations of the two spins being acted on in $|E_a\rangle$ and $|E_b\rangle$ respectively, and columns 3 to 11 give the matrix elements corresponding to the different coupling choices indicated at the top of the columns.

As such, we see in table 6.1, that XX-, YY- and XY-terms included between vertices in G_0 result in couplings within S_0 while including the couplings between vertices in G_1 results in couplings within S_1. Placing these couplings anywhere else results in no internal couplings. Whether XZ- or YZ-couplings result in internal couplings within S_0, S_l or neither depend on which sub-graph the X or Y coupling is placed in. ZZ-couplings result in every state becoming coupled to itself. Therefore, they result in internal couplings within both S_0 and S_l regardless of where they are placed.

From the arguments we have made, we expect minimal effect on the size of the gap minimum when there are no internal couplings within either S_0 or S_l. This is the case for the following catalyst choices:

- $G_2 \leftrightarrow G_2$: XX, YY, XY, XZ, YZ,
- $G_0 \leftrightarrow G_1$: XX, YY, XY, YX,
- $G_0 \leftrightarrow G_2$: XX, YY, XY, YX, ZX, ZY,
- $G_1 \leftrightarrow G_2$: XX, YY, XY, YX, ZX, ZY.

Looking at Figures 6.1 and 6.2, we see that these catalysts do indeed have a limited effect on the gap size for both positive and negative $J^{\text{cat}}_{\alpha\beta}$. There are, however, other catalyst choices that result in comparatively small gap enhancement despite the terms in the catalyst resulting in perturbations entering the sums in equations 6.4 and 6.5. To differentiate between these cases and the catalysts that do perform well, we must consider the matrix elements, $f^{\alpha\beta,ij}_{ab}$. These are presented in table 6.2.

Each of the catalysts that we examine consist of a sum of couplings of one type (XX, XY, *etc*...) acting on different pairs of qubits. Table 6.2 displays the matrix elements, $f^{\alpha\beta,ij}_{ab}$, associated with the individual couplings as follows. The first two columns give the spin orientations of the two qubits being acted on by the coupling term in the states $|E_b\rangle$ and $|E_a\rangle$ respectively – b displayed before a since $|E_b\rangle$ is the state being operated on. Columns 3 through 11 then give the value of $f^{\alpha\beta,ij}_{ab}$ for the different choices of α and β. From the way in which we have laid out the table, there are some repeat entries since *e.g.* XZ and ZX are the same coupling with the indices swapped. We have chosen to include all the columns, despite the repeats, to make it easier to refer between the tables and the numerical results in Figures 6.1 and 6.2.

As a note, for the XX-, YY- and XY-catalysts, each coupled pair of states will be associated with only one of the terms in the catalyst – the term which couples the two spins that $|E_a\rangle$ and $|E_b\rangle$ differ by. As such, there is only one term, $(f_{ab}^{\alpha\beta,ij} + f_{ab}^{\alpha\beta,ij})$, per coupled pair of states. For the XZ- and YZ-catalysts, each coupled pair of states is associated with many different terms in the catalyst – any term where either the X or the Y operator is acting on the spin which the two states differ by. For the ZZ-catalysts, each state is coupled to itself by every term in the catalyst.

We will now go through the different coupling choices, referring to tables 6.1 and 6.2, to gain some insight into our numerical observations from the previous section. Overall, we find that our perturbative analysis is able to elucidate the numerical results for most of the 45 cases (listed towards the end of section 6.1). In many cases, we see effects at larger $|J_{\alpha\beta}^{\text{cat}}|$ that are not predicted by perturbation theory – which is not surprising since the perturbations must be sufficiently small for this approach to be valid.

6.2.1 ZZ

We will start by considering the different ZZ-catalysts. These catalysts result in internal couplings within S_0 and S_l regardless of where they are placed since they couple each state to itself. $f_{aa}^{\alpha\beta,ij} = 1$ for ZZ-terms acting on spins that are aligned in $|E_a\rangle$ and $f_{aa}^{\alpha\beta,ij} = -1$ if the spins are misaligned in $|E_a\rangle$. This means, for a catalyst consisting of ZZ-terms placed within G_0, that the values of $f_{aa}^{\alpha\beta,ij}$ associated with couplings within S_0 are a mixture of ±1. This is because the ZZ-terms will be acting on states associated with different subsets on G_0. Meanwhile, the

$f_{aa}^{\alpha\beta,ij}$ values associated with the couplings within S_l will all be positive since the spins in G_0 are all down for the states in this set.

Thus, we can expect that introducing this catalyst with a positive $J_{\alpha\beta}^{\text{cat}}$ will result in a positive perturbation to $|E_l(\lambda)\rangle$. The effect that the catalyst might have on $|E_0(\lambda)\rangle$ is less clear due to the mixture of values that $f_{aa}^{\alpha\beta,ij}$ can take. One possible conclusion is that the catalyst has a limited effect on $|E_0(\lambda)\rangle$ since the perturbations can cancel each other out. This being the case, we might expect the catalyst introduced with $J_{\alpha\beta}^{\text{cat}} > 0$, to enhance the size of the gap minimum by lifting $|E_l(\lambda)\rangle$ away from $|E_0(\lambda)\rangle$. Looking at Figure 6.1(a) we see that this is indeed the case.

The situation is reversed for ZZ-couplings placed within G_1. That is, the $f_{aa}^{\alpha\beta,ij}$ values associated with the internal couplings within S_l have mixed values while those associated with the internal couplings within S_0 are all equal to $+1$. To remove the perurbative crossing we would want to lower the energy of $|E_0(\lambda)\rangle$. And indeed we see that, in this case, the gap size is enhanced when $J_{\alpha\beta}^{\text{cat}} < 0$ – see Figure 6.1(b). For ZZ-couplings placed within G_2, $f_{aa}^{\alpha\beta,ij} = 1$ for all the states in both S_0 and S_l. As such, both $|E_0(\lambda)\rangle$ and $|E_l(\lambda)\rangle$ will receive similar perturbations from the catalyst so that we might expect minimal effect on the perturbative crossing. This is reflected in the numerical results in Figure 6.1(c).

Let us now consider the couplings placed between the sub-graphs. With the ZZ-couplings placed between G_0 and G_1, the $f_{aa}^{\alpha\beta,ij}$ values associated with internal couplings within both sets are a mixture of ± 1. Correspondingly, we see in Figure 6.2(a) that the catalyst has very little effect on the gap size. When the couplings are placed between G_0 and G_2, the $f_{aa}^{\alpha\beta,ij}$ values associated with internal couplings within S_0 are again a mixture of ± 1. The matrix elements associated with the couplings within S_l, however, will all be positive. Therefore, we expect $J_{\alpha\beta}^{\text{cat}} > 0$ to enhance the gap minimum for this catalyst. Looking at Figure 6.2(b) we see that this is indeed the case. The situation is reversed when the couplings are placed between G_1 and G_2 and, once again, we see the expected results in Figure 6.2(c).

Before moving on, we would like to briefly mention that diagonal catalysts have been considered previously in the literature [68] through the idea of biasing towards the solution state. In [68], the catalysts consisted of local Z-fields such that biasing towards the solution state is a matter of orienting the Z-fields in the correct direction. With the catalysts discussed here being sums of ZZ-couplings, the situation is a little different. The catalysts could potentially be thought of as biasing towards states where the coupled spins are aligned or misaligned depending on the sign of J_{zz}^{cat}. However, we found that thinking about whether or not these ZZ-catalysts biased towards or away from the optimal solution did not help elucidate their performance. For instance, biasing towards states in which all the spins in one of the sub-graphs are aligned would bias towards both the global and local optima – which does not help us understand why a negative ZZ-catalyst performs well when introduced into G_1.

6.2.2 XZ

Now let us go through the XZ-catalysts, beginning with those where the couplings are introduced within the sub-graphs. As with the XX-catalyst, placing the XZ-terms within G_0 will result in couplings within S_0 and not S_l while the opposite is true if we place the couplings within G_1. However, unlike for the XX-catalysts, the values of $f_{ab}^{\alpha\beta,ij}$ can be ± 1 depending on the orientation of the spin acted on by σ^z. This helps explain why XZ performs much worse than XX when placed in either G_0 or G_1 even though similar couplings within S_0 and S_l are introduced – see Figures 6.1(a-b). XZ-terms within G_2 do not result in couplings within either S_0 or S_l.

Turning our attention to the cases where the XZ-terms are placed between the sub-graphs, we see that there is significant gap enhancement for any case where σ^x is acting on a spin in either G_0 or G_1 – *i.e.*, any case where couplings exist between states in either S_0 or S_l. The reason these placements work, while those internal to $G_{0/1}$ don't, is because the spin that σ^z is acting on will always be down for the states the couplings form between such that $f_{ab}^{\alpha\beta,ij}$ is always equal to -1.

Thus, for the two cases where σ^x is acting on a spin in G_0, we expect $J_{\alpha\beta}^{\text{cat}} > 0$ to lower the energy of $|E_0(\lambda)\rangle$ while leaving $|E_l(\lambda)\rangle$ unchanged – thereby removing the perturbative crossing and enhancing the size of the gap minimum. Conversely, when σ^x is acting on a spin in G_1, we expect $J_{\alpha\beta}^{\text{cat}} < 0$ to increase the size of the gap minimum since since this will increase the energy of $|E_l(\lambda)\rangle$ while leaving $|E_l(\lambda)\rangle$ unchanged. Looking at Figure 6.2, we see the numerical results are in line with these predictions.

6.2.3 YY

The YY-catalysts result in the same couplings between states as the XX-catalysts. However, for the cases where the XX-catalyst did have a significant impact on the size of the gap minimum (*i.e.*, couplings placed either within G_0 or G_1) we see that the effect of the YY-catalyst is quite different – see Figures 6.1(a) and (b). Looking at table 6.2 we see that $f_{ab}^{\alpha\beta,ij} = -1$ if the spins being acted on are aligned and 1 if the spins are anti-aligned. As with the ZZ-couplings, this means that $E_0(\lambda)$ (if the YY-terms are placed in G_0) or $E_l(\lambda)$ (if the YY-terms are placed in G_1) will receive a mixture of positive and negative perturbations form the YY-catalyst.

Looking at Figure 6.1(b), we see that indeed there is very minimal gap enhancement when the couplings are placed within G_1. In the case where they are placed in G_0 however, we do still see some gap enhancement – although the maximum enhancement here is only around 50% of that when the XX-catalyst is used, and a higher catalyst strength is needed to achieve it. It is possible that the explanation for this lies in the details of how many and which of the $f_{ab}^{\alpha\beta,ij}$ values are positive compared to negative. It is also worth nothing that the perturbative approach that we have taken to understanding the effects of these catalysts won't be capturing all of the physics. It will also become increasingly inaccurate as $|J^{\mathrm{cat}}_{\alpha\beta}|$ increases. In any case, a precise understanding of the performance of the YY-catalyst lies outside the scope of this work.

6.2.4 XY and YZ

The XY- and YZ-catalyst differ from those we have discussed so far in that the matrix elements, $f^{\alpha\beta,ij}_{ab}$, are imaginary and have the property $f^{\alpha\beta,ij}_{ab} = f^{\alpha\beta,ij}_{ba}$. Equations 6.4 and 6.5 thus imply that the catalyst does not perturb the energies at all. However, looking at Figures 6.1 and 6.2 we see that this is not reflected in the numerical results – with many of the catalysts associated with these couplings achieving the maximum possible gap enhancement.

The pattern that we observe is that, for the coupling choices that result in internal couplings within S_0 or S_l (see table 6.1), significant gap enhancement is observed for the cases where σ^y is acting on a spin in G_0. These are the XY-, (YX-) and YZ- (ZY-) $G_0 \leftrightarrow G_0$ catalysts, the YZ- $G_0 \leftrightarrow G_1$ catalyst and the YZ- $G_0 \leftrightarrow G_2$ catalyst. Something else to note is that the results from these catalysts are insensitive to the sign of $J^{\mathrm{cat}}_{\alpha\beta}$.

The fact that these catalysts perform well in many cases and are insensitive to the sign with which the couplings are introduced, suggests that they have potential for alleviating algorithmic bottlenecks in QA. However, unlike for the other catalysts that were found to significantly enhance the gap minimum, it is unclear to us where this enhancement is coming from and so it is also unclear where and how they could be implemented. These questions lie outside the scope of this work however we believe this may be an avenue worth future exploration.

6.3 Discussion

For the cases where the matrix elements, $f^{\alpha\beta,ij}_{ab}$, associated with the catalyst were real, we have observed generally good agreement between the theory and the numerical results – both with regards to what catalysts would result in gap enhancement and also what sign $J_{\alpha\beta}^{\text{cat}}$ should take to receive this gap enhancement. Firstly, any catalyst that did not result in internal couplings within either S_0 or S_l, and therefore was not expected to result in a any perturbations to either $|E_0(\lambda)\rangle$ or $|E_l(\lambda)\rangle$, had minimal effect on the gap size. For the cases where internal couplings *were* present, we were able differentiate fairly well between the cases where the catalyst *did* have a significant effect and where it did not by considering whether the perturbations that entered equations 6.4 and 6.5 were of a consistent sign or not.

For the catalysts that only resulted in couplings within one of the sets (S_0 or S_l), the catalyst was able to significantly enhance the gap minimum when the perturbations had a consistent sign and were therefore cumulative. The catalysts had a minimal effect when they resulted in a mixture of positive and negative perturbations. The one exception to this was the catalyst consisting of YY-couplings placed within G_0. However, the gap enhancement in this case was still only around 50% of the maximum gap enhancement possible while the catalysts with perturbations of consistent signs achieved, or at least came very close to, this maximum gap enhancement.

For the catalysts which resulted in internal couplings within both S_0 and S_l, which were the ZZ-catalysts, a significant effect on the gap minimum was only seen when the perturbations to exactly one of the states, $|E_0(\lambda)\rangle$ or $|E_l(\lambda)\rangle$, had consistent signs. For the case where we found that both states received consistent perturbations of the same sign there was limited gap enhancement – which we took to be a result of both states receiving a similar perturbation. There were no cases where one state received only negative perturbations while the other received only positive. For such a setting we might also expect the catalyst to have the capacity to remove a perturbative crossing.

For the catalysts with imaginary matrix elements, which were the XY- and ZY-catalysts, the theoretical predictions from equations 6.4 and 6.5 did not match up to our numerical results. While we have not been able to understand the effect of these catalysts in this brief study, we suggested that they may be of interest for future work since (a) they resulted in significant gap enhancement in multiple settings and (b) their effect was insensitive to the sign with which the couplings were introduced.

Let us now consider what the results in this section mean for how the different couplings could be used more generally – rather than just for a problem with the particular structure used here. The first take-away is that most couplings performed well when placed between vertices in G_0. However, this may not be the most useful result since this may relate to the fact that such catalysts resulted in couplings between states in S_0 – see table 6.1. In a general MWIS problem, or indeed another problem setting altogether, it is likely to be difficult (if not impossible) to identify the suitable qubits to couple together to achieve this – since this would require knowledge of the global optimum of the problem QA is trying to solve.

The fact that the ZZ-, YZ-, and in particular the XZ-catalysts performed well when placed between the sub-graphs (see Figure 6.2) may seem to indicate that these couplings could be useful for random catalysts. However, this may be misleading and more an artifact of our graph construction. In each of these coupling locations, one of the operators is still acting on one of the spins associated with either the global or local optimum participating in the perturbative crossing. Though it is likely to depend on the specific problem type and setting, we suggest that the couplings within G_2 may be a better indicator of what introducing catalyst terms at random coupling locations may be able to achieve. This is because a larger, more generic, problem may include many more vertices that are not associated with the states responsible for a perturbative crossing than our streamlined toy problems which are designed to have these perturbative crossings for small system sizes.

In the previous chapter, we suggested that an XX-catalyst which removed a perturbative crossing bottle-necking the quantum annealing algorithm could be constructed with knowledge of the state corresponding to the local optimum, $|E_l\rangle$, responsible for the perturbative crossing. This would be achieved through identifying states that are close in energy and Hamming distance to $|E_l\rangle$ (*i.e:*, the states that make up S_l) and introducing positive XX-terms that couple these states to each other. In this section, the catalysts that performed well as a result of coupling states within S_l together were the XZ-catalysts with the σ^x operator placed within G_1. The mechanism behind this is much the same as for the XX-catalysts but with the states within S_l being coupled by single spin flips – rather than two spins being flipped at once.

However, that the XZ-catalysts were able to achieve this gap enhancement was contingent on the spins acted on by the σ^z operators having the same orientations for the different states in S_l. It may well be straightforward to identify couplings with this property depending on the problem setting and if, for some hardware setting, these couplings would be easier to implement and their locations more suited to the hardware graph, these could provide an alternative to the XX-couplings we will be discussing for most of this thesis. In some sense though, the theoretical results (and their consistency with the numerical results) highlight an inherent benefit of using XX-couplings in that the signs of $f_{ab}^{\alpha\beta,ij}$ will always be consistent. As such, targeted coupling terms can be selected solely based on what problem states they couple, without the added complexity of the signs of $f_{ab}^{\alpha\beta,ij}$.

6.4 Summary

In this chapter we have examined the effect of different 2-local catalysts on the presence of a perturbative crossing. We tested the different catalysts on an example tri-partite graph for which a single perturbative crossing was present between the low energy states in the original annealing spectrum. The catalysts considered consisted of couplings of a single type (either XX, YY, ZZ, XY, XZ or YZ) and connected either all vertices within one sub-graph together or all vertices from one sub-graph to all the vertices in another. Utilising the same theoretical framework as we did in section 5.1.2, we were able to shed some light on our numerical observations.

Outside of the cases where the catalysts contained imaginary components (which were the XY and YZ catalysts), there was generally good agreement between theory and numerics, lending validity to the approach that we have taken to understanding the effects of different catalyst Hamiltonians. However, it is clear that this approach will need to be extended if we want to use it to elucidate the behaviour of catalysts with complex matrix elements.

We have drawn particular attention to an inherent benefit of using XX-couplings when it comes to designing catalyst Hamiltonians which is that the signs of the perturbations they introduce are consistent. We also highlighted a number of other 2-local catalysts that performed well with regards to enhancing the gap minimum. However, it was not as clear if and how appropriate coupling locations could be identified in general problem settings. Nevertheless, these results do suggest that many different 2-local coupling choices may have the capacity for gap enhancement when placed appropriately within a problem graph.

Chapter 7

Effect of the catalyst on a single perturbative crossing

In chapter 5, we identified a particular form of targeted XX-catalyst that had the potential to enhance the gap minimum at a perturbative crossing. We also suggested that constructing such a catalyst may be possible using accessible knowledge about the problem. This chapter examines the scaling behaviour of this form of catalyst as well as its sensitivity to the properties of the perturbative crossing it is targeting. Our aim is to understand how well this catalyst will alleviate the algorithmic bottleneck introduced by a single perturbative crossing rather than how well the catalyst will perform in more general settings where there may be many different competing local optima. As such, the graphs that we examine in this section are bi-partite to produce spectra which contain a single perturbative crossing. We explore the performance of the catalyst when multiple local optima are involved in chapter 9.

We scale our graphs as described in section 3.3 and illustrated in figure 3.1(b). Something to note is that by scaling the graphs in this way we are not growing the difference in sub-graph sizes ($n_1 = n_0 + 1$). By the arguments made in chapter 3 this implies the difference between the perturbations to the problem GS and 1ES does not proportionally increase with system size. This does not mean that this way of scaling our problem graphs is invalid with respect to producing a setting that is hard for QA – as we have seen, the resultant gap scaling is exponential with the system size. However, this point will be relevant to our discussion. Another option for the graph scaling would have been to keep the ratio of vertices contained in the two sub-graphs constant such that the difference between the two sub-graph sizes grows linearly with the system size. While do not perform a rigorous study using this graph scaling, we will reference a few preliminary results introducing the catalyst to this setting in our discussion.

The catalysts consist of XX-terms between vertices in G_1 such that the problem state corresponding to the local optimum becomes coupled to other low energy states. We first, in section 7.1, examine the performance of a catalyst containing all-to-all couplings within G_1. This results in the number of couplings included in the catalyst scaling quadratically with the system size. Section 7.2 then examines the other extreme of including just one XX-coupling in the catalyst. Finally, in section 7.3, we consider the effectiveness of the catalyst when the number of couplings included scales linearly with the problem size. In each case, we examine the effect of the catalyst on instances with parameters that result in the comparatively stronger and weaker (although still exponential) gap scaling – as discussed in chapter 3 and demonstrated in figure 3.4.

Overall, we find that if sufficient couplings are included the catalyst is able to remove the perturbative crossing completely. However, if not enough couplings are used, the effect of the catalyst is less stable – particularly when the original gap scaling is more severe. In most of these cases, we saw that the catalyst was still able to reduce the impact of the perturbative crossing but that additional effects were also observed.

7.1 All-to-all XX-couplings

The catalyst used in this section has the form

$$\hat{H}_c = J_{xx} \sum_{i,j \in G_1,\, j>i} \hat{\sigma}_i^x \hat{\sigma}_j^x. \tag{7.1}$$

We do not scale the catalyst strengths by the number of couplings used as we did for our results in chapter 5 – see equation 5.15. This normalisation of the catalyst strength was introduced to allow us to better compare the catalysts containing all-to-all couplings and single couplings in a way that allowed us to better understand the physics behind our results. Here we are simply interested in how effective the catalyst is in removing a perturbative crossing and how the coupling strengths required to achieve this scale with system size.

Note that all-to-all in this context refers to all possible couplings within G_1 – *i.e.*, all the XX-terms that create the desired couplings between problem states. The result is that the number of couplings is scaling quadratically with the problem size. For this specific setting, the fraction of possible pairs of qubits that end up XX-coupled is $(n+1)/4n$ which tends to a quarter as $n \to \infty$.

Other than some small system size effects, we find that the catalyst performs similarly regardless of the severity of the gap scaling. This is in contrast to the results using the catalyst with a single XX-coupling as we will see in the following section. Overall, we find that the all-to-all catalyst, introduced with J_{xx} within a suitable range, results in the maximum possible gap enhancement – in that the minimum gap between the instantaneous GS and 1ES is the spectral gap at the end of the anneal, ΔE_{01}. We also find that the value of J_{xx} needed to achieve this enhancement decreases with the system size.

We begin this section by examining the results for the WGS setting.

7.1.1 Weak gap scaling regime

The magnitude of the gap minimum, $\Delta E_{01}^{\mathrm{min}}$, as a funtion of problem size is plotted in figure 7.1(a). The results with and without the catalyst are plotted in purple and black respectively. We also include the energy difference between the problem GS and 1ES, ΔE_{01}, in grey to indicate the maximum gap enhancement that can be achieved – however, this is

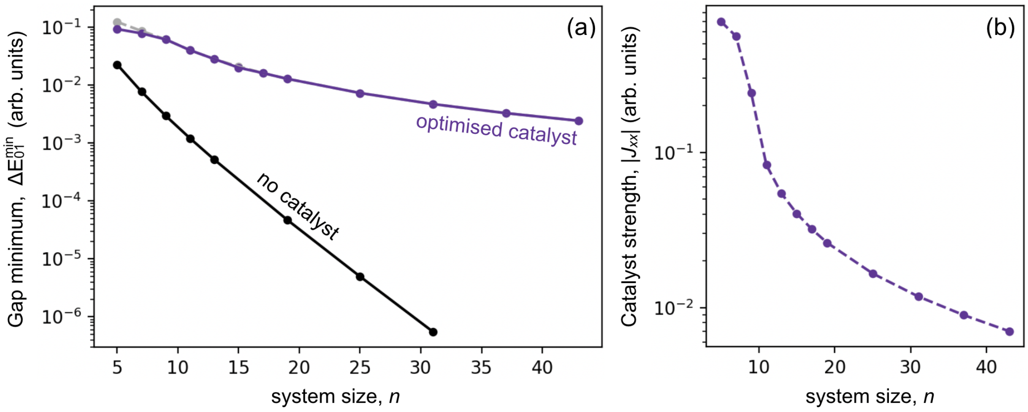

Figure 7.1: Plots showing the scaling behaviour of a catalyst on a bipartite problem with the parameters $\delta W = 0.37$ and $J_{zz} = 37.5$. The sub-graph sizes are scaled as $n_0 = (n-1)/2$ and $n_1 = (n+1)/2$. The catalyst consists of all-to-all couplings within G_1 and is introduced with the opposite sign to the driver. Plot (a) shows the size of the gap minimum in the catalyst free setting (black) and the size of the gap minimum when the optimal catalyst strength is used (purple). We also include the size of the spectral gap between the GS and 1ES at the end of the anneal (grey). Plot (b) shows the $|J_{xx}|$ values which maximise the gap minimum. The lines are a guide to the eye.

only visible for $n < 10$. For each system size, we performed a sweep over a range of catalyst strengths to find the value of J_{xx} that optimised the gap enhancement due to the catalyst. The resultant J_{xx} values are presented in figure 7.1(b). Looking at figure 7.1(a), we see that, for $n > 7$, the catalyst results in the maximum possible gap enhancement in that $\Delta E_{01}^{\min} = \Delta E_{01}$. Furthermore, in figure 7.1(b), we see a decrease in the catalyst strength required to achieve this optimal enhancement.

While the catalyst has achieved the maximum gap enhancement, the reader may observe that the resultant scaling still appears exponential – albeit with a significantly reduced exponent. That ΔE_{01} closes exponentially with the problem size is a result of states corresponding to sub-sets of G_0. The weights of the two sub-graphs, W_0 and W_1, are far enough apart that some sub-sets of G_0 have larger weights than G_1. As a result, the problem states corresponding to these sub-sets have lower energies than the state corresponding to the local optimum associated with G_1. Ending the anneal in one of these sub-sets is not a concern for the algorithm since the global optimum can be found from here through

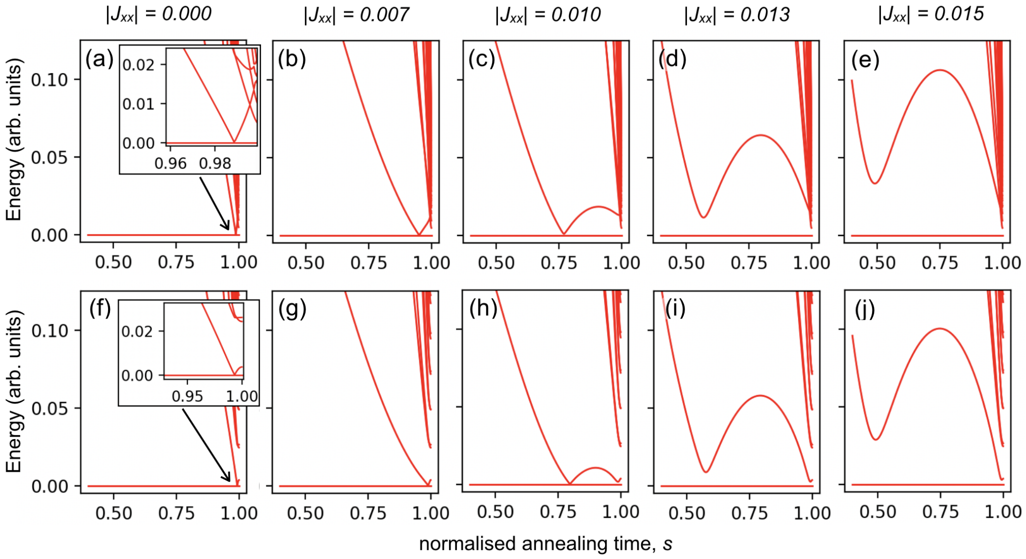

Figure 7.2: Gap spectra corresponding to two 31-spin problem instances. The top row (a-e) shows the results for a problem instance with $\delta W = 0.37$, $J_{zz} = 37.5$ and the bottom row (f-j) shows the results for an instance with $\delta W = 0.01$, $J_{zz} = 5.33$. A catalyst is introduced that consists of all couplings within G_1. The catalyst strength, $|J_{xx}|$, associated with each plot increases from left to right and its value is given above the plots.

simple gradient-descent.

The preceding discussion implies that, rather than focusing only on whether or not $\Delta E_{01}^{\min}$ is equal to ΔE_{01}, we should be considering the effect the catalyst has on the gap minimum associated with the perturbative crossing. To discuss this, let us separate out the different gap sizes we wish to refer to in our notation. We have been using s_{x} to denote the location of the gap minimum associated with the perturbative crossing. As such we can denote its size as $\Delta E_{01}(s_{\mathrm{x}})$ to differentiate it from the minimum gap in the annealing spectrum, $\Delta E_{01}^{\min}$. $\Delta E_{01}^{\min}$ can, by definition, not exceed the spectral gap at the end of the anneal, ΔE_{01}. On the other hand, we find that the catalyst can bring $\Delta E_{01}(s_{\mathrm{x}}) \gg \Delta E_{01}$.

This is illustrated in plots (a-e) in figure 7.2. These plots show the gap spectrum of the $n = 31$ instance with the WGS parameters for five catalyst strengths. An inset is included in plot (a) showing the details of the spectrum at the very end of the anneal. In this setting, $\Delta E_{01} \approx 0.006$.

Prior to the introduction of the catalyst, $\Delta E_{01}^{\min} = \Delta E_{01}(s_\mathrm{x} = 0.99)$ is on the order of 10^{-6}. As J_{xx} is increased, $\Delta E_{01}(s_\mathrm{x})$ becomes significantly larger than ΔE_{01}. We see, in plot (e), that for $J_{xx} = 0.015$ its magnitude is on the order of 10^{-2}. The magnitude of $\Delta E_{01}(s_\mathrm{x})$ can in fact be increased further than this – with the maximum enhancement being closer in magnitude to the spectral gap associated with the driver Hamiltonian, $\Delta E_{01}(s = 0) = E_1(s = 0) - E_0(s = 0)$, which closes polynomially with system size. While we have chosen to use $\Delta E_{01}(s_\mathrm{x}) = \Delta E_{01}$ as our threshold for extracting the J_{xx} scaling presented in figure 7.1(b), we note that the J_{xx} values needed to achieve higher levels of gap enhancement appear to experience a comparable drop-off in magnitude – however we have not rigorously studied this scaling.

Let us now consider whether or not the catalyst is in fact removing the perturbative crossing. We find that, while the catalyst significantly enhances $\Delta E_{01}(s_\mathrm{x})$, the catalyst never fully removes the presence of this gap minimum in the spectrum – *i.e.*, there always exists a local gap minimum in the spectrum that evolves smoothly with varying J_{xx} from the location of the perturbative crossing in the catalyst free setting. However, associating this gap minimum with an avoided level crossing becomes increasingly nonsensical as we increase J_{xx}. Firstly, the theory behind the formation of perturbative crossings (see section 3.1) implies that such an AC can only form towards the end of the anneal. Looking at figure 7.2(d) however, we can see that s_x has been pushed towards the centre of the anneal ($s_\mathrm{x} \approx 0.6$). Increasing J_{xx} further, s_x ends up approaching $s = 0$.

As for whether the gap minimum can still be associated with an AC (even if not the original perturbative crossing) – let us consider the size of the gap minimum, $\Delta E_{01}(s_\mathrm{x})$, and also the evolution of the ground state vector at s_x. While the gap still reaches a minimum, the size of this minimum is no longer small in comparison to the rest of the spectral gaps. Another signature of an AC is a sharp exchange in the make-up of the two instantaneous states between which the AC occurs. We

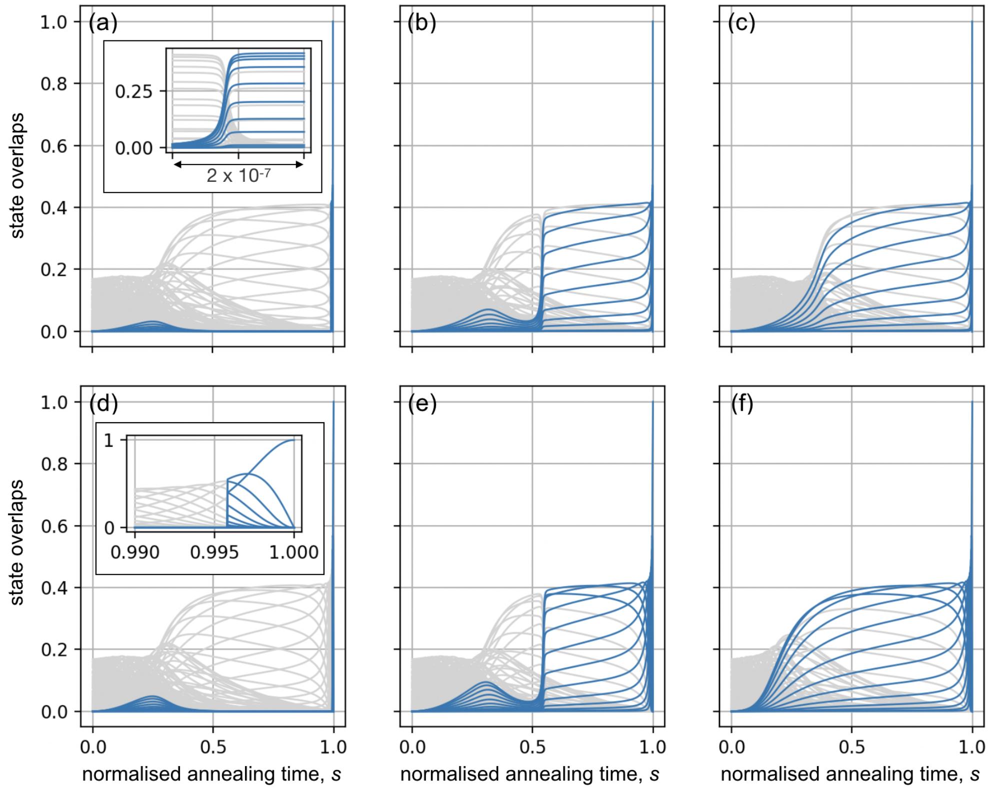

Figure 7.3: Data for the evolution of the instantaneous GS vector for two 43-spin problem instances. The plots on the top row (a-c) correspond to an instance with $\delta W = 0.37$, $J_{zz} = 37.5$ (WGS parameters) and the plots on the bottom row (d-e) to an instance with $\delta W = 0.01$, $J_{zz} = 5.33$. The left hand plots, (a) and (d), show the catalyst free results. Plots (b), (c), (e) and (f) show the results with a catalyst containing all-to-all XX-couplings within G_1. The catalysts strengths in each plot are as follows: (b) $J_{xx} = 0.00704$, (c) $J_{xx} = 0.01000$, (e) $J_{xx} = 0.00704$, (f) $J_{xx} = 0.0975$. The insets in plots (a) and (d) show a closer view of the evolution around the location of the perturbative crossing in the catalyst free cases. The inset in (a) in centred on $s = 0.9922088$.

present the smoothing out of the evolution of the instantaneous GS vector as further evidence for the removal of the AC present in the original spectrum. Plots (a-c) in figure 7.3 show the evolution of the magnitudes of the instantaneous GS vector components for the 43-spin system. We highlight, in blue, some of the key vector components present in the GS after the AC to help illustrate the change in the GS evolution. Plot (a) shows the evolution prior to the introduction of the catalyst. There a sharp exchange of vector components at $s \approx 0.9922$ over a range $\mathcal{O}(\Delta s) = 10^{-7}$. In (b) we show the evolution when the catalyst is introduced with $J_{xx} = 0.00704$ – the catalyst strength that results in $\Delta E_{01}(s_\mathrm{x}) = \Delta E_{01}$. We see that a similar exchange in vector components is still present but now it occurs towards the mid-point of the anneal and happens over a range $\mathcal{O}(\Delta s) = 10^{-2}$. The rate of this exchange is comparable to that which happens at the very end of the anneal as the mixed GS localises into the problem GS. Finally, plot (c) shows the results when $J_{xx} = 0.01000$. Here we see a gradual increase in the presence of the highlighted vector components over the course of the first half of the anneal.

