

---

# Computational Methods in Quantum Annealing

---

April 24, 2024

*Daniel T. O'Connor*

A dissertation submitted in partial fulfillment  
of the requirements for the degree of  
**Doctor of Philosophy**  
of  
**University College London.**

Department of Electronic and Electrical Engineering  
University College London

April 24, 2024

I, Daniel T. O'Connor, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

---

# Abstract

---

Quantum annealing is a non-universal sub-field of quantum computing used for solving combinatorial optimisation problems found in industries such as telecommunications, supply-chain networks, and finance. However, whether current quantum annealers can provide additional commercial value compared to state-of-the-art classical optimization algorithms is still an open question, given that some of the quantum nature of the computation is lost to noise sources plaguing modern qubits.

The work presented addresses a range of problems faced when solving optimisation problems on quantum annealers, with experimental validation of the novel methods proposed where possible. We compare parameter setting methods, demonstrating that the choice of method is crucial to the success of minor-embedding, something which is typically overlooked in the literature when solving large-scale optimisation problems that cannot be directly embedded onto the quantum processing unit (QPU) topology. It is also shown how the mapping that exists between NP problems is useful in reducing qubit overheads in constrained optimisation problems, with the reduction of the graph-colouring problem into several maximal-independent set problems evaluated.

Reverse-annealing for optimisation was explored for its potential as a local-search algorithm that can improve existing suboptimal results. It was found that thermal effects on the QPU contribute to the performance of this algorithm, hinting that non-adiabatic algorithms can be successful in finding solutions on modern quantum annealers. The usefulness of thermally assisted computation is investigated using a Hamiltonian with tuneable hardness called the perturbed ferromagnetic chain, which was also used to demonstrate that the thermalization mechanisms observed on

the QPU are distinctly different from classical thermalization mechanisms.

Finally, the framework of the locally-suppressed transverse-field diabatic quantum annealing is extended to optimisation problems and compared to traditional adiabatic quantum annealing and classical analogues. The protocol is also adapted into an interferometer and used as a test of coherence for multi-qubit systems.

---

# Impact Statement

---

The key methods and insights presented in this thesis are intended to demonstrate to the users of quantum annealers how to improve the computational success rate given the limitations of current devices. The users of modern quantum annealers span both academic and industrial professions, of which the latter is primarily concerned with commercial optimisation problems. This includes those faced by British Telecom, the co-sponsor of this work, where the promise of quantum-enhanced network optimisation could lead to improved services from the United Kingdom's internet infrastructure operator. Other industries that stand to benefit from quantum-enhanced optimization include, but are not limited to, financial services, logistics, and the pharmaceutical industry.

The technical findings of this thesis also probed the physical processes of quantum annealers to understand the usefulness of quantum dynamics in the annealing computation. This has been disseminated to the academic community through publications and presentations at conferences such as the American Physical Society March Meeting. Efforts have also been made to make this thesis and associated codebases publicly available to allow the wider community to adopt the research methods employed throughout. This includes the novel quantum optimization protocols presented here that are feasible on current hardware but not yet accessible due to the device constraints. The impact of enabling these methods on modern quantum annealers will contribute to the diversification of quantum algorithms used to solve optimization problems.

---

# UCL Research Paper Declaration Form

---

1. **1. For a research manuscript that has already been published** (if not yet published, please skip to section 2):

(a) **What is the title of the manuscript?**

Perturbed ferromagnetic chain: Tunable test of hardness in the transverse-field Ising model

(b) **Please include a link to or doi for the work:**

<https://doi.org/10.1103/PhysRevA.105.022410>

(c) **Where was the work published?**

Physical Review A

(d) **Who published the work?**

American Physical Society

(e) **When was the work published?**

10 February 2022

(f) **List the manuscript's authors in the order they appear on the publication:**

Daniel T. O'Connor, Louis Fry-Bouriaux, and Paul A. Warburton

(g) **Was the work peer reviewed?**

Yes.

(h) **Have you retained the copyright?**

No.

(i) **Was an earlier form of the manuscript uploaded to a preprint server (e.g. medRxiv)? If 'Yes', please give a link or doi**

<https://doi.org/10.48550/arXiv.2106.11019>

2. **For a research manuscript prepared for publication but that has not yet been published** (if already published, please skip to section 3):

(a) **What is the current title of the manuscript?**

Already published, skip to section 3.

3. **For multi-authored work, please give a statement of contribution covering all authors** (if single-author, please skip to section 4):

- Daniel O'Connor designed the Hamiltonian and performed the analytical and numerical analysis of the problem.
- Louis Fry-Bouriaux assisted in defining the methodologies used throughout.
- Paul Warburton reviewed the publication and proposed additional analyses that were included.

4. **In which chapter(s) of your thesis can this material be found?**

Chapter 3 Section 3.1.

**e-Signatures confirming that the information above is accurate** (this form should be co-signed by the supervisor/ senior author unless this is not appropriate, e.g. if the paper was a single-author work):

**Candidate:** Daniel O'Connor

**Date:** 13/11/2023

**Supervisor/Senior Author signature** Paul Warburton

**Date:** 14/11/2023

---

# Acknowledgements

---

I would first and foremost like to thank my supervisor Professor Paul Warburton for his absolute support and patience throughout this journey, whose guidance has been invaluable throughout both within and beyond my studies. I am also extremely thankful for the support and advice from my industrial supervisor Catherine White, and her colleagues from British Telecom, Keith Briggs and Nigel Walker, for sharing their precious insights and expertise. I would like to thank Professor Ed Romans and Professor Sugato Bose for their help in my studies, and to thank my peers from the London Centre for Nanotechnology, Louis Fry-Bouriaux, Gioele Consani, Nadeen Alsharif, Robert Banks and Natasha Feinstein, for their support and discussions that taught me so much and helped spur the ideas in the work presented here.

I am forever grateful to Walter Vinci for bringing me to the Quantum Artificial Intelligence Laboratory (QuAIL) in NASA Ames Research Center, where I met and learnt from the most extraordinary people I had ever encountered, including Eleanor Rieffel, Gianni Mossi, Filip Wudarski, Jeffrey Marshall and Max Wilson. I would also like to thank Stephen Eidenbenz and Brandi Brown from Los Alamos National Laboratory, whose approval and help to access the D-Wave quantum annealer enabled a large majority of my work to occur in the first place. I am thankful to Federico Spedalieri from the Information Sciences Institute of the University of Southern California for also providing important access to a range of D-Wave quantum annealers, and to Daniel Lidar, David Ferguson and Tameem Albash for the valuable discussions regarding my research and understanding.

I would like to dearly thank all my friends and family for always believing in me, in particular my mother, Sophia, who is the sole reason that I was able to believe



I could strive to achieve this feat, and my grandmother, Wilhelmina, for providing a home and support during the unprecedented time of lockdown and COVID-19, something that I will forever treasure. Finally, I would like to thank my loving partner Amanda, who has given me unwavering support throughout.

This work was supported by the Engineering and Physical Sciences Research Council and British Telecom. Other institutions that have supported this work include Los Alamos National Laboratory, Information Sciences Institute, and NVIDIA. The research presented is also based upon work (partially) supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), and the Defense Advanced Research Projects Agency (DARPA), via the U.S. Army Research Office Contract No. W911NF-17-C-0050. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the ODNI, IARPA, DARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

---

# Contents

---

<b>1</b>	<b>Adiabatic Quantum Computing &amp; Quantum Annealing</b>	<b>21</b>
1.1	Adiabatic Quantum Computation . . . . .	24
1.2	Quantum Annealing . . . . .	30
1.2.1	Formalism . . . . .	30
1.2.2	Applications . . . . .	36
1.2.3	Reverse Annealing . . . . .	38
1.3	Quantum Annealing using Superconducting Hardware . . . . .	42
1.3.1	Experimental Implementation . . . . .	42
1.3.2	Minor Embedding . . . . .	48
1.4	Simulating Quantum Annealing . . . . .	52
1.4.1	Adiabatic Master Equation . . . . .	52
1.4.2	Quantum Monte Carlo . . . . .	56
1.4.3	Classical Dynamics . . . . .	60
<b>2</b>	<b>Quantum Annealing Optimization</b>	<b>67</b>
2.1	Parameter Setting . . . . .	68
2.2	Graph Colouring . . . . .	81
2.3	Reverse Annealing . . . . .	92
2.3.1	Experimental . . . . .	99
2.3.2	Simulation . . . . .	106
<b>3</b>	<b>Thermalization in Quantum Annealing</b>	<b>116</b>
3.1	Perturbed Ferromagnetic Chain . . . . .	118
3.1.1	Classical Model . . . . .	119

3.1.2	Transverse-Field Ising Model . . . . .	122
3.1.3	Semi-Classical Analysis . . . . .	127
3.2	Simulated Thermalization . . . . .	129
3.2.1	Quantum & Classical Methods . . . . .	129
3.2.2	Dynamical Simulations . . . . .	133
3.3	Experimental Thermalization . . . . .	140
3.3.1	Model Comparisons . . . . .	141
3.3.2	Thermalization with Large Systems . . . . .	144
<b>4</b>	<b>Diabatic Quantum Annealing</b>	<b>152</b>
4.1	Diabatic Quantum Annealing for Optimization . . . . .	154
4.1.1	Maximal Independent Set . . . . .	157
4.1.2	Max Cut . . . . .	165
4.2	Testing Coherence using Diabatic Quantum Annealing . . . . .	169
4.2.1	Two-qubit Perturbed Ferromagnetic Chain . . . . .	170
4.2.2	LTSF-DQA Interferometry . . . . .	175
<b>5</b>	<b>Conclusion</b>	<b>185</b>
5.1	General Conclusions . . . . .	185
5.2	Future Work . . . . .	189
	<b>Appendices</b>	<b>191</b>
<b>A</b>	<b>Monte Carlo Algorithms</b>	<b>191</b>
<b>B</b>	<b>List of Publications</b>	<b>194</b>
<b>C</b>	<b>Colophon</b>	<b>195</b>
	<b>Bibliography</b>	<b>196</b>

---

# List of Figures

---

1.1	Illustration of an anti-crossing in a two-level system, with a minimum gap $\Delta_{10}$ that forms between the ground state and first-excited state. . . . .	28
1.2	Illustration of classical and quantum dynamics traversing the potential energy landscape. . . . .	31
1.3	Illustration of the open system thermalization effects that occur as a system approaches a minimum gap. . . . .	34
1.4	Schematic of the transitions induced between the ground and first-excited state in RFQA. . . . .	37
1.5	Static phase diagram for the $p$ -spin model using adiabatic reverse annealing. . . . .	40
1.6	Plot of the annealing control parameter, $s$ , in a reverse anneal schedule on a D-Wave quantum annealer. . . . .	41
1.7	Schematic of <b>a</b> the superconducting flux qubit and <b>b</b> its potential energy diagram. . . . .	43
1.8	Plot of the D-Wave 2000Q annealing schedule at the Los Alamos National Laboratory. . . . .	44
1.9	D-Wave's proprietary Chimera and Pegasus quantum processing unit architectures. . . . .	45
1.10	An Illustration of a minor-embedded logical qubit with degree 6 to multiple physical qubits (green) with degree 3. . . . .	50
1.11	Examples of (a) one-hot and (b) domain wall encodings for a single vertex colouring constraint of a four colouring problem. . . . .	51

2.1	Illustration of the minor-embedding of a three spin problem, where the central logical spin consists of three physical spins in the embedding. . . . .	73
2.2	Data from the LANL D-Wave 2000Q without post-processing showing a) the median time-to-solution (TTS) and b) the count of whether the single ground-state solution was found out of the 20 random instances generated for each logical problem size for each parameter setting method. . . . .	76
2.3	Data with post-processing showing a) the median time-to-solution (TTS) and b) the count of whether the single ground-state solution was found out of the 20 random instances generated for each logical problem size for each parameter setting method. . . . .	77
2.4	The best performing parameter-setting method, illustrated by win percentage out of all 140 problem instances, where a win for a method is when it has the lowest TTS for a given instance. . . . .	77
2.5	Plot of a) average minimum non-locality $\tilde{C}_{\min}$ , as a function of system size, and b) a fit of TTSR (Equation (2.11)) against $\tilde{C}_{\min}$ to illustrate the correlation between the variables. . . . .	79
2.6	Illustration of a 7-vertex problem graph $G = (V, E)$ correctly coloured with three colours. . . . .	82
2.7	Times to the correct colouring of Turán graphs (order $k$ , with 12 vertices) as a function of the number of computations. . . . .	88
2.8	Illustration of the optimal TTS scaling as a function of Turán graph colouring for the Lucas, RMIS, and GMIS (Alg. 1) graph colouring methods . . . . .	89
2.9	Plots of a) TTS as a function of problem size for the RMIS and GMIS colouring methods, and b) embedding ratio (Equation 2.20) between the Lucas and MIS methods. . . . .	91
2.10	Graphical illustration of the four spin chain problem Hamiltonian used for the reverse annealing experiments. . . . .	93

- 2.11 Plots of *a)* the first five instantaneous energy eigenstates  $E_i$  of the four spin chain as a function of the anneal parameter  $s$ , *b)* the scaling of the minimum gap  $\Delta_{10}$  as a function of the tuneable parameter  $d$ , *c)* the Hamming weights of the first five states according to the Hamming weight operator as a function of  $s$ , and *d)* the negativity (entanglement) of the ground state with respect to the inner qubits of the four spin chain. . . . . 94
- 2.12 Plots of the semi-classical potential as a function of  $s$  for a four spin chain with  $d = 0.05$ . with the red marker indicating the global minimum of the landscape. . . . . 96
- 2.13 Ground state probability after reverse annealing with a pause of *a)*  $0\mu s$  and *b)*  $2\mu s$  from each of the four degenerate first-excited-states on the LANL D-Wave 2000Q quantum annealer. . . . . 100
- 2.14 Plots of *a)* the time-to-solution as a function of  $s^*$  for the reverse anneal with a  $2\mu s$  pause, and *b)* the minimum TTS at  $s^*_{\text{opt}}$  for various pause times when reverse annealed from each of the four degenerate first-excited-states on the LANL D-Wave 2000Q quantum annealer. . . . . 101
- 2.15 The probability of being in the  $|0\rangle$  state after a reverse anneal as a function of  $s^*$  (the minimum value of  $s$  annealed to) for a single unbiased spin on the D-Wave 2000Q. Starting in either of the initial states biases the qubit differently for both annealing rates of *a)*  $r = 0.5$  (i.e., shorter reverse anneal) and *b)*  $r = 0.05$  (i.e., longer reverse anneal) . . . . . 102
- 2.16 Plots of *a)* the oscillations in ground state probability as a function of anneal time  $t_{\text{anneal}}$ , such that  $t_{\text{anneal}} = 2t_1 + t_p$  where  $t_1 = 2(1 - s^*)\mu s$  and  $t_p = 2\mu s$ , and *b)* its corresponding power spectrum. . . . . 104
- 2.17 Plots of ground state probability after a reverse quantum anneal from each of the initial states *a)* without cross-talk and *b)* with cross-talk, which was simulated using the AME in the singular coupling limit. . 108

2.18	Reverse simulated anneals of the four spin chain from each of the computational first-excited states for <i>a</i> ) no pause and <i>b</i> ) a pause of $2\mu s$ . . . . .	109
2.19	Reverse spin-vector Monte Carlo (SVMC) anneals of the four spin chain from each of the computational first-excited states for <i>a</i> ) no pause and <i>b</i> ) a pause of $2\mu s$ at a system temperature of 12.26 mK. . . . .	110
2.20	Plots of reverse annealing the four spin chain with <i>a</i> ) no pause and <i>b</i> ) a pause of $2\mu s$ when simulated by SVMC with transverse-field dependent updates (SVMC-TF) starting in each of the computational first-excited states. . . . .	111
2.21	PIMC reverse anneals of the four spin chain from each of the computational first-excited states for <i>a</i> ) no pause and <i>b</i> ) a pause of $2\mu s$ at a system temperature of 12.26 mK. . . . .	112
2.22	Reverse anneal simulations of the four spin chain with a pause of $2\mu s$ and using PIMC with dephasing. The initial states are from each of the computational first-excited states, and the extent of dephasing is represented here by three regimes of the bath-coupling, parameterized by $\alpha$ at a system temperature of 12.26 mK. . . . .	113
3.1	Illustration of the perturbed ferromagnetic chain Hamiltonian in general form. . . . .	119
3.2	(a) Intensity plot of the average qubit magnetization in the instantaneous ground state for a PFC in the presence of a transverse field with $M = 2$ and $R = 1.0$ . (b) Cross sections of (a) showing the average magnetization during an anneal. . . . .	123
3.3	Plot of how the minimum gap $\Delta_{10}$ scales with system size $M$ in the TFIM for various values of the perturbative parameter $d$ . . . . .	125
3.4	Plots of the semiclassical energy potential (top row) and the energy along the hyperplane (white dashed line) passing through the landscape (bottom row) at the specified values of normalized time, $s$ . . . . .	126

- 3.5 The trace-norm distance between the spin-coherent state and the instantaneous ground state (Equation (3.16)) in the vicinity of the minimum gap at  $s = 0.841$  for a PFC with  $M = 2$ ,  $R = 1.0$  and  $d = 0.09$ . . . . . 129
- 3.6 Probability of being in the ground state ( $P_G$ ) for (a) both SVMC (solid lines) and spherical SVMC-TF (dotted lines), and (b) SVMC-TF (solid lines) and spherical SVMC (dotted lines), as the system scales in size,  $M$ , for a PFC with  $d = 0.1$ . . . . . 134
- 3.7 State probability at  $s = 0.83$  ( $= t/t_{\text{anneal}}$ ), just after passing through the minimum gap ( $s = 0.8227$ ) for (a) the AME with  $t_{\text{anneal}} = 200ns$  and (b) spherical SVMC-TF with  $t_{\text{anneal}} = 10,000$  sweeps. . . . . 136
- 3.8 Plots of state probability for being in either the ground state (solid lines) or any of the  $2^M$ -degenerate first excited states (dashed lines) at the end of an anneal. A PFC with system size  $M = 3$  was evolved using quantum and classical dynamics. . . . . 137
- 3.9 Plot of the evolution of ground state probability from the AME (black solid) and the Gibbs state (black dotted), as well as the transition rate (blue). . . . . 138
- 3.10 Plots of D-Wave annealing coefficient schedules for the *a*) DW\_2000Q\_6, *b*) Los Alamos National Laboratory DW\_2000Q, and *c*) Advantage\_system4.1 systems. . . . . 141
- 3.11 The example energy-gap spectrum for an  $M = 6$  perturbed ferromagnetic chain with  $d = 0.25$ , where the  $d < \frac{1}{N-1}$  condition is violated. 142



- 3.12 Comparison between the simulated classical and quantum models to the D-Wave Advantage\_system4.1 annealer for the short anneal time regime between  $500ns - 1\mu s$  for an  $M = 3, d = 0.05$  PFC. Plot *a*) illustrates simulated models with no integrated control errors in the anneal, whilst *b*) introduces cross-talk into the Hamiltonian with  $\chi = -0.035$  for all simulated models, and parameter setting errors for both biases and coupling with a standard deviation of 0.03 for SVMC-TF simulations. . . . . 143
- 3.13 State probabilities after a forward anneal on an  $M = 15, d = 0.1$  PFC using *a*) the DW\_2000Q\_6 and *b*) the Advantage\_system4.1 D-Wave annealing systems. . . . . 145
- 3.14 Illustration of the forward anneal with a pause for time  $t_p$  at a constant value of  $s = s_p$ . . . . . 146
- 3.15 Ground state probability after a forward anneal using *a*) the DW\_2000Q\_6 and *b*) SVMC-TF for various energy scales  $R$  and pause positions  $s_p$ . . . . . 147
- 3.16 Diagram of an intermediate state configuration that could provide a route to the ground state configuration. . . . . 148
- 3.17 Plots of *a*) time-to-solution ( $TTS$ ) for a  $d = 0.19$  PFC for various anneal times and system sizes,  $M$ , and *b*) the optimum time-to-solution found for each system size. . . . . 149
- 4.1 Example of a diabatic annealing spectrum where instantaneous state energy is a function of normalized annealing time  $s$ . The green points  $1 \rightarrow 2$  and  $3 \rightarrow 4$  indicate diabatic transitions at exponentially scaling avoided-crossings, whilst all other spectral gaps are polynomially scaling. . . . . 153
- 4.2 Plots of *a*) the energy spectrum and *b*) the annealing schedules for both adiabatic quantum annealing (Equation 4.1) and LSTF-DQA (Equation 4.2). . . . . 156

- 4.3 The graphs of the only two non-isomorphic 3-regular graphs with 6 vertices, with an example of a maximum independent set for each graph indicated by the red coloured vertices. . . . . 158
- 4.4 Plots of the expected ground-state probability as a function of anneal time for *a*) the prism and *b*) the bipartite isomorphisms simulated using closed quantum system,  $O(2)$  and  $O(3)$  spin-vector dynamics for both adiabatic quantum annealing (AQA) and LSTF-DQA protocols. . . . . 161
- 4.5 Difference in expected ground-state probability  $\langle P_{GS}(t_{\text{anneal}}) \rangle$  when quantum LSTF-DQA anneals chose the target qubit based on whether it possesses the largest or lowest local frustration, i.e.,  $k = \arg \max f(k)$  or  $k = \arg \min f(k)$  respectively. . . . . 164
- 4.6 Illustration of the 8-spin max-cut problem and the respective spectra, showing the two-lowest energy levels in the adiabatic quantum annealing and LSTF-DQA settings. . . . . 167
- 4.7 Plot of the average ground-state probability for the 8-qubit problem where local bias is added to the target qubit of degree  $\delta$  and all adjacent qubits. . . . . 167
- 4.8 Plots of *a*) the two lowest energy states in the unperturbed setting, and *b*) the comparison between the numerical and analytical results for both  $\bar{\Delta}(s)$  and  $s_g$ . . . . . 172
- 4.9 Plot of *a*) shows the two avoided-level crossings in the perturbed setting with  $c_0 = 2.0 \times 10^{-3}$ ,  $c_1 = 2.75 \times 10^{-3}$ ,  $d = 0.3$  and  $s_x = 0.2$ . Plot *b*) illustrates the closed-system evolution of the Hamiltonian defined in plot *a*) for a 500 ns anneal. . . . . 174
- 4.10 The numerical values of  $c_1$  that result in both avoided-level crossing gaps being equal in size compared to the estimated value of  $c_0 = c_1 = 2 \times 10^{-3}$  as a function of  $d$ . . . . . 175

4.11	Results of simulated LSTF-DQA interferometry for the ground state of the perturbed ferromagnetic chain at $N = 1$ and $d = 0.75$ as a function of the anneal time, $t_{anneal}$ . . . . .	177
4.12	Plots of <i>a</i> ) the ground state probability $P_{GS}$ oscillation frequency changing as a function of the perturbative parameter $d$ and anneal time $t_{anneal}$ , and <i>b</i> the oscillation frequency derived from the numerical methods. . . . .	178
4.13	Demonstration that LSTF-DQA interferometry occurs for PFC systems with $M > 1$ and $d = 0.75$ . . . . .	179
4.14	Open-system simulations of LSTF-DQA interferometry using the adiabatic master equation illustrating oscillations in ground state probability $P_{GS}$ as a function of anneal time, $t_{anneal}$ . . . . .	181
4.15	Plots of the parameters fitted to the results in Figure 4.14 using Equation 4.39 for various temperatures, $T$ , and bath-coupling constants, $\eta$ . . . . .	182

---

# List of Tables

---

2.1	Outline of the various parameter-setting methods used in testing and their given time complexity. . . . .	75
2.2	Reverse annealing experiments used to reproduce oscillations in ground state probability $P_G$ , where various schedules are used to deduce the physical premise for these oscillations. . . . .	105

# Chapter 1

---

## Adiabatic Quantum Computing & Quantum Annealing

---

Simulating nature and its phenomena was the original inspiration behind the conception of the quantum computer proposed by Feynmann, Benioff and Manin [1–3] in the early 1980s due to the classical Turing machine being unable to simulate quantum processes efficiently. The quantum computer’s ability to solve problems (not just in quantum physics) was formalized by Deutsch [4], after which a host of algorithms demonstrating speed-ups over classical computers were developed [5–9]. It is this ability of being able to solve ‘hard’ problems that drives research in quantum computing to develop algorithms that exhibit a “quantum advantage”. This is where the problem size increases for a specific class of problem (e.g., discrete logarithm problem), the runtime of the quantum algorithm scales better (i.e., lower computational complexity) compared to any known classical algorithm. Experimental realizations however face huge challenges that prevent quantum computers being large enough to solve industrially relevant problems. Despite this, there are sufficiently large quantum computers that demonstrate a quantum advantage over modern classical computing [10].

Quantum computing has several domains that differ in the methods used to explore the Hilbert space where the problem solution lies, and are equivalent with some polynomial overhead [11]. The most prominent domain is that of *gate-based*

*quantum computing*, where an input state is guided through the Hilbert space to the solution of the problem by applying sets of unitary gates. However, the domain that will be the focus of this work will be that of *adiabatic quantum computation* (AQC), where the adiabatic algorithm proposed by Farhi *et al.* [12] can be used to find solutions to optimization problems provided that we meet the conditions of the adiabatic theorem [13] by remaining in the ground state (the lowest energy/cost state) at all times. As a result, the system in AQC is evolved continuously through the Hilbert space, rather than discretely as in gate-based quantum computing.

The adiabatic algorithm is primarily used to find solutions to specific combinatorial optimization problems that lie in the complexity class NP, where on a classical computer it is not possible to find a solution to these problems in polynomial time. Adiabatic quantum computing can provide a general speed-up if the adiabatic condition (evolving the system slow enough to remain in the ground state at all times) is maintained, but this speed-up is uncertain for finite temperature systems that undergo decoherence due to the presence of noise. These systems are referred to as open-quantum systems, and all physical manifestations of quantum computers fall under this category of quantum system. Using the adiabatic algorithm at finite temperature and not insisting on adiabaticity or universality is commonly known as quantum annealing. This is where we initialize the system in a quantum state and anneal to a classical state, which is comparable to traditional annealing, where one anneals from a hot to a cold state. If the adiabatic condition is maintained, then the final classical state would correspond to the optimal solution to the optimization problem. However, the solution found using quantum annealing is typically suboptimal due to the violation of the adiabatic condition, which tends to occur when a problem possesses an avoided-level crossing between the ground state and higher excited states. The scaling of a problem therefore depends on how the crossing gap scales in the system size, with larger gapped problems being easier to solve.

The technological leader in quantum annealing hardware is D-Wave Systems Inc., where one can solve optimization problems with up to 5640 binary variables [14–25]. The practical performance of this hardware remains an area of active research

because of the detrimental effects that the experimental compromises have on the quantum system. In addition to noise and avoided crossings, there are restrictions on the number of qubits and their connectivity on the hardware, and techniques such as problem decomposition and embedding are needed to overcome this at the cost of it negatively impacting computation [26, 27]. The culmination of these issues has meant that no computational advantage for quantum annealing optimization has been explicitly shown [28], and there are classical algorithms such as quantum Monte Carlo that exhibit similar scaling (but with larger prefactors) to quantum annealing [29]. However, there are hints of a scaling advantage when using a quantum annealer to simulate quantum systems [30], and there are some theoretical problem cases where quantum annealing can outperform its classical counterparts [31, 32].

In this work, we look to address and present methods to help circumvent the three main limitations that hinder successful computation in quantum annealing, which consist of:

1. Hardware restrictions that prevent large and highly connected problems from being implemented.
2. Noise in experimental quantum annealers that causes decoherence and prevents coherent quantum dynamics from occurring.
3. Minimum gaps at avoided-level crossings that close exponentially in problem size and force the system into excited states.

The first limitation is an area of continual development that drives industry to produce more and better connected qubits, and in Chapter 2 various approaches are presented that attempt to mitigate the detrimental effects of embedding and better utilize the current hardware to increase the size of problems that can be solved on a quantum annealer. The second limitation to quantum annealing is also a hardware problem but is focused on the quality of qubits, such that they are more resilient to the multitude of noise sources that disrupt quantum processes, from thermal noise to cosmic rays [33–36]. On the D-Wave processors, the significant noise levels result in coherence times on the order of tens of nanoseconds, which is typically shorter than

the available minimum anneal time. However, tunnelling events are still shown to be present within this timescale in quantum annealing [29, 37, 38], inferring that there is some quantum computational value to be had despite the system experiencing decoherence. The value of incoherent tunnelling is explored in Chapter 3 with a gadget that can also be used to contrast the differences between classical and quantum dynamics in an incoherent quantum system.

The final limitation to computation with quantum annealing is the minimum gap problem, which is an entirely theoretical problem and limits that kinds of problems that can be solved using quantum annealing. The adiabatic theorem places a “speed limit” on the time it takes to pass through these gaps, such that if a gap reduces in size exponentially as the problem scales in size, then the adiabatic run time also increases exponentially. In Chapter 4, we explore a possible solution to this problem by removing the adiabatic condition completely and use excited states in our computation in what is called diabatic quantum annealing (DQA). Specifically, we look at a near-term DQA protocol presented in Ref. [39] for optimization, which can also be used as a test of coherence. Throughout the rest of this chapter, a more in depth introduction to adiabatic quantum computing, quantum annealing and these limitations are given, highlighting the advantages and disadvantages of this form of quantum computation, as well as exploring the current state-of-the-art methods used to solve optimization problems.

## 1.1 Adiabatic Quantum Computation

Both the gate-based and adiabatic models for quantum computation rely on transitioning away from the binary variable known as a *bit*, to a quantum variable known as a qubit (quantum-bit), which can be thought of as a nuclear spin that can have any value within a Bloch sphere [40]. The  $z$ -axis of the Bloch sphere is the computational axis, where the values  $-1$  (spin down) and  $1$  (spin up) are equal to the binary values of 0 and 1 given a change of variables. The simplest matrix notation of the qubit



system and its three dimensions are given by the Pauli matrices,

$$\sigma^x = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad \sigma^y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix} \quad \sigma^z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \quad (1.1)$$

where all matrices are non-commuting (i.e.,  $[\sigma^x, \sigma^y] = \sigma^x \sigma^y - \sigma^y \sigma^x \neq 0$ ) because each has its own unique eigenbasis. The Pauli matrices can be used to construct the Hamiltonian of a quantum system, where adding a qubit to your system increases the size of the Hamiltonian (and therefore Hilbert space) exponentially. For example, a rudimentary Hamiltonian describing two independent qubits with a  $Z$ -component is  $(\sigma^z \otimes \mathbb{1}) + (\mathbb{1} \otimes \sigma^z)$ , where  $\mathbb{1}$  is the  $2 \times 2$  identity matrix and  $\otimes$  is the Kronecker product. Additionally, the state-vector  $|\psi\rangle$  of two qubits in the spin up configuration, where the spin up eigenstate of  $\sigma^z$  is  $|0\rangle = [1, 0]^T$  (the spin down state is  $|1\rangle = [0, 1]^T$ ), is  $|\psi\rangle = |00\rangle = |0\rangle \otimes |0\rangle$ . A further pedagogical overview of the linear algebra describing quantum systems can be found in Ref. [41].

Instead of applying unitary quantum gates to the initial state-vector as is used in the gate model, adiabatic quantum computation evolves an easy to prepare ground state-vector to an unknown ground state using a time-dependent quantum Hamiltonian. The success of AQC relies on the adiabatic theorem [13] which guarantees that when the system is initialized in the ground state, it will remain in the ground state at all times if the Hamiltonian is evolved slowly enough with respect to the system minimum energy gap. This was formalized for optimization in [12, 14, 42] where the adiabatic algorithm evolves a linear combination of non-commuting Hamiltonians, which in its simplest form consists of interpolating between  $H_{initial}$  which has the easy to prepare initial ground state, to  $H_{problem}$  with an unknown ground state that encodes the problem of interest. This total time-dependent Hamiltonian is written as

$$H(t) = A(t)H_{initial} + B(t)H_{problem}, \quad (1.2)$$

where  $A(t)$  and  $B(t)$  are monotonic time-dependent coefficients that allow us to interpolate between the Hamiltonian terms (i.e.,  $\arg \max(A(t)) = 0$  and  $\arg \max(B(t)) =$

$t_{\text{final}}$ ). It was also later shown that AQC is an extension of universal quantum computation [11], where both the gate model and AQC can simulate each other with no more than a polynomial resource overhead. Therefore, there exist adiabatic variants of well known gate model algorithms, such as Grover and Deutsch-Jozsa [43, 44] whose speed-up is equivalent to that in the gate model. There is also an attempt to implement the quantum Fourier-transform on quantum annealers [45], essential to Shor's algorithm [9].

AQC is naturally framed to solve combinatorial optimization problems, with the earliest variants of AQC taking inspiration from adding quantum dynamics to the established simulated annealing algorithm [46–48]. One of the most common models of AQC to use for optimization is the transverse-field Ising model (TFIM), which is a quantum extension of the classical Ising model [49]. The optimization problem is encoded in the longitudinal fields of the classical Ising model to form  $H_{\text{problem}}$ , and then transverse field components are added to form  $H_{\text{initial}}$ . Taking  $\sigma^z$  to be the computational basis when encoding the Ising model in the problem Hamiltonian, the total Hamiltonian is

$$H(t) = -A(t) \sum_{i=1}^N \sigma_i^x + B(t) \left[ \sum_{i=1}^N h_i \sigma_i^z + \sum_{\langle i,j \rangle} J_{ij} \sigma_i^z \sigma_j^z \right] \quad (1.3)$$

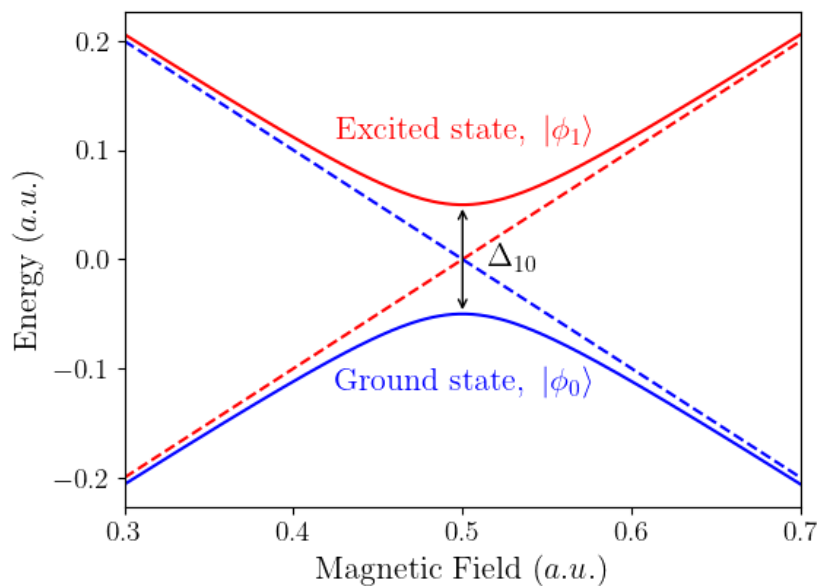
where  $\sigma^x$  represents the transverse-field component of the model,  $h_i$  is the spin bias on the  $i^{\text{th}}$  qubit, and  $J_{ij}$  is the coupling strength between two qubits  $i$  and  $j$ . Also note that because this is an  $N$ -qubit system,  $\sigma_i^x = (\otimes^{i-1} \mathbb{1}) \otimes \sigma^x \otimes (\otimes^{N-i} \mathbb{1})$ ,  $\forall i \geq 1$  and where  $\otimes^k \mathbb{1}$  is Kronecker product of  $k$  lots of identity matrices. The initial ground state for this Hamiltonian at  $t = 0$  (i.e.,  $B(0) = 0$  and  $A(0) \neq 0$ ) is  $|\psi\rangle = |+\rangle^{\otimes N}$ , where  $|+\rangle = [1, 1]^T / \sqrt{2}$  and is an eigenstate of  $\sigma^x$ . The initial state  $|\psi\rangle$  in the TFIM is interpreted as a superposition of all possible solutions to the optimization problem in the  $\sigma^z$  basis.

This class of AQC Hamiltonian is defined as being *stoquastic* (quantum stochastic), where the Hamiltonian always has real non-positive off-diagonal matrix elements with respect to a given basis [50] (typically the computational basis). In the TFIM,

the off-diagonal contributions come from the transverse-field component, which are non-positive in Equation (1.3). If a Hamiltonian were to violate the stoquastic condition, it is known as *non-stoquastic*. Given this definition of stoquasticity, one may naïvely say that changing the sign in front of the transverse-field term in Equation (1.3) will make the off-diagonal elements positive and therefore non-stoquastic, but there also is a unitary transformation that can be applied to rotate the system such that the Hamiltonian will be stoquastic once again. Determining whether a set of local unitary transforms can “cure” an arbitrary non-stoquastic Hamiltonian to be stoquastic again is also known to be NP-complete [51].

If stoquastic AQC is made to follow the adiabatic theorem and remain in the ground state, then it is widely believed that stoquastic AQC is not universal [52], and the ground state can be associated with a classical probability distribution, which is generally attributed to a Boltzmann distribution [53–57]. The non-universality of ground state stoquastic AQC has led to widespread scrutiny over whether a quantum advantage exists for this class of computation. Stoquasticity also has important implications to the simulability of a quantum system, as it has been demonstrated that ground state stoquastic AQC can be efficiently simulated by quantum Monte Carlo (QMC), such that both methods scale similarly with some overhead between them, both theoretically and experimentally [29, 58–61]. However, exceptions exist that possess either topological obstructions (such as a particle restricted to a circular potential) [32, 62] that prevent QMC algorithms from equilibrating, and/or large differences between L1 and L2 normalized wavefunctions [63] that can cause long convergence times in QMC.

In general though, the simulability of stoquastic Hamiltonians is attributed to there not being a sign problem [64]. If a problem is sign-problem free then for all terms in the partition function  $Z = \sum_c W_c$ , it holds that  $W_c \geq 0 \forall c$ , where  $c$  is a configuration of the Hamiltonian  $H$ . This is derived from the common form of the partition function  $Z = \text{Tr}\{e^{-\beta H}\}$  for a system at inverse temperature  $\beta$  (see Ref. [65, 66] for derivation). The simplest example of where  $W_c < 0$  can arise, and therefore a sign problem exists, is when simulating fermionic systems, as the



**Figure 1.1:** Illustration of an anti-crossing in a two-level system, with a minimum gap  $\Delta_{10}$  that forms between the ground state and first-excited state. The solid lines represent a system whose states are coupled via some perturbation, and the dashed lines represent the system without this perturbation. In AQC, the unperturbed system  $H_{\text{problem}}$  is perturbed by the non-commuting Hamiltonian  $H_{\text{initial}}$  (Equation (1.2)).

Pauli exclusion principle states that an exchange of fermions is antisymmetric, i.e.,  $\psi(x_1, x_2) = -\psi(x_2, x_1)$ , and negative weight configurations can arise. Non-stoquastic Hamiltonians also can possess this sign problem, and therefore become NP hard to simulate [65], making them a potential route to quantum speed-ups over classical computation [67, 68] and a potential route to universal adiabatic quantum computation [11, 69]. However, there do exist some instances where non-stoquastic Hamiltonians are simulable [70–72], caused by a vanishing geometric phase of that Hamiltonian.

In addition to the non-universality and simulability of stoquastic AQC, the problem of gap scaling remains to be the most pertinent hindrance to all approaches searching for ground-state (optimal) solutions to optimization problems. Many problems in AQC can possess an avoided-level crossing with a minimum gap  $\Delta_{10}$ , where two energy levels become relatively close to one another but can never cross (see Figure 1.1). Starting in the ground state, if one were to pass through this gap too fast (i.e., non-adiabatically) then Landau-Zener transitions [73, 74] transfer

probability density from the ground state to the excited states, violating the adiabatic condition. The time,  $t_a$ , required to pass through this gap adiabatically [13] with no Landau-Zener transitions is bounded by  $t_a \gg 1/\Delta_{10}^3$  [75] (the worst case) or by  $t_a \gg 1/\Delta_{10}^2$  [76] (the best case if certain conditions are met). Ref. [77] also generalizes the quantum adiabatic theorem for open quantum systems where there is finite temperature and bath coupling. Therefore, the time-to-solution of the adiabatic algorithm relies on  $\Delta_{10}$  scaling polynomially for an exponential speed-up over classical algorithms, but this is not guaranteed [12].

If a problem were to have a minimum gap that closes exponentially (i.e.,  $\Delta_{10} \propto e^{-cN}$ ,  $c > 0$ ), then annealing adiabatically incurs an exponential slow-down within stoquastic AQC, and can be the result of a first-order quantum phase transition [42, 78–84]. This typically occurs when a localized ground state (i.e., a state whose probability density is concentrated on a few computational states) has an avoided-level crossing with a localized first-excited state, whose magnetization phase is distinctly different. However, there are some examples where only polynomial scaling of the energy gap occurs in the presence of a first-order phase transition [85, 86].

Despite the issues that face stoquastic AQC, there have been provable quantum speed-ups that come from removing the restriction of adiabatic evolution. The most well known is the glued-trees problem [31, 87], where the proof exploits the fact that no classical algorithm can find the solution, with high probability, with less than a sub-exponential number of oracles, whereas stoquastic AQC can find a solution in polynomial time. To achieve this, the stoquastic Hamiltonian is initialized in the ground state and evolved towards the first avoided-level crossing where a diabatic transition causes the state-vector to enter the first excited state (breaking the adiabatic condition), and continues to the second avoided-level crossing and then diabatically transitions back to the ground state corresponding to the optimal solution. This approach utilizes the exponential gap scaling problem to its advantage in order to gain a quantum speed-up.

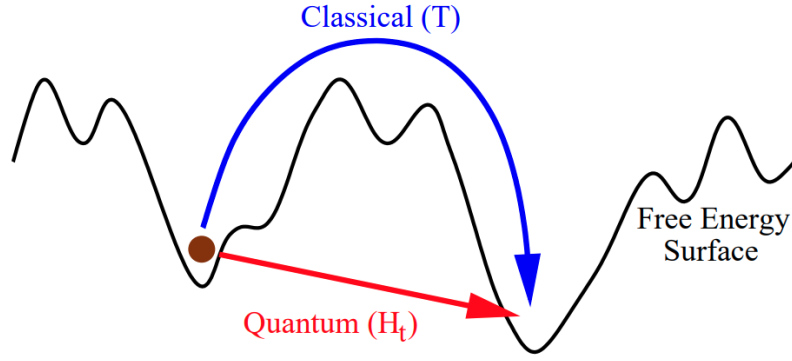
The relaxation of the adiabatic condition to permit evolution with excited states

has also been used to give speed-ups to problems seen in Refs. [67, 70, 88, 89], and for the case of a 3-local stoquastic Hamiltonian that evolves non-adiabatically, it is shown to be equivalent to general AQC and thus universal [90]. Non-adiabatic (diabatic) evolution is a promising approach to counter slow-downs within AQC, but there exist other methods that also focus on navigating the minimum gap efficiently to mitigating its detrimental effects, such as adding extra terms to the Hamiltonian such as a catalyst [67, 70, 81, 91–95], counter-diabatic driving [96–103], inhomogeneous driving of the qubits [39, 104–109], and optimizing evolution schedules [110–112]. In the next section, we focus on a prominent intersection to AQC known as quantum annealing, but for a further overview of AQC and more details about techniques used to overcome slow-downs within AQC, see Ref. [28].

## 1.2 Quantum Annealing

### 1.2.1 Formalism

The key to universality for adiabatic quantum computing (AQC) relies on ensuring adiabatic evolution of the quantum system, which for open-quantum systems cannot be guaranteed. Quantum annealing is a framework that relaxes the adiabatic condition such that the entire spectrum can be used for computation, but at the cost of losing general universality [113]. Quantum annealing (QA) adiabatically at zero temperature and noise is equivalent to AQC, but this is not possible to achieve experimentally. Therefore, it is unknown whether experimental realisations of quantum annealing can provide a quantum advantage. The first instances of QA were considered part of AQC and attempted to build upon the simulated annealing (SA) model developed by S. Kirkpatrick *et al.* [114], where a system is annealed slowly from a “hot” state to a “cold” state such that the system settles in a good (but generally suboptimal) solution to a combinatorial optimization problem (see Appendix A, Alg. 2 for more details). This method relies on using thermal fluctuations to traverse the energy landscape of the optimization problem, whereas in QA, the idea is to use quantum tunnelling to pass through energy barriers that may exist (Figure 1.2). Analogous to SA, the hot state would instead be a fully quantum state, which is annealed to a classical state



**Figure 1.2:** Illustration of classical and quantum dynamics traversing the potential energy landscape, where the landscape changes as a function of the solution to the optimization problem ( $x$ -axis). The displayed mechanisms are analogous to the dynamics used by simulated and quantum annealing, where the objective is to reach the minimum energy (optimal) solution in the landscape. Simulated annealing uses thermal fluctuations parameterized by temperature  $T$ , and quantum annealing uses quantum tunnelling through the energy barriers parameterized by the transverse-field  $H_t$ . Figure taken from Brooke *et al.* [115].

that corresponds to a solution of the optimization problems [46–48, 115]. The first direct comparison of QA to SA hinted that QA performs better than SA [47], and it has since been further verified that a quantum scaling advantage exists by using QA, but not one that quantifies a definitive quantum speed-up (see review [28]).

Both SA and QA can be framed within the Ising model [49], such that optimization problems are formulated as an ensemble of discrete spins,  $s_i \in \{-1, 1\}$ , that mathematically represent complex magnetized systems. The Ising model is related to classical binary optimization by using a change of variables from spins to binary values, such that the Ising problem now constitutes a quadratic unconstrained binary optimization (QUBO) problem. Given a set of spins, the general cost (energy) function for an Ising system with 2-local interactions is

$$\mathcal{E}(\mathbf{s}) = \sum_i^N h_i s_i + \sum_{\langle i,j \rangle} J_{ij} s_i s_j, \quad (1.4)$$

where  $\mathbf{s} \in \{-1, 1\}^N$ ,  $h_i$  is the bias of spin  $s_i$ , and  $J_{ij}$  is the coupling strength between two spins such that  $J_{ij} < 0$  ( $J_{ij} > 0$ ) corresponds to ferromagnetic (antiferromagnetic) couplings. SA minimizes this cost function by using Metropolis-Hastings updates to move to new solutions,  $\mathbf{s}'$ , that incur a lower cost. The formulation of optimization

problems within the Ising model has been generally shown for all of Karp's 21 NP-complete problems [116], as well as some NP-hard instances such as Steiner trees, maximum independent set, and max-cut. These problems are typically defined in QUBO notation using binary variables  $x_i \in \{0, 1\}$ , which form a set of vertices  $V$  on a graph  $G = (V, E)$ . The QUBO cost function is

$$\mathcal{Q}(\mathbf{x}) = \sum_{i \in V} c_i x_i + \sum_{(i,j) \in E} W_{ij} x_i x_j, \quad (1.5)$$

where  $N = |V|$  is the number of variables in our problem, and  $c_i$  and  $W_{ij}$  are the problem weights. Substituting  $x_i = (s_i + 1)/2$  into Equation (1.5) translates the QUBO into the Ising model.

By substituting spin  $s_i$  in Equation (1.4) for Pauli  $Z$  matrices, we create the problem Hamiltonian

$$H_{\text{problem}} = \sum_i^N h_i \sigma_i^z + \sum_{\langle i,j \rangle} J_{ij} \sigma_i^z \sigma_j^z, \quad (1.6)$$

which is a diagonal matrix of size  $2^N \times 2^N$ , whose eigenvalues correspond to the energies,  $e_i$ , of each computational solution,  $|i\rangle$ , such that  $H_{\text{problem}} |i\rangle = e_i |i\rangle$ . The primary goal of QA is to find the eigenstate of  $H_{\text{problem}}$  with the lowest eigenenergy; this is achieved by adding a non-commuting Hamiltonian  $H_{\text{initial}}$  (also known as a driver Hamiltonian) that has a ground state of which is easy to prepare and introduces quantum fluctuations into the system. The simplest model that can be realized is the transverse-field Ising model (TFIM), where  $H_{\text{initial}}$  is the transverse-field driver term composed of Pauli  $X$  matrices local to each qubit, with an easily prepared ground state of  $|+\rangle^{\otimes N}$ . The combination of these terms forms the total TFIM Hamiltonian,

$$H(t) = -A(t) \sum_i^N \sigma_i^x + B(t) \left[ \sum_i^N h_i \sigma_i^z + \sum_{\langle i,j \rangle} J_{ij} \sigma_i^z \sigma_j^z \right], \quad (1.7)$$

where the annealing aspect is driven by the Hamiltonian coefficients  $A(t)$  and  $B(t)$  that monotonically decrease and increase in time, respectively. Quantum annealing




is not restricted to the TFIM Hamiltonian, but its simplicity has meant that it is easier to design hardware for this model compared to other models with more complex terms, which has led to quantum annealing in the TFIM being often simplified to quantum annealing. This simplicity is reflected by the number of qubits that are commercially available from D-Wave Systems Inc. ( $\sim 5640$ ) compared to other quantum computing systems, which is discussed further in the Section 1.3.