Finally, we comment on the scaling with system size of the value of the catalyst strengths in figure 7.1(b). That the catalyst strength required for the optimal gap enhancement decreases with system size is perhaps unsurprising given that the number of couplings included in the catalyst increases quadratically. In the perturbative argument we used to motivate the catalyst (section 5.1.2) we saw that the XX-terms included had a cumulative effect – such that we could decrease the catalyst strength if we include more couplings. Recall also that the aim of the catalyst is to reduce the negative perturbations to the problem state corresponding to the local optimum such that it does not cross the perturbed ground state. In the introduction to this chapter, we suggested that the difference between the perturbations to these states does not proportionally increase with system size for the graph scaling used in this section. It could be argued therefore that, in some sense, the demand on the cat-

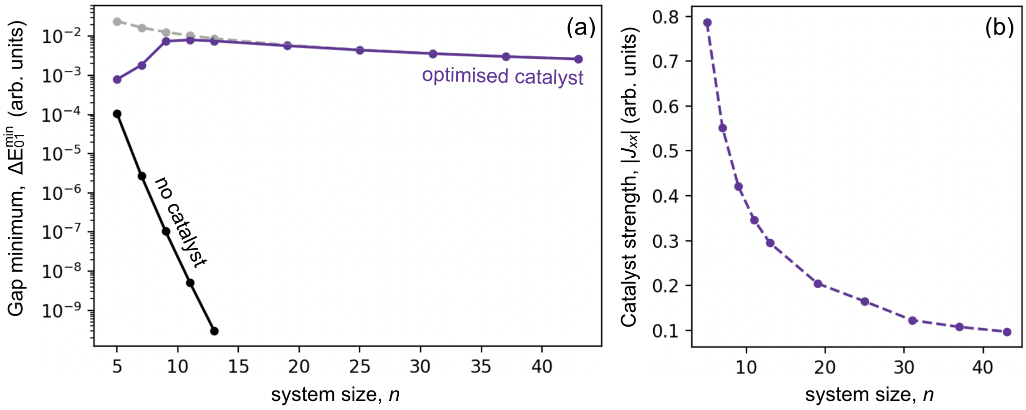

Figure 7.4: Plots showing the scaling behaviour of a catalyst on a bipartite problem with the parameters $\delta W = 0.01$ and $J_{zz} = 5.33$. The sub-graph sizes are scaled as $n_0 = (n-1)/2$ and $n_1 = (n+1)/2$. The catalyst consists of all-to-all couplings within G_1 and is introduced with the opposite sign to the driver. Plot (a) shows the size of the gap minimum in the catalyst free setting (black) and the size of the gap minimum when the optimal catalyst strength is used (purple). We also include the size of the spectral gap between the GS and 1ES at the end of the anneal (grey). Plot (b) shows the $|J_{xx}|$ values which maximise the gap minimum.

alyst is not increasing with the system size. We will comment on this further in our discussion at the end of this chapter.

We now consider the effect of the all-to-all catalyst Hamiltonian when applied to the setting where the parameters result in the comparatively stronger exponential gap scaling.

7.1.2 Strong gap scaling regime

As in the preceding section, we begin by presenting results for the gap scaling with and without the catalyst in figure 7.4(a). The catalyst strengths used are plotted in 7.4(b). We see very similar results to those for the WGS regime with the catalyst resulting in the maximum achievable gap enhancement for all but the very smallest system sizes and the required values of J_{xx} decreasing with system size.

One difference is that the catalyst strengths in figure 7.4(b) are around one order of magnitude larger than those in the WGS setting. This is potentially an unimportant detail. It is however still worth pointing out that this mostly relates to where the state corresponding to the local optimum sits in the problem spectrum. Plots (f-j) in figure 7.2 show example gap spectra for the SGS setting with five different catalyst strengths. These are the equivalent plots to those for the WGS setting in (a-e). Looking at the inset in (f) we see that the state corresponding to the local optimum that crosses the ground state is the problem 1ES. This is not the case for the WGS setting. There, as seen in figure 7.2, this is the 3ES. This is a result of the weight separation between the two sub-graphs being large enough that sub-sets of G_0 have a larger weight than G_1. In the SGS plots (f-j) we see that, as the gap is enhanced by the catalyst, a "lip" remains at the end of the anneal which becomes the smallest gap in the spectrum. This "lip" is also present in the WGS case for the 3ES. Because it takes much larger catalyst strengths to fully remove this, the extracted J_{xx} values end up much larger for the SGS setting. However, comparing the top and bottom rows of figure 7.2 we can see that enhancement to $\Delta E_{01}(s_{\mathrm{x}})$ is similar in the WGS and SGS cases for the same catalyst strengths.

Let us turn to the crucial question we want to explore in this section: do our results suggest that the catalyst is able to remove the perturbative crossing present in this system when the initial gap scaling is more severe? Plots (d-f) in figure 7.3 show how the evolution of the GS vector changes with increasing J_{xx} for the 43-spin system – these are the same sort of results that we present for the WGS setting in the plots above. The results for the catalyst-free case, presented in plot (d), show a sharp exchange at $s \approx 0.996$. This exchange happens over too short a timescale for us to observe; however it is certainly over a range of s less than 10^{-11}. Plot (f) shows the results when the catalyst is introduced with the strength at which $\Delta E_{01}^{\mathrm{min}} = \Delta E_{01}$, which in this case is $J_{xx} = 0.00975$. We see a similar smoothing out of the evolution of the highlighted vector components as we did for the WGS setting for a comparable catalyst strength, indicative of the removal of the perturbative crossing. We also include results for an intermediate catalyst strength which, for ease of comparison with the WGS setting, is chosen to be $J_{xx} = 0.00704$. We see similar results to the WGS setting with the exchange in the GS vector components still present, but occurring now over a range on the order of 10^{-2}.

Much like in the WGS setting, $\Delta E_{01}(s_{\mathrm{x}})$ can be enhanced to the point where it is closer to the spectral gap between the ground and first excited state of the driver Hamiltonian – such that it is no longer small in comparison to the other spectral gaps. We do note that, as a result of how competitive we have made the local optimum in this setting, the smallest gap in the spectrum, $\Delta E_{01}^{\mathrm{min}} = \Delta E_{01}$, *is* comparatively small. This however becomes less and less the case as we increase the problem size; for $n = 65$ we find that ΔE_{01} is on the order of other spectral gaps.

Finally, we highlight one additional difference that we observe between this setting and the WGS setting from the previous section which is the effect of the catalyst on the very smallest systems. For $n < 9$ we observe a closing of the gap minimum associated with the perturbative crossing as well as the appearance, closing and re-opening of an additional gap minimum in the spectrum. These closing gaps will be discussed further in section 7.2.2 where we find that this phenomenon persists to larger system sizes when only one XX-term is included in the catalyst.

With regards to the use of the all-to-all XX-catalyst examined in this section, we have found that the proposed catalyst appears to remove the perturbative crossing present in our toy system. We also found this success to be stable to changes in the severity of the associated exponential gap scaling and that the required catalyst strength decreased with the system size. We will discuss what general conclusions can be drawn from these findings at the end of this chapter.

7.2 Single XX-coupling

So far we have examined the effectiveness of the proposed catalyst for removing a perturbative crossing when all appropriate XX-couplings were included. This amounted to a quadratic scaling in the number of XX-terms with system size with the fraction of possible pairs of qubits that end up XX-coupled tending to a quarter. To understand better if such catalysts will be useful in practice, it will be important to consider their performance when the number of couplings included in the catalyst is varied. We have already discussed, in section 5.4, the problem of identifying suitable couplings – which is one thing that may limit the number of couplings that could be included in the catalyst. For a practical implementation, another consideration would be the connectivity of the hardware.

In this section, we examine the impact of a catalyst that contains just one XX-coupling such that it has the form

$$\hat{H}_c = J_{xx}\hat{\sigma}_i^x\hat{\sigma}_j^x. \tag{7.2}$$

The indices, i and j, correspond to vertices in G_1. Considering this extreme case, where only one XX-term term is introduced, provides some insight into how the catalyst may perform when fewer couplings are present. We also observe some interesting effects as a result of this single coupling catalyst which we believe are worth exploring.

Overall, we find the effects of this catalyst to be less predictable than what we saw in the previous section. These effects include some gap enhancement, but also the appearance of additional gap minima in the spectrum. Interestingly, we find that the catalyst continues to have a significant effect on the annealing spectrum as the system size is increased despite the fact that we are increasing neither the number of XX-terms nor the catalyst strength, J_{xx}.

It is less clear if the effects we observe in this section will be beneficial with regards to helping the system reach a high overlap with the final GS for shorter annealing times. Certainly, the resultant spectra suggest that the catalyst does not significantly reduce the annealing time needed for adiabatic evolution. However, it *is* possible that the spectra may facilitate

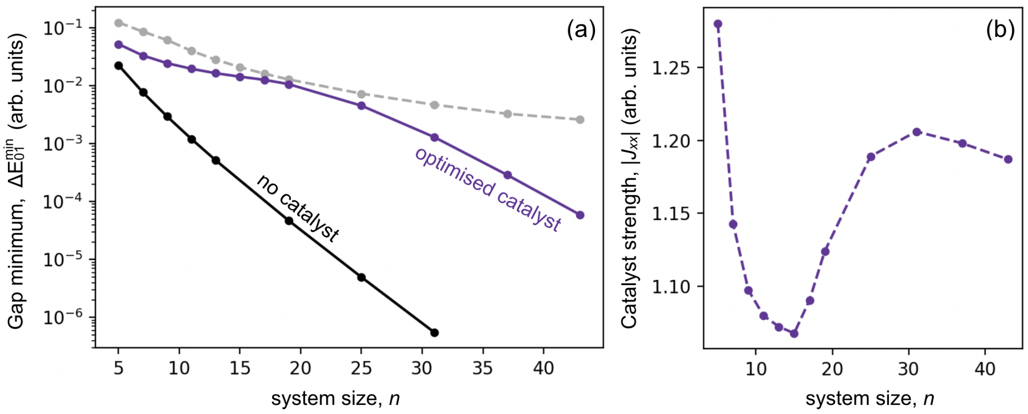

Figure 7.5: Plots showing the scaling behaviour of a catalyst on a bipartite problem with the parameters $\delta W = 0.37$ and $J_{zz} = 37.5$. The sub-graph sizes are scaled as $n_0 = (n-1)/2$ and $n_1 = (n+1)/2$. The catalyst consists of a single coupling within G_1 and is introduced with the opposite sign to the driver. Plot (a) shows the size of the gap minimum in the catalyst free setting (black) and the size of the gap minimum when the optimal catalyst strength is used (purple). We also include the size of the spectral gap between the GS and 1ES at the end of the anneal (grey). Plot (b) shows the $|J_{xx}|$ value which maximises the gap minimum.

a faster route to the GS through a *diabatic* anneal. We will point out where and how we believe this may be the case throughout this section; the majority of our discussion on the potential for utilising these spectra for DQA will however occur in the final section of this chapter.

As before, we examine the effect of the catalyst in the WGS and SGS settings in turn. We will also compare the response of the instantaneous GS evolution to the introduction of the catalyst to better understand the differing effects that we see between the two regimes.

7.2.1 Weak gap scaling regime

We begin by looking at the same kind of scaling data that we presented for the all-to-all catalyst. Figure 7.5(a) shows the catalyst free gap scaling in black and the gap scaling with the catalyst in purple. The associated J_{xx} values are plotted in figure 7.5(b). As before, the spectral gaps between the problem ground and first excited states are included in figure 7.5(a) in grey, indicating the maximum possible gap enhancement. For

$n < 20$ the gap scaling data follows a similar pattern to what we saw in the previous section for the very smallest system sizes – with the gap enhancement from the catalyst tending towards the maximum possible enhancement. However, for $n > 20$, we see that the gap minimum begins to close more rapidly.

That there is a drop-off in the catalyst's capacity to suppress the exponential gap scaling is perhaps unsurprising given that the catalyst contribution to the Hamiltonian is proportionally decreasing with the system size. Interestingly though, looking at the annealing spectra themselves, we find that the catalyst continues to significantly impact the spectrum for catalyst strengths on the order of unity even for the very largest systems that we examined ($n = 71$). However, what we see for $n > 20$, is the appearance of additional gap minima in the spectrum. In particular, we see the formation of a double gap-minimum reminiscent of what was found in [100]. Let us discuss the manifestation of this double minimum in more detail.

Figure 7.6 shows example spectra at three different system sizes. Plots (a-d) are for the 11-spin system, (e-j) correspond to the 21-spin system, and (k-p) correspond to the 31-spin system. Each set of plots show the annealing spectra for that system size at different values of J_{xx} to illustrate the evolution of the spectrum with catalyst strength. Looking at the 11-spin system, we see that the behaviour reflects what we observed in section 5.2.1. That is, we see an initial enhancement to the gap minimum before it begins to close again. As for the very smallest systems when the all-to-all catalyst was used, this enhancement does not bring $\Delta E_{01}^{\min} = \Delta E_{01}$.

Next, let us look at plots (k-p) corresponding to the 31-spin system. The coloured arrows included on these plots indicate which gap minima evolve smoothly into each other. As J_{xx} is increased from 0 to 1.45 (and indeed beyond this point) we find that the gap minimum associated with the perturbative crossing (indicated with a purple arrow) is enhanced by the catalyst as desired. Furthermore, this enhancement does

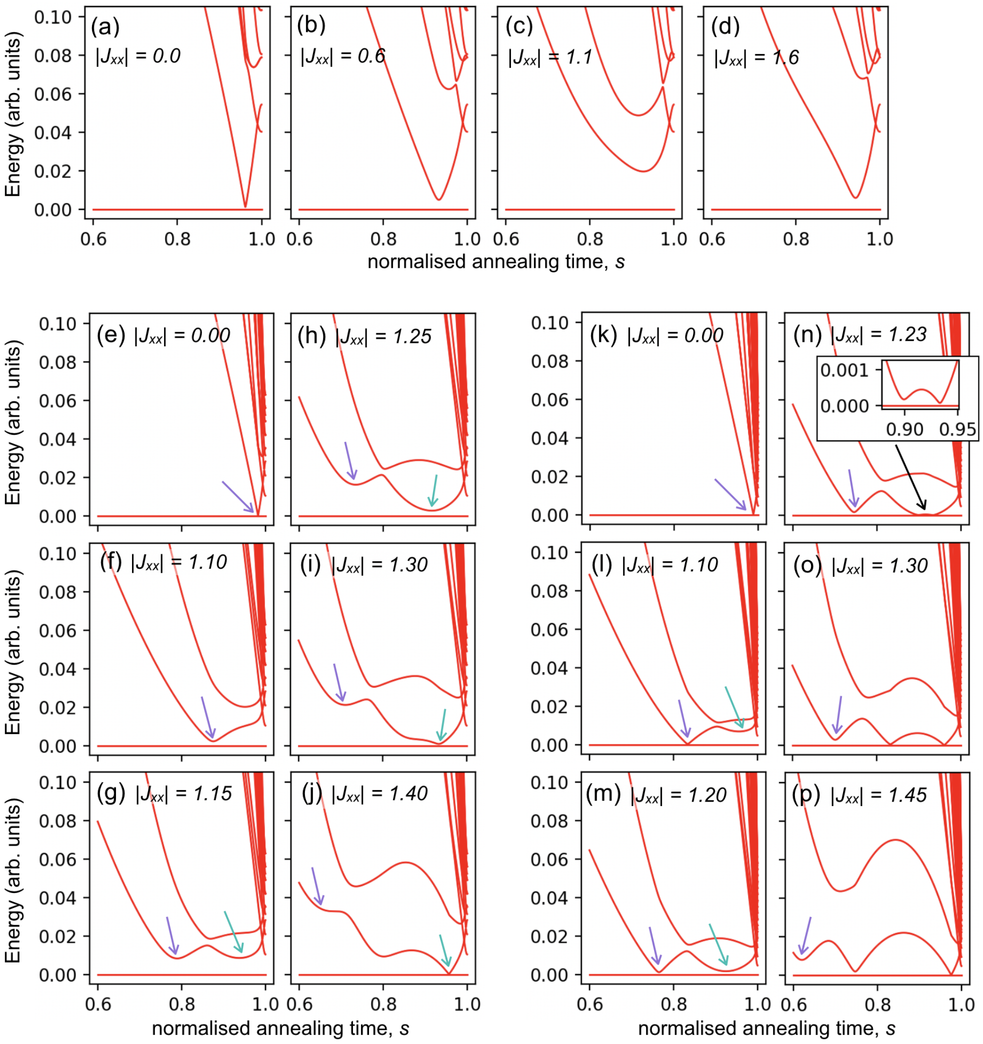

Figure 7.6: Gap spectra corresponding to bipartite problem instances with the parameters $\delta W = 0.37$ and $J_{zz} = 37.5$. The plots are grouped into three sets. (a-d), (e-j) and (k-p) show the results for an 11, 21 and 31-vertex instance respectively. A catalyst is introduced that consists of a single coupling within G_1. The catalyst strength, $|J_{xx}|$, associated with each plot is given on the plots. Where it may be unclear, coloured arrows are used to indicate gap minima that continuously evolve into one another as the catalyst strength is increased.

bring $\Delta E_{01}^{\mathrm{min}} > \Delta E_{01}$ – see plot (p). However, from around $J_{xx} = 1.10$ we observe the appearance of a new gap minimum in the spectrum (indicated with the turquoise arrow). Looking at plots (l-n), we see that as J_{xx} is increased further, the size of this new gap minimum decreases until eventually, at around $J_{xx} = 1.23$, the instantaneous 1ES effectively "dips below" the GS creating a pair of avoided level crossings. Increasing J_{xx} beyond this point, the two new ACs move further apart – see plots (n-p). Eventually the AC at the lower s value begins to lift while the gap minimum associated with the AC at the higher s value shrinks. This behaviour is precisely what was is observed in [100]. We also note that, considering the couplings between problem states introduced by $\hat{H}_c$ and the evolution and signs of the GS vector components, this behaviour seems to stem from the same phenomenon. As such, we will not linger on the theory behind the appearance of the two additional ACs and instead refer the reader to [100].

Plots (e-j) in figure 7.6, corresponding to the 21-spin system, show the intermediate behaviour. As for the 31-spin example we see that the gap minimum associated with the perturbative crossing can be significantly enhanced by the catalyst and we also see the appearance of the new gap minimum. In this case however, only one small gap minimum with the ground state is ever present.

Let us turn to the question of whether the single-coupling catalyst is removing the perturbative crossing present in the original spectrum – even if additional effects are present. We find that, as for the all-to-all catalyst, $\Delta E_{01}(s_{\mathrm{x}})$ can be enhanced past the value of the spectral gap between the GS and 1ES at the end of the anneal. This goes hand in hand with a similar "smoothing out" of the GS vector components to what was observed in the previous section. In this sense it would seem that the single coupling catalyst is still successfully removing the perturbative crossing. However, unlike when all-to-all couplings within G_1 were used, we find that the value of J_{xx} needed to reach the same level of enhancement increases with the system size – this can be seen comparing the plots in figure 7.6. However, we cannot fully comment on this scaling behaviour as we have only extracted the J_{xx} values associated with the smallest spectral gap in the anneal, $\Delta E_{01}^{\text{min}}$, reaching its maximum value – rather than $\Delta E_{01}(s_{\text{x}})$.

The focus of this chapter is primarily to understand the effect that the catalyst has on the annealing gap spectrum and the evolution of the instantaneous GS vector – rather than what this might mean for the dynamics. For context however, and to guide us to the important features of our findings, we have included some discussion throughout on what the results imply for the dynamics of an anneal. As such, let us briefly consider what conclusions can be drawn regarding the effect of the catalyst on the anneal time needed to end the anneal with a high GS fidelity. While we do not see the same gap enhancement as we did for the all-to-all catalyst, the spectra that the catalyst produces appear amenable to diabatic annealing – as proposed in [100]. Consider the spectrum presented in figure 7.6(p). If the anneal was run with a suitable annealing time, we might expect the system to transition into the 1ES at the first of the new gap minima ($s \approx 0.75$) and then back into the GS at the second ($s \approx 0.98$). The possibility of this for larger systems will depend on how the catalyst strength that results in the double AC scales in relation to the catalyst strength that results in sufficient enhancement to $\Delta E_{01}(s_{\text{x}})$.

7.2.2 Strong gap scaling regime

Once again, we start by considering how the gap scaling is affected by the catalyst – the results for which are presented in figure 7.7(a). We see that, while the catalyst does enhance the size of the gap minimum for each of the system sizes plotted, the exponential scaling associated with the closing of the gap minimum is largely unchanged. As with the WGS setting, however, the catalyst is still having a significant impact on the annealing gap spectra and GS evolution – even as the system size is increased. However, this impact is quite different to what we observed in the WGS setting.

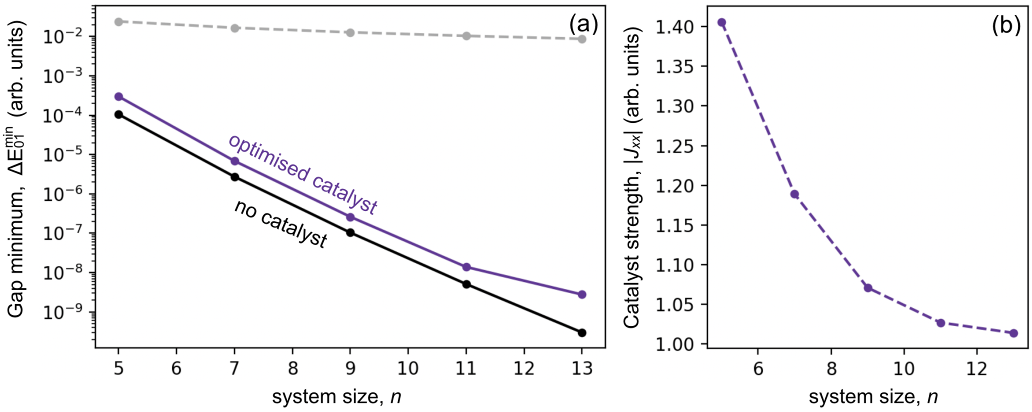

Figure 7.7: Plots showing the scaling behaviour of a catalyst on a bipartite problem with the parameters $\delta W = 0.01$ and $J_{zz} = 5.33$. The sub-graph sizes are scaled as $n_0 = (n-1)/2$ and $n_1 = (n+1)/2$. The catalyst consists of a single coupling within G_1 and is introduced with the opposite sign to the driver. Plot (a) shows the size of the gap minimum in the catalyst free setting (black) and the size of the gap minimum when the optimal catalyst strength is used (purple). We also include the size of the spectral gap between the GS and 1ES at the end of the anneal (grey). Plot (b) shows the $|J_{xx}|$ value which maximises the gap minimum.

One key difference is that the catalyst does not appear to remove the perturbative crossing in that we do not see a smoothing out of the exchange in the GS vector components, with s_{x} being pushed earlier in the anneal, and we do not see a significant enhancement to $\Delta E_{01}(s_{\mathrm{x}})$. Furthermore, while we do observe the appearance of new gap minima in the spectrum, they appear to have a different origin to those from the previous section. We also observed these closing gaps for the smallest SGS systems when the all-to-all catalyst was used. Here, we observe them for the higher system sizes as well. We will now discuss these closing gaps in more detail.

Much like in all the previous settings, the gap minimum associated with the perturbative crossing is found to reach some maximum value with increasing J_{xx} before eventually decreasing again – although in this case this maximum does not differ substantially from its starting value. In addition, we observe a closing of this gap minimum for a value of J_{xx} less than that for which it reaches its maximum. We appear to be able to make this gap arbitrarily small through increasingly fine tuning of the catalyst strength. This suppression to $\Delta E_{01}(s_{\mathrm{x}})$ was only observed here and for the very smallest instances of the SGS setting in other contexts. We also observe the appearance of a new gap minimum in the spectrum for some value of J_{xx} larger than that for which $\Delta E_{01}(s_{\mathrm{x}})$ was observed to go to zero. As with $\Delta E_{01}(s_{\mathrm{x}})$, we find that we can also bring this new gap minimum arbitrarily close to zero. We will refer to the location of this new gap minimum as $s_{\mathrm{x}'}$ such that its size is denoted with $\Delta E_{01}(s_{\mathrm{x}'})$.

We note that there is some difficulty in confirming whether or not these gap minima can truly be made arbitrarily small. Certainly, the size of the gap minima can be suppressed to zero to the numerical precision that we are able to achieve. Potentially more compelling evidence is that we do not observe any curvature in the variation of the gap minimum with J_{xx} which would indicate a minimum being reached. Towards the end of this section we will offer one further observation that suggests that these gap minima do indeed go to zero. First though, let us return to the effect of the catalyst on the annealing gap spectrum.

In figure 7.8 we present annealing spectra that illustrate the manifestation of this additional gap minimum in the spectrum for different system sizes. Plots (a-c), (d-f) and (g-i) correspond to the 15-, 25- and 35-spin systems respectively. The plots in the middle row show the spectrum when the catalyst strength is selected such that $\Delta E_{01}(s_{\mathrm{x}'}) = 0$ as far as our numerical precision allows. The top and bottom row plots show the results when J_{xx} is 0.95 and 1.05 times this value respectively. For the smaller systems, we find that the gap minimum which goes to zero is the only new gap minimum that appears – as illustrated in plots (a-c). As we increase our system size however, we observe the appearance of a further gap minimum in the spectrum. This further gap minimum does not share the behaviour of the first two in that we cannot arbitrarily suppress its magnitude through our selection of J_{xx}.

We will not focus too much on the specifics of how the gap spectrum evolves under changing J_{xx} as we do not feel that conclusions can be drawn from these details that will help us understand how this catalyst

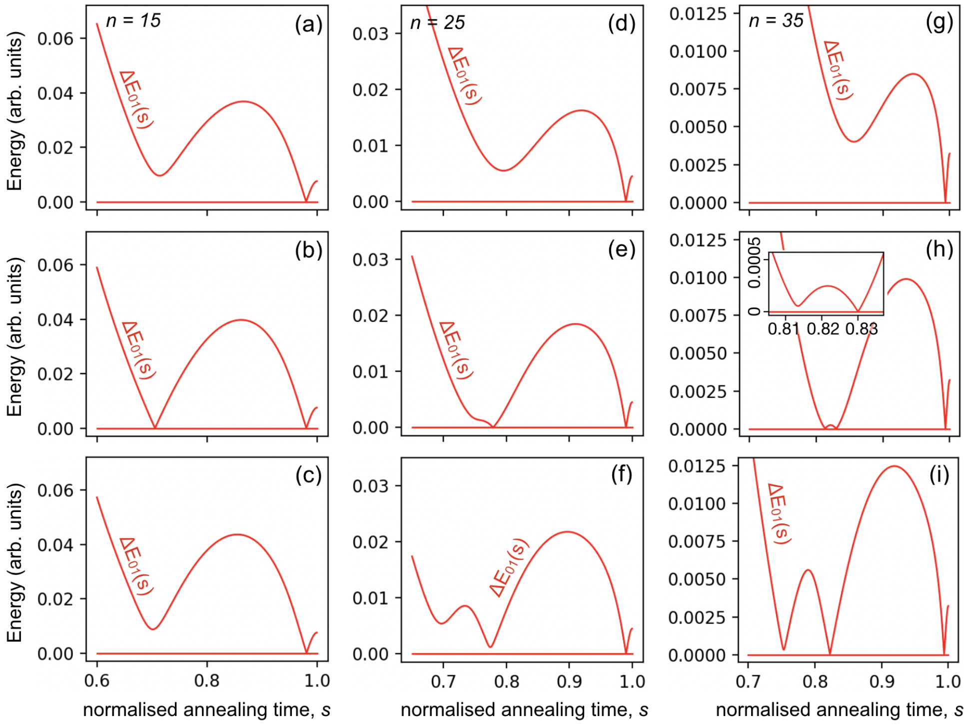

Figure 7.8: Gap spectra corresponding to bipartite problem instances with the parameters $\delta W = 0.01$ and $J_{zz} = 5.33$ with the presence of a catalyst consisting of one XX-coupling between two vertices in G_1. Plots (a-c), (d-f) and (g-i) correspond to 11, 21 and 31-vertex instance respectively. The middle plots in each column show the results when J_{xx} is chosen to minimise the gap minimum created by the catalyst. The plots on the top and bottom rows show the results when the the catalyst strength is chosen to be 0.95 and 1.05 of this value respectively.

may impact annealing spectra in more general settings. This is in part due to the variation in the response of the spectrum to the catalyst with varying system size and because, unlike for the effects from the catalyst in the previous sections, we do not have a clear picture of the physics behind the observations we have made here. However, we feel the closing gap minima are an interesting phenomenon and so we will give some attention to their formation. We reserve most of this discussion for the following section where we compare the effect of the single coupling catalyst in the two regimes. However, let us briefly highlight two key points.

First, we find that these closing gaps occur for the value of J_{xx} for which

s_n coincides with the location of the respective gap minimum – where s_n is the point in the anneal at which some of the vector components become negative. Second, as J_{xx} is increased past this value, we observe what appears to be a discontinuous change in which vector components cross zero. This discontinuous change suggests to us some kind of phase transition, implying that there may indeed be a vanishing gap at this point.

7.2.3 Comparing the two regimes

Unlike for the all-to-all catalyst examined in section 7.1, we have found the effect of the single-coupling catalyst to be strongly dependent on the severity of the exponential scaling associated with the perturbative crossing present in the original annealing spectrum. In the WGS setting, the single-coupling catalyst resulted in a similar lifting of the perturbative crossing to what we saw for the all-to-all catalyst. That is, we observed a smoothing out of the sharp exchange in the instantaneous GS vector and an enhancement to the gap minimum at s_x. We also saw the formation of two *new* ACs as the 1ES effectively dipped below the GS for a portion of the anneal. We attributed this behaviour to the same mechanism discussed in [100]. In the SGS setting, we saw limited enhancement to $\Delta E_{01}(s_\mathrm{x})$ from the single-coupling catalyst and also that this gap minimum could be brought arbitrarily close to zero by increasingly fine tuning of J_{xx}. Furthermore, we observed the appearance of a an additional gap minimum in the spectrum which could also be arbitrarily suppressed.

The physics behind these vanishing gaps is not entirely clear to us. However, we are able to make a few notes on some of the features associated with their occurrence as well as the differences that we observe between the SGS setting, where they do occur, and the WGS setting, where they do not. To facilitate this, we include data for a 5-spin system with the WGS and SGS parameters in figures 7.9 and 7.10 respectively. The 5-spin system is a suitable choice to aid our discussion since the response of the smaller systems to the catalyst are the least involved. For the WGS we do not see the appearance of the double AC and only one new gap minimum appears in the SGS setting. However, the key features we wish to highlight are already present.

Figure 7.9(a) shows the size of the gap minimum, $\Delta E_{01}(s_{\mathrm{x}})$, and its location, s_{x}, with varying J_{xx} in solid and dashed purple lines respectively. We also include the s value at which the first negative GS vector components are seen, s_n, with a dashed grey line. Figures (b-f) show the evolution of the instantaneous GS for different catalyst strengths – with the vector components corresponding to the GS and 1ES highlighted in blue and orange respectively. The location of the gap minimum, s_{x}, is indicated with a dashed purple line on each plot. We indicate the J_{xx} values that these five plots correspond to on figure 7.9(a) with dotted grey lines.

Figure 7.10(a) shows the same data as figure 7.9(a) but for the SGS setting. Below this, in figure 7.9(b), we plot the size, $\Delta E_{01}(s_{\mathrm{x}'})$, and location, $s_{\mathrm{x}'}$, of the the new gap minimum that forms in this setting. Note that plots (a) and (b) in figure 7.9 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. The location of the sign change, s_n, is therefore the same in each plot. As for the WGS setting, we include plots showing the evolution of the instantaneous GS for different catalyst strengths which are marked on plots (a) and (b) with dashed grey lines.