Whether the introduction of quantum phenomena holds any general computational benefit is an unresolved question within the field, as it was shown in Ref. [89] that quantum tunnelling is neither necessary nor sufficient for speed-ups solving the perturbed Hamming weight problem. In an open quantum system setting where decoherence has already manifested itself, thermally-assisted quantum tunnelling has been shown to play an important computational role [29, 37, 38, 59, 117] (explored more in Chapter 3). When this phenomenon is shown to occur between energy levels [117], it is therefore beyond what is described by the AQC framework and is specific to quantum annealing. Whether this mechanism can also provide a general quantum advantage is still subject to debate, but there have been improved scaling advantages for specific problem instances [30].

The notion of thermally-assisted tunnelling is distinctly different from that of coherent tunnelling, where a coherent state of qubits can tunnel through energy barriers to remain in the ground state whilst also retaining its phase. When the system is coupled to a bath, noise processes damp the coherent oscillations causing the system to undergo decoherence, and prevents this pure form of quantum tunnelling from occurring. However, the coupling to the environment can provide additional degrees of freedom that are not available in a closed system setting, such that if the coupling to the environment can be seen as a perturbation, then the transition rate between eigenstates is given by Fermi's golden rule

$$\Gamma_{1 \rightarrow 0}(t) \approx \gamma(\omega_{10}(t)) \sum_{i=1}^N \left| \langle \psi_0(t) | \sigma_i^z | \psi_1(t) \rangle \right|^2, \quad (1.8)$$

where constants are set to unity,  $\gamma$  is the noise spectrum, and  $\omega_{10}(t)$  is the en-



**Figure 1.3:** Illustration of the open system thermalization effects that occur as a system approaches a minimum gap. Both excitations and relaxations to and from the ground state follow Equation (1.8), and allow the system to thermally equilibrate. Figure taken from Ref. [37].

ergy gap between states  $|\psi_0(t)\rangle$  and  $|\psi_1(t)\rangle$ . This represents the thermal relaxation rate between states, and the equivalent of the thermal excitation is given by  $\Gamma_{0\rightarrow 1}(t) = e^{-\beta\omega_{10}(t)}\Gamma_{1\rightarrow 0}(t)$ , where  $\beta$  is the inverse temperature. If the thermal relaxation process is shown to pass through a classically forbidden region, then this introduces the notion of thermally-assisted tunnelling (also known as incoherent tunnelling). However, whether the relaxation process (Figure 1.3) in experimental quantum annealers is quantum or classical in nature is still under question, as both quantum [118] and classical [119] dynamics can be used to replicate the processes.

Another quantum phenomenon that can enable computation beyond what is classically possible is entanglement, where in the gate-model it is an important resource for technologies such as quantum cryptography [120]. It is also present in quantum annealing, where qubit tunnelling spectroscopy [121] was used to demonstrate the existence of entanglement [122, 123] using negativity [124, 125] and an entanglement witness [126–128]. However, the computational role that entanglement plays has so far proved to be inconclusive, with increasing entanglement not correlating to changes in success probability [129, 130]. In fact, if the Hamiltonian at the end of

the anneal is diagonal then there exists an analytical upper-bound on ground state probability, which states that if entanglement persists at the end of an anneal then it is detrimental to the ground state probability [129]. Nonetheless, entanglement is present and necessary for many quantum annealing evolutions, and will likely be a key resource in any potential quantum advantage found in quantum annealing.

The TFIM Hamiltonian in Equation (1.7) is also classed as stoquastic and can therefore be simulated efficiently by quantum Monte Carlo [52, 65]. Using path-integral Monte Carlo (PIMC) methods to simulate QA within the Ising model [131–133] is also called simulated quantum annealing (SQA), and in specific circumstances it is an efficient emulator of most stoquastic quantum systems, even those exhibiting quantum tunnelling [58, 134, 135]. Therefore, the performance of most quantum annealing protocols are comparable in scaling to SQA [29, 38, 136–138], with exception to simulating geometrically frustrated magnets where quantum annealing demonstrated a scaling advantage over PIMC [30]. SQA can also be adapted to include open system effects such as dephasing [139–141], which is a form of quantum noise that affects physical implementations of QA. Open system effects can also be explored more precisely with numerical simulations of master equations [142–145], but the intractable scaling of these equations means that only small systems can be analysed. These methods are defined and further discussed in Section 1.4.

From a theoretical perspective, many believe that the current state of quantum annealing cannot be advantageous over classical methods, and several limitations need to be addressed before it can be considered to be potentially successful [28, 146–148]. This includes increased connectivity between qubits (explored in Section 1.3.2), increased coupler order (e.g., 2-local  $\sigma_i^z \sigma_j^z \rightarrow$  3-local  $\sigma_i^z \sigma_j^z \sigma_k^z$ ), noise control (explored in Section 1.3), and error correction. The theoretical success of the gate-model stems from proposals for *fault-tolerant* quantum computing, where error-correction can be used to preserve the quantum state even in the presence of open-system effects, otherwise it cannot function on a large scale [149]. Given that quantum annealing is an analogue protocol, error-correction is also of significant importance for it to be resilient to the ubiquitous open-system effects. Unsurprisingly, there are methods

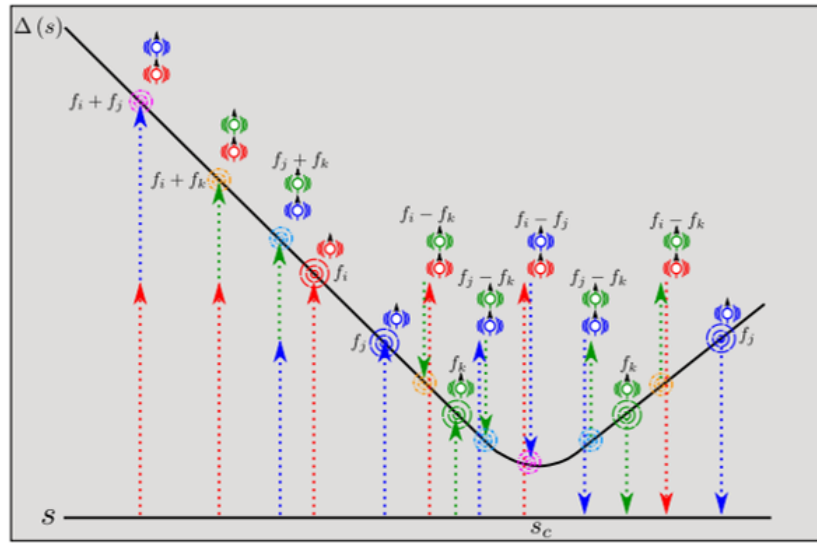
for error correction that exist for adiabatic quantum computation [150, 151] and quantum annealing [152, 153], which reduce the number of available variables for computation, but it is seen to have an improvement on success probability [154].

### 1.2.2 Applications

Despite the numerous problems that hinder quantum annealing (QA), the number of applications that have used QA for computation continues to grow. There are three classes of problems that QA can be applied to: optimization, sampling, and simulation. Optimization was the primary purpose of QA upon its inception [12, 47], but the break-down of the adiabatic condition means that probability density is not restricted to the ground-state, and therefore QA can be thought of sampling states according to some probability distribution. For long anneal times that are quasistatic, the sampled distribution from a quantum annealer was shown to be similar to a Boltzmann distribution [53], where the optimal solution would be the most probable state, i.e.,  $p(s_0) = e^{-\beta E(s_0)} / Z$ . However, QA experiences nonequilibrium dynamics that impedes fair sampling [155–158], such that some states are exponentially suppressed including those that are part of a degenerate ground state [137, 159].

Nonetheless, there have been many successful attempts in finding solutions to a broad range of classical combinatorial optimization problems. These applications range from computer science [14–16], flight gate assignment [18], air traffic management [160], radio networks [19], traffic optimization [20, 22], automated guided vehicles [161], finance [17, 162–164], mathematics [165–168], scheduling [169–171], circuit fault diagnosis [148, 172], chemistry [23–25], and machine learning [173–176]. However, the ground states of certain optimization problems are highly sensitive to the parameters, like the knapsack problem [116], and therefore are detrimentally affected by noise sources.

A promising area of application that is beyond the domain of AQC is sampling, whereby the goal is not to find the ground state precisely, but sample a range of low-energy states that represent a distribution that is difficult to sample from. This is seen as a way towards a quantum advantage with QA [177], an example being when samples from a Boltzmann distribution are required for training restricted



**Figure 1.4:** Schematic of the transitions induced between the ground and first-excited state in RFQA, where spins  $\{i, j, k\}$  have an oscillating driver term (Equation (1.9a)) with frequencies  $\{f_i, f_j, f_k\}$  applied to them. Combinations of the frequencies cause resonance tunnelling events that allow the two states to thermally equilibrate. Figure taken from Ref. [108].

Boltzmann machines (RBM) [178] in machine learning. Sampling from an RBM uses a method called persistent contrastive divergence [179, 180], which when sampled from classically, typically has highly-correlated samples, leading to long mixing times. A quantum annealer can be used to generate less correlated samples to speed up mixing times, which can then be used to speed up the training of RBMs. This premise has been extended to quantum Boltzmann machines as well, and has been used to train a host of deep-learning models [56, 181–187]. Interestingly, the temperature of the Boltzmann distribution generated by a quantum annealer suggests an instance dependent temperature that is higher than the physical system temperature [54, 188], and this is attributed to sampling a system that has frozen before the anneal has finished [53]. Additionally, the training of deep-learning models with QA samples is not as affected by noise sources, largely because noise is typically seen as beneficial to the training process in order to avoid over-fitting.

However, as previously mentioned, having a representative distribution typically relies on fair sampling from that distribution, which is something that quantum annealing cannot achieve in its present form. To remedy this, an extension to quantum annealing, named random-field quantum annealing or radio-frequency

quantum annealing (RFQA) [108, 109], was developed to allow for micro-canonical thermalization of the system during the entire anneal rather than near the minimum gap (see Figure 1.4). This is enabled by an oscillating driver term which is of the form

$$H_{\text{Driver}}(t) = -A(t) \sum_j \left[ \cos(\theta_j(t)) \sigma_j^x + \sin(\theta_j(t)) \sigma_j^y \right] \quad (1.9a)$$

$$\theta_j(t) = \alpha \sin(2\pi f_i t), \quad (1.9b)$$

where  $f_i$  is randomly chosen, and is small enough (i.e.,  $\lesssim \Delta_{10}$ ) to prevent access to higher energy states. It is shown to have a speed-up over traditional adiabatic QA in terms of time-to-solution, particularly near critical points in quantum phase transitions, demonstrating that quantum thermalization can be used as a computational resource.

The final application that QA is used for is simulating both quantum and classical physical systems, which can be seen as a combination of both optimization and sampling (depending on the application). These experiments include simulating the Kibble-Zurek mechanism [189–192], Griffiths-McCoy singularity [193], Shastry-Sutherland model [194], Kosterlitz-Thouless phase-transition [30, 136], and spin-glasses and their phase transitions [195–197]. The theory behind these well-established models also means that experimental data can be compared to the predictions of a purely quantum model. A model that exhibits this clearly is the Kibble-Zurek mechanism, where the kink density (the density of misaligned spins in the system),  $\bar{n}$ , of a ferromagnetic quantum Ising model is proportional to  $t^{-0.5}$ , where  $t$  is the annealing time. If open system effects are prevalent and the system decoheres, then the scaling factor increases (e.g.,  $\bar{n} \propto t^{-0.4}$ ), therefore meaning it can be used as proof of coherence in a quantum system [189, 198]. This idea of using QA to simulate quantum physics aligns with Feynman’s idea of what would be the primary purpose of a quantum computer [1].

### 1.2.3 Reverse Annealing

The traditional QA protocol initializes in a quantum state before annealing to a classical state, but as discussed previously it can struggle in the presence of exponentially

small gaps and noise. Catalysts and tailored coefficient schedules can be used to improve the competitive performance of vanilla QA over classical algorithms, but one variant of QA that can be used in conjunction with these methods is adiabatic reverse annealing (ARA) [199–202]. In this protocol, the system is initialized in a classical state that has been derived by other means, and then has quantum fluctuations applied to it by a driver, with the parameters of the original problem Hamiltonian slowly introduced. For a given anneal time  $t_a$ , the ARA Hamiltonian is defined as


$$H(t) = (1 - s(t))(1 - \lambda(t))H_{initial} + (1 - s(t))\lambda(t)H_{driver} + s(t)H_{problem}, \quad (1.10)$$

where  $s(0) = \lambda(0) = 0$ ,  $s(t_a) = \lambda(t_a) = 1$ ,  $H_{initial}$  encodes an initial classical solution,  $H_{driver}$  introduces the transverse field (quantum fluctuations), and  $H_{problem}$  encodes the cost function we are trying to minimize. If the Hamiltonian is evolved within the adiabatic limit, then it is shown that first order phase transitions are avoided if the initial solution is close in Hamming distance to the ground state (see Figure 1.5), causing the gap to scale polynomially rather than exponentially. The example used in Ref. [200] is the  $p$ -spin model

$$H_p = -N \left( \frac{1}{N} \sum \sigma_i^z \right)^p, \quad (1.11)$$

where  $p \geq 3$ . This Hamiltonian possesses a first-order phase transition if  $s(t)$  and  $\lambda(t)$  are applied incorrectly, but the path restriction is relaxed the closer the initial solution is to the ground state.

In an attempt to realize this physically, D-Wave systems adjusted their quantum annealers to carry out this protocol in the TFIM. However, the Hamiltonian used by D-Wave is still of the same form as the Hamiltonian in Equation (1.7), which is distinctly different from Equation (1.10). Instead, the initial state-vector is set by biasing the qubits with nearby SQUIDs (superconducting quantum interference devices) according to the initial solution. This biasing is then switched off, and the problem Hamiltonian fields are turned on instantly, contrary to ARA where the problem and initial Hamiltonians are slowly increased and decreased respectively




**Figure 1.5:** Static phase diagram for the  $p$ -spin model (Equation 1.11) using adiabatic reverse annealing. The coloured lines indicate the position of a first-order phase transition for a given value of  $s$  and  $\lambda$  (see Equation (1.10)). The value  $c \in [0, 1]$  is a measure of how close the initial state is to the ground state in terms of Hamming distance, with  $c = 1$  being the ground state solution itself. If the initial solution is close enough to the ground state ( $c = 0.8, 0.9$ ), then the first-order phase transition is “softened” to a second-order phase transition for some values of  $s$  and  $\lambda$ , such that the gap no longer scales exponentially in that regime. Figure taken from Ref. [201].

during the anneal. We refer to this method as reverse annealing (RA), and if this method were made to follow the adiabatic condition, then no ground state would ever be found as no transitions between eigenstates would occur. However, quantum annealing typically breaks the adiabatic condition, and given that ground state/lower energy solutions are found experimentally [17, 30, 57, 136, 168, 203], it is evident that diabatic transitions are used for computation in RA.

An illustration of the RA schedule implemented on a D-Wave annealer is seen in Figure 1.6, whereby  $s = 1$  corresponds to a classical system,  $s = 0$  corresponds to a quantum system, and  $s^*$  is the value of  $s$  that the reverse anneal is annealed to before then returning to a classical state. The lower the value of  $s^*$ , the more transverse-field is applied to the system, which is represented by the ratio of the coefficients in Equation (1.7),  $Q(s) = A(s)/B(s)$ . The value of  $s^*$  can also be interpreted as the range of the local-search. However, it was suggested in Ref. [203] that the range of the local search is sensitive to the noise present on a quantum annealer, with





**Figure 1.6:** Plot of the annealing control parameter,  $s$ , in a reverse anneal schedule on a D-Wave quantum annealer. The anneal begins and finishes at  $s = 1$ , and linearly interpolates to different values of  $s$ , with the minimum denoted as  $s^*$ . A pause mid-anneal is used to maintain a value of  $Q(s^*)$  for time  $t_{\text{pause}}$ , where  $Q(s) = A(s)/B(s)$  is a proxy for the magnitude of quantum fluctuations induced in the system using the coefficients from Equation (1.7).

a lower noise system having a larger search radius. This was explored with two experimental annealers each with different noise profiles, but it is likely that an optimal amount of noise for RA exists that is out of the range of these annealers, where too little noise would not allow diabatic transitions, and too much would obstruct the relaxation process. The schedule in Figure 1.6 also features a pause mid-anneal, which is a simple method of observing thermalization effects at different points in an anneal as the relaxation rate to the ground state is dependent on the value of  $s^*$  (see Section 2.3). If the system is paused near the minimum gap, the ground state probabilities are seen to increase due to thermalization effects driving the diabatic transitions [57], and therefore pausing with reverse annealing can be seen as a method of detecting the minimum gap position during an anneal.

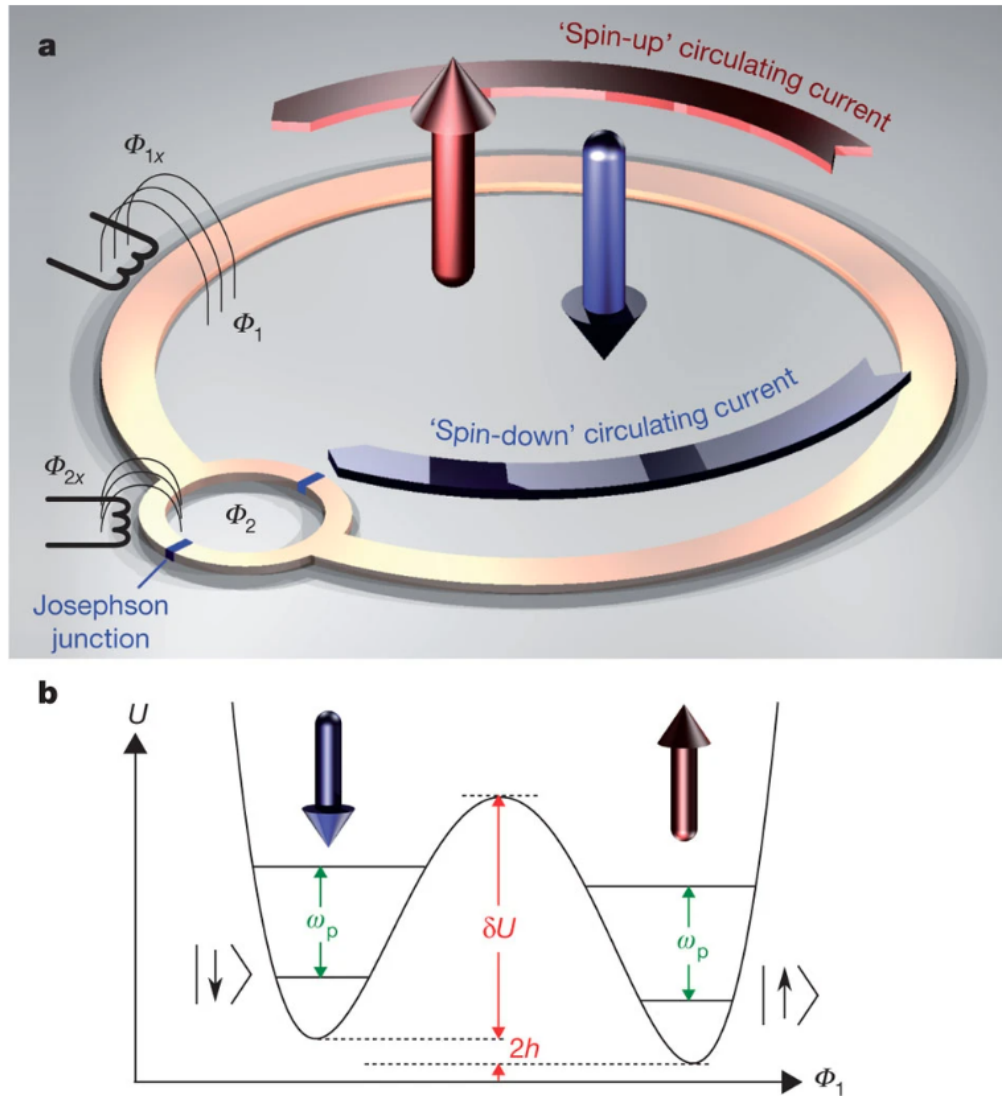
## 1.3 Quantum Annealing using Superconducting Hardware

### 1.3.1 Experimental Implementation

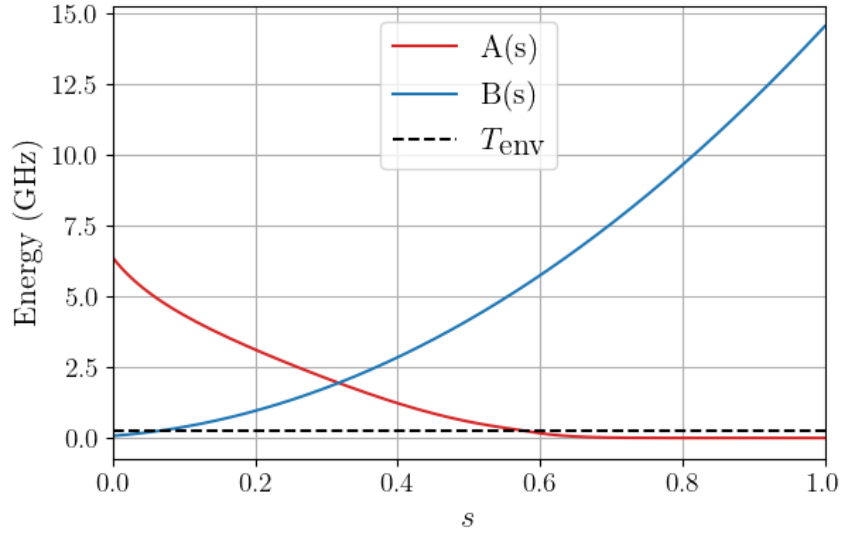
Artificial Ising spin systems can be realized within a host of different quantum models, such as trapped ions [204], but the most common is the analogue superconducting flux qubit model [205–221]. This is where the spin of a qubit is encoded into the flux of a superconducting loop, such that a spin up state  $|\uparrow\rangle$  corresponds to an anticlockwise circulating (persistent) current, and vice versa for a spin-down state  $|\downarrow\rangle$ . A schematic of the superconducting flux qubit is given in Figure 1.7a, where the two loops  $\phi_1$  and  $\phi_2$  correspond to the  $Z$  and  $X$  components of the qubit, respectively. At the beginning of an anneal where the transverse-field (flux through  $\phi_{2x}$ ) is large, the potential in Figure 1.7b is unimodal, allowing for a superposition of both  $|\uparrow\rangle$  and  $|\downarrow\rangle$  states. As the transverse-field decreases,  $\delta U$  increases and a double-well potential starts to form, each well corresponding to the computational spin-state of the qubit.

To use the qubits for computation, longitudinal fields and couplers must be introduced that can be easily parameterized. The local bias parameter  $h_i$  for qubit  $i$  is controlled by the external field  $\phi_{1x}$  in Figure 1.7, which lowers the energy of one of the potential wells with respect to the other by  $2h_i$ . The coupling between qubits is introduced by a programmable inductive coupler that is almost identical to the qubit in Figure 1.7a, but there is only control over  $\phi_{2x}$ , which controls the coupler strength  $J_{ij}$  and whether it is ferromagnetic or antiferromagnetic. As it is inductive, a change in persistent current in one qubit will induce a current in the coupler, affecting the current in another inductively connected qubit.

This framework of using superconducting flux qubits with inductive couplers is what is made commercially available to the public by the quantum annealing



**Figure 1.7:** Schematic of **a** the superconducting flux qubit and **b** its potential energy diagram. Here, the direction of the flux in  $\phi_1$  and  $\phi_2$  corresponds to the spin sign in the Z and X bases, respectively. The circulating (persistent) current in  $\phi_1$  can encode the spin-up  $|\uparrow\rangle$  and spin-down states  $|\downarrow\rangle$  simultaneously, creating the necessary superposition for quantum annealing. External biases  $\phi_{1x}$  and  $\phi_{2x}$  control the values of the bias  $2h$  and barrier height  $\delta U$  in the potential energy diagram, respectively, such that increasing  $\phi_{1x}$  increases the energy difference between the two wells, and an increased  $\phi_{2x}$  decreases  $\delta U$ . The intra-well energy spacing is  $\omega_p$ , and the qubit is measured by its magnetization, which should not be affected by intra-well excitations, improbable as they are. Figure taken from Ref. [207].



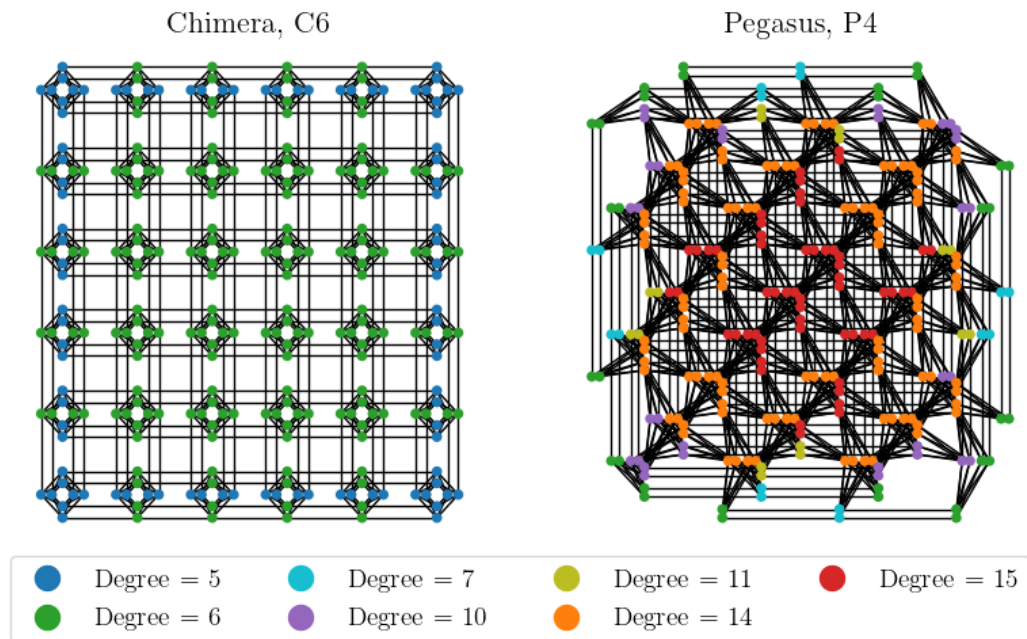
**Figure 1.8:** Plot of the D-Wave 2000Q annealing schedule at the Los Alamos National Laboratory for the Hamiltonian coefficients in Equation (1.12), as a function of the control parameter  $s$ . The black dashed is the environmental temperature of the qubits, which is 12.26mK.

company D-Wave Systems Inc. The form of this Hamiltonian is given by

$$H(s) = -\frac{A(s)}{2} \sum_i \sigma_i^x + \frac{B(s)}{2} \left[ \sum_i h_i \sigma_i^z + \sum_{i,j} J_{ij} \sigma_i^z \sigma_j^z \right], \quad (1.12)$$

which is equivalent to the transverse-field Ising model defined in Equation (1.7), but with coefficients scaled by a factor of a half and parameterized by a control parameter  $s$ . The local fields and couplers are also bounded by  $h_i \in [-2, 2]$  and  $J_{ij} \in [-2, 1]$ . The coefficient functions are unique to each D-Wave quantum annealer because of qubit calibration, and an exemplar is given in Figure 1.8 for the D-Wave 2000Q formerly housed at Los Alamos National Laboratory. In every schedule, the longitudinal field coefficient  $B(s)$  is quadratic in order to give a linear change in qubit persistent current as  $s$  changes. Furthermore, the value of  $s$  is further parameterized by the normalized annealing time  $\tau = t/t_a$  given the total annealing time  $t_a$ . An example of how  $s$  can be parameterized with respect to time is given in Figure 1.6, but a vanilla quantum anneal follows  $s(\tau) = \tau$ .

Topological constraints in superconducting hardware mean that all-to-all connectivity between many qubits is unattainable with what is currently technologically



**Figure 1.9:** A diagram of D-Wave’s proprietary Chimera and Pegasus quantum processing unit architectures. Each graph has a maximum qubit degree of 6 and 15, respectively, which yields more usable couplers for the Pegasus 4 (264 qubits and 1604 couplers) than the Chimera 6 (288 qubits and 816 couplers) despite similar numbers of qubits.

feasible. As a result, the arrangement of qubits and their couplers follow either the Chimera or Pegasus architectures in D-Wave quantum annealers [222]. Chimera was the first architecture to be implemented on physical quantum annealers (Figure 1.9), and consists of repeating  $K_{4,4}$  bipartite cells that are connected to adjacent cells, such that a qubit has a maximum degree (number of couplers) of 6. There have been four major iterations of the D-Wave annealer with the chimera topology - D-Wave one (128 qubits), D-Wave two (512 qubits), D-Wave 2X (1024 qubits) and the D-Wave 2000Q (2048 qubits). The next and latest topology that is commercially available is the Pegasus, which is used on the D-Wave Advantage system. Not only is there an increase in the number of qubits to 5640, but the maximum qubit degree has increased to 15, with odd-cycles now also permitted in the graph. See Figure 1.9 for a comparison.

Since January 2019 (when research included in this work commenced), there have been three distinct D-Wave devices that have been used, the first being the D-Wave 2000Q, the second being the low noise variant of the D-Wave 2000Q that

saw an enhancement in tunnelling rates by a factor of  $\sim 7.4$  [223], and the third being the D-Wave Advantage, which can solve the largest problems to date on a quantum annealer. Generally, each iteration of the D-Wave has shown better performance than the last [224]. The Advantage however was recently seen to be noisier than the low-noise 2000Q [225], and this contributes to the debate on how quantum the D-Wave is due to the prevalent quantum and classical noise sources that can obscure important quantum processes [143, 226, 227]. To address this, D-Wave recently demonstrated that the low-noise D-Wave 2000Q can perform coherent quantum annealing for anneal times  $\sim 25ns$  [189], where the experimental results were consistent with coherent Kibble-Zurek and Landau-Zener processes. The next generation topology due to be released in 2023/2024 is Zephyr on the Advantage 2 [228], which is promised to have 7440 qubits with a maximum degree of 20 and features  $K_{8,8}$  bipartite and  $K_4$  complete graph configurations.

Using superconducting qubits for quantum annealing also comes with a number of integrated control errors that must be accounted for when modelling any D-Wave device, due to the fact that they can distort the programmed values of  $h$  and  $J$ , thus changing the problem being solved. The first of these is cross-talk, whereby the flux of nearby qubits can induce additional ferromagnetic couplings and biases that were not explicitly programmed. The model accounting for cross-talk ( $XT$ ) is

$$h_i^{XT} = h_i - \chi \sum_{k \in \mathcal{A}(i)} J_{ik} h_k \quad (1.13a)$$

$$J_{ij}^{XT} = J_{ij} + \chi \sum_{k \in S(i,j)} J_{ik} J_{jk}, \quad (1.13b)$$

where  $\mathcal{A}$  is the neighbourhood function of the original problem,  $S(i, j) = \{\mathcal{A}(i) \cap \mathcal{A}(j)\}$ ,  $\chi$  is the background susceptibility multiplied by the mutual inductance [159],  $h_i^{XT}$  is the adjusted bias with contributions from the bias leaked from adjacent qubits, and  $J_{ij}^{XT}$  is the induced next-nearest neighbour coupling such that if  $J_{ij}$  does not exist in the original problem graph, then  $J_{ij} = 0$ .

The next source of noise is low frequency ( $1/f$ ) noise that is ubiquitous on superconducting flux qubits and is considered a major hindrance to their develop-

ment [220]. This type of noise is used to generally describe the complex random noise sources that are prominent at low frequencies. On quantum devices,  $1/f$  noise can originate from thermal noise processes but also other quantum phenomena such as particle scattering and interference. For example, in quantum annealing this can be significantly detrimental if fields are near-zero for  $h$  and  $J$  due to the increased likelihood of field sign changes, altering the problem being computed. Another source of noise is introduced by the manufacturing inconsistencies between qubits that produce systematic errors to bias and coupling values. Therefore, setting the values of  $h$  and  $J$  on the qubits can be modelled as a stochastic process

$$h_i \sim \mathcal{N}(h_i + \delta_{h_i}, \sigma_{h_i}) \quad (1.14a)$$

$$J_{ij} \sim \mathcal{N}(J_{ij} + \delta_{J_{ij}}, \sigma_{J_{ij}}), \quad (1.14b)$$

where  $\mathcal{N}$  is the normal distribution with arguments of the mean and standard deviation, and  $\delta$  and  $\sigma$  are the systematic error and deviation, respectively. The values for these errors and susceptibilities can be found in the annealer parameter manuals [229]. Negating these systematic errors is achieved through a technique known as *shimming*, which was introduced in Refs. [30, 136]. This is a classical gradient descent method used to calibrate a system of qubits by adding additional flux biases to qubits (with no previously set bias) that ensure the average magnetization of a qubit is centred about zero as much as possible. This calibration step is problem dependent and is necessary for the successful simulation of quantum systems.

Another prevalent source of error is spin-bath polarization, which introduces additional biases and correlations in the superconducting system. This is caused by the parameters of a quantum anneal polarizing the qubit substrate, therefore introducing additional bias mid-anneal. If anneal times are long and sufficient time is not left between anneals, the spin-bath polarization can persist, causing subsequent anneals to be correlated. This is particularly prevalent in reverse-annealing, where the biasing SQUID causes spin-bath polarization that affects the results significantly (see Section 2.3). An effective way to average out the analogue errors in superconducting

flux qubits is by applying a *spin-reversal transform* (SRT), which is a form of gauge transformation that changes the bias and coupling signs whilst still preserving the ground state. For an  $N$  qubit problem, a random transform  $g \in \{-1, 1\}^N$  is generated that changes the parameters

$$h_i^{SRT} = h_i g_i, \quad J_{ij}^{SRT} = J_{ij} g_i g_j, \quad (1.15)$$

such that the spin in the unaltered solution  $s_i$  is recovered from the transformed solution  $s_i^{SRT}$  by  $s_i = s_i^{SRT} g_i$ . Therefore, there are at most  $2^N$  transforms that one can apply to the problem, but in practice only a fraction of anneals have a transform applied (e.g., a different transform for every 10 anneals) due to the increased run-time of using this error-mitigation technique. The final source of error to be discussed is readout error, where noise present on the QPU can impact the spin-value measured in the readout procedure. This is mitigated by introducing a readout thermalization time before the readout procedure begins to allow qubits time to equilibrate with the system to reduce the impact of thermal noise fluctuations.

### 1.3.2 Minor Embedding

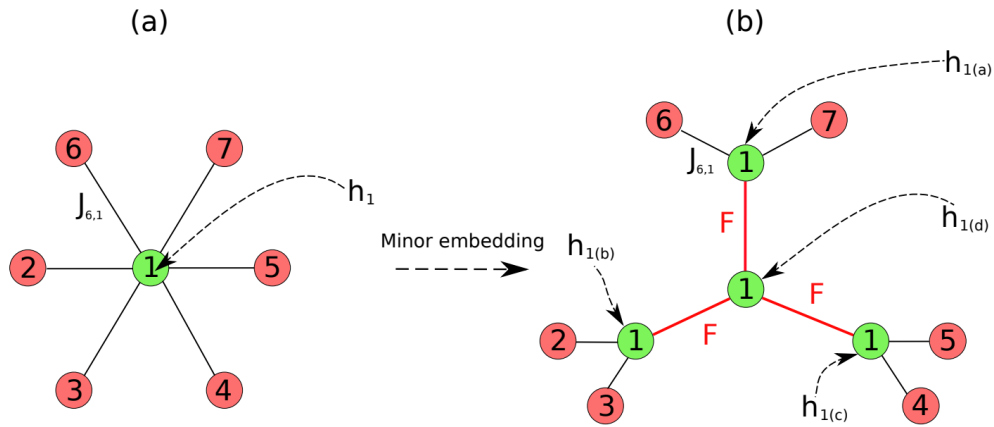
So far, the parameters defined have been for the physical qubits rather than the actual problem, and relating the logical problem parameters defined by the problem Hamiltonian to the physical qubits requires an embedding. If the logical problem graph is a sub-graph of the hardware graph, then a direct embedding exists such that the logical problem qubit is a single physical qubit, and the graph parameters are placed directly onto the QPU if they are within the allowed parameter ranges of  $h_i \in [-2, 2]$  and  $J_{ij} \in [-2, 1]$ . Otherwise, the problem parameters are globally rescaled to meet these ranges. This generally is not a problem unless the parameters are too far out of this range, in which case the differences between the rescaled values can be lost to the noise on the QPU. Embedding to specific subsections of the QPU may also bias the states measured due to the quality of qubits varying across the entire QPU. Averaging across many embeddings can be used mitigate this, and if the embedding is small enough then space-division multiplexing can be used. This



method divides the QPU into many subsections to place many embeddings across it to maximise the qubit usage across the entire QPU, increase the number of measured samples per quantum anneal, and also mitigate embedding biases.

Unfortunately, many logical problem graphs (particularly those with constraints [116]) cannot be directly embedded onto the hardware graph, either because they are too large or because no sub-graph exists. If a logical problem graph  $G$  is not a sub-graph of the hardware  $H$ , a process called minor-embedding is required to find a mapping  $I$  that will create a new graph  $I(G)$  that is a sub-graph of  $H$ . This is done by inserting multiple physical qubits to represent one logical qubit, instead of the one-one mapping in a direct embedding. Finding a minor-embedding is NP-hard [230], but heuristic algorithms such as *minorminer* [230] mean that finding an embedding is tractable. The multiple physical qubits that represent one logical qubit are ferromagnetically coupled to act as a collective spin, with an illustration shown in Figure 1.10. For large problems, this can result in many physical qubits forming one logical qubit, incurring large coupling strengths relative to the logical problem parameters. Instances that require minor embedding typically are dense graphs, have logical qubits with large degrees of connectivity, and have odd cycles in a graph (Chimera architecture only). To avoid minor-embedding completely, one would have to turn to novel architectures such as the Lechner-Hauke-Zoller (LHZ) annealer, which has all-to-all connectivity using quasi-local interactions [231, 232]. This is technologically beyond current quantum annealers as it requires a 4-local  $ZZZZ$  coupler in order to operate.

Not only does minor embedding increase the problem formulation overhead, but it is also extremely detrimental to the scaling performance of quantum annealing [27], and can also lead to exponential suppression of logical solutions in sampling problems [26]. The impact of embedding increases as chains get longer (i.e., more physical qubits in an embedded logical qubit), where chains can become broken (misaligned physical qubits) which means some form of post-processing is needed to choose a logical spin value. This is why there is a drive to not only increase the number of qubits in a quantum annealer, but to increase connectivity between



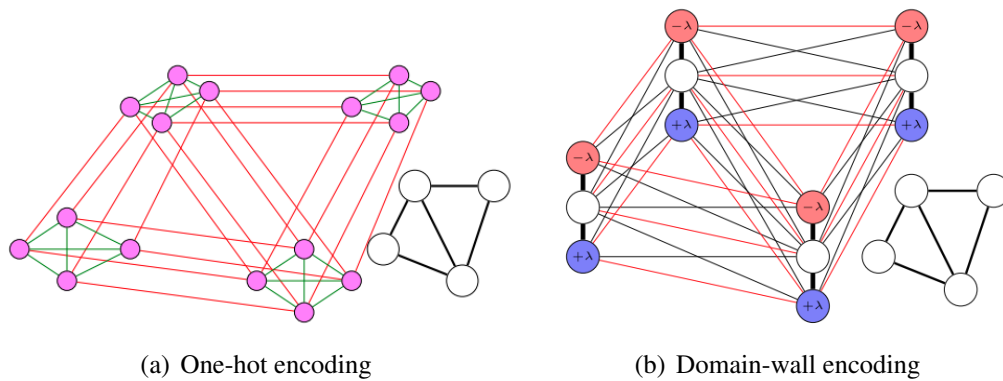
**Figure 1.10:** An Illustration of a minor-embedded logical qubit with degree 6 in (a) to multiple physical qubits (green) with degree 3 in (b). The logical qubit becomes is composed of multiple spins in the embedding (typically called a chain) coupled by a strong ferromagnetic coupling,  $F$ , such that all spins in the logical qubit act as an effective qubit (i.e., aligned at all times). Coupler values external to the embedding remain the same, but the logical qubit bias is distributed amongst the physical qubits. Figure taken from Ref. [233].

qubits to minimize the amount of minor-embedding. Once a minor-embedding is found, the next question is how do we parameterize the additional physical qubits that were added to the problem such that it preserves the ground state of the logical problem. Theoretical bounds to the parameter setting problem have been explored in some detail [233–236], but do not match the experimentally derived result for the Sherrington-Kirkpatrick problem [237] for large problem instances. This is further explored in Section 2.1.

Particularly difficult problems to solve using quantum annealing because of the minor-embedding overhead are those with constraints, which in the Ising model manifest as complete graphs within the original problem graph. A simple example of this is the graph colouring problem, where we want to know whether we can colour a graph (assign colours each vertex in  $V$ ) with  $k$  colours such that no adjacent vertex in the graph  $G = (V, E)$  share the same colour. In QUBO formulation, this has the form [116]

$$Q(\mathbf{x}) = \sum_{v \in V} \left( 1 - \sum_{i=1}^k x_{v,i} \right)^2 + \sum_{(u,v) \in E} \sum_{i=1}^k x_{u,i} x_{v,i}, \quad (1.16)$$

and the constraint here is stipulated by the quadratic term that penalizes the solution



**Figure 1.11:** Examples of a) one-hot and b) domain wall encodings for a single vertex colouring constraint of a four colouring problem (Equation (1.16)). In (a), all couplers are anti-ferromagnetic, and in (b) red edges are anti-ferromagnetic and black are ferromagnetic, with the original graph shown for comparison. The constant  $\lambda > 0$  is suitably large to enforce the domain-wall constraint. Both figures are taken from Ref. [238].

that tries to colour a vertex with more than one colour. This encoding is also called a one-hot constraint, but generally there can be  $K$ -hot constraints. When the quadratic term is expanded, the logical problem to solve now consists of  $K_k$  complete graph at each vertex (see Figure 1.11a), and increases the minor-embedding overhead making it harder to solve.

In order to reduce the impact that  $K$ -hot constraints have on minor-embedding, two alternative constraint encoding schemes have been proposed. The first is a method that encodes the constraint as a domain wall [238], which is shown to have improved performance by comparison with the one-hot method [239], largely due to reduction in constraint variables from  $k$  to  $k - 1$  and there being linear connectivity rather than a complete graph in the domain-wall encoded constraint (see Figure 1.11). The other approach is to use hybrid quantum-classical methods that focus on reducing the quadratic constraint term to a linear term by decomposing the problem to other sub-problems without constraints in polynomial time (see Section 2.2). A general example of this reduction method uses the Hubbard-Stratonovich transformation to reduce the quadratic constraint to a new parameter that can be found using classical optimization, which, when set correctly, will return the ground state solution [240]. These methods should not be confused with *D-Wave Hybrid* methods that use a combination of classical and quantum solvers to find the ground state solution.

## 1.4 Simulating Quantum Annealing

Benchmarking quantum annealing against algorithms that can be performed on a classical computer is an important aspect of identifying whether a quantum advantage can exist for quantum annealing, and in its current state, it is widely believed that no quantum advantage can exist. Despite its success, Quantum annealing is outperformed by parallel tempering Monte Carlo [241–245], the Hamze-de Freitas-Selby heuristic [246, 247] (for Chimera architectures only), and the hybrid cluster method [237]; but it can outperform simulated annealing [60, 114] and population annealing [248, 249]. In this section, we do not focus on methods used to benchmark quantum annealing as an optimization algorithm, but instead survey a range of numerical and stochastic methods that are used to simulate quantum annealing in order to better understand the dynamics that occur on a quantum annealer.

### 1.4.1 Adiabatic Master Equation

To understand how quantum dynamics can be used for computation in quantum annealing, it is important to simulate the system in both open (i.e., coupled to an external bath and with limited coherence) and closed (i.e., perfect quantum system with no decoherence) settings using numerical methods. All quantum systems are typically described by a wavefunction  $|\psi\rangle$ , whose mechanics are described by the Schrödinger equation [250]. The Schrödinger equation can also be written in density-matrix formalism, such that it contains information about the probability densities of each eigenstate in the system. The wavefunction can be composed as a sum of many eigenstates of the system  $|\psi\rangle = \sum_i a_i |i\rangle$ , such that the density matrix of this system is simply  $\hat{\rho} = \sum_i p_i |i\rangle\langle i|$ , where  $p_i = |a_i|^2$  is the probability of being in computational state  $|i\rangle$ . For a time-dependent Hamiltonian  $\hat{H}(t)$  and therefore time-dependent state  $|\psi(t)\rangle$ , the dynamics are described by the von-Neumann equation

$$\hbar\dot{\hat{\rho}}(t) = -i [\hat{H}(t), \hat{\rho}(t)] , \quad (1.17)$$

which is equivalent to the Schrödinger equation in the dynamics it describes. From now on all constants such as  $\hbar$  are set to unity. When a system Hamiltonian such as

the transverse-field Ising model (Equation (1.7)) is the total Hamiltonian, it describes a closed quantum system which can be numerically approximated for small system sizes using

$$\hat{\rho}(t) = U\hat{\rho}(0)U^\dagger \quad (1.18a)$$

$$U = T_+ \exp\left\{-i \int_0^{t_f} H(t) dt\right\} \quad (1.18b)$$

where  $T_+$  is the forward time-ordering operator. However, when a bath is added to the system, the total Hamiltonian  $\hat{H}$  now includes  $\hat{H} = \hat{H}_S + \hat{H}_I + \hat{H}_B$ , where the terms are the system, interaction and bath Hamiltonians, respectively. The total Hamiltonian can be simplified by transforming  $\hat{H}$  into a rotating frame defined by unitary  $U(t)$  such that,  $\tilde{H} = U^\dagger(t)\hat{H}U(t)$ , where  $U(t)$  acts non-trivially on the bath Hamiltonian term to remove it, leaving  $\tilde{H} = \tilde{H}_S + \tilde{H}_I$ . The Redfield equation [251] is a well-established master equation that chooses this rotating frame to be  $U(t) = U_S(t) \otimes U_B(t)$ , where  $U_S(t)$  is defined in Equation (1.18b) and  $U_B(t) = \exp\{-iH_B t\}$ . However, this master equation is computationally intractable for systems of even a few qubits, which makes it impracticable to simulate quantum annealing with unless the problem can be approximated by a two-level system [38].

Attempts to maintain computational tractability or to capture different coupling regimes in open quantum systems has resulted in a plethora of master equations, which all make various approximations or assumptions [144, 252–259]. An appropriate and relatively tractable method of simulating open-system quantum dynamics in quantum annealing is the adiabatic master equation (AME) [142, 143]. It is a Davies-style master equation [255], where coupling to the bath is introduced through sets of Lindblad operators [256] in the weak-coupling limit. In order to reduce complexity, both the Markov assumption and rotating-wave approximation are also applied to yield

$$\begin{aligned} \dot{\hat{\rho}}_S(t) = & -i [\hat{H}_S(t) + \hat{H}_{LS}(t), \hat{\rho}_S(t)] + \sum_{\alpha\beta} \sum_{\omega} \gamma_{\alpha\beta}(\omega) \left[ L_{\omega,\beta}(t) \hat{\rho}_S(t) L_{\omega,\alpha}^\dagger(t) \right. \\ & \left. - \frac{1}{2} \left\{ L_{\omega,\alpha}^\dagger(t) L_{\omega,\beta}(t), \hat{\rho}_S(t) \right\} \right], \end{aligned} \quad (1.19)$$

where  $H_{LS}$  is the Lamb-shift term (see Ref. [145]),  $\gamma_{\alpha\beta}(\omega)$  is the noise spectral density, and  $L_{\omega,\alpha}(t)$  is the Lindblad operator define by

$$L_{\omega,\alpha}(t) = \sum_{\omega=E_b-E_a} \langle E_a|A_\alpha|E_b\rangle |E_a\rangle\langle E_b|. \quad (1.20)$$

The  $A_\alpha$  term here denotes the coupling of the sub-system  $\alpha$  to the bath. To define the noise spectrum and the  $A_\alpha$  term denoting the operator coupling the sub-system  $\alpha$  to the bath, further assumptions based on the quantum system must be made. Generally, the model of decoherence in open quantum annealing systems typically assumes that all qubits are coupled equally to the bath independently, qubits experience decoherence by dephasing, and that decoherence takes place in (or close to) the instantaneous energy eigenbasis. Therefore, the dephasing on each qubit occurs through  $A_\alpha = \sigma_i^z$ . The noise spectrum is also generally assumed to take the form of a bosonic Ohmic bath

$$\gamma(\omega) = 2\pi\eta g^2 \frac{\omega \exp(-|\omega|/\omega_c)}{1 - \exp(-\beta\omega)}, \quad (1.21)$$

where  $\beta = T^{-1}$  is the inverse temperature,  $\omega_c$  is the cut-off frequency,  $\eta g^2$  is the dimensionless bath coupling strength. The AME presented so far is in the weak-coupling limit, where  $\tilde{H}_S$  dominates  $\tilde{H}_I$  such that decoherence occurs in the energy eigenbasis. However, the other extreme limit where  $\tilde{H}_I$  dominates  $\tilde{H}_S$  is called the singular coupling limit (SCL). This system is represented by the master equation

$$\begin{aligned} \dot{\hat{\rho}}_S(t) = & -i [\hat{H}_S(t) + \hat{H}_{LS}(t), \hat{\rho}_S(t)] + \sum_{\alpha\beta} \sum_{\omega} \gamma_{\alpha\beta}(0) \left[ A_\beta \hat{\rho}_S(t) A_\alpha^\dagger(t) \right. \\ & \left. - \frac{1}{2} \{ A_\alpha^\dagger(t) A_\beta(t), \hat{\rho}_S(t) \} \right], \end{aligned} \quad (1.22)$$

where now the Lindblad operators are simply the system coupling operators  $A_\alpha$ , and  $\gamma_{\alpha\beta}(0) = 2\pi\eta g^2/\beta$  from Equation (1.21). Decoherence for this master equation now occurs in the computational basis [143], and results in a loss of instantaneous quantum ground state coherence, making adiabatic quantum computation almost

impossible due to the final state being fully mixed. This results in the state probability being fully random (i.e.,  $p_i = 1/2^N$ ) for the singular coupling limit, compared to the weak coupling limit that follows a Gibbs distribution (i.e.,  $p_i = e^{-\beta E_i}/Z$ ).

The AME in the weak-coupling limit has served as a good approximation of the open-system dynamics on quantum annealers [39, 117, 119, 159, 260], however, when the energy gap  $\omega$  is sufficiently small, the weak-coupling assumption in the AME starts to break down and therefore its validity fails. This is the case for many hard instances, since the energy gap can be orders of magnitude smaller than the bath temperature. In such a regime the energy levels become broadened due to the stronger coupling to the bath, such that the discrete energy levels should emulate a more continuous potential, similar to the semiclassical picture. Therefore, despite the AME not being able to describe these strong-coupling regimes as accurately as more sophisticated models like the Redfield equation, it serves as a reasonable approximation of the open-system dynamics.

Another limiting factor of both the Redfield equation and AME is the increasing error with respect to the bath coupling strength. To overcome this, models such as the non-interacting blip approximation (NIBA) [252, 259], also called polaron-transformed Redfield equation (PTRE) [144, 145], are designed to operate in these stronger coupling limits, which is thought to better simulate the dynamics of experimental quantum annealing [38, 145, 261]. In the polaron transform, the bath and interaction terms become

$$\hat{H}_B = \sum_{i,k} \omega_{i,k} b_{i,k}^\dagger b_{i,k}, \quad (1.23)$$

$$\hat{H}_I = \sum_i g_i \sigma_i^z \otimes B_i \quad (1.24a)$$

$$B_i = \sum_k \lambda_k (b_{i,k}^\dagger - b_{i,k}), \quad (1.24b)$$

such that the Polaron transformed TFIM Hamiltonian becomes

$$H(t) = A(t) \left[ \sum_i^N \sigma_i^+ \otimes \xi_i^+(t) + \sigma_i^- \otimes \xi_i^-(t) \right] + B(t) H_{\text{Problem}} \quad (1.25a)$$

$$\xi_i^\pm = U_B^\dagger(t) \exp \left\{ \pm \sum \frac{2g_i \lambda_i}{\omega_k} (b_{i,k}^\dagger - b_{i,k}) \right\} U_B(t), \quad (1.25b)$$

where  $\sigma_i^x = \sigma_i^+ + \sigma_i^-$ . If the system is rotated again with a Markov approximation applied, then the AME can be re-derived with a different Lindblad operator [145]

$$L_i^{\omega,\alpha}(t) = A(t) \sum_{\omega=E_b-E_a} \langle E_a | \sigma_i^\alpha | E_b \rangle | E_a \rangle \langle E_b |, \quad (1.26)$$

where  $A(t)$  is the time-dependent transverse-field coefficient. The most notable part of the AME-PTRE is that the noise spectrum now consists of high and low frequency parts that represent a hybrid Ohmic bath [262, 263]. In the polaron frame, the noise spectrum is simply a convolution of high and low frequency components  $G_L(\omega)$  and  $G_H(\omega)$ , respectively, such that

$$\gamma_p(\omega) = \frac{1}{2\pi} \int_{-\infty}^{\infty} G_L(\omega-x) G_H(x) dx \quad (1.27a)$$

$$G_L(\omega) = \sqrt{\frac{\pi}{2W^2}} \exp \left\{ -\frac{\omega - 4\varepsilon_L}{8W^2} \right\} \quad (1.27b)$$

$$G_H(\omega) = \frac{4\gamma(\omega)}{\omega^2 + 4\gamma^2(0)}, \quad (1.27c)$$

where  $\gamma(\omega)$  is the Ohmic noise spectrum (Equation (1.21)),  $W$  is the macroscopic resonant tunnelling linewidth, and  $\varepsilon_L$  is the reorganization energy, which is related to  $W$  via  $W^2 = 2\varepsilon_L T$  by the fluctuation-dissipation theorem. This hybrid model of noise can therefore simulate quantum systems with low-frequency noise in the strong coupling limit, which is ubiquitous in superconducting qubits.

## 1.4.2 Quantum Monte Carlo

The intractable simulability of master equations has impelled the development of classical analogues to large quantum systems, a field commonly known as quantum Monte Carlo (QMC). There exist many variants of QMC, and they derive their formulations from the quantum partitions function defined as  $Z = \text{Tr} \left\{ e^{-\beta H} \right\}$ , where the expectation of an operator at thermal equilibrium is  $\langle A \rangle = \text{Tr} \left\{ A e^{-\beta H} \right\} / Z$ . In this section, we focus on simulating the transverse-field Ising model (Equation (1.7))



with path-integral Monte Carlo, off-diagonal expansion Monte Carlo, and subtypes within each method. Path-integral Monte Carlo (PIMC) [131, 133] is known to be efficient in simulating quantum annealing [29, 30, 58, 134–137], and is derived from a Suzuki-Trotter approximation [264] of an exponentiated Hamiltonian containing non-commuting terms. The partition function of the transverse-field Ising model containing terms  $H_D$  (transverse-field) and  $H_P$  (problem) can be rewritten as

$$\begin{aligned}
Z &= \text{Tr} \left\{ e^{-\beta H} \right\} \\
&= \text{Tr} \left\{ e^{-\beta H/N_\tau} \right\}^{N_\tau} \\
&= \text{Tr} \left\{ e^{-\beta(H_D+H_P)/N_\tau} \right\}^{N_\tau} \\
&= \sum_{\{s^1\}} \dots \sum_{\{s^{N_\tau}\}} \langle s^1 | e^{-\beta(H_D+H_P)/N_\tau} | s^2 \rangle \dots \langle s^{N_\tau} | e^{-\beta(H_D+H_P)/N_\tau} | s^1 \rangle,
\end{aligned} \tag{1.28}$$

where  $\{s^\tau\}$  is the set of all possible spin configurations in the  $k$ th Trotter slice that follow the completeness relation  $\mathbb{1} = \sum_{\{s^\tau\}} |s^\tau\rangle\langle s^\tau|$ . Note that a cyclic boundary condition is imposed such that  $s^{N_\tau+1} = s^1$ . Applying the Trotter break-up formula  $\exp\{-\beta(H_D + H_P)/N_\tau\} \approx \exp\{-\beta H_D/N_\tau\} \exp\{-\beta H_P/N_\tau\}$  introduces an error that is proportional to  $O((\beta/N_\tau)^2)$ . For  $H_D = -A(t) \sum_{i=1}^n \sigma_i^x$  and  $A(t) > 0$ , each term in Equation (1.28) therefore becomes

$$\begin{aligned}
\langle s^\tau | e^{-\beta(H_D+H_P)/N_\tau} | s^{\tau+1} \rangle &\approx \langle s^\tau | e^{-\beta H_D/N_\tau} e^{-\beta H_P/N_\tau} | s^{\tau+1} \rangle \\
&= \langle s^\tau | e^{-\beta H_D/N_\tau} | s^{\tau+1} \rangle e^{-\beta E^{\tau+1}/N_\tau} \\
&= e^{-\beta E^{\tau+1}/N_\tau} \prod_{i=1}^n \langle s_i^\tau | e^{\beta A(t) \sigma_i^x / N_\tau} | s_i^{\tau+1} \rangle.
\end{aligned} \tag{1.29}$$

It can be shown that [131]  $\langle \uparrow | e^{a\sigma^x} | \uparrow \rangle = \langle \downarrow | e^{a\sigma^x} | \downarrow \rangle = \cosh a$ , and  $\langle \uparrow | e^{a\sigma^x} | \downarrow \rangle = \langle \downarrow | e^{a\sigma^x} | \uparrow \rangle = \sinh a$ , and given that the target form of  $\langle s | e^{a\sigma^x} | s' \rangle$  is  $C e^{J_\perp s s'}$ , where  $s$  and  $s'$  are individual spins, it can be shown that [131]

$$\langle s^\tau | e^{-\beta H_D/N_\tau} e^{-\beta H_P/N_\tau} | s^{\tau+1} \rangle = C^n \exp \left\{ J_\perp \sum_i s_i^\tau s_i^{\tau+1} \right\} \exp \left\{ -\frac{\beta}{N_\tau} E^{\tau+1} \right\} \tag{1.30}$$

where  $J_{\perp}$  is a ferromagnetic coupling between  $N_{\tau}$  Trotter replicas along an additional dimension named the Trotter/imaginary time dimension. Both the prefactor  $C$  and  $J_{\perp}$  are given by

$$J_{\perp}(t) = -\frac{1}{2} \ln \tanh \frac{\beta A(t)}{N_{\tau}}, \quad (1.31)$$

$$C^2(t) = \frac{1}{2} \sinh \frac{2\beta A(t)}{N_{\tau}}, \quad (1.32)$$

which then for the full partition function yields

$$Z \approx Z_{\text{PIMC}} = C^{nN_{\tau}} \sum_{\{s^1\}} \dots \sum_{\{s^{N_{\tau}}\}} e^{-\beta \mathcal{E}_{d+1}}. \quad (1.33)$$

The  $d+1$  refers to the additional Trotter dimension added to simulate the properties of the original problem. The total cost function that can be minimized using Monte Carlo is

$$\mathcal{E}_{d+1}(t) = \frac{B(t)}{N_{\tau}} \sum_{\tau=1}^{N_{\tau}} \left[ \sum_{i=1}^n h_i s_{i,\tau} + \sum_{\langle i,j \rangle} J_{i,j} s_{i,\tau} s_{j,\tau} \right] - \frac{J_{\perp}(t)}{\beta} \sum_{i,\tau} s_{i,\tau} s_{i,\tau+1}. \quad (1.34)$$

This method simulates finite-temperature quantum annealing efficiently and scales better than the numerical methods employed in Section 1.4.1. Additional open system effects such as dephasing can also be introduced into this model by including a phenomenological noise spectrum of the form  $\gamma(\omega) = 2\pi\alpha\omega$  [139–141] where  $\alpha$  is the system-bath coupling, which introduces a coupling between all Trotter slices. This supplements Equation (1.34) with

$$\mathcal{E}_B = -\frac{\alpha\pi^2}{2\beta N_{\tau}^2} \sum_{i=1}^n \sum_{\tau=1}^{N_{\tau}} \sum_{\tau'=\tau+1}^{N_{\tau}} \frac{s_{i,\tau} s_{i,\tau+1}}{\sin^2\left(\frac{\pi}{N_{\tau}}|\tau - \tau'|\right)}. \quad (1.35)$$

Despite the success of this algorithm, the error of this method does scale with  $\beta$  and the total Trotter replicas, but it is found in practice that more Trotter slices does not necessarily lead to better results, with there being optimal values for  $N_{\tau}$  [131, 133]. To remove the error, continuous-time PIMC (CT-PIMC) takes the  $N_{\tau} \rightarrow \infty$  limit such that the Trotter dimension is no longer composed of discrete imaginary time

replicas for each individual spin, but now a continuous timeline of length  $\beta$  [132]. A spin is updated in CT-PIMC by inserting domains into its timeline by a Poisson process, and the updates are accepted or rejected by a chosen update process such as Metropolis-Hastings.

The success of PIMC is highly dependent on the method used to update the  $d + 1$ -system. The simplest method is using single-spin Metropolis-Hastings updates, to sweep through all  $nN_\tau$  spins. It is known however that near phase transitions this method suffers from increased autocorrelation times, drastically increasing the time to reach thermal equilibrium that is required to calculate the critical exponents of the system [265–267]. This is remedied by updating clusters of spins instead of single spins, using methods such as the Swendsen-Wang [268] and Wolff cluster updates [269].

The next method to be discussed is the off-diagonal expansion Monte Carlo (ODE) [270, 271], which comes away from path-integral approaches that introduce discretization error by instead approximating the partition function using a Taylor expansion of the exponential. It builds upon the stochastic series expansion Monte Carlo (SSE) [272, 273], by calculating generalized Boltzmann weights instead of sampling traces. Both methods rely on the off-diagonal operators of the Hamiltonian being permutation operators with respect to the computational basis, i.e.,  $\sigma^x |\uparrow\rangle = |\downarrow\rangle$ . For the transverse-field Ising model Hamiltonian  $H = -\Gamma \sum_i \sigma_i^x + H_P$ , the partition function can be expanded using the Taylor series to give

$$\begin{aligned}
 Z &= \text{Tr} \left\{ e^{-\beta(-\Gamma \sum_i \sigma_i^x + H_P)} \right\} \\
 &= \sum_{\{s\}} \langle s | e^{-\beta(-\Gamma \sum_i \sigma_i^x + H_P)} | s \rangle \\
 &= \sum_{\{s\}} \sum_{k=0}^{\infty} \frac{\beta^k}{k!} \langle s | (\Gamma \sum_i \sigma_i^x - H_P)^k | s \rangle \\
 &= \sum_{\{s\}} \sum_{k=0}^{\infty} \sum_{\{V_k\}} \frac{\beta^k}{k!} \langle s | V_k | s \rangle,
 \end{aligned} \tag{1.36}$$

where  $\{V_k\}$  is the set of all product sequences of length  $k$  between  $H_P$  and  $\sigma_i^x$ . The next steps (where we refer the reader to Ref. [270]) involve separating the diagonal

and off-diagonal terms such that  $\{V_k\}$  no longer has any contributions from  $H_P$ . The final form of the partition function becomes

$$Z = \sum_{\mathcal{C}} W_{\mathcal{C}}, \quad (1.37)$$

where the configuration  $\mathcal{C}$  is a set of permutation operators  $\{V_k\}$  that satisfy  $\langle s | \prod_i V_{k,i} | s \rangle = 1$ , such that the generalized Boltzmann weight  $W_{\mathcal{C}}$  is

$$W_{\mathcal{C}} = \frac{(\beta\Gamma)^k}{k!} e^{-\beta E_{\mathcal{C}}}, \quad (1.38)$$

where  $E_{\mathcal{C}}$  is the classical effective energy calculated from the divided differences of the multiset of energies.

The approach of ODE is not to flip spins as we would in PIMC or simulated annealing, but to generate sets of permutation operators that would then permute  $s$  to other spin configurations within  $\langle s | \prod_i V_{k,i} | s \rangle = 1$ . The procedure of generating sets of permutation operators is continued until ODE Monte Carlo has observed convergence for some expectation value such as total internal energy  $\langle H \rangle$ . Note that the ODE approach is typically used to retrieve equilibrium statistics from the quantum system instead of emulating quantum anneals. A more recent extension to ODE is the permutation matrix representation Monte Carlo (PMR) [66], which generalizes the forms presented here to capture more complex models, including those with non-stoquastic terms [274]. Not only does PMR converge faster than SSE for the problems tested (especially for low temperatures), but PMR is interesting from the perspective that when a generalized Boltzmann weight is negative, the problem has a sign-problem. Therefore, PMR can be used to help determine whether a problem is non-stoquastic.

### 1.4.3 Classical Dynamics

Identifying quantum dynamics within quantum annealing relies on ruling out possible classical mechanisms that can explain the observed dynamics. Since quantum annealing was inspired from the simulated annealing algorithm, the first dynamical

comparisons were between simulated and quantum annealing [47, 115], but the distributions of ground states were shown to be uniquely different due to the ground state in quantum annealing following a different potential landscape compared to the fixed simulating annealing landscape [159]. A more representative approach to simulating quantum annealing with classical dynamics is by taking a semi-classical approximation of the potential landscape, in which a classical particle can follow the potential described by the transverse-field Ising model (TFIM) Hamiltonian  $H$  (Equation (1.7)). A semiclassical ansatz  $|\Omega\rangle$  can be derived from the spin-coherent path integral formalism [275], where the amplitude  $\mathcal{A}$  between the state at the end of an anneal at time  $t_a$  and at  $t = 0$  is defined as

$$\begin{aligned}\mathcal{A} &= \langle \Omega(t_a) | T_+ e^{-\frac{i}{\hbar} \int_0^{t_a} H(t) dt} | \Omega(0) \rangle \\ &= \int \mathcal{D}\Omega(t) e^{\frac{i}{\hbar} S[\Omega(t)]},\end{aligned}\quad (1.39)$$

where the action  $S[\Omega(t)]$  is given by

$$S[\Omega(t)] = \int dt \mathcal{L} = \int dt (i\hbar \langle \Omega(t) | \partial_t | \Omega(t) \rangle - \langle \Omega(t) | H(t) | \Omega(t) \rangle), \quad (1.40)$$

and  $\mathcal{L}$  is the Lagrangian equation. The ansatz  $|\Omega\rangle$  is described by the spin coherent state [276]

$$|\Omega(t)\rangle \equiv |\theta(t), \phi(t)\rangle \equiv \bigotimes_{j=1}^N \cos\left(\frac{\theta_j(t)}{2}\right) |0\rangle + e^{i\phi_j(t)} \sin\left(\frac{\theta_j(t)}{2}\right) |1\rangle. \quad (1.41)$$

This formalism parameterizes each spin in the system with two angles  $\theta$  and  $\phi$ , which are the polar and azimuthal angles in a spherical coordinate system, respectively, and provides the basis to the classical approaches used to simulate quantum annealing in the TFIM.

#### 1.4.3.1 Spin-Vector Monte Carlo

The spin-vector Monte Carlo algorithm replaces the Pauli matrices in the TFIM Hamiltonian with  $O(3)$  rotors in the Bloch-sphere. The energy function minimized

using Metropolis-Hastings updates is

$$\mathcal{E}(s) = -A(s) \sum_{j=1}^N \cos \phi_j \sin \theta_j + B(s) \left[ \sum_{j=1}^N h_j \cos \theta_j + \sum_{\langle j,k \rangle} J_{jk} \cos \theta_j \cos \theta_k \right]. \quad (1.42)$$

Restriction to the  $XZ$ -plane of the Bloch sphere (i.e.,  $\phi = 0$ ) recovers the  $O(2)$  SVMC algorithm typically used to simulate quantum annealing [277]. The SVMC algorithm attempts to update all rotor angles,  $\theta$ , in every sweep (i.e. every increment of time). There are two update methods typically used for SVMC:

$$\theta_j^t \in [0, \pi], \quad \theta_j^{t+1} \sim \text{Uniform}(0, \pi), \quad (1.43)$$

$$\theta_j^t \in [0, \pi], \quad \theta_j^{t+1} = \theta_j^t + \min\left(\frac{A(s)}{B(s)}, 1\right) u, \quad u \sim \text{Uniform}(-\pi, \pi). \quad (1.44)$$

The traditional method of update in SVMC is described by Equation (1.43), whereby the new angle is a sample from a uniform distribution from zero to  $\pi$ . A more recent update used to capture additional annealing artefacts such as freeze-out is the transverse-field dependent update method (Equation (1.44)), which we refer to as SVMC-TF [119].

### 1.4.3.2 Spin-Vector Dynamics

Despite Monte Carlo methods being efficient simulators of classical dynamics in the TFIM, it does not have a direct relationship to evolution time and is an approximation to the semiclassical system. A more exact representation of the dynamics can be retrieved from the Euler-Lagrange equations of motion, which extremize the action

$$-\frac{d}{dt} \left( \frac{\partial \mathcal{L}}{\partial \dot{\theta}_i} \right) + \frac{\partial \mathcal{L}}{\partial \theta_i} = 0. \quad (1.45)$$

We look to derive the equations of motion for both  $O(2)$  and  $O(3)$  semiclassical models, and set all constants to unity. Starting with the  $O(2)$  model [278] where  $\phi = 0$  in Equation (1.41), it can be shown that  $\langle \theta | \partial_t | \theta \rangle = 0$ , simply leaving the potential component to the Lagrangian. The Euler-Lagrange equation therefore

yields

$$-\frac{\partial}{\partial \theta_i} \langle \theta(t) | H(t) | \theta(t) \rangle = 0, \quad (1.46)$$

where  $V_{SC} = \langle \theta | H | \theta \rangle$  plays the role of the semi-classical potential. This can be interpreted as the path of least action through the potential. Rotating the system from the  $XZ$ -plane to the  $XY$ -plane changes the spin-vector to  $\vec{\theta}_i = (\sin(\theta_i), \cos(\theta_i))$ . This results in the system of non-homogeneous equations to be solved numerically

$$\frac{d\theta_i}{dt} = \omega_i \quad (1.47a)$$

$$\frac{d\omega_i}{dt} = -\frac{dV_i}{d\theta_i} = A(t) \cos \theta_i + B(t) \sin \theta_i \left[ h_i + \sum_{i,j:i \neq j} J_{ij} \cos \theta_j \right]. \quad (1.47b)$$

To include temperature dependence, the Langevin equation is used to add Brownian motion into the system. The Langevin equation for a particle in a potential is defined as

$$\frac{d\theta_i}{dt} = \omega_i \quad (1.48a)$$

$$\frac{d\omega}{dt} = -\frac{dV}{dx} - \frac{\lambda}{m} \omega + \frac{1}{m} \zeta(t), \quad (1.48b)$$

where  $\zeta$  is a Gaussian random process with moments

$$\langle \zeta(t_1) \rangle = 0 \quad (1.49a)$$

$$\langle \zeta(t_1) \zeta(t_2) \rangle = 2\lambda k_B T \delta(t_1 - t_2). \quad (1.49b)$$

The over-damped Langevin equation for an  $O(2)$  system at finite temperature is therefore

$$\frac{d\theta_i}{dt} = \dot{\theta}_i = \omega_i \quad (1.50a)$$

$$\frac{d\omega_i}{dt} = -\frac{dV_i}{dx} - \lambda \omega_i + \sqrt{2\lambda T} W_t, \quad (1.50b)$$

where  $W_t$  is a Wiener process,  $\lambda$  is the damping constant, and all constants are set to unity. The first term is the force of the potential returning the particle to the minimum, the second is a damping term on the speed of the particle, and the last term

contributes the Brownian motion of the particle according to the system temperature.

The final equations of motion to be derived is for the  $O(3)$  model (Equation (1.41)), where the ansatz is not restricted in the Bloch sphere. As a result,  $\langle \Omega | \partial_t | \Omega \rangle \neq 0$ , and it can be shown that

$$\langle \Omega | \partial_t | \Omega \rangle = -\frac{1}{2} \sum_i (1 - \cos \theta_i) \frac{d\phi_i}{dt}. \quad (1.51)$$

When applied to the Euler-Lagrange equations, this yields

$$\frac{1}{2} \sin \theta_i \frac{d\theta_i}{dt} - \frac{\partial}{\partial \phi_i} \langle \Omega(t) | H(t) | \Omega(t) \rangle = 0 \quad (1.52a)$$

$$-\frac{1}{2} \sin \theta_i \frac{d\phi_i}{dt} - \frac{\partial}{\partial \theta_i} \langle \Omega(t) | H(t) | \Omega(t) \rangle = 0, \quad (1.52b)$$

where each equation can be interpreted as the effective torque with respect to each angle in the spherical basis. This system can instead be converted into a set of Bloch equations by recognizing the following relation

$$\frac{d\vec{M}_i}{dt} = \gamma \frac{d\vec{L}_i}{dt} = \gamma \vec{\tau}_i = \gamma \begin{bmatrix} 0 \\ \dot{\theta}_i \\ \sin(\theta_i) \dot{\phi}_i \end{bmatrix}, \quad (1.53)$$

where  $\vec{M}_i$  is magnetization vector of spin  $i$ ,  $\gamma$  is the gyromagnetic ratio,  $\vec{\tau}$  is torque and  $\vec{L}$  is the angular momentum. Note that there is no dependence on the radius, which therefore has no contribution to angular momentum, i.e.,  $\gamma = 1$ . Rearranging Equations 1.52 yields the substitutions

$$\dot{\theta}_i = 2A(t) \sin \phi_i \quad (1.54a)$$

$$\sin(\theta_i) \dot{\phi}_i = 2A(t) \cos \theta_i \cos \phi_i + 2B(t) \sin \theta_i E_i, \quad (1.54b)$$

where  $E_i = h_i + \sum_{i,j} J_{ij} \cos \theta_j$ . Currently, the magnetization in Equation (1.53) is in the spherical coordinate basis, and a change of basis to the Cartesian equivalents



results in the retrieval of the Bloch equation

$$\frac{d\vec{M}_i}{dt} = \begin{bmatrix} -2B(t)E_i \sin \theta_i \sin \phi_i \\ 2A(t) \cos \theta_i + 2B(t)E_i \sin \theta_i \cos \phi_i \\ -2A(t) \sin \theta_i \sin \phi_i \end{bmatrix} = \vec{H}_i \times \vec{M}_i, \quad (1.55)$$

where  $\times$  denotes the cross product,  $\vec{M}_i = (\sin \theta_i \cos \phi_i, \sin \theta_i \sin \phi_i, \cos \theta_i)$  and

$$\vec{H}_i = -2A(t)\hat{x} + 2B(t) \left[ h_i + \sum_{i,j} J_{ij} \vec{M}_j \cdot \hat{z} \right] \hat{z}. \quad (1.56)$$

In order to include thermal noise contributions in the time-dependent potential, the Landau-Lifshitz-Gilbert equation [279–281] must be used

$$\frac{d\vec{M}_i}{dt} = \left[ \vec{H}_i - \lambda \frac{d\vec{M}_i}{dt} + \sqrt{2\lambda T} W_t \right] \times \vec{M}_i. \quad (1.57)$$

This is a non-linear stochastic differential equation which is hard to numerically approximate. In order to circumvent this complexity, Equation (1.55) is perturbatively inserted to yield a Landau-Lifshitz friction term [279, 282] resulting in the following equation of motion

$$\frac{d\vec{M}_i}{dt} = \left[ \vec{H}_i - \lambda \vec{H}_i \times \vec{M}_i + \sqrt{2\lambda T} W_t \right] \times \vec{M}_i. \quad (1.58)$$

The equation of motion can be separated into the form  $d\vec{M}_i = f(u, t)dt + g(u, t)dW$ , which is the scalar stochastic differential equation to be solved. We now denote  $\vec{M}_i = (m_x, m_y, m_z)$  such that

$$d\vec{M}_i = \begin{bmatrix} -2\lambda A(t)(m_y^2 + m_z^2) - 2B(t)E_i(\lambda m_x m_z + m_y) \\ 2A(t)(\lambda m_x m_y + m_z) + 2B(t)E_i(m_x - \lambda m_y m_z) \\ 2A(t)(\lambda m_x m_z - m_y) - 2\lambda B(t)E_i(m_x^2 + m_y^2) \end{bmatrix} dt + \sqrt{2\lambda T} \begin{bmatrix} m_y + m_z \\ m_x + m_z \\ m_x + m_y \end{bmatrix} dW \quad (1.59)$$

Although these equations of motion have not been as successful as SVMC in simulating the dynamics seen in quantum annealing [138, 159], they still serve as exact

models whose dynamics can be easily understood.

In this chapter, the theoretical framework of adiabatic quantum computing, quantum annealing and the transverse-field Ising model (TFIM) has been introduced, alongside the schema to encode quadratic unconstrained binary optimisation (QUBO) problems to the TFIM Hamiltonian. This enables the translation of Karp's 21 NP-complete problems, a collection of hard problems that scale exponentially in computational complexity, to the TFIM [116] and is explored in more depth for the graph-colouring problem in Chapter 2.2. The experimental procedures and constraints of quantum annealing using D-Wave annealers are detailed, including the challenges posed by limited inter-qubit connectivity and sources of noise and error on the quantum processing unit (QPU). The former challenge is solved by minor-embedding, which enables generic problem graphs to be embedded onto the QPU with the cost of an increased qubit overhead and the question of how to best set the parameters on the additional qubits, which is investigated further in Chapter 2.1. Several models of quantum annealing that can include the ubiquitous error and noise processes are also presented, which attempt to describe the potential dynamical processes in experimental quantum annealing, such as thermalization explored in Chapter 3. These models are represented by both master equation and Monte-Carlo simulations of classical and quantum systems, which are compared throughout the subsequent chapters to contrast between quantum and classical dynamical processes that could impact quantum annealing.

## Chapter 2

---

# Quantum Annealing Optimization

---

To solve an optimization problem using quantum annealing, the adiabatic theorem [13] must be adhered to in order to remain in the ground state at all times during the anneal. However, in experimental quantum annealing this is typically not possible due to decoherence in the quantum system and the existence of exponentially scaling minimum gaps for the hardest of problems. Therefore, despite initializing in the ground state of a quantum system, diabatic transitions occur to excited energy states which correspond to suboptimal solutions to the optimization problem. This has meant that when solving optimization problems on a quantum annealer, the time-to-solution scales exponentially in problem size [29, 60]. Nonetheless, the search continues for quantum annealing algorithms that can scale better than their classical counterparts, despite the exponential scaling. This has been demonstrated in the field of quantum simulation, with hints of a scaling advantage observed [30], but this success requires intimate knowledge of the quantum annealer machine to correct integrated control errors and adjust the problem to meet the hardware constraints, meaning that a quantum annealer cannot generally be treated as a black box.

For most optimization problems, several transformations are required (see Section 1.3.2) to reach the level where the problem can be implemented on the physical quantum processing unit (QPU). This includes finding the values for the additional parameters introduced in the transformation that preserve the problem. The transform that is particularly detrimental to performance is that of minor-embedding [27],

where the additional qubit overhead (particularly due to constraints) increases problem hardness [26]. In this section, we look to remedy these issues that quantum annealing optimization faces. Firstly, in Section 2.1 we perform a comparison of various parameter-setting methods, including the proprietary D-Wave method, for solving random optimization problems and look at their performance measured by time-to-solution. Secondly, in Section 2.2 we focus on reducing the number of additional qubits (overhead) introduced by minor-embedding by removing constraints through hybrid quantum-classical methods, in particular we focus on reducing the graph colouring problem by transforming the cost function in Ref. [116] to several maximal independent set problems.

Finally, in Section 2.3 we look at a new form of heuristic quantum annealing that uses diabatic transitions for computation, called reverse annealing. The protocol acts as a local search algorithm by initializing the system in a previously found solution (normally by classical means), and then partially turning on the transverse field to allow for transitions to better solutions. We explore the properties of this method by measuring how likely it is for an initial state to reach the ground state as a function of their Hamming distance using both experimental and numerically simulated reverse anneals. We find that reverse annealing can be used for computation, but is plagued by experimental noise sources that bias the system. For a problem that is sensitive to errors in its parameters, it is unlikely that reverse annealing can be used for computation.

## 2.1 Parameter Setting

In quantum annealing, the logical problem graph we want to solve typically does not directly translate (embed) onto the hardware due to topological constraints (see figure 1.9 for hardware topologies). To overcome this, a method called minor-embedding is used to make the problem graph translate onto the hardware by adding additional qubits to the problem (section 1.3.2). The introduction of new qubits and couplers can produce bottlenecks in practical quantum annealing [27] as dense problems typically need far more additional qubits in an embedding, limiting the logical

problem size. The additional physical qubits introduced through an embedding must act as a collective spin to represent the logical qubit in the original problem, meaning that all the new embedded couplers within a logical spin must be ferromagnetic. This leaves the open the question of how do we set the parameters of these new qubit biases and couplers that will be physically implemented on the hardware? This question has not been widely explored in the literature, largely because the heuristic parameter setting normally used (i.e., uniform distribution of the original bias amongst qubits and strong ferromagnetic couplers of the same strength) is seen as enough to preserve the ground state solution in an embedding. However, optimal bounds for the ferromagnetic coupling strength were found to exist experimentally for the Sherrington-Kirkpatrick (SK) model [237] by finding what ferromagnetic coupling strength yields the highest ground state probability for the minor-embedded problem. The SK model does not contain any qubit biases, so the bound does not extend to problems with biases.

Bounds on the ferromagnetic coupling strength exist because the minimum coupling strength should be large enough to ensure that the fully aligned state within an embedding is always energetically favourable. Theoretical bounds to the parameter-setting values were first produced by V. Choi [234], with a tighter bound closer to the experimental result later found by Y. L. Fang [233]; however, neither fully match the experimental bound set in Ref. [237], particularly for large problems. The basis for the bound set by Y.L. Fang [233] exhaustively compares the energies of all configurations in the minor embedding to set the coupling strength large enough to ensure the fully aligned state is the minimum energy state in the embedding. The Fang method relies on the logical qubit bias already being distributed across the physical qubits in the minor-embedding, of which the optimal distribution is not known. The alternative method by Choi [234] does provide a method of setting biases onto the new qubits, and determines this from calculating a quantity that describes whether a qubit is locally determinable. What we will present in the rest of this section is an assessment of the performance (in terms of time-to-solution) of various parameter setting methods finding solutions to Erdős-Rényi graphs with

randomly set logical parameters. The Fang method with uniform biases proves to be the best parameter setting method overall, but given its complexity, it is intractable if the embedding is large (i.e., the logical vertex contains more than 20 spins in the embedding).

Let us define the logical problem graph  $G = (V(G), E(G))$ , to be embedded onto a hardware graph  $U = (V(U), E(U))$ , and if  $G$  is not a sub-graph of  $U$ , then a minor-embedding is required to map  $G$  onto a sub-graph of  $U$ . If a minor-embedding mapping  $I$  exists, then  $I := (\nu, \varepsilon)$ , where

- $\nu : V(G) \rightarrow V(U)$ , where each vertex  $i \in V(G)$  is mapped to a set of vertices  $\nu(i) \subset V(U)$ ,
- $\varepsilon : V(G) \times V(G) \rightarrow V(U)$ , which is the mapping of edges such that for edge  $ij \in E(G)$ ,  $\varepsilon(i, j) \in \nu(i)$  and  $\varepsilon(j, i) \in \nu(j)$ .

If  $|\nu| > 1$  (i.e., there is no direct mapping of the vertex in  $G$  to a vertex in  $U$ ), there must also exist a set of edges  $\tau(i)$  that connect the set of vertices in  $V(U)$  that represent the logical vertex  $i$  such that  $\tau(i) = E(\nu(i)) \subset E(U)$ . It is the parameters within  $(\nu, \tau)$  that we are concerned with when it comes to the parameter setting bounds, where the edges in  $\tau(i)$  are the ferromagnetic couplers used to create the aligned spins in the logical qubit. The bounds depend on the problem parameters in the logical graph, therefore if the logical problem Hamiltonian to be embedded is

$$H_G = \sum_{i \in V(G)} h_i \sigma_i^z + \sum_{ij \in E(G)} J_{ij} \sigma_i^z \sigma_j^z, \quad (2.1)$$

then the Hamiltonian for the mapped graph  $I(G) \subset U$  is

$$H_{I(G)} = \sum_{i \in V(G)} \left( \sum_{k \in \nu(i)} h_{i(k)} \sigma_k^z + \sum_{pq \in \tau(i)} F_i^{pq} \sigma_p^z \sigma_q^z \right) + \sum_{ij \in E(G)} J_{ij} \sigma_{\varepsilon(i,j)}^z \sigma_{\varepsilon(j,i)}^z, \quad (2.2)$$

where  $F_i^{pq}$  is the ferromagnetic coupling strength between vertices  $p, q \in V(U)$  that represent logical qubit  $i$  (to ensure  $p, q$  behave as a collective vertex), and where  $h_i = \sum_{k \in \nu(i)} h_{i(k)}$  such that the energy of the logical Hamiltonian is conserved with

only some offset of  $c = \sum_{i \in V(G)} \sum_{pq \in \tau(i)} |F_i^{pq}|$ . The optimal distribution of  $h_{i(k)}$  is an open question that we will explore further in the section. The exception to this is Choi's method, whereby the biases need to be set in a certain way for the definition of the optimal bound (for Choi's second theorem). This approach requires us to define a classical measure of non-locality  $C(i)$  on qubit  $i$ ,

$$C(i) = -|h_i| + \sum_{j \in \mathcal{A}(i)} |J_{i,j}| \quad (2.3)$$

where  $h_i$  and  $J_{i,j}$  are the qubit biases and adjacent couplers respectively, and  $\mathcal{A}(i)$  is the neighbourhood function that returns a set of nodes adjacent to  $i$  in the graph  $G$ . If  $C(i) < 0$  then the bias dictates the spin of the qubit such that it is locally determined, otherwise the spin of the qubit is determined by the global state (non-local). To define  $h_{i(k)}$ , we also define the function  $\Delta(i, E)$  as the degree of vertex  $i$  within a given set of edges  $E$ , the number of leaves  $l(i)$  in the sub-graph defined by the edge set  $\tau(i)$ , and the original neighbourhood function,  $O(i(k))$ . This function will return a set of adjacent nodes that are in a different logical node within the embedding, i.e., not other nodes within a logical node. Without loss of generality, it is assumed that  $C(i) \geq 0$  (non-local), such that if the local biases are defined as

$$h_{i(k)} = \text{sign}(h_i) \begin{cases} -\frac{C(i)}{l(i)} + \sum_{j \in O(i(k))} |J_{ij}| & \text{if } \Delta(k, \tau(i)) = 1, \\ \sum_{j \in O(i(k))} |J_{ij}| & \text{otherwise} \end{cases} \quad (2.4)$$

then taking from Choi's first theorem  $F_i \leq -|h_i| - \sum_{j \in \mathcal{A}(i, E(G))} |J_{i,j}|$  means that a tighter bound can be set to define Choi's second theorem

$$F_i \leq -\frac{l(i) - 1}{l(i)} C(i). \quad (2.5)$$

The method developed by Fang however yields a bound that is equal or tighter than


Equation (2.5) if and only if the same biases are used as defined in Equation (2.4)

$$\frac{l(i) - 1}{l(i)} C(i) \geq \frac{C(i)}{|\partial W_i|} \quad \forall C(i) \geq 0, \quad (2.6)$$

where  $|\partial W_i|$  is the number of possible domain walls that can form in the minor-embedding of vertex  $i$ , which is always  $|\partial W_i| \geq 1$ . To derive the bound for an embedding using the Fang method, all parameters except for the minor embedding ferromagnetic coupling strengths must be set, from which the energies of all possible spin configurations within the embedding are calculated. To ensure that the fully aligned states are always energetically favourable, it holds that the energy of the all up (down) configuration  $e_\uparrow$  ( $e_\downarrow$ ) should be less than every other spin configuration  $e_i$ , such that  $e_i - e_\uparrow > 0$  and  $e_i - e_\downarrow > 0 \forall i$ . By using PYTHON's *SymPy* library, a bound for the ferromagnetic coupling strength  $F$  can be found generally for any minor-embedded qubit in the problem, as long as there are 20 or fewer spins that form the logical spin (see Ref [233] for a more detailed proof). The Fang method also allows for different bias distributions to be implemented, but can only guarantee that the bound holds if the biases set follow the condition for admissibility, which is defined as guarantee that the parameters used for an embedding do not exclude any possible spin configuration. Examples of possible non-admissible embeddings occur if the condition  $C(i) \geq 0$  is true, and if a single bias distribution is used [236], where a single qubit in the embedding possesses the logical bias  $h_i$ , with all others set to zero.

Both the Choi and Fang methods are more sophisticated than the proprietary D-Wave embedding method, which sets all minor-embedding chain strengths to  $F = -1$ , and uniformly spreads a bias among all qubits that constitute a logical qubit. This heuristically makes it unfavourable to have anti-aligned spins within a minor-embedded vertex, but not with absolute energetic certainty. An example of the methods evaluated is shown in Figure 2.1, where a single logical spin is composed of three physical spins in the embedded system. The sum of the new biases  $h_{1_0} + h_{1_1} + h_{1_2} = h_1$ , and the minor-embedding strengths are calculated for both Choi and Fang methods. The non-locality is  $C(1) = 1.1$  (i.e., an admissible





**Figure 2.1:** Illustration of the minor-embedding of a three spin problem, where the central logical spin consists of three physical spins in the embedding. The local field biases on the embedded spin are  $h_{1_0} = -0.3$ ,  $h_{1_1} = -0.3$ ,  $h_{1_2} = -0.3$  for the uniform bias, and  $h_{1_0} = -0.45$ ,  $h_{1_1} = 0.0$ ,  $h_{1_2} = -0.45$  for the Choi and weighted bias methods. The minor-embedding ferromagnetic coupling strength for D-Wave is  $F = -1.0$ , and  $F = -0.55$  for both Choi and Fang (Choi and weighted bias) coupling strengths. The Fang coupling strength with uniform bias yields  $F = -0.4$ .

embedding) and  $l(1) = 2$ , resulting in  $F = -0.55$  for the Choi method. The Fang methods exhaustively test all combinations in the embedding once biases are set, yielding embedding strengths equal to the Choi method, with exception to the uniform bias setting method, which yielded the lowest minor-embedding strength ( $F = -0.4$ ) out of all methods for this problem.

To evaluate the performance of each method, we explore experimentally (using Erdős-Rényi graph instances) which out of the D-Wave, Fang or Choi methods yields the best time to solution (TTS), defined as

$$\text{TTS} = t_f \max \left( \frac{\ln(1 - p_d)}{\ln(1 - p_G)}, 1 \right), \quad (2.7)$$

where  $t_f$  is the annealing time,  $p_d$  is the desired probability threshold (typically set to 0.99 confidence that the ground state has been found), and  $p_G$  is the measured ground state probability. This measure is commonly used to compare quantum anneals of different annealing times, as longer anneals can yield higher ground state probabilities but not necessarily lower times-to-solution. The ground state solution(s)

for these systems are found using exhaustive searches, thus we do not look at graphs with size greater than 20 vertices. The question of what is the optimal bias setting method was left as an open question by Fang [233], so we will also compare various methods such as Choi, uniform, and weighted bias settings, as each affects the non-locality of the embedded problem differently. The bias of a spin in the weighted method is defined as

$$h_{i(k)} = \frac{h_i}{\sum_{j \in \mathcal{A}(i)} |J_{ij}|} \sum_{j \in \mathcal{O}(i(k))} |J_{ij}|, \quad (2.8)$$

where bias is weighted on qubits with more logical couplers. If the original logical spin has a  $C(i) \geq 0$ , then for all embedded spins, the weighted bias also ensures that

$$C(i(k)) = \left( 1 - \frac{|h_i|}{\sum_{j \in \mathcal{A}(i)} |J_{ij}|} \right) \sum_{j \in \mathcal{O}(i(k))} |J_{ij}| + \sum_{p \in \mathcal{I}(i(k))} |F_i^{pi(k)}| \geq 0, \quad (2.9)$$

where  $\mathcal{I}$  is the internal neighbourhood function of an embedded spin. We exclude the single bias distribution method [236] from the experiments due to the high likelihood of non-admissibility [233]. The final TTS reported here does not include the embedding pre-processing time, the same way that minor-embedding mappings are not included into the TTS, because it is considered part of the problem formulation process and not part of the solution computation.

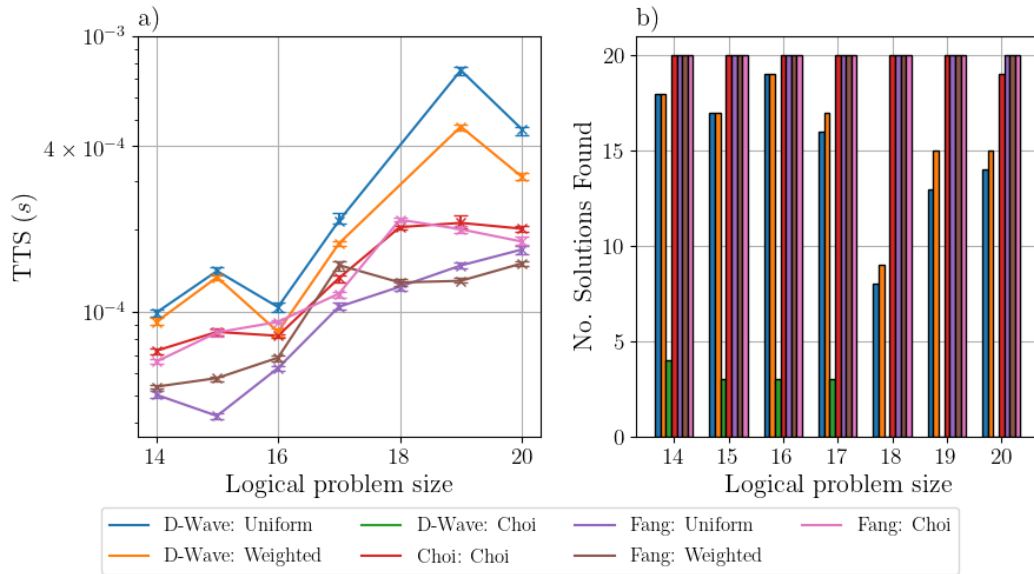
To illustrate the problem instances that can be embedded onto D-Wave architectures, Erdős-Rényi graphs (randomly generated graphs, where an edge is placed between a pair of nodes with some probability  $p$ ) of various sizes with an average node degree of 6 were generated (i.e.,  $p = 6/(N - 1)$ ). The bias and coupler value are randomly chosen from a normal distribution of zero mean and unit variance, and are then normalized to the standard D-Wave parameter ranges of  $h_i \in [-2, 2]$  and  $J_{ij} \in [-2, 1]$ . For logical problem sizes of 14 to 20 vertices, twenty Erdős-Rényi instances were generated (140 instances in total) and then minor-embedded onto the Los Alamos National Laboratory (LANL) D-Wave 2000Q quantum annealer. Ferromagnetic chain strengths and qubit biases were calculated for each of the seven methods for a given embedding. The details and time complexity for each method

Method	Coupling Strength	Bias	Time Complexity	Average $F$
1	D-Wave	Uniform	$\mathcal{O}(1)$	-1.0
2	D-Wave	Weighted	$\mathcal{O}(DL)$	-1.0
3	D-Wave	Choi	$\mathcal{O}(DL)$	-1.0
4	Choi	Choi	$\mathcal{O}(DL)$	-2.008
5	Fang	Uniform	$\mathcal{O}(D2^L)$	-1.960
6	Fang	Weighted	$\mathcal{O}(D2^L)$	-1.854
7	Fang	Choi	$\mathcal{O}(D2^L)$	-2.004

**Table 2.1:** Outline of the various parameter-setting methods used in testing and their given time complexity, where  $L$  is the number of physical qubits in an embedding, and  $D$  is the degree of the logical qubit. Both Choi and Fang methodologies are outlined in Ref. [234] and Ref. [233], respectively, with the D-Wave method simply setting coupling strengths to -1, irrespective of problem information. The ferromagnetic coupling strength  $F$  within an embedding is averaged across all nodes in all 140 instances is also shown to illustrate the differences between each method.

are illustrated in Table 2.1. The additional overheads incurred by each method are not included in the TTS metric, as we focus on the effects of embedding on the quantum computation rather than the classical resources required to run the quantum computation (which are typically significant). All problems have a non-degenerate ground-state and were annealed for  $4\mu\text{s}$ , with no spin-reversal transforms applied to directly compare the naïve implementations of each parameter setting method due to the potential biases introduced from transforming the embedding. For example, the Choi method may have frustration across a ferromagnetic coupler as biases can have opposite signs, and this frustration may be relieved through a gauge transform. Majority-vote post-processing is also applied, where in that case the spins forming the logical spin in the embedding are not all be aligned, the sign of the logical spin is taken to be the most common spin sign in the logical spin.