We first turn our attention to the SGS results in figure 7.10. Looking at plots (a) and (b) we observe the aforementioned closing of the two gap minima for $J_{xx} \approx 0.35$ and 1.90 respectively. As mentioned in the preceding section, the closing gaps are seen to occur for the J_{xx} values for which s_n coincides with the location of the gap minima – s_{x} and $s_{\mathrm{x}'}$. Associating the closing gaps with their locations coinciding with s_n does nothing however to account for why we see in phenomenon here and not in the WGS setting. Looking at the WGS results in figure 7.9(a), we see

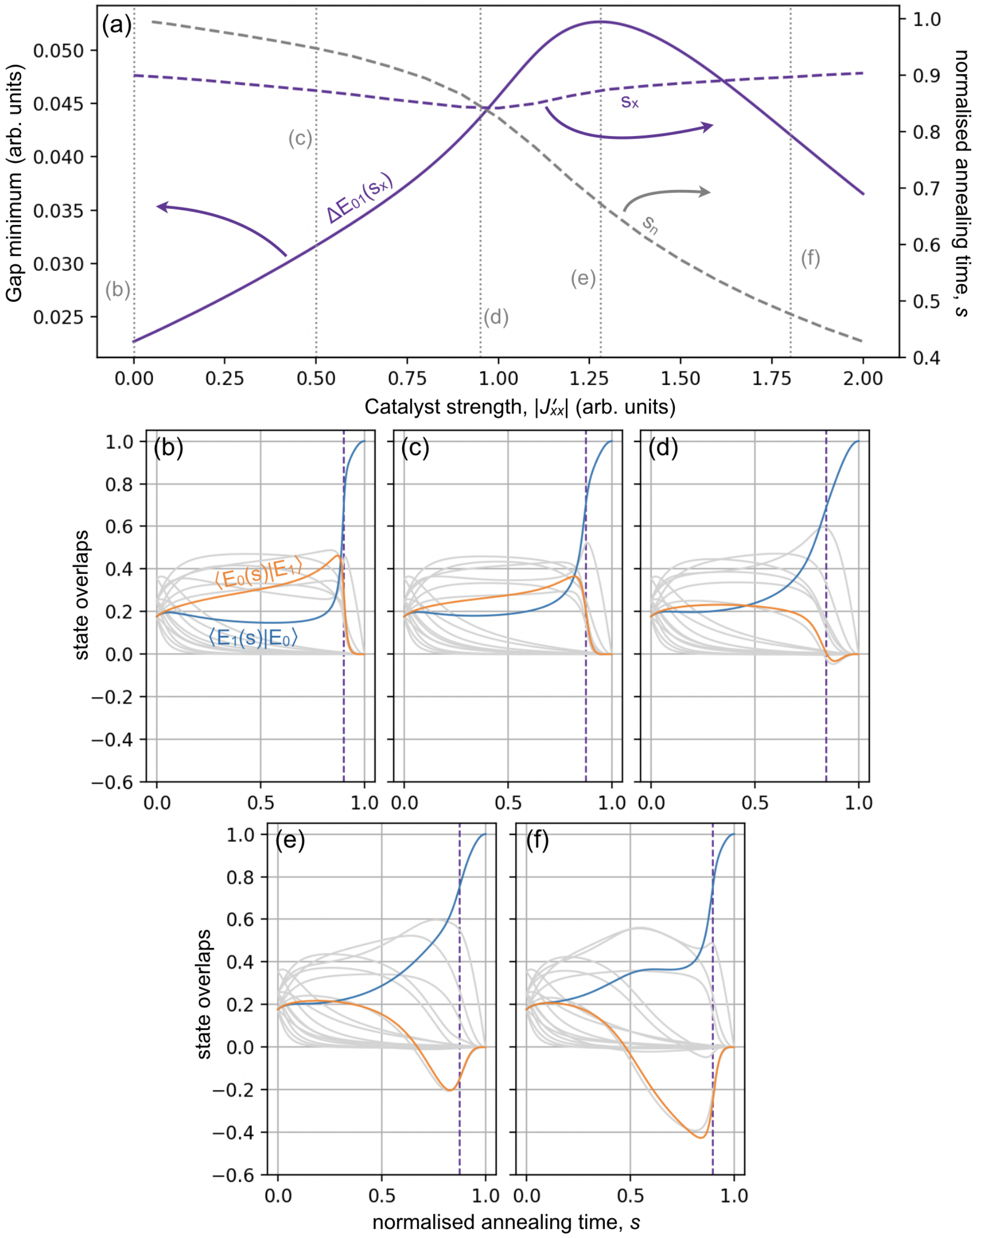

Figure 7.9: Numerical results for a problem instance with $n_0 = 2$, $n_1 = 3$, $\delta W = 0.37$ and $J_{zz} = 37.5$ (*i.e.*, the parameters associated with the WGS) and a catalyst consisting of a single XX-coupling between two vertices in G_1. (a) shows, for increasing catalyst strength, results for 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 either $\langle E_0(s)|E_0\rangle$ or $\langle E_0(s)|E_1\rangle$ becomes negative, s_n (dashed grey). The evolution of the instantaneous ground-state for different catalyst strengths is shown in (b-f). These plots have s_x marked with purple dashed lines. The catalyst strengths for which we show the evolution are marked on (a) with vertical grey dashed lines.

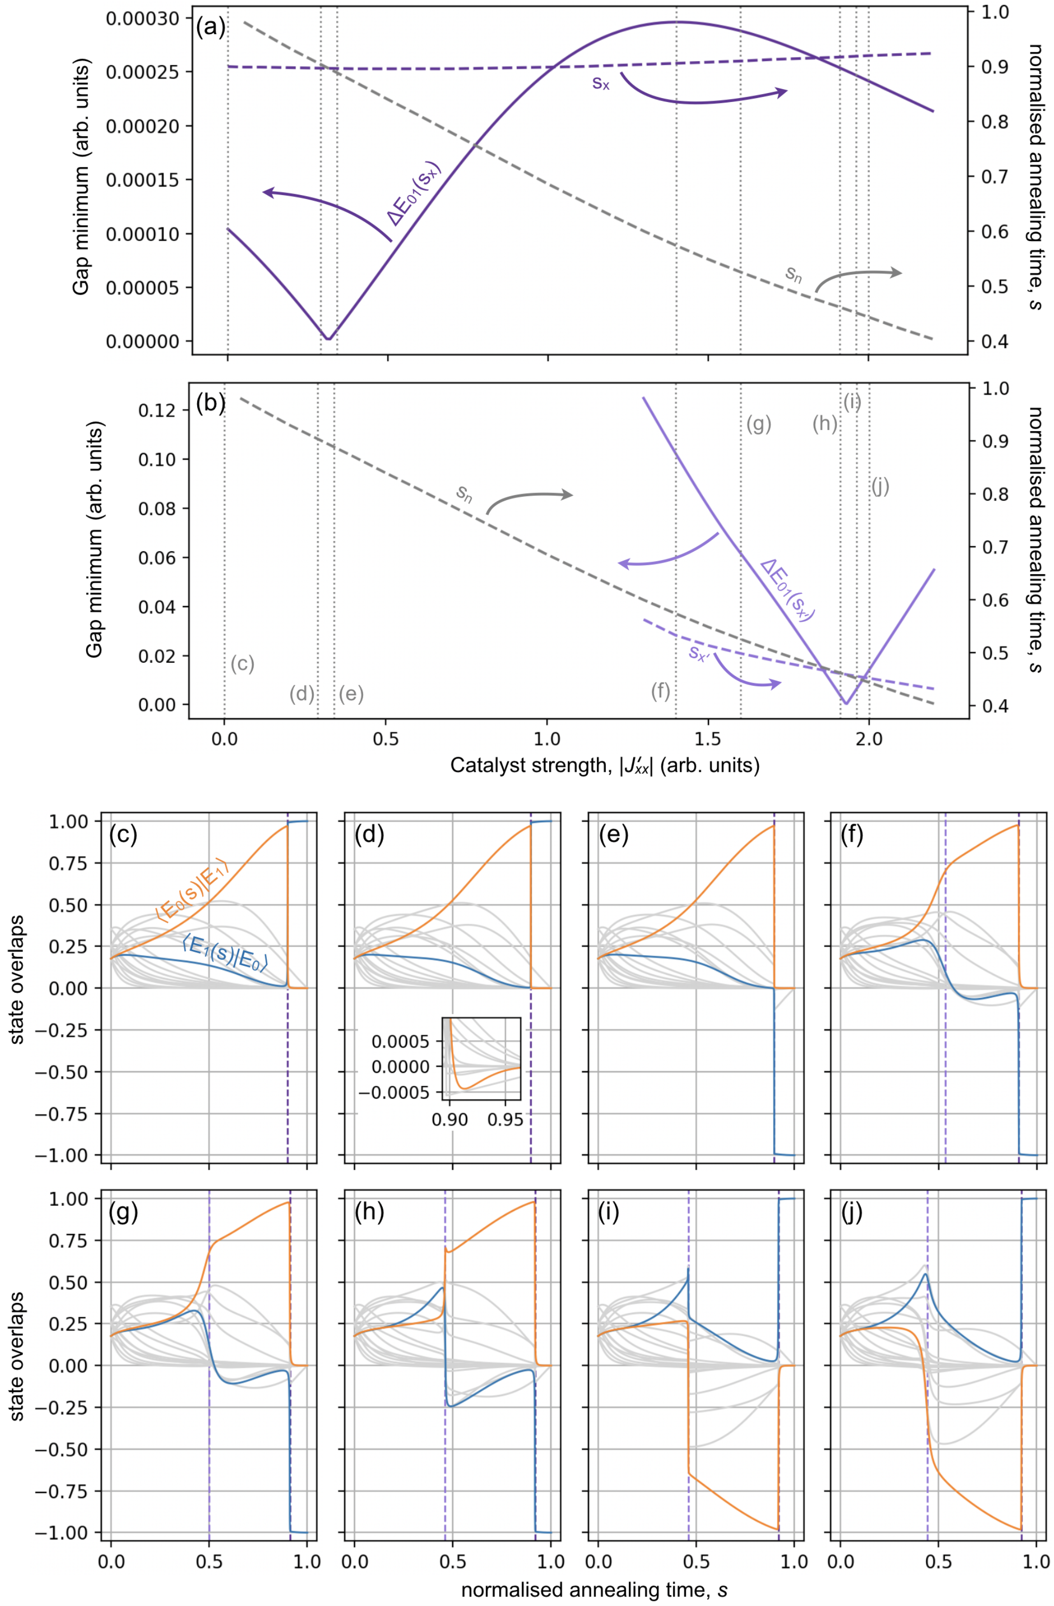

Figure 7.10: As for figure 7.9 but for the parameters associated with the SGS setting, $\delta W = 0.01$, $J_{zz} = 5.33$. (b) shows the size $\Delta E_{01}(s_{x'})$, and location $s_{x'}$ of the new gap minimum that forms in this setting and $s_{x'}$ is marked on plots (c-j) with the same lighter purple used in (b).

that there is also a value of J_{xx} for which $s_n = s_\text{x}$. However, no closing of the gap at the perturbative crossing is observed. To better understand the differences in the behavior, let us compare how the evolution of the GS vector changes with J_{xx}.

First note that, for all the settings and variations on the catalyst examined in this chapter, the GS vector component always takes the opposite sign to the vector component associated with the state representing the local optimum, for $s > s_n$. (For the 5-spin systems the state representing the local optimum is the 1ES.) This is unsurprising given the Hamming distance between states and our discussion in section 5.3. Looking at figures 7.9(b-f), we see that it is consistently the vector component associated with the 1ES that becomes negative when the catalyst is introduced to the WGS setting. In figures 7.10(c-j) however, we see that, in the SGS setting, the component that becomes negative depends on J_{xx}. We find that the relative signs between the vector components remain consistent with varying catalyst strength (and with our discussion in section 5.3) but that our choice of J_{xx} dictates the global signs. Given this fact, it may be tempting to think that there is no physical change occurring. However, as noted in section 4.1.1, when 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 components become negative will result in a measurable difference in the evolution of a system following the ground state.

Let us now continue considering plots (f-h) in figure 7.10. In particular, let us make some observations on the rate of change of the vector components. As J_{xx} is increased from zero, we see that there is a sharpening of this rate of change around $s_{\text{x}'}$ as s_n approaches $s_{\text{x}'}$. Then, as s_n passes $s_{\text{x}'}$, which happens between plots (h) and (i), there appears to be a discontinuous change in which vector components are crossing zero. After this point, the change in the vector components around $s_{\text{x}'}$ begins to soften again. We observe similar behaviour around the value of J_{xx} at which the gap minimum associated with the AC approaches zero – between plots (c) and (d). This however is harder to see since the change is always rapid enough to appear instantaneous for the range over which we are plotting. We do not see this same sharpening of the change in vector components around $s_n = s_{\mathrm{X}}$ in plots (b-f) in the WGS setting – see figure 7.9.

As we have mentioned a few times previously, the rate of the change of the instantaneous GS vector at a given point is intrinsically connected with the spectral gap separating the GS from the 1ES at this point. In some sense then these closing gaps are explained by the sharpening of the rate of change in the GS vector associated with the change in their signs that occurs when $s_n = s_{\mathrm{X}}, s_{\mathrm{X}'}$. However, it does not explain the physics behind this behaviour or why we have only observed it under certain conditions. To recap, we have seen these closing gaps in the SGS setting and when fewer couplings are included in the catalyst – *i.e.*, for the single-coupling catalyst or the all-to-all catalyst on the smallest system sizes. One comment we can make is that, from the preceding discussion, it seems clear that the closing gaps are intrinsically connected to the appearance of negative vector components in the instantaneous GS. We have seen that, and offered an explanation as to why, the appearance of negative vector components is suppressed when additional XX-couplings are introduced. As such it makes sense that we do not see these closing gaps when more couplings are included in the catalyst.

The question still remains though as to why these closing gaps should only appear when the original gap scaling is more severe. We do not have a concrete answer to this but one observation we can make is, in the WGS setting, the AC has already become significantly smoothed out for the values of J_{xx} at which $s_n = s_{\mathrm{X}}$. This can be seen for the 5-spin example in figure 7.9(a) and (d). In contrast, 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 in the SGS setting for the J_{xx} value at which $s_n = s_{\mathrm{X}}$. 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 find that, smoothly transitioning from the SGS to the WGS regime, the J_{xx} values associated with the two closing gaps get closer together. *i.e.*, the range of J_{xx}, for which it is the set of vector components containing $\langle E_0(s)|E_0\rangle$ that crosses zero, shrinks as we move towards the WGS setting. Eventually this region disappears entirely and we no longer see the closing of the gap minima. We have not studied how this cut-off scales with the system size.

7.3 Linearly scaling number of couplings

Section 7.1 examined the effect of the proposed catalyst when the number of XX-couplings introduced scaled quadratically with the system size. We found that such a catalyst seemed able to fully remove the perturbative crossing present in the annealing spectrum and that its performance was insensitive to changes in the severity of the gap scaling associated with the AC. We noted however that implementing this many couplings may not be feasible in practice. We then investigated the use of a catalyst containing only one XX-coupling in section 7.2. Here, we observed additional effects from the catalyst that differed depending on the severity of the gap scaling. While some of the resultant spectra appeared to facilitate faster routes to the GS, the results were less consistent and, certainly in the SGS setting, the catalyst did not have the intended effect of removing the AC. These results perhaps give an indication of how we might expect the catalyst to perform if insufficient couplings are introduced.

As the beginnings of an investigation into what would be a sufficient number of couplings, we will now examine the effect the catalyst has when the number of couplings included scales *linearly* with the system size. Unlike for the all-to-all or single-coupling catalyst, there are many different ways a for linear number of coupling terms to be introduced into G_1 and it is entirely possible the effects from the catalyst would differ depending on the specific graph structure of the XX-couplings present. In this section we investigate only one potential form linearly scaling catalyst. Without any other indicator for a suitable graph structure to use, we have simply chosen one that allows us to continue to utilise the way in which we build our Hamiltonians to reduce the size of the Hilbert space.

The catalyst used in this section has the form

$$\hat{H}_c = J_{xx} \sum_{i \neq j,k \in G_1} (\hat{\sigma}_j^x \hat{\sigma}_i^x + \hat{\sigma}_k^x \hat{\sigma}_i^x), \tag{7.3}$$

where j and k are indices corresponding to two vertices in G_1. That is, two vertices in G_1 are connected to every other vertex in G_1 resulting in a catalyst that contains $2(n_1 - 2) = n - 3$ couplings. The fraction of the total number of couplings included in the catalyst decreases as $1/n$.

Overall, the results closely resemble those for the all-to-all coupling catalyst, with the catalyst seemingly being able to remove the perturbative crossing and the required catalyst strength decreasing with system size.

7.3.1 Weak gap scaling regime

Figure 7.11 shows the same scaling data we have presented for each of the catalyst and parameter setting combinations we have examined so far – with the data colour coded as before. Looking at the results for the gap scaling, we see that the catalyst with a linearly increasing number of couplings is able to enhance the gap such that $\Delta E_{01}^{\min} = \Delta E_{01}$ – seemingly performing as well as the all-to-all catalyst. We also observe that, as for the all-to-all catalyst, the J_{xx} values required for this enhancement decrease with the system size. Interestingly, we find these values do not differ significantly from those needed when the number of couplings included in the catalyst was scaling quadratically with the system size.

We noted, in section 7.1, that we would want the catalyst to be able to enhance $\Delta E_{01}(s_{\mathrm{x}})$ beyond ΔE_{01} to avoid transitions out of the ground state for a polynomial-time anneal. We find that the catalyst containing

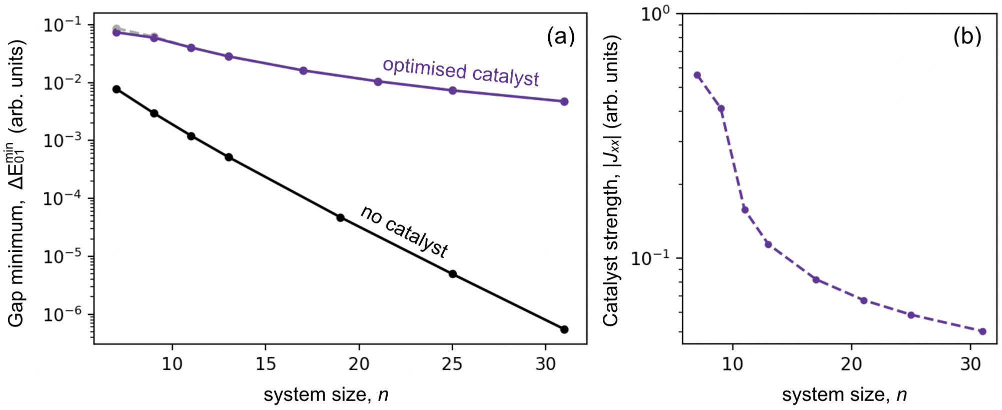

Figure 7.11: Plots showing the scaling behaviour of a catalyst on a bipartite problem with the parameters $\delta W = 0.37$ and $J_{zz} = 37.5$. The sub-graph sizes are scaled as $n_0 = (n-1)/2$ and $n_1 = (n+1)/2$. The catalyst consists of a $n-3$ couplings within G_1 introduced with the opposite sign to the driver. Plot (a) shows the size of the gap minimum in the catalyst free setting (black) and the size of the gap minimum when the optimal catalyst strength is used (purple). We also include the size of the spectral gap between the GS and 1ES at the end of the anneal (grey). Plot (b) shows the $|J_{xx}|$ value which maximises the gap minimum.

a linear number of couplings is indeed able to achieve this for catalyst strengths on the same order as those needed to achieve $\Delta E_{01}^{\min} = \Delta E_{01}$. In addition to the gap enhancement, we observe a similar smoothing out of the exchange in the GS vector components to what we saw for the all-to-all catalyst, indicative of the catalyst successfully removing the perturbative crossing.

Now, let us examine the performance of this catalyst in the SGS setting.

7.3.2 Strong gap scaling regime

The scaling results associated with the introduction of the catalyst to the SGS setting are presented in figure 7.12. We see similar results to what we saw for the all-to-all catalyst, with the smallest spectral gap, $\Delta E_{01}^{\min}$, approaching the maximum possible enhancement, ΔE_{01}, as the system size is increased. We also see, in 7.12(b), a decrease with system size in the catalyst strength needed to achieve this enhancement and, as for the WGS setting, that these J_{xx} values are not dissimilar to those needed

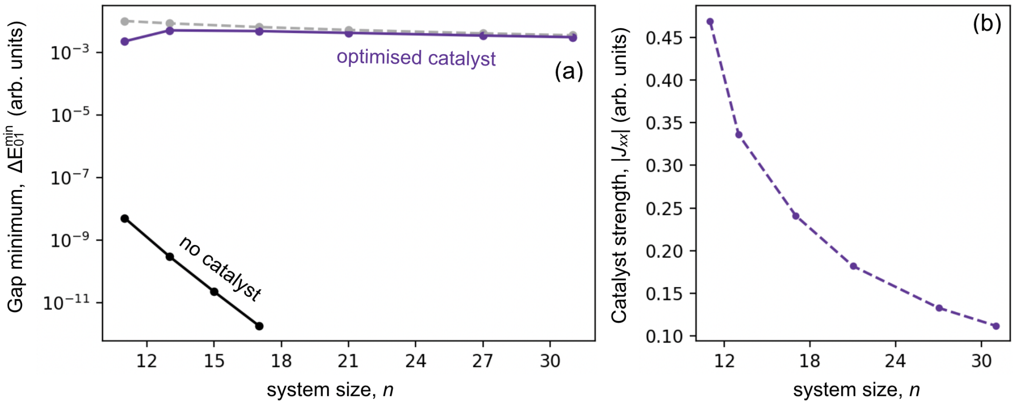

Figure 7.12: Plots showing the scaling behaviour of a catalyst on a bipartite problem with the parameters $\delta W = 0.01$ and $J_{zz} = 5.33$. The sub-graph sizes are scaled as $n_0 = (n-1)/2$ and $n_1 = (n+1)/2$. The catalyst consists of a $n-3$ couplings within G_1 introduced with the opposite sign to the driver. Plot (a) shows the size of the gap minimum in the catalyst free setting (black) and the size of the gap minimum when the optimal catalyst strength is used (purple). We also include the size of the spectral gap between the GS and 1ES at the end of the anneal (grey). Plot (b) shows the $|J_{xx}|$ value which maximises the gap minimum.

when the all-to-all catalyst was used.

Investigating the individual annealing spectra themselves we find that, for the larger system sizes examined, we begin to see the manifestation of the behaviour that we observed when introducing the single-coupling catalyst to the WGS setting. That is, the catalyst results in the appearance of a new gap minimum in the spectrum that eventually leads to two new ACs forming between the GS and 1ES. As such, we may expect the optimal gap enhancement from the catalyst to drop off for the larger system sizes – however this will depend on how the catalyst strength needed to bring $\Delta E_{01}(s_x) = \Delta E_{01}$ compares to the scaling of the value of J_{xx} that results in the new gap minimum dropping below ΔE_{01}. We also noted, in section 7.2.2, that this new double AC may not present a particular problem with regards to ending the anneal in the GS with high fidelity since it could potentially be traversed diabatically. This being the case, the scaling behaviour specifically of $\Delta E_{01}(s_x)$ in the presence of the catalyst, rather than that of $\Delta E_{01}^{\min}$, will be what dictates the success of the catalyst in removing the run-time bottleneck associated with the perturbative crossing.

7.4 Discussion

We have found that the more couplings that are included in the catalyst, the more stable that its effects on the annealing spectrum are to changes in the severity of the initial gap scaling. For the all-to-all catalyst, we observed a similar effect on the annealing spectrum in both the WGS and SGS settings with the catalyst appearing to fully remove the perturbative crossing for all system sizes that we examined. On the other hand, the effects of the single coupling catalyst on the annealing spectra were strikingly different between the two regimes. In the WGS setting, the catalyst was found to remove the perturbative crossing and replace it with two new ACs between the GS and 1ES. In the SGS setting, however, the catalyst was unable to remove the perturbative crossing. Furthermore, it introduced new gap minima and resulted in some of the gap minima going to zero for certain values of J_{xx}. When the number of couplings in the catalyst was scaled linearly with the system size, the effect of the catalyst was relatively similar to when all-to-all couplings were included. However we began to see the appearance of a new gap minimum in the SGS setting for the larger systems that we looked at.

The additional effects that we observed (*i.e.*, anything that was not the intended smoothing out of the perturbative crossing) tended to manifest when the gap scaling in the catalyst free annealing spectrum was more severe. They also appeared to be intrinsically tied to the presence of negative vector components. In section 5.3, we found that including more XX-couplings in the catalyst suppressed the presence of these negative vector components. This may explain why including more couplings in the catalyst resulted in greater stability to these aditional effects. Recall, however, that our suggestion for the reason behind this suppression, was that the additional XX-couplings effectively introduced frustration into equation 5.16. It is possible to select couplings such that limited frustra-

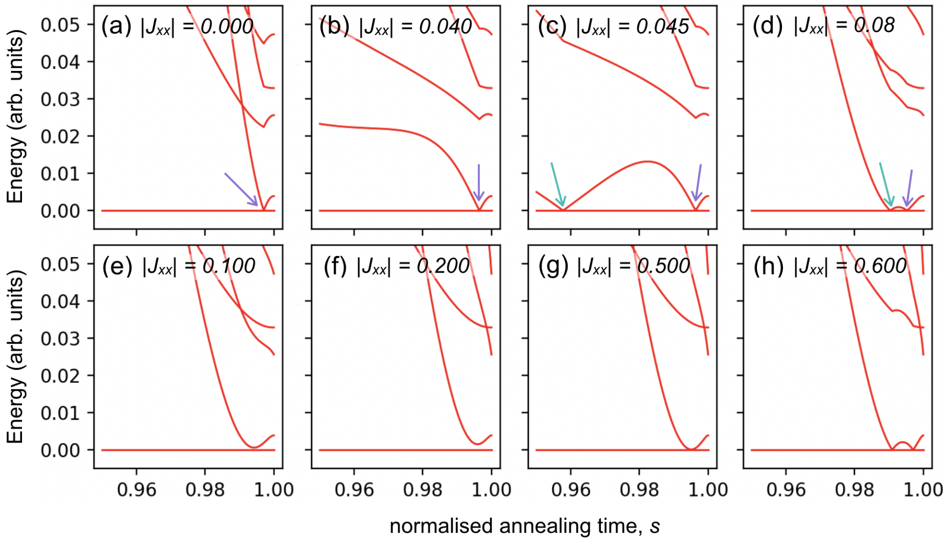

Figure 7.13: Gap spectra corresponding to a problem instance with $G_0 = 12$, $G_1 = 18$, $\delta W = 0.01$ and $J_{zz} = 5.33$. A catalyst is introduced that consists of all couplings within G_1. The catalyst strength, $|J_{xx}|$, associated with each plot is given above the plots. Coloured arrows are used to indicate gap minima that continuously evolve into one another as the catalyst strength is increased.

tion is introduced, as is specifically aimed for in [100], so this may not always hold true. Certainly though, introducing more couplings will increase the possibility of frustration in equation 5.16.

We briefly mention, at this point, a few limited results we have obtained when introducing the catalyst to the setting where we maintain a constant ratio between the two sub-graphs, G_0 and G_1. That is, the subgraph sizes are scaled as $n_0 = 2n/5$ and $n_1 = 3n/5$. Scaling the graph this way results in more severe gap scaling for both the WGS and SGS parameter settings. Furthermore, because we are increasing $n_1 - n_0$, we can expect a corresponding increase in the difference between the driver perturbations to the problem states corresponding to the global and local optima – potentially placing a greater demand on the catalyst. When the WGS parameters were used, the all-to-all catalyst had a similar effect to what was observed for our original graph scaling – with the catalyst seemingly removing the perturbative crossing for system sizes up to $n = 45$ which was the largest system size examined. Furthermore, we see a similar decrease with system size in the J_{xx} values required to achieve this. When the SGS parameters were used, however, we once again observed the appearance of additional gap minima minima in the spectrum – example spectra illustrating this are shown in figure 7.13 for the $n = 30$ system. This, again, suggests to us that the more severe the initial gap scaling, the more likely we are to see the appearance of new gap minima in the spectrum.

The motivation behind the catalyst was to enhance the size of the gap minimum by removing the perturbative crossing present in an annealing spectrum – thereby allowing the evolution to remain adiabatic for shorter anneal times. However, it is worth mentioning that some of the gap minima that appear as a result of the catalyst seem as though they may help the anneal reach the final GS with high fidelity – or at least not prove a hindrance. Consider the gap spectra in figures 7.6(o-p). Here the catalyst was able to remove the perturbative crossing but also resulted in the appearance of two new ACs between the GS and 1ES. If transitions were to occur at these ACs, and limited amplitude exchange occurred during the rest of the anneal, we would expect the system to end in the GS with high fidelity.

Similar comments can be made for the resultant spectra when the all-to-all catalyst was introduced to the SGS setting with the altered graph scaling – *i.e.*, $n_0 = 2n/5$ and $n_1 = 3n/5$. Looking at figure 7.13, we can see that the evolution of the spectrum under increasing catalyst strength is very different to what we have seen in the settings discussed so far. However, the spectra once again appear amenable to a diabatic anneal for a range of J_{xx} values.

In this sense, the appearance of the new ACs do not significantly affect the capacity for the anneal to find the global optimum – with a few caveats. With regards to what we observed in section 7.2.1, the double AC that forms only exists for a certain range of J_{xx}. For J_{xx} just below this range, there exists a single small gap minimum between the GS and 1ES. And, if J_{xx} is increased beyond this point, we see the lifting of one of these ACs such that, again, we are left with just one small gap minimum. How this J_{xx} range relates to the J_{xx} value needed to sufficiently lift the original perturbative crossing will dictate whether a suitable spectrum for a diabatic anneal can be produced. Another point is that transitions out of the ground state mean that spectral gaps beyond $\Delta E_{01}(s)$ must be considered. We might also expect some more complicated dynamics to impact the anneal in the open system setting.

Let us now turn our attention to the closing gap minima that we observed when the single-coupling catalyst was introduced to the SGS setting. It is not clear to us if the effects observed here could be exploited in some way. Looking at figure 7.8(b), which corresponds to the 11-spin SGS system, we see that the catalyst has produced a spectrum that appears amenable to diabatic annealing. That is, the system could transition into the 1ES at the new gap minimum created by the catalyst, and then back into the GS at the perturbative crossing. However, the additional gap minima that appear for the larger systems result in spectra that appear less amenable to diabatic annealing. Nevertheless, in chapter 9, we explore the potential of the single-coupling catalyst to create diabatic paths in settings with multiple perturbative crossings. The manifestation of the single, tunable additional gap minimum in the spectrum also makes the small system sizes of this toy model an ideal test-bed for diabatic annealing protocols. We will be utilising this system, in chapter 8, to investigate the robustness of diabatic annealing to changes in the annealing parameters.

By exploiting the symmetries in our systems, we have been able to reduce the size of the Hilbert space, allowing us to access larger system sizes in our simulations. However, we are still looking at very small systems with regards to the kind of scales one might hope quantum annealing to eventually be applied to. One could optimistically say that our results suggest that a catalyst of this form, containing a linear number of couplings, is sufficient to remove the bottleneck associated with a perturbative crossing – with or without the appearance of the double AC. However, we cannot rule out the possibility of the closings gaps, observed in section 7.2.2, manifesting for larger system sizes when the number of couplings included in the catalyst scales linearly or even quadratically with the system size. A better understanding of the physics behind these closing gaps will help elucidate this question.

That the J_{xx} values necessary to remove the perturbative crossings were found to decrease with system size seems to us a promising result as it suggests the required coupling strengths won't blow up for bigger problem instances. Note that we still saw this decrease in J_{xx} when we switched to the graph scaling for which $n_1 - n_0$ increased with the system size. That being said, we only investigated the all-to-all catalyst in this setting, for which the number of XX-couplings scales quadratically with the system size.

7.5 Summary

In this chapter we have explored the effect of the proposed catalyst on annealing instances constructed to contain a single perturbative crossing. Our aim was to understand how the capacity of the catalyst to remove this perturbative crossing scaled with the problem size as well as how this capacity was affected by the number of couplings included in the catalyst and the severity of the exponential gap scaling of the AC. We first investigated the effects of the catalyst when all possible appropriate XX-couplings were included in the catalyst (*i.e.*, all-to-all couplings within G_1) such that the number of couplings scaled quadratically with the system size. We then examined the other extreme of only including a single coupling in the catalyst. Finally, we examined the performance of the catalyst when the number of couplings included scaled linearly with the system size. In each case, we obtained results for two different parameter settings which affected the exponent of the closing gap minimum present in the original annealing spectrum.

Overall our results suggest that the less severe the initial gap scaling, and the more couplings that are included in the catalyst, the more likely the catalyst is to remove the AC without any additional effects. And the *more* severe the initial gap scaling, and the *fewer* couplings that we include in the catalyst, the more likely we are to see additional closing gaps of some kind. The data we have obtained so far indicates that a linearly scaling number of couplings may be sufficient to remove a perturbative crossing without the creation of any new gap minima and may be able to do so with a catalyst strength that decreases with system size. However, as we have discussed, confirming this for general and larger systems will require further investigation.

Something we have not addressed in this chapter is how suitable values for J_{xx} could be identified in practice and how precise these values would need to be. Such considerations will be crucial for the implementation of catalysts of this nature in more general settings. One potential route to begin investigating this could be to further develop the theoretical framework introduced in section 5.1.2. This, however, we leave for future work.

Chapter 8

Robustness of diabatic annealing

In this chapter, we utilise one of the settings from the preceding chapter to study the robustness of diabatic quantum annealing (DQA). The setting we will be making use of is that from section 7.2.2, in which the catalyst-free gap-scaling was more severe and our catalyst consisted of a single XX-coupling. (We will recap the specifics of the problem setting and catalyst at the start of section 8.1.) Here we found that the catalyst did not remove the single perturbative crossing present in the original annealing spectrum but instead introduced additional gap minima, the size of which depended on the strength with which the catalyst was introduced, J_{xx}.

For $n \lesssim 30$, the catalyst resulted in the formation of just one additional gap minimum between the instantaneous ground and first excited states. Examples of the resultant spectra for two different system sizes can be seen in figures 7.8(b) and (e). The spectra appear amenable to DQA in that a path to the final GS exists which involves the system transitioning into the 1ES at the new gap minimum created by the catalyst, and then back into the GS at the perturbative crossing.

For $n \gtrsim 30$, we observe the appearance of *two* additional gap minima as a result of the catalyst – as seen in figure 7.8(h). Whether or not the resultant spectra associated with these larger systems offer any advantage over the catalyst-free spectra is an open question that we do not address in this work. Certainly though, the straightforward diabatic path outlined above is no longer present. Our focus in this chapter, however, is not to examine the performance of the single XX-catalyst in this specific setting. Rather, we want to leverage the simplicity of the diabatic path present for the smaller system sizes, and that fact that we can easily tune the size of the new gap minimum with J_{xx}, as a tool to study DQA.