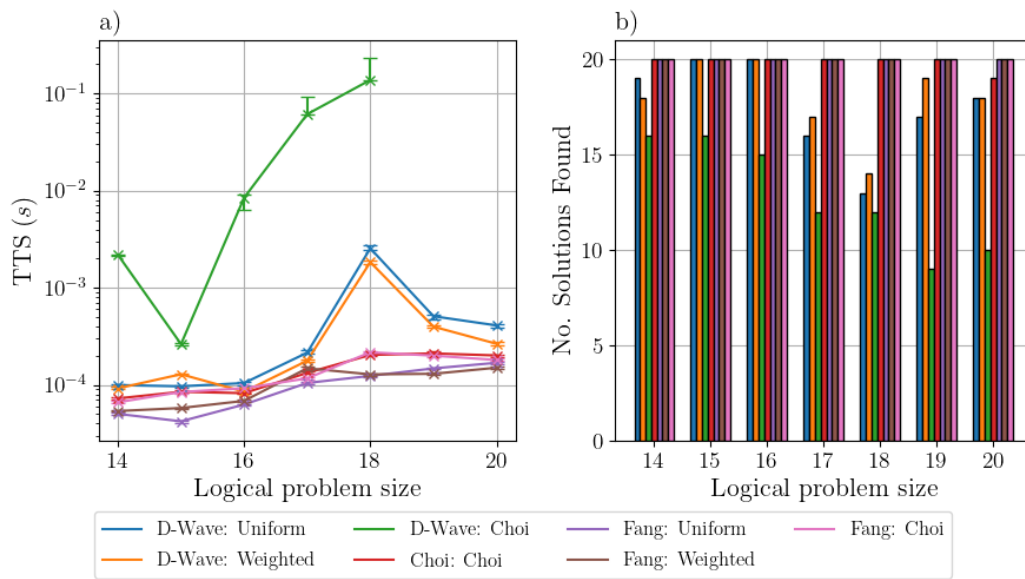
The times-to-solution as a function of logical problem size for the raw (no post-processing) and the majority-vote post-processed solutions for each method in Table 2.1 are shown in Figure 2.2 and Figure 2.3, respectively. The number of solutions found for each logical problem size are also shown to further highlight the differences in each method. The worst performing method was the constant D-Wave chain strength of  $-1$  which set biases with the Choi method, which had an infinite TTS if there was no post-processing. However, upon post-processing there is a large



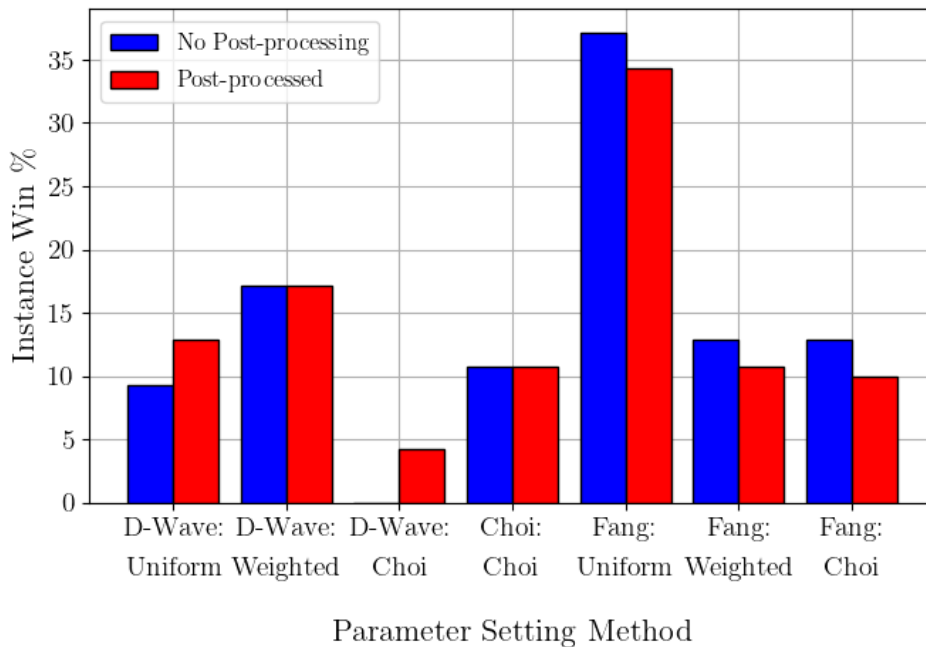
**Figure 2.2:** Data from the LANL D-Wave 2000Q without post-processing showing a) the median time-to-solution (TTS) and b) the count of whether the single ground-state solution was found out of the 20 random instances generated for each logical problem size for each parameter setting method. For all 140 problem instances, the probability of finding the ground-state was calculated from the median of 21 sets of 10,000  $4\mu\text{s}$  quantum anneals with no spin-reversal transforms. The error-bars plotted are the 95% confidence bounds of the median TTS values. A different minor-embedding with random placement on the LANL D-Wave 2000Q was used for each set of anneals. For the TTS,  $p_d = 0.99$  and the median over all problem instances for a given logical problem size is plotted in a). Note that the median TTS for the D-Wave: Choi parameter setting method could not be defined in a), as too few ground-state solutions could be found for the defined problem instances, i.e., less than 11 solutions in b) results in no finite median existing.

recovery in the number of solutions found (Figure 2.3b), which indicates that many of the embeddings are misaligned (broken). This is typical of the Choi bias setting method, as despite the sum of biases of the embedded qubits equalling the original bias on the logical qubit, having both positive and negative biases can frequently occur within an embedding. This creates frustration within the minor-embedding due to all couplers being ferromagnetic, and increases the likelihood of a domain wall forming, which is only remedied through majority-vote post-processing. Using the chain strength set by the Choi method does not suffer from this due to the coupling strength typically being large enough to prevent domain walls from occurring, but it is still out-performed by the other methods.

To better illustrate the performance of all methods, for each of the 140 instances



**Figure 2.3:** Data from Figure 2.2 with post-processing showing a) the median time-to-solution (TTS) and b) the count of whether the single ground-state solution was found out of the 20 random instances generated for each logical problem size for each parameter setting method. Instead of discarding solutions with broken embeddings as in Figure 2.2, majority-vote post-processing was used to assign logical-spin values to broken embeddings. The error-bars plotted are the 95% confidence bounds of the median TTS values.



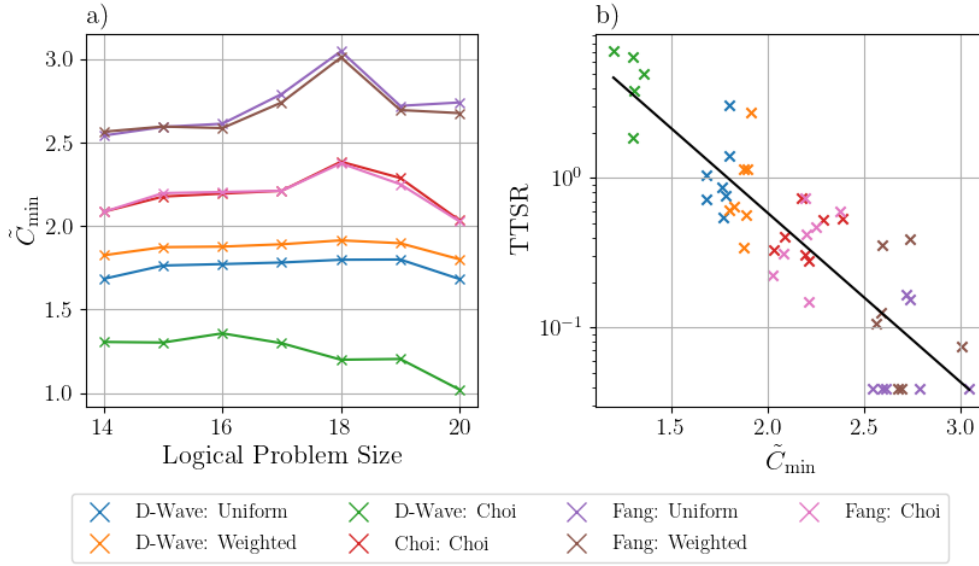
**Figure 2.4:** The best performing parameter-setting method, illustrated by win percentage out of all 140 problem instances, where a win for a method is when it has the lowest TTS for a given instance. Both raw and post-processed data sets are shown.

the best performing method “wins” if it has the lowest TTS out of all methods. The win percentage out of all 140 instances is shown in Figure 2.4 for both raw and post-processed data. Here, the best performing method is the Fang method with uniform biases, with all others performing relatively similarly (apart from method 3 as discussed previously). As seen in Table 2.1, the lowest average coupling strength  $F$  that energetically ensures an aligned embedding is method 6, the weighted bias Fang method. Typically, the lower the value of the coupling strength, the better we expect it to perform due to the problem being effected less by rescaling of the problem to fit the D-Wave parameter range whilst also being considered a tighter bound to what may be experimentally derived [237]. Therefore, to assist in the explanation as to why the uniform Fang method performs best, we look at using the non-local measure  $C(i)$  stated in Equation (2.3) as a proxy for admissibility in the embedded graph,  $I(G)$ . If an embedding has a  $C(i(k)) < 0$ , then its local bias determines the spin-sign and therefore excludes a spin-configuration, hence violating the admissibility condition. To gauge how admissible the embedding parameters are, the average minimum non-locality is found, which we define as

$$\tilde{C}_{\min} = \frac{1}{|V(G)|} \sum_{i \in V(G)} \min \{C(i(0)), C(i(1)), \dots, C(i(|v(i)|))\}. \quad (2.10)$$

The median  $\tilde{C}_{\min}$  across all 20 instances for each method and problem size is shown in Figure 2.5a, and illustrates a consistent ordering between methods. The ordering of  $\tilde{C}_{\min}$  for the D-Wave methods in Figure 2.5a also corresponds to the instance win % order in Figure 2.4. Furthermore, the largest  $\tilde{C}_{\min}$  also corresponds to the uniform Fang method, which has the highest win %. Relating  $\tilde{C}_{\min}$  of each problem to the post-processed TTS data in Figure 2.3 requires a new metric to fairly compare TTS across different system sizes. This will be defined as the time-to-solution ratio (TTSR), where for some parameter-setting method  $m$  out of the 7 methods and for a graph of size  $N$ , it is defined as

$$\text{TTSR}(m, N) = \log A \frac{\text{TTS}_{m,N}}{\min \{\text{TTS}_{1,N}, \text{TTS}_{2,N}, \dots, \text{TTS}_{7,N}\}} \quad (2.11)$$



**Figure 2.5:** Plot of a) average minimum non-locality  $\tilde{C}_{\min}$ , as a function of system size, and b) a fit of TTSR (Equation (2.11)) against  $\tilde{C}_{\min}$  to illustrate the correlation between the variables. The  $\tilde{C}_{\min}$  used is the median out of the 20 graph instances for each problem size, and the TTS data used to calculate the TTSR is taken from the post-processed data in Figure 2.3. The correlation has  $R^2 = 0.750$  and is statistically significant, with a  $p$ -value of  $1.33 \times 10^{-9}$ , which was found from a two-sided hypothesis test whose null hypothesis is that the slope has a zero gradient, using a Wald Test with the t-distribution of the test statistic.

where  $A$  is some scaling factor that ensures  $TTSR > 0$ , as to not infer that the minimum TTS out of all the methods tested is the optimum parameter-setting method generally. Setting  $A = 1.04$ , the TTSR is plotted in Figure 2.5b against  $\tilde{C}_{\min}$  and illustrates the negative correlation of  $R^2 = 0.750$  between the two variables. The statistical significance of this correlation is measured using a two-sided hypothesis test whose null hypothesis is that the slope has a gradient of zero, and a  $p$ -value of  $1.33 \times 10^{-9}$  was found. Note that any apparent missing data points in Figure 2.5b are due to non-finite TTS's.

It must be realized that simply increasing  $\tilde{C}_{\min}$  to reduce the TTSR as suggested by the correlation in Figure 2.5b would be limited in its success as increasing ferromagnetic coupling strength increases  $\tilde{C}_{\min}$  (Equation (2.3)) but also minimizing the ferromagnetic coupling strength is crucial to the experimental success probability [237] as too large a coupling strength will be detrimental to the computation. Considering this interplay, finding bias setting methods that can increase  $\tilde{C}_{\min}$  for a given ferromagnetic coupling strength is shown to be a good proxy for finding

optimal parameters for an embedded random problem. For the proprietary D-Wave method this is shown to be true, as despite it not energetically ensuring the alignment of an embedding, the weighted bias method has the best win rate in Figure 2.4 out of the other D-Wave methods with other bias parameters. This relates back to equation 2.9 that ensures  $C(i(k)) \geq 0, \forall k$  if  $C(i) \geq 0$ , whereas the uniform bias method does not ensure this condition, meaning values of  $C(i(k)) < 0$  can occur, which is non-admissible and lowers the value of  $\tilde{C}_{\min}$  (Figure 2.5a).

In conclusion, a good choice of parameter setting method that takes into account the specific embedding of the problem can lead to better times-to-solution over the proprietary D-Wave approaches, therefore mitigating the detrimental effects of minor-embedding. However, what parameter-setting method performs best is highly dependent on the optimization problem and the embedding itself. Therefore, there is no “one size fits all” optimal parameter setting method, as seen by the general spread of win rates in Figure 2.4. However, out of the parameter-setting methods tested on the 140 Erdős-Rényi instances with randomly set biases and couplers, the Fang method with uniform biases proved to perform the best in terms of having the largest win rate in Figure 2.4, and having consistently low times-to-solution relative to the other methods in both raw and post-processed quantum anneals in Figure 2.2 and Figure 2.3, respectively. However, as seen from Table 2.1, the complexity of the Fang method,  $\mathcal{O}(D2^L)$ , means that it is only computationally accessible if the number of embedded spins  $L$  is less than 21. The heuristic D-Wave methods that are generally used for quantum annealing because of the low computational overhead, often performed worse compared to theoretical methods that ensure aligned spin configurations in an embedding are always the energetically favourable state.


This poor performance was notable when the D-Wave method was combined with the Choi bias setting method, where the majority of solutions couldn't be found without post-processing. This is largely because the Choi biasing method introduces frustration into the embedding by sometimes placing biases of opposite signs across the ferromagnetic coupler within the logical spin, which detrimentally affects all methods that used the Choi biasing method. For the random problems



tested, it was also found that the time-to-solution ratio of all parameter-setting methods correlated with the average minimum non-locality parameter  $\tilde{C}_{\min}$ . This correlation with an  $R^2 = 0.750$  provides a possible reason as to why the uniform Fang method generally performs well despite having a higher average coupling strength compared to the weighted Fang method (Table 2.1). Therefore,  $\tilde{C}_{\min}$  of a parameter setting method can be used as a guide to picking the best method for an embedding, but it should be noted that the interplay with the optimal embedded coupling strength also contributes to performance. Therefore, future work would look to test other classes of problems, as well explore new biasing schemes that ensure admissibility and maximize the minimum non-locality in a way that keeps the embedded coupling strength minimized. Further evaluation is also required in the context of quantum advantage, where the cost of the classical resources when parameter setting will be significant compared to the quantum computation, especially for the Fang parameter setting method.

## 2.2 Graph Colouring

The formulations of Karp's 21 NP-complete problems [283] in quadratic unconstrained binary optimization (QUBO) form (which can then be translated to the Ising model) are presented by Lucas in Ref. [116], and are an important basis for many quantum-annealing optimization problems. However, many of Karp's problems require solutions to fulfil some constraints, which can be achieved by using  $k$ -hot encoding in QUBO form. This encoding stipulates that out of a set of binary variables, only  $k$  variables can have the value 1. This encoding normally comes at the cost of having a quadratic increase in the number of bits as we encode the computational problem into the logical problem, but classically this results in an exponential increase in the difficulty of solving the problem. Furthermore, the variables that form a  $k$ -hot encoding are fully-connected, creating regions of high density within the logical problem. This last feature proves to be detrimental in the final process of embedding the logical problem onto a quantum annealer, which as discussed in the previous section already poses a major obstacle to the success of practical quantum



**Figure 2.6:** Illustration of a 7-vertex problem graph  $G = (V, E)$  correctly coloured with three colours. The minimum colouring of this graph (i.e., the chromatic number) is with two colours.

annealing.

Graph colouring is one of Karp's 21 NP-complete problems (see Figure 2.6), which in QUBO form uses one-hot encoding to ensure that each node in the computational problem is coloured with only one colour. Many optimization problems (particularly within the context of networks) can be formulated as a graph colouring problem, such as frequency/time division multiple access networks [284] and wavelength division multiplexing [285]. Specifically, the NP-complete version of the graph colouring problem asks whether we can colour a graph with  $k$ -colours such that no two adjacent vertices have the same colour; whereas finding the chromatic number of a graph (the minimum number of colours needed to colour it) is an NP-hard problem. In this section, we will show that the one-hot encoded graph colouring (GC) problem can be decomposed into several maximal independent set (MIS) problems, which we refer to as greedy MIS colouring. It is acknowledged that a similar version was recently presented by Kwok and Pudenz in Ref. [286], and we compare our work against that method as well. Overall, this approach to graph colouring uses fewer variables than the encoded version and solves the NP-hard version of this problem, and therefore the NP-complete problem. Removing the one-hot encoding is shown to lead to better scaling as problem size increases, as well as have the practical advantage of being able to implement larger problems on

experimental quantum annealers. The NP-complete version of the graph colouring problem presented by Lucas [116] (hence now called the Lucas method) states that if there is a  $k$  colouring solution  $\mathbf{x}$  for a problem graph  $G = (V, E)$ , the QUBO cost function

$$\mathcal{Q}(\mathbf{x}) = \sum_{v \in V} \left( 1 - \sum_{i=1}^k x_{v,i} \right)^2 + \sum_{(u,v) \in E} \sum_{i=1}^k x_{u,i} x_{v,i} \quad (2.12)$$

equals zero. The squared term represents the one-hot encoding that penalizes any vertices having a colouring  $\neq 1$ . Expanding out this term allows us to write the full QUBO, which represents the logical version of the computational problem,  $\mathcal{L}$ , which has  $k|V|$  binary variables

$$\mathcal{Q}(\mathbf{x}) = \sum_{v \in V} \left( 1 - \sum_{i=1}^k x_{v,i} + 2 \sum_{i < j} x_{v,i} x_{v,j} \right) + \sum_{(u,v) \in E} \sum_{i=1}^k x_{u,i} x_{v,i}. \quad (2.13)$$

Using the equation  $s_i = 2x_i - 1$ , a change of variables is used to turn the QUBO cost function into an energy cost function that can be solved in the Ising model

$$\mathcal{E}(\mathbf{s}) = \sum_{v \in V} \left( \sum_{i=1}^k \frac{2k + \Delta_v - 4}{4} s_{v,i} + \frac{1}{2} \sum_{i < j} s_{v,i} s_{v,j} \right) + \frac{1}{4} \sum_{(u,v) \in E} \sum_{i=1}^k s_{u,i} s_{v,i} + c, \quad (2.14)$$

where  $\Delta_v$  is the degree of node  $v$  and  $c = (|V|(4 + k(k - 3)) + k|E|)/4$ . The quadratic increase in spins from  $|V|$  to  $k|V|$  stems from each vertex in  $G$  now constituting a complete graph of order  $k$ , created by the quadratic constraint.

The sets of colours for a valid graph colouring solution are all disjoint, such that they are by definition independent sets. Therefore, by taking the greedy approach to finding the colour sets, we would want to find multiple maximal independent sets (MIS) to complete the colouring. The MIS QUBO cost function is defined as

$$\mathcal{Q}(\mathbf{x}) = -A \sum_{v \in V} x_v + B \sum_{(u,v) \in E} x_u x_v, \quad (2.15)$$

where  $B > A$  to always maintain the independent set condition. In Ising form,

equation 2.15 becomes

$$\mathcal{E}(\mathbf{s}) = \sum_{i \in V} \left( -\frac{A}{2} + \sum_{j \in \mathcal{A}(i)} \frac{B}{4} \right) s_i + \sum_{(i,j) \in E} \frac{B}{4} s_i s_j + \left( \frac{B|E| - 2A|V|}{4} \right). \quad (2.16)$$

The values of  $A = 1$  and  $B = 2$  are used throughout the rest of this section.

Re-framing the graph colouring problem to the MIS problem is a well known reduction in graph theory, but it is not widely used because finding solutions to MIS problems can be as hard, if not harder in some instances, than finding the colouring of that graph. However, the converse is true for quantum annealers where hardware constraints exist, where it is harder to find solutions to colouring Hamiltonians with quadratic constraints than it is to find solutions to MIS Hamiltonians without quadratic constraints because of minor-embedding issues. The method used to colour a graph with MIS's is defined in algorithm 1, and assigns an MIS to constitute a colour set, removing that set from the graph, and repeating until there are no vertices left in the graph. The number of variables required to find a  $k$  colouring with Alg. 1 is always less than or equal to  $k|V|$ .

---

**Algorithm 1** Greedy MIS Colouring

---

```

1: Graph  $G = (V, E)$ 
2: Colours =  $\emptyset$ 
3: while  $G \neq \emptyset$  do
4:    $M \leftarrow$  Vertices with degree 0  $\in G$ 
5:    $G = G - M$ 
6:   Find maximal independent sets of  $G$ 
7:    $I \leftarrow$  maximal independent set with most adjacent edges in  $G$ 
8:   Colours = Colours +  $\{M + I\}$ 
9:    $G = G - I$ 
10: end while

```

---

If Alg. 1 is to finish in the fewest iterations possible, which not only would return a colouring close to the chromatic number but also reduce computational effort, an additional greedy protocol is required to removing the maximal independent set with the most adjacent edges upon a graph update. The Kwok method [286] is similar in its formulation to Alg. 1 but does not set a condition for choosing the MIS to be removed and does this randomly. The reason for the greedy protocol is given by the

lower and upper bounds on the size of maximal independent sets (equation 2.18), of which proofs are now given. We first start by defining the number of vertices and edges in the problem graph  $G$  as  $n = |V|$  and  $e = |E|$ , respectively, and the number of edges on any graph is upper bounded by Turán's theorem defined in Ref [287]

$$e \leq \left(1 - \frac{1}{r}\right) \frac{n^2}{2}, \quad (2.17)$$

where  $r$  is the size of the largest clique (fully-connected sub-graph) in the  $G$ . The bounds on the MIS size for any graph are now given.

**Theorem 2.2.1.** *Let  $|I|$  be the cardinality of the maximal independent set  $I$  in graph  $G$ , it holds that  $|I|$  is bounded by*

$$\frac{n^2}{2e+n} \leq |I| \leq \frac{1 + \sqrt{1 + 4n^2 - 4n - 8e}}{2} \quad \forall e > \frac{n}{2}. \quad (2.18)$$

*Proof.* Upper bound: The number of edges between vertices that are not in  $I$  are at most  $\binom{n-|I|}{2}$ , and the number of edges that connect vertices in  $I$  to those not in  $I$  are at most  $|I|(n - |I|)$ . There are no edges between vertices in  $I$ . Therefore,

$$\begin{aligned} e &\leq \binom{n-|I|}{2} + |I|(n - |I|), \\ 2e &\leq (|I|^2 + |I|(1 - 2n) + (n^2 - n)) + 2|I|(n - |I|), \\ |I|^2 - |I| - (n^2 - n - 2e) &\leq 0, \end{aligned}$$

and given that  $|I| \geq 0$ , solving the quadratic yields

$$|I| \leq \frac{1 + \sqrt{1 + 4n^2 - 4n - 8e}}{2}.$$

Lower bound: Let  $\bar{G}$  be the compliment for the problem graph  $G$ , it holds that the maximal clique size  $\bar{r}$  in  $\bar{G}$  is equal to the maximal size of an independent set in  $G$ . Therefore, using equation 2.17,

$$\bar{e} = \frac{n(n-1)}{2} - e \leq \left(1 - \frac{1}{\bar{r}}\right) \frac{n^2}{2} = \left(1 - \frac{1}{|I|}\right) \frac{n^2}{2},$$

where  $n(n-1)/2$  is the number of edges in a complete graph. The maximal independent set is therefore bounded by

$$\frac{n^2}{2e+n} \leq |I|.$$

□

The bounds defined in equation 2.18 dictate the range of MIS sizes that are possible, and both bounds increase if the relative number of vertices to edges increases (i.e., average vertex degree decreases). Therefore, after each iteration in Alg. 1, it is necessary to remove the MIS that will cause the largest decrease in the number of edges in the subsequent graph, which is equivalent to choosing a maximal independent set that has the maximum number of edges between the vertices in  $I$  and those not in  $I$  ( $V \setminus I$ ). We refer to Alg. 1 with this greedy protocol as greedy MIS colouring (GMIS), and we also refer to the Kwok method which does not have this protocol as random MIS colouring (RMIS).

We now move to onto a comparison between the Lucas, RMIS, and GMIS methods of graph colouring (Alg. 1), using a quantum annealer to compute the combinatorial optimization components of each method, such that the RMIS and GMIS methods can be thought of as *quantum accelerated* algorithms (sometimes also referred to as *quantum-classical hybrid* algorithms). Two problem classes are used in this comparison, the first being the extremal case of Turán's theorem (equation 2.17) named the Turán graph, and the second being a set of random graphs that are chosen to be sparse, planar, and leafless such that they hold some resemblance to modern day telecoms networks, such that we refer to them as random network graphs. Turán graphs of  $n$  vertices are defined to be  $G = T(n, k)$  and are complete  $k$ -partite graphs that have an equal number of vertices in each partite set (if not equal, the remaining nodes are shared equally among the sets). The chromatic number of a Turán graph is always its partite number  $k$ , and the special cases of the Turán graph includes the complete graph ( $T(n, n)$ ) and the complete bipartite ( $T(2n, 2)$ ). The Turán graph is generally considered to be a dense graph, as it has a density  $\geq 50\% \forall k \geq 2$ .

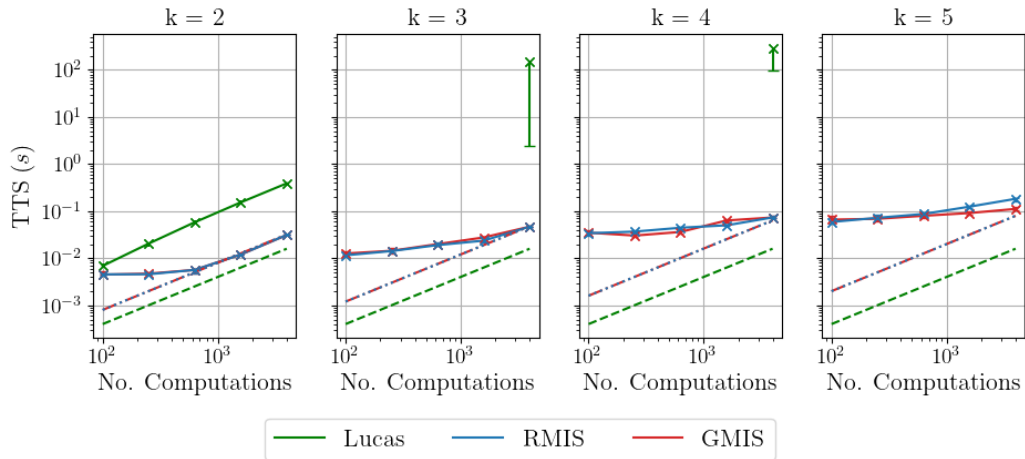
The random network graphs are defined here as  $G = R(n, p)$ , where  $p \in [0, 1]$  is the fraction of additional edges to randomly add to an already randomly generated 2-regular graph, such that the total number of edges is  $n + \lceil pn \rceil$ . All network graphs in this work are generated with  $p = 0.25$ . To ensure planarity,  $n \geq 6$ , and this also ensures that for any random network graph it has a colouring of  $k \leq 4$  [288, 289].

The metric used to compare the performance of the Lucas, RMIS, and GMIS colouring methods is time-to-solution (TTS, see equation 2.7), but the iterative and unconstrained nature of GMIS and RMIS means that a new metric for TTS is needed for a fair comparison. For a  $k$  colouring, the MIS methods use the quantum annealer  $k$  times compared to the single usage from the Lucas method, therefore to represent this, we extend equation 2.7 to include the number of computations,

$$TTS_k = n_c t_f \max \left( \frac{\ln(1 - p_d)}{\ln(1 - p_k)}, 1 \right), \quad (2.19)$$

where  $p_k$  is the probability of finding the colouring  $k$ , and  $n_c$  is the number of computations. For the Lucas method, the number of computations used is simply  $n_c = n_a$ , where  $n_a$  is the number of quantum anneals used to find solutions to the constrained Hamiltonian. However, there are three stages of computations needed for the RMIS and GMIS methods. The first is the  $k$  sets of quantum anneals to find a  $k$  colourings, the second is the number of anneals  $n_r$  used within each iteration to find a maximal independent set  $n_r$ , and the third is how many times the former two steps are run in order to find  $p_k$  with reasonable precision, which we refer to as the number of runs  $r$ . Therefore, for each value of TTS there are  $n_c = krn_r$  computations used in Alg. 1, where the number anneals per iteration is  $n_a = rn_r$ .

Using the D-Wave Advantage quantum annealer housed in Burnaby, Canada, the experimental performance of each method's ability to find the colouring of Turán graphs is assessed, where we focus on finding the optimal TTS as a function of the number of computations,  $n_c$ , for Turán graphs of various colourings with 12 vertices. The annealing time is kept constant at a value of  $4\mu\text{s}$ , with the number of spin-reversal transforms equal to the number of anneals. The logical graphs for each method are embedded onto the Advantage annealer using *minorminer* and



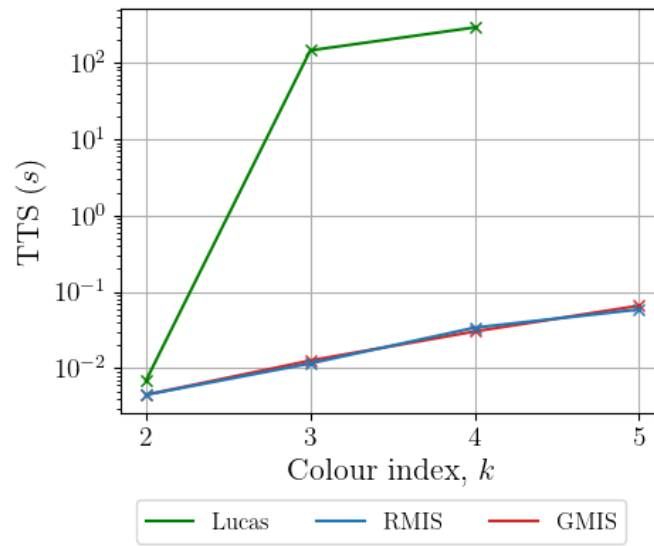
**Figure 2.7:** Times to the correct colouring of Turán graphs (order  $k$ , with 12 vertices) as a function of the number of computations. Lucas, RMIS [286], and GMIS (Alg. 1) graph colouring solutions were computed using the D-Wave Advantage quantum annealer with  $4\mu\text{s}$  anneals and majority-vote post-processing. Each data point represents the bootstrapped median of 20 times-to-solution (TTS) with 95 % confidence bounds. Solid lines are interpolations between data points to approximate the functional form of TTS against number of computations, and the dashed and dotted lines represent the minimum bounds of the TTS for each method.

the proprietary D-Wave parameter setting methods (i.e., uniform bias setting and chain strengths of -1). For RMIS and GMIS, we set  $r = 100$  to estimate  $p_k$ , and  $n_r = \lfloor n_a/r \rfloor$ , with the values of  $k$  taken to be a constant overhead.

Figure 2.7 illustrates the performance of each method when colouring Turán graphs, where the measured TTS is given by the solid line and the minimum possible TTS is given by the dashed line. The minimum TTS for all methods scale linearly in the number of computations, but RMIS and GMIS minimum TTS is always a factor of  $k$  larger than the Lucas method due to the iterative nature of Algorithm 1. Nonetheless, all MIS methods have a consistently lower TTS than the Lucas one-hot constraint method. The MIS methods also reach the minimum TTS in many instances, indicating that the correct maximum independent sets are found in the algorithm when a sufficient number of computations are used in each iteration.

When colouring Turán graphs, the Lucas method is unable to find correct colouring for a low number of computations  $k > 3$  even with majority-vote post-processing, indicating that the probability of finding the colouring solution is very low. For each value of  $k$ , the number of edges in the graph is given by  $e = 72 -$





**Figure 2.8:** Illustration of the optimal TTS scaling as a function of Turán graph colouring for the Lucas, RMIS [286], and GMIS (Alg. 1) graph colouring methods. The Turán graphs are of logical size 12, and each data point is the optimal TTS derived from interpolated data in Figure 2.7.

$\lceil 72/k \rceil \forall k > 1$ , which therefore increases the graph density and therefore the level of minor-embedding required for all methods. The poor performance of the Lucas method is also compounded by the quadratic increase of qubits in  $k$  that results in increased hardness. Using interpolation methods, we find the number of reads that give the optimal TTS and plot against the colouring index to demonstrate the scaling of each method. Despite the additional computational overhead in the RMIS and GMIS methods, the optimal TTS for each Turán graph colouring in Figure 2.8 confirms that the MIS methods outperform the Lucas method, but no significant difference is seen between the RMIS and GMIS methods for these dense problems.

Finally, we also test the methods on the random network graphs, where instead of increasing the chromatic number and keeping system size constant, we keep the chromatic number bounded by  $k \leq 4$  and increase the system size. Starting from a graph size of 32 vertices and up to 2048 vertices, the largest logical problem for the Lucas (GMIS) method that could be embedded onto the D-Wave Advantage annealer was 512 logical (2048 physical) vertices. Again, standard proprietary D-Wave embedding methods and an annealing time of  $4\mu\text{s}$  were used, and spin-reversal transforms were applied once every 50 anneals instead of every anneal. For RMIS


and GMIS, we set  $r = 100$  and  $n_r = 100$ , and for the Lucas method  $n_a = 10,000$ . Given that the number of anneals  $n_a$  is the same for both methods, we remove this as a constant overhead from equation 2.19 by setting  $n_a = 1$ , and focus on the colouring overhead for the MIS methods, such that  $n_c = \bar{k}$ , where  $\bar{k}$  is the average  $k \forall k \leq 4$  colouring solutions out of the  $r = 100$  runs. Interestingly, RMIS and GMIS were both able to find colouring solutions for all problems in Figure 2.9a, whereas the Lucas method was not able to find solutions to any of the problems for  $k = 4$ . In an attempt to remedy this, different parameter setting methods such as weighted biases and longer anneal times were used, but the embeddings were still too broken to yield valid colouring solutions for the Lucas method. The GMIS method is also seen to outperform RMIS for these sparse problems, with an improvement over RMIS seen for 4 out of the 7 problem sizes tested.

To illustrate the detrimental effect of embedding for the Lucas method and why it is unable to find solutions in Figure 2.9 using a quantum annealer, we look at the embedding overhead of each method, which is defined to be the ratio of the number of vertices in the minor-embedded problem to the number of vertices in the original problem. This value will always be  $\geq 1$ , and we can compare the ratio of overheads between the Lucas method and the MIS method using the embedding ratio  $\zeta$ , which is defined as

$$\zeta = \frac{n_I^{\text{Lucas}} n_G^{\text{GMIS}}}{n_G^{\text{Lucas}} n_I^{\text{GMIS}}} \quad (2.20)$$

where  $n$  is the number of vertices. Figure 2.9b shows two sets of ratios, where the first (solid line) is the embedding ratio with respect to the original problem of size  $N$ , and the second being the adjusted embedding ratio (dashed line). The adjusted ratio compares the logical graph sizes, where for the Lucas method it is always has  $k = 4$  times greater number of logical vertices than the MIS methods. As expected from a one-hot encoding, the overhead from the Lucas method is always greater than the GMIS method for both ratios, and increasing problem size compounds this further.

Overall, the greedy MIS approach presented here has proven itself to be an efficient alternative to the one-hot encoded graph colouring in equation 2.12, and is an example of how reformulating the problem can increase the problem sizes one can




**Figure 2.9:** Plots of *a)* TTS as a function of problem size for the RMIS and GMIS colouring methods, and *b)* embedding ratio (Equation 2.20) between the Lucas and MIS methods for the original problem (solid) and the logical problem (dashed). A valid solution is one that returns a colouring  $\leq 4$ , and a D-Wave Advantage annealer was used to find solutions using an annealing time of  $4\mu\text{s}$ . Data points in *a)* are bootstrapped medians from 20 estimates of the TTS, and the error bars are the 95% confidence intervals.

solve, circumventing the hardware restrictions, as well as reducing the embedding and therefore detrimental effects this has on performance. Experimentally, we have seen that the GMIS method (Alg. 1) is able to correctly colour all the dense Turán and sparse random network graphs tested, with correct colourings found for sizes equal to 2048, demonstrating that this method is a viable route to graph colouring with industrial size problems of sizes  $> 10,000$  variables. The method presented by Kwok, which we refer to as random MIS colouring, also performs well against the Lucas method of colouring. It also performs comparably to the greedy MIS method for dense graphs, but was seen to have a worse TTS than the greedy MIS method when looking at large sparse graphs. The Lucas method is only able to find colourings for small graphs with a low chromatic number, as demonstrated by increasing TTS and need for higher number of computations in Figure 2.8, indicating that increasing the number of options in the one-hot constraint is particularly detrimental to performance. The GMIS method can also be tested in future work against alternative encoding schemes such as domain wall encoding [238, 239], as well as different parameter setting methods as explored in Section 2.1.

## 2.3 Reverse Annealing

As previously introduced in Section 1.2.3, reverse annealing can be viewed as a local search algorithm that will improve upon solutions already found by some optimization algorithm [17, 199, 202, 203]. The protocol starts by initializing the system in a state with zero transverse-field that corresponds to a previously found solution, instead of a superposition state with maximum transverse-field in traditional quantum annealing. A transverse-field is then introduced, which is equivalent to introducing quantum fluctuations into the system that can allow the initial solution to explore the local configuration space in the hope that a better solution exists nearby. This is shown to avoid first-order phase transitions in adiabatic reverse annealing [200, 201] (Equation (1.10)), but is heuristically implemented in D-Wave annealers such that diabatic transitions are needed to reach the ground state. It was shown in Ref. [57] that reverse annealing had the largest probability of reaching the ground state from a non-optimal solution if the system was held before the minimum gap and did not pass through it. However, knowing the position of the minimum gap and therefore when to stop increasing transverse field in a reverse anneal is a computationally hard question, but there do exist some methods to approximate its position using perturbation theory [290].

Reverse annealing therefore opens up a paradigm of adiabatic quantum computation where we do not have to pass through the minimum gaps, at the cost of it being restricted to a being a local-search protocol. Taking quantum tunnelling to be the primary dynamical mechanism in quantum annealing, we look at the reverse annealing protocol on the D-Wave 2000Q quantum annealer and ask how effective is reverse annealing at exploring the solution space? To help answer this, we consider a Hamiltonian with a quadruply degenerate first excited state, where the Hamming distances from the ground state are  $\{2, 3, 3, 4\}$ , i.e., small enough to permit tunnelling to the ground state. By reverse annealing from each of the four first excited states, we can compare experimental and simulated data to determine whether the Hamming distance to the ground state affects the ground state probability, and try to match a model that explains the observed phenomena best. However, it found that there



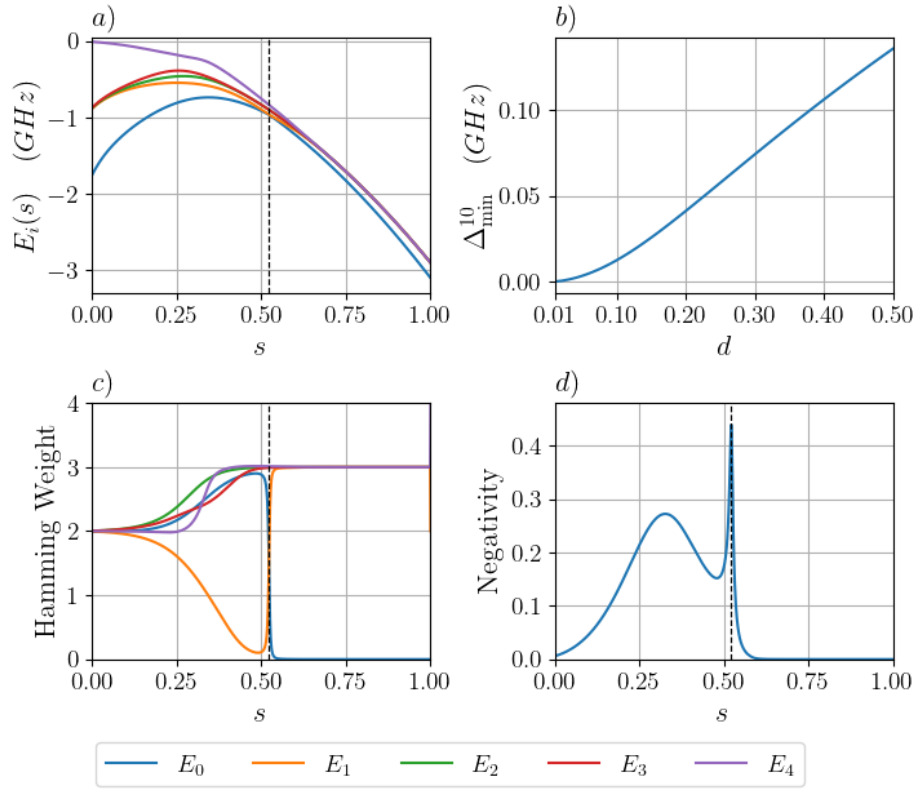
**Figure 2.10:** Graphical illustration of the four spin chain problem Hamiltonian used for the reverse annealing experiments. The tuneable parameter,  $d > 0$ , tunes the size of the minimum gap, and  $R > 0$  tunes the energy scale of the problem. The sign convention for this problem follows the Hamiltonian defined in Equation 2.21.

are hardware artefacts such as *spin-bath polarization* that bias reverse anneals, such that none of the models tested match the experimental data. It also means that if an optimization problem is highly sensitive to its parameters (i.e., a large condition number) then spin-bath polarization can change the problem you are trying to solve. Nonetheless, the phenomena affecting reverse anneals are probed to give a better understanding of their effects.

The problem Hamiltonian of interest for our reverse annealing experiments is shown in Figure 2.10, and consists of a tunable spectral gap (controlled by the value of  $d$ ), and has a classical energy spectrum with a single isolated ground state and quadruply degenerate first excited-states whose Hamming distances from the ground state increase from 2 to 4. For ease of analytical and numerical calculations, the Hamiltonian in Figure 2.10 is kept to four spins, and therefore will be called the four spin chain (FSC); however larger versions can exist and are explored in Chapter 3. The form of the full Hamiltonian in Figure 2.10 in the transverse field Ising model (Equation (1.3)) can be written as

$$H = -A(s) \sum_{i=1}^4 \sigma_i^x + RB(s) \left[ (1-d) (\sigma_2^z + \sigma_3^z) - (\sigma_1^z + \sigma_4^z) - \sum_{i=1}^3 \sigma_i^z \sigma_{i+1}^z \right]. \quad (2.21)$$

The properties of the FSC are illustrated in Figure 2.11 using the eigenstates and eigenenergies of the diagonalized Hamiltonian in Equation 2.21. The energy spec-



**Figure 2.11:** Plots of *a)* the first five instantaneous energy eigenstates  $E_i$  of the four spin chain (Figure 2.10) as a function of the anneal parameter  $s$ , *b)* the scaling of the minimum gap  $\Delta_{\min}^{10}$  as a function of the tuneable parameter  $d$ , *c)* the Hamming weights of the first five states according to the Hamming weight operator in Equation 2.22 as a function of  $s$ , and *d)* the negativity (entanglement) of the ground state with respect to the partial transposes of the inner qubit sub-spaces of the four spin chain. The global energy scale is set to  $R = 1.0$  GHz throughout, and a value of  $d = 0.05$  was used for plots *a)*, *c)*, and *d)*. The schedules  $A(s)$  and  $B(s)$  used follow the D-Wave 2000Q schedule defined in Figure 1.8. The black dashed line shows the position of the minimum energy gap.

trum of the five lowest energy states is shown in Figure 2.11a, and the minimum gap between  $E_0$  (the ground state) and  $E_1$  (first excited state) is shown as a function of  $d$  in Figure 2.11b. Note that the first-excited-state degeneracy of the FSC problem is broken for a non-zero transverse-field. Two other properties of the FSC also exhibit interesting features. The first is that the Hamming weight is given by the operator

$$\mathbf{W} = \frac{1}{2} \sum_i^N (\mathbb{1}_i - \sigma_i^z), \quad (2.22)$$

where  $\langle E_i | \mathbf{W} | E_i \rangle$  is the expectation value of the Hamming weight for eigenstate  $|E_i\rangle$  in Figure 2.11c. It demonstrates the presence of a quantum phase transition at the

minimum gap (the black dashed line) due to the large change in Hamming weight (and therefore magnetization). Also note how the four excited states superimpose to all have a Hamming weight equal to 3, which is caused by the outer qubits in the FSC being in a superposition of spin  $-1$  and spin  $1$  simultaneously (i.e., Hamming weight = 0.5) and inner qubits have Hamming weight of 1. This then collapses onto their computational Hamming weights of  $\{2, 3, 3, 4\}$  as soon as the transverse-field is zero.

Secondly, the negativity of the instantaneous ground state (Figure 2.11d), which is a measure of how entangled the state is, defined as

$$\mathcal{N} = \frac{1}{2} (1 - \|\rho^{TA}\|_1), \quad (2.23)$$


where  $\|\cdot\|_1$  is the trace norm,  $\rho$  is the ground state density matrix, and  $A$  is the subsystem which is partially transposed in  $\rho$  (taken to be the inner qubits in this case). The spike in negativity also confirms the presence of a quantum phase transition between the ground and first excited state, and increases further if  $d$  is decreased (i.e., smaller minimum gap).

The final part of the analysis involves looking at the semi-classical potential  $V$  of the FSC, calculated using the spin-coherent ansatz in the  $O(2)$  model [275] (see Section 1.4.3), which is defined as

$$|\theta\rangle = \bigotimes_k^N \left[ \cos \frac{\theta_k}{2} |0\rangle + \sin \frac{\theta_k}{2} |1\rangle \right] \quad (2.24)$$

$$V(t, \theta, \phi) = \langle \theta, \phi | H(t) | \theta, \phi \rangle, \quad (2.25)$$

where  $\theta_k$  is the polar angle describing the  $k^{\text{th}}$  qubit. The potential equation is restricted to the  $XZ$ -plane of the Bloch sphere (removing the azimuthal component), and can be described using only two angles due to the permutation symmetry of the FSC, such that  $\theta_1$  and  $\theta_2$  represent the inner and outer qubits, respectively. This restriction reflects the fact that the Hamiltonian (Equation 2.21) only has energy contributions in the  $XZ$ -plane, and therefore this form can be used to emulate



**Figure 2.12:** Plots of the semi-classical potential from Equation 2.26 as a function of  $s$  for a four spin chain with  $d = 0.05$  and  $R = 1.0$ . The red marker indicates the global minimum of the landscape, and the inner and outer spin angles are denoted by  $\theta_1$  and  $\theta_2$ , respectively.

quantum annealers [277]. The simplified form is therefore

$$V(s, \theta_1, \theta_2) = \langle \theta_2 \theta_1 \theta_1 \theta_2 | H(s) | \theta_2 \theta_1 \theta_1 \theta_2 \rangle, \quad |\theta_k\rangle = \cos \frac{\theta_k}{2} |0\rangle + \sin \frac{\theta_k}{2} |1\rangle. \quad (2.26)$$

The FSC semi-classical potential is illustrated in Figure 3.4, with a discontinuous change in position of the red marker (global minimum of the potential) exhibited. This occurs about the minimum gap, and signifies the presence of a phase transition and that annealing guides particles to a false minimum. It also confirms the ground state is isolated away from the other first excited states, with the first excited states all being found along the  $\theta_1 = \pi$  and  $-\pi$  axis. Using these properties, one can use D-Wave's reverse annealing (as outlined in Figure 1.6) to initialize the system in each of the first excited states  $|0110\rangle$ ,  $|1110\rangle$ ,  $|0111\rangle$  and  $|1111\rangle$ , that are of different Hamming distances to the ground state  $|0000\rangle$ , and measure the ground state probability after a reverse anneal. To support the experimental data, open system Monte Carlo and numerical simulations of RA are used to emulate the experimental dynamics. Both the Monte Carlo and numerical models selected emulate quantum annealing in a regime where there are strong decoherence effects. These models are also tractable to simulate whilst also being able to model any incoherent dynamics



that could explain the experimental observations.

For the numerical simulations, the adiabatic master equation (AME) outlined in [143] is used to simulate open system RA in the singular coupling limit (see Section 1.4.1, Equation (1.22) for more details). However, given that analytical simulations of the FSC are expensive and restricted to coupling limit regimes, we use both classical and quantum Monte Carlo (MC) to further support the experimental data. For the classical approaches, Metropolis-Hastings updates are used to minimize the following energy functions

$$\mathcal{E}(t) = \beta(t) \left[ \sum_i^N h_i s_i + \sum_{i,j} J_{i,j} s_i s_j \right] \quad (2.27)$$

$$\mathcal{E}(\theta, t) = -A(t) \sum_i^N \sin(\theta_i) + B(t) \left[ \sum_i^N h_i \cos \theta_i + \sum_{i,j} J_{i,j} \cos \theta_i \cos \theta_j \right], \quad (2.28)$$

where Equation (2.27) and Equation (2.27) correspond to classical (simulated) annealing and spin-vector Monte Carlo (SVMC), respectively. SVMC explores the Bloch sphere in the  $XZ$ -plane only, and has varying degrees of success in replicating the results of the D-Wave quantum annealer [89, 119, 159, 277]. Secondly, we use path-integral Monte Carlo (PIMC) that applies the Trotter break-up formula to the Hamiltonian, such that a discrete imaginary time dimension with  $N_\tau$  Trotter slices is introduced to approximate the quantum system using a classical one (see Section 1.4.2). PIMC can be used to successfully model incoherent quantum annealing as well as other dynamical effects such as tunnelling. Single spin Metropolis-Hastings updates are used to minimize the energy function that represents a closed quantum system with no sign problem,

$$\beta \mathcal{E}(t) = \frac{\beta}{N_\tau} B(t) \sum_\tau \left[ \sum_i h_i s_{i,\tau} + \sum_{i,j} J_{i,j} s_{i,\tau} s_{j,\tau} \right] - J_\perp(t) \sum_{i,\tau} s_{i,\tau} s_{i,\tau+1}, \quad (2.29)$$

where  $J_\perp(t)$  is the coupling strength between Trotter slices along the imaginary time direction such that,

$$J_\perp(t) = -\frac{1}{2} \ln(\tanh(\beta A(t)/N_\tau)). \quad (2.30)$$

This method can be used to emulate QA efficiently and scales better than the numerical methods employed, and it can be extended to include quantum noise such as dephasing by adding an explicit bath term to Equation (2.29). This term takes the form

$$\beta \mathcal{E}_B = -\frac{\alpha}{2} \left( \frac{\pi}{N_\tau} \right)^2 \sum_{i=1}^N \sum_{\tau=1}^{N_\tau} \sum_{\tau'=\tau+1}^{N_\tau} \frac{s_{i,\tau} s_{i,\tau+1}}{\sin^2 \left( \frac{\pi}{N_\tau} |\tau - \tau'| \right)} \quad (2.31)$$

where  $\alpha$  is the effective system-bath coupling strength. This can be used to explore intermediate coupling regimes that D-Wave quantum annealing systems are likely to be operating in, but it is a regime that the master equations cannot access [143].

Throughout the rest of this section, we use  $d = 0.05$  to create an FSC Hamiltonian whose minimum energy gap value,  $\epsilon_G$ , is an order of magnitude less than the system temperature (taking the system temperature to be 15mK [229],  $\epsilon_G$  is  $\approx 12 \times$  smaller). This value of  $d$  also is large enough to ensure that if integrated control errors change the value of  $d$ , then  $d > 0$  holds for more than 95% of the time for a standard deviation of  $\sigma = 0.03$ . Reverse annealing within this regime will allow us to probe in greater detail the role of thermal dynamics in the D-Wave annealer for optimization. It is also expected that for a coherent or equilibrium system, there is no ground state probability dependence on the initial states Hamming weight when finite transverse-field is introduced. This stems from the observations in Figure 2.11c, as for any non-zero transverse-field where  $s^* < s < 1$  the excited states all have an expected Hamming Weight of 3. However, the D-Wave annealers possess incoherent systems and nonequilibrium dynamics, therefore it is unknown whether the Hamming weight of the initial state would affect the ground state probabilities. We explore this experimentally and with simulations in Section 2.3.1 and 2.3.2 respectively.

To summarize, we confirm that the ground state is isolated and is not energetically accessible during the anneal by analysing the semi-classical energy potential. The evolution of the potential landscapes as a function of  $s$  is shown in Figure 2.12, and demonstrates the creation of two separate potential wells at  $s = 0.4$ . This is indicative of a false minimum, as for low values of  $s$  the energetically preferred

route is to the eigenstates that will finish being the degenerate first excited states at  $s = 1$ . At the phase transition, the isolated potential well becomes the ground state, confirming the fact that either thermal and/or quantum dynamics are required to reach the ground state during an anneal.


### 2.3.1 Experimental

Using the D-Wave 2000Q at Los Alamos National laboratory, reverse anneals using the FSC were performed with the same schedule defined in Figure 1.6, whereby the annealing rate  $r \in (0, 1]$  and the minimum value of  $s$ ,  $s^*$ , define the reverse anneal time, such that

$$t_1 = \frac{1 - s^*}{r}, \quad (2.32)$$

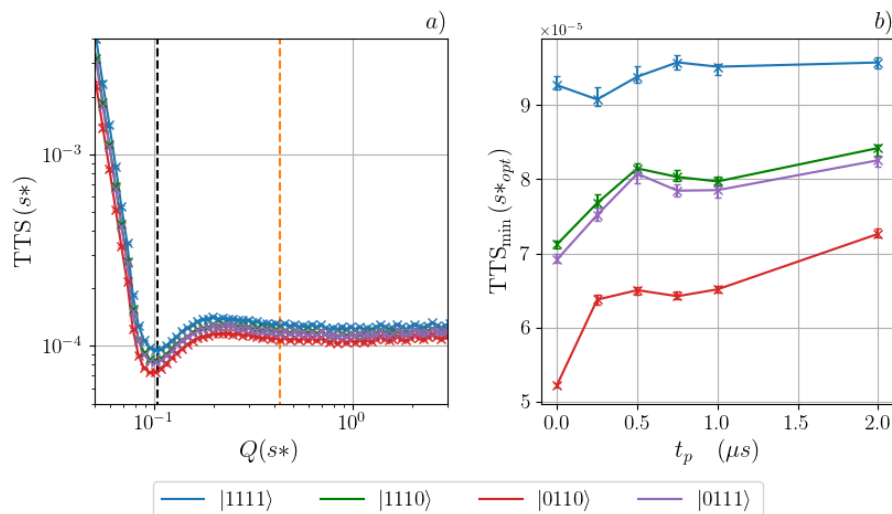
and total reverse anneal time is  $t_{\text{anneal}} = 2t_1 + t_p$ , where  $t_p$  is the time spent pausing, allowing the system to thermalize. Throughout the experiments, the annealing rate is set to  $r = 0.5\mu s^{-1}$ , and various pause times are used to explore the effect of pausing on the time-to-solution defined in Equation (2.7). The magnitude of quantum fluctuations introduced into the reverse anneal is parameterised by  $Q(s^*) = A(s^*)/B(s^*)$ , where  $A(s^*)$  and  $B(s^*)$  are the maximum transverse and minimum longitudinal field strengths, respectively, in the reverse anneal (see FSC Hamiltonian in Equation 2.21). Figure 2.13 illustrates the amplitude of  $Q(s^*)$  needed to enable diabatic transitions from the initial state to the ground state with and without pausing the system. For all  $Q(s^*)$ , the reverse anneals with initial states closer in Hamming distance to the ground state had a higher probability of reaching the ground state at the end of the anneal, inconsistent with what is expected from an equilibrium system, but consistent with reverse annealing being a local search algorithm [201–203]. The additional pause time is shown to increase ground state probability about the minimum gap, consistent with thermalization phenomena previously observed in literature [57, 118, 119].

From an optimization perspective, this peak of probability about the minimum gap would also correspond to a minimum time-to-solution (TTS) due to the reverse-anneal times being shorter for smaller values of  $Q(s^*)$ . A translation of Figure 2.13b



**Figure 2.13:** Ground state probability after reverse annealing with a pause of *a)*  $0\mu s$  and *b)*  $2\mu s$  from each of the four degenerate first-excited-states on the LANL D-Wave 2000Q quantum annealer. The maximum relative magnitude of transverse-field applied to the FSC is defined as  $Q(s^*) = A(s^*)/B(s^*)$ , where  $s^*$  is the minimum value of the annealing parameter in a reverse anneal. A ramp rate (Equation 2.32) of  $0.5\mu s^{-1}$  is used for both plots, the black dashed line marks the position of the minimum gap, and the orange dashed line marks the value of  $Q(0.4)$ , where to the right of this line the potential starts to become unimodal (Figure. 2.12). Each reading is the bootstrapped median from 100 ground state probabilities, each calculated from 125,000 measurements on the QPU with a readout thermalization of  $1\mu s$ . Using space-division multiplexing, 125 FSC's were fitted onto the D-Wave 2000Q device, such that only 1000 runs were needed to retrieve 125,000 measurements.

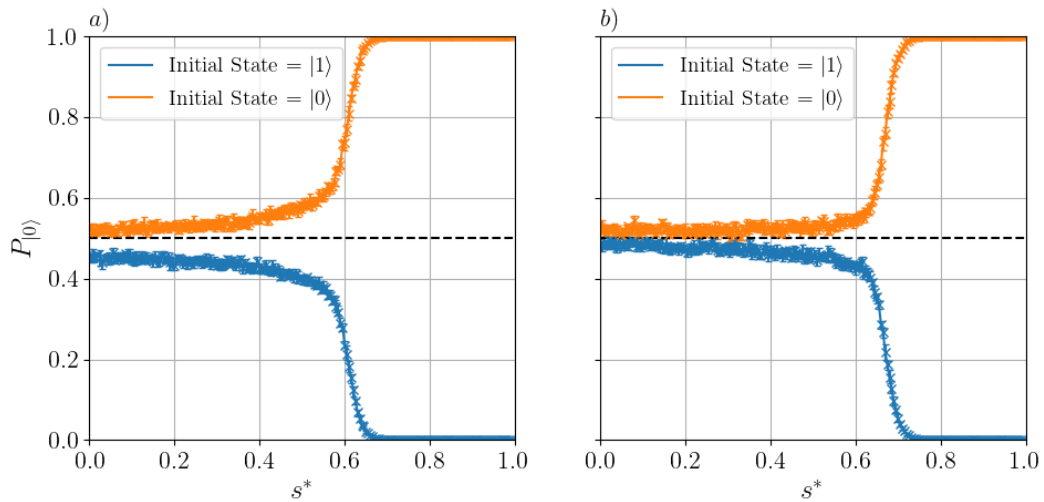
to TTS is shown in Figure 2.14a, where the minimum TTS is shown to occur when  $s^*$  is at the minimum gap, i.e., significant system thermalization when paused at this position. However, an increase in ground state probability using thermalization is offset by an increase the total computation time, resulting in no computational benefit to the TTS. Figure 2.14b illustrates the minimum TTS that occurs at  $s^*_{opt} \forall s^*$  for reverse anneals anneals with various pause times. The states closest in Hamming distance to the ground state have a minimum TTS at  $t_p = 0\mu s$  and therefore do not benefit from any pausing, with exception to initial state furthest in Hamming distance (i.e., the hardest state to initialise in), where some pausing was seen to be potentially beneficial. To confirm this effect, larger problems would be needed to be reverse annealed in order to increase the hardness and Hamming distances between initial problem states.



**Figure 2.14:** Plots of *a*) the time-to-solution as a function of  $s^*$  for the reverse anneal with a  $2\mu s$  pause shown in Figure 2.13b, and *b*) the minimum TTS at  $s^*_{opt}$  for various pause times when reverse annealed from each of the four degenerate first-excited-states on the LANL D-Wave 2000Q quantum annealer. A ramp rate (Equation 2.32) of  $0.5\mu s^{-1}$  is used for both plots, the black dashed line marks the position of the minimum gap, and the orange dashed line marks the value of  $Q(0.4)$ , where to the right of this line the potential starts to become unimodal (Figure. 2.12). The experimental procedures follow those defined in Figure 2.13 for all results shown. The error-bars plotted in *b*) are the 95% confidence bounds of the median TTS values.

There are two regimes of interest when reverse annealing the FSC, the first being the dynamics about the minimum gap (the black dashed line in Figure 2.13) where the thermalization phenomena occur, and secondly is when the semiclassical potential starts to become unimodal again as seen in Figure 2.12. For  $s^* < 0.4$ , the reverse anneal will start to enter a potential that is no longer bimodal, and the orange dashed line marks the beginning of reverse annealing entering the unimodal distribution for  $Q(s^*) > Q(0.4)$ . Despite reverse annealing into a region where the potential is unimodal and therefore losing most memory of the initial state magnetization, there still appears to be some memory of the initial state due to the consistent ordering of ground state probabilities. The separation is less prominent for longer anneal times, as seen in the longer pause-time regime in Figure 2.13b, and therefore reverse annealing maybe affected by some integrated control error on the QPU.

A single unbiased qubit is used to illustrate to what extent this control error in reverse annealing affects state probability in Figure 2.15, where we initialize a



**Figure 2.15:** The probability of being in the  $|0\rangle$  state after a reverse anneal as a function of  $s^*$  (the minimum value of  $s$  annealed to) for a single unbiased spin on the D-Wave 2000Q. Starting in either of the initial states biases the qubit differently for both annealing rates of *a*)  $r = 0.5$  (i.e., shorter reverse anneal) and *b*)  $r = 0.05$  (i.e., longer reverse anneal). Each data point is a bootstrapped median of 100 probabilities, with error bars representing the 95% confidence intervals. Each probability is found from a randomly placed spin on the D-Wave 2000Q that is sampled 1000 times.

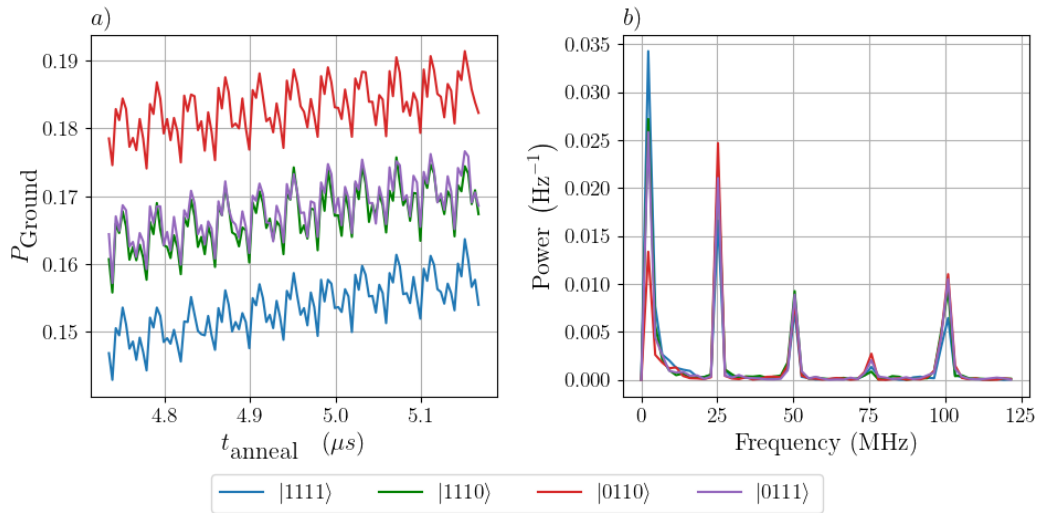
random qubit on the QPU in either the spin up ( $|0\rangle$ ) or spin down ( $|1\rangle$ ) state and then measure the probability of being in the spin up state ( $P_{|0\rangle}$ ) for different values of  $s^*$ . We expect that the random placement of qubits across the QPU should yield an average result of  $P_{|0\rangle} = 0.5$ , but as clearly observed in Figure 2.15 there is a biasing of  $P_{|0\rangle}$  in towards the initial state configuration. Two different annealing rates were used for this demonstration, where in both cases the system is frozen for high values of  $s^*$  before then entering a regime where quantum dynamics allows convergence to  $P_{|0\rangle} = 0.5$ . Figure 2.15a uses an annealing rate of  $r = 0.5\mu s^{-1}$ , and there is a clear deviation from  $P_{|0\rangle} = 0.5$ , which interestingly is larger if the system is initialized in the  $|1\rangle$  state. For a smaller annealing rate of  $r = 0.05\mu s^{-1}$  in Figure 2.15b (i.e., longer annealing times), the deviation is smaller in both cases, with only a marginally larger deviation for the  $|1\rangle$  state.

The likely explanation for this effect can be attributed to a biasing SQUID used in the D-Wave reverse anneal protocol, whereby the initial state is encoded onto qubits by an adjacent SQUID that induces a persistent current before the reverse

anneal. If this SQUID polarizes the substrate, it can retain flux after the reverse anneal starts, and therefore keep a persistent memory of the initial state on the qubit. This effect of polarizing the substrate is known as spin-bath polarization, and is known to occur on D-Wave annealers (see Section 1.3.1). This introduces additional bias to the qubits, which corresponds to state probabilities biased away from the expected result. This effect is also seen to reduce in magnitude over time, as seen in Figure 2.15b, and is consistent with the residual flux in the adjacent SQUID being lost over time. Therefore, the separation of states in Figure 2.13 is hypothesized to be the artefact of spin-bath polarization, and not the result of reverse annealing being a local-search algorithm.

To confirm that separation of states is not the result of some other dynamical process, we rule out other models in Section 2.3.2 to further support that spin-bath polarization is responsible. Despite the presence of this effect, it can also be mitigated to some extent, such that reverse annealing has been used in the literature for quantum simulations of phenomena such as the Kosterlitz-Thouless phase-transition [30, 136]. The effect is mitigated using two methods, where firstly for a set of consecutive reverse anneals the final state found from the previous reverse anneal is left to be the initial state of the next reverse anneal. This avoids the reinitialization of the initial state used for the first reverse anneal, and therefore avoids using the biasing SQUIDs that inflict spin-bath polarization. The second method is using a technique called “shimming”, which is a classical gradient descent method used to apply flux bias to unbiased qubits such that they have an expected magnetization as close to zero as possible before reverse annealing. This can be applied to Figure 2.15 in order to counter-bias the qubits such that they actually meet the expected value of  $P_{|0\rangle} = 0.5$ , and is key to the success of the work presented in Refs. [30, 136].

Another artefact that was also observed when the system had passed through the minimum gap are oscillations in ground state probability (Figure 2.16a) as a function of reverse anneal time. The reverse anneal used was one with a constant anneal rate of  $r = 0.5\mu s^{-1}$  and with a  $2\mu s$  pause, such that the anneal time was parameterized by  $t_{\text{anneal}} = 2(3 - 2s^*)\mu s$ . A fundamental harmonic of 25 MHz for the ground state



**Figure 2.16:** Plots of *a*) the oscillations in ground state probability as a function of anneal time  $t_{\text{anneal}}$ , such that  $t_{\text{anneal}} = 2t_1 + t_p$  where  $t_1 = 2(1 - s^*)\mu\text{s}$  and  $t_p = 2\mu\text{s}$ , and *b*) its corresponding power spectrum. Readings were taken using the same methodology used in Figure 2.13.

probability oscillations is illustrated in the power spectrum in Figure 2.16b, with other subsequent resonances also shown to be present. In a pure closed system, we would expect oscillatory behaviour in the ground state probability for a low enough value of  $s^*$  due to the interference caused by Mach-Zehnder interferometry [291]. However, the experimental quantum annealer is not in a coherent limit on these timescales tested, and it is therefore unlikely that the oscillations are a result of Mach-Zehnder interferometry.

To further explore this, additional reverse annealing schedules are experimentally implemented with different combinations of fixed or variable annealing rates, pause times,  $s^*$ , annealing times, and time past the minimum gap  $t_{\text{gap}}$ . The schedule combinations are presented in Table 2.2, which tries to identify different annealing regimes possessing oscillations. Schedule 1 is the original reverse anneal schedule used to generate oscillations in Figure 2.16, with the only other schedule being able to replicate ground state probability oscillations being schedule 3. This schedule is uniquely different to schedule 1 as there is no time-domain due to the fixed the annealing time and  $t_{\text{gap}}$ , but oscillations are still seen as a function of  $s^*$ . Dependence on time is also ruled out by other annealing schedules, particularly in schedule 2



Schedule	Fixed $r$	Fixed $t_p$	Fixed $s^*$	Fixed $t_a$	Fixed $t_{\text{gap}}$	Oscillating $P_G$
1	✓	✓	×	×	×	✓
2	✓	×	✓	×	×	×
3	✓	×	×	✓	✓	✓
4	×	×	✓	×	✓	×
5	×	✓	×	×	✓	×
6	×	✓	✓	✓	×	×

**Table 2.2:** Reverse annealing experiments used to reproduce oscillations in ground state probability  $P_G$  as a function of pause time  $t_p$  and/or annealing time  $t_a$  and/or time past the minimum gap  $t_{\text{gap}}$ . The schedules also fix or vary the annealing rate,  $r$ , and minimum annealing parameter,  $s^*$ . Schedule 1 is the schedule used to produce the oscillations in Figure 2.16.

where we anneal down with a constant rate to a fixed value of  $s^*$ , pause for various times before coming back to  $s = 1$ .

One could still propose that despite the constant time oscillations in schedule 3, there were still different accumulations of phase in each state because the scheduled annealed down to different values of  $s^*$ . Therefore, we would want to attempt to change the amplitude of these oscillations by annealing at different fixed rates using schedules 1 and 3. In both cases, there were no changes in oscillation amplitude, and also no dependence on frequency with respect to the initial state, it is therefore likely that this phenomena the result of some artefact/interference effect caused by integrated control errors on the D-Wave QPU. This artefact occurs when we reverse anneal to different values of  $s^*$  (and therefore transverse-field) at a fixed annealing rate, which indicates that it is possibly the result of quantization error cause by the digital to analogue conversion of the schedule. Therefore, minor deviations from the true values of  $A(s^*)$  and  $B(s^*)$  may occur systematically, leading to these apparent oscillations in ground state probability.

Overall, reverse annealing with the D-Wave 2000Q quantum annealer exhibited how quantum fluctuations and thermalisation effects can be used to access the ground state from initial state of varying Hamming distance. The reverse anneals with initial states closer in Hamming distance to the ground state consistently had a higher probability of reaching the ground state at the end of the anneal. Pausing in a reverse anneal resulted in a reduced separation in ground state probability between

initial states, and also a peak in probability about the minimum gap induced by thermalisation. To see whether this had a beneficial impact on the time-to-solution (TTS), the optimal time to solution was calculated for various pause times. Reverse anneals with no pause had the lowest TTS, with exception to the initial state furthest in Hamming distance (i.e., the hardest reverse anneal), indicating that pausing could be computationally beneficial for harder problem instances.

The persistent separation in ground state probability between initial states even when the energy potential was unimodal suggested that spin-bath polarization was impacting reverse anneals. This was demonstrated in Figure 2.15, by randomly selecting 1000 unbiased qubits and reverse annealing from either a spin-up or spin-down initial state. It was found that a persistent bias was found even for long anneal times that depended on the initial state, confirming that spin-bath polarization is caused by the setting of the initial state in reverse annealing. Minor oscillations in ground-state probability for all initial states were also observed for this problem if a specific reverse anneal schedule was used. This was also attributed to hardware control errors in the reverse anneal, where quantization error in the digital-to-analogue conversion of the anneal schedule parameters can introduce systematic errors leading to the apparent oscillations. Therefore, reverse annealing as an optimisation method holds promise as a local-search algorithm, but the system errors and biases may detrimentally impact problems that are parameter sensitive and where there is a poor choice of initial state.


### 2.3.2 Simulation

This section focuses on reverse annealing with several models generally used to emulate D-Wave annealers and whose dynamics are well understood. The models explored are not expected to be accurate in modelling the D-Wave annealer, but are used to rule-out certain dynamical models in order to better explain the experimental observations. We first look at the adiabatic master equation (AME) [143] in the singular coupling limit (Equation (1.22)), which is a model of decoherence specifically in the computational basis. In this limit, quantum dynamics are quickly lost due to strong decoherence, leading to a fully mixed final state (i.e., one in an infinite temper-

ature limit) in a standard quantum anneal. This is not found to be an accurate model of experimental quantum annealing [137], but is a computationally tractable model that can be compared against the experimental data in Figure 2.13. For a reverse anneal with a pause time of  $0\mu s$  (for tractability purposes), Figure 2.17 illustrates how the FSC performs when modelled by the AME in the singular coupling limit.


Both models with (Figure 2.17b) and without (Figure 2.17a) cross-talk (see Equation (1.13)) possess a separation in ground state probability between the  $|1111\rangle$  initial state and all other initial states, contrary to what is seen experimentally. A peak in probability also occurs as the reverse anneal approaches the unimodal distribution of the potential, before then decreasing to the probability of a random guess ( $= 1/2^4$ ) as we increase transverse field to reach the  $|+\rangle$  state which becomes maximally mixed in the singular limit. Cross-talk also does little to effect the ground state probabilities, and increases the separation between the  $|1111\rangle$  state and other states, as the  $|1111\rangle$  state will experience the largest favourable contribution from the cross-talk due to the number of aligned spins in the state. We are unable to use AME in the weak coupling limit for this problem Hamiltonian, due to the intractable convergence times caused by the small minimum gap.

For more tractable simulations, Monte Carlo methods are used (see Section 1.4), where we first look at the classical reverse annealing models. All simulations employed have noisy bias and coupling values set before the anneal begins, where the noise is a sample from the normal distribution parameterized by the mean and standard deviation set to  $\mu = 0$  and  $\sigma = 0.03$ , respectively, after which cross-talk is applied to the system with a background susceptibility of  $\chi = -0.035$ . The temperature of the simulations is set to be the minimum operating temperature of  $12.26mK$ , which in Ising energy units for the LANL D-Wave 2000Q annealer is  $T = 0.018$ . The first classical model used is simulated annealing, where in the context of reverse annealing the initial state begins from a cold system (few thermal excitations) and is annealed to a hotter system parameterized by maximum temperature  $T^*$ , then back to the cold system. The cold temperature is set to be  $1.226mK$ , which is a tenth of the minimum system temperature of the LANL D-Wave



**Figure 2.17:** Plots of ground state probability after a reverse quantum anneal from each of the initial states *a)* without cross-talk and *b)* with cross-talk, which was simulated using the AME in the singular coupling limit [143]. Both the  $|1110\rangle$  and  $|0111\rangle$  state perfectly overlap in the plots. The maximum relative magnitude of transverse-field applied to the FSC is defined as  $Q(s^*) = A(s^*)/B(s^*)$ , where  $s^*$  is the minimum value of the annealing parameter in a reverse anneal. The system was evolved using the same time schedules used experimentally in Figure 2.13 for a pause time of  $0\mu\text{s}$ . An Ohmic bath (Equation (1.21)) is used with a coupling strength of  $\eta g^2 = 10^{-3}$ ,  $T = 12.26$  mK, and a bath cut-off frequency of  $\omega_c = 8\pi$  GHz. The background susceptibility for the cross-talk was set to  $\chi = -0.035$ . The black dashed line marks the position of the minimum gap, the orange dashed line marks where the semiclassical energy potential starts to become unimodal at  $Q(0.4)$ , and the dotted black line indicates the probability of a random guess. The reverse anneal parameter schedules were the same as those used in Figure 1.6, which parameterize the Hamiltonian coefficients defined in the Figure 1.8 for the LANL D-Wave 2000Q quantum annealer.

2000Q QPU. The maximum temperature is a value where the Boltzmann factor gives a 50% acceptance probability for the largest energy change in the problem (i.e.,  $\exp\{-E_{\text{max}}/T_{\text{hot}}\} = 0.5$ ), and both temperatures are also scaled to Ising energy units. The temperature is related to the minimum annealing parameter  $s^*$  by a geometric progression from the hottest temperature at  $s = 0$  to the coldest temperature at  $s = 1$ , and the same schedule used for Figure 2.13 is also used here. The results of the reverse simulated anneal are shown in Figure 2.18, where the magnitude of ground state probability is greater than that seen in an experimental reverse quantum anneal, but possesses no separation of states at any temperature range (due to the degenerate first excited state being in thermal equilibrium). The reverse simulated anneal does feature a thermalization peak when the system is paused at a fixed temperature, and




**Figure 2.18:** Reverse simulated annealing of the four spin chain from each of the computational first-excited states for *a*) no pause and *b*) a pause of  $2\mu s$ . The simulated annealing cold and hot temperature limits are 1.226mK and 3K, respectively, where  $T^*$  is the temperature at which the system is held in the reverse simulated anneal. The black dashed line represents the temperature that equals the energy difference between the ground and first excited state. Each data point is a median of 50 bootstrapped probabilities with 95% confidence intervals, calculated from 1000 simulated anneals. The reverse anneal schedules were the same as those used in Figure 1.6, and the number of increments (sweeps) in the simulated anneal is related to the total annealing time  $t_a = 2(3 - 2s^*)$  by  $1000 \text{ steps } \mu s^{-1}$ .

occurs about the point in the anneal where the maximum temperature  $T^*$  equals the energy difference between the ground and first excited state ( $\Delta_{10} = 0.2$ ).

The next method is spin-vector Monte Carlo (SVMC) which explores the semi-classical potential of the transverse-field Ising model using thermal excitations, and is thought to emulate quantum annealing in some respects [89, 119, 277]. However, the data shown in Figure 2.19, deviate from what is observed experimentally in Figure 2.13 both in the magnitude of ground state probability and there not being any separation of states. Pausing the system also exhibits some peak in probability from thermalization effects, but it is centred after the minimum gap, which is more akin to the simulated annealing thermalization model rather than what is seen experimentally. After SVMC reverse anneals into a unimodal distribution, the ground state probability also becomes constant, and is a probability less than a random guess ( $p_{\text{random}} = 0.0625$ ).


For Figure 2.19b, timescales are long enough for SVMC to yield non-zero



**Figure 2.19:** Reverse spin-vector Monte Carlo (SVMC) anneals of the four spin chain from each of the computational first-excited states for *a*) no pause and *b*) a pause of  $2\mu\text{s}$  at a system temperature of 12.26 mK. The maximum relative magnitude of transverse-field applied to the FSC is defined as  $Q(s^*) = A(s^*)/B(s^*)$ , where  $s^*$  is the minimum value of the annealing parameter in a reverse anneal. The black dashed line marks the position of the minimum gap, and the orange dashed line marks where the semiclassical energy potential starts to become unimodal at  $Q(0.4)$ . Each data point is a median of 50 bootstrapped probabilities with 95% confidence intervals, calculated from 1000 anneals. The reverse anneal parameter schedules were the same as those used in Figure 1.6, which parameterize the Hamiltonian coefficients defined in the Figure 1.8 for the LANL D-Wave 2000Q quantum annealer. The number of increments (sweeps) in the SVMC anneal is related to the total annealing time  $t_a = 2(3 - 2s^*)\mu\text{s}$  by 1000 steps  $\mu\text{s}^{-1}$ .


probability for low values of  $Q(s^*)$ , which is known to be a region of system freeze-out experimentally. In order to remedy this, we use a SVMC with transverse-field dependent updates (SVMC-TF), which was first presented in Ref. [119] (also see Section 1.4.3 for more details). This is where the spin rotors used to explore the semiclassical potential become more restricted as transverse-field decreases, such that SVMC-TF can emulate freeze-out effects. This is shown to be true when reverse annealing with SVMC-TF in Figure 2.20, and also results in a slight increase in ground-state probability, but not enough to reach the random probability threshold. Again a peak due to thermalization occurs, but it is less pronounced compared to SVMC, and again the probability also becomes constant after the potential becomes unimodal.

Finally, we look at quantum Monte Carlo methods to simulate reverse annealing,



**Figure 2.20:** Plots of reverse annealing the four spin chain with *a*) no pause and *b*) a pause of  $2\mu s$  when simulated by SVMC with transverse-field dependent updates (SVMC-TF) starting in each of the computational first-excited states. The simulation uses the same settings as those defined in Figure 2.19.


specifically looking at path-integral Monte Carlo (PIMC) which is shown to emulate quantum annealing [29, 30, 38, 60]. The FSC has noise and cross-talk applied to it in the same fashion as the classical methods had, and the PIMC cost function in Equation (2.29) follows the schedule used by the LANL D-Wave 2000Q. Single spin-flips are used to update the system in PIMC, and therefore results in PIMC simulations that exhibit nonequilibrium dynamics. Dynamics closer to thermal equilibrium can be achieved by using cluster update methods [26, 141], but seeing as we are not concerned with calculating exponents, and because the features in the experimental data are likely caused by nonequilibrium behaviour, we do not look to implement these more advanced update methods. The results of reverse annealing with PIMC are shown in Figure 2.21, whereby the range of  $Q(s^*)$  simulated is extended to show the full range of dynamics up to  $s = 0$  (maximum transverse-field). Remarkably, PIMC achieves a separation of states is in order of Hamming weight for low values of  $Q(s^*)$ , and a ground state probability similar to what is experimentally observed. As expected, the separation of states converge for large values of  $Q(s^*)$  unlike the experimental data, but this simulation does not take into account spin-bath polarization.



**Figure 2.21:** PIMC reverse anneals of the four spin chain from each of the computational first-excited states for *a*) no pause and *b*) a pause of  $2\mu s$  at a system temperature of 12.26 mK. The maximum relative magnitude of transverse-field applied to the FSC is defined as  $Q(s^*) = A(s^*)/B(s^*)$ , where  $s^*$  is the minimum value of the annealing parameter in a reverse anneal. The black dashed line marks the position of the minimum gap, and the orange dashed line marks where the semiclassical energy potential starts to become unimodal at  $Q(0.4)$ . Each data point is a median of 50 bootstrapped probabilities with 95% confidence intervals, calculated from 1000 anneals. The reverse anneal parameter schedules were the same as those used in Figure 1.6, which parameterize the Hamiltonian coefficients defined in the Figure 1.8 for the LANL D-Wave 2000Q quantum annealer. The number of Trotter slices used is  $N_\tau = 20$ , and the total number of increments (sweeps) in the PIMC anneal is related to the total annealing time  $t_a = 2(3 - 2s^*)\mu s$  by 50 steps  $\mu s^{-1}$ .

The dynamics of PIMC with and without pauses become active about the point in a reverse anneal where the potential becomes unimodal, instead of it being about minimum gap as seen in SVMC and the experimental data. To see whether this is corrected by emulating quantum noise sources such as dephasing (i.e., adding Equation (2.31) to the PIMC cost function), PIMC is simulated under three different bath-coupling regimes that control the extent of dephasing in the system. The tunability of the bath-coupling allows this version of PIMC to enter intermediate coupling regimes that are inaccessible to master equations such as the AME [143]. Reverse anneals with a  $2\mu s$  pause are illustrated in these regimes in Figure 2.22, where a lower value of  $\alpha$  represents a weaker coupling to the bath. For  $\alpha = 0.2$ , dephasing is not substantial enough to effect the system, and therefore resembles PIMC without an explicit dephasing term. As  $\alpha$  is increased to 0.6, the system





**Figure 2.22:** Reverse anneal simulations of the four spin chain with a pause of  $2\mu s$  and using PIMC with dephasing. The initial states are from each of the computational first-excited states, and the extent of dephasing is represented here by three regimes of the bath-coupling, parameterized by  $\alpha$  at a system temperature of 12.26 mK. The simulations use the same settings as those defined in Figure 2.21.

dynamics start to become frozen along the Trotter dimension of the system such that it acts as a collective spin, therefore leading to dynamics that resemble those seen from simulated annealing, as shown by the increase in probability.

For  $\alpha = 1.0$ , the system starts to freeze-out completely along the Trotter dimension as dephasing dominates dynamics. This results in the largest separation of states seen so far, but means that dynamics are only accessible for the largest values of  $Q(s^*)$  possible on the LANL D-Wave 2000Q schedule. Despite there being no peak in probability that is representative of thermalization dynamics, PIMC supports the physical intuition that initial state Hamming distance does affect ground state probability for low values of  $Q(s^*)$ , making it the model closest to what is observed experimentally. The absence of consistent state separation for all simulated models supports the hypothesis that spin-bath polarization is affecting experimental reverse anneals, as other integrated control errors such as cross-talk and noise do not replicate the experimental observations.

In conclusion, the FSC problem Hamiltonian (Figure 2.10) used to assess reverse annealing as an optimization method has demonstrated that it does behave as a local search method, with ground state probabilities being larger for an initial state

that is closer in Hamming distance to the ground state for an experimental reverse anneal on the LANL D-Wave 2000Q quantum annealer. Initial states further in Hamming distance also correspond to a lower ground state probability, and therefore longer times-to-solution. Pauses were used mid anneal to increase ground state probability and decrease TTS near the minimum gap (Figure 2.13b), which is a result of thermalization effects driving diabatic transitions, and is explored further in Chapter 3. The increase of ground state probability was particularly prominent about the minimum gap, and is largest before passing through the minimum gap, suggesting that reverse annealing can be used to successfully compute solutions to optimization problems without passing through the minimum gap. However, in order to achieve this the initial state needs to be in one of the low-energy states of the original problem, and it is hard to find a priori the position of the minimum gap.

There are also some caveats to be had with this method of optimization, as despite the fact that pausing increases ground state probability, the benefit to TTS is nullified by the increase in computation time. Additionally, all reverse anneal times-to-solution are longer than if one performed the corresponding forward anneal (Figure 2.14). The apparent local-search behaviour seen in the experimental data is also not something that a system at equilibrium would exhibit, and by excluding other integrated control errors such as parameter noise and cross-talk using several simulated models of quantum annealing, the local-search behaviour is likely a result of spin-bath polarization. This effect introduces erroneous biases into reverse annealing due to the adjacent biasing SQUID used to initialize the state. Therefore, for an optimization problem with a large condition number (i.e., one that is sensitive to errors), reverse annealing may alter the solution space to the extent that the ground state of the original problem is no longer the ground state being found experimentally. This can be mitigated using techniques such as shimming, but it requires many prior reverse anneals to find the set of counter-biases to apply to each qubit.

Another artefact that was found to be the cause of some integrated control error is the presence of oscillations in ground state probability. This is found to occur when there was a reverse anneal with a constant annealing rate and a varying  $s^*$  parameter,

and is thought to be the result of some quantization error in the transverse-field magnitude that affects the ground state probability. Neither these oscillations nor constant probability differences due to the initial state are seen in the simulation models tested. PIMC held the closest resemblance to the experimental data in terms of ground state probability and correct separation of states for low values of  $Q(s^*)$  (relative maximum magnitude of transverse field). Other models such as simulated annealing and SVMC did not exhibit probability comparable to the experimental data or any initial state dependence on the ground state probability, but did have peaks in probability when the system was paused and left to thermalize, something that did not feature in PIMC simulations. The AME in the singular coupling limit also did not resemble the reverse anneal, and despite it having probability differences, it did not agree with what was seen experimentally. Future work with reverse annealing would involve finding regimes where TTS may be lower than what can be found using forward annealing. Larger optimization problems would be needed to find this regime, and would also require the use of shimming in order to mitigate spin-bath polarization effects that bias reverse annealing.

## Chapter 3

---

# Thermalization in Quantum Annealing

---

Modern realizations of quantum computers all must find methods to mitigate the noise sources that cause quantum systems to decohere, as coherence is key to the theoretical foundation on which quantum computation was established. However, noise will always manifest itself after some timescale in superconducting qubits, and without error-correction protocols it is the major limiting factor to computation. In the context of D-Wave quantum annealers, coherence lifetimes are on the order of tens of nanoseconds due to the rate at which noise permeates into the system dynamics. Therefore, the quantumness of the dynamics used for computation on quantum annealers is still subject to debate due to the prevalent quantum and classical noise sources that can obscure coherent quantum processes [143, 225–227]. However, it has been shown that incoherent processes can aid computation on quantum annealers where the anneal times are orders of magnitude larger than the single-qubit decoherence time [37, 38, 53]. This is largely in the form of thermally driven transitions between the energy eigenstates, such that if this transition is seen to occur between states separated by a classically forbidden region, then this process is also known as thermally-assisted tunnelling, inferring that tunnelling for computation does not necessarily require coherent dynamics.

These effects notably affect the ground state probability when the system is paused mid-anneal and allow the system to thermalize near the minimum gap [57, 118, 119]. However, to what extent thermalization is computationally beneficial

is still an active area of research, with lower noise quantum annealers (i.e. those with lower thermalization rates) being seen to improve tunnelling ranges in local search (reverse) quantum anneals [203]. A theoretical foundation as to why ground state probability increases when a system is paused is given in Ref. [118], and finds that this thermalization phenomenon only occurs when (i) the thermalization rate is decreasing after the minimum gap, (ii) the thermalization rate is large near the minimum gap, (iii) the thermalization rate is small at the end of an anneal, and (iv) the ground state at the end of the anneal is sub-thermal (i.e., is lower in probability than is predicted by the Boltzmann distribution). If these criteria are met, then thermalization phenomenon driven by quantum dynamics are seen. However, it was also shown in Ref. [119] that the thermalization signature of a quantum anneal with a pause can be replicated with spin-vector Monte Carlo (SVMC) [277], which is a classical heuristic used to mimic the behaviour of physical quantum annealers [60,89]. Therefore, it is possible that classical dynamics could be responsible for the thermal transitions observed in experimental quantum anneals.


In order to differentiate between the classical and quantum mechanisms driving the thermal transitions on a quantum annealer, it is necessary to use artificial gadgets that are hard to solve in the transverse-field Ising model [29, 31, 38, 60, 292]. These gadgets possess properties such as false minima within the energy potential that classical evolutions would follow to suboptimal solutions. Increasing the hardness to better contrast between dynamics generally relies on increasing the system size, and therefore thermalization cannot be simulated using master equations after the Hamiltonian exceeds a few qubits. In Section 3.1, a new Hamiltonian, the perturbed ferromagnetic chain, is presented where the degree of hardness can be controlled by a tuneable parameter at a fixed system size (published in Ref. [117]). The properties of the perturbed ferromagnetic chain are studied, and it is shown that it possesses both a false minimum and an exponentially large (in system size) first-excited-state manifold. Both classical and open quantum system dynamical models are simulated in Section 3.2, and demonstrate that ground state probability differs by orders of magnitude between the models, suggesting that quantum thermalization

is driving computation. Finally, experimental evidence of quantum thermalization is presented in Section 3.3 for different D-Wave quantum annealers and various systems sizes of the perturbed ferromagnetic chain, showing that thermalization is most prevalent in time-scales many orders of magnitude larger than the coherence lifetime of the quantum annealers. We also present evidence that suggests that thermal transitions that appear to happen between states that have a large Hamming distance use intermediate states in order to complete this transition.

### **3.1 Perturbed Ferromagnetic Chain**

Many classical problems can in fact become harder to solve when placed in the transverse-field Ising model (TFIM) compared to a purely classical models, such as simulated annealing, due to the presence of harsher quantum phase transitions. An example of this is the random energy model, which develops a first order quantum phase transition in the TFIM compared to one of second order phase transition in the classical model [82]. The perturbed ferromagnetic chain (PFC) [117] possess this feature in the TFIM, and is inspired from other quantum consistency tests presented in Refs. [137, 159], where these cyclic gadgets have highly degenerate ground states. If annealed classically, then probability would be equally shared amongst the ground states, but in a quantum anneal it is shown to populate all states apart from one. By adding a perturbative parameter, one can in fact break the degeneracy of this ground state into a single ground state and a highly degenerate first excited state. Additionally, the minimum gap energy,  $\Delta_{10}$ , can be tuned by the perturbative parameter, providing access to different minimum gap regimes e.g., a gap smaller than the environmental temperature.

Throughout the rest of this section, we present the theoretical properties of the PFC and compare classical spin-vector Monte Carlo (SVMC) variants with the adiabatic quantum master equation. We demonstrate that SVMC methods get trapped in the exponentially large first-excited computational state manifold when solving this frustrated problem, whereas evolution using quantum dynamics remains in the lowest energy eigenstates. This results in significant differences in ground-state



**Figure 3.1:** Illustration of the perturbed ferromagnetic chain Hamiltonian in general form. In yellow are the auxiliary qubits with biases of  $-R$ , and in turquoise are the backbone qubits with biases  $R(1-d)$ . The dark blue edges are the ferromagnetic couplers of strength  $-R$ . The total number of qubits,  $N$ , in this problem is  $2M$ , where  $M$  denotes the number of two-qubit subsystems (indexed by  $i$ ) in the Hamiltonian. The system energy is scaled by  $R$ , and  $d$  describes the magnitude of perturbation. The properties of this model hold generally for  $M \geq 2$ ,  $R > 0$ , and  $1 > d > 0$ .

probability when using either classical or quantum annealing dynamics in the TFIM. It is also shown that the mechanism contributing to these significant differences is that of thermally-assisted tunnelling in a system that experiences decoherence at a finite temperature, such that we can observe its computational role when annealing with the PFC.

### 3.1.1 Classical Model

The PFC (Figure 3.1) is a ferromagnetically coupled chain of  $M$  frustrated subsystems, each composed of two qubits, and is designed to be an easily embeddable gadget that can be used as an experimental test of hardness in the TFIM. The PFC Hamiltonian is given by

$$\hat{H}_P = R \left[ \sum_{i=1}^M (1-d) \hat{\sigma}_{b,i}^z - \hat{\sigma}_{a,i}^z - R \hat{\sigma}_{a,i}^z \hat{\sigma}_{b,i}^z \right] - \sum_{i=1}^{M-1} \hat{\sigma}_{b,i}^z \hat{\sigma}_{b,i+1}^z \quad (3.1)$$

where  $M$  is the number of subsystems,  $R$  scales the energy of the problem, and the magnitude of the perturbation is characterized by the parameter  $0 < d < 1$ . The value

of  $d$  is proportional to the size of the energy gap between the ground and first-excited state, and therefore controls the problem hardness. The parameter  $d$  is referred to as the perturbative parameter throughout. The auxiliary and backbone qubits, depicted by the yellow and turquoise circles shown in Figure 3.1 respectively, are denoted by Pauli Z matrices  $\hat{\sigma}_a^z$  and  $\hat{\sigma}_b^z$  respectively.

The ground state of this Hamiltonian is the  $|0_b^{\otimes M}, 0_a^{\otimes M}\rangle$  (all up) state, and given that  $d < (M-1)^{-1}$  the first excited state is a degenerate manifold whose size grows exponentially in  $M$ . This validity range exists due to the minimum gap between the manifold states and the ground state being  $\Delta_{10} = 2RMd$ , which states that gap size is proportional to system size, leading to the manifold state no longer being the first excited state for large  $M$ . This manifold always has the backbone qubits in the  $|1_b\rangle^{\otimes M}$  (all down) configuration, and the auxiliary qubits are isoenergetic with respect to their spin state, creating a  $2^M$ -degenerate manifold of “floppy” auxiliary qubits. This creates an energy gap between the ground state and exponential manifold of  $\Delta_{10} = 2RMd$ , where the states in the manifold have a Hamming distance between  $M$  and  $2M$  from the ground state.

Classically the PFC is exactly solvable via the transfer matrix method [293], where at an inverse temperature,  $\beta$ , the partition function,  $\mathcal{Z}$ , can be found in polynomial time. The partition function in transfer matrix form is expressed as

$$\mathcal{Z} = \mathbf{v} \mathbf{W}^{M-1} \mathbf{v}^T, \quad (3.2)$$

where,

$$\mathbf{v} = \left( e^{\frac{1}{2}\beta R(d+1)}, e^{\frac{1}{2}\beta R(d-3)}, e^{\frac{1}{2}\beta R(1-d)}, e^{\frac{1}{2}\beta R(1-d)} \right) \quad (3.3)$$

handles the boundary subsystems of the chain, and



$$\mathbf{W} = \begin{bmatrix} e^{\beta R(d+2)} & e^{\beta R d} & 1 & 1 \\ e^{\beta R d} & e^{\beta R(d-2)} & e^{-2\beta R} & e^{-2\beta R} \\ 1 & e^{-2\beta R} & e^{\beta R(2-d)} & e^{\beta R(2-d)} \\ 1 & e^{-2\beta R} & e^{\beta R(2-d)} & e^{\beta R(2-d)} \end{bmatrix}, \quad (3.4)$$

handles the inner subsystems of the chain. We can then find the magnetization of a subsystem at thermal equilibrium using

$$\langle \sigma_i^z \rangle = \frac{1}{\mathcal{Z}} [\mathbf{v} \mathbf{W}^{i-1} \sigma_i^z \mathbf{W}^{M-i} \mathbf{v}^T], \quad (3.5)$$

where  $\sigma_i^z = \frac{1}{2} (\hat{\sigma}_{a,i}^z + \hat{\sigma}_{b,i}^z)$ . The average magnetization of the PFC is then an average over all contributions,  $\langle \sigma^z \rangle = \frac{1}{M} \sum_i^M \langle \sigma_i^z \rangle$ . From this formalism, the free-energy of the PFC can be derived, which can then be used to derive further thermodynamic properties. The transfer-matrix representation of the PFC partition function (Equation (3.2)) involves singular, noncommuting matrices  $\mathbf{W}$  (Equation (3.4)) and  $\mathbf{V} = \mathbf{v}^T \mathbf{v}$  (where  $\mathbf{v}$  is defined in Eq. (3.3)), which does not allow for an obvious reduction to an analytical free energy that is generally defined as

$$F = - \lim_{N \rightarrow \infty} \frac{1}{\beta N} \ln \mathcal{Z}, \quad (3.6)$$

where  $\mathcal{Z}$  is the partition function. We begin by redefining the partition function in Equation (3.2) to

$$\mathcal{Z} = \text{Tr}(\mathbf{W}^{M-1} \mathbf{V}), \quad (3.7)$$

where  $M = N/2$ , such that we now take the limit in  $M \rightarrow \infty$  due to  $\mathbf{W}$  containing information about the subsystem rather than a single qubit. Performing an eigen-decomposition on  $\mathbf{W}$  yields a diagonal matrix of eigenvalues,  $\mathbf{D}$ , and a matrix of eigenvectors,  $\mathbf{P}$ , in the form  $\mathbf{W} = \mathbf{P} \mathbf{D} \mathbf{P}^{-1}$ . When  $\mathbf{W}$  is raised to any power, the decomposition simply becomes  $\mathbf{W}^M = \mathbf{P} \mathbf{D}^M \mathbf{P}^{-1}$ . Given the cyclic invariance of the

trace, the partition function therefore becomes

$$\mathcal{Z} = \text{Tr}(\mathbf{P}\mathbf{D}^{M-1}\mathbf{P}^{-1}\mathbf{V}) = \text{Tr}(\mathbf{D}^{M-1}\mathbf{P}^{-1}\mathbf{V}\mathbf{P}). \quad (3.8)$$

Taking  $\mathbf{A} = \mathbf{P}^{-1}\mathbf{V}\mathbf{P}$ , and the largest absolute eigenvalue of  $\mathbf{W}$  to be  $\lambda_1$  (the spectral radius), the trace summation will yield

$$\begin{aligned} \text{Tr}(\mathbf{D}^{M-1}\mathbf{A}) &= \sum_{i=1}^4 A_{ii}\lambda_i^{M-1} = \sum_{i=1}^4 \left( \frac{A_{ii}\lambda_i}{\sqrt{A_{ii}}} \right)^{M-1} \\ &= \left( \frac{A_{11}\lambda_1}{\sqrt{A_{11}}} \right)^{M-1} \sum_{i=1}^4 \left( \frac{A_{ii}\lambda_i}{\sqrt{A_{11}\lambda_1}} \right)^{M-1}. \end{aligned} \quad (3.9)$$

Given that  $\mathbf{W}$  is positive semidefinite, in the limit of  $M \rightarrow \infty$  the trace is simply left with  $\lambda_1^{M-1}$  as a nonvanishing term, such that we find the analytical form of the free energy to be

$$F = - \lim_{M \rightarrow \infty} \frac{1}{\beta M} \ln \text{Tr}(\mathbf{D}^{M-1}\mathbf{A}) = - \frac{1}{\beta} \ln \lambda_1. \quad (3.10)$$


This eigenvalue can be found symbolically using PYTHON's *SymPy* library, such that

$$\begin{aligned} \lambda_1 &= e^{2\beta R} \cosh \beta R d + \cosh \beta R (2-d) \\ &+ \sqrt{(e^{2\beta R} \cosh \beta R d + \cosh \beta R (2-d))^2 - 4 \sinh 4\beta R}. \end{aligned} \quad (3.11)$$


The validity of this free energy relies on the fact that our spectral radius is  $\geq 1$ , otherwise  $\mathbf{W}$  would converge to zero in the limit of  $M \rightarrow \infty$ . However, the PFC has bounds of  $R > 0$  and  $0 < d < 1$ , such that our eigenvalue  $\lambda_1 \geq 4$ , and is therefore finite everywhere except for  $\beta R = \infty$ .