In section 2.3, we outlined a number of examples from the literature in which the possibility of DQA was discussed. In particular, an important consideration that has been raised is that there may be more stringent demands on the annealing time chosen for DQA to maintain its advantage (where it has one) over adiabatic annealing [104, 105]. This is because, in DQA, the final GS overlap is a result of a complex interplay of transitions which depend on the speed at which the anneal is run – compared to adiabatic annealing in which one understands that longer run-times will yield equivalent or better results.

For the most part, the discussion around DQA has been with regards to settings where it is facilitated in some way by the original annealing spectrum as dictated by the driver and the problem Hamiltonian [19, 87, 101, 102, 104]. More recently though, there have been proposals for *manipulating* the annealing spectrum to facilitate DQA. Such a possibility was discussed in [100] in relation to the introduction of a catalyst Hamiltonian and in [106] with inhomogeneous driving. And indeed we have seen hints of the possibility of diabatic paths being created in our work as well.

In all these cases, there are one or more parameters that are associated with the manipulation of the annealing spectrum. As such, it will be important to understand the robustness to changes in these parameters of any GS fidelity enhancement that can be obtained – as well as the interplay between the precision needed in these parameters and the annealing time. It is this that we will be investigating using our toy system, in which the manipulation of the spectrum consists of a single tunable gap minimum, controlled by a single J_{xx} value.

We begin this chapter with a numerical study in section 8.1, starting by more fully characterising the response of the annealing spectrum to the introduction of the catalyst in this setting. This will serve as a recap to the setting we will be using for this investigation as well as introducing some quantities that will be needed during our analysis and discussion. We will then, in sections 8.1.1 and 8.1.2, examine results for the closed system dynamics.

To do this, we must first pick an energy scale for our Hamiltonians. Reflective of the energies used in D-Wave Hardware [109], we choose to give our Hamiltonians the units GHz. More specifically, we set

$$\hat{\sigma}^x = \begin{pmatrix} 0 & 1\text{GHz} \\ 1\text{GHz} & 0 \end{pmatrix} \tag{8.1}$$

and

$$\hat{\sigma}^z = \begin{pmatrix} 1\text{GHz} & 0 \\ 0 & -1\text{GHz} \end{pmatrix}. \tag{8.2}$$

Substituting $\hat{\sigma}^x$ into equation 1.2 gives us our driver. Substituting $\hat{\sigma}^z$ into equation 3.8, along with the problem specific parameters (normalised to the driver as described in section 3.4 and appendix A), gives us our problem Hamiltonian. Finally, our catalyst Hamiltonian is defined by substituting $\hat{\sigma}^x$ into equation 5.2 along with a value for J_{xx}. The total Hamiltonian is then given by equation 5.1.

We first, in section 8.1.1, confirm that the system is indeed able to follow the diabatic path that we have described and that this allows a near unity GS fidelity to be obtained for significantly shorter anneal-times than in the catalyst free case. We then, in section 8.1.2, examine the sensitivity of this fidelity enhancement to changes in t_a and J_{xx} and observe a trade-off between the precision needed in the two parameters – with greater precision in one resulting in greater robustness to imprecision in the other.

In section 8.2 we use a Landau-Zener (LZ) model to shed some light on the physics behind these results – allowing us to understand what these findings imply for other settings in which the spectrum is manipulated to facilitate DQA. We discuss these implications in section 8.3. The LZ analysis was done jointly with Ivan Shalashilin.

8.1 Numerical results

Let us start by recapping the problem setting from the previous chapter that we will be using and characterising the behaviour of the spectrum in response to varying J_{xx}.

As for all the problem settings examined in the previous chapter, this was a scalable bipartite graph with the sub-graph sizes $n_0 = (n-1)/2$ and $n_1 = (n+1)/2$. The setting for which we observed the additional, tunable gap minimum was the SGS-setting for which the parameters are $W_0 = 1.01$, $W_1 = 1.00$ and $J_{zz} = 5.33$. The additional gap minimum was produced by a catalyst Hamiltonian that consisted of a single XX-coupling introduced between two vertices in G_1 with a positive catalyst strength, J_{xx}. For plots that show the manifestation of this additional gap minimum, we refer the reader to figures 7.8(a-c).

This new gap minimum seemed, from our numerical results, to go to zero for some critical value of J_{xx} which varied with system size. We will refer to this value as J^*_{xx}. Figure 8.1 shows how these J^*_{xx} values scale with system size – as obtained through numerical sweeps of J_{xx}. We see that J^*_{xx} decreases with system size; we have however not been able to discern any particular scaling behaviour. This is perhaps unsurprising since, as noted in section 7.2.2 and at the start of this chapter, additional effects on the spectrum start to creep in for system sizes approaching $n = 30$.

Decreasing or increasing J_{xx} away from J^*_{xx} results in an increase in the size of the new gap minimum produced by the catalyst. We define the quantity $\Delta J_{xx} = (J_{xx} - J^*_{xx})/J^*_{xx}$ and plot the dependence of $\Delta E_{01}(s_{x'})$ on ΔJ_{xx} (where $s_{x'}$ is the location of the new gap minimum) for different

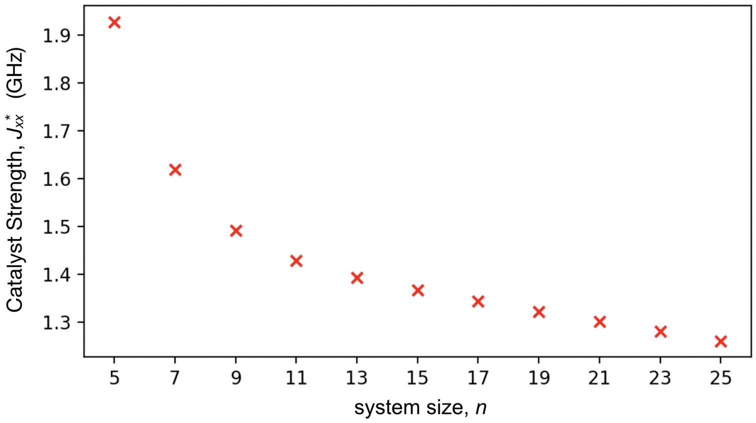

Figure 8.1: Plot showing the numerically obtained catalyst strength, J_{xx}^*, that minimises the gap minimum introduced by the single-coupling XX-catalyst when introduced to problem settings with $n_0 = (n-1)/2$, $n_1 = (n+1)/2$, $W_0 = 1.01$, $W_1 = 1.00$ and $J_{zz} = 5.33$. The units are defined as discussed in paragraph 8 of chapter 8.

system sizes in figure 8.2(a). An approximately linear dependence on $|\Delta J_{xx}|$ is observed with the rate of increase differing slightly depending on whether ΔJ_{xx} is positive or negative. This difference becomes more pronounced for larger system sizes. Note that this linear dependence breaks down for sufficiently large $|\Delta J_{\mathrm{xx}}|$.

In figure 8.2(b) we plot the second derivative of the ground (blue) and first excited state (orange) energies with varying ΔJ_{xx} for the 9-spin system. These second derivatives will become relevant during our Landau-Zener analysis of the dynamics in section 8.2. For now, we simply note that the second derivative of the energy appears to diverge for $\Delta J_{xx} = 0$, further suggestive of some kind of critical point.

Let us now examine the closed system annealing dynamics.

8.1.1 Final GS fidelity enhancement

We will begin by looking at the dynamics when the catalyst is introduced with $J_{xx} = J_{xx}^*$. Our aim in this section is to confirm that the proposed diabatic path does indeed allow the system to reach the final GS for faster anneal times. Throughout this and the following section, we

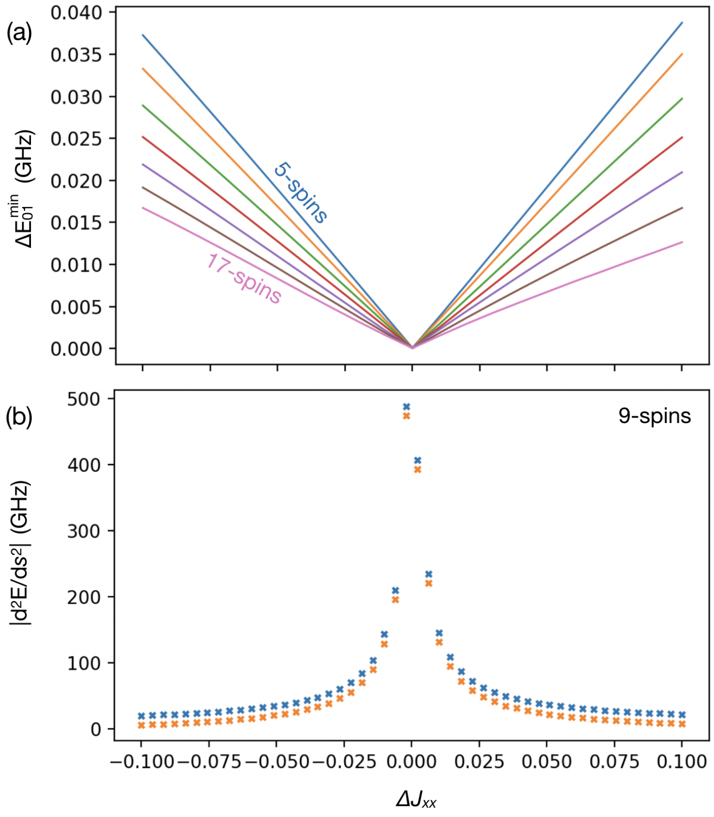

Figure 8.2: Numerical data corresponding to bipartite annealing instances with $n_0 = (n-1)/2$, $n_1 = (n+1)/2$, $W_0 = 1.01$, $W_1 = 1.00$ and $J_{zz} = 5.33$. The plots show different spectral properties varying with ΔJ_{xx}. Plot (a) shows how the size of the new gap minimum varies for system sizes ranging from 5 to 17 in steps of 2. Plot (b) shows the magnitude of the second derivative of the ground (blue) and first excited state (orange) energies at the gap minimum for the 9-spin system. The units are defined as discussed in paragraph 8 of chapter 8.

supplement our discussion with numerical results for the 9-spin system. Similar results were also obtained for the other system sizes examined – $n = 5, 7, 11, 13, 15$ and 17.

Figure 8.3(a) shows the final GS (blue) and 1ES (orange) fidelities with increasing anneal time for the 9-spin system. The results without the catalyst are shown in dotted lines and with the catalyst in solid lines. Without the catalyst, we see that the final GS fidelity is negligible for all the anneal times plotted and that the system ends in the 1ES. Looking at Figures 8.3(b) and (d), which show gap spectrum and dynamics for a $t_a = 2\mu$s anneal, we see that this is a result of the expected transition into the 1ES at the location of the gap minimum.

We were not able to reach the adiabatic limit in our simulations and so were unable to directly extract the t_a values needed to reach a high final GS fidelity in the catalyst free case. A full understanding of the performance of this catalyst with regards to its effect on the time scaling lies outside the scope of this chapter. However, to give us some kind of comparison, we can obtain an estimate for the time needed to reach the adiabatic limit by performing a linear fit to the fidelities obtained for shorter annealing times.

For the 9-spin system, for which we present our example data, the size of the gap minimum associated with the perturbative crossing is small enough that even for a 1ms anneal the final GS fidelity is $< 10^{-8}$. Considering, instead, the 5-spin system, we are able to estimate the adiabatic limit to be on the order of 10ms. The fit was performed to results obtained for $t_a = 10$, 50, 100, 200, 300, 400 and 500μs and had an average squared error of 27μs. The data corresponding to this fit is presented in figure 8.4.

Returning to figure 8.3(a) we see that, by introducing the catalyst with J^*_{xx} for the 9-spin system, the anneal is able to approach a final GS fidelity of unity for t_a on the order of 1μs. Given that the time needed to reach a near unity overlap in the 5-spin system was estimated to be 10ms, this suggests an improvement over the catalyst free case of at least four orders of magnitude – though given the exponentially closing gap minimum, as shown with the lighter purple in figure 3.4(d), the improvement is likely much greater than this.

Looking at figures 8.3(c) and (e), we can see that this reduction in the necessary run-time is indeed a result of the system transitioning to the 1ES at the new gap minimum created by the catalyst, and then back into the GS at the perturbative crossing.

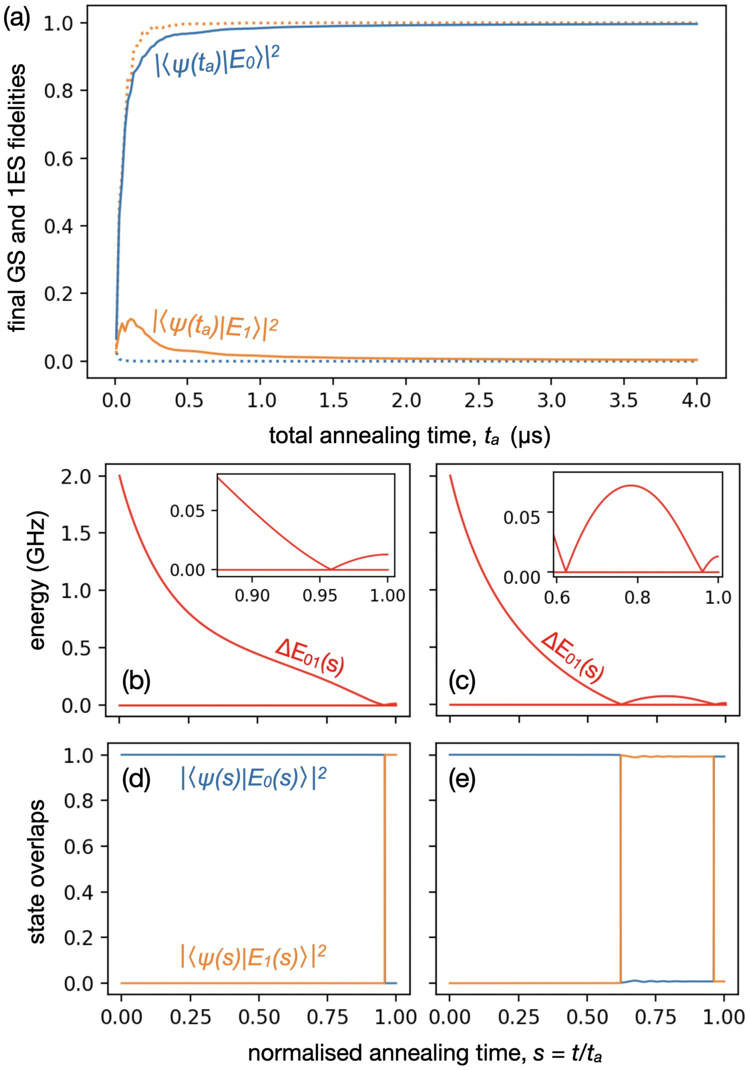

Figure 8.3: Numerical results corresponding to the 9-spin problem instance with $n_0 = 4$, $n_1 = 5$, $W_0 = 1.01$, $W_1 = 1.00$, $J_{zz} = 5.33$ and a catalyst containing a single XX-coupling within G_1. The units are defined as discussed in paragraph 8 of chapter 8. Plot (a) shows how the final GS and 1ES fidelities vary with the total annealing time, t_a, in blue and orange respectively. The results without a catalyst are shown with dotted lines and the results with the catalyst introduced with $J_{xx} = J^*_{xx}$ are shown with solid lines. Plots (b) and (c) show the gap spectrum for the 9-vertex instance without and with a catalyst respectively. Plots (d) and (e) show the corresponding dynamics for a $t_a = 2\mu$s anneal with the state of the system represented in terms of its overlap with the instantaneous ground and first excited states in blue and orange respectively.

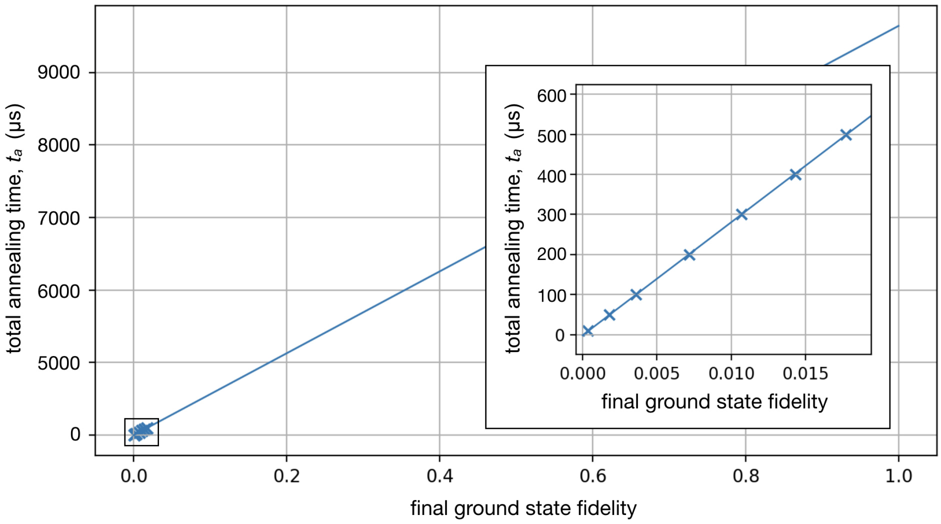

Figure 8.4: Plot showing the time needed to reach different final GS fidelities for the catalyst-free 5-spin problem instance with $n_0 = 2$, $n_1 = 3$, $W_0 = 1.01$, $W_1 = 1.00$ and $J_{zz} = 5.33$. The crosses show the seven numerically sampled data points. A least squares polynomial fit was then obtained. The Hamiltonians are introduced with an energy scale of GHz as defined in paragraph 8 of chapter 8.

8.1.2 Robustness of enhancement

We have established, in the preceding section, that the system is able to follow the proposed diabatic path and that this facilitates an enhancement to the final GS fidelity. We will now explore the *robustness* of this enhancement to changes in the annealing time, t_a, and to deviations in the catalyst strength, ΔJ_{xx}, which changes the size of the first gap minimum involved in the diabatic path.

A colour plot is presented in figure 8.5(a) showing the final GS fidelity for the 9-spin system with varying ΔJ_{xx} and t_a. We first note that we observe a GS fidelity close to unity for a range of t_a and ΔJ_{xx} values. Notably (as can also be seen from figure 8.3 discussed in the previous section), the final GS fidelity remains unity as t_a is increased if $\Delta J_{xx} = 0$. However, for finite $|\Delta J_{xx}|$, we observe a drop-off in fidelity with increasing t_a.

To see this more clearly, we plot the final GS fidelity against t_a for different values of ΔJ_{xx} in figure 8.5(b). The slices of figure 8.5(a) that the

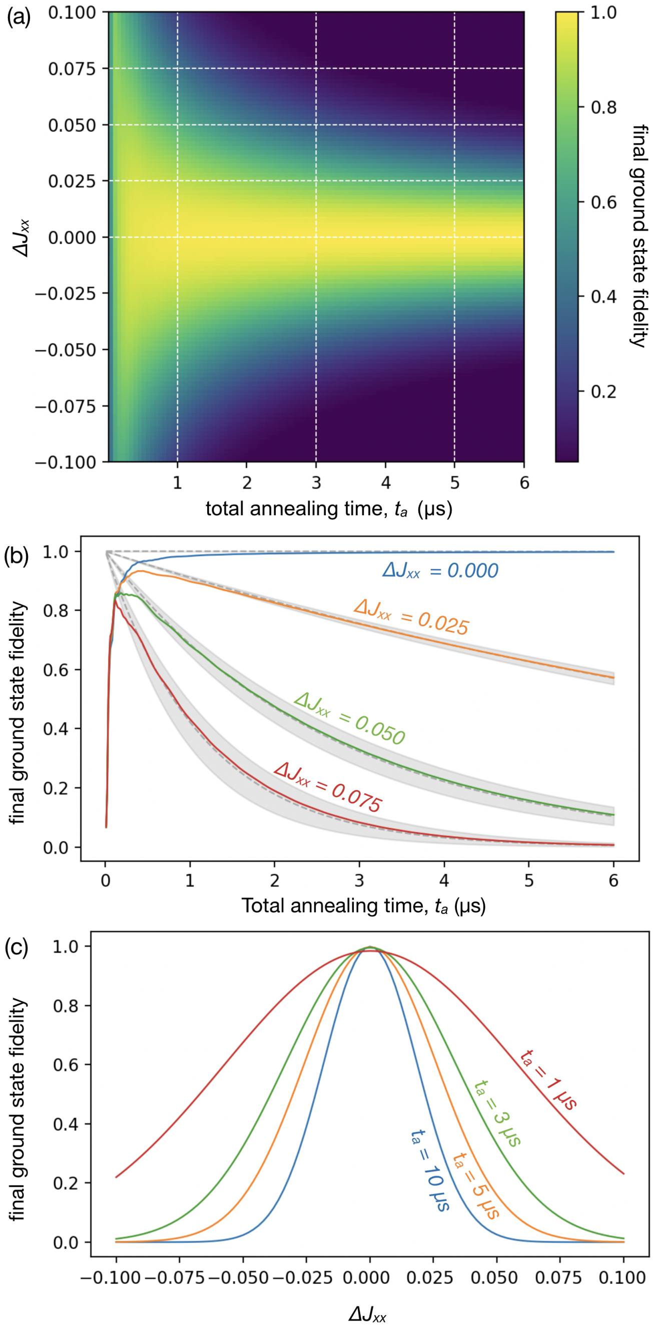

Figure 8.5: Numerical results for the final GS fidelity with varying t_a and ΔJ_{xx} for the 9-spin system discussed in section 8.1. Figure (a) shows a grid plot and Figures (b) and (c) show the slices of this grid indicated with dashed white lines. (b) shows the final GS fidelity with increasing anneal time, t_a, for different values of ΔJ_{xx}. The coloured curves show the numerical results and the grey shaded areas and dashed curves show the Landau-Zener predictions – obtained as described in Section 8.2. (c) shows the results for the final GS fidelity with varying ΔJ_{xx} for different values of t_a. The Hamiltonians are introduced with an energy scale of GHz as defined in paragraph 8 of chapter 8.

curves in this plot correspond to are indicated with white dashed lines. We find that this decrease in fidelity is well fit by an exponential decay, with the decay rate increasing with $|\Delta J_{xx}|$. That there is a drop-off in fidelity for very short anneal times for all ΔJ_{xx} is a result of the dynamics being too fast for the system to follow the instantaneous states such that the system becomes distributed over the state space. In the limit of $t_a = 0$, where the system has no time to evolve at all, the system ends in the equal superposition state and so the final GS fidelity would be $1/2^n$.

These results tell us that if J_{xx} is selected with perfect precision, such that the new gap minimum created by the catalyst is minimised, the enhancement to the final GS fidelity is robust to changes in the total anneal time – given that t_a is sufficiently long for negligible amplitude to leak into other eigenstates. However, as we increase ΔJ_{xx}, *i.e.*, increase the size of the gap minimum created by the catalyst, a greater demand is placed on the value chosen for t_a. If t_a is too large, the fidelity will have decayed significantly from its maximum value and if t_a is too small we will be in the regime where the dynamics are too fast which will also reduce the fidelity. As such, we can say that the less precision we have in J_{xx}, the more precision we need in our selection of t_a.

We briefly note that, for the 5-spin system, oscillations in the final GS fidelity with increasing t_a were observed. These oscillations have little impact on the fidelity in comparison to the decays and are also not present for the larger system sizes. We therefore do not take the time to discuss them here. We suspect, however, that these oscillations are the result of the same interference discussed in [117].

We can also consider how the precision in t_a affects robustness to imprecision in J_{xx}. Looking at $t_a = 6\mu$s in figure 8.5(b), we can see that by choosing an unfavourable anneal time, the final GS fidelity has a strong dependence on ΔJ_{xx} – with its value ranging from unity to zero depending on the catalyst strength chosen. Whereas for $t_a \approx 0.5\mu$s, the fidelity does not drop below 0.7 for the largest ΔJ_{xx}. As such, we can say that the greater our precision in choosing a favourable anneal time, the more robust the final GS fidelity becomes to imprecision in J_{xx}.

To see this more clearly we plot the final GS fidelity with varying ΔJ_{xx} for three different values of t_a in figure 8.5(c). The slices of figure 8.5(a) that these curves correspond to are again indicated with white dashed lines – with the exception of the $t_a = 10\mu$s curve which lies outside the range of figure 8.5(a). For $t_a = 3$, 5 and 10μs, a fidelity of unity is obtained for $\Delta J_{xx} = 0$. The full width half maxima (FWHM) of the curves decreases as t_a is increased, indicating a greater demand on the precision needed in J_{xx}. While the curves continue to increase in width for shorter anneal times, the peaks begin to decrease, reflecting the system delocalising across the state space as a result of the dynamics being too fast.

8.2 Landau-Zener transitions

In the previous section we observed a trade-off between the precision needed in the anneal time, t_a, and the catalyst strength, J_{xx}, in order to preserve the enhancement to the final GS fidelity that could be achieved through the creation of the diabatic path. Let us now discuss the physics behind our numerical results. We begin by making some observations on the dynamics of the individual anneals to allow us to build a qualitative picture of the mechanisms behind our results. These observations will also serve as assumptions when, in the second part of this section, we use a Landau-Zener (LZ) model to lend more rigour to this description. The predictions from the LZ-model are in excellent agreement with the numerically obtained fidelities.

The first observation we make is that, so long as we restrict ourselves to sufficiently long anneal times (for a given system size), no amplitude exchange happens outside of the locations of the two gap minima. As such, if we restrict ourselves to this regime, we only need to consider the dynamics at the location of the perturbative crossing and the tunable gap minimum created by the catalyst. In addition to this, we need only concern ourselves with the subspace spanned by the instantaneous ground and first excited states.

Next, we note that size of the gap minima at the perturbative crossings are such that the amplitude is fully exchanged at this point – *i.e.*, any amplitude in the first excited state will be transferred to the ground state and *vice versa*. For this to not be the case, we would need to increase the annealing time by several orders of magnitude – as indicated by the large anneal times needed to reach the adiabatic limit.

Putting these two observations together, we can expect the final GS fidelity to be precisely the amplitude transferred to the first excited state at the tunable gap minimum created by the catalyst. Thus, when this gap approaches zero, (*i.e.*, we have a small enough $|\Delta J_{xx}|$), we can expect the fidelity to be unity for a range of annealing times – since all the amplitude will be transferred out of the ground state. However, as we increase the gap size by increasing $|\Delta J_{xx}|$, the same run-times will result in less amplitude being transferred to the 1ES, and hence a lower final GS fidelity – as observed in Figure 8.5(c). Correspondingly, keeping ΔJ_{xx} the same and increasing the run-time will result in a lower fidelity – as seen in Figure 8.5(b).

Let us now make this description more quantitative by using the Landau-Zener (LZ) formula to obtain an expression for the diabatic transition rate at the new gap-minimum created by the catalyst. From the preceding observations, we can assume that the system is completely in the ground state prior to this gap minimum and that we can restrict ourselves to the subspace spanned by $|E_0(s)\rangle$ and $|E_1(s)\rangle$ – treating the system as a two level system. The LZ formula for the probability of a diabatic transition is

$$P_D = \exp\left(-2\pi \frac{C^2}{|\frac{d}{dt}(E_a - E_b)|}\right), \tag{8.3}$$

where C is the off-diagonal element coupling the states of the two-level system and $E_{a,b}$ are two linear crossing energy levels between which a

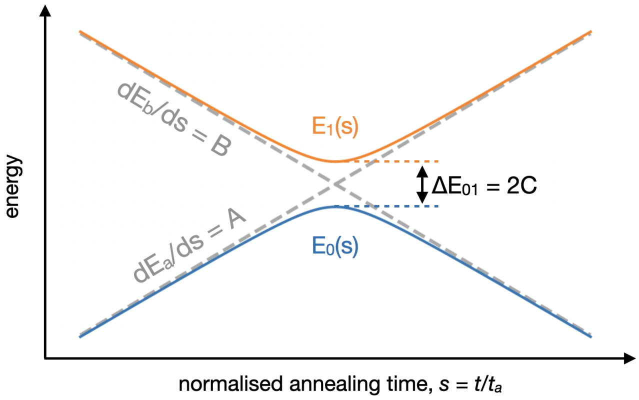

Figure 8.6: Cartoon of an avoided level crossing, illustrating the key components that enter the Landau-Zener formula.

finite gap emerges due to level repulsion. Substituting in $t = s \times t_a$, we can rewrite this expression as

$$P_D = \exp\left(-2\pi \frac{C^2}{|A-B|} t_a\right). \tag{8.4}$$

where $A = dE_a/ds$ and $B = dE_b/ds$. A cartoon illustrating this is presented in Figure 8.6. We have argued that the final GS fidelity will be equal to P_D and so we expect the decay rate of this fidelity with t_a to be

$$\Gamma = 2\pi \frac{C^2}{|A-B|}. \tag{8.5}$$

The following derivation, as well as the numerical extraction of second order derivatives (which will become relevant shortly), was performed by Ivan Shalashilin. To obtain expressions for A, B and C we introduce the 2-level Hamiltonian resulting from the following expansion around the location of the gap minimum, $s_{\mathrm{x}'}$, which we use to first order:

$$H^{\text{2-level}}(s) = \bar{E}(s_{\mathrm{x}'})\mathbb{I} + \begin{pmatrix} A(s - s_{\mathrm{x}'}) & C \\ C & B(s - s_{\mathrm{x}'}) \end{pmatrix} + \mathcal{O}(s^2). \tag{8.6}$$

Here, $s_{\mathrm{x}'}$ denotes the location of the gap minimum, $\mathbb{I}$ is the identity matrix and $\bar{E}(s)$ is the average energy of the instantaneous ground and first excited state, $\frac{1}{2}(E_0(s) + E_1(s))$.

The only effect of the $\bar{E}(s_{X'})\mathbb{I}$ term is to introduce a global phase to the time-dependent unitary. Since this does not affect the dynamics of the system, we may discard it in our analysis. Subtracting $\bar{E}(s_{X'})\mathbb{I}$ from equation 8.6, we obtain

$$\bar{H}^{\text{2-level}}(s) = \begin{pmatrix} A(s - s_{X'}) & C \\ C & B(s - s_{X'}) \end{pmatrix} + \mathcal{O}(s^2)$$

where the energy levels are now centred on zero. This will simplify the following calculations. We can further simplify things by centering the gap minimum on zero with the change of variables $X = s - s_{X'}$ such that we now have

$$\bar{H}^{\text{2-level}}(X) = \begin{pmatrix} AX & C \\ C & BX \end{pmatrix} + \mathcal{O}(X^2). \tag{8.7}$$

The energies of the two eigenstates of this Hamiltonian are

$$E_{0,1}^{\text{2-level}}(X) = \frac{(A+B)X \pm \sqrt{(A-B)^2X^2 + 4C^2}}{2}. \tag{8.8}$$

By taking $X \to \pm\infty$, we find that the energies are linear in this limit and that their gradients are A and B – as desired. C can be read off Equation 8.7 as the off-diagonal element.

Setting $X = 0$ in Equation 8.8 we get:

$$E_{0,1}^{\text{2-level}}(X = 0) = \pm C,$$

$$C = E_1^{\text{2-level}} = -E_0^{\text{2-level}} = \frac{1}{2} \times \Delta E_{01}^{\text{2-level}}, \tag{8.9}$$

where we have dropped the $X = 0$ for readability. Differentiating equation 8.8 twice, with respect to X, and then setting $X = 0$ one obtains

$$\left.\frac{dE_{0,1}^{\text{2-level}}}{dX}\right|_{X=0} = \frac{1}{2}(A+B) \tag{8.10}$$

and

$$\left.\frac{d^2E_{0,1}^{\text{2-level}}}{dX^2}\right|_{X=0} = \pm\frac{(A-B)^2}{4C}. \tag{8.11}$$

Using Equations 8.10 and 8.11, we can obtain the following expressions for A and B:

$$A = E_{0,1}^{\text{2-level}\prime} + \sqrt{\frac{1}{2}\Delta E_{01}^{\text{2-level}} \times |E_{0,1}^{\text{2-level}\prime\prime}|} \tag{8.12}$$

$$B = E_{0,1}^{\text{2-level}\prime} - \sqrt{\frac{1}{2}\Delta E_{01}^{\text{2-level}} \times |E_{0,1}^{\text{2-level}\prime\prime}|}, \tag{8.13}$$

where we have switched to the prime derivative notation and, as in Equation 8.9, we have dropped the $X = 0$ for readability.

Substituting Equations 8.9-8.13 into Equation 8.5 one finds:

$$\Gamma = \frac{\pi}{\sqrt{2}}\frac{(\Delta E_{01}^{\text{2-level}})^2}{\sqrt{\Delta E_{01}^{\text{2-level}} \times |E_{0,1}^{\text{2-level}\prime\prime}|}}. \tag{8.14}$$

The 2-level energies and their derivatives can then be equated with the corresponding numerical results for the full Hamiltonian at the location of the gap minimum. *i.e.*, $\Delta E_{01}^{\text{2-level}} = \Delta E_{01}(s_{X'})$ and $E_{0,1}^{\text{2-level}\prime\prime} = E_{0,1}''(s_{X'})$. We can now write the final decay rate as

$$\Gamma = \frac{\pi}{\sqrt{2}}\frac{\Delta E_{01}(s_{X'})^2}{\sqrt{\Delta E_{01}(s_{X'}) \times |E_{0,1}''(s_{X'})|}}. \tag{8.15}$$

The values for $\Delta E_{01}(s_{X'})$ can be trivially extracted from the numerical data and the second derivatives, $E_{0,1}''(s_{X'})$, are computed by finite differences. That is, we calculate the the second derivatives as

$$E_{0,1}''(s_{X'}) = \frac{E_{0,1}(s_{X'}+\Delta s) - 2E_{0,1}(s_{X'}) + E_{0,1}(s_{X'}-\Delta s)}{2(\Delta s)^2}, \tag{8.16}$$

where the Δs values used are on the order of 10^{-2}.