### 3.1.2 Transverse-Field Ising Model

Translating the PFC into the TFIM involves the addition of a noncommuting transverse-field term, composed of  $\hat{\sigma}^x$  operators, which introduces the driver of quantum fluctuations that can potentially be used to aid computation [14, 47]. The



(a)



(b)

**Figure 3.2:** (a) Intensity plot of the average qubit magnetization in the instantaneous ground state for a PFC in the presence of a transverse field (Equation (3.12)) with  $M = 2$  and  $R = 1.0$ . The solid red line shows the boundary between the quantum paramagnetic and negative magnetization phases. The white dashed line indicates the position of the minimum gap. (b) Cross sections of (a) showing the average magnetization during an anneal.

TFIM Hamiltonian of the PFC is given by

$$\hat{H}(s) = -A(s) \sum_{j=1}^N \hat{\sigma}_j^x + B(s) \hat{H}_P, \quad (3.12)$$


where the classical PFC is encoded into  $\hat{H}_P$  [Eq. (3.1)]. The coefficients are taken to be  $A(s) = 3(1 - s)$  GHz and  $B(s) = 3s$  GHz throughout this work, where  $s$  is the normalized annealing time  $s = t/t_{\text{anneal}}$ .

For sufficiently small values of  $d$ , the ground state of this Hamiltonian is seen to undergo a quantum phase transition, illustrated in Figure 3.2 by the change in average qubit magnetization from negative to positive phases. Therefore, in the TFIM, the perturbative parameter  $d$  can be used to tune the problem hardness. The average qubit magnetization is defined as

$$\langle \sigma^z \rangle = \frac{1}{N} \sum_{j=1}^N \langle E_0(s) | \hat{\sigma}_j^z | E_0(s) \rangle, \quad (3.13)$$

where  $|E_0(s)\rangle$  is the instantaneous ground state from the diagonalized Hamiltonian of Eq. (3.12) at some value of  $s$ , and  $N$  is the number of qubits in the PFC. The formation of the negative phase before the minimum gap is indicative of the ground-state qubits becoming magnetized to resemble the exponentially large degenerate first-excited-state manifold (further illustrated in Sections 3.1.3 and 3.2.2). After passing through the minimum gap, the instantaneous ground state enters the positive phase and then goes on to finish in the computational ground state. It must be noted that even if the  $d < (M - 1)^{-1}$  condition is broken such that the exponentially degenerate manifold is no longer the computational first-excited-state, for small  $d$  the instantaneous ground state maintains its resemblance to the exponentially degenerate manifold before the minimum gap.

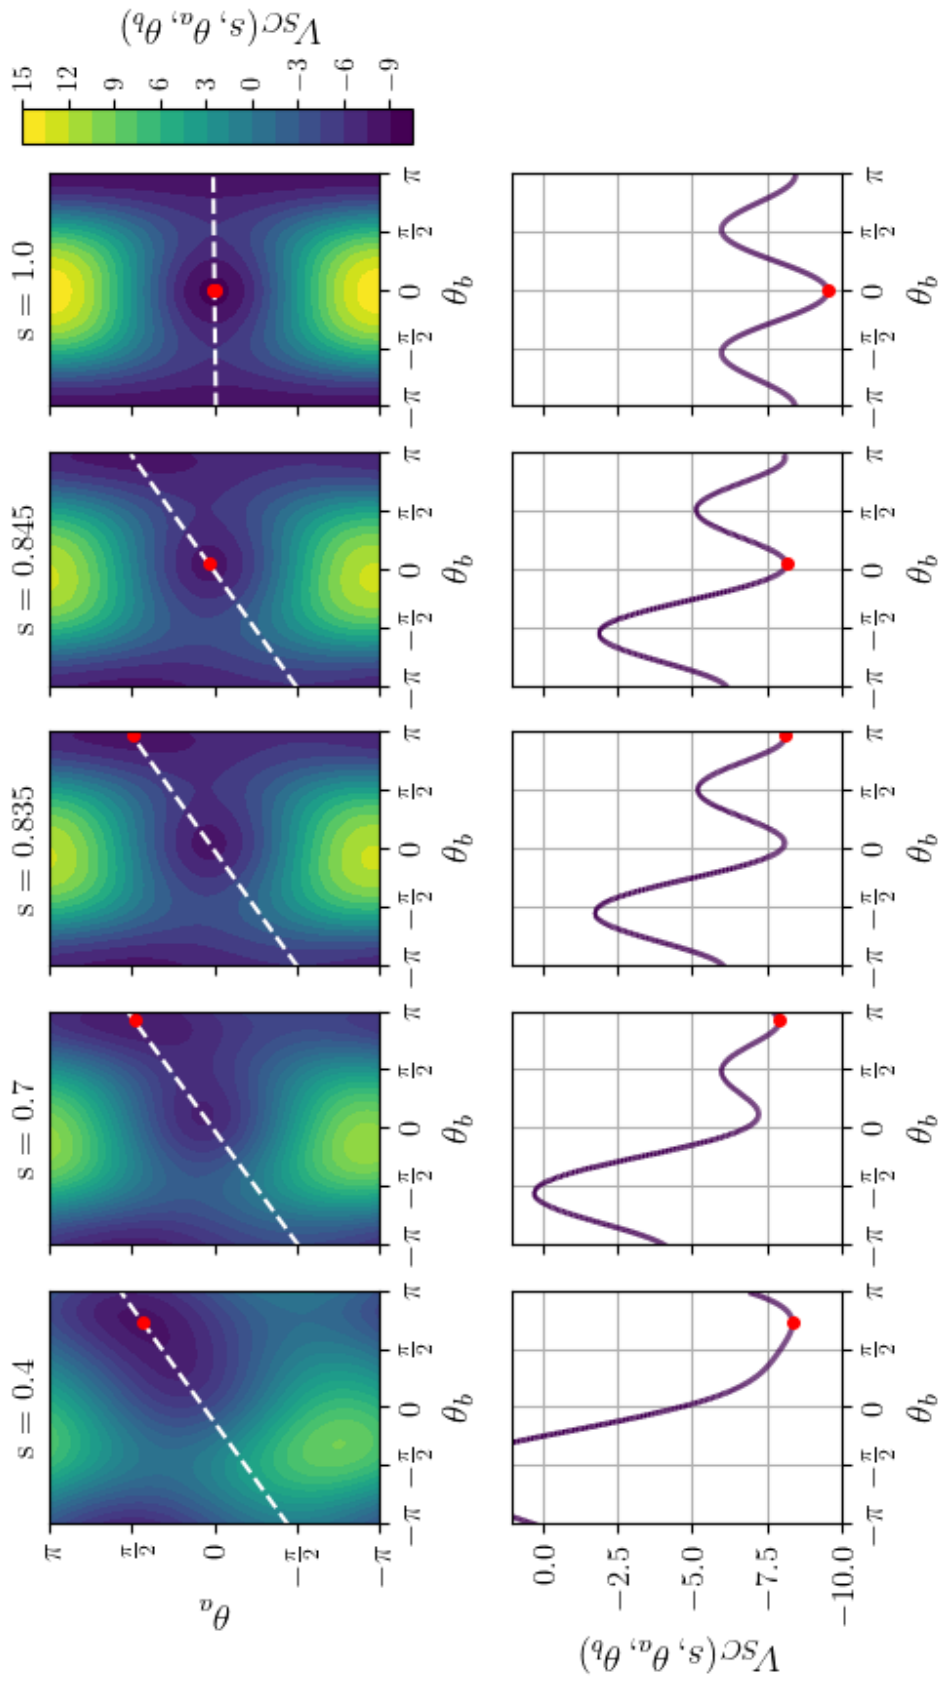
Figure 3.3 illustrates the exponential scaling of the minimum gap between the ground and first-excited state,  $\Delta_{10}$ , in system size,  $M$ , for various perturbative parameters,  $d$ . The gap sizes were determined through numerical diagonalization of the TFIM Hamiltonian (Equation 3.12) for all  $s$ , and follow exponential fits of



**Figure 3.3:** Plot of how the minimum gap  $\Delta_{10}$  scales with system size  $M$  in the TFIM for various values of the perturbative parameter  $d$ . The PFC used  $R = 1.0$ , and the lines plotted are fits to the equation  $Ae^{-bM}$ .

the form  $Ae^{-bM}$ . The exponential minimum gap scaling in  $M$  is also consistent with first-order quantum phase transitions, therefore making  $M$  another method of increasing the problem hardness. However, large quantum systems lead to intractable computational times for some of the simulation methods explored in the results, making  $d$  a more desirable tunable hardness parameter because of the minimum gap  $\Delta_{10} \rightarrow 0$  as  $d \rightarrow 0$ .

In summary, the PFC becomes hard in the TFIM for small values of  $d$  due to the presence of a quantum phase transition, and for large values of  $M$  where the minimum gap exponentially reduces in size. In the next section, we perform a semiclassical analysis to show that when the PFC is translated into the TFIM a false minimum exists. The interplay between the false and true minima is particularly prominent when annealing through the region where the ground state is in the negative magnetization phase shown in Figure 3.2.



**Figure 3.4:** Plots of the semiclassical energy potential (top row) and the energy along the hyperplane (white dashed line) passing through the landscape (bottom row) at the specified values of normalized time,  $s$ . This potential is for a PFC with  $M = 2$ ,  $R = 1.0$ , and  $d = 0.09$ . The red marker indicates the global minimum of the landscapes. The backbone and auxiliary spins in the PFC are parametrized by angles  $\theta_b$  and  $\theta_a$ , respectively.

### 3.1.3 Semi-Classical Analysis

To further explore the behaviour of the PFC in the TFIM, a semiclassical approximation can be made by using the spin-coherent ansatz [275] (see Section 1.4.3) to calculate the semiclassical effective potential landscape as a function of  $s$ . The magnetization expectation values of the auxiliary qubits are almost identical, and the same is true for the backbone qubits. We can therefore approximate the states of the PFC in the spin-coherent ansatz as

$$|\theta_a, \theta_b\rangle = \left[ \bigotimes_{i=1}^M \cos\left(\frac{\theta_a}{2}\right) |0\rangle + \sin\left(\frac{\theta_a}{2}\right) |1\rangle \right] \otimes \left[ \bigotimes_{i=1}^M \cos\left(\frac{\theta_b}{2}\right) |0\rangle + \sin\left(\frac{\theta_b}{2}\right) |1\rangle \right]. \quad (3.14)$$

Here,  $\theta_a$  and  $\theta_b$  are the angles of the states in the  $XZ$  plane of the Bloch spheres for all of the auxiliary and backbone qubits respectively. Here, we assume that the azimuthal angle  $\phi_j$  is equal to zero as there are no  $Y$ -contributions to the TFIM Hamiltonian eigenvalues. This semiclassical approximation without the  $Y$ -component has been used to simulate specific problems in the TFIM [277]. The semiclassical potential is then given by

$$V_{SC}(s, \theta_a, \theta_b) = \langle \theta_a, \theta_b | \hat{H}(s) | \theta_a, \theta_b \rangle. \quad (3.15)$$

The visual representation of the potential at various stages of an anneal (Figure 3.4) shows the PFC initially taking a path to the first excited states ( $\theta_a \in [-\pi, \pi], \theta_b = \pi$ ). This is then followed by a discontinuous change in the position of the global energy minimum about the minimum gap (at  $s = 0.841$ ) to the computational ground state ( $\theta_a = 0, \theta_b = 0$ ). This is referred to as a false minimum, which is where an initial single minimum forms two or more minima mid-anneal, where the lowest minimum at this point (i.e., the most probably route) leads to an excited state at the end of the anneal. This false minimum will always be separate from the minimum that leads to the ground-state solution past this point. The false minimum can be visualised by taking a hyperplane that passes through the global minimum (and the local minimum where applicable). It is clear that as we evolve

from a unimodal to a bimodal potential, the all-down ( $\theta_b = \pi$ ) configuration of the backbone qubits is energetically preferable until the minimum gap is traversed. Additionally, the computational ground state is energetically isolated from the low-energy excited states, meaning that further dynamical evolution is still needed to reach the ground state after the minimum gap. If the system is evolved under an adiabatic, coherent regime [13] then the dynamical process is quantum tunnelling. In a classical model (like SVMC) we can only use thermal excitations to traverse these energy barriers and escape the false minimum.


If tunnelling were to occur in the instantaneous ground state, then this would result in delocalization about the bistable potential. Using the trace-norm distance,

$$D(s, \theta_a, \theta_b) = \sqrt{1 - |\langle E_0(s) | \theta_a, \theta_b \rangle|^2}, \quad (3.16)$$

we can quantify the distance between the instantaneous ground state and the spin-coherent ansatz to determine the extent to which the ansatz accurately describes the instantaneous ground state. In Figure 3.5 we show the trace-norm distance for a four-qubit ( $M = 2$ ) instance of the PFC in the vicinity of the minimum gap. This is the same instance whose potential landscape is shown in Figure 3.4 and whose minimum gap occurs at  $s = 0.841$ . At  $s = 0.835$  the trace-norm distance has a global minimum whose location in  $(\theta_a, \theta_b)$  space closely corresponds to the global minimum of the semiclassical potential (i.e., near  $\theta_b = \pi$ ). Nevertheless, there is a local minimum in the trace-norm distance which extends along the indicated hyperplane, showing that, prior to the minimum gap, tunnelling enables a finite probability amplitude in the local minimum of the potential landscape near  $\theta_b = 0$ . After the minimum gap, as shown for the case  $s = 0.845$  in the right panel of Figure 3.5, the global minimum of the trace-norm distance now closely corresponds to the global minimum of the semiclassical potential near  $\theta_b = 0$ . This minimum continuously evolves into the global minimum of the problem Hamiltonian at  $s = 1$  as shown in Figure 3.4.

Using semiclassical analysis, we have shown the existence of a false global minimum before the minimum gap, when the value of  $d$  is small enough. Under quantum evolution, the transition from the false minimum to the true minimum





**Figure 3.5:** The trace-norm distance between the spin-coherent state and the instantaneous ground state (Equation (3.16)) in the vicinity of the minimum gap at  $s = 0.841$  for a PFC with  $M = 2$ ,  $R = 1.0$  and  $d = 0.09$ . The backbone and auxiliary spins in the PFC are characterized by angles  $\theta_b$  and  $\theta_a$ , respectively. The red cross marker indicates the global minimum of the trace-norm landscapes, and the red circle indicates the minimum of the potential landscape in Figure 3.4. The hyperplane (white dashed line) from Figure 3.4 is also plotted.

exploits tunnelling, and this is visualized by measuring the trace-norm distance between spin-coherent states and the instantaneous ground state to show delocalization across the potential barrier (Figure 3.5). Under classical evolution, thermal excitation of multiple qubits is needed to traverse the barrier to reach the computational ground state. Additionally, the manifold along  $\theta_b = \pi$  (corresponding to the exponentially large computational first-excited-state manifold) is equally accessible under classical dynamics. Therefore, classical algorithms that explore this energy landscape, such as SVMC, can remain in this manifold instead of reaching the ground state. We explore in Section 3.2.2 the extent to which this hinders computation in both classical and quantum evolution.

## 3.2 Simulated Thermalization

### 3.2.1 Quantum & Classical Methods

The two main methods to simulate the TFIM with coupling to a thermal bath are through the use of spin-vector Monte Carlo (SVMC) and master equations. The former is a fully classical representation of the TFIM, which can be thought of as

a classical particle following the continuous semiclassical potential. This model uses stochastic temperature dependent updates to the spin vector (parameterised by angles) to traverse the energy potential and reach the ground state. It has been used to simulate experimental quantum annealing and its thermalization effects [89, 119, 159, 277], and has therefore raised the question as to how quantum the D-Wave quantum annealer is if it can be simulated by a classical emulator. The master equations are fully quantum representations of the dynamical evolution, where both coherent and incoherent evolutions through a quantized energy potential are modelled. Entanglement between qubits and tunnelling through barriers in the energy potential is possible in this regime, as well as thermalisation between energy levels for master equations that model incoherent processes.

To explore this with the PFC, we look at SVMC and several variants as the classical simulator of open-system dynamics. Numerical approaches such as spin-vector dynamics (SVD) that use the Langevin equation (Section 1.4.3) could also be used for the classical simulation, but in the dynamical simulations we are interested in the distribution of states generated, which is better represented by SVMC. The dynamics simulated by SVD have also not been accurate representations of what is seen in experimental quantum annealing [138, 159]. The variants of SVMC to be explored include the original SVMC proposed in Ref. [277], and SVMC with transverse-field dependent updates (SVMC-TF) proposed in Ref. [119], details of which are in Section 1.4.3.1. Both SVMC and SVMC-TF are dynamically restricted to only operating in the  $XZ$ -plane of the Bloch sphere, however, this restriction is removed by including the azimuthal angle  $\phi_j$  such that SVMC now has access to the entire Bloch sphere. We will refer to these variants as spherical SVMC and spherical SVMC-TF. This coordinate extension does not affect the  $Z$  (polar) components of the energy function. It does, however, affect the transverse-field component, such that the new energy function becomes

$$\mathcal{E}(s) = -A(s) \sum_{j=1}^N \cos \phi_j \sin \theta_j + B(s) \left[ \sum_{j=1}^N h_j \cos \theta_j + \sum_{\langle j,k \rangle} J_{jk} \cos \theta_j \cos \theta_k \right], \quad (3.17)$$

where the azimuthal angle,  $\phi_j \in [-\pi, \pi]$ , is also updated in the same way as the polar angle,  $\theta_j$ . For all variants, we also take the annealing functions to be  $A(s) = 3(1 - s)$  GHz and  $B(s) = 3s$  GHz, where the anneal runs from  $s = 0$  to  $s = 1$  at a temperature of 12 mK. All algorithms update spins individually in a randomly permuted order and thus cannot capture any simultaneous multiqubit moves, unlike those that may occur in systems evolved using quantum dynamics.

The second approach to simulating quantum systems coupled to a thermal-bath is through the use of master equations, all of which are extensions of the von Neumann equation (see Section 1.4.1). In particular, we look at the adiabatic master equation (AME) in the weak coupling limit [142, 143], due to it being far from the semiclassical limit and computationally tractable for the system sizes we are interested in. This model is a reduction of the Redfield equation, where noise is assumed to be Markovian and the rotating-wave approximation made by assuming there is weak coupling between the system energy levels to the bath. In this model, we assume that coupling to the bath introduces decoherence by dephasing all qubits independently, and that the bath spectral density takes the form of a Bosonic Ohmic bath,

$$\gamma(\omega) = 2\pi\eta g^2 \frac{\omega \exp(-|\omega|/\omega_c)}{1 - \exp(-\beta\omega)}, \quad (3.18)$$

where  $\beta = (k_B T)^{-1}$  is the inverse temperature,  $\omega_c$  is the cut-off frequency,  $\eta g^2$  is the dimensionless bath coupling strength and  $k_B$  is the Boltzmann constant. Throughout the rest of the paper we specify the bath parameters to be  $T = 12$  mK,  $\omega_c = 4$  GHz and  $\eta g^2 = 10^{-3}$ . Bath coupling strengths on quantum annealing hardware are thought to be  $\eta g^2 \geq 10^{-3}$ , but higher coupling strengths would violate the weak-coupling approximation used to derive the AME and therefore yield unphysical dynamics.

The decoherence by dephasing manifests itself in the energy eigenbasis through the time-dependent Lindblad operators

$$\hat{L}_{j,\omega_{kl}}(t) = \langle E_l(t) | \sigma_j^z | E_k(t) \rangle | E_l(t) \rangle \langle E_k(t) |. \quad (3.19)$$

This describes how the  $j^{\text{th}}$  qubit couples to the environment with respect to the energy gap,  $\omega_{kl} = E_k - E_l$ , between the instantaneous energy eigenstates,  $|E_k(t)\rangle$  and  $|E_l(t)\rangle$  of the system Hamiltonian (Equation (3.12)). Coupling to the energy eigenstates models how probability density moves between the energy levels near to the ground state, and therefore how quantum annealing can still perform quantum computation after the system has decohered. Coupling to the computational basis would result in strong decoherence to the maximally mixed state, which is not observed to occur in quantum annealing.

This form of the AME is also used to model thermally-assisted adiabatic quantum computation near the adiabatic limit. Assuming that most of the ground state population is lost to the first excited state after passing through the minimum gap, the re-population of the ground state via thermal relaxation can be related to the transition rate via

$$\Gamma_{1\rightarrow 0}(t) \propto \gamma_{1\rightarrow 0}(t) = \gamma(\omega_{10}(t)) \sum_j \left| \langle E_0(t) | \sigma_j^z | E_1(t) \rangle \right|^2. \quad (3.20)$$

Here the temperature dependence of the transition rate is introduced by the Ohmic spectral density function,  $\gamma$  in Equation (3.18). However, when the gap,  $\omega_{10}$ , is sufficiently small, the weak-coupling assumption in the AME starts to break down. This is the case for hard PFC instances, since the gap can be very small relative to the bath temperature. In such a regime the energy levels become broadened due to the stronger coupling to the bath, such that the discrete energy levels should emulate a more continuous potential, similar to the semi-classical picture. Therefore, despite the AME not being able to describe these strong-coupling regimes as accurately as more sophisticated models like the Redfield equation, it serves as a reasonable approximation of an open-system model of the PFC.

For all simulations, the initial state is the pure ground state of the system at  $s = 0$ , which when using Equation (3.12) is  $\hat{\rho}(0) = |+\rangle\langle +|$ . To identify where thermally-assisted dynamics may be affecting computation, all AME simulations are compared against closed system simulations using the von Neumann equation. Both the AME and von Neumann equation are simulated using the Hamiltonian Open


Quantum System Toolkit [145] (HOQST), a simulation library written for the JULIA language.

### 3.2.2 Dynamical Simulations


We begin the dynamical analysis of the PFC by observing the performance of the SVMC variants when scaling in  $M$ . The systems chosen meet the  $d < (M - 1)^{-1}$  condition such that the first excited state is the exponentially large manifold which is at least a Hamming distance of  $M$  away from the ground state. The combined effects of an exponentially scaling gap and manifold are observed in Figure 3.6(a) by measuring the ground-state probability at the end of the anneal with respect to the number of incremental sweeps used in all SVMC variants. Typically, we expect an increasing number of sweeps to correspond to an increasing ground-state probability, but here there are three distinct regimes when annealing the PFC. For low sweep numbers, where the semiclassical potential is evolved in large steps, we have a relatively high ground-state probability as the false minimum is not well resolved but SVMC still guides the spin vector to the low-energy states.

For medium sweep numbers we see reduction in ground-state probability, caused by the quasicontinuous evolution of the semiclassical potential now leading the SVMC algorithm to the false minimum. This guides SVMC to the  $\theta_b = \pi$  manifold (Figure 3.4) corresponding to the degenerate computational first excited states, causing SVMC to spread out into this manifold and into states that are potentially further in Hamming distance from the computational ground state (see Figure 3.7 for further evidence of this). Finally, for high sweep numbers, SVMC starts to thermally equilibrate and ground-state probability begins to return.

To confirm the detrimental role that the exponential manifold has on SVMC, the state probabilities just after the minimum gap are measured in Figure 3.7 for both spherical SVMC-TF (for 10,000 sweeps) and a system evolved using AME (for 200 ns). The probability distribution is measured at  $s = 0.83$  for a PFC ( $M = 3$ ,  $R = 1.0$ ,  $d = 0.1$ ) with a minimum gap at  $s = 0.8227$ . It can be seen that the probability density spreads into the computational first-excited-state manifold when evolving using spherical SVMC-TF, which is caused by following the false minimum



(a)




(b)

**Figure 3.6:** Probability of being in the ground state ( $P_G$ ) for (a) both SVMC (solid lines) and spherical SVMC-TF (dotted lines), and (b) SVMC-TF (solid lines) and spherical SVMC (dotted lines), as the system scales in size,  $M$ , for a PFC with  $d = 0.1$ . Here, all probabilities are found from 20,000 samples, which we repeat 50 times and bootstrap to find the median and 95% confidence intervals for the data point and error bars, respectively.

and accessing the  $\theta_b = \pi$  manifold using classical dynamics. However, a system evolved using the AME remains in the lowest instantaneous eigenstates after the minimum gap, where  $E_0$  also has a large overlap with the computational ground state,  $\langle E_0(s = 0.83) | 0^{\otimes N} \rangle = 0.98$ . After this point in the anneal, both the AME and spherical SVMC-TF experience freeze-out, which prevents any more dynamical evolution that could affect ground-state probability. This is illustrated by the ground-state probability of the AME (spherical SVMC-TF) evolution at  $s = 0.83$  being  $\sim 0.50$  ( $\sim 3.4 \times 10^{-3}$ ), compared to that at  $s = 1$  being  $\sim 0.51$  ( $\sim 3.4 \times 10^{-3}$ ) (see Figure 3.8). It is also worth noting that we measure in the computational basis for SVMC since it is a classical algorithm with no other analogous discrete states to compare against the instantaneous states used by the AME.

Finally, we explore how the tunable hardness parameter,  $d$ , affects the PFC in both quantum and SVMC simulations for  $M = 3$ , in Figure 3.8. We measure the probability at the end of the anneal of being in the ground state as well as any of the first excited states for different annealing durations. The value of  $d$  also determines the size of the minimum gap, such that we span  $d$  to capture various regimes at a fixed system temperature of 12 mK. At  $d \sim 0.227$  we have a minimum gap that approximately equals the system temperature. The form of all SVMC variants when scaling in  $d$  in Figure 3.8 is similar to what is also seen in Figure 3.6(a) when scaling in  $M$ , which is a result of increasing hardness and the variants preferably annealing to the first-excited-state manifold (following the false minimum). Additionally, spherical SVMC-TF consistently outperforms all other variants for the hardest problems ( $d \leq 0.15$ ), which is a result of the spherical component on average reducing the transverse-field in Equation (3.17) and slowed dynamics from the transverse-field dependent updates preventing the spin vector from getting lost in the first-excited-state manifold.

The probabilities at the end of a closed-system quantum anneal (for time  $t_{\text{anneal}}$ ) are of a similar form, with a bump for short anneal times (the diabatic bump [294]) and then reaching the adiabatic limit (where the ground-state probability tends towards 1) at longer anneal times (e.g.,  $t_{\text{anneal}} \simeq 20ns$  for  $d = 0.3$ ). The closed-system

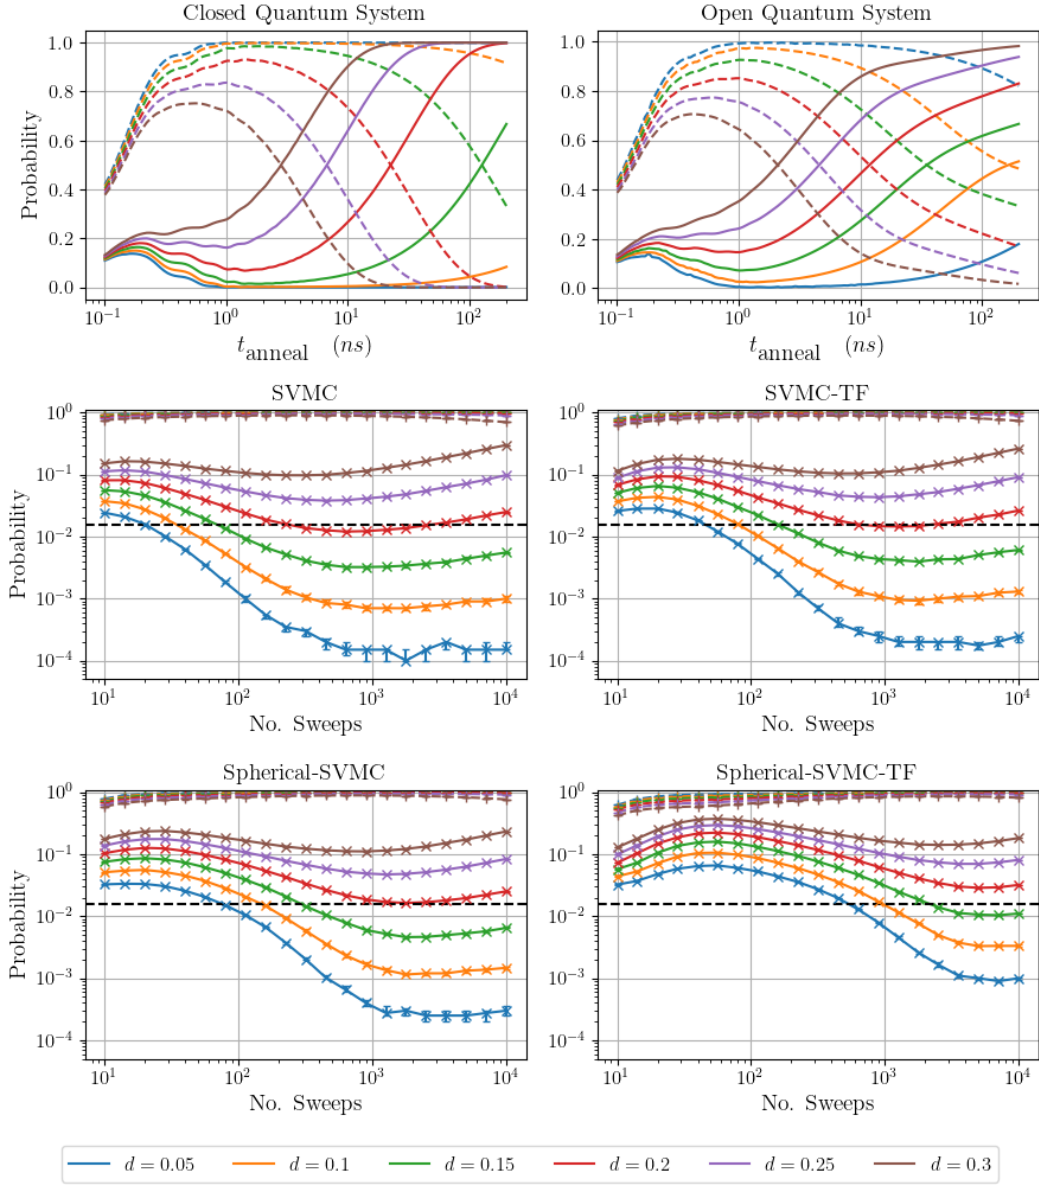


**Figure 3.7:** State probability at  $s = 0.83$  ( $= t/t_{\text{anneal}}$ ), just after passing through the minimum gap ( $s = 0.8227$ ) for (a) the AME with  $t_{\text{anneal}} = 200ns$  and (b) spherical SVMC-TF with  $t_{\text{anneal}} = 10,000$  sweeps. We measure a PFC with  $M = 3$  and  $d = 0.1$ . For the AME we measure the probability of being in the  $j$ th instantaneous state,  $E_j$ , and for spherical SVMC-TF we take a classical measurement of being in either the ground state or any of the first excited states. The spherical-SVMC-TF probabilities are found from 20,000 samples, of which we repeat 50 times and bootstrap to find a median and 95% confidence intervals for the data point and error bars, respectively.


dynamics describe quantum evolution at a temperature of 0 K with no dephasing, and therefore requires a run time of  $t_{\text{anneal}} = O(1/\Delta_{10}^2)$  to run adiabatically. For example, at  $d = 0.05$  the approximate adiabatic run time is  $\sim 163\mu s$  and we therefore see a probability  $\sim 0$  due to the short run times.

However, for the open-system simulations, we see a nonzero ground-state probability at  $d = 0.05$  for  $t_{\text{anneal}} = 200ns$ . This means that the addition of a thermal bath to the system is aiding computation, and since this transition passes through a potential barrier, it can be attributed to thermally-assisted tunnelling governed by the relaxation rate (Equation (3.20)) being nonzero about the minimum gap





**Figure 3.8:** Plots of state probability for being in either the ground state (solid lines) or any of the  $2^M$ -degenerate first excited states (dashed lines) at the end of an anneal. A PFC with system size  $M = 3$  was evolved using quantum and classical dynamics. The horizontal black dashed line indicates the probability with random guessing, i.e.,  $1/64$ . The quantum simulations are plotted against anneal time in nanoseconds, while the SVMC simulations are plotted against the number of sweeps. The closed- and open-system dynamics are evolved according to the von Neumann and adiabatic master equation, respectively. The probabilities for all SVMC variants are found from 20,000 samples, of which we repeat 50 times and bootstrap to find a median and 95% confidence intervals for the data point and error bars, respectively.



**Figure 3.9:** Plot of the evolution of ground-state probability from the AME (black solid) and the Gibbs state (black dotted), as well as the transition rate (red) (Equation 3.20). A PFC of  $M = 3$ ,  $R = 1.0$ , and  $d = 0.05$  was simulated at a system temperature of 12 mK. The AME was evolved for  $t_{\text{anneal}} = 200\text{ns}$  and the minimum gap occurs at  $s = 0.9059$ .

(Figure 3.9). This can occur because the system is subthermal [118] (i.e., the ground-state probability is less than that at thermal equilibrium) immediately after the minimum gap, and the energy gap is still small enough to allow significant thermalization from the first excited state to the ground state. Additionally, this transition involves all backbone qubits changing their magnetization simultaneously.

Therefore, thermalization in a quantum system is seen to be of some computational use, as has been seen in other literature [37,57,118,119,203]. However, we see that on the time scales tested that classical thermalization mechanisms represented in SVMC play a far less significant role, and results in a ground-state probability (after 10,000 sweeps) two orders of magnitude lower than a 100-ns AME evolved anneal. The marked difference between the two types of dynamical simulations highlights the effect of the minimum gap and the exponential manifold on ground-state probability, making it a gadget of interest for when differentiating between quantum and classical evolutions in the TFIM. For a better contrast between dynamics, annealing larger versions of the PFC ( $M > 10$ ) would result in negligible statistical contribution

from random state sampling, and an extremely large first-excited-state manifold that would likely result in SVMC failing to find the ground state.

In this section, we have introduced the perturbed ferromagnetic chain (PFC) [117], a gadget with an exponentially large first-excited-state manifold and an isolated ground state, whereby problem hardness and frustration is tuned by the perturbative parameter,  $d$ . When annealed in the transverse field Ising model (TFIM), the PFC develops computationally hard characteristics such as an exponentially small minimum gap (in  $N$ ), a quantum phase transition, and a false minimum. The evolution of the PFC in the TFIM was assessed with quantum dynamics using the adiabatic master equation (AME), and classically using both spin-vector Monte Carlo (SVMC) and its variants (see section 3.1.3 for more information). For quasi-continuous evolution of the PFC with the SVMC methods, the false minimum is followed to the computational first-excited-state manifold. This results in probable transitions to other low energy states further in Hamming distance away from the computational ground state (Figure 3.7), and therefore reduces the probability of reaching the ground state significantly. This is compounded by increasing problem size (Figure 3.6(a)) and by tuning  $d$  (Figure 3.8).

For a PFC evolved using the AME, the system mostly remains in the lowest two eigenstates (Figure 3.7) instead of accessing the exponential manifold, such that a  $100ns$  open-quantum-system anneal results in a ground state probability two-orders of magnitude larger than a 10,000 sweep spherical SVMC-TF anneal for the hardest comparative problem simulated (Figure 3.8). The AME evolution permits thermalization to the ground state (Figure 3.9) in time-scales too short for adiabatic evolution, indicating that thermalization aids computation of the PFC in the TFIM. The distinct differences seen between classical and quantum evolutions therefore makes the PFC a useful gadget in differentiating dynamical systems, something which other gadgets cannot always exhibit [60, 89, 119]. Comparative tests at larger PFC system sizes of  $M > 10$  would also allow for better differentiation and insight into the computational extent of thermalization in quantum and classical evolutions.

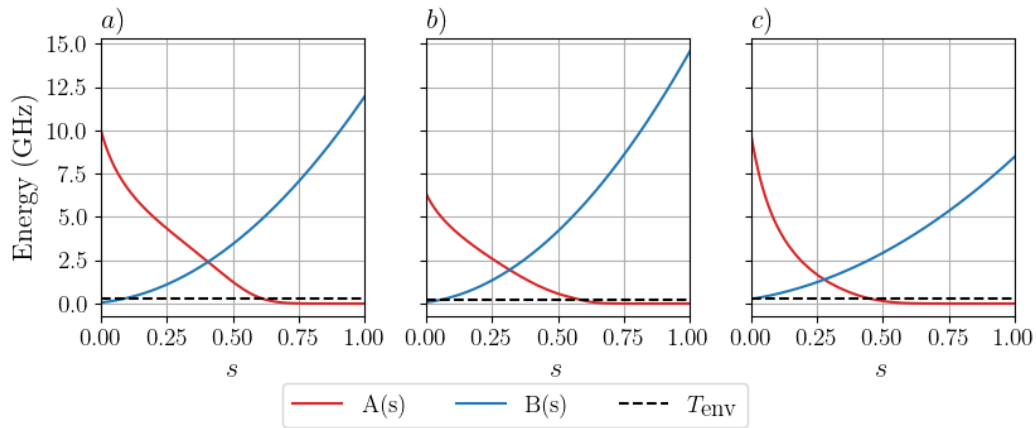
### 3.3 Experimental Thermalization

In the previous section, it was demonstrated with simulations how the perturbed ferromagnetic chain (PFC) [117] is able to distinguish between evolutions evolved using thermally-assisted quantum dynamics, and those that are classical or purely quantum by means of ground state probability measured. It is known that D-Wave quantum annealers use thermalization dynamics to aid computation in some instances [37, 57, 119, 203], but whether these dynamics are classical or quantum in nature is still a debated question. Therefore, the PFC can be used to shed light on this question, as well as whether incoherent quantum dynamics can be used for computation. In this section, we compare the simulated models to the experimental results obtained using various D-Wave quantum annealers, which yielded insights into the questions posed and allowed for the characterization of quantum dynamics on the various D-Wave annealers. The Hamiltonian for all annealers tested is

$$H(s) = -\frac{1}{2}A(s) \sum_{i=1}^N \sigma_i^x + \frac{1}{2}B(s) \left[ \sum_{i=1}^N h_i \sigma_i^z + \sum_{i,j} J_{ij} \sigma_i^z \sigma_j^z \right] \quad (3.21)$$

where the parameter schedules  $A(s)$  and  $B(s)$  are defined in Figure 3.10 for a D-Wave 2000Q housed at Los Alamos National Laboratory (LANL), D-Wave Advantage\_system4.1 hosted by D-Wave over cloud services, and a low-noise D-Wave 2000Q (named DW\_2000Q\_6) also hosted by D-Wave over cloud services. The simple connectivity of the PFC means that no minor-embedding is required and no parameter rescaling is required when placing the PFC on the quantum processing unit (QPU).

In Section 3.3.1, the  $M = 3$  PFC with  $d = 0.05$  is run on the D-Wave Advantage annealer and compared to the simulation models used in Section 3.2, using the shorter minimum anneal time of  $500ns$  available on the D-Wave Advantage QPU compared to the minimum of  $1\mu s$  on 2000Q annealers. Larger system sizes beyond what is possible to simulate using the adiabatic master equation are then explored and compared to classical SVMC in Section 3.3.2, where the effects of thermalization are more distinct. For larger system sizes, the condition of  $d > \frac{1}{M-1}$  required for the




**Figure 3.10:** Plots of D-Wave annealing coefficient schedules for the *a*) DW\_2000Q\_6, *b*) Los Alamos National Laboratory DW\_2000Q, and *c*) Advantage\_system4.1 systems, with operating temperatures  $T_{\text{env}}$  of  $13.5 \pm 1.0$  mK ( $0.279 \pm 0.021$  GHz),  $12.3$  mK ( $0.256$  GHz, no error provided) and  $15.4 \pm 0.1$  mK ( $0.321 \pm 0.002$  GHz), respectively. The schedules parameterize Equation 3.21 as a function of the annealing parameter  $s$ .

exponentially-degenerate manifold state,  $\mathcal{M}$ , to be the first-excited state is broken in order to preserve the PFC properties in case of integrated control errors (i.e.,  $d > 0$  at all times). An example is shown in Figure 3.11 for  $d \not\gg \frac{1}{M-1}$ , whereby despite this condition being broken the avoided-level crossing is still between the  $\mathcal{M}$ -state and the ground state, with the other energy state not part of the exponential manifold (called the intermediate-state) crossing the  $\mathcal{M}$ -state such that the intermediate  $I$ -state becomes the new first-excited state. Here there are 4 degenerate intermediates states caused by a bit-flip on a terminal backbone qubit, that can happen on either end of the PFC, and whereby the auxiliary qubit is isoenergetic to its spin-state.

### 3.3.1 Model Comparisons

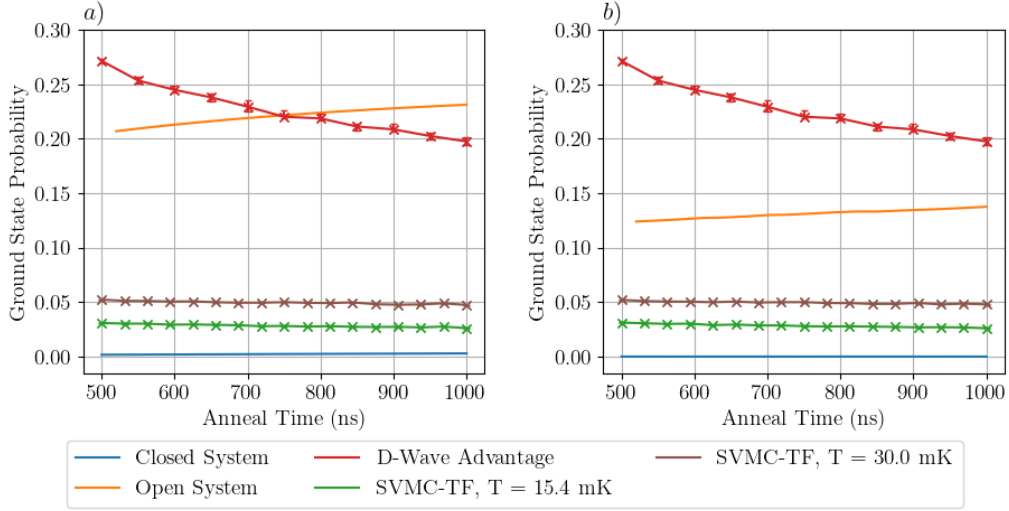
For tractability of simulations in the model comparison, the shorter  $500\text{ns}$  minimum anneal time Advantage annealer, Advantage\_system4.1, is used in this section with a system temperature of  $15.4 \pm 0.1$  mK. Experiments are compared against SVMCTF, closed quantum system, and AME simulations. The roles of two prominent QPU noise sources are also explored in the comparison, the first being cross-talk, which is the leakage of magnetic fields to neighbouring qubits characterised by the magnetic susceptibility of qubits  $\chi$ , and the second being parameterization noise, where the



**Figure 3.11:** The example energy-gap spectrum for an  $M = 6$  perturbed ferromagnetic chain with  $d = 0.25$ , where the  $d < \frac{1}{M-1}$  condition is violated. In this regime, the manifold states,  $\mathcal{M}$ , is no longer the first-excited state at the end of the anneal, and are replaced by the intermediate states,  $I$ . The system is annealed according to the Hamiltonian defined in Equation 3.21 with the DW\_2000Q\_6 coefficient schedule, and then rescaled such that at  $s = 1$  the energy difference is in the Ising energy units of the problem Hamiltonian.

defined bias and coupler values are the means of a normal distribution with some standard deviation  $\sigma$ . The latter error is only applied to the Monte-Carlo simulations, and values of  $\chi = -0.035$  and  $\sigma = \sqrt{0.03}$  are used throughout, in accordance with values stated in QPU datasheets. Experimentally, these noise sources are mitigated using methods such as spin-reversal transforms, longer intervals between samples to reduce inter-sample correlation, and random direct embedding to average over the error across the entire QPU.

Comparisons of simulation models against the Advantage\_system4.1 annealer is shown in Figure 3.12 for models with and without QPU noise emulation. Ground-state probability was adversely impacted in both open and closed quantum simulations by the addition of cross-talk, whereas there was negligible impact on the SVMC-TF simulations. Furthermore, no direct match between the experimental data and models tested is observed, and for short anneal times an unusually high ground state probability is observed, which then decreases as anneal-time increases. Quantum models were only observed to have decreasing ground-state probability



**Figure 3.12:** Comparison between the simulated classical and quantum models to the D-Wave Advantage<sub>system4.1</sub> annealer for the short anneal time regime between  $500\text{ns}$  –  $1\mu\text{s}$  for an  $M = 3$ ,  $d = 0.05$  PFC. Plot *a*) illustrates simulated models with no integrated control errors in the anneal, whilst *b*) introduces cross-talk (see Equation 1.13) into the Hamiltonian with  $\chi = -0.035$  for all simulated models, and parameter setting errors for both biases and coupling (see Equation 1.13) with a standard deviation of 0.03 for SVMC-TF simulations. The probabilities for the D-Wave Advantage and SVMC-TF was obtained from the median of 50 sets of 10,000 samples that were bootstrapped to yield error bars to two standard deviations. The number of Monte-Carlo steps per  $\mu\text{s}$  for SVMC-TF were chosen to be  $1000\mu\text{s}^{-1}$ . Open quantum system simulations assumed an Ohmic bath with  $\eta g^2 = 1 \times 10^{-3}$ ,  $\omega_c = 4.0\text{ GHz}$  at temperature  $T = 15.4\text{ mK}$ .

in anneal time after the “diabatic bump”, observed in Figure 3.8, however the QPU is annealed at a timescale where this phenomenon has already passed. Given that we are also not at a timescale to observe adiabaticity, the decreasing ground-state probability observed for the QPU can be attributed to mechanisms similar to the classical SVMC due to the consistency in the decreasing ground-state probability justified in Section 3.2.2.