For these second derivatives, we have a choice of using either $E''_0(s_{X'})$ or $E''_1(s_{X'})$ when substituting into equation 8.15. Equation 8.11 suggests that the two values will have the same magnitude however this will only be true for an actual 2-level system. Figure 8.2(b) shows the numerically obtained values for $|E''_0(s_{X'})|$ and $|E''_1(s_{X'})|$ with varying ΔJ_{xx} for the 9-spin system. We see that the values obtained for $E''_0(s_{X'})$ are larger in magnitude than those for $E''_1(s_{X'})$. As such, using the values associated with the GS will result in a smaller prediction for the decay rate.

The LZ predictions for the decay rates corresponding to the 9-spin system are compared to the numerical results in Figure 8.5(b). We show the area enclosed by the decay predictions obtained using either $E''_0(s_{X'})$ or $E''_1(s_{X'})$ in grey. The decay rates obtained by averaging the two results are shown in dashed grey curves. We see that there is excellent agreement between these decays and the numerical results. That the numerical results don't follow the LZ predictions for $t_a < 0.5\mu$s is because the assumption that no amplitude exchange happens away from the two gap minima no longer holds.

8.3 Discussion

So far, we have studied the closed system dynamics associated with small instances of a toy model in which a single-coupling XX-catalyst created an additional, tunable, gap minimum in the spectrum. We then used an LZ-model shed some light on the physics behind our results.

We numerically determined that if J_{xx} is selected such that the new gap minimum is suppressed to near zero, the final GS fidelity is robust to changes in the anneal time. That is, so long as the anneal is run slowly enough that the system does not delocalise across the state space, a GS fidelity of unity is obtained. If we move away from this value of J_{xx}, such that the gap minimum is larger, we see an exponential decrease in the final GS fidelity with t_a – with the decay rate increasing with $|\Delta J_{xx}|$.

Thus, as we decrease our precision in J_{xx}, we increase the demand on the precision in t_a that is needed to maintain the enhancement to the final GS fidelity. Similarly, by selecting an appropriate anneal time, the effect that imprecision in J_{xx} has on the final GS fidelity can be minimised. However, as we increase our anneal time away from this point, we reduce the robustness of the final GS fidelity to imprecision in J_{xx}. These results indicate a trade-off between the precision needed in the choice of anneal time, t_a, and catalyst strength, J_{xx}.

We attributed this trade-off to the fact that, when $J_{xx} = J^*_{xx}$, the gap minimum created by the catalyst is small enough that all the amplitude is transferred to the 1ES at this point for a range of t_a, such that it is all returned to the GS at the perturbative crossing. However, when $\Delta J_{xx} \neq 0$, the gap becomes larger and a sufficiently short anneal time must be chosen to ensure that the amplitude transferred to the 1ES is close to unity – but not so short that the dynamics are too fast for the system to follow them at all. As such, we expect these findings to apply to other settings where there is some parameter that controls the gap size at a point where we want a diabatic transition to occur. With this in mind, let us discuss what our findings may imply for DQA strategies in general.

8.3.1 Implications for DQA

Let us consider what, in practice, might lead to deviations from the optimal values of J_{xx} (or other parameters that may be used to alter the spectrum, *e.g.* driver fields) and t_a. Firstly, we may have a lack of knowledge of what the optimal parameters actually are. For the catalyst examined in [100], the author proposes a method for obtaining an estimate for a range in which an appropriate J_{xx} value can be found – an appropriate J_{xx} value being one that results in a double AC between the instantaneous GS and 1ES such that a diabatic anneal is possible. In this setting, selecting a J_{xx} value that is too low will result in there being only a single gap minimum between the GS and 1ES. Selecting too large a J_{xx} value will result in an enhancement to gap size at one of the two ACs – *i.e.*, a similar scenario to what we observe here. It is suggested in [100] that QA could be run a polynomial number of times for evenly spaced J_{xx} values within this range. Understanding the viability of this strategy will require the introduction of the selection of annealing time into the discussion since this will impact the precision needed in J_{xx}.

In [106], the diabatic path is created by reducing the driver field on a single target qubit to produce the new gap minimum. Here, the value of the parameter that minimises the gap is known to be zero and so there is no uncertainty in the optimal value that will make the final GS fidelity robust to changes in t_a. However, even if we can expect to know the optimal value of these parameters, there will also be hardware limitations to achievable precision. In D-Wave quantum annealers, which consist of superconducting flux qubits, the precision of the local fields and coupling strengths is limited by, among other things, quantisation errors, flux noise and the fact that the qubits will not be manufactured to be perfectly identical. This can result in errors on the order of 0.01 to the field and coupling strengths.

As an example of what this would mean, Figure 8.5(c) tells us, for our 9-spin example system, that such a deviation in J_{xx} would not reduce the final GS fidelity to below 0.8 so long as $t_a \leq 10\mu s$. What a given hardware precision would mean for a specific system or diabatic annealing strategy would need to be considered on a case by case basis based on the sensitivity of the spectrum to changes in the parameters used to alter it.

There is also a further issue that would need to be considered. As noted, in order for the desired diabatic path to be followed, the annealing time must be sufficiently large to avoid the system transitioning between unwanted energy levels. As such the spectral gaps away from the minima that we want diabatic transitions to occur at also play a key role in the precision required in annealing parameter selection. Depending on their scaling, there is a possibility that the increase in t_a required to prevent

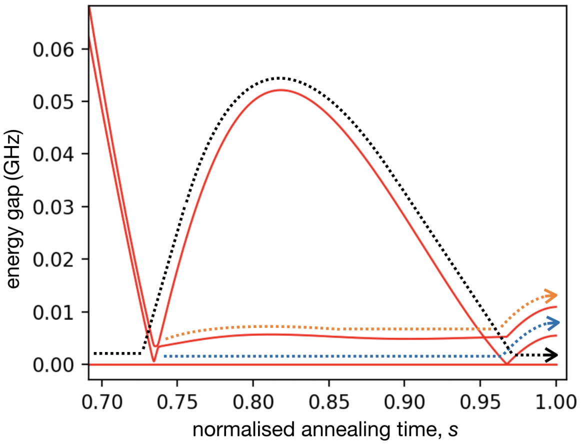

Figure 8.7: Gap spectrum of an annealing instance with parameters $n_0 = 2$, $n_1 = n_2 = 3$,$W_0 = 1.010$, $W_1 = 1.005$, $W_2 = 1.000$ and $J_{zz} = 5.33$. The units are defined as discussed in paragraph 8 of chapter 8. A catalyst is introduced which contains a single XX-coupling between two vertices in G_1 and another between two vertices in G_2. Each of these couplings is introduced with $J_{xx} = 1.125$. The arrows indicate the different paths the system could take, as discussed in section 8.3.

the system from delocalising across state space is such that intractable demands are placed on the parameters used to manipulate the annealing spectrum – even if this gap scaling is sub-exponential. It is also possible that usable parameter ranges will not exist for larger system sizes.

The setting we have chosen for our investigation is, intentionally, very simple. In practice, however, manipulations to the annealing spectrum may be such that the diabatic path created involves many transitions between energy levels. Some such examples can be found in [100]. We will also see, in chapter 9 of this dissertation, examples where an extension of the kind of XX-catalyst used in this chapter produces more complex diabatic paths that involve multiple instantaneous states. We present a single example here in figure 8.7 – though we will leave the details of the formation of this spectrum for the following chapter. Looking at the spectrum in figure 8.7 we see there are two small gaps towards the end of the spectrum, which are also present in the catalyst free case, and two gap minima around $s = 0.74$ which are created by the catalyst.

As with the examples studied in this chapter, the gaps at the ACs at the end of the spectrum are small enough that, unless very long anneal times are chosen, there will be a complete exchange in amplitude at these points. The two new gap minima at $s \approx 0.74$ both increase as J_{xx} is decreased or increased away from some "critical" value. We therefore expect the effect of imprecision in the catalyst strength to be much the same as our findings here – but with the effects compounding at each gap minimum since any amplitude not transferred to the next energy level will not be returned to the GS. This is illustrated in figure 8.7 with coloured arrows.

There are of course many other potential spectra that could be utilised for DQA and numerous potential avenues that may result in techniques for creating such spectra. There will, therefore, likely be many other considerations that will need to be made when implementing such strategies – even when restricting ourselves to the closed system setting. However, for any parameterised change to the annealing spectrum that aims to create new gap minima for a diabatic path, the issues we have highlighted here will be relevant.

8.3.2 Implications for the single-coupling XX-catalyst

So far, we have discussed what we consider to be some general conclusions that can be drawn from this study and what they imply for the successful implementation of strategies where the annealing spectrum is manipulated to facilitate DQA. And indeed this was the intended focus of this investigation. We will however now make some limited comments on the scaling behaviour we have observed for our system in particular. While it seems likely that the behaviour we have seen here will not extend to larger system sizes, we feel this discussion is still useful as an example setting.

Let us consider how the precision needed in J_{xx} and t_a scales with system size. We characterise the precision needed in t_a by the decay rate of the fidelities and the precision needed in J_{xx} with the FWHM of the

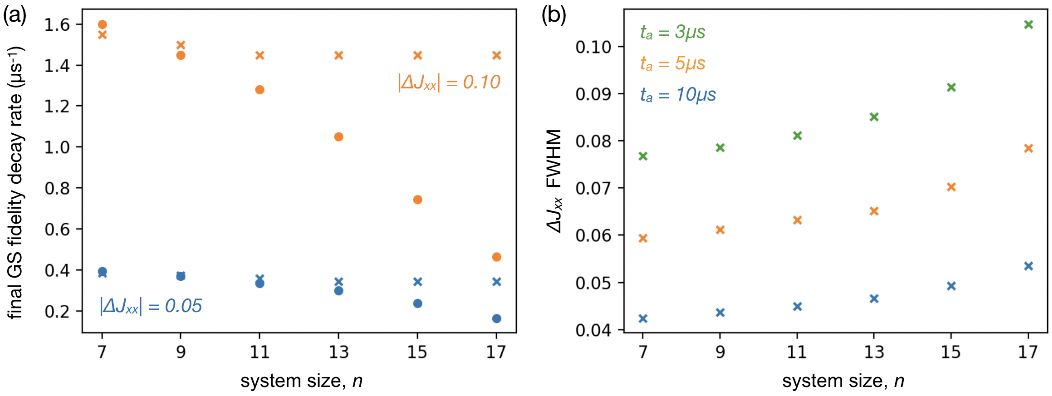

Figure 8.8: Plot (a) shows the decay rate of the final GS fidelity against system size when $|\Delta J_{xx}| = 0.05$ in blue and when $|\Delta J_{\mathrm{xx}}| = 0.10$ in orange. Results for positive and negative ΔJ_{xx} are plotted with dots and crosses respectively. Plot (b) shows the system size scaling of the ΔJ_{xx} FWHM of the final GS fidelity for three different anneal times – 3, 5 and 10μs. The values in this figure are obtained from the numerical data associated with the problem instances with parameters $n_0 = (n-1)/2$, $n_1 = (n+1)/2$, $W_0 = 1.01$, $W_1 = 1.00$ and $J_{zz} = 5.33$. The Hamiltonians are introduced with an energy scale of GHz as defined in paragraph 8 of chapter 8.

curves shown in Figure 8.5(c). Figure 8.8(a) shows the fidelity decay rates when $\Delta J_{xx} = \pm 0.05$ and ± 0.10. The results for positive ΔJ_{xx} values are plotted with solid dots and the results for negative ΔJ_{xx} with crosses. For the positive values we observe a clear decrease in the decay rate with increasing system size – indicating that less precision is needed for the larger systems. For the negative ΔJ_{xx} values, shown with crosses, there is also some reduction in the decay rate. The values quickly plateau, however. In Figure 8.8(b) we see that the FWHM increases with the system size for the three anneal times plotted – also indicating that less precision is needed as the system size increases.

This reduction in the required precision can be attributed to a decrease in the rate at which the size of the gap minimum created by the catalyst increases with ΔJ_{xx} – as seen in 8.2(a). Recall that it is the increase in this gap size that places higher demands on the selection of the annealing time since t_a must be short enough for the system to transition into the 1ES at the gap minimum. As such, a decrease in the sensitivity of $\Delta E_{01}(s_{x'})$ on ΔJ_{xx} will lower the demand on the precision needed in both

J_{xx} and t_a. We also note that the increase in the discrepancy between positive and negative ΔJ_{xx} results we see in 8.2(a) is explained by a corresponding increase in discrepancy in the dependence of $\Delta E_{01}(s_{X'})$ on ΔJ_{xx}.

That the required precision in the t_a and J_{xx} appears to decrease with system size would seem to be a positive result for the performance of the catalyst we have used in this chapter. However, even outside of the additional behaviour that takes shape for the larger systems, the situation is not so simple. That the rate of gap increase with $|\Delta J_{xx}|$ becomes smaller for larger system sizes goes hand in hand with the fact that the spectral gaps generally decrease with increasing system size. This means that larger anneal times will be needed to prevent amplitude leaking between the instantaneous states away from the two gap minima. As discussed in the preceding section, and as can be seen explicitly for this system in figure 8.2(b), an increase in t_a results in an increase to the precision needed in J_{xx}. Understanding how the required precision in J_{xx} scales in practice would require us to more carefully study the interplay between the general gap behaviour in the spectrum and the data presented in Figure 8.2(b).

8.4 Summary

In this chapter, we have performed a numerical and analytical study on the robustness of GS fidelity enhancement through diabatic quantum annealing. Specifically, we have examined robustness of this enhancement to changes in a parameter controlling the annealing spectrum – in our toy model, the strength with which a catalyst Hamiltonian is introduced – and the time for which the anneal is run. Our numerical simulations revealed a trade-off between the precision needed in these two parameters, with greater precision in one resulting in greater robustness to imprecision in the other. We were able to understand this trade-off through the theory of Landau-Zener transitions. Through this understanding, we determined that this finding would also apply to other settings in which the spectrum was being similarly manipulated to facilitate DQA.

In the preceding sections, we have highlighted some of the key implications of our results for the viability of DQA and discussed some of the considerations that we believe will be vital in developing such strategies moving forward. In particular, we have highlighted the importance of introducing the selection of annealing time into the discussion of the precision needed in any parameters being used to modify the annealing spectrum. We have also highlighted a potential issue of usable parameter ranges disappearing for larger system sizes if the scaling of the annealing time required to remain in the desired low energy subspace is too severe – even if this time scaling is sub-exponential.

The findings in this chapter do not, by themselves, give any indication as to whether DQA facilitated through manipulation of the annealing spectrum will be a viable route to improving GS fidelity in quantum annealing. Rather, we hope that they will help guide future investigations into such strategies.

Chapter 9

Removal of multiple perturbative crossings

This chapter explores the use of a multi-part catalyst to target multiple perturbative crossings.

In chapter 7, we examined the performance of a targeted XX-catalyst (as motivated in chapter 5) on annealing instances where the spectrum contained a single perturbative crossing. Overall, our results suggested that the catalyst may be able to alleviate the annealing time bottleneck associated with this perturbative crossing. For most of the cases we considered, this was through gap enhancement facilitating adiabaticity for shorter annealing times – which was in line with the motivation behind the catalyst.

In some cases, though, the effect of the catalyst on the annealing spectrum was more involved, with additional gap minima forming between the instantaneous ground- and first excited-states. The occurrence of these gap minima appeared to relate to additional non-stoquastic effects that could be suppressed by the introduction of more couplings to the catalyst Hamiltonian. While these new gap minima limit the capacity for *adiabatic* annealing, we observed that, for suitably chosen catalyst strengths, some of the resultant spectra appeared amendable to *diabatic* annealing meaning that a faster time to solution could potentially still be achieved in these cases.

In this chapter, we aim to address how these findings translate to cases where there are *multiple* perturbative crossings involving the GS in the original annealing spectrum. While we are also interested in exploiting potential diabatic paths through the spectrum, let us start by considering a catalyst consisting of sufficient couplings to remove a perturbative crossing without producing any new gap minima in the spectrum. This reflects the theoretical motivation behind the catalysts and will thus better aid our discussion. In chapter 5 we saw hints that such a catalyst would result in less gap enhancement when multiple perturbative crossings were present since the perturbative crossing being removed was effectively being replaced by another. (We will recap, with an example, what exactly we mean by this at the start of section 9.1.)

The question then naturally arises: can we simply add another component to the catalyst targeting the next perturbative crossing? Motivated by this idea, we extend the XX-catalysts from equation 5.2 to include multiple components which can target different perturbative crossings. We can write a general catalyst of this form as

$$H_c = \sum_a J_{xx}^{(a)} \sum_{(i,j)\in S_{\text{edges}}^{(a)}} \sigma_x^i \sigma_x^j, \tag{9.1}$$

where each set of edges, $S_{\text{edges}}^{(a)}$, enters with its own catalyst strength, $J_{xx}^{(a)}$. The hope here is that such a catalyst could remove all the perturbative crossings with the GS such that the gap enhancement is comparable to what we observed when only a single perturbative crossing was present. We will explore this in section 9.1 of this chapter.

We are, however, also interested in whether or not catalysts containing fewer couplings may be able to produce spectra amenable to DQA in settings with multiple perturbative crossings. The formation of these spectra was not something we predicted when introducing the motivation behind these catalysts in chapter 5 and so it is less obvious how these effects might extend to such settings. Recall that the diabatic annealing spectra that we highlighted in chapter 7 formed in two distinct ways.

When the initial gap scaling was less severe, the single-coupling catalyst was able to enhance the gap minimum associated with the perturbative crossing but also resulted in the formation of a correlated double AC between the GS and 1ES – see figures 7.6(k-p). We noted that the double AC that formed appeared to relate to the findings in [100]. We therefore choose to focus on the *other* case where we observed the formation of a diabatic path, which we describe in the next paragraph.

When the single coupling-catalyst was introduced into the setting where the initial gap scaling was *more* severe, we saw a negligible effect on the gap minimum associated with the perturbative crossing. For the smaller system sizes examined, the catalyst resulted in the formation of a single additional gap minimum in the spectrum such that a diabatic path to the GS was created – see figures 7.8(a-c). It is currently not clear if this effect will prove useful in the larger systems since we start to see the appearance of a further gap minimum for the larger system sizes that we examined – see figures 7.8(g-i). However, given that the formation of these spectra appear to be a novel effect, we wanted to investigate how it would translate to settings where there was more than one local optimum bottlenecking the algorithm.

We continue to utilise the graph structure introduced in chapter 3 for this investigation, allowing us to easily control the number of local optima that have perturbative crossings with the ground state. This chapter is structured as follows. First, in section 9.1, we examine the effectiveness of the catalyst given in equation 9.1 at removing multiple perturbative crossings. Then, in section 9.2, we explore the capacity of the catalyst to create diabatic paths through a spectrum. In both cases, we will start by going through an example where two perturbative crossings are present in the original annealing spectrum. Then we will look at the same kind of scaling results as we did in chapter 7 but where we are scaling up a tri-partite graph – as shown in figure 9.1. We will then examine the effect of the catalyst when we increase the number of local optima that have perturbative crossings with the ground state by

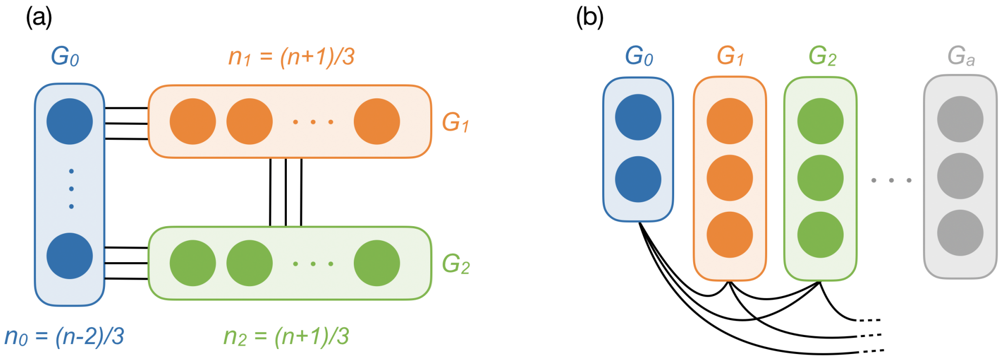

Figure 9.1: In (a) we show how the sub-graph sizes scale for the tri-partite instances examined in sections 9.1.1 and 9.2.1. In (b) we illustrate the n-partite graph constructions used in sections 9.1.2 and 9.2.2

introducing additional sub-graphs to the problem – this is illustrated in figure 9.1(b). The specifics of the parameter choices will be discussed in the relevant sections.

The results in this chapter, and in particular those for the creation of diabatic paths, are rather preliminary. Overall, however, they suggest that the proposed catalyst has the capacity to remove the annealing time bottleneck associated with multiple perturbative crossings. We discuss this in section 9.3 and also comment on how the different components of the catalyst could potentially be introduced recursively based on information obtained from successive annealing runs. We will then highlight what we believe to be the key take-aways in section 9.4.

9.1 Gap enhancement

We begin this section by looking at a single tri-partite example with 11 spins and consider whether or not the catalyst given in equation 9.1 is able to remove the two perturbative crossings that form in the annealing spectrum. The graph sizes associated with this example are $n_0 = 3$, $n_1 = 4$ and $n_2 = 4$ such that both the perturbed states associated with the two local optima cross the perturbed ground state – as illustrated in the cartoon inset in figure 9.2(a).

Our interest in this section is specifically on whether the catalyst can be extended to target multiple perturbative crossings – rather than testing the limits of the settings in which it has the capacity to do so. As such, we choose the graph weights and edge penalty to be $W_0 = 1.6$, $W_1 = 1.3$, $W_2 = 1.0$ and $J_{zz} = 35$ which results in relatively weak ACs. Note, however, that if we scale this graph up as illustrated in figure 9.1 the scaling of the gap minimum remains exponential.

Data for the end portion of the anneal of this 11-spin example in the catalyst free setting is presented in figure 9.2(a). The top plot shows the gap spectrum and the bottom plot shows the evolution of the instantaneous GS vector – with its overlap with the problem ground, first and third excited states highlighted in blue, orange and green respectively. For this example, the state that corresponds to the local optimum associated with sub-graph G_2 is the problem 3ES. The overlaps with other states are shown in grey. Looking at the gap spectrum we see a small gap minimum between the instantaneous GS and 1ES indicative of an AC. At the same value of s, we see an exchange in the problem state overlaps, with the presence of $|E_1\rangle$ decreasing and the instantaneous GS becoming dominated by $|E_0\rangle$ – reflective of a perturbative crossing between these states.

To avoid confusion, let us comment briefly on the fact that some of the state overlaps shown in grey have a greater presence in the GS that the highlighted "key" states. This is because, having moved to the Dicke representation as discussed in section 4.2, the states being plotted encompass all symmetrically equivalent states – such that these overlaps can be thought of as the sum of their contributions.

Let us now consider introducing an XX-catalyst which targets the state responsible for the perturbative crossing identified in figure 9.2(a) – which is the problem 1ES, $|E_1\rangle$. The expected effect of such a catalyst is illustrated in the cartoon inset in figure 9.2(b). *i.e.*, the perturbed 1ES no longer crosses the perturbed GS, a secondary perturbative crossing still remains however. This is reflected in the numerical results. Figure

9.2(b) shows the gap spectrum and instantaneous GS evolution when a catalyst containing all-to-all XX-couplings within G_1 is introduced with $J_{xx}^{(1)} = 0.31$ – which was the value that maximised the gap enhancement as obtained through a numerical sweep.

Looking at the gap spectrum we see that, while the catalyst has enhanced the size of the minimum gap, a distinct gap minimum still remains between the instantaneous ground and first excited state. Tracing through the gap spectrum, we see that a transition at this gap minimum no longer leads to the problem 1ES but to the problem 3ES. The 3ES is the state with which we now expect there to be a perturbative crossing. Looking at the evolution of the instantaneous GS we see that $|E_1\rangle$ no longer has a significant presence prior to the location of the gap minimum whereas the presence of $|E_3\rangle$ has increased – again indicating that the previous perturbative crossing (between $|E_0\rangle$ and $|E_1\rangle$) has been replaced by another, albeit milder, perturbative crossing (between $|E_0\rangle$ and $|E_3\rangle$).

Finally, figure 9.2(c) shows the results when a second component is introduced to the catalyst that targets $|E_3\rangle$. Again, we perform a sweep and select the catalyst strength that maximises the gap enhancement – which we found to be $J_{xx}^{(2)} = 0.12$. For clarity, the final catalyst being used is

$$H_c = 0.31 \sum_{i \in G_1} \sum_{j > i \in G_1} \sigma_x^i \sigma_x^j + 0.12 \sum_{i \in G_2} \sum_{j > i \in G_2} \sigma_x^i \sigma_x^j. \tag{9.2}$$

Once again, the cartoon inset illustrates the intended effect of the catalyst, with the remaining perturbative crossing removed.

Looking at the gap spectrum, we see that the size of the minimum gap in the spectrum has been significantly enhanced over the catalyst free case, as well as the case where the catalyst only targets one of the perturbative crossings. In fact, the smallest gap between the instantaneous GS and 1ES is now much closer in magnitude to the spectral gap at the end of the anneal, $E_{01} = E_1 - E_0$. Turning our attention to the evolution of the instantaneous GS, we see that the presence of $|E_3\rangle$ has been suppressed

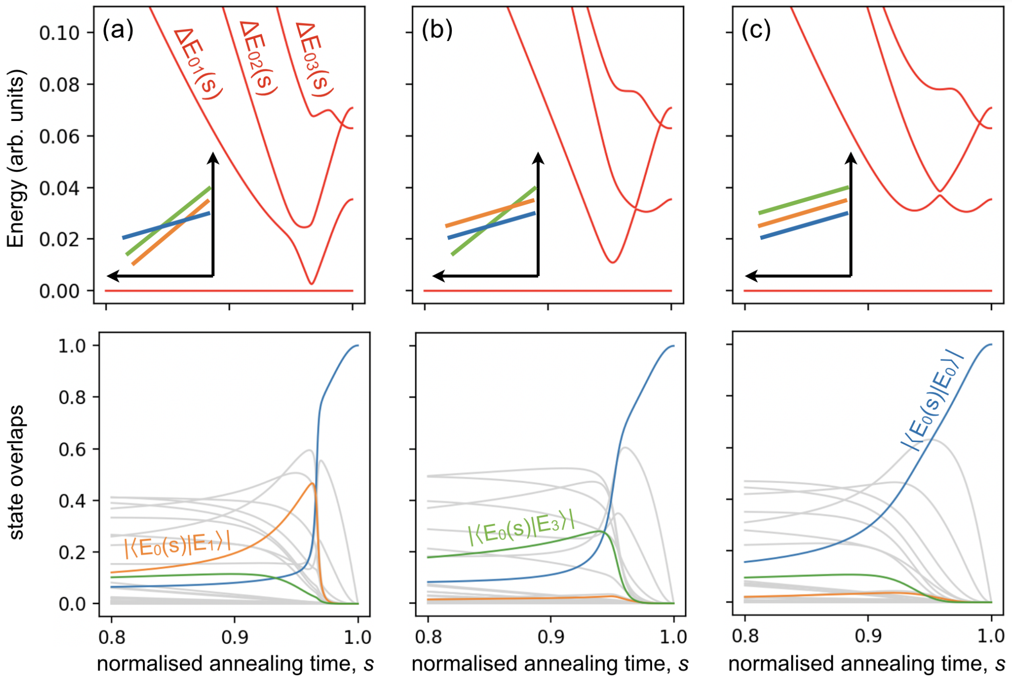

Figure 9.2: Plots showing the effect of a 2-component catalyst on the annealing spectrum corresponding to a tripartite problem graph with the parameters $n_0 = 3$, $n_1 = 4$, $n_2 = 4$, $W_0 = 1.6$, $W_1 = 1.3$, $W_2 = 1.0$ and $J_{xx} = 35$. Gap spectra are presented in the top row and plots showing the evolution of the instantaneous GS vector are presented on the bottom row. The cartoons in the top row illustrate the expected perturbative crossing structure. The overlaps with problem GS, 1ES and 3ES are highlighted in blue, orange and green respectively. (a) shows the catalyst-free case. (b) shows the results when a catalyst containing all-to-all XX-couplings within G_1 is introduced with $J_{xx}^{(1)} = 0.31$. (c) shows the results when we *also* include an additional component to the catalyst consisting of all-to-all couplings within G_2 with $J_{xx}^{(2)} = 0.12$.

over what we saw in figure 9.2(b) and that there is a smoother evolution towards $|E_0\rangle$.

Now that we have seen an example of a piece-wise catalyst targeting more than one perturbative crossing, let us examine the performance of such a catalyst when we increase the system size and the number of perturbative crossings.

9.1.1 Catalyst scaling with two perturbative crossings

We now examine the same sort of scaling behaviour as we did in chapter 7 but for the case where we have two perturbative crossings being removed. In chapter 7, we saw that the catalyst was able to remove a perturbative crossing (without the appearance of additional gap minima) both when all-to-all XX-couplings were included within the relevant sub-graph and when the number of XX-couplings included in the catalyst scaled linearly with the system size. This was the case regardless of the severity of the gap scaling in the catalyst free case – at least within the parameter ranges we have investigated.

We are interested in understanding if the scaling behaviour is as favourable when the catalyst is removing two perturbative crossings. Rather than re-investigate *all* of the above cases for the tri-partite setting, we have chosen to examine just one example setting. We have opted to use the all-to-all catalyst as this reduces the computational complexity of the numerical diagonalisation in comparison to other catalyst choices since it preserves the symmetries we use to reduce the size of the Hilbert space. Reducing the computation complexity is especially crucial when we introduce additional local optima since the more sub-graphs we split our problem into, the larger the reduced-size Hilbert space (as discussed in section 4.2) will be. The problem parameters used are the same as those associated with the example depicted in figure 9.2 which are $W_0 = 1.6$, $W_1 = 1.3$, $W_2 = 1.0$ and $J_{zz} = 35$. We have, however, checked a few example systems with parameters that result in smaller gap minima in the catalyst free case and will also mention these briefly at the end of the section.

Our results for the parameter settings and catalyst choice outlined above are presented in figure 9.3. The left hand plot shows the gap scaling behaviour and the right hand plot shows the catalyst strengths used. The graph is scaled as shown in figure 9.1(a) such that G_1 and G_2 always contain one vertex more than G_0. The scaling of the minimum spectral gap *prior* to the introduction of the catalyst is shown in figure 9.3(a) in

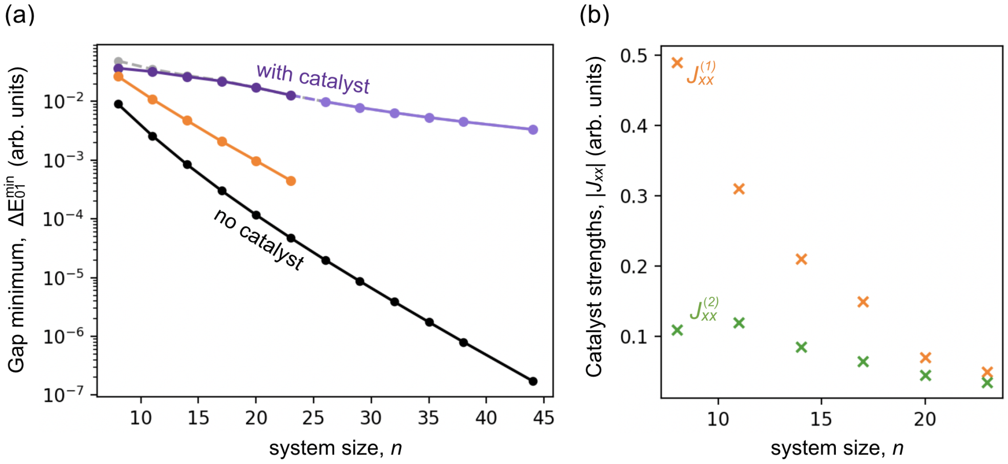

Figure 9.3: Plots showing numerical results for the problem instances scaled as shown in figure 9.1(a) with the problem parameters $W_0 = 1.6$, $W_1 = 1.3$, $W_2 = 1.0$ and $J_{zz} = 35$. Plot (a) shows the scaling of the gap minimum without the presence of a catalyst in black. The gap scaling when a two-component catalyst is used is plotted in purple. The darker purple shows the results when the optimised parameters for each system size, as shown in plot (b), are used. The results shown in the lighter purple use the optimal parameters for the $n = 23$ system. In orange, we plot the results when only one component of the catalyst is included – *i.e.*, the values plotted for $J_{xx}^{(1)}$ in (b) are used and $J_{xx}^{(2)}$ is set to zero. Finally, we plot the spectral gap at the end of the anneal for each system size in grey. The majority of the grey points are obscured by the purple points.

black and we can see that the gap scaling is exponential as expected.

We now introduce a catalyst of the form given in equation 9.1 with two sets of couplings. $S_{\text{edges}}^{(1)}$ and $S_{\text{edges}}^{(2)}$ consist of all-to-all couplings within G_1 and G_2 respectively. We first perform parameter sweeps in a grid for systems up to $n = 23$ spins to obtain the optimal $J_{xx}^{(1)/(2)}$ values. These are presented in figure 9.3(b). The resultant gap minima are presented in figure 9.3(a) in the darker purple. As with the bipartite systems, we see that (a) the gap enhancement from the catalyst quickly tends to the maximum possible gap enhancement as the system size increases and (b) the catalyst strengths decrease with system size.

As for the corresponding bipartite setting, in section 7.1.1, we see that this maximum gap enhancement results in a scaling that still appears to be exponential, albeit with a significantly weaker exponent. Considering the same factors as we did in the bipartite case, we conclude that (a) this scaling does not bottleneck the algorithm since it is associated with problem states corresponding to sub-sets of the global optimum and (b) the perturbative crossings have been fully removed by the catalyst for system sizes ≥ 20. We will not repeat this discussion here but refer the reader to paragraphs 2-5 in section 7.1.1.

We can make sense of the fact that $J_{xx}^{(1)} > J_{xx}^{(2)}$ by considering the properties of the two perturbative crossings that the corresponding components of the catalyst are targeting. Both perturbative crossings are between the problem GS and an excited state associated with one of the local optima. The unperturbed energy of the excited state responsible for the crossing that $J_{xx}^{(2)}$ is targeting is larger than that of the excited state that $J_{xx}^{(1)}$ is associated with. From the construction of the problem graph, the driver perturbation to these problem states will be comparable. We therefore expect, from the arguments made in chapter 3, the crossing with which $J_{xx}^{(2)}$ is associated to be comparatively weaker. As a result, we would expect a smaller perturbation from the *catalyst* to be able to remove this perturbative crossing. Since $S_{\text{edges}}^{(1)}$ and $S_{\text{edges}}^{(2)}$ contain the same number of edges, we can expect the required $J_{xx}^{(2)}$ to be smaller than $J_{xx}^{(1)}$.

As we increase the size of our problem instances, the two dimensional optimal parameter sweeps become quickly intractable. Our initial plan was to perform a fit to the optimal J_{xx} values we obtained for system sizes up to $n = 23$ and use this to estimate the values needed for the larger system sizes. In practice though, we found that simply using the J_{xx} values obtained for the $n = 23$ spin instance gave better results with regards to gap enhancement. Assuming that the J_{xx} scaling behaviour observed in figure 9.3(b) persists, however, we would expect smaller values of $J_{xx}^{(1)}$ and $J_{xx}^{(2)}$ to achieve the same enhancement to the minimum spectral gap. The gap scaling results for $26 \leq n \leq 44$, using the J_{xx} parameters obtained for $n = 23$, are plotted with the lighter purple in figure 9.3(a). We see that the catalyst continues to achieve the maximum gap enhancement for these larger system sizes as well.

We also show, in figure 9.3(a), the gap-scaling results when just the first component of the catalyst is included – this is plotted in orange for system sizes up to $n = 23$. In this case, we expect the stronger perturbative crossing to have been removed while the weaker one remains. And indeed we see that, while there is some gap enhancement, the scaling is not as significantly altered as when both components of the catalyst are included. Looking at the individual gap spectra, we see similar results to what we observed for the example in figure 9.2(b) – with the exchange in the problem state coefficients sharpening as the system size increases.

As we noted at the start of this section, we have not obtained corresponding results when parameters closer to those of the SGS-setting from chapter 7 are used. However, we have tested a few example instances and found that the two-part catalyst does appear to remove both perturbative crossings in these instances as well.

Now, let us examine the effect of the catalyst as the number of perturbative crossings is increased above two.

9.1.2 Two or more perturbative crossings

The problem instances examined in this section are on N-partite graphs where the sub-graph associated with the global optimum has two vertices while the other sub-graphs (corresponding to the different local optima) all have three vertices – as illustrated in figure 9.1(b). We assign G_0 a total weight of 1.3 and G_1 a weight of 1.0. Each additional sub-graph that we include is assigned a weight 0.01 smaller than that of the previous sub-graph such that the final weights are $W_0 = 1.3$, $W_1 = 1.0$, $W_2 = 0.99$, $W_3 = 0.98$ *etc.* We take $J_{zz} = 35$ as in the previous section. The reason we have assigned the sub-graphs associated with the local optima similar weights, rather than spacing the weights of the sub-graphs equally between the highest and lowest weighted sub-graph, is so that introducing additional local optima does not result in the problem 1ES energy being pushing increasingly close to the GS energy.

Note that we do *not* expect an exponential closing of the minimum spectral gap in the catalyst-free setting when we use the graph scaling illustrated in figure 9.1(b). This is because the exponential scaling typically associated with a perturbative crossing is a result of an increase in Hamming distance between the two problem states which are involved in the perturbative crossing (as discussed in section 3.1). While this is generally expected to increase with system size, here the Hamming distance between any of the local optima and the global optimum is always five. This expected lack of exponential scaling is reflected in the numerical results.

We have examined systems with between three and six sub-graphs meaning that the largest system we consider here has $n = 17$ spins. The resultant annealing spectra then have between two and five perturbative crossings. We can still use the Hilbert space reductions from section 4.2. However, because we are increasing the number of sub-graphs rather than their sizes, the size of Hilbert space still scales exponentially with the system size. This severely limits the system sizes for which we are able to obtain the annealing spectra.

For each of the systems, the number of components we include in the catalyst Hamiltonian (as given in equation 9.1) will equal the number of perturbative crossings, meaning we need to select as many J_{xx} values. Performing parameter sweeps to identify the optimal values becomes quickly intractable since, not only is the Hilbert space scaling exponentially, but the dimension of the search space is also increasing. The parameters we used were obtained as follows. We first performed a parameter sweep for the tripartite instance to obtain the optimal $J_{xx}^{(1)}$ and $J_{xx}^{(2)}$ values needed for the catalyst in this setting. We then kept the value obtained for $J_{xx}^{(1)}$ for the 4-partite system and performed a 2-dimensional sweep to select values for $J_{xx}^{(2)}$ and $J_{xx}^{(3)}$. We then kept this value of $J_{xx}^{(2)}$ for the next system and optimised the other two parameters and so on. The resultant parameter sets used are shown in figure 9.4. That each suc-

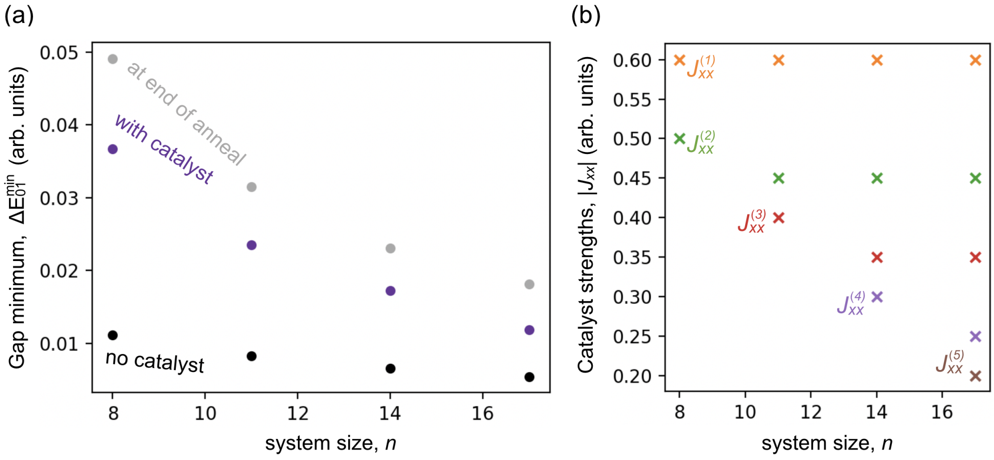

Figure 9.4: Numerical results for the problem instances scaled as shown in figure 9.1(b) with the problem parameters $W_0 = 1.30$, $W_1 = 1.00$, $W_2 = 0.99$, $W_3 = 0.98...$ and $J_{zz} = 35$. Plot (a) shows the minimum spectral gap with and without the presence of a catalyst in purple and black respectively. The spectral gap at the end of the anneal for each system size is plotted in grey. Plot (b) shows the J_{xx} values associated with the catalyst at each system size.

cessive J_{xx} value is lower than the last (*i.e.*, $J_{xx}^{(1)} > J_{xx}^{(2)} > J_{xx}^{(3)}...$) can be explained in much the same way as the J_{xx} magnitudes in the previous section.

Figure 9.4(a) shows the sizes of the minimum spectral gap, $\Delta E_{01}^{\text{min}}$, with and without the presence of the catalyst in purple and black respectively. As always, we include the spectral gap at the end of the anneal, ΔE_{01}, in grey. For all the system sizes examined, we see that the gap is enhanced by the catalyst – although not to the maximum possible value. This is perhaps unsurprising since this maximum enhancement was also not reached for the smallest bi- and tripartite systems. We related this to the number of couplings included in the catalyst being insufficient to suppress additional effects from the non-stoquastic catalyst that resulted in non-monotonic gap enhancement. While the total number of couplings being included in the catalyst increases as we add more subgraphs to the problem, the number of couplings that we expect to effect each particular perturbed problem state is always three.

It is hard to draw any concrete conclusions regarding how the performance of the catalyst scales with the number of perturbative crossings. We do observe what appears to be a slight decrease in the performance of the catalyst in that $\Delta E_{01} - \Delta E_{01}^{min}$ in the presence of the catalyst grows in relation to $\Delta E_{01} - \Delta E_{01}^{min}$ in the catalyst free case. However, this may simply be a result of the J_{xx} values becoming increasingly removed from the optimal values. To understand the actual scaling behaviour we would need to perform a full optimisation of the parameters, either through running extremely computationally-intensive sweeps or by employing some kind of search algorithm.

9.2 Diabatic path

Let us now investigate how the effects of the catalyst that we saw in section 7.2.2 extend to settings with multiple local optima. To recap, when we introduced the single-coupling XX-catalyst to the SGS-setting from chapter 7, we saw a limited impact on the size of the gap minimum associated with the perturbative crossing and instead observed the appearance of new gap minima in the spectrum. We noted the creation of a *diabatic* path through the spectrum in that the system could transition into the 1ES at a gap minimum created by the catalyst and then back down into the GS at the perturbative crossing. As the system size increased, we observed the appearance of a further gap minimum in the spectrum as a result of the catalyst, disrupting this diabatic path. Nevertheless, we feel this is an interesting effect that is worth exploring for settings with more than one perturbative crossing.

For the investigation carried out in the first half of this chapter, the motivation behind introducing a catalyst component for each perturbative crossing was clear – if a catalyst targeting one local optimum could remove a single perturbative crossing, introducing a piece-wise catalyst to remove *multiple* perturbative crossings was a natural next step. However, as we noted in the introduction to this chapter, it is less clear if the findings in section 7.2.2 will translate to anything useful with the catalysts extended in the same way. Since we are not clear on the mechanism behind the formation of the gap minimum that created the diabatic path, we did not necessarily have any reason to believe that introducing additional components to the catalyst would create further gap minima – and, if they did, that the new gap minima would create a spectrum amenable to DQA. What we have, nevertheless, found in this section is that this form of catalyst can in fact produce such a diabatic path through the spectrum.

Let us go through an example with two perturbative crossings to illustrate this. Figure 9.5 shows numerical results for the gap spectra corresponding to annealing instances of a problem with $G_0 = 2$, $G_1 = 3$, $G_2 = 3$, $W_0 = 1.02$, $W_1 = 1.01$, $W_2 = 1.00$ and $J_{zz} = 5$. (We also include results for the evolution of the instantaneous GS. We will not however be discussing these results until the end of the chapter.) This instance has the same perturbative crossing structure as the instance associated with figure 9.2 in the catalyst-free case. That is, the perturbed states associated with the two local optima cross the perturbed ground state as shown in the cartoon inset in figure 9.2(a). Here, the problem states associated with the two local optima are the first and second excited states.

Figure 9.5(a) shows the results in the catalyst-free case where we see that indeed there is a small spectral gap at around $s = 0.95$ which corresponds to the perturbative crossing between the problem GS and 1ES. Figure 9.5(b) then shows the result of introducing a single XX-coupling chosen to target this perturbative crossing – *i.e.*, an XX-coupling within G_1. As in section 7.2.2, we see that the original small spectral gap remains but now there is a new gap minimum in the spectrum. However, in this case, this new gap minimum does not result in a diabatic path to the GS. Assuming the annealing time chosen is insufficient to avoid a transition at the new gap minimum created by the catalyst, we might expect the system to end in the problem 2ES with high fidelity. Finally,

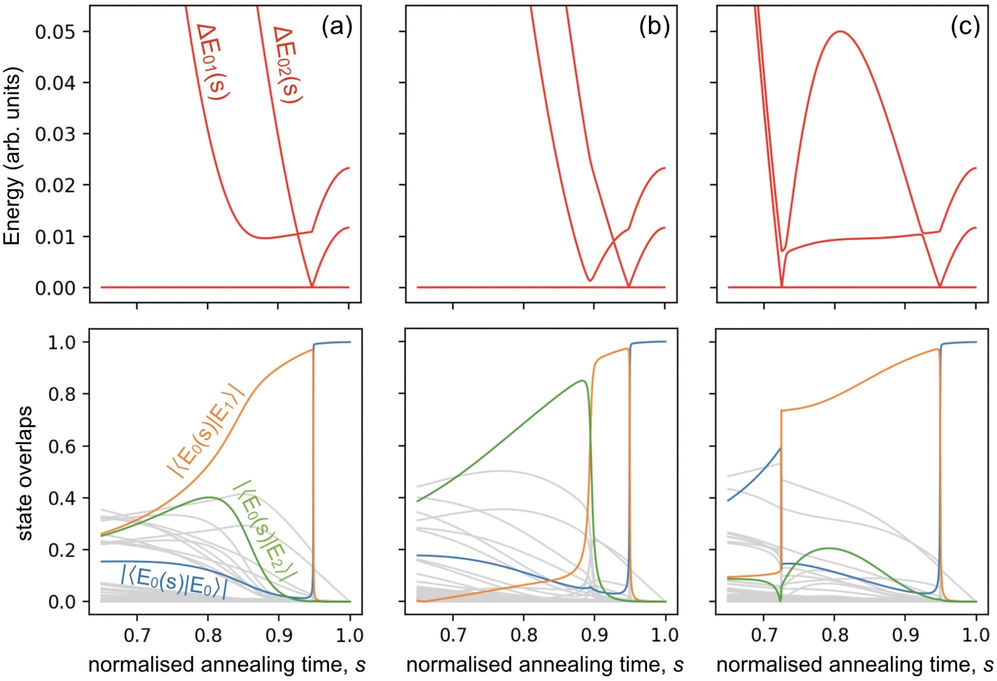

Figure 9.5: Plots showing the effect of a 2-component catalyst on the annealing corresponding to a tripartite problem graph with the parameters $n_0 = 2$, $n_1 = 3$, $n_2 = 3$, $W_0 = 1.2$, $W_1 = 1.1$, $W_2 = 1.0$ and $J_{xx} = 5$. Gap spectra are presented in the top row and plots showing the evolution of the instantaneous GS vector are presented on the bottom row. The overlaps with problem GS, 1ES and 2ES are highlighted in blue, orange and green respectively. (a) shows the catalyst free case. (b) shows the results when a catalyst containing a single XX-coupling within G_1 with $J_{xx}^{(1)} = 2.14$ is introduced. (c) shows the results when we *also* include an additional component to the catalyst consisting of a single coupling within G_2 with the same catalyst strength.

figure 9.5(c) shows the result of including a second XX-coupling that targets the other perturbative crossing – *i.e.*, an XX-coupling within G_2. Now we see that, if the system were to transition between the instantaneous states at the locations of the gap minima, it would enter the 1ES, then the 2ES, before transitioning back into the 1ES and then the GS.

To confirm that this hypothetical path through the spectrum can indeed be exploited diabatically, we simulate the closed system dynamics. As described in chapter 8, we assign the operators $\hat{\sigma}^x$ and $\hat{\sigma}^y$ the units GHZ, emulating the energy scaled of the D-Wave devices. Figure 9.6(a) shows the overlap of the system with the instantaneous GS, 1ES and 2ES over

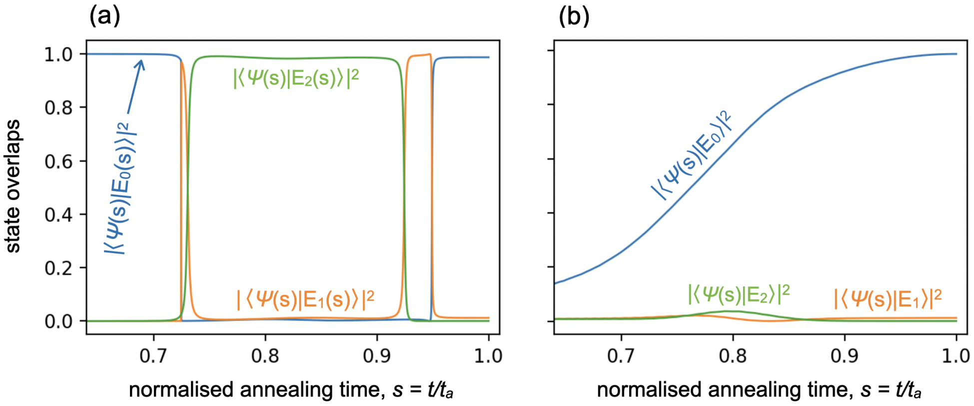

Figure 9.6: Plots showing results for the closed system dynamics corresponding to the setting with the gap spectrum depicted in figure 9.5 (c). The Hamiltonians are introduced with an energy scale of GHz (as defined in paragraph 8 of chapter 8) and the annealing time is chosen to be $t_a = 1\mu s$. Plots (a) and (b) show the evolution in terms of the system's overlap with the instantaneous states and problem states respectively – with the overlap with the ground, first, and second excited state plotted in blue, orange and green respectively.

the course of a 1μs anneal. We see that the expected transitions do indeed take place and that the system ends having a high overlap with the GS.

We will now go on to look at results for larger system sizes, scaling our graphs in the same two ways as in the previous sections, illustrated in figure 9.1. When we were examining the capacity of the catalyst to remove multiple perturbative crossings, we were able to use the size of the minimum gap to characterise the success of the catalyst which also meant we could use it as the parameter which we were optimising to select our J_{xx} values. Here, however, we are not trying to maximise the gap enhancement but rather create the best diabatic path through the spectrum. As such, the gap size is no longer a useful metric. Instead we will be looking at the enhancement to the final GS fidelity in order to select our J_{xx} values. This is of course much more computationally intensive since it requires us to simulate the closed system dynamics rather than just diagonalising the instantaneous Hamiltonians. Furthermore, introducing only a single coupling within each relevant sub-graph (rather than all-to-all couplings) disrupts the symmetry that allows us to reduce the size of the Hilbert space. As a result, the Hilbert space for a given system size is larger in this setting.

Fortunately, after performing sweeps for the smallest system sizes, we found that the optimal $J_{xx}^{(a)}$ values associated with each of the couplings were approximately equal. As such, we chose to sweep over one global J_{xx} value for the larger systems – greatly reducing the computational demands on the parameter sweeps. Considering the resultant spectra, we have since realised that the optimisation of J_{xx} could be done purely through consideration of the gap spectrum – without the need for dynamic simulations. The optimal spectrum for a diabatic anneal, as we discuss in chapter 8, is one in which the spectral gaps at points where we *want* a transition to occur are minimised. In this setting, these are the gaps that allow the system to cascade up to the highest excited state associated with the local optima so that it can transition back to the GS through the perturbative crossings. Since all these gaps are minimised around the same J_{xx} value, we could optimise J_{xx} through the minimisation of the gap minimum introduced by the catalyst between the instantaneous GS and 1ES. This would likely allow us to access larger systems sizes than we have been able to so far. This, however, we leave for future work.

When examining the capacity for the catalyst to remove multiple perturbative crossings, we were able to make some comments regarding the change in anneal time needed for a single anneal to return the GS with high fidelity by using the gap sizes as a proxy for the expected dynamics. That is, since adiabatic annealing requires $t_a \propto 1/[\Delta E_{01}^{\min}]^2$, we can compare the gap sizes, before and after the introduction of the catalyst, in order to understand the effect on the time scaling. When we are considering *diabatic* annealing however, the gap size is no longer a direct indicator of the annealing time required for a high GS overlap at the end of the anneal. Primarily, our focus in this section will be to see if these diabatic paths can still be created when more perturbative crossings are present – rather than trying to understand the scaling behaviour. We will however still make some comments on this where we can.

Let us begin by examining the formation of the diabatic path as we increase the system size of the tripartite example.

9.2.1 Catalyst scaling with two perturbative crossings

First, let us recap what we observed in section 7.2.2 as we increased the system size of the setting with a single perturbative crossing. For the smallest system sizes, a single new gap minimum formed that could be brought arbitrarily close to zero by increasingly fine tuning of J_{xx}. From around $n = 25$, however, a secondary additional gap minimum begins to form. By $n = 35$, there is a clear manifestation of two new gap minima as a result of the catalyst. Spectra illustrating these three cases can be found in figure 7.8. As for the smaller bipartite instances, introducing the catalyst to the 8-spin *tripartite* example from the previous section created a diabatic path that the system could follow to the ground state – see figures 9.5 and 9.6. We now scale up this example, as illustrated in figure 9.1(a), keeping the same problem parameters, *i.e.*, $W_0 = 1.02$, $W_1 = 1.01$, $W_2 = 1.00$ and $J_{zz} = 5$.

For each system size, we first perform a parameter sweep to identify the J_{xx} value that results in the optimal diabatic path. As discussed in the previous section, we search for this optimal J_{xx} by simulating the closed system dynamics and selecting the value that maximises the final GS overlap. To do this, we want to select a long enough annealing time to ensure that minimal state transfer happens away from the perturbative crossings and the small gaps created by the catalyst. We began with $t_a = 2\mu s$ and used $t_a = 4\mu s$ for the largest systems as the spectral gaps became generally smaller. In practice, what we did was to perform the sweep and if the final GS overlap did not surpass 0.95 for any J_{xx} we would increase the annealing time. (Recall that here we are not trying to optimise the annealing time but rather use the final GS overlap as

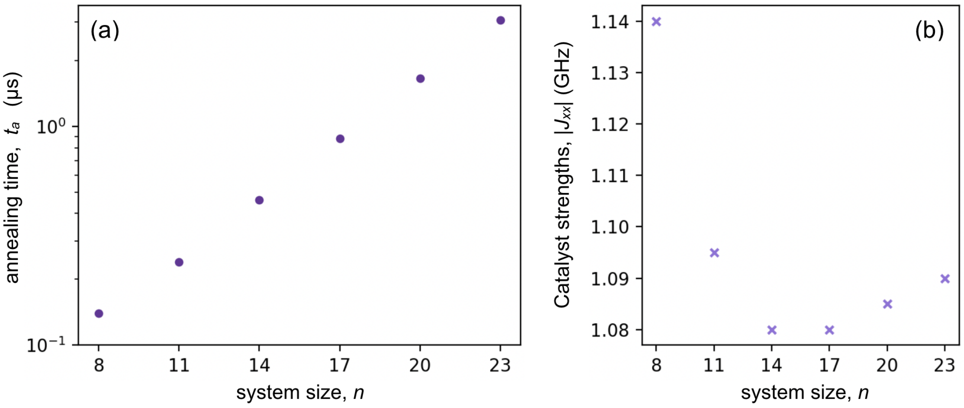

Figure 9.7: Plots showing numerical results for the problem instances scaled as shown in figure 9.1(a) with the problem parameters $W_0 = 1.02$, $W_1 = 1.01$, $W_2 = 1.00$ and $J_{zz} = 5$. The Hamiltonians are introduced with an energy scale of GHz as defined in paragraph 8 of chapter 8. A catalyst is introduced which contains a single coupling within G_1 and another within G_2. Both of these couplings are introduced with the same catalyst strength, J_{xx}, which varies between system sizes. Plot (a) shows the scaling of the annealing time, t_a, needed to reach a GS fidelity of 0.9 and plot (b) shows the J_{xx} values.

an indicator for the formation of a particular structure in the spectrum.) The J_{xx} values obtained this way are plotted in figure 9.7(b).

The resultant annealing spectra associated with the 11-, 17-, and 23-spin systems are presented in figure 9.8. In each case, we see that a diabatic path, similar to what we observed for the 8-spin system, has formed between the instantaneous ground, first and second excited states. We do not, for these system sizes, observe the appearance of any further gap minima in the spectra. It is likely, however, that we have simply not gone up to large enough system sizes. The additional gap minimum in the bipartite setting only began to take shape for system sizes of around $n = 25$ which we have not been able to reach here.

In figure 9.8(a), we plot the annealing time, t_a, needed to reach a final GS overlap of 0.9 when the catalyst is introduced with the optimised J_{xx} values plotted in figure 9.8(b). These results show a clear exponential scaling in t_a for the system sizes which we have examined. Looking at the dynamics of individual anneals, we found that amplitude was being lost to the higher excited states over the portion of the anneal where

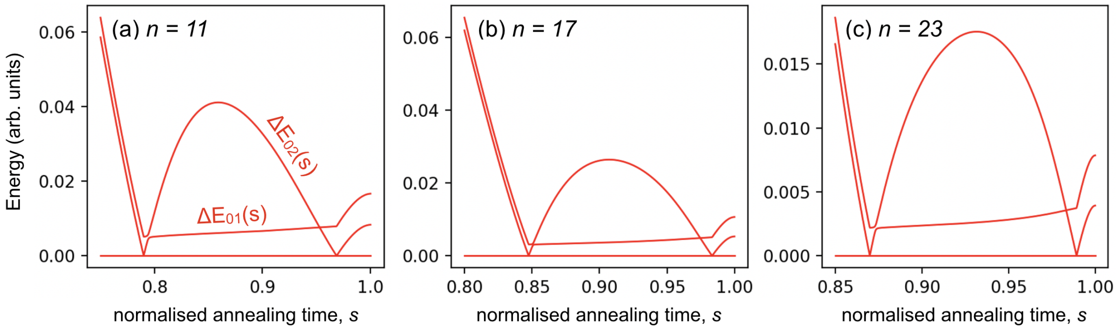

Figure 9.8: Gap spectra showing the optimised diabatic path, as created by a catalyst, for three different system sizes. Each instance is a tripartite graph with $W_0 = 1.2$, $W_1 = 1.1$, $W_2 = 1.0$ and $J_{xx} = 5$. The sub-graph sizes in each case are $n_0 = (n-2)/3$ and $n_1 = n_2 = (n+1)/3$. The catalyst in each case contains two couplings – one between two vertices in G_1 and another between two vertices in G_2. The coupling strengths used in each case are (a) $J_{xx} = 1.095$, (a) $J_{xx} = 1.080$ and (a) $J_{xx} = 1.090$.

the diabatic path had been created – with the majority of the amplitude being lost while the system predominantly occupied the instantaneous 2ES. There is only a very small amount of amplitude transferred between the instantaneous ground, first and second excited states away from the gap minima that make up the diabatic path and a total exchange in amplitude at these points. *i.e.*, the diabatic path is being followed by the portion of the system that remains in the desired sub-space. The bottleneck is therefore the closing of the gap separating this sub-space from the rest of the energy spectrum. An example corresponding to the 11-spin system is shown in figure 9.9.

While the scaling has remained exponential, we can deduce that there has still been a reduction to the annealing time of several orders of magnitude. We are not able to simulate annealing times (for even the smallest system size examined here) that approach the adiabatic limit. However, taking the same approach as we did in chapter 8, we can estimate that it would take on the order of $10^4 \mu s$ to approach this limit for the catalyst-free 14-spin system for instance. The catalyst brings the required annealing time below $1 \mu s$.

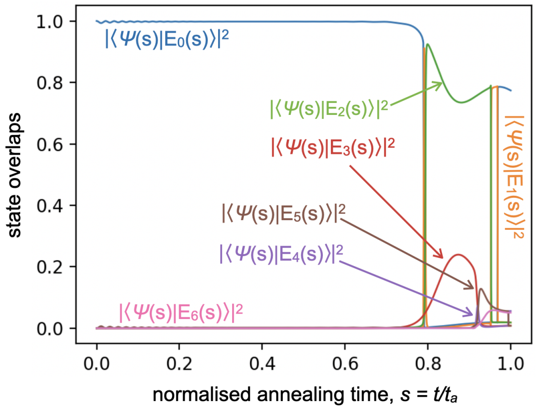

Figure 9.9: Closed system dynamics of an anneal of duration $t_a = 0.12\mu$s for an annealing instance with $G_0 = 3$, $G_1 = 4$, $G_2 = 4$, $W_0 = 1.02$, $W_1 = 1.01$, $W_2 = 1.00$ and $J_{zz} = 5$. The Hamiltonians are introduced with an energy scale of GHz as defined in paragraph 8 of chapter 8.

9.2.2 Two or more perturbative crossings

We end this chapter by briefly examining the formation of the diabatic paths when further perturbative crossings are introduced. The problem instances used in this section have the structure illustrated in figure 9.1(b). The lowest weighted sub-graph is given a weight $W_a = 1.00$ and the largest weighted sub-graph a weight of $W_0 = 1.02$. The remaining sub-graph weights are evenly spaced between these two and we set $J_{zz} = 5$. (Note that the 8-spin system in this section is therefore equivalent to that of the previous section.)

We perform the J_{xx} parameter sweeps as we did in the previous section, using closed system dynamics to optimise one global catalyst strength for each problem instance. As we have already discussed, introducing additional sub-graphs expands the Hilbert space much more rapidly than increasing the size of the sub-graphs. As a result, we were only able to perform these sweeps for two additional systems – the 11-spin system with three perturbative crossings and the 14-spin system with four perturbative crossings.

We find that in both cases, the catalyst can produce spectra that result

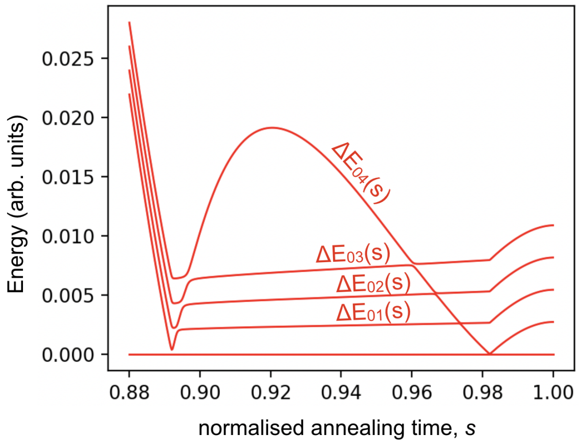

Figure 9.10: Gap spectrum corresponding to a setting with the parameters $G_0 = 2$, $G_1 = 3$, $G_2 = 3$, $G_3 = 3$, $G_4 = 3$, $W_0 = 1.020$, $W_1 = 1.015$, $W_2 = 1.010$, $W_3 = 1.005$, $W_4 = 1.00$ and $J_{zz} = 5$. A catalyst has been introduced that includes a coupling within each of the sub-graphs G_a, $a = 1, 2, 3, 4$. The catalyst is introduced with a catalyst strength $J_{xx} = 0.842$.

in near unity overlap for annealing times $t_a < 1\mu s$. Looking at the annealing spectra, we confirm that this is indeed a result of the formation of a diabatic path that involves a number of excited states equal to the number of perturbative crossings. The spectrum for the 14-spin instance is shown in figure 9.10. The J_{xx} values to produce these spectra are on the same order of magnitude as those from the previous section.

9.3 Discussion

In this chapter we have explored the use of multi-component targeted catalysts in settings where there are multiple perturbative crossings that bottleneck the annealing algorithm.

We first considered a setting similar to one in which the catalyst removed a single perturbative crossing (section 7.1) but where *two* perturbative crossings were present. In this case, the catalyst had two separate components that targeted the two local optima responsible for the perturbative crossings. We observed similar behaviour to that seen in the single perturbative-crossing case, with the catalyst being able to remove both perturbative crossings and the catalyst strengths required to achieve this decreasing with the system size.

We then expanded the problem graph structure to introduce additional perturbative crossings, including a further set of couplings in the catalyst to target each of the new local optima. We examined the effect of this catalyst on systems with up to five perturbative crossings. In each case the system was able to enhance the size of the minimum spectral gap but was not able to fully remove the presence of an AC between the instantaneous ground- and first excited-states. It is worth noting that, since we have not fully optimised the J_{xx} values here, it is likely that we are not seeing the maximum gap enhancement that could be achieved with a catalyst of this structure. However, even if this were not the case, these results are perhaps unsurprising since we also saw that the catalyst was not able to fully remove the presence of an AC in the smallest bi- and tripartite cases either.

Apart from the numerical indications that the gap enhancement would tend to the maximum possible value if we were to increase the sub-graph sizes for any of the n-partite systems, we also have theoretically motivated reasons to believe this would be the case. In section 5.3, we concluded that the reason the gap enhancement from the catalyst reached a maximum, rather than monotonically increasing, was related to certain vector components becoming negative as a result of the introduction of the non-stoquastic catalyst. We argued that the appearance of these negative vector components would be suppressed by the inclusion of additional XX-couplings through effectively introducing frustration into equation 5.16. We confirmed numerically that additional couplings did indeed suppress the presence of the negative vector components. As we increase the system size, more potential couplings become available, such that we can expect the catalyst to be more successful at enhancing the gap minimum. There is no immediate reason to believe that this logic would not also apply in the setting where we are using a multi-part catalyst to remove multiple perturbative crossings.

The problem parameters for the example settings we have used for this investigation bear closer resemblance to those associated with the WGS setting from chapter 7. When we were considering settings with just one perturbative crossing, the performance of the all-to-all catalyst (which we have used here) was largely unaffected by the severity of the gap scaling. We may hope that we would see a similar consistency in performance when multiple perturbative crossings are being targeted – and indeed this appeared to be the case from the preliminary investigations which we made into these settings. However, a logical next step in this investigation may be to properly examine the performance of the catalyst in a setting where the problem parameters result in perturbative crossings with a more severe associated gap scaling.

In section 9.2 of this chapter we examined a setting similar to that in which the catalyst did *not* remove the perturbative crossing (section 7.2.2) but instead created a diabatic path to the ground state. The number of system sizes which we were able to simulate here was limited by the fact that a catalyst containing only a single coupling per sub-graph does not totally preserve the symmetry present prior to the introduction of the catalyst. As a result, the Hilbert space associated with a given system size in this setting is significantly larger than for the same system size in the previously discussed setting. We were therefore not able to answer the all of the questions that we were hoping to in this investigation. However, we can still make some comments on the results we did obtain.

The main question we wanted to address was: how would the results which we saw in the single perturbative crossing case extend to a setting where we introduced a catalyst intended to target multiple perturbative crossings? Since we do not have a complete physical picture explaining the appearance of the additional gap minimum we observed in chapter 7, it was unclear what we should expect to see. What we found was that, for the small system sizes which we have been able to simulate, the catalyst was able to produce a diabatic path through the annealing spectrum that included an additional energy level for each perturbative crossing that the catalyst targeted.

In the bipartite case we saw that the diabatic path started to break down as a result of the appearance of an additional gap minimum that began to take shape for $n > 25$. It seemed likely that something similar would be observed in a setting with multiple perturbative crossings but, again, it was unclear how this would manifest. Unfortunately we were not able to access system the sizes for which we expected these additional effects to arise. A natural next step would be to take advantage of the fact that minimising the new gap minimum appearing between the ground and first excited state seemed to also minimise the other new gap minima that formed the diabatic path. Hence we could use diagonalisation over a small portion of the anneal to optimise J_{xx} rather than simulating the closed system dynamics.