SVMC-TF ground-state probability increases when simulated at a temperature higher than the defined QPU temperature of 15.4 mK, and the largest probabilities observed occur between  $T = 30 - 50\text{ mK}$ , however, all SVMC-TF ground-state probabilities measured do not match the experimental observations. Other causes that result in the larger than expected experimental probabilities such as initialising in high temperature thermal state at  $t_{\text{anneal}} = 0$  are also improbable due to the large separation between energy states at the beginning of the anneal. Therefore, it is likely that the

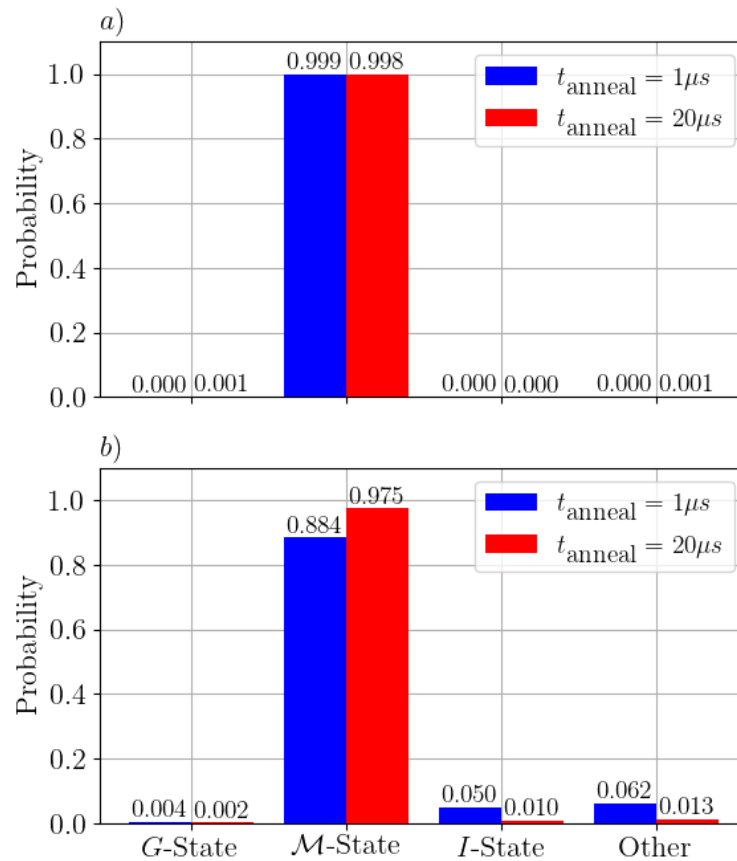
noise-model and/or system-bath coupling for the QPU are beyond the Ohmic bath in the weak-coupling limit approximations, and closer to that described by the polaron-transformed Redfield equation [144, 145] which tends to a classical description in the infinite coupling strength limit. Therefore, in this setting it is not possible to determine whether incoherent quantum processes are driving computation, and we extend to larger system sizes to exponentially increase the number of manifold states, which is known to hinder a system following a semiclassical potential (Section 3.2.2) to better determine whether incoherent quantum processes are present.

### 3.3.2 Thermalization with Large Systems

The indication that the Advantage\_system4.1 annealer evolution is closer to simulated classical evolutions meant that thermalization experiments with large PFC systems were conducted with the low-noise DW\_2000Q\_6 annealer to mitigate non-trivial effects beyond the adiabatic master equation. To illustrate the key differences between the two annealers, a forward anneals of  $1\mu s$  and  $20\mu s$  were conducted using a PFC of  $M = 15$  and  $d = 0.1$ . The ground( $G$ )-state, manifold( $\mathcal{M}$ )-state, intermediate( $I$ )-state, and other excited state probabilities were measured in Figure 3.13 for both annealers, with the DW\_2000Q\_6 placing the majority of probability density on the  $\mathcal{M}$ -states as expected by both the classical and quantum models. Additionally, the DW\_2000Q\_6 is also seen to thermalise to the ground state for longer anneal times, consistent with the thermalization models over long anneals. In contrast, the Advantage\_system4.1 annealer is seen to place probability density on a variety of excited states away from the  $\mathcal{M}$ -states for short anneal times, and begins to correct itself for the longer anneal time of  $20\mu s$ . The disparity in profiles demonstrates how it is possible to characterise noise in annealers with the PFC.


Progressing with the DW\_2000Q\_6 annealer to demonstrate thermalization in large PFC systems, we explore the effect on ground state probability for annealing schedules with a pause, which is known to show peaks in ground-state probability if the schedule pauses at the correct combination of transverse and longitudinal fields [57]. Figure 3.14 illustrates the pause schedule we use for the  $M = 15$ ,  $d = 0.1$  PFC as a function of normalised annealing time  $s$ , which parametrises the





**Figure 3.13:** State probabilities after a forward anneal on an  $M = 15$ ,  $d = 0.1$  PFC using a) the DW\_2000Q\_6 and b) the Advantage\_system4.1 D-Wave annealing systems, where  $G$ -state,  $M$ -state,  $I$ -state and other refer to the ground, manifold, intermediate and higher excited states, respectively. Probabilities were collected from 10,000 samples that underwent 100 spin-reversal gauge transformations.


Hamiltonian coefficients defined in Figure 3.10. The schedule includes a pause time of  $t_p = 19\mu s$  at a specific value of  $s = s_p$ , and keeps the time when the fields are non-static at  $1\mu s$ , such that the total anneal time is always equal to  $20\mu s$  and the rate which we evolve dynamics is kept constant throughout to emphasise the role that the pause has on system thermalization. Instead of sweeping the temperature of the system, which is fixed in the experimental setup, the classical energy-levels of the PFC are scaled by  $R$ . For  $R < 0.047$  is where the energy scale of the PFC biases and couplers start to be less than the energy of the system temperature in the DW\_2000Q\_6 annealer, and where we can expect thermal noise to disrupt the system completely.



**Figure 3.14:** Illustration of the forward anneal with a pause for time  $t_p$  at a constant value of  $s = s_p$ . In this example, the dynamic periods  $t_1$  and  $t_2$  are a total of  $1\mu s$  in duration, with a pause of  $t_p = 19\mu s$ .

Figure 3.15 directly compares ground-state probability measured at various values of  $R$  and pause position  $s_p$  for both the DW\_2000Q\_6 annealer and the SVMC-TF model, which is the classical model known to emulate quantum annealing thermalization signatures in Ref. [119]. The PFC at this system size demonstrates that the classical model cannot access the ground-state using classical thermalization mechanisms in Figure 3.15b, whilst the quantum annealer displays a clear thermalization signature. The effect of spin-bath polarization potentially contributing to the difference in profiles has also been mitigated with random placement of the PFC on the QPU, and spin-reversal transforms. Additionally, the annealing timescale of  $20\mu s$  is shorter than the  $1ms$  timescale that the spin-bath polarization effect has been previously reported to manifest at [119], and is therefore not considered to significantly contribute to the measured ground-state probability.


Given that in Figure 3.13 the system is primarily in the exponentially large manifold in both the quantum and classical systems, pausing in the quantum system demonstrates that the mechanism used to access the ground-state is distinctly different from what can be classically modelled using SVMC-TF. Therefore, thermalization in the quantum annealer is distinctly different to what is described by the classical model. However, to thermalize from the  $\mathcal{M}$ -state to the ground-state in this



**Figure 3.15:** Ground state probability after a forward anneal following the schedule defined in Figure 3.14 using *a)* the DW\_2000Q\_6 and *b)* SVMC-TF for various energy scales  $R$  and pause positions  $s_p$ . A PFC with  $M = 15$  and  $d = 0.1$  was used for these experiments, which was randomly embedded onto the QPU upon for each  $20\mu s$  run where 10,00 samples were collected with 100 spin-reversal gauge transforms. Probabilities obtained for both plots are from the median of 10 sets of runs, and the QPU system temperature and coefficient schedules were used in the simulations (see Figure 3.10). The number of SVMC-TF Monte-Carlo steps per  $\mu s$  for this comparison with the QPU were chosen to be at  $1000 \mu s^{-1}$ .

problem requires the system to tunnel a Hamming distance  $\geq 15$ , which involves all 15 qubits in the backbone that are negatively magnetized (see Figure 3.2) to become positively magnetized simultaneously. The quantum thermal transition rate according to Fermi's golden rule in Equation 3.20 will tend to zero as Hamming distance between states increase due to the term summing the vanishing off-diagonal elements tending to zero in this case, therefore making a thermal transition highly improbable.

To provide an explanation for this transition, we consider the intermediate states that occur between the  $\mathcal{M}$ -state and the ground-state, which for  $N = 15$  there are 5 distinct energy eigenvalues corresponding to degenerate  $I$ -states. The  $I$ -states as a classical Ising spin-state represent domain-walls along the backbone of the PFC, and are closer in Hamming distance and energy to the  $\mathcal{M}$ -states compared to the ground state. Given that domain-walls commonly occur in ferromagnetic systems [238, 239] on quantum annealers (e.g., minor-embedding), it is hypothesised that the  $I$ -states provide an alternative and potentially lower action thermalization route through the Hilbert space to the ground state, compared to a single multi-qubit thermal transition that requires a large change in Hamming weight. Figure 3.16 illustrates that a

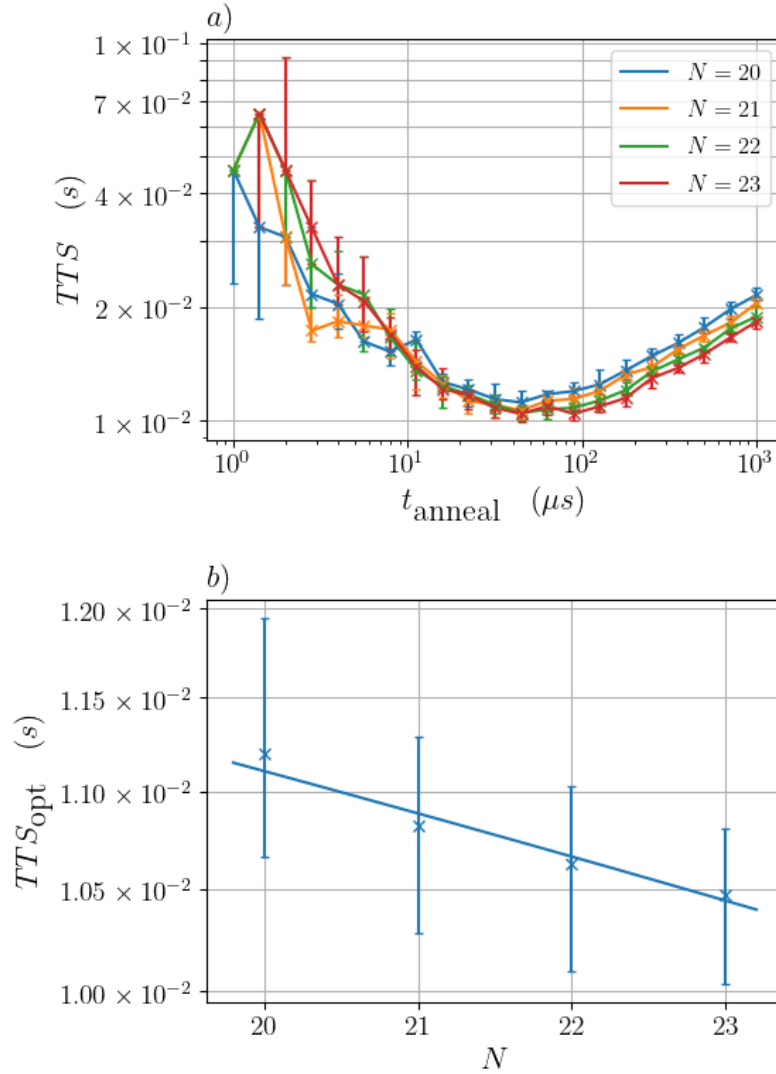


**Figure 3.16:** Diagram of a possible intermediate state configuration that could provide a route to the ground state configuration, where  $s^a$  and  $s^b$  are the auxiliary and backbone spins, respectively. Starting from the first-excited state closest in Hamming distance to the ground state, a spin flip on a terminal backbone spin, such as  $s_M^b$ , can result in spins further down the backbone also flipping (i.e., to other lower energy intermediate states) to reach the ground state configuration of all spins equal to 1.

terminal backbone qubit  $s^b$  that has undergone a spin flip can induce subsequent backbone spins to also flip to the all spin-up configuration (i.e., the ground state). Experimental observations in Figure 3.13a have shown that few to no  $I$ -states were measured despite being energetically lower than the manifold states, and in a system that experiences thermalization, we should expect non-negligible probability in these states.

To test whether the intermediate states are playing a role in accessing the ground-state, PFC's of increasing size and therefore an increasing number of intermediate states between the exponential manifold and ground state were evaluated on the Los Alamos National Laboratory D-Wave 2000Q quantum annealer. Given that sufficient time is required for the PFC to access the ground-state via thermalization, the time-to-solution (TTS) metric is used to measure performance of the PFC over various anneal times and find an optimum TTS. The time-to-solution is defined as

$$TTS = t_{\text{anneal}} \max \left( 1, \frac{\ln(1 - P_d)}{\ln(1 - P_G)} \right) \quad (3.22)$$



**Figure 3.17:** Plots of *a*) time-to-solution ( $TTS$ , Equation 3.22) for a  $d = 0.19$  PFC for various anneal times and system sizes,  $M$ , and *b*) the optimum time-to-solution found for each system size extracted using univariate spline interpolation from the PYTHON package *SciPy*. These results were obtained using the Los Alamos National Laboratory D-Wave 2000Q quantum annealer with 100 spin-reversal transforms, reduced inter-sample correlation, and a readout thermalization time of  $20\mu s$ . Probabilities were obtained from the median of 50 sets of 10,000 samples that were bootstrapped to yield the 95% confidence intervals. Space-division multiplexing was used to reduce the number of runs required to reach 10,000 samples, from which a different random QPU embedding was used on each iteration.

where  $P_d$  is the desired probability threshold that is set to  $P_d = 0.99$ , and  $P_G$  is the probability of measuring the ground state. For sufficiently small PFC systems, the optimum TTS is at a time shorter than  $1\mu s$  and is inaccessible on the QPU, therefore, large PFC's are required to find an optimum TTS. The TTS and corresponding optima are plotted for PFC's of size  $M \geq 20$  in Figure 3.17, where for an increasing system size a constant/decreasing optimal TTS is observed. Furthermore, for anneal times  $> 20\mu s$ , the larger PFC's have a higher probability of measuring the ground state than smaller PFC's, as indicated by the lower TTS in Figure 3.17a. Given the literature reviewed, this is the first problem found with an exponentially small avoid-level crossing gap whose experimentally derived optimal TTS is observed to not scale exponentially as system size increases.

The explanation for this scaling supports the hypothesis that the intermediate states are playing a significant role in the thermalization process, as despite the closing gap with increasing PFC system size, the number of intermediate states increase between the ground and  $\mathcal{M}$ -state as the manifold states are pushed to higher energies (gap size  $\Delta_{\mathcal{M}0}$  between the manifold and ground state is  $\Delta_{\mathcal{M}0} = 2RMd$ ). For all system sizes measured, the  $I$ -state closest in Hamming distance to the  $\mathcal{M}$ -states is always a distance of 6 away throughout. Provided that the  $\mathcal{M}$ -states can access the closest  $I$ -state via thermalization, this provides the pathway to the ground state as the Hamming distance interval between intermediate-states is 1. Consequently, the time-scales that were annealed to are long enough to also incur spin-bath polarization, potentially contributing to the larger than expected ground-state probabilities. However, this phenomenon has not been observed to be significant enough to cause non-exponential scaling in system size, but does likely contribute to the appearance of the small negative scaling in Figure 3.17b, instead of a constant scaling if there is a constant thermalization rate between the  $\mathcal{M}$ -states and  $I$ -state. Further confirmation of this hypothesis for large PFC systems would require analysis of quenched anneals, in an attempt to observe a distribution of intermediate states being accessed progressively to reach the ground state. However, D-Wave QPUs do not allow for accurate system quenching, as the rate at which fields can

be changed in time is two orders of magnitude lower than is required for accurate quenching.

Overall, in this section we have compared different models of thermalization in short-time scales to experimental measurements from the D-Wave Advantage\_system4.1 annealer, which demonstrated that the QPU possesses a thermalization model that is beyond the weak-coupling approximation of an open quantum system and the classical approximation of the annealer. The PFC was then used to experimentally confirm this difference by demonstrating that the D-Wave Advantage\_system4.1 is a noisier system compared to the known lower-noise DW\_2000Q\_6 annealer due to the distribution of states away from the manifold states for short anneal times. Using the low-noise annealers, it was demonstrated that the thermalization mechanism on the QPU is distinctly different from that described solely by classical dynamics, due to the clear thermalization signature present in Figure 3.15 for the annealer. However, the Hamming distance from the manifold-states to the ground-state in the thermal transition is 15, which is highly improbable given a thermal transition governed by Fermi's golden rule.

A hypothesis that intermediate-states are being used to reach the ground-state was therefore proposed to provide an explanation as to how the system can make a large Hamming distance transition. An experiment that demonstrated how increasing PFC system size, and therefore Hamming distance of the effective thermal transition, resulted in flat scaling of the optimal TTS, which has not been observed yet in literature for problems with exponentially scaling gaps such as the PFC. To confirm the presence of intermediate states in the thermal transition, quenching experiments are proposed in order to observe a distribution of intermediate states being used in the transition. Quenching with D-Wave QPUs cannot happen on a timescale that freezes the dynamics fast enough to take a snapshot of the state distribution during the anneal, and therefore more specialist hardware is required to further confirm this thermalization mechanism.

## Chapter 4

---

# Diabatic Quantum Annealing


---

The objective of most adiabatic algorithms is to finish in the ground state of some objective Hamiltonian at the end of the anneal, but exponential closing of avoided-level crossings in system size for some problems make the ground-state effectively inaccessible for large problems, even when evolved in a zero-temperature quantum system. In an attempt to generalize, algorithms have been designed to use avoided-crossings to temporarily break the adiabatic condition and use excited states mid-anneal before then returning to the ground state at the end of the anneal [31, 70, 95].

Given that the system is prepared to be the ground state of some initial Hamiltonian, the protocol that purposely uses excited states during an anneal is referred to as a diabatic quantum annealing (DQA). To have a successful DQA algorithm, there must be a path using excited energy states that returns to the ground state at the end of the anneal. The simplest example of a path is when there are two avoided-crossings between the ground and first-excited state, which are both separated from higher excited states. If the scaling of the avoided-crossings and higher excited state separation is exponential and polynomial, respectively, then this provides a quantum speed-up (see Figure 4.1), but this has not yet been generalized and is only applicable to specific problems. The hardware requirements to run many DQA algorithms are also beyond what is currently technologically feasible, as they make use of  $XX$  couplers or other specialized Hamiltonians [31, 95].

More recently, an algorithm proposed by Fry-Bouriaux *et al.* [39], named the





**Figure 4.1:** Example of a diabatic annealing spectrum taken from Somma *et al.* [31]. The spectrum plotted is instantaneous state energy as a function of normalized annealing time  $s$ , whereby green points  $1 \rightarrow 2$  and  $3 \rightarrow 4$  indicate diabatic transitions at exponentially scaling avoided-crossings, whilst all other spectral gaps are polynomially scaling. It illustrates how the first excited state is used temporarily during an anneal (from  $s = 0$  to  $s = 1$ ) to reach the ground state in a way that provides a provable quantum speed-up over a purely adiabatic approach and classical methods such as quantum Monte Carlo.

locally suppressed transverse-field diabatic quantum annealing protocol (LSTF-DQA), manages to create the spectra necessary for successful DQA within current hardware constraints. In this algorithm, a single qubit in the system has its transverse-field component suppressed to freeze its local dynamics and guide the rest of the qubits towards the optimal solution. This method uses the unitary dynamics of a quantum system to guide the system to the final ground state solution. However, the protocol does not guarantee that every problem will form the correct spectra for DQA, and relies on the problem having inherent magnetic frustration due to local biases.

Given the potential promise of DQA being a solution to the problem of exponentially scaling avoided-crossings in quantum annealing, the rest of this chapter details research conducted to address open questions:

1. What problems can we use locally suppressed transverse-field diabatic quantum annealing for?
2. Is there a way of identifying a priori which qubit has to have a suppressed transverse-field that ensures the necessary double avoided crossing for locally suppressed transverse-field diabatic quantum annealing?

3. Can classical spin-vector dynamics emulate locally suppressed transverse-field diabatic quantum annealing?
4. Can DQA be adapted to become an interferometer and therefore a test of coherence?

Section 4.1 addresses the first three questions within the context of combinatorial optimisation, focusing on specific problem classes such as maximum independent set (Section 4.1.1) and max-cut (Section 4.1.2) problems. We demonstrate how maximum independent set problems can benefit from the LSTF-DQA algorithm due to the inherent frustration present in the problem Hamiltonians, and how the unitary quantum evolution provides the necessary dynamics to reach the ground state efficiently compared to classical evolutions. A framework is also presented that introduces local frustration to Hamiltonians without any local fields, such as max-cut, whilst preserving the ground-state. Finally, in Section 4.2, the LSTF-DQA algorithm is repurposed into an interferometer by interpreting the avoided-crossings as beam splitters and conducting Mach-Zehnder style interferometry experiments [291, 295, 296]. The perturbed ferromagnetic chain defined in Chapter 3 is used to provide examples of how interferometry can be conducted and how it is sensitive to the level of decoherence experienced by the system.

## 4.1 Diabatic Quantum Annealing for Optimization

Locally suppressed transverse-field diabatic quantum annealing (LSTF-DQA) is a near-term implementation of DQA algorithms on quantum annealing hardware. The current hardware that is commercially available commonly uses the transverse-field Ising model Hamiltonian defined by

$$\hat{H}(s) = -A(s) \sum_i^N h_i^x \sigma_i^x + B(s) \left[ \sum_i^N h_i^z \sigma_i^z + \sum_{ij} J_{ij}^z \sigma_i^z \sigma_j^z \right], \quad (4.1)$$

where the first term represents the transverse-field components, and the second term contains the longitudinal fields and couplers that encode the optimization problem of interest. Schedules  $A(s) = 1 - s$  and  $B(s) = s$  in normalized time  $s = t/t_{\text{anneal}}$  are


used consistently throughout the rest of the chapter. For a more detailed overview of this Hamiltonian, readers are referred to Section 1.2.

The LSTF-DQA protocol adapts Equation 4.1 by choosing a qubit in the Hamiltonian, which we refer to as the target qubit, and suppressing its transverse-field component (i.e., the local suppression of the transverse-field). An independent longitudinal field schedule is also applied to the target qubit to initialize the spin state of the qubit, which prepares the qubit in either a spin-up or spin-down state. The LSTF-DQA Hamiltonian is defined as

$$\begin{aligned} \hat{H}(s) = & -a(s) \sum_{i \neq k}^N \sigma_i^x + b(s) \left[ \sum_{i \neq k}^N h_i^z \sigma_i^z(s) + \sum_{ij} J_{ij}^z(s) \sigma_i^z \sigma_j^z \right] \\ & - c_x(s) \sigma_k^x + c_z(s) h_k^z \sigma_k^z, \end{aligned} \quad (4.2)$$

where the target qubit  $k$  has a longitudinal and transverse field schedule independent of the rest of the system. The spectrum necessary for DQA is created using this Hamiltonian if the problem possesses an avoided-level crossing induced by localised frustration, which can be relieved if the transverse-field of the correct choice of target qubit is suppressed [39]. Suppressing the transverse-field component and choosing the correct longitudinal-field bias on the target qubit guides the system to the ground-state solution. This is analogous to providing the system with part of the solution, which then enables it to find the ground-state. This independent tuning of fields required for LSTF-DQA is not made available on quantum annealing systems such as D-Wave, and therefore means the work presented here is entirely based on numerical simulations and analytical results.

The schedules presented in equation 4.2 are composed of two stages, an initialization phase and an annealing phase, which end and begin at  $s = s_x$ , respectively, for  $s \in [0, 1]$ . The initialization phase is where the ground-state is slowly made unfavourable, whilst the rest of the system remains stationary with constant fields, to force probability density into the excited state at the end of the initialization phase. The annealing phase is the period during which the fields of the entire system are evolved from  $\hat{H}(s_x)$  to the problem Hamiltonian,  $\hat{H}(s = 1)$  (Equation 4.2). The



**Figure 4.2:** Plots of *a)* the energy spectrum and *b)* the schedules respective annealing schedules for both adiabatic quantum annealing (Equation 4.1) and LSTF-DQA (Equation 4.2) for a  $M = 1$  perturbed ferromagnetic chain (presented in Chapter 3) with  $d = 0.3$ . The DQA schedules presented in *b)* are defined in Equations 4.3-4.6 and are parameterized with  $s_x = 0.2$ ,  $c_0 = c_1 = 0$ , and  $b = 1$ .

schedules in equation 4.2 are plotted in Figure 4.2b and are defined as

$$a(s) = \min \left( 1, 1 - \frac{s - s_x}{1 - s_x} \right) \quad (4.3)$$

$$b(s) = \max \left( 0, \frac{s - s_x}{1 - s_x} \right) \quad (4.4)$$

$$c_x(s) = \begin{cases} \frac{c_0 s}{s_x} & \text{if } s < s_x \\ c_0 + (c_1 - c_0) \frac{s - s_x}{1 - s_x} & \text{otherwise} \end{cases} \quad (4.5)$$

$$c_z(s) = \begin{cases} b \frac{s - s_x}{1 - s_x} & \text{if } s < s_x \\ \frac{s - s_x}{1 - s_x} & \text{otherwise} \end{cases} \quad (4.6)$$

The magnitudes of transverse field on the target qubit are set by values  $c_0$  and  $c_1$ , which are the initial and final values of transverse-field on the target qubit respectively. Please note that both Equation 4.5 and Equation 4.6 are different to what is defined in Ref. [39] in order to make  $c_x(s = 0) = 0$  and introduce a custom choice of initial spin-state ( $b \in \{-1, 1\}$  for spin-down and spin-up initial state preparation, respectively) on the target qubit. The double-crossing between the ground and first excited states

required for LSTF-DQA are achieved when  $c_0 = c_1 = 0$ , which corresponds to zero transverse field on the target qubit. This is illustrated in Figure 4.2 alongside the adiabatic equivalent of the problem. If we were to use a non-zero transverse field on the target qubit and maintain  $c_0 \ll 1$  and  $c_1 \ll 1$ , then the double crossings we have from  $c_0 = c_1 = 0$  become double avoided crossings, whose gaps we can tune with  $c_0$  and  $c_1$  (see Section 4.2.1.2 Figure 4.9 for an example).


In this section, we build on the work presented in Ref. [39] by using LSTF-DQA to solve the maximal independent set (MIS) and max-cut combinatorial optimisation problems. For the MIS problem in Section 4.1.1, we look at two non-isomorphic three-regular graphs that have a highly degenerate ground state, where the degeneracy can be broken through random perturbations. The LSTF-DQA protocol is simulated using both unitary quantum and classical spin-vector dynamics, where the successful anneals are analysed to identify any key properties that a target qubit possesses that would allow for an a priori selection of the target qubit. In Section 4.1.2, a framework is introduced to address the issue that LSTF-DQA has with problems that possess no local-qubit biases, such as the max-cut problem. This creates local-frustration and the spectra necessary for LSTF-DQA for max-cut problems with uniform coupling weights.

### 4.1.1 Maximal Independent Set

The maximal independent set (MIS) problem is an NP-hard problem that aims to find the largest set of vertices on a graph, such that none of the vertices in the set are adjacent to one another. This problem class is particularly useful in combinatorial optimization because of the reduction of all satisfiability problems to the 3-SAT problem [283], which can then be reduced to the MIS problem [16]. The cost function for the MIS problem [116] is

$$Q(\mathbf{x}) = -A \sum_{i \in V} x_i + B \sum_{(i,j) \in E} x_i x_j, \quad (4.7)$$

where  $A < B$  such that it is never favourable to violate the constraint that no two adjacent vertices can have  $x_i = 1$  for  $x_i \in \{0, 1\}$ . This cost function is also equivalent



**Figure 4.3:** The graphs of the only two non-isomorphic 3-regular graphs with 6 vertices, with an example of a maximum independent set for each graph indicated by the red coloured vertices. The left and right graphs have two and four maximum independent sets, respectively.

to the set-packing problem, which asks for the largest number of subsets  $V_i$  which are all disjoint. Using a change of variables, Equation 4.7 can be put into an Ising spin formulation where  $s_i = \{-1, 1\}$  and  $x_i = \frac{s_i + 1}{2}$ . The cost function with spin variables is therefore,

$$Q(\mathbf{s}) = \sum_{i \in V} \left( -\frac{A}{2} + \frac{B|\mathcal{A}(i)|}{4} \right) s_i + \frac{B}{4} \sum_{(i,j) \in E} s_i s_j + \left( \frac{B|E| - 2A|V|}{4} \right), \quad (4.8)$$

where in an annealing setting the biases  $h_i$  and coupler  $J_{ij}$  values are equal to the coefficients of the  $s_i$  and  $s_i s_j$  terms, respectively.

LSTF-DQA lends itself well to solving MIS problems, as it satisfies the criteria of possessing local biases and local frustration caused by all couplers being anti-ferromagnetic and biases that are generally positive on every spin. However, making an a priori determination on the choice of target problem is still not exactly clear, therefore trial and error is the only appropriate method of determining a suitable candidate. To obtain insights into the characteristics of what is a good choice of target qubit (i.e., the qubit that will create the spectra necessary for LSTF-DQA), 3-regular graphs with 6 vertices are investigated, of which there are only 2 non-isomorphic graphs that are presented in Figure 4.3. These graphs are referred to as the bipartite and prism isomorphisms throughout. Furthermore, by setting  $A = 1$  and  $B = 2$  the

general MIS cost function for 3-regular graphs is,

$$Q(\mathbf{s}) = \sum_{i \in V} s_i + \frac{1}{2} \sum_{(i,j) \in E} s_i s_j \quad (4.9)$$

which results in degenerate ground-states for both 6-vertex MIS graph problems. For LSTF-DQA, uniform noise is added to perturb  $A$  to break the degeneracy and produce a single ground state whilst maintaining  $B > A$ , i.e.,  $A \in (1, 2)$ . With the correct target qubit chosen, LSTF-DQA will produce double crossings and finish in the ground state at the end of the anneal. For each isomorphism, 20 perturbations of the MIS problem were generated, each with a single ground-state solution (found using a brute-force search of all  $2^6$  possible solutions) and the valid spectra necessary for DQA. For all perturbations generated, it was found that the target qubits resulting in the correct DQA spectra were part of the maximal independent set, i.e., the ground-state solution for that perturbation. Therefore, given that the MIS for the prism and bipartite graphs are of size 2 and 3, respectively, all valid combinations of perturbation and target qubit were tested, totalling 100 DQA anneals for a given annealing time.

Whether quantum dynamics are required in LSTF-DQA anneals is first addressed by comparing quantum and classical simulations in a closed system setting. Quantum simulations are modelled using the von-Neumann equation

$$\dot{\hat{\rho}}(t) = -i [\hat{H}(t), \hat{\rho}(t)] , \quad (4.10)$$

using the Hamiltonian defined in Equation 4.2 with  $c_0 = c_1 = 0$ , and an initial state of all qubits in the  $|+\rangle$  state except the target qubit which is initialized in the spin-up state,  $|0\rangle$ . It is noted that the spin state of the target qubit is important in LSTF-DQA, but because the MIS nodes are identified by spin-up states in the solution, we know a posteriori to consistently initialise the target qubit to in the spin-up state. The classical spin-vector dynamics (SVD) simulations are derived using the Euler-Lagrange equations of motion for both a single angle ( $O(2)$ ) rotor and a full Bloch sphere double angle ( $O(3)$ ) rotor. Dynamics for SVD  $O(2)$  simulations are

kept in the spherical coordinate basis but restricted to the  $XZ$ -plane, i.e.,  $\phi = 0$ , and are described by a system of non-homogeneous equations

$$\frac{d\theta_i}{dt} = \omega_i \quad (4.11a)$$

$$\frac{d\omega_i}{dt} = -\frac{dV_i}{d\theta_i} = A(t) \cos \theta_i + B(t) \sin \theta_i \left[ h_i + \sum_{i,j:i \neq j} J_{ij} \cos \theta_j \right]. \quad (4.11b)$$

where  $V = \langle \theta | H | \theta \rangle$  plays the role of the semi-classical potential described by the Hamiltonian. The simulations for SVD  $O(3)$  dynamics were also derived in the spherical basis but changed to Cartesian basis for numerical stability and its equivalence to the Bloch equation, given by

$$\frac{d\vec{M}_i}{dt} = \vec{H}_i \times \vec{M}_i = \begin{bmatrix} -2B(t)E_i \sin \theta_i \sin \phi_i \\ 2A(t) \cos \theta_i + 2B(t)E_i \sin \theta_i \cos \phi_i \\ -2A(t) \sin \theta_i \sin \phi_i \end{bmatrix}, \quad (4.12)$$

where  $\times$  denotes the cross product,  $\vec{M}_i = (\sin \theta_i \cos \phi_i, \sin \theta_i \sin \phi_i, \cos \theta_i)$  and

$$\vec{H}_i = -2A(t)\hat{x} + 2B(t) \left[ h_i + \sum_{i,j} J_{ij} \vec{M}_j \cdot \hat{z} \right] \hat{z}. \quad (4.13)$$

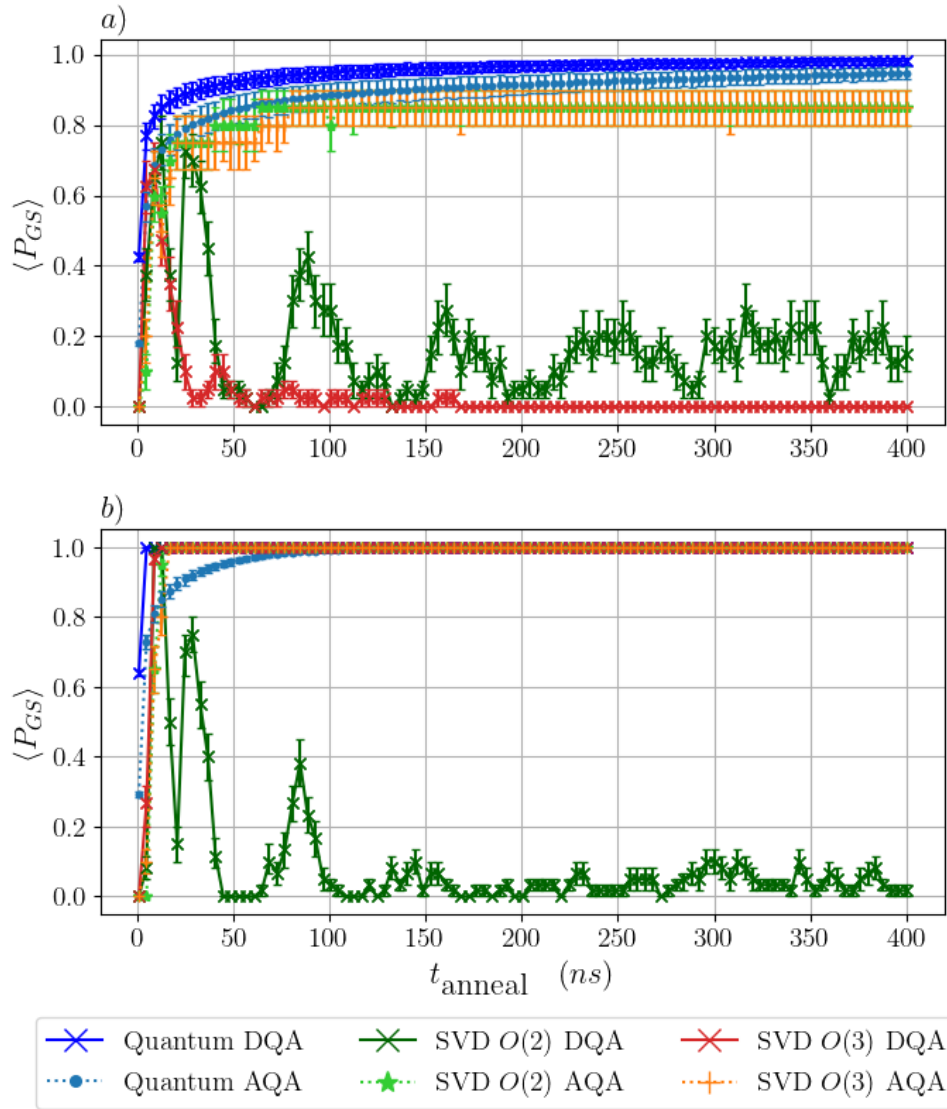
Further details on the derivation of these equations are found in Section 1.4.3.2. For each method, the expected ground-state probability is measured, which is defined as

$$\langle P_{GS}(t_{\text{anneal}}) \rangle = \frac{1}{n|S|} \sum_{i=1}^n \sum_{k=1}^{|S|} P_{GS}(t_{\text{anneal}}, i, k) \quad (4.14)$$

where  $n$  is the number of perturbed problems,  $|S|$  is the cardinality of the MIS,  $S$ , and  $k$  is the target qubit in the set  $S$ . This probability measure evaluates the success probability of the algorithm at a given time across all perturbations and target qubits.

A comparison between the quantum and classical dynamics for both LSTF-DQA and adiabatic quantum annealing (AQA) is shown in Figure 4.4, where for all anneal times evaluated the  $\langle P_{GS}(t_{\text{anneal}}) \rangle$  for the quantum DQA protocol is seen





**Figure 4.4:** Plots of the expected ground-state probability  $\langle P_{GS}(t_{\text{anneal}}) \rangle$  as a function of anneal time for *a*) the prism and *b*) the bipartite isomorphisms simulated using closed quantum system,  $O(2)$  and  $O(3)$  spin-vector dynamics (SVD) for both adiabatic quantum annealing (AQA) and LSTF-DQA protocols. The expected ground-state probability was calculated using Equation 4.14 and the 1 standard deviation error-bars were calculated from bootstrapped sets of ground-state probability.

to be greater than or equal to all other simulated dynamics tested. The bipartite graph transpired to be an easy problem to solve for all dynamical models, as they were able to find the ground-state with 100% probability for  $t_{\text{anneal}} > 100\text{ns}$ , except for SVD  $O(2)$  DQA simulations that failed to consistently find solutions for both instances. In the harder prism graph, both classical DQA models failed whilst their AQA variants performed better, but plateaued at  $t_{\text{anneal}} > 75\text{ns}$  meaning that the even the AQA classical instance could not find solutions to all perturbed instances, unlike the quantum models. This confirms that quantum unitary dynamics are crucial to the success of LSTF-DQA solving MIS problems. Furthermore, the quantum DQA protocol performed better at all time-scales, whereas the success of the SVD DQA protocols is highly dependent on the rate at which the classical state-vector is evolved through the semi-classical potential. After sufficient time, the SVD  $O(3)$  DQA simulation converged for both graphs, whereas the SVD  $O(2)$  DQA expected probability continued to fluctuate significantly, which is caused by the restriction to the  $XZ$ -plane of the Bloch sphere.

In six out of the twenty prism instances, the ground-state probability in quantum DQA was also found to be dependent on which target qubit was selected, with the protocol converging to  $\langle P_{GS}(k, t_{\text{anneal}}) \rangle \rightarrow 1$  faster if the correct target qubit was selected. To assess whether local-frustration contributes to this difference, a classical measure of frustration is defined for spin  $i$  with respect to spin  $j$  as

$$f(h_i, h_j, J_{ij}) = \tanh \left( \text{sign}(h_j) J_{ij} \frac{h_i}{||h_i| - |h_j||} \right) \quad (4.15)$$

where  $f(h_i, h_j, J_{ij}) > 0$  and  $f(h_i, h_j, J_{ij}) < 0$  corresponds to a frustration and alignment across a coupler  $J_{ij}$ , respectively, e.g., ferromagnetic frustration  $f(h_i = 1, h_j = -1, J_{ij} = -1) = 1$  and anti-ferromagnetic alignment  $f(h_i = 1, h_j = -1, J_{ij} = 1) = -1$ . Note that this measure is directional from the qubit defined in the first argument, i.e.,  $f(h_i, h_j, J_{ij})$  is with respect to qubit  $i$ , and  $f(h_i, h_j, J_{ij}) \neq f(h_j, h_i, J_{ji})$  unless  $h_i = h_j$ . It is also a limited measure of frustration that looks only at the local frustration between qubits, rather the wider system. Furthermore,  $J_{ij} = 0.5$  for couplers

in all problems, and therefore the shorthand  $f(h_i, h_j, 0.5) = f(h_i, h_j)$  is used. The local-frustration of a target qubit is defined as the average frustration over all adjacent couplers,


$$f(k) = \frac{1}{3} \sum_{j \in \mathcal{A}(k)} f(h_k, h_j), \quad (4.16)$$

where  $\mathcal{A}$  is the neighbourhood function returning the set of adjacent qubits, where  $|\mathcal{A}(k)| = 3$  for all 3-regular graphs.

Out of the two possible target qubits for the prism instances, the Pearson correlation,  $C_{X,Y}$ , between the ground-state probability and frustration was first measured. There was no significant correlation between ground-state probability and target-qubit frustration across all instances and target qubits, however in cases where a target qubit yielded a higher probability than the other possible target qubit, a correlation of  $C_{P_{GS}, f(k)} = 0.956$  was found between local-frustration and ground-state probability. If a problem possesses a target qubit with this property is not obvious a priori, and only indicates that problems with this property can benefit from quantum LSTF-DQA if it contains high local-frustration that can be relieved through the correct choice of target qubit.

To determine what choice of target qubit should be used in LSTF-DQA, the effect of target-qubit frustration on quantum DQA ground-state probability in prism instances is shown in Figure 4.5. The performance of two heuristics are presented, where the choice of target qubit is based on whether it has the highest or lowest local-frustration, i.e,  $k = \arg \max f(k)$ , or  $k = \arg \min f(k)$ , respectively. For long anneals, a target qubit with lower local frustration tended to 100% probability faster than target qubits with higher frustrations on average. For shorter anneal times, the converse was true, resulting in a cross-over between the two heuristics at  $\sim 60ns$ . The result of lower local frustration tending to give higher expected ground-state probability with LSTF-DQA is also consistent with the correct choice of target qubit for the two-qubit system in Ref. [39], where the frustrated system only formed the correct spectra for the less frustrated qubit.

In this section, it has been shown for MIS problems on 6-qubit 3-regular graphs with random perturbations that quantum dynamics are necessary for the LSTF-DQA



**Figure 4.5:** Difference in expected ground-state probability  $\langle P_{GS}(t_{\text{anneal}}) \rangle$  when quantum LSTF-DQA anneals chose the target qubit based on whether it possesses the largest or lowest local frustration, i.e.,  $k = \arg \max f(k)$  or  $k = \arg \min f(k)$  respectively (see Equation 4.16). The expected ground-state probability was calculated using Equation 4.14 and the 1 standard deviation error-bars were calculated from 1000 bootstrapped sets of ground-state probability.

protocol to run successfully when compared to dynamical spin-vector equations for both  $O(2)$  and  $O(3)$  rotors in Figure 4.4, which failed in several instances to find solutions above 50% probability. Furthermore, quantum LSTF-DQA was also observed to consistently out-perform the adiabatic protocol for all variants tested. However, it was observed that in some instances the choice of target qubit resulted in different ground-state probabilities, and this was analysed against a measure of local frustration defined in Equation 4.15. In these instances, the local frustration of the target qubit was found to positively correlate with ground-state probability ( $C_{P_{GS},f(k)} = 0.956$ ), hinting that problems with this property and high-frustration could benefit from quantum LSTF-DQA. Further analysis of local-frustration in Figure 4.5 for 6-qubit 3-regular graphs indicated that by choosing target qubits with lower local frustration can lead to higher expected ground-state probabilities for

long anneal times. This hypothesis is consistent with previous work [39], supporting the heuristic of choosing a target qubit with the lowest local frustration. Extending this heuristic to other MIS problems graphs without high symmetry and constant coupling strengths is necessary to further test generality.

### 4.1.2 Max Cut

The max-cut of a graph  $G = (V, E)$  is the partitioning of vertices such that the total weight of the edges (which indicate the capacity/flow) between the partition is maximised. A vertex is assigned to a partition depending on whether it is spin up and spin down. The problem is known to be NP-hard, and is described concisely by the antiferromagnetic Ising cost function,

$$Q(\mathbf{s}) = \sum_{i,j \in E} J_{ij} \left( \frac{s_i s_j - 1}{2} \right), \quad (4.17)$$

where  $J_{ij}$  is the edge weight/capacity in the problem graph. The absence of local biases in this problem means that the solution to this problem is always (at least) doubly degenerate, and also does not possess the local frustration required by LSTF-DQA [39]. Therefore, in this section, we provide a framework to solve max-cut when all  $J_{ij} = R$  with the LSTF-DQA protocol by including local frustration into the problem by using the properties of Hamiltonians that always form the correct DQA spectra (e.g., MIS), whilst also maintaining a ground-state solution to max-cut.

Adding a single bias  $\in \mathbb{R}$  to any spin in the max-cut Hamiltonian will break the degenerate-state and result in a single ground-state set that is aligned in the direction of the bias. Note that adding a single bias will always preserve the ground-state. To create local frustration, further local biases must be added, where the extra energy contribution has to be always less than the energy separation between the ground state and all other excited states, which can be bounded for max-cut with uniform weights. Defining  $E^0$  as the ground-state cost/energy,  $n$  biases of constant value  $h$  are added to the problem to ensure frustration in the antiferromagnetic system. To

preserve the ground state, it must hold that


$$E^k - E^0 > \sum_i^n h s_i^0 - h s_i^k, \quad \forall k > 0, \quad (4.18)$$

where  $E^k$  is an excited energy state such that  $E^k > E^0$ , and  $s_i^0$  ( $s_i^k$ ) is the  $i^{\text{th}}$  spin in the ground (excited) state. The worst-case scenario for the summation in equation 4.18 will result in  $E^k - E^0 > 2nh$ . Given that all  $J_{ij} = R$ , the lowest-energy difference between excited and ground states is achieved by swapping vertices in the partitions that yield a slightly lower max-cut, for which the lowest energy change is  $E^k - E^0 \geq 2R$ . The bias value is therefore bounded by,


$$h < \frac{R}{n}, \quad (4.19)$$

which guarantees the preservation of the max-cut ground state whilst being able to increase magnetic frustration artificially in the problem, necessary for successful LSTF-DQA. This bound asymptotically tends to 0 as the number of local biases added increases, and is not valid for problems with non-uniform  $J_{ij}$  weights.

An 8-spin planar graph with  $R = 1$  (shown in Figure 4.6) is used to assess how much artificial local-frustration is required to add to the max-cut cost function in order to yield the correct DQA spectra. It was found that for a given choice of target qubit, a double-crossing was observed if all neighbouring qubits were also biased, i.e., frustration on every coupler connected to the target qubit. Therefore, the bias applied to the target qubit  $k$  and all adjacent qubits is  $h = -\frac{0.99R}{\delta_k + 1}$ , where  $\delta_k$  is the degree of the target qubit. The pre-factor of 0.99 ensures that the bias value does not exceed the bound and preserves the problem ground state. The spectra for one of the lowest degree qubits is seen in Figure 4.6, where the energy level separation in-between the two DQA crossings is small, and is compounded further for qubits of higher degree as the local bias applied tends to zero. This DQA framework for max-cut is therefore anticipated to not perform well in open quantum systems where energy-level broadening occurs, therefore resulting in large transition rates between the ground and first-excited states due to the small energy level separations.



**Figure 4.6:** Illustration of the 8-spin max-cut problem and the respective spectra, showing the two-lowest energy levels in the adiabatic quantum annealing and LSTF-DQA settings. All couplers in this problem have value  $J_{ij} = 1$ , and the biased spins have value  $h = -\frac{0.99}{3}$ . The target qubit is indicated by the square node and the biased adjacent spins in orange. The annealing schedules are  $A(s) = 1 - s$  and  $B(s) = s$ , with  $c_x(s) = c_z(s) = 0$  for the LSTF-DQA protocol.



**Figure 4.7:** Plot of the average ground-state probability for the 8-qubit problem (Figure 4.6) where local bias is added to the target qubit,  $k$ , of degree  $\delta_k$  and all adjacent qubits. The DQA and AQA variants are plotted with solid and dashed lines, respectively. The biased spins have value  $h = -\frac{0.99}{\delta_k + 1}$ , and are annealed with schedules  $A(s) = 1 - s$ ,  $B(s) = s$ , and  $c_x(s) = c_z(s) = 0$ .

The dynamical performance of this framework is evaluated for closed quantum systems in Figure 4.7, where both adiabatic and diabatic quantum annealing protocols are compared for the max-cut problems with artificial local-frustration. If the low-degree target qubits are chosen, then the average DQA ground-state probability  $\overline{P_{GS}}$  is observed to converge to 1 faster than the adiabatic variants. For  $\delta_k = 2$ , the anneal time required to reach 99% probability was  $\sim 5.4ns$  and  $\sim 17.4ns$  for LSTF-DQA and AQA, respectively, resulting in a  $\sim 3.2\times$  speed-up. Note that the offset in anneal time between AQA and DQA anneals is due to the additional initialization time required by DQA (e.g., if  $s_x = 0.2$  then DQA anneals are  $1.25\times$  longer) but the rate at which we anneal through the gaps and crossings is equal. The ground-state probability is averaged over target qubits with the same degree, i.e., problems with the same number of local biases and bias  $h = -\frac{0.99}{\delta_k+1}$ , where target qubits with the lower degrees perform better overall. The one exception being  $\delta_k = 4$ , where the target qubit  $k = 1$  formed a triple crossing instead of the double-crossing seen for all other target qubits. This therefore resulted in  $P_{GS} = 0$  for  $k = 1$ , averaging to  $\overline{P_{GS}} = 0.5$  due to  $k = 5$  DQA ground-state probability tending to 1 for long anneal times.