Over the course of this dissertation, where we have identified potential for a catalyst to reduce the demands on the annealing time, we have considered if and how suitable couplings could be identified in practice. In the case with a single perturbative crossing, we noted that an anneal which ran too quickly to remain in the ground-state, but not so fast that the the system became totally de-localised across the state space, was likely to return the state responsible for the perturbative crossing with high fidelity. In the final section of chapter 5, we discussed how knowledge of this state could be used to select appropriate couplings that effectively targeted the associated perturbative crossing – certainly for the MWIS problem and possibly for others as well.

Let us consider what this would look like for a setting with multiple perturbative crossings that involve the problem GS. Assuming the anneal can be run at a speed for which the majority of amplitude exchange happens at the perturbative crossings, we would expect a catalyst free annealing run to return, with high fidelity, the state responsible for the perturbative crossing that manifests as an AC between the instantaneous GS and 1ES. The identification of this state allows for the selection of couplings that target this first perturbative crossing – such that a catalyst can be introduced to remove it. We then expect to be left with a different perturbative crossing manifesting as an AC with the instantaneous GS such that the annealing spectrum effectively leads to the next problem state that would want to identify. In this way, we can imagine recursively identifying and removing the various perturbative crossings present.

Interestingly, while the resultant gap spectra look quite different at each step, the same kind of recursive process could be envisioned for the settings in which we produced the diabatic paths. That is, we find that the spectrum at each step leads us to the next state for which we would want to select targeted XX-couplings in order to construct the diabatic path that eventually leads back to the GS.

Going hand in hand with this is the more general comment that, despite the resultant spectra being so different, the actual dynamics between the two cases are not dissimilar. Let us return to the dynamics of the example setting presented in figure 9.6 in which the system is following a diabatic path through the spectrum. Figure 9.6(a) shows the system transitioning through the instantaneous 1ES, into the 2ES, then back into the 1ES before reaching the GS again. Looking at figure 9.6(b) however, which shows the overlap of the system with the *problem* ground, first and second excited states, we see that system is smoothly evolving towards the problem GS in much the same way as it would in the setting where adiabatic annealing has been facilitated by the removal of the perturbative crossings.

This is perhaps unsurprising since diabatic transitions are generally not indicative of a sharp change in the properties of the system but rather it is the properties of the *instantaneous eigenstates* that are being exchanged. And indeed, the trajectory of the system, observed in figure 9.6(b), is reflected in the make-up of the instantaneous eigenstates that the system transitions between. Looking specifically at the evolution of the instantaneous GS for this example, which is presented in figure 9.5(c), we see that there is a large problem GS contribution prior to the occurrence of the first gap minimum. In some ways then the effect of the catalyst can be thought of, in both cases, as helping the system end in the GS by guiding the *instantaneous* GS towards the *problem GS*. For the settings examined in section 9.2 there is also a rapid re-ordering of the energy levels that happens towards the end of the anneal – as the problem GS amplitude passes up through the instantaneous excited states then back to the instantaneous GS.

9.4 Summary

Overall, the results in this section suggest that the catalysts motivated in chapter 5 can be extended to target multiple perturbative crossings in the same annealing spectrum. Our first indication for this is the fact that the scaling behaviour when two perturbative crossings are present is largely reflective of the single perturbative crossing case examined in chapter 7. Furthermore, when additional perturbative crossings were introduced to the annealing spectrum, we continued to see the same qualitative effect from the extended catalyst Hamiltonians – both in the setting where the catalyst resulted in significant gap enhancement (section 9.1) and the setting where a diabatic annealing path was produced (section 9.2). This, combined with the theoretically motivated arguments presented in section 9.3, lead us to be optimistic with regard to the use of this form of catalyst in more general settings.

We also, in the preceding section, offered a brief description of how a suitable catalyst could be recursively constructed in practice – based off states returned by previous annealing runs. We will expand on this idea of an algorithmic approach to catalyst construction in section 10.2.3 of our discussion chapter.

Chapter 10

Discussion

In this thesis, we have explored the use of catalyst Hamiltonians for the removal of perturbative crossings in quantum annealing. Motivated by results in the literature which suggest that the effect a catalyst has on the annealing spectrum is highly dependent upon the specifics of the problem Hamiltonian, we have taken a targeted approach in which we have sought to incorporate accessible, problem specific, information into the catalyst construction.

Using perturbation theory, we took two analytical approaches to understanding how the effect of different XX-terms relates to the couplings they introduce between the eigenstates of the problem Hamiltonian. We used these findings to motivate the following two options for removing a perturbative crossing between two problem states, $|E_a\rangle$ and $|E_b\rangle$ ($b > a$), associated with two local optima of the problem:

Either

- Introduce XX-terms with the *same* sign as the driver that introduce couplings between $|E_a\rangle$ and low energy states that are close in Hamming distance to $|E_a\rangle$.

Or

- introduce XX-terms with the *opposite* sign to the driver that introduce couplings between $|E_b\rangle$ and low energy states that are close in Hamming distance to $|E_b\rangle$.

The first of these approaches results in a total Hamiltonian that is stoquastic for the entirety of the anneal. In the second case, the total Hamiltonian is non-stoquastic for $s \neq 0, 1$ (where s is the annealing parameter, t/t_a).

Over the course of this dissertation, we explored these ideas numerically. To facilitate our numerical investigations, we introduced structured MWIS instances that allowed us to controllably produce spectra with different perturbative crossing structures and tune the severity of the associated exponential gap scaling. By utilising symmetries present in these systems to reduce the size of the Hilbert space, we were able to obtain data for systems with up to 73 spins. This exploration was the primary focus of this dissertation and we will spend the majority of this final chapter discussing what our main findings were in this area, their implications, and possible avenues for further research.

Outside of this, we found that the theoretical framework we introduced to understand the effect of different XX-coupling choices may be a useful tool for catalyst design more generally. Extending the analysis to general 2-local coupling choices, we found that there was good agreement between the theoretical predictions and the numerical results for all the non-complex catalyst choices. Let us briefly discuss these results before returning to the main focus of this work.

10.1 A perturbative framework for catalyst design

In chapter 6 we examined the effect of general 2-local catalysts (XX, YY, ZZ, XY, XZ and YZ) on an example graph containing a single perturbative crossing, considering the performance of different placements of the various coupling choices. We found that 2-local terms of any almost type were able to enhance the minimum spectral gap if their placement was such that they "targeted" the problem GS – *i.e.*, they introduced couplings between the GS and other low-energy states. While the consistency of this is appealing, it may not be a particularly useful result as far as motivating routes for catalyst design, since being able to target the GS presupposes that we already have the solution to our problem.

2-local catalysts that targeted the *excited* state involved in the perturbative crossing were less consistently helpful. We were able to determine that this related to the different phases that were introduced if Z or Y operators formed part of the couplings. These phases resulted in perturbations of different signs effectively cancelling each other out. Where such couplings *were* able to significantly enhance the gap minimum, this could be understood by considering precisely how the different signs entered the different relevant perturbations. However, it was difficult to draw any helpful conclusions for more general settings.

This perhaps highlights a particular benefit of using XX-couplings when designing targeted catalysts: it is easier to predict the sign with which the couplings will introduce perturbations to the desired states. Furthermore, we feel that the observed agreement between the theory and simulation is an important result in and of itself since it offers further validity to this approach to catalyst design. The only coupling choices for which our theoretical analysis did not agree with the numerical results were those which included a single Y operator and thus had imaginary components. Understanding where this discrepancy comes from and how the analysis could be extended to mitigate this is a potential avenue for future work.

While XX catalysts would appear to be the most promising choice from what we have investigated, we have far from exhausted all the potential forms a catalyst could take. For instance, a catalyst could consist of a mixture of different coupling types and/or higher order interactions. For this latter case, the perturbative analysis could be easily extended to describe the effect of different n-body terms.

Let us now move on to the main focus on this dissertation, which is the use of targeted XX-catalysts.

10.2 Targeted XX-catalysts

The only perturbative crossings that bottleneck adiabatic quantum annealing are those that involve the problem GS. As such, it is on the removal of these types of crossing that we have focused our attention. We showed, in section 5.1, that one option for removing such a perturbative crossing is a catalyst that effectively targets the GS through the inclusion of XX-terms, sharing the same sign as the driver, that introduce couplings between the GS and other low energy states. Introducing such catalysts to a selection of example settings in section 5.2.1, we confirmed that they were indeed able to enhance the gap minimum over the catalyst-free case. Examining the evolution of the instantaneous states in section 5.2.2, we confirmed that this gap enhancement was a result of the intended effect on the perturbed problem states. We found that this gap enhancement could be achieved with varying numbers of couplings being included in the catalyst, so long as the strength with which the catalyst was introduced was sufficiently large.

The other catalyst proposed (in section 5.1) was one that targeted the *excited* state responsible for the perturbative crossing through the inclusion of XX-terms, with the *opposite* sign to the driver, which introduced couplings between this excited and other low-energy states. In this case, the situation was a little more complicated, with the same gap enhancement only achieved if there were no other perturbative crossings present in the annealing spectrum. Otherwise, a small gap-minimum corresponding to a different perturbative crossing would remain.

Furthermore, unlike when we were targeting the GS, the maximum gap enhancement that could be reached depended on the number of couplings included in the catalyst, with more couplings allowing for greater gap enhancement. We attributed this difference in behaviour to additional complexities associated with the non-stoquastic nature of the total Hamiltonian when the XX-couplings were introduced with the opposite sign to the driver. Specifically, in section 5.3, we extended our theoretical analysis from section 5.1 and were able to relate these differences to changes in the signs of the instantaneous GS vector components. (These can only be negative in the non-stoquastic setting.) We found, numerically and analytically, that the appearance of these negative vector components could be generally suppressed by including more XX couplings in the catalyst.

Finally, to further confirm the importance of the coupling placement, we also looked at the effect of introducing XX-terms that were targeting a local optimum *not* responsible for a perturbative crossing with the ground state. As expected, we found that this had little effect on the size of the gap minimum between the instantaneous ground and first excited state.

10.2.1 Catalyst performance

Despite the additional complications present in the non-stoquastic setting, we chose to focus our investigation on catalysts that target the *excited*-state, rather than the ground-state. While a catalyst targeting the GS may more reliably be able to enhance the minimum spectral gap, its construction would require more information of the GS that we would likely have access to. On the other hand, constructing a catalyst utilising information of the local optima in which the algorithm is likely to become trapped is a much more feasible route to catalyst design. We will discuss this further in section 10.2.3.

Numerically investigating the performance of the catalyst with increasing system size, in chapter 7 we continued to find that the maximum gap enhancement the catalyst could achieve depended on the number of couplings included in the catalyst. We *also* found that it was dependent on the severity of the gap scaling associated with the perturbative crossing in the catalyst-free case, with greater gap enhancement being possible for cases where the initial exponential scaling was less severe. For the scalable instances on which we focused our investigation in chapter 7, a catalyst containing a linearly scaling number of couplings was sufficient to remove a perturbative crossing regardless of the severity of the gap scaling. Furthermore, we found that the required catalyst strength decreased with system size.

In the extreme case where only a single coupling was included in the catalyst, the response of the annealing spectrum to the catalyst was much more involved and also differed significantly depending on the severity of the initial gap scaling. Not only was less gap enhancement observed, but we also saw the appearance of additional gap minima in the spectrum. Once again, these additional effects were associated with sign changes in the instantaneous GS vector components. We will discuss their manifestations and the consequences thereof in section 10.2.2.

The question of whether or not a linearly increasing number of XX-terms in the catalyst will be sufficient to remove a perturbative crossing in more general settings requires further investigation. Firstly, while Hilbert-space reductions have allowed us to access larger system sizes than we would have been able to otherwise, we have no guarantee that what we are seeing is reflective of the large-system limit. (Note that, due to this catalyst disrupting the symmetries in the system, we were only able to obtain results for system sizes up to $n = 31$ spins.) If we consider the effects of the *single*-coupling catalyst in the WGS setting, the additional gap minima only began to manifest at around $n = 20$. As such, we cannot rule out the possibility of the closings gaps, observed for the single-coupling catalyst, manifesting for larger system sizes when the number of couplings included in the catalyst scales linearly or even quadratically with the system size.

Furthermore, even if the observed behaviour persists to arbitrarily large systems, we have by no means considered the performance of the catalyst under general parameter settings. It is always possible to alter the parameters to further heighten the severity of the gap scaling which, as we have seen, can introduce unwanted additional effects from the catalyst. Not only this, but there are of course many different ways in which we could choose to scale the graph structure. Indeed, in our brief look at an alternative graph scaling, which naturally resulted in a more severe gap scaling in the catalyst-free setting, we saw the appearance of additional gap minima even when a quadratically scaling number of couplings was used. It is potentially relevant, however, that the new gap minima that formed in this setting did seem to facilitate a diabatic path to the ground state. As such, the catalyst could still be considered to be removing the bottleneck associated with the perturbative crossing. We will return to this point in section 10.2.2.

We suspect that an understanding of the number of catalyst terms needed to remove a perturbative crossing in different settings may lie in extending the theoretical work from section 5.1.2, considering more rigorously the conditions that need to be met to remove a perturbative crossing and how these relate its properties. Combining such findings with an investigation into the typical properties of the perturbative crossings which we might expect to manifest for a given problem type will help us understand what demands we can expect to be placed on the catalyst.

Finally, in section 9.1 we explored the capacity of these catalysts to remove multiple perturbative crossings from the annealing spectrum, with separate sets of XX-terms targeting the different local optima responsible for the different perturbative crossings. We found that constructing a catalyst in this way was indeed able to address multiple perturbative crossings as intended. Obtaining clear scaling results was more difficult here. This is because introducing more perturbative crossings required us to break up the problem into more sub-graphs which limits the extent to which we can reduce the size of the Hilbert space. Nevertheless, our numerical results, combined with our theoretical understanding of the catalyst, led us to believe that catalysts constructed in this way are likely to have the capacity to fully remove all the perturbative crossings present in the annealing spectrum, provided sufficient XX-terms are included.

10.2.2 New gap minima

In chapter 5 we noted that the appearance of negative GS vector components, owing to the non-stoquasticity of the catalyst, limited the maximum gap enhancement that could be achieved. As we began looking at larger systems and settings with more severe exponential gap scaling in chapter 7, we identified instances for which the presence of negative vector components resulted in significantly more involved effects on the annealing spectra from the catalyst. In particular, we observed the introduction of additional gap minima to the spectrum which manifested in a number of different ways.

While the appearance of new gap minima is of course a hindrance for adiabatic QA, we noted that some of the resultant spectra appeared amenable to *diabatic* annealing. For instance, when we introduced the single-coupling catalyst to the WGS setting, we observed a response from the spectrum reminiscent of that observed in [100]. That is, we saw the appearance of a correlated double AC between the ground and first excited state *in addition* to a lifting of the gap minimum associated with the perturbative crossing. Since the system can transition into the 1ES at the first of the two ACs before transitioning back into the GS at the second, these gap minima likely do not pose a bottleneck to the algorithm.

The catalyst resulted in annealing spectra with a similar structure when introduced to settings associated with our alternative graph scaling – see figure 7.13. In this case, all possible relevant couplings were included in the catalyst. The fact that this was the only manifestation of new gap minima in the spectrum when multiple-coupling catalysts were used, and the result appeared to facilitate a diabatic anneal, seems to us a promising result. However, a better understanding of the formation of these gap minima is still required for us to understand if we can rely on their manifesting in a way that facilitates DQA.

We also observed the formation of a diabatic path in the smaller SGS instances when a single-term catalyst was used. In this case, the enhancement to the gap minimum associated with the perturbative crossing was minimal. Instead a new gap minimum appeared in addition to the still present perturbative crossing. Interestingly, we appeared to be able to set this new gap minimum arbitrarily close to zero through increasingly fine tuning of J_{xx}. While we were able to link this closing gap to changes in the signs of the GS vector components, it is not clear to us what the physical mechanism behind their appearance is. Over the course of the project, we took a number of approaches to better understand this behaviour. These included finite-size-scaling analysis and searching for an emergent symmetry which could have resulted in a vanishing gap between two states in different subspaces. As of yet, however, the exact physics behind the manifestation of this gap minimum is still unclear to us.

While the introduction of this single additional gap minimum created a diabatic path, we found that this breaks down for larger system sizes as a further gap minimum begins to manifest. Nevertheless, given that this was an unexpected and, to our knowledge, novel means of manipulating the spectrum, we were interested in understanding a potential extension to this effect. Specifically, we wanted to see how this effect would manifest if were were to introduce a catalyst targeting multiple perturbative crossings – as we did for the case where the catalyst *did* remove the perturbative crossings. Doing this (in section 9.2), we found that a complex diabatic path through the spectrum was produced which the system could follow to reach a high problem-GS fidelity for significantly shorter annealing times than in the catalyst free case. Without a more thorough theoretical understanding of how these spectra form, however, it is not clear if this these effects will prove useful outside of our toy models.

What we *can* say, however, is that the smaller SGS instances, with a single-term catalyst, may be a useful tool in studying DQA owing to a simple diabatic path with a tunable gap minimum. Furthermore, these spectra manifest for systems as small as $n = 5$ and would only require a single XX-coupler to implement. As such, should XX-couplers become available, these systems may provide a suitable test-bed for experimental, as well as numerical, investigations of DQA.

To conclude this section, we note that while there are indeed cases in which the appearance of additional gap minima facilitate DQA, any attempts to utilise these more complex effects from the catalyst will require much more care and consideration than simple gap enhancement. In chapter 8 of this thesis, we did in fact perform a numerical study on the robustness of DQA using the aforementioned small instances of the SGS setting. In that study, we observed a trade-off between the precision in J_{xx} and the annealing time, t_a, that is needed to find the GS with high fidelity. By applying the theory of LZ-transitions we determined that similar trade-offs are likely to be observed for any setting in which there are some parameters that control the gap size at points where we want diabatic transitions to occur.

10.2.3 An algorithmic approach

Throughout this dissertation, we have discussed the idea of being able to construct an appropriate catalyst based on *accessible problem information* obtained from local optima that the algorithm is likely to become trapped in. More specifically, if we assume that the only exponentially closing gaps in the annealing spectrum are those associated with the perturbative crossings, then a polynomial-time anneal is likely to return the eigenstate of H_p associated with the perturbative crossing responsible for the gap minimum between the instantaneous ground and first excited states. One would then want to select XX-terms that create couplings between this excited state and other low energy states and introduce them with the opposite sign to the driver.

Assuming this can be done successfully, we will either be left with a clear route to the ground state or, perhaps, the presence of further perturbative crossings which could be removed with the same procedure.

That is, QA could be run with the inclusion of the constructed catalyst such that the next relevant eigenstate of H_p can be obtained. This then allows for the construction of an updated catalyst Hamiltonian that includes couplings that target this newly identified excited state. In theory, this process could be repeated until no more perturbative crossings remain such that a polynomial-time anneal is able to obtain the GS with high fidelity.

There are a number of open questions that would need to be addressed in order to understand if and how such an algorithmic approach could be viable – some of which we have discussed already in the final section of chapter 5. First, there must be a way of selecting appropriate XX-couplings. For the MWIS problem this is straightforward. One simply wants to couple vertices within the relevant local optimum to each other – thereby coupling the corresponding excited state to states corresponding to other independent sets. For other problem types, one would either want to determine some similar rule-set for XX-term selection, or one could simply perform some kind of a search for suitable XX-couplings. Generally, it is not difficult to determine the quality of a potential solution given the problem's cost function. Furthermore there are only a quadratically scaling number of potential XX-placements. Thus, it may be possible to evaluate the cost function associated with the optimisation problem on each potential solution associated with the $n(n-1)/2$ states that we could couple the relevant excited state to.

There are then some somewhat more involved questions. For instance: is it likely, by selecting XX-terms based solely on the states to which they couple the relevant excited state, that we will avoid introducing couplings that alter the perturbations to the problem states in less favourable ways? For instance, will we avoid also coupling the *GS* to a comparable number of low energy states, such that the presence of the perturbative crossing is likely to be unchanged? Another open question that we have already discussed is: how many XX-terms must be included in the catalyst? This then leads to the related question of: is it guaranteed, or at least likely, that there is a sufficient number of suitable XX-placements to reach this target? While we have offered some tentative answers to some of these questions throughout this work, a far more rigorous understanding is still required.

There is then also the question of how the appearance of additional gap minima, which we observed in some cases, would impact the algorithmic construction of the catalyst. From the point of view of the proposed procedure, we are not interested in (and in fact we would not have knowledge of) the path which the system has taken to the final state it ends up in. In many cases, the appearance of additional gap minima created a diabatic path that led to the same state that removing the targeted AC would have led the system to. In such cases, the proposed algorithm may be largely unaffected. However, as we have already noted, introducing diabaticity comes with its own set of complications that would need to be taken into account. Furthermore, given that we have seen additional gap minima manifest in a number of different ways, it is likely that our investigations have not accounted for all the subtleties that would need to be considered. All in all, a clearer picture of how and when these additional effects manifest is still required.

Something we have addressed minimally in this work is the question of catalyst strength selection. With regards to the settings in which the catalyst simply enhanced the size of the minimum gap, we found that the catalyst strength required to maximise this gap enhancement decreased with system size. Assuming this scaling persists, this suggests that we will not have to contend with impractically large catalyst strengths being required for larger problems. Beyond this, we have also observed that the gap enhancement is relatively insensitive to the precise value of the catalyst strength when sufficient couplings are included in the catalyst. We have taken only a preliminary look at this but so far, for the larger systems, we have found that simply introducing the catalyst terms with the same magnitude as the driver fields yields optimal results. It is therefore possible that a careful selection of J_{xx} values need not be part of the algorithmic procedure for catalyst construction.

The presence of closing gaps complicates this somewhat since the diabatic paths that we observed were only present for particular ranges of J_{xx}. Once again, a better characterisation of how they form is required. Specifically, an understanding is needed of how the usable range of J_{xx} changes with system size and how this relates to other relevant energy scales of J_{xx} – for instance, the catalyst strength required to remove the original perturbative crossing and/or the implementable coupling strengths on the hardware. There is also the question of how an appropriate annealing time would be selected. In the adiabatic case a longer annealing time will always lead to a better GS fidelity. As such, one can simply select the longest time that is compatible with other constraints – *e.g.* the coherence time of the hardware. When it comes to DQA, however, we found that longer anneal times may place greater demands on the precision needed in J_{xx}.

Finally, there are considerations that relate to implementation on actual hardware. The first, obvious, point to make is that, as of yet, XX-couplings have not been properly implemented on annealing hardware. However, assuming that such couplings will at some point be accessible, there are also the limitations of graph connectivity to consider. In order for the proposed catalysts to be implementable, there must be a sufficient number of appropriate XX-coupling placements that are implementable on a given hardware graph.

10.3 Conclusion

All of the questions we have posed in this chapter are likely to be complicated problems to solve in their own right. If these issues could be tackled, however, the implication would be a quantum annealing algorithm where the number of calls to the annealer scales linearly with the number of local optima, and where the length of each annealing run is polynomial in the problem size. Many difficult problems will have an exponential number of local optima meaning that this would by no means be a polynomial time algorithm for NP-hard problems. However, it would reduce the demand on hardware coherence times and could potentially help achieve a quantitative speed-up over classical algorithms.

As it stands, we have introduced a potential route to mitigating a key bottleneck of the quantum annealing algorithm. Crucially, while we have used specific instances of the MWIS problem to facilitate our investigation, the physics behind the proposed catalysts hinges on structure in the problem Hamiltonian that is relevant to any problem setting. As such, we expect the conclusions reached in this thesis to be applicable to the quantum annealing of general Ising Hamiltonians.

Appendix A

Normalisation factor derivation

Here, we derive the normalisation factor, F, introduced in section 3.4. The aim behind this normalisation factor is to ensure a constant energy scale for a given system size and also to maintain a consistent relationship between the energy scales of the driver and the problem Hamiltonians with *increasing* system size. We begin with the following expression.

$$F = \frac{E_{2^n-1}(0) - E_0(0)}{E_{2^n-1} - E_0}. \tag{A.1}$$

The energies in the numerator of equation A.1 are the highest excited state and the ground state of the driver Hamiltonian,

$$\hat{H}_d = -\sum_{i=1}^{n} \hat{\sigma}_i^x. \tag{A.2}$$

Working in the X-basis, it is clear that the GS of this Hamiltonian is the state with all spins pointed up and that the highest excited state is the state with all spins pointed down. The energies of these states are straightforwardly

$$E_0(0) = \langle E_0(0)|H_d|E_0(0)\rangle = -n,$$

and

$$E_{2^n-1}(0) = \langle E_{2^n-1}(0)|H_d|E_{2^n-1}(0)\rangle = n,$$

such that

$$E_{2^n-1}(0) - E_0(0) = 2n = 2\sum_{a=0}^{k-1} n_a, \tag{A.3}$$

where n_a are the sizes of the different sub-graphs that make up the problem.

Now let us move on to the energies in the denominator of equation A.1. These are the highest excited state and the ground state of the problem Hamiltonian. For the problems we are considering, H_p can be written

$$\hat{H}_p = \sum_{a=0}^{k-1}\left[\left(\underbrace{J_{zz}\sum_{b\neq a} n_b}_{d_i \forall i\in G_a} - 2\underbrace{\frac{W_a}{n_a}}_{w_i \forall i\in G_a}\right)\sum_{i\in G_a}\hat{\sigma}_i^z + J_{zz}\sum_{i\in G_a}\sum_{b>a}\sum_{j\in G_b}\hat{\sigma}_i^z\hat{\sigma}_j^z\right]. \tag{A.4}$$

To recap the notation we are using: We have k sub-graphs. Each sub-graph, G_a ($0 \leq a < k$), contains n_a vertices and has a total weight of W_a assigned to it which is split evenly between the vertices. J_{zz} is an edge penalty that is introduced between vertices in different sub-graphs.

As already discussed in section 3.4, the weights and edge penalty are chosen so that the GS of this Hamiltonian is the state with all spins pointed down (in the Z-basis) other than those in G_0 – which are pointed up. Using this knowledge, let us obtain an expression for E_0. To do this, let us first split off the $a = 0$ term from the outermost sum in equation A.4 to obtain

$$\begin{aligned}\langle E_0|\left[\left(J_{zz}\sum_{b>0} n_b - 2\frac{W_0}{n_0}\right)\sum_{i\in G_0}\hat{\sigma}_i^z + J_{zz}\sum_{i\in G_0}\sum_{b>0}\sum_{j\in G_b}\hat{\sigma}_i^z\hat{\sigma}_j^z\right]|E_0\rangle \\ = \left(J_{zz}\sum_{b>0} n_b - 2\frac{W_0}{n_0}\right)n_0 - J_{zz}n_0\sum_{b>0} n_b = -2W_0,\end{aligned} \tag{A.5}$$

where we have used fact that all the spins in G_0 are pointed up and are misaligned with the spins in all the other sub-graphs G_b, $b > 0$.

Now let us consider the other terms. For $a > 0$, we can write each term in the outer sum of equation A.4 as

$$\begin{aligned}
&\langle E_0| \left[\left(J_{zz} \sum_{b \neq a} n_b - 2\frac{W_a}{n_a} \right) \sum_{i \in G_a} \hat{\sigma}_i^z + J_{zz} \sum_{i \in G_a} \sum_{b > a} \sum_{j \in G_b} \hat{\sigma}_i^z \hat{\sigma}_j^z \right] |E_0\rangle \\
&\quad = - \left(J_{zz} \sum_{b \neq a} n_b - 2\frac{W_a}{n_a} \right) n_a + J_{zz} n_a \sum_{b > a} n_b = 2W_a - J_{zz} n_a \sum_{b < a} n_b, \qquad \text{(A.6)}
\end{aligned}$$

where we have used the fact that all the spins in G_a ($a \neq 0$) are pointed down and aligned with all spins in G_b ($b \neq 0$). Putting together equations A.5 and A.6 we can write the final energy as

$$E_0 = \langle E_0|H_p|E_0\rangle = \sum_{a=1}^{k-1} \left(2W_a - J_{zz} n_a \sum_{b < a} n_b \right) - 2W_0. \qquad \text{(A.7)}$$

Now we need to obtain an expression for E_{n^2-1}. From the discussion in section 3.2, we know that the highest excited state will be that with all the spins in every sub-graph pointed up. The energy associated with this state is

$$\begin{aligned}
&E_{2^n-1} = \langle E_{2^n-1}|H_p|E_{2^n-1}\rangle = \\
&\langle E_{2^n-1}| \sum_{a=0}^{k-1} \left[\left(J_{zz} \sum_{b \neq a} n_b - 2\frac{W_a}{n_a} \right) \sum_{i \in G_a} \hat{\sigma}_i^z + J_{zz} \sum_{i \in G_a} \sum_{b > a} \sum_{j \in G_b} \hat{\sigma}_i^z \hat{\sigma}_j^z \right] |E_{2^n-1}\rangle \\
&\quad = \sum_{a=0}^{k-1} \left[\left(J_{zz} \sum_{b \neq a} n_b - 2\frac{W_a}{n_a} \right) n_a + J_{zz} n_a \sum_{b > a} n_b \right] \\
&\quad = \sum_{a=0}^{k-1} \left[-2W_a + J_{zz} n_a \left(\sum_{b < a} n_b + 2 \sum_{b > a} n_b \right) \right]. \qquad \text{(A.8)}
\end{aligned}$$

Putting equations and together, we obtain the denominator in equation

A.1 as

$$\begin{aligned}E_{2^n-1} - E_0 &= \sum_{a=0}^{k-1} \left[-2W_a + J_{zz} n_a \left(\sum_{0 \leq b < a} n_b + 2 \sum_{b > a} n_b \right) \right] \\ &\quad + \sum_{a=1}^{k-1} \left[-2W_a + J_{zz} n_a \sum_{b<a} n_b \right] + 2W_0 \\ &= 2J_{zz} n_a \sum_{b>0} n_b + \sum_{a=1}^{k-1} \left[-4W_a + 2J_{zz} n_a \sum_{b \neq a} n_b \right] . \qquad \text{(A.9)}\end{aligned}$$

Re-indexing the first sum we can incorporate it into the second, such that our final expression is

$$E_{2^n-1} - E_0 = \sum_{a=1}^{k-1} \left[2J_{zz} n_a n_0 + 2J_{zz} n_a \sum_{b \neq a} n_b - 4Wa \right] . \qquad \text{(A.10)}$$

Finally, we can put the numerator, from equation A.3, and the denominator, from equation A.10, together to obtain the final normalisation factor:

$$F = \frac{\sum_{a=0}^{k-1} n_a}{\sum_{a=1}^{k-1} \left[J_{zz} n_a n_0 + J_{zz} n_a \sum_{b \neq a} n_b - 2Wa \right]} . \qquad \text{(A.11)}$$

Appendix B

Normalised parameters

The problem parameters quoted throughout the main body of this dissertation are the *un*-normalised weights (W_a and w_i) and edge penalities (J_{zz}). Here we quote the associated *normalised* parameters as calculated as outlined in section 3.4 and appendix A.

We do not list the full *Hamiltonian* parameters that would be obtained substituting the normalised weights and edge penalty into equation 3.9. However, we briefly note that the local fields, h_z^i, are all on the order of unity – *i.e.*, comparable to the driver field strengths. The coupling strength that enters the Hamiltonian is equivalent to the normalised edge penalty, J'_{zz}.