Therefore, this section has presented an initial framework of generalising LSTF-DQA to problems without local-frustration. It has been shown possible to create the double crossings necessary for DQA by artificially adding localised frustration to max-cut Hamiltonians, whilst preserving the ground-state. For target qubits of low degree, LSTF-DQA is successful in achieving higher ground-state probabilities compared to AQA over all annealing times (i.e., lower times to solution) for the 8-qubit problem tested. However, this framework is only general to max-cut problems with equal weights on the graph edges, and will have biases tending to zero for dense problem graphs. Nonetheless, this framework increases the problem scope in which it is possible to use the LSTF-DQA protocol to find solutions to combinatorial optimization problems more efficiently than AQA can. Further exploration is required to generalise LSTF-DQA to more combinatorial optimization problems, and to investigate to what extent vanishing biases in this framework are impacted by open-



system thermalization effects.

## 4.2 Testing Coherence using Diabatic Quantum Annealing

Coherence in quantum systems is vital to the success of many quantum algorithms, as it allows for tunnelling and interference effects to take place between quantum states. Therefore, it is seen as a key metric for the quality of a qubit used for computation. There are two timescales that can be used to assess the coherence of a quantum system, these are  $T1$  and  $T2$  times that measure the relaxation and dephasing times, respectively. The  $T1$  relaxation time defines the timescale of which a spin undergoes a spontaneous bit flip, e.g. decay from  $|1\rangle$  to the  $|0\rangle$  state, due to an energetic relaxation process. Similarly, the  $T2$  dephasing time defines the timescale in which the phase of a qubit deviates over time such that the initial and final states become uncorrelated. For an ensemble of qubits,  $T2$  can also measure how the mutual phase between qubits is lost over time, which is important for many phenomena such as coherent tunnelling between states.

Both  $T1$  and  $T2$  capture different aspects of decoherence within quantum systems, but a generic method of assessing whether a system possesses coherence is by using interferometry. Within the context of quantum annealing, interference can occur between populated ground and first excited states in a coherent regime due to the difference in the accumulated phases in each state. The splitting and recombination of state probability density takes place at avoided-level crossings, which is analogous to a beam splitter in traditional Mach-Zehnder interferometry, due to the Landau-Zener transitions that occur between states. Interference effects have been exhibited numerically and experimentally in quantum annealing [291, 295, 296], where the standard approach is to initialize the system in the ground state, pass through an avoided-level crossing to split probability density between the ground and first excited state, evolve the system for some time  $t$  to let phase difference accumulate, and finally pass back through another avoided-level crossing to cause interference. The measured interference pattern manifests itself as oscillating ground-

state probability in time (see Section 4.2.2.1 for examples).

In open quantum systems, both relaxation and dephasing decoherence mechanisms will disrupt this interference process. Thermal relaxation and excitation can occur between the ground and first excited state away from the avoided-level crossing as the open system attempts to reach the Gibbs state (the thermal equilibrium state), and dephasing causes damping of state phase. In this section, we will show how the LSTF-DQA protocol can be converted into an interferometry experiment for a multi-qubit system (Section 4.2.1), and provide analytical interferometry expressions for the perturbed ferromagnetic chain (PFC) Hamiltonian at  $M = 1$  (also used in Ref. [39]). How interferometry is affected by decoherence when introducing open-system quantum noise is explored in Section 4.2.2.2, as well as discussing how LSTF-DQA can be used as a test of coherence for quantum systems.

## 4.2.1 Two-qubit Perturbed Ferromagnetic Chain

### 4.2.1.1 Unperturbed System

The first closed-system to be explored using LSTF-DQA interferometry is the perturbed ferromagnetic chain (PFC) at  $M = 1$ , where the Hamiltonian is defined as

$$H_S(s) = -\frac{1}{2} [c_x(s)\sigma_1^x + a(s)\sigma_2^x - c_z(s)(1-d)\sigma_1^z + b(s)(\sigma_2^z + \sigma_1^z\sigma_2^z)] , \quad (4.20)$$

where  $d$  is the perturbative parameter in the system, the target qubit is  $k = 1$ , and the coefficient schedules are

$$c_z(s) = \frac{s - s_x}{1 - s_x} \quad (4.21)$$

$$c_x(s) = \begin{cases} \frac{c_0 s}{s_x} & \text{if } s < s_x \\ c_0 + (c_1 - c_0)c_z(s) & \text{otherwise} \end{cases} \quad (4.22)$$

$$a(s) = \min(1, 1 - c_z(s)) \quad (4.23)$$

$$b(s) = \max(0, c_z(s)) , \quad (4.24)$$

where  $s_x$  marks the beginning of the annealing phase,  $c_0$  is the transverse-field at  $s = s_x$ , and  $c_1$  is the transverse-field magnitude at  $s = 1$ , which should therefore be much less than  $b(s = 1)$  for accurate system measurements in the  $Z$  basis at the end of the anneal.

So far in this chapter, we have specified  $c_x(s) = 0$  in order to achieve the double crossings required for DQA. In this section, we refer to this as the unperturbed system. as when  $c_0 \neq 0$  and  $c_1 \neq 0$  the crossings change into avoided-level crossings (i.e., the beam splitters), which are required for interferometry. To extract analytical equations for the perturbed case, we first derive the eigenenergy functions for the ground and first-excited-state for Equation 4.20 in the unperturbed case (i.e.,  $c_x(s) = 0$ ). Due to the frozen dynamics of the target qubit, the Hamiltonian can be decomposed into two single qubit Hamiltonians with a constant offset, where the first offset is one where the target qubit  $\sigma_1^z$  is spin up ( $H_S^\uparrow$ ), and the second is with the spin down ( $H_S^\downarrow$ ) offset.

$$H_S^\uparrow = \frac{1}{2} \begin{pmatrix} c_z(s)(1-d) - 2b(s) & -a(s) \\ -a(s) & c_z(s)(1-d) + 2b(s) \end{pmatrix} \quad (4.25)$$


$$H_S^\downarrow = \frac{1}{2} \begin{pmatrix} -c_z(s)(1-d) & -a(s) \\ -a(s) & -c_z(s)(1-d) \end{pmatrix}. \quad (4.26)$$

The four energy levels from the two decomposed Hamiltonians form the full spectrum of the original two-qubit Hamiltonian, where the two lowest energy levels are the states that cross one another and are defined as

$$E_0^\uparrow = \frac{1}{2} \left[ (1-d)c_z(s) - \sqrt{4b(s)^2 + a(s)^2} \right] \quad (4.27)$$

$$E_0^\downarrow = -\frac{1}{2} [(1-d)c_z(s) + a(s)], \quad (4.28)$$

where  $E_0^\uparrow$  begins and ends as the ground state, and  $E_0^\downarrow$  is the ground state between the two crossings. Substituting in the coefficient functions and expressing  $E_0^\uparrow$  and  $E_0^\downarrow$



**Figure 4.8:** Plots of *a*) the two lowest energy states in defined in Equations 4.29 and 4.30 in the unperturbed setting for  $d = 0.3$  and  $s_x = 0.2$ , and *b*) the comparison between the numerical (solid lines) and analytical (dotted lines) results for both  $\bar{\Delta}_{10}$  and  $s_g$  as a function of the perturbative parameter,  $d$ . The integral was evaluated numerically using the *QuadGK* package in the *Julia* language.

in terms of the annealing parameter,  $s$ , yields

$$E_0^\uparrow(s) = \begin{cases} \frac{1}{2} \left[ (1-d) \frac{s-s_x}{1-s_x} - 1 \right] & \text{if } s < s_x \\ \frac{(1-d)(s-s_x) - \sqrt{5s^2 - (8s_x+2)s + 4s_x^2 + 1}}{2(1-s_x)} & \text{otherwise} \end{cases} \quad (4.29)$$

$$E_0^\downarrow(s) = \begin{cases} -\frac{1}{2} \left[ (1-d) \frac{s-s_x}{1-s_x} + 1 \right] & \text{if } s < s_x \\ \frac{sd + s_x(1-d) - 1}{2(1-s_x)} & \text{otherwise} \end{cases} \quad (4.30)$$

where we can observe that  $E_0^\downarrow$  is piecewise linear in  $s$  (see Figure 4.8a).

For the rest of the section, we define  $s_x = 0.2$  in order to interpret the effect that the perturbative parameter  $d$  has on the spectra, and therefore its effects on interferometry. The most important variable in interferometry is the relative phase acquired between beam splitters, for which the position of the first beam-splitter is always set by the value of  $s_x$ , which marks the start of the annealing phase. The position of the second crossing (beam splitter),  $s_g$ , for  $s_x = 0.2$  is defined as

$$s_g = \frac{d^2 + 3d - 5}{5(d^2 - d - 1)}, \quad 0 < d < 1. \quad (4.31)$$

The energy difference between the ground and first-excited states in the interval

$[s_x, s_g]$  for the unperturbed case is defined as

$$\Delta_{10}(s) = \frac{1}{8} \left[ 5(1 - 2d)s + 2d + 3 - \sqrt{125s^2 - 90s + 29} \right], \quad (4.32)$$

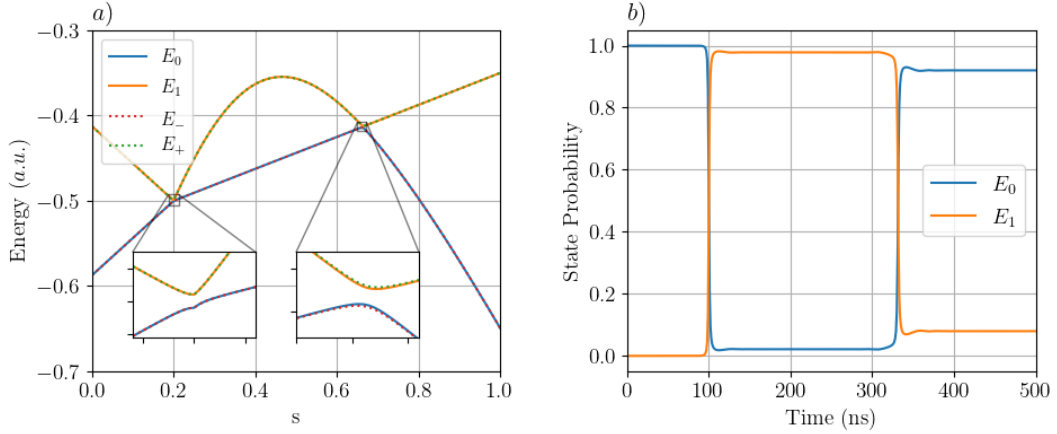
where the integral of  $\Delta_{10}(s)$  in this interval is equal to the phase difference, which is given by

$$\begin{aligned} \bar{\Delta}_{10} &= \int_{s_x=0.2}^{s_g} \Delta_{10}(s) ds \\ &= -\frac{1}{2000} \left[ 64\sqrt{5} \ln \left| \sqrt{5} \sqrt{125s^2 - 90s + 29} + 25s - 9 \right| \right. \\ &\quad \left. + 5(25s - 9) \sqrt{125s^2 - 90s + 29} \right. \\ &\quad \left. + 625(2d - 1)s^2 - 250(2d + 3)s \right]_{s_x=0.2}^{s_g}, \end{aligned} \quad (4.33)$$

Comparison of the numerical and analytical values for  $s_g$  and  $\bar{\Delta}_{10}$  are given in Figure 4.8b as a function of the perturbative parameter  $d$ . The integral was evaluated numerically using the *QuadGK* package in the Julia language, and exact overlap between the analytical and numerical values are observed. Interestingly,  $d$  controls the magnitude of the relative phase difference, which in Section 4.2.2.1 is shown to be equal to the frequency of oscillations. Tuning the interferometer frequency with a single parameter perturbative parameter is simpler from an experimental quantum annealing perspective instead of adjusting coefficient schedules or global energy scales.

#### 4.2.1.2 Perturbed System

In order to perform interferometry, a small perturbation is introduced to the target qubit (i.e.,  $c_0 \neq 0$  and  $c_1 \neq 0$ ) to create avoided-level crossings that act as the beam splitters of state probability between the ground and first-excited states. For initial state preparation and measurement accuracy in the computational basis, it is necessary to have  $c_0 \ll a(s = 0)$  and  $c_1 \ll b(s = 1)$ , respectively, which means that the unperturbed derivations defined in Section 4.2.1.1 can be approximated to the perturbed system if this condition is met. Furthermore, a two-level approximation



**Figure 4.9:** Plot of  $a)$  shows the two avoided-level crossings from the diagonalized Hamiltonian in Equation 4.20 (solid line) and from Equation 4.35 (dotted line) in the perturbed setting with  $c_0 = 2.0 \times 10^{-3}$ ,  $c_1 = 2.75 \times 10^{-3}$ ,  $d = 0.3$  and  $s_x = 0.2$ . Plot  $b)$  illustrates the closed-system evolution of the Hamiltonian defined in plot  $a)$  for a 500 ns anneal.

of the system using the two unperturbed energy levels  $E_0^\downarrow$  and  $E_0^\uparrow$  is represented by the perturbed Hamiltonian


$$H_S^*(s) = \begin{pmatrix} E_0^\uparrow(s) & -\frac{c_x(s)}{2} \\ -\frac{c_x(s)}{2} & E_0^\downarrow(s) \end{pmatrix}, \quad (4.34)$$

where its eigenenergies are given by

$$E_\pm(s) = \frac{1}{2} \left( E_0^\uparrow(s) + E_0^\downarrow(s) \right) \pm \frac{1}{2} \sqrt{\Delta_{10}(s)^2 + |c_x(s)|^2}. \quad (4.35)$$

Figure 4.9 compares Equation 4.35 to the energies found from diagonalizing the full system Hamiltonian (Equation 4.20), and illustrates how state probability is no longer kept solely in the ground state during an anneal with avoided-level crossings. For consistency, the second gap size parameterized by  $c_1$  is chosen to be equal to the size of the first gap, which is always equal to  $c_0$ . Using Equation 4.35, it can be found that  $c_1 = c_0$  results in equal gap sizes  $\forall d$ , however, deriving  $c_1$  numerically (see Figure 4.10) finds that the value of  $c_1$  is dependent on the value of  $d$  used in the Hamiltonian.

This discrepancy can be attributed to the two-level Hamiltonian  $H_S^*(s)$  neglecting higher order perturbative terms from states beyond the two lowest energy states,



**Figure 4.10:** The numerical values of  $c_1$  that result in both avoided-level crossing gaps being equal in size compared to the estimated value of  $c_0 = c_1 = 2 \times 10^{-3}$  (dashed line) as a function of  $d$ .

resulting in minor deviations in the gap size of the second gap,  $\Delta_{10}(s_g)$ , which are visible in the second inset in Figure 4.9a. In order to achieve  $\Delta_{10}(s_g) = c_0$  (i.e., equal gap sizes at each avoided-level crossing),  $c_1 > c_0$  is required to offset the higher order terms making  $\Delta_{10}(s_g)$  smaller than predicted when  $c_1 = c_0$ . The magnitude of this offset is also roughly proportional to  $d$  (Figure 4.10), which also defines the problem hardness in the adiabatic quantum anneal setting, where a smaller  $d$  corresponds to smaller minimum gaps [117]. The oscillation frequency can also be considered to be independent of  $c_0$  and  $c_1$  due to their negligible contribution if the approximation condition is met, but the magnitudes of  $c_0$  and  $c_1$  can impact the oscillation amplitudes observed over time.

## 4.2.2 LTSF-DQA Interferometry

### 4.2.2.1 Coherent Dynamics

Given that each avoided-level crossing acts as a beam-splitter, closed-system interferometry can be conducted using the Hamiltonian defined in Equation 4.20. The interference between the ground and first-excited states are demonstrated by oscillations in the ground state probability as a function of anneal time. By keeping the two-level approximation of the Hamiltonian, the approximate final ground state

probability for a given annealing time is defined as


$$P_{GS} = \left| c_{00}^0 c_{00}^1 + e^{-i\zeta} c_{01}^0 c_{10}^1 \right|^2 \quad (4.36)$$

where  $c_{ij}^k$  is the  $j^{\text{th}}$  state amplitude after passing through the  $k^{\text{th}}$  avoided-level crossing from the  $i^{\text{th}}$  state, and  $\zeta$  is the phase accumulated by amplitude  $c_{01}^0$  in the first excited state between the two avoided-level crossings. Analytical expressions cannot be found for the amplitudes, therefore  $P_{GS}$  is calculated using numerical methods. Using parameters that allow for near-term implementation on hardware, closed-system interferometry simulations of the von-Neumann equation were conducted using the *Hamiltonian Open Quantum System Toolkit* [145] package in the Julia language.

LSTF-DQA interferometry is demonstrated in Figure 4.11, with ground-state probability oscillating due to constructive and destructive interference at the second minimum gap,  $s_g$ . The amplitude of oscillations is observed to have three distinct regimes, the first being that for short anneal times diabatic Landau-Zener transitions are highly probable at both avoided-level crossings (i.e., minimal beam-splitting) resulting in minimal interference at the second avoided-level crossing. The dashed line in Figure 4.11 plots the probability of being in the ground state after the first avoided-level crossing, and shows that it is low for short anneal times. The second regime is one where there is an equal division of probability between the ground and first-excited state after the first avoided-level crossing, resulting in maximum oscillation amplitude; and the final regime is where the system anneal time is long enough such that adiabatic transitions start to dominate at the avoided-level crossings, reducing the probability in the first-excited state and again lowering the oscillation amplitudes.

How quickly we can enter any one of these three regimes is highly dependent on the system parameters used in the interferometry. The division of probability density between the ground and first excited state, and therefore the oscillation amplitude, can be tuned by using the perturbative  $d$  parameter, the first gap position  $s_x$ , initial transverse-field magnitude on the target qubit  $c_0$ , and the global energy scale of the system, which has so far been kept constant. However, these parameters also affect






**Figure 4.11:** Results of simulated LSTF-DQA interferometry for the ground-state probability of the perturbed ferromagnetic chain at  $M = 1$  and  $d = 0.75$  as a function of the anneal time,  $t_{\text{anneal}}$ . The black dashed line is the ground state probability at the midpoint between the two avoided-level crossings. LSTF-DQA interferometry was conducted according to the Hamiltonian defined in Equation 4.20 with  $s_x = 0.2$ ,  $c_0 = 2.0 \times 10^{-3}$ ,  $c_1 = 2.29 \times 10^{-3}$ .

the frequency of oscillations, and therefore careful tuning is required to obtain a system with the desired oscillation amplitude and frequency for a given annealing timescale.

By keeping  $s_x = 0.2$  and  $c_0 = 2.0 \times 10^{-3}$ , we can sweep values of the perturbative parameter  $d$  and anneal time,  $t_{\text{anneal}}$ , to observe the effect on frequency and amplitude of the ground state probability  $P_{GS}$ , as illustrated in Figure 4.12a. The value of  $c_1$  is determined numerically throughout to ensuring that both avoided-level crossings consistently have the same gap size. It may also be noted that so far, experiments have focused on the traditional approach of fixing the system parameters and sweeping the anneal-time, but Figure 4.12a demonstrates that interferometry experiments can also be conducted for a fixed annealing time whilst varying  $d$ . However, the latter is not explored here due to the non-linear oscillation frequency as a function of  $d$ .


Given that the magnitude of transverse-field on the target qubit is sufficiently small, the accumulated phase and therefore the frequency can also be approximated using the analytical solution of the unperturbed system. From Equation 4.36, the



**Figure 4.12:** Plots of *a*) the ground state probability  $P_{GS}$  as a function of the perturbative parameter  $d$  and anneal time  $t_{\text{anneal}}$ , and *b*) the oscillation frequency derived from the numerical methods compared to what is predicted using Equation 4.33. The system has parameters  $s_x = 0.2$  and  $c_0 = 2.0 \times 10^{-3}$ , with  $c_1$  numerically set to always keep the second minimum gap the same size as the gap.

accumulated phase is defined as  $\zeta = 2\pi t_{\text{anneal}} \int_{s_x}^{s_g} \Delta_{10}^*(s) ds$  for the energy gap  $\Delta_{10}^*(s)$  of the perturbed problem. Approximating  $\Delta_{10}^*(s)$  to the unperturbed energy gap  $\Delta_{10}(s)$  results in  $\zeta \approx 2\pi t_{\text{anneal}} \bar{\Delta}_{10}$ , where  $\bar{\Delta}_{10}$  is defined in Equation 4.33 for  $s_x = 0.2$ . Figure 4.12b illustrates the accuracy of this approximation by comparing both the numerically derived frequency and the predicted approximation to the unperturbed analytical solution across various values of  $d$ .

So far, we have only applied interferometry to the simplest case of the perturbed ferromagnetic chain of  $M = 1$ , and have kept all simulations closed-system (i.e., infinite quantum coherence lifetime). To use this method as a test of quantum coherence for multi-qubit systems, we extend this approach to larger perturbed ferromagnetic chain system sizes. In Figure 4.13, the same protocol is applied to both  $M = 2$  and  $M = 3$  case of the perturbed ferromagnetic chain, which also exhibit the same oscillations but with different frequency and amplitude to the  $M = 1$  case. The same control qubit is used in all cases, with  $c_1$  determined numerically to ensure equal sized avoided-level crossings. Therefore, LSTF-DQA interferometry can also be used as a possible test of coherence for multi-qubit systems, and the tunable nature of the Hamiltonian should accommodate for most experimental setups. In the next section, we shall explore to what extent quantum coherence is necessary and



**Figure 4.13:** Demonstration that LSTF-DQA interferometry occurs for PFC systems with  $M > 1$  and  $d = 0.75$ . The target qubit for all problems was kept constant, with additional qubit subsystems always appended to the right of the target qubit. LSTF-DQA interferometry was conducted according to the Hamiltonian defined in Equation 4.20 with  $s_x = 0.2$  and  $c_0 = 2.0 \times 10^{-3}$  for all problems. The value of  $c_1$  was determined numerically for all problems in order to maintain equal gap sizes, which for  $M = 1$ ,  $M = 2$  and  $M = 3$  were found to be  $c_1 = 2.29 \times 10^{-3}$ ,  $c_1 = 2.66 \times 10^{-3}$  and  $c_1 = 2.67 \times 10^{-3}$ , respectively.

how it is possible to extract a measure of decoherence.

#### 4.2.2.2 Measuring Decoherence in Open Systems

LSTF-DQA interferometry has so far only been simulated with full quantum coherence, and to determine whether coherence is necessary we explore how decoherence impacts LSTF-DQA interferometry. To extend the von-Neumann equation used for the coherent simulations, we use the adiabatic master equation (AME) that is commonly used to model quantum noise in open-system quantum annealing [119, 142, 143]. The AME used is defined in Section 1.4.1, where in this model the decoherence of the system occurs in the instantaneous energy eigenbasis. The level of decoherence is defined through the Ohmic bath spectral density

$$\gamma(\omega) = \eta \frac{\omega \exp(-|\omega|/\omega_c)}{1 - \exp(-\beta\omega)}, \quad (4.37)$$

where  $\eta$  is the bath-coupling strength,  $\omega_c$  the bath cut-off frequency and  $\beta$  is the inverse temperature. The type of decoherence introduced into the energy eigenbasis is dephasing noise, which is used to emulate quantum noise in experimental annealing [159].

The sensitivity of LSTF-DQA interferometry to dephasing will determine how the oscillations are damped, from which we can model approximately using the semi-empirical equation defined in Ref. [296]. Defining the dephasing frequency of the ohmic bath to be


$$\gamma_d(s) = \frac{1}{2} \gamma(\Delta_{10}(s)) (1 + \exp(-\beta \Delta_{10}(s))), \quad (4.38)$$

and the average thermalization rate as  $\bar{\gamma}_d = \int_0^1 ds \gamma_d(s)$ , the semi-empirical equation is defined as

$$P'_{GS}(t_f) = (P_{GS}(t_f) - P_E(\beta)) e^{-\bar{\gamma}_d t_f} + P_E(\beta), \quad (4.39)$$

where  $t_f$  is the annealing time,  $P'_{GS}$  is the ground-state probability of the decoherent system,  $P_{GS}$  is the ground state probability of the coherent system, and  $P_E(\beta)$  is the probability of being in the ground state at thermal equilibrium at the end of the anneal, given by  $P_E(\beta) = e^{-\beta E_0} / Z$ , with  $Z = \sum_i e^{-\beta E_i}$ . This expression was derived for a two-level system and describes how the closed-system ground-state probability decays to the steady-state solution (the Gibbs state) as  $t_f \rightarrow \infty$ , where the decay rate is the thermalization rate  $\bar{\gamma}_d$ . The validity of Equation 4.39 holds for interferometry if the two-level approximation made for the perturbed system in Section 4.2.1.2 remains true in the open-system setting, i.e., state probability above the first-excited state is negligible.


LSTF-DQA interferometry simulations in a system governed by the AME with an Ohmic bath is shown in Figure 4.14 for various temperatures and bath-coupling strengths. Given the numerical closed-system solutions, a curve-fit of the open-system measurements to the semi-empirical equation was performed with the Python package *SciPy*, where the thermal ground-state probability,  $P_{\text{Gibbs}}(\beta)$ , and thermalization rate,  $\bar{\gamma}_d$ , were extracted from the fit. The fitted curves plotted



**Figure 4.14:** Open-system simulations of LSTF-DQA interferometry using the adiabatic master equation [142, 143] illustrating oscillations in ground state probability  $P_{GS}$  as a function of anneal time,  $t_{\text{anneal}}$ , for a PFC with  $M = 1$  (Equation 4.20) and  $d = 0.75$ . Solid lines represent the numerical simulation results and the dashed lines are fits to the semi-empirical equation defined in Equation 4.39, with the fitted parameters shown in Figure 4.15. Plot *a*) is for a temperature varying open-system with a bath-coupling strength fixed at  $\eta = 5 \times 10^{-5}$ , and plot *b*) is for an open-system at a fixed temperature of  $T = 8\text{mK}$  with varying bath-coupling strength. The LSTF-DQA schedule parameters are  $s_x = 0.2$ ,  $c_0 = 2.0 \times 10^{-3}$ , and  $c_1 = 2.29 \times 10^{-3}$ .

in Figure 4.14 show close agreement to the numerical result for systems with a bath-coupling strengths of  $\eta < 10^{-4}$ , and for simulations above this threshold, large deviations were observed indicating that the semi-empirical equation approximations are no longer valid.

The average thermalization rates and steady-state probabilities derived from the curve-fits are also presented in Figure 4.15 for various temperatures and bath-coupling strengths. Note that the error-bars to one-standard deviation from fitting to Equation 4.39 are on average 5 orders of magnitude smaller than the values plotted, and are therefore omitted from the plots. However, this does not mean that the curve-fit accurately fitted the data to Equation 4.39. The fitted  $P_E(\beta)$  values are observed to be consistently lower than the true ground state probability at thermal equilibrium for that given temperature, implying that the measured system temperature is consistently higher than the true system temperature. The cause for the higher temperature observations is due to the system not being in thermal equilibrium for the time-scales the fit was performed over, and therefore the system dynamics



**Figure 4.15:** Plots of the parameters fitted to the results in Figure 4.14 using Equation 4.39 for various temperatures,  $T$ , and bath-coupling constants,  $\eta$ . The plot *a*) shows the predicted steady-state probability,  $P_E$ , including the true ground-state probability at thermal equilibrium (black-dashed line) for some temperature  $T$ , and the plot *b*) shows the predicted average thermalization rate  $\bar{\gamma}_d$  from the curve-fit.

can freeze earlier than anticipated. This is commonly known as freeze-out [57], and because this occurs when the gap-sizes between the ground and first-excited states are smaller than the gap at the end of the anneal, the system temperature appears to be larger than expected, also called the effective temperature of the system. This freeze-out effect is also seen for the PFC at  $M = 2$  in Section 3.2.2. A system that possesses late freeze-out will therefore have a lower effective temperature that is closer to the true temperature of the system.

Both the thermalization rate and spectral broadening of energy levels can have competing effects on freeze-out, which makes  $P_E$  a poor choice for measuring system decoherence using LSTF-DQA interferometry. This is demonstrated in Figure 4.15a for a system with  $\eta = 5 \times 10^{-5}$  appearing to be at a higher effective temperature than a system with  $\eta = 1 \times 10^{-4}$  for  $T > 15mK$ . This is due to a higher thermalization rate at larger temperatures shown in Figure 4.15b causing freeze-out to occur later in the anneal. For low temperatures, most of the transitions at the avoided-level crossing as diabatic, with thermal excitation and relaxation rates effected more by the broadening of the energy levels at larger system-bath coupling strengths.

For all problems, the average thermalization rate  $\bar{\gamma}_d$  proved to be the most consistent measure of decoherence, even when curve-fits to the semi-empirical

equation (Equation 4.39) have large deviations from the numerical solutions (e.g., for large system bath coupling strengths). The interesting aspect of this method is that it can be used to characterise noisy multi-qubit quantum systems without complex schedules or any exotic hardware changes. This can be achieved by measuring the oscillating ground-state probability and fitting to the semi-empirical equation to obtain  $\bar{\gamma}_d$ , which is a measure of the multi-qubit coherence time of the system. In fact,  $\bar{\gamma}_d$  can be used to define the T2 time for the adiabatic master equation in the weak coupling limit in Ref [143], and therefore the inverse of the average thermalization rate can be thought of as a multi-qubit extension to the T2 time defined there.

Overall, LSTF-DQA can be used as a scalable multi-qubit interferometer in a coherent quantum system. The implementation is experimentally accessible using current hardware, i.e., no exotic hardware or schedules are required for implementation, and it requires no exotic topology to be directly embedded onto the hardware. We have shown that the system can be reduced to a two-level approximation (Section 4.2.1.2), and consequently the interferometry oscillation frequency can be approximated to an analytical expression (Equation 4.33). The number of parameters in the schedule and problem means that there is a large degree of tunability when it comes to choosing the desired oscillation amplitudes and frequencies. However, in order to ensure that both gap-sizes are the same in LSTF-DQA interferometry, the  $c_1$  schedule parameter had to be determined numerically due to higher-order interactions beyond the two-level approximation creating a smaller than expected gap at the second avoid-level crossing. This effect was quantified in Figure 4.10 to demonstrate that this deviation increases as the perturbative parameter gets smaller, i.e., a harder problem and less separation between the first-excited state and higher-excited states.

In the open-system setting, LSTF-DQA interferometry proved to be sensitive to decoherence, as observed by the damping of oscillations as both temperature and/or bath-coupling strength increased when simulated in an Ohmic bath using the adiabatic master equation. By fitting the numerical results to the semi-empirical equation (Equation 4.39), an additional coherence time can be extracted from the average thermalization rate, even when there were large deviations between the

fitted and simulated result. LSTF-DQA interferometry therefore has the potential to be a good measure of decoherence in multi-qubit systems, provided that the bath-coupling strength and system temperature remain sufficiently low. Establishing this for experimental realizations, as well as for larger systems (since the compounding of decoherence in multi-qubit systems is generally non-trivial), would mean that LSTF-DQA interferometry could serve as a benchmark for multi-qubit decoherence, which is not provided by the single qubit  $T_1$  and  $T_2$  benchmark times typically used today.



# Chapter 5

---

## Conclusion

---

### 5.1 General Conclusions

The computational methods presented here demonstrate that whilst noise and integrated control errors are ubiquitous in modern quantum annealers, it is possible to still extract solutions to combinatorial optimization problems using quantum dynamics. This is achieved by addressing the pertinent issues limiting the success of quantum annealing, first being the hardware restrictions that prevent quantum annealers from finding solutions to large and/or dense problems, second whether incoherent quantum dynamics are of computational value in quantum annealing, and finally can the minimum gaps at avoided-level crossings that close exponentially in problem size be solved using diabatic quantum annealing protocols. Chapter 2 provides solutions to the hardware issues in Section 2.1 and Section 2.2, through better minor-embedding via parameter-setting and reformulation of the optimization problem, respectively. By comparing several parameter-setting methods on random problem instances on the D-Wave 2000Q quantum annealer, the detrimental effects of minor-embedding were shown to be lessened when using schemas that guaranteed it was energetically unfavourable to break a minor-embedding rather than heuristically setting parameters such as those seen in proprietary methods used by D-Wave. A parameter measuring classical minimum non-locality is also identified as a heuristic indicator of parameter setting performance, however the interplay

with optimal embedding coupling strengths (i.e., as small as possible) must also be considered with this measure.

Reformulating the constrained graph-colouring cost-function into several maximal-independent set (MIS) problems in Section 2.2 demonstrated the large reduction in hardware overheads required to solve the graph colouring problem. The reduction in overheads is attributed to the removal of the one-hot constraint in the original graph colouring problem that stipulates the solution must have  $k$  colours, and which therefore also reduces the amount of minor-embedding required. A greedy MIS colouring approach was proposed that reduces the problem size by choosing the MIS that removes the most density from the problem graph. This is in contrast to another in the literature (Ref. [286]) that had a random selection approach, but both rely on improving performance through the removal of the constraint. An incremental performance gain was observed using the greedy approach on the D-Wave Advantage annealer, for planar graph problems of sizes 32 to 2048 nodes, whilst graph colouring with one-hot constraints failed for very small problem sizes.

The computational value of incoherent quantum annealing for solving optimization problems is first addressed in Section 2.3, where reverse-annealing is used to find the ground-state of a small frustrated problem. Given that reverse-annealing is not an adiabatic algorithm, it was demonstrated experimentally on the D-Wave 2000Q that the ground state was accessed using diabatic transitions driven by incoherent dynamics. Ground-state probability was also found to be higher as the Hamming distance between the initial solution and ground-state decreases. This was consistently observed, whilst other quantum, classical, and quantum inspired models did not demonstrate a similar ordering. It was hypothesised that the initialisation phase used to pre-program the computational state in reverse-annealing triggered a phenomenon called spin-bath polarization, which altered the optimization problem being solved. To better contrast between hardware errors and computational processes driven by incoherent quantum dynamics, Chapter 3 introduced the perturbed ferromagnetic chain (PFC), which is a Hamiltonian whose problem hardness (i.e., gap size) is controlled by a tuneable perturbative parameter and not just problem size. In the

transverse-field Ising model, it possesses the property of having a false minimum which is exponentially large in system size, and therefore requires additional dynamics to access the ground state if the system passes through the exponentially scaling minimum gap non-adiabatically.

The computational value of incoherent dynamics is tested by comparing numerical solutions to closed and open quantum system equations, and several spin-vector Monte Carlo (SVMC) methods, which have been proposed as classical mechanisms that D-Wave annealers could follow. Closed system simulations failed to achieve any significant ground-state probability for the hardest problem tested, whereas open-system dynamics provided a thermalisation route to the ground state not accessible in the closed-system setting. The ground-state probability using the open-quantum dynamics for this problem was also found to be two orders of magnitude larger than all classical SVMC models tested, demonstrating that incoherent quantum dynamics are important computationally for the PFC, and can also be used to distinguish between quantum and classical thermalisation mechanisms. This was experimentally tested using various D-Wave annealers, where no model was observed to emulate the dynamics of the D-Wave advantage annealer. However, the PFC was used to show that the Advantage annealer was a noisier quantum system than anticipated due to the distribution of states away from the ground and exponential manifold states. Using the low-noise D-Wave 2000Q annealer, thermalisation was observed for large PFCs where the analogue classical system failed to reach the ground state completely, therefore demonstrating that incoherent quantum dynamics can be useful for computation. A mechanism for the large Hamming weight change in the thermal transition was proposed using intermediate states that exist between the manifold and ground-state for large problem sizes. The effect of this mechanism on time-to-solution by increasing problem size was not observed to scale with the exponential closing minimum gap size on the D-Wave 2000Q. Instead, the time-to-solution had (close to) zero scaling as problem size increased.

The final question regarding diabatic quantum annealing is addressed in Chapter 4, where the locally suppressed transverse-field diabatic quantum annealing

(LSTF-DQA) protocol was extended from the random frustrated problems in Ref. [39] to select MIS and max-cut combinatorial optimisation problem instances. This is to utilise the faster annealing that is possible in LSTF-DQA due to double crossings created by the protocol. Using 3-regular graphs, MIS problems with inherent local frustration (required by LSTF-DQA) demonstrated that quantum unitary dynamics are required due to the failure of classical dynamical models using LSTF-DQA. Measures of local frustration were also defined as heuristics to choose the target qubit, and demonstrated that target qubits with lower local frustration yielded higher ground-state probabilities for the problems tested. A framework to introduce local frustration into the max-cut problem was also presented in Section 4.1.2, by introducing local fields to qubits that preserved the ground-state and create the necessary LSTF-DQA spectra. It was demonstrated that for a small max-cut problem, the LSTF-DQA protocol converged to 100% probability faster than the adiabatic quantum annealing variant when the target qubit has a low degree ( $\sim 3.2\times$  faster for the lowest degree target qubit). However, this method is only applicable to max-cut problems with constant coupling values, limiting the range of optimisation problems that can be solved using LSTF-DQA.

Finally, in Section 4.2 an interferometer using LSTF-DQA is proposed as a test of coherence in quantum annealing. Analytic approximations of the interferometer oscillation frequency are given for a special case of the PFC and are shown to match the numerically derived solutions. The interferometer is extended to larger PFC sizes, such that coherent oscillations were observed for 6-qubit systems and could be extended further. Numerical simulations of LSTF-DQA interferometry performed in open-quantum systems exhibited oscillation damping, and oscillations at low system-bath coupling strengths could be fitted to a semi-empirical equation. The values extracted from fitting this equation to the numerical results were measures of average thermalization rates, and therefore gave a measure of the noise and coherence time in the quantum system. LSTF-DQA is therefore a promising optimization protocol that can be used to solve combinatorial optimisation problems with local-frustration, and also serve as a tool to benchmark the decoherence in many-qubit systems.

Overall, quantum annealing still has the potential to demonstrate a scaling advantage over classical methods for specific optimisation problem classes, however, major challenges such as embedding and hardware errors need to be still overcome when using quantum annealers (such as those produced by D-Wave). Therefore, a black-box approach to quantum annealing is not likely to result in successful computation, and only approaches tailored to the problem can mitigate these issues, such as those presented in this work. This also means that a general quantum optimisation advantage with annealing unlikely, but hybrid or novel quantum optimisation algorithms still hold promise and can still yield good solutions to specific optimisation problems. To help achieve this near term, researchers require increased analogue control of the state-of-the-art quantum annealers to test the many promising algorithms, such as LSTF-DQA. In the longer term, users of quantum annealers will still require improvements to qubit count, inter-qubit connectivity and coherence times to compete with state-of-the-art optimisation methods.

## 5.2 Future Work

The most promising future research area given the work presented lies in further generalisations of LSTF-DQA to optimisation problems if it is able to mitigate/prevent the closing gap problem experienced in adiabatic variants. This can only be tested currently on small problems due to the intractable simulations for larger problems. Given that the LSTF-DQA protocol was designed to be easily implemented on existing hardware, it is therefore of benefit to quantum annealing researchers and algorithm developers to make annealers with higher degrees of qubit control accessible. In the case of LSTF-DQA, individual field control of a single target qubit is required to enable this algorithm. With this additional control, larger optimisation problems can be tested and compared experimentally to adiabatic variants. LSTF-DQA interferometry and therefore tests of coherence can also be conducted on low-noise quantum annealers, allowing for the coherence of many-qubit systems to be evaluated.

In addition to extra qubit controls, rapid quenching of quantum annealers for

state analysis would also help confirm the possible role of intermediate states in the thermalisation mechanism of the PFC. The use of intermediate-states to reach low-energy configurations from excited states would strengthen the connection between quantum annealing and physical systems that also use intermediate-states, e.g., lasers. Furthermore, additional control would also assist in conclusively verifying that the observed incoherent thermalisation mechanism used by the PFC cannot be attributed to effects such as spin-bath polarization, which is observed for long anneal times and in reverse annealing.

# Appendix A

---

## Monte Carlo Algorithms

---

A variety of Markov-chain algorithms are used to simulate quantum annealing and find solutions to optimization problems throughout this document, and here we provide pseudocode to the algorithms implemented. For the complete GitHub repository of the algorithms presented, please refer to Ref. [297]. The first is the simulated annealing (Alg. 2) [114], where the cost of accepting a proposed spin-flip is only accepted given that it yields a lower-energy state or where the Boltzmann factor is greater than a sample from a uniform distribution,  $\text{Uniform}(0, 1)$ . The anneal starts in a hot temperature environment, i.e., many proposed changes accepted, before finishing at a low temperature, freezing the dynamics. Spin-vector Monte Carlo

---

**Algorithm 2** Simulated Annealing

---

```
1: Spins =  $\{-1, 1\}^N$ 
2: TemperatureSchedule = Interpolate( $T_{\text{hot}}, T_{\text{cold}}, \text{NumIncrements} = \text{sweeps}$ )
3: for T in TemperatureSchedule do
4:   RandSpins  $\leftarrow$  RandomPermutation(1, ..., N)
5:   for i in RandSpins do
6:     CurrentSpin  $\leftarrow$  Spinsi
7:     NewSpin  $\leftarrow$  Flip(CurrentSpin)
8:      $\Delta E \leftarrow$  EnergyDifference(i, CurrentSpin, NewSpin)
9:     if  $\exp\{-\Delta E/T\} > \text{Uniform}(0, 1)$  then
10:       Spinsi  $\leftarrow$  NewSpin
11:     end if
12:   end for
13: end for
```

---

(Alg. 3) [277] is a classical analogue to quantum annealing, where the system is annealed in normalised time,  $s$ , which parameterizes coefficient schedules  $A(s)$  and  $B(s)$ . The condition for accepting a new configuration is the same as simulated annealing, but instead discrete spin states are represented as  $O(2)$  rotors given by a continuous angle  $\theta \in (0, \pi)$ , where the discrete spin state is retrieved by  $\text{sign}(\cos \theta)$ . Spin-vector Monte Carlo with transverse-field dependent updates (Alg. 4) [119] is

---

**Algorithm 3** Spin-vector Monte Carlo

---

```

1: SpinVector =  $\{\frac{\pi}{2}\}^N$ 
2: Schedule = Interpolate(0, 1, NumIncrements = sweeps)
3: for  $s$  in Schedule do
4:   RandSpins  $\leftarrow$  RandomPermutation(1, ..., N)
5:   for  $i$  in RandSpins do
6:      $\theta^i \leftarrow$  SpinVector $_i$ 
7:      $\theta^{i+1} \sim$  Uniform(0,  $\pi$ )
8:      $\Delta E \leftarrow$  EnergyDifference( $i$ ,  $\theta^i$ ,  $\theta^{i+1}$ ,  $s$ )
9:     if  $\exp\{-\Delta E/T\} >$  Uniform(0, 1) then
10:       SpinVector $_i \leftarrow \theta^{i+1}$ 
11:     end if
12:   end for
13: end for

```

---

an extension where the angle updates include transverse-field dependence to emulate freeze-out when  $A(s) \ll B(s)$ . The angles are still bounded,  $\theta \in (0, \pi)$ , but are updated additively such that as  $A(s)/B(s) \rightarrow 0$ ,  $\Delta\theta \rightarrow 0$ . The final algorithm listed is path-integral Monte Carlo (Alg. 5) [131, 133], where the quantum dynamics are Trotterized into  $n_\tau$  replicas of the original  $N$ -spin problem. The slices are coupled ferromagnetically according to  $J_\tau = -\frac{1}{2} \ln(\tanh(\beta A(s)/n_\tau))$ , where  $\beta = 1/T$  and contributes to the energy difference calculated in the spin-flip proposals.



---

**Algorithm 4** Spin-vector Monte Carlo – Transverse Field
 

---

```

1: SpinThetas =  $\{\frac{\pi}{2}\}^N$ 
2: Schedule = Interpolate(0, 1, NumIncrements = sweeps)
3: for s in Schedule do
4:   RandSpins  $\leftarrow$  RandomPermutation(1:N)
5:   for i in RandSpins do
6:      $\theta^t \leftarrow$  SpinThetasi
7:      $u \sim$  Uniform( $-\pi, \pi$ )
8:      $\theta^{t+1} \leftarrow \theta^t + \min\left(1, \frac{A(s)}{B(s)}\right) u$ 
9:      $\theta^{t+1} \leftarrow \max(\min(\theta^{t+1}, \pi), 0)$ 
10:     $\Delta E \leftarrow$  EnergyDifference(i,  $\theta^t, \theta^{t+1}, s$ )
11:    if  $\exp\{-\Delta E/T\} >$  Uniform(0, 1) then
12:      SpinThetasi  $\leftarrow \theta^{t+1}$ 
13:    end if
14:  end for
15: end for

```

---



---

**Algorithm 5** Path Integral Monte Carlo
 

---

```

1: Spins =  $\{-1, 1\}^{N \times n_\tau}$ 
2: Schedule = Interpolate(0, 1, NumIncrements = sweeps)
3: for s in Schedule do
4:   for  $\tau$  in  $\{1, \dots, n_\tau\}$  do
5:     RandSpins  $\leftarrow$  RandomPermutation(1, ..., N)
6:     for i in RandSpins do
7:       CurrentSpin  $\leftarrow$  Spinsi,  $\tau$ 
8:       NewSpin  $\leftarrow$  Flip(CurrentSpin)
9:        $\Delta E \leftarrow$  EnergyDifference(i,  $\tau$ , CurrentSpin, NewSpin)
10:      if  $\exp\{-\Delta E/T\} >$  Uniform(0, 1) then
11:        Spinsi,  $\tau$   $\leftarrow$  NewSpin
12:      end if
13:    end for
14:  end for
15: end for

```

---

# Appendix B

---

## List of Publications

---

A list of publications related to the work presented:

1. **D. T. O'Connor**, L. Fry-Bouriaux, and P. A. Warburton, Perturbed Ferromagnetic Chain: Tunable Test of Hardness in the Transverse-Field Ising Model, *Phys. Rev. A* 105, 022410 (2022).
2. 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, 052616 (2021).

A list of publications beyond the scope of the work presented here:

1. **D. O'Connor** and W. Vinci, RBM-Flow and D-Flow: Invertible Flows with Discrete Energy Base Spaces, *ArXiv:2012.13196 [Cs, Stat]* (2021).

# Appendix C

---

## Colophon

---

This document was created using the Overleaf editor using the template provided by the Research Computing team, Information Services Division, University College London. Times Roman typeface using L<sup>A</sup>T<sub>E</sub>X and BibT<sub>E</sub>Xis used throughout, and all references to the literature were collected using the reference management software, Zotero. All figures that were created by the author used the PYTHON packages:

- *matplotlib* v3.7.1 [298]
- *NetworkX* v2.8.8 [299]
- *D-Wave Ocean Software Development Kit* v6.3.0

Numerical simulations and analysis plotted were made accessible in tractable times using the languages JULIA v1.6 [300] and CYTHON v0.26 [301]. Important libraries used in the work presented include *Hamiltonian Open Quantum System Toolkit* v0.7.3 [145] and *QuadQK* v2.6.0 in JULIA, as well as *SciPy* v1.10.1 [302] and *NumPy* v1.24.3 [303] in PYTHON. Furthermore, simulation tools that were created for this work were published to the GitHub repository in Ref. [297].

---

# Bibliography

---

- [1] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6/7), 1981.
- [2] Paul Benioff. The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines. *Journal of Statistical Physics*, 22(5):563–591, May 1980.
- [3] Yuri Manin. Computable and uncomputable. *Sovetskoye Radio, Moscow*, 128, 1980.
- [4] David Deutsch and Roger Penrose. Quantum theory, the Church-Turing principle and the universal quantum computer. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 400(1818):97–117, 1985.
- [5] David Deutsch and Richard Jozsa. Rapid Solution of Problems by Quantum Computation. *Proceedings of the Royal Society of London Series A*, 439:553–558, December 1992.
- [6] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, page 212–219, New York, NY, USA, 1996. Association for Computing Machinery.
- [7] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997.

- [8] Daniel R. Simon. On the power of quantum computation. *SIAM Journal on Computing*, 26(5):1474–1483, 1997.
- [9] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.
- [10] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, October 2019.
- [11] Dorit Aharonov, Wim van Dam, Julia Kempe, Zeph Landau, Seth Lloyd, and Oded Regev. Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM Review*, 50(4):755–787, 2008.
- [12] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. Quantum Computation by Adiabatic Evolution. *arXiv:quant-ph/0001106*, January 2000.
- [13] Tosio Kato. On the Adiabatic Theorem of Quantum Mechanics. *Journal of the Physical Society of Japan*, 5(6):435–439, November 1950.
- [14] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, Joshua Lapan, Andrew Lundgren, and Daniel Preda. A Quantum Adiabatic Evolution Algorithm Applied to Random Instances of an NP-Complete Problem. *Science*, 292(5516):472–475, April 2001.
- [15] Luca Asproni, Davide Caputo, Blanca Silva, Giovanni Fazzi, and Marco Magagnini. Accuracy and minor embedding in subqubo decomposition with fully connected large problems: A case study about the number partitioning problem. *Quantum