B.1 Chapter 3

Figure 3.3:

$W_0 = 1.04$, $W_1 = 1.02$, $W_2 = 1.00$ and $J_{zz} = 2$.

n_0, n_1, n_2	$w'_i,\ i \in G_0$	$w'_i,\ i \in G_1$	$w'_i,\ i \in G_2$	J'_{zz}
3, 2, 2	0.04047	0.05954	0.05837	0.2335
3, 4, 2	0.03121	0.02296	0.04502	0.1801
3, 2, 4	0.03121	0.04592	0.02251	0.1801

Figure 3.4:

$W_0 = 1.01$, $W_1 = 1.00$ and $J_{zz} = 5.33$.

n_0, n_1	$w'_i,\, i \in G_0$	$w'_i,\, i \in G_1$	J'_{zz}
3, 2	0.02717	0.04035	0.4301
37, 36	0.0001404	0.0001428	0.02741
2, 3	0.04075	0.02690	0.4301
8, 9	0.002804	0.002467	0.1184

$W_0 = 1.37$, $W_1 = 1.00$ and $J_{zz} = 37.5$.

n_0, n_1	$w'_i,\, i \in G_0$	$w'_i,\, i \in G_1$	J'_{zz}
2, 3	0.007645	0.003720	0.4185
34, 35	3.115×10^{-5}	2.209×10^{-5}	0.02899

B.2 Chapter 5

Figure 5.2:

$W_0 = 1.70$, $W_1 = 1.35$, $W_2 = 1.00$ and $J_{zz} = 35$.

n_0, n_1, n_2	$w'_i,\, i \in G_0$	$w'_i,\, i \in G_1$	$w'_i,\, i \in G_2$	J'_{zz}
3, 4, 2	0.002809	0.001673	0.002479	0.1735

Figures 5.3, 5.5 and 5.6:

$W_0 = 1.70$, $W_1 = 1.35$, $W_2 = 1.00$ and $J_{zz} = 35$.

n_0, n_1, n_2	$w'_i,\, i \in G_0$	$w'_i,\, i \in G_1$	$w'_i,\, i \in G_2$	J'_{zz}
2, 3, 4	0.004214	0.002231	0.001239	0.1735

Figure 5.7:

$W_0 = 1.37$, $W_1 = 1.00$ and $J_{zz} = 37.5$.

n_0, n_1	$w'_i,\, i \in G_0$	$w'_i,\, i \in G_1$	J'_{zz}
3, 4	0.003560	0.001949	0.2923

Figure 5.8:

$W_0 = 1.37$, $W_1 = 1.00$ and $J_{zz} = 37.5$.

n_0, n_1	$w'_i,\ i \in G_0$	$w'_i,\ i \in G_1$	J'_{zz}
4, 5	0.002058	0.001202	0.2253

B.3 Chapter 6

Figures 6.1 and 6.2:

$W_0 = 1.70$, $W_1 = 1.35$, $W_2 = 1.00$ and $J_{zz} = 35$.

n_0, n_1, n_2	$w'_i,\ i \in G_0$	$w'_i,\ i \in G_1$	$w'_i,\ i \in G_2$	J'_{zz}
4, 5, 3	0.001552	0.0009862	0.001218	0.1278

B.4 Chapter 7

Figures 7.1, 7.2, 7.3, 7.5, 7.6, 7.9 and 7.11:

$W_0 = 1.37$, $W_1 = 1.00$ and $J_{zz} = 37.5$.

n_0, n_1	$w'_i,\ i \in G_0$	$w'_i,\ i \in G_1$	J'_{zz}
2, 3	0.007645	0.003720	0.4185
5, 6	0.001341	0.0008155	0.1835
10, 11	0.0003488	0.0002315	0.09548
15, 16	0.0001573	0.0001077	0.06459
21, 22	8.096×10^{-5}	5.641×10^{-5}	0.04654

Figures 7.2, 7.3, 7.4, 7.7, 7.8, 7.10 and 7.12:

$W_0 = 1.01$, $W_1 = 1.00$ and $J_{zz} = 5.33$.

n_0, n_1	$w'_i,\ i \in G_0$	$w'_i,\ i \in G_1$	J'_{zz}
2, 3	0.04075	0.02690	0.4301
7, 8	0.003638	0.003151	0.1344
12, 13	0.001267	0.001158	0.08022
15, 16	0.0008165	0.0007579	0.06463
17, 18	0.0006379	0.0005965	0.05722
21, 22	0.0004201	0.000397	0.04656

Figure 7.13:

$W_0 = 1.01$, $W_1 = 1.00$ and $J_{zz} = 5.33$.

n_0, n_1	$w'_i,\ i \in G_0$	$w'_i,\ i \in G_1$	J'_{zz}
12, 18	0.001098	0.0007245	0.06950

B.5 Chapter 8

Figures 8.1, 8.2, 8.3, 8.4 and 8.5:

$W_0 = 1.01$, $W_1 = 1.00$ and $J_{zz} = 5.33$.

n_0, n_1	$w'_i,\ i \in G_0$	$w'_i,\ i \in G_1$	J'_{zz}
2, 3	0.04075	0.02690	0.4301
4, 5	0.01076	0.008523	0.2271
12, 13	0.001267	0.001158	0.08022

Figure 8.7:

$W_0 = 1.010$, $W_1 = 1.005$, $W_2 = 1.000$ and $J_{zz} = 5.33$.

n_0, n_1, n_2	$w'_i,\ i \in G_0$	$w'_i,\ i \in G_1$	$w'_i,\ i \in G_2$	J'_{zz}
2, 3, 3	0.01838	0.01219	0.01213	0.1940

B.6 Chapter 9

Figures 9.2 and 9.3:

$W_0 = 1.6$, $W_1 = 1.3$, $W_2 = 1.0$ and $J_{zz} = 35$.

n_0, n_1, n_2	$w'_i,\ i \in G_0$	$w'_i,\ i \in G_1$	$w'_i,\ i \in G_2$	J'_{zz}
3, 4, 4	0.002099	0.001279	0.0009838	0.1377
14, 15, 15	0.0001114	8.447×10^{-5}	6.498×10^{-5}	0.03411

Figure 9.4:

$W_0 = 1.3$, $W_{a>0} = 1 - 0.01(a-1)$ and $J_{zz} = 35$.

$n_0 = 2$ and $n_{a>0} = 3$.

k = total number of sub-graphs.

k	$w'_i,\ i \in G_0$	$w'_i,\ i \in G_1$	$w'_i,\ i \in G_2$	$w'_i,\ i \in G_3$	$w'_i,\ i \in G_4$	$w'_i,\ i \in G_5$	J'_{zz}
3	0.001705	0.0009835	0.0009737	-	-	-	0.1377
4	0.001107	0.0006384	0.0006320	0.0006256	-	-	0.08938
5	0.0008174	0.0004716	0.0004669	0.0004622	0.0004575	-	0.06602
6	0.0006476	0.0003736	0.0003699	0.0003661	0.0003624	0.0003587	0.05231

Figures 9.5, 9.6, 9.7 and 9.8:

$W_0 = 1.02$, $W_1 = 1.01$, $W_2 = 1.00$ and $J_{zz} = 5$.

n_0, n_1, n_2	$w'_i,\ i \in G_0$	$w'_i,\ i \in G_1$	$w'_i,\ i \in G_2$	J'_{zz}
2, 3, 3	0.01981	0.01308	0.01295	0.1942
3, 4, 4	0.009445	0.007014	0.006945	0.1389
5, 6, 6	0.003628	0.002993	0.002964	0.0889
7, 8, 8	0.001909	0.001654	0.001637	0.0655

Figure 9.10:

$W_0 = 1.020$, $W_1 = 1.015$, $W_2 = 1.010$, $W_3 = 1.005$, $W_4 = 1.000$ and $J_{zz} = 5$.

$n_0 = 2$, $n_1 = 3$, $n_2 = 3$, $n_3 = 3$ and $n_4 = 3$

k	$w'_i,\ i \in G_0$	$w'_i,\ i \in G_1$	$w'_i,\ i \in G_2$	$w'_i,\ i \in G_3$	$w'_i,\ i \in G_4$	J'_{zz}
5	0.009249	0.0061361	0.0061058	0.00607560	0.006045	0.09068

Appendix C

Perturbative corrections to the problem states

Here we list the first three perturbative corrections to the low energy, non-degenerate problem states associated with the local optima:

$$E_i^{(1)} = \langle E_i | H_d | E_i \rangle = 0,$$

$$|E_i^{(1)}\rangle = \sum_{j \neq i} \frac{\langle E_j | H_d | E_i \rangle}{(E_i - E_j)} |E_j\rangle ,$$

$$E_i^{(2)} = \sum_{j \neq i} \frac{| \langle E_j | H_d | E_i \rangle |^2}{(E_i - E_j)},$$

$$|E_i^{(2)}\rangle = \sum_{j \neq i} \sum_{k \neq i} \frac{\langle E_k | H_d | E_j \rangle \langle E_j | H_d | E_i \rangle}{(E_i - E_k)(E_i - E_j)} |E_k\rangle - \frac{1}{2} \sum_{j \neq i} \frac{\langle E_j | H_d | E_i \rangle}{(E_i - E_j)} |E_i\rangle ,$$

$$E_i^{(3)} = \sum_{k \neq i} \sum_{j \neq i} \frac{\langle E_i | H_d | E_k \rangle \langle E_k | H_d | E_j \rangle \langle E_j | H_d | E_i \rangle}{(E_i - E_j)(E_i - E_k)},$$

$$
\begin{aligned}
|E_i^{(3)}\rangle = \sum_{l\neq i}\sum_{k\neq i}\sum_{j\neq i} & \frac{\langle E_l|H_d|E_k\rangle \, \langle E_k|H_d|E_j\rangle \, \langle E_j|H_d|E_i\rangle}{(E_i - E_j)(E_i - E_k)(E_i - E_l)} \, |E_l\rangle \\
& - \sum_{k\neq i}\sum_{j\neq i} \frac{\langle E_k|H_d|E_j\rangle \, |\, \langle E_j|H_d|E_i\rangle \, |^2}{(E_i - E_j)(E_i - E_k)^2} \, |E_k\rangle \\
& \qquad - \sum_{k\neq i}\sum_{j\neq i} \frac{\langle E_i|H_d|E_k\rangle \, \langle E_k|H_d|E_j\rangle \, \langle E_j|H_d|E_i\rangle}{(E_i - E_j)(E_i - E_k)^2} \, |E_i\rangle \, .
\end{aligned}
$$

Except for in the first order energy correction, where we have explicitly included the only term in the correction despite it vanishing, we have omitted any terms that equal zero or that we have found to be cancelled by terms from higher order corrections.

Bibliography

[1] N. Feinstein, L. Fry-Bouriaux, S. Bose, and P. A. Warburton, *Effects of xx catalysts on quantum annealing spectra with perturbative crossings*, Phys. Rev. A, 110 (2024), doi:10.1103/PhysRevA.110.042609. Article number: 042609.

[2] N. Feinstein, I. Shalashilin, S. Bose, and P. Warburton, *Robustness of diabatic enhancement in quantum annealing*, Quantum Science and Technology, 10 (2025), doi:10.1088/2058-9565/adab14. Article number: 025011.

[3] 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 (2016), doi:10.1088/1367-2630/aa5079. Article number: 123034.

[4] B. Apolloni and D. D. Falco, *Quantum Stochastic Optimization*, Stochastic Processes and their Applications, 33 (1989), pp. 233–244, doi:https://doi.org/10.1016/0304-4149(89)90040-9.

[5] 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 (1994), pp. 343–348, doi:10.1016/0009-2614(94)00117-0.

[6] T. Kadowaki and H. Nishimori, *Quantum annealing in the transverse Ising model*, Phys. Rev. E, 58 (1998), pp. 5355–5363, doi:10.1103/PhysRevE.58.5355.

[7] ÁLVARO RUBIO-GARCÍA, J. J. GARCÍA-RIPOLL, AND D. PORRAS, *Portfolio optimization with discrete simulated annealing*, 2022. https://arxiv.org/abs/2210.00807.

[8] E. GABBASSOV, *Transit facility allocation: Hybrid quantum-classical optimization*, PLOS ONE, 17 (2022), doi:10.1371/journal.pone.0274632. Article identifier: e0274632.

[9] A. WITT, C. KÖRBER, A. KIRSTÄDTER, AND T. LUU, *Tactile network resource allocation enabled by quantum annealing based on ilp modeling*, 2022. https://arxiv.org/abs/2212.07854.

[10] 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. https://arxiv.org/abs/2212.01578.

[11] J. F. A. SALES AND R. A. P. ARAOS, *Adiabatic quantum computing impact on transport optimization in the last-mile scenario*, Front. Comput. Sci., 15 (2023), doi:10.3389/fcomp.2023.1294564.

[12] C. H. BENNETT, E. BERNSTEIN, G. BRASSARD, AND U. VAZIRANI, *Strengths and weaknesses of quantum computing*, SIAM Journal on Computing, 26 (1997), pp. 1510–1523, doi:10.1137/s0097539796300933.

[13] D. S. ABRAMS AND S. LLOYD, *Nonlinear quantum mechanics implies polynomial-time solution for NP-complete and # P problems*, Phys. Rev. Lett., 81 (1998), pp. 3992–3995, doi:10.1103/physrevlett.81.3992.

[14] T. F. RØNNOW, Z. WANG, J. JOB, S. BOIXO, S. V. ISAKOV, D. WECKER, J. M. MARTINIS, D. A. LIDAR, AND M. TROYER, *Defining and detecting quantum speedup*, Science, 345 (2014), pp. 420–424, doi:10.1126/science.1252319.

[15] R. AU-YEUNG, N. CHANCELLOR, AND P. HALFFMANN, *NP-hard but no longer hard to solve? using quantum computing to tackle*

optimization problems, Front. Quantum Sci. Technol., 2 (2023), doi:10.3389/frqst.2023.1128576.

[16] J. Preskill, *Quantum computing in the nisq era and beyond*, Quantum, 2 (2018), p. 79, doi:10.22331/q-2018-08-06-79.

[17] P. W. Shor, *Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer*, SIAM Journal on Computing, 26 (1997), pp. 1484–1509, doi:10.1137/s0097539795293172.

[18] L. K. Grover, *A fast quantum mechanical algorithm for database search*, in ACM symposium on Theory of Computing, 1996.

[19] R. D. Somma, D. Nagaj, and M. Kieferová, *Quantum speedup by quantum annealing*, Phys. Rev. Lett., 109 (2012), doi:10.1103/PhysRevLett.109.050501. Article number: 050501.

[20] M. B. Hastings, *The power of adiabatic quantum computation with no sign problem*, Quantum, 5 (2021), p. 597, doi:10.22331/q-2021-12-06-597.

[21] J. Roland and N. J. Cerf, *Quantum search by local adiabatic evolution*, Phys. Rev. A, 65 (2002), doi:10.1103/physreva.65.042308. Article number: 042308.

[22] E. J. Crosson and D. A. Lidar, *Prospects for quantum enhancement with diabatic quantum annealing*, Nature Reviews Physics, (2021), pp. 466–489, doi:10.1038/s42254-021-00313-6.

[23] A. Villanueva, P. Najafi, and H. J. Kappen, *Why adiabatic quantum annealing is unlikely to yield speed-up*, Journal of Physics A: Mathematical and Theoretical, 56 (2023), doi:10.1088/1751-8121/ad0439. Article number: 465304.

[24] P. Hauke, H. G. Katzgraber, W. Lechner, H. Nishimori, and W. D. Oliver, *Perspectives of quantum annealing: methods and implementations*, Reports on Progress in Physics, 83 (2020), doi:10.1088/1361-6633/ab85b8. Article number: 054401.

[25] M. BERNASCHI, I. GONZÁLEZ-ADALID PEMARTÍN, V. MARTÍN-MAYOR, AND G. PARISI, *The quantum transition of the two-dimensional ising spin glass*, Nature, 631 (2024), pp. 749–754, doi:10.1038/s41586-024-07647-y.

[26] T. ALBASH AND D. A. LIDAR, *Adiabatic quantum computation*, Reviews of Modern Physics, 90 (2018), doi:10.1103/revmodphys.90.015002. Article number: 015002.

[27] A. LUCAS, *Ising formulations of many NP problems*, Frontiers in Physics, 2 (2014), doi:10.3389/fphy.2014.00005.

[28] T. M. MCCORMICK, Z. KLAIN, I. HERBERT, A. M. CHARLES, R. B. ANGLE, B. R. OSBORN, AND R. L. STREIT, *Multiple target tracking and filtering using bayesian diabatic quantum annealing*, 2022. https://arxiv.org/abs/2209.00615.

[29] J. DING, G. SPALLITTA, AND R. SEBASTIANI, *Effective prime factorization via quantum annealing by modular locally-structured embedding*, Scientific Reports, 14 (2024), doi:10.1038/s41598-024-53708-7. Article number: 3518.

[30] E. PELOFSKE, *4-clique network minor embedding for quantum annealers*, Phys. Rev. Applied, 21 (2024), doi:10.1103/physrevapplied.21.034023. Article number: 034023.

[31] H. M. NGO, T. KAHVECI, AND M. T. THAI, *ATOM: An efficient topology adaptive algorithm for minor embedding in quantum computing*, 2023. https://arxiv.org/abs/2307.01843.

[32] E. LOBE AND V. KAIBEL, *Optimal sufficient requirements on the embedded ising problem in polynomial time*, Quantum Information Processing, 22 (2023), doi:10.1007/s11128-023-04058-2. Article number: 305.

[33] D-WAVE, *D-WAVE solutions and products, systems*. `https://www.dwavequantum.com/solutions-and-products/systems/`. Accessed: 2025-06-23.

[34] T. P. Orlando, J. E. Mooij, L. Tian, C. H. van der Wal, L. S. Levitov, S. Lloyd, and J. J. Mazo, *Superconducting persistent-current qubit*, Phys. Rev. B, 60 (1999), pp. 15398–15413, doi:10.1103/physrevb.60.15398.

[35] W. Oliver and P. Welander, *Materials in superconducting quantum bits*, MRS Bulletin, 38 (2013), pp. 816–825, doi:10.1557/mrs.2013.229.

[36] A. J. Kerman, *Superconducting qubit circuit emulation of a vector spin-1/2*, New Journal of Physics, 21 (2019), doi:10.1088/1367-2630/ab2ee7. Article number: 073030.

[37] I. Ozfidan, C. Deng, A. Smirnov, T. Lanting, R. Harris, L. Swenson, J. Whittaker, F. Altomare, M. Babcock, C. Baron, A. Berkley, K. Boothby, H. Christiani, P. Bunyk, C. Enderud, B. Evert, M. Hager, A. Hajda, J. Hilton, S. Huang, E. Hoskinson, M. Johnson, K. Jooya, E. Ladizinsky, N. Ladizinsky, R. Li, A. MacDonald, D. Marsden, G. Marsden, T. Medina, R. Molavi, R. Neufeld, M. Nissen, M. Norouzpour, T. Oh, I. Pavlov, I. Perminov, G. Poulin-Lamarre, M. Reis, T. Prescott, C. Rich, Y. Sato, G. Sterling, N. Tsai, M. Volkmann, W. Wilkinson, J. Yao, and M. Amin, *Demonstration of a nonstoquastic hamiltonian in coupled superconducting flux qubits*, Phys. Rev. Applied, 13 (2020), doi:10.1103/physrevapplied.13.034037. Article number: 034037.

[38] R. J. Banks, N. Feinstein, R. Ghosh, S. Bose, and P. A. Warburton, *Gadgets for simulating a non-native xx interaction in quantum annealing*, 2025. https://arxiv.org/abs/2503.16663.

[39] 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 2,000 qubit ising chain*, Nature Physics, 18 (2022), pp. 1324–1328, doi:10.1038/s41567-022-01741-6.

[40] C. Rigetti, J. M. Gambetta, S. Poletto, B. L. T. Plourde, J. M. Chow, A. D. Córcoles, J. A. Smolin, S. T. Merkel, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen, *Superconducting qubit in a waveguide cavity with a coherence time approaching 0.1 ms*, Phys. Rev. B, 86 (2012), doi:10.1103/physrevb.86.100506. Article number: 100506.

[41] X. Jin, A. Kamal, A. Sears, T. Gudmundsen, D. Hover, J. Miloshi, R. Slattery, F. Yan, J. Yoder, T. Orlando, S. Gustavsson, and W. Oliver, *Thermal and residual excited-state population in a 3d transmon qubit*, Phys. Rev. Lett., 114 (2015), doi:10.1103/physrevlett.114.240501. Article number: 240501.

[42] A. Somoroff, Q. Ficheux, R. A. Mencia, H. Xiong, R. V. Kuzmin, and V. E. Manucharyan, *Millisecond coherence in a superconducting qubit*, Phys. Rev. Lett., 130 (2023), doi:10.1103/PhysRevLett.130.267001. Article number: 267001.

[43] M. Saffman, *Quantum computing with atomic qubits and rydberg interactions: progress and challenges*, Journal of Physics B: Atomic, Molecular and Optical Physics, 49 (2016), doi:10.1088/0953-4075/49/20/202001. Article number: 202001.

[44] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho, S. Choi, S. Sachdev, M. Greiner, V. Vuletić, and M. D. Lukin,

Quantum phases of matter on a 256-atom programmable quantum simulator, Nature, 595 (2021), pp. 227–232, doi:10.1038/s41586-021-03582-4.

[45] A. W. Glaetzle, R. M. W. van Bijnen, P. Zoller, and W. Lechner, *A coherent quantum annealer with rydberg atoms*, Nature Communications, 8 (2017), doi:10.1038/ncomms15813.

[46] 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 (2022), pp. 1209–1215, doi:10.1126/science.abo6587.

[47] A. G. de Oliveira, E. Diamond-Hitchcock, D. M. Walker, M. T. Wells-Pestell, G. Pelegrí, C. J. Picken, G. P. A. Malcolm, A. J. Daley, J. Bass, and J. D. Pritchard, *Demonstration of weighted-graph optimization on a rydberg-atom array using local light shifts*, PRX Quantum, 6 (2025), doi:10.1103/prxquantum.6.010301. Article number: 010301.

[48] J. Kombe, G. Pelegrí, A. J. Daley, and J. D. Pritchard, *A quantum wire approach to weighted combinatorial graph optimisation problems*, 2025. https://arxiv.org/abs/2503.17115.

[49] W. Lechner, P. Hauke, and P. Zoller, *A quantum annealing architecture with all-to-all connectivity from local interactions*, Science Advances, 1 (2015), doi:10.1126/sciadv.1500838.

[50] 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 (2001), pp. 472–475, doi:10.1126/science.1057726.

[51] T. Hogg, *Adiabatic quantum computing for random satisfiability problems*, Phys. Rev. A, 67 (2003), doi:10.1103/physreva.67.022314. Article number: 022314.

[52] V. N. Smelyanskiy, U. V. Toussaint, and D. A. Timucin, *Simulations of the adiabatic quantum optimization for the set partition problem*, 2001. https://arxiv.org/abs/quant-ph/0112143.

[53] A. P. Young, S. Knysh, and V. N. Smelyanskiy, *Size dependence of the minimum excitation gap in the quantum adiabatic algorithm*, Phys. Rev. Lett., 101 (2008), doi:10.1103/physrevlett.101.170503. Article number: 170503.

[54] A. P. Young, S. Knysh, and V. N. Smelyanskiy, *First-order phase transition in the quantum adiabatic algorithm*, Phys. Rev. Lett, 104 (2010), doi:10.1103/PhysRevLett.104.020502. Article number: 020502.

[55] I. Hen and A. P. Young, *Exponential complexity of the quantum adiabatic algorithm for certain satisfiability problems*, Phys. Rev. E, 84 (2011), doi:10.1103/physreve.84.061152. Article number: 061152.

[56] M. Ž nidarič and M. Horvat, *Exponential complexity of an adiabatic algorithm for an NP-complete problem*, Phys. Rev. A, 73 (2006), doi:10.1103/physreva.73.022329. Article number: 022329.

[57] M. H. Amin, *Effect of local minima on adiabatic quantum optimization*, Phys. Rev. Lett., 100 (2008), doi:10.1103/PhysRevLett.100.130503. Article number: 130503.

[58] T. Jörg, F. Krzakala, J. Kurchan, and A. C. Maggs, *Simple glass models and their quantum annealing*, Phys. Rev. Lett., 101 (2008), doi:10.1103/physrevlett.101.147204. Article number: 147204.

[59] T. Jörg, F. Krzakala, J. Kurchan, and A. C. Maggs, *Quantum annealing of hard problems*, Progress of Theoretical Physics, 184 (2010), pp. 290–303, doi:10.1143/ptps.184.290.

[60] 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 (2012), p. 435301, doi:10.1088/1751-8113/45/43/435301.

[61] S. Knysh, *Zero-temperature quantum annealing bottlenecks in the spin-glass phase*, Nature Communications, 7 (2016), p. 12370, doi:10.1038/ncomms12370.

[62] B. Altshuler, H. Krovi, and J. Roland, *Anderson localization makes adiabatic quantum optimization fail*, Proceedings of the National Academy of Sciences, 107 (2010), pp. 12446–12450, doi:10.1073/pnas.1002116107.

[63] M. H. Amin and V. Choi, *First-order quantum phase transition in adiabatic quantum computation*, Phys. Rev. A, 80 (2009), doi:10.1103/PhysRevA.80.062326. Article number: 062326.

[64] N. G. Dickson and M. H. Amin, *Does adiabatic quantum optimization fail for NP-complete problems?*, Phys. Rev. Lett., 106 (2011), doi:10.1103/PhysRevLett.106.050502. Article number: 050502.

[65] V. Choi, *The effects of the problem Hamiltonian parameters on the minimum spectral gap in adiabatic quantum optimization*, Quantum Information Processing, 19 (2020), p. 90, doi:10.1007/s11128-020-2582-1.

[66] 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 (2011), pp. 181–214, doi:10.26421/qic11.3-4-1.

[67] T. Albash, *Role of nonstoquastic catalysts in quantum adiabatic optimization*, Phys. Rev. A, 99 (2019), doi:10.1103/PhysRevA.99.042334. Article number: 042334.

[68] T. ALBASH AND M. KOWALSKY, *Diagonal catalysts in quantum adiabatic optimization*, Phys. Rev. A, 103 (2021), doi:10.1103/PhysRevA.103.022608. Article number: 022608.

[69] A. DEL CAMPO, *Shortcuts to adiabaticity by counterdiabatic driving*, Phys. Rev. Lett., 111 (2013), doi:10.1103/physrevlett.111.100502. Article number: 100502.

[70] G. PASSARELLI, V. CATAUDELLA, R. FAZIO, AND P. LUCIGNANO, *Counterdiabatic driving in the quantum annealing of the p-spin model: A variational approach*, Phys. Rev. Research, 2 (2020), doi:10.1103/physrevresearch.2.013283. Article number: 013283.

[71] L. PRIELINGER, A. HARTMANN, Y. YAMASHIRO, K. NISHIMURA, W. LECHNER, AND H. NISHIMORI, *Two-parameter counter-diabatic driving in quantum annealing*, Phys. Rev. Research, 3 (2021), doi:10.1103/physrevresearch.3.013227. Article number: 013227.

[72] A. HARTMANN, G. B. MBENG, AND W. LECHNER, *Polynomial scaling enhancement in the ground-state preparation of ising spin models via counterdiabatic driving*, Phys. Rev. A, 105 (2022), doi:10.1103/physreva.105.022614. Article number: 022614.

[73] I. ČEPAITĖ, A. POLKOVNIKOV, A. J. DALEY, AND C. W. DUNCAN, *Counterdiabatic optimized local driving*, PRX Quantum, 4 (2023), doi:10.1103/prxquantum.4.010312. Article number: 010312.

[74] N. O. GJONBALAJ, D. K. CAMPBELL, AND A. POLKOVNIKOV, *Counterdiabatic driving in the classical β-fermi-pasta-ulam-tsingou chain*, Phys. Rev. E, 106 (2022), doi:10.1103/physreve.106.014131. Article number: 014131.

[75] F. P. BARONE, O. KISS, M. GROSSI, S. VALLECORSA, AND A. MANDARINO, *Counterdiabatic optimized driving in quantum phase sensitive models*, New Journal of Physics, 26 (2024), doi:10.1088/1367-2630/ad313e. Article number: 033031.

[76] M. Ohkuwa, H. Nishimori, and D. A. Lidar, *Reverse annealing for the fully connected p-spin model*, Phys. Rev. A, 98 (2018), doi:10.1103/physreva.98.022314. Article number: 033031.

[77] D. Venturelli and A. Kondratyev, *Reverse quantum annealing approach to portfolio optimization problems*, Quantum Machine Intelligence, 1 (2019), pp. 17–30, doi:10.1007/s42484-019-00001-w.

[78] Y. Yamashiro, M. Ohkuwa, H. Nishimori, and D. A. Lidar, *Dynamics of reverse annealing for the fully connected p-spin model*, Phys. Rev. A, 100 (2019), doi:10.1103/physreva.100.052321. Article number: 052321.

[79] R. Haba, M. Ohzeki, and K. Tanaka, *Travel time optimization on multi-agv routing by reverse annealing*, Scientific Reports, 12 (2022), doi:10.1038/s41598-022-22704-0. Article number: 17753.

[80] M. S. Jattana, *Reverse quantum annealing assisted by forward annealing*, Quantum Reports, 6 (2024), pp. 452–464, doi:10.3390/quantum6030030.

[81] R. Ghosh, L. A. Nutricati, N. Feinstein, P. A. Warburton, and S. Bose, *Exponential speed-up of quantum annealing via n-local catalysts*, 2024. https://arxiv.org/abs/2409.13029.

[82] J. Dziarmaga and M. M. Rams, *Dynamics of an inhomogeneous quantum phase transition*, New Journal of Physics, 12 (2010), doi:10.1088/1367-2630/12/5/055007. Article number: 055007.

[83] 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 (2018), pp. 4–7, doi:10.7566/JPSJ.87.023002.

[84] Y. Susa, Y. Yamashiro, M. Yamamoto, I. Hen, D. A. Lidar, and H. Nishimori, *Quantum annealing of the p-spin model un-*

der inhomogeneous transverse field driving, Phys. Rev. A, 98 (2018), doi:10.1103/physreva.98.042326. Article number: 042326.

[85] 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 (2020), doi:10.1088/2058-9565/ab935a. Article number: 035011.

[86] E. Farhi, J. Goldston, D. Gosset, S. Gutmann, H. B. Meyer, and P. Shor, *Quantum adiabatic evolution algorithms with different paths*, Quantum Info. Comput., 11 (2011), pp. 181–214, doi:10.26421/QIC11.3-4-1.

[87] E. Crosson, E. Farhi, C. Y.-Y. Lin, H.-H. Lin, and P. Shor, *Different strategies for optimization using the quantum adiabatic algorithm*, 2014. https://arxiv.org/abs/1401.7320.

[88] L. Zeng, J. Zhang, and M. Sarovar, *Schedule path optimization for adiabatic quantum computing and optimization*, Journal of Physics A: Mathematical and Theoretical, 49 (2016), doi:10.1088/1751-8113/49/16/165305. Article number: 165305.

[89] L. Hormozi, E. W. Brown, G. Carleo, and M. Troyer, *Nonstoquastic Hamiltonians and quantum annealing of an Ising spin glass*, Phys. Rev. B, 95 (2017), doi:10.1103/PhysRevB.95.184416. Article number: 184416.

[90] V. Mehta, F. Jin, H. De Raedt, and K. Michielsen, *Quantum annealing with trigger hamiltonians: Application to 2-satisfiability and nonstoquastic problems*, Phys. Rev. A, 104 (2021), doi:10.1103/physreva.104.032421. Article number: 032421.

[91] E. Crosson, T. Albash, I. Hen, and A. P. Young, *De-signing Hamiltonians for quantum adiabatic optimization*, Quantum, 4 (2020), p. 334, doi:10.22331/Q-2020-09-24-334.

[92] Y. Seki and H. Nishimori, *Quantum annealing with antiferromagnetic fluctuations*, Phys. Rev. E, 85 (2012), doi:https://doi.org/10.1103/PhysRevE.85.051112. Article number: 051112.

[93] H. Nishimori and K. Takada, *Exponential enhancement of the efficiency of quantum annealing by non-stoquastic Hamiltonians*, Frontiers in ICT, 4 (2017), doi:10.3389/fict.2017.00002.

[94] Y. Susa, J. F. Jadebeck, and H. Nishimori, *Relation between quantum fluctuations and the performance enhancement of quantum annealing in a nonstoquastic hamiltonian*, Phys. Rev. A, 95 (2017), doi:10.1103/physreva.95.042321. Article number: 042321.

[95] 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*, Phys. Rev. Research, 3 (2021), doi:10.1103/physrevresearch.3.043013. Article number: 043013.

[96] 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 (2020), pp. 1–25, doi:10.7566/JPSJ.89.044001.

[97] M. Mohseni, J. Strumpfer, and M. M. Rams, *Engineering non-equilibrium quantum phase transitions via causally gapped hamiltonians*, New Journal of Physics, 20 (2018), doi:10.1088/1367-2630/aae3ed. Article number: 105002.

[98] R. Juhász and G. Roósz, *Reducing defect production in random transverse-field ising chains by inhomogeneous driving fields*, Phys. Rev. B, 108 (2023), doi:10.1103/physrevb.108.224203. Article number: 224203.

[99] N. G. Dickson and M. H. Amin, *Algorithmic approach to adiabatic quantum optimization*, Phys. Rev. A, 85 (2012), pp. 1–7, doi:10.1103/PhysRevA.85.032303.

[100] V. Choi, *Essentiality of the Non-stoquastic Hamiltonians and Driver Graph Design in Quantum Optimization Annealing*, 2021. https://arxiv.org/abs/2105.02110.

[101] D. Wecker, M. B. Hastings, and M. Troyer, *Training a quantum optimizer*, Phys. Rev. A, 94 (2016), doi:10.1103/physreva.94.022309. Article number: 022309.

[102] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, *Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices*, Phys. Rev. X, 10 (2020), doi:10.1103/physrevx.10.021067. Article number: 021067.

[103] A. Rakcheev and A. M. Läuchli, *Diabatic quantum and classical annealing of the sherrington-kirkpatrick model*, Phys. Rev. A, 107 (2023), doi:10.1103/physreva.107.062602. Article number: 062602.

[104] S. Muthukrishnan, T. Albash, and D. A. Lidar, *Tunneling and speedup in quantum optimization for permutation-symmetric problems*, Phys. Rev. X, 6 (2016), doi:10.1103/physrevx.6.031010. Article number: 031010.

[105] L. T. Brady and W. van Dam, *Necessary adiabatic run times in quantum optimization*, Phys. Rev. A, 95 (2017), doi:10.1103/physreva.95.032335. Article number: 032335.

[106] L. Fry-Bouriaux, D. T. O'Connor, N. Feinstein, and P. A. Warburton, *Locally suppressed transverse-field protocol for diabatic quantum annealing*, Phys. Rev. A, 104 (2021), doi:10.1103/physreva.104.052616. Article number: 052616.

[107] J. Côté, F. Sauvage, M. Larocca, M. Jonsson, L. Cincio, and T. Albash, *Diabatic quantum annealing for the frustrated ring model*, Quantum Science and Technology, 8 (2023), doi:10.1088/2058-9565/acfbaa. Article number: 045033.

[108] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev, *Adiabatic quantum computation is equivalent to standard quantum computation*, in IEEE Symposium on Foundations of Computer Science, 2004, pp. 42–51, doi:10.1109/FOCS.2004.8.

[109] D-WAVE, *D-WAVE qpu solver datasheet, annealing implementation and controls*. `https://docs.dwavesys.com/docs/latest/c_qpu_annealing.html`. Accessed: 2023-02-07.

[110] A. Braida and S. Martiel, *Anti-crossings and spectral gap during quantum adiabatic evolution*, Quantum Information Processing, 20 (2021), doi:10.1007/s11128-021-03198-7. Article number: 260.

[111] NumPy, *numpy.linalg.eigh*. `https://numpy.org/doc/2.0/reference/generated/numpy.linalg.eigh.html`. Accessed: 2024-10-14.

[112] LAPACK, *Lapak routines*. `https://www.netlib.org/lapack/`. Accessed: 2024-10-14.

[113] QuTiP, *qutip.sesolve*. `https://qutip.org/docs/4.0.2/modules/qutip/sesolve.html`. Accessed: 2024-10-17.

[114] QuTiP, *qutip.mesolve*. `https://qutip.org/docs/4.0.2/modules/qutip/mesolve.html`. Accessed: 2024-10-17.

[115] QuTiP, *qutip.options*. `https://qutip.org/docs/4.0.2/guide/dynamics/dynamics-options.html`. Accessed: 2024-10-17.

[116] L. Gupta and I. Hen, *Elucidating the interplay between non-stoquasticity and the sign problem*, Advanced Quantum Technologies, 3 (2019), p. 1900108, doi:https://doi.org/10.1002/qute.201900108.

[117] H. Munoz-Bauza, H. Chen, and D. Lidar, *A double-slit proposal for quantum annealing*, npj Quantum Information, 5 (2019), doi:10.1038/s41534-019-0160-0. Article number: 51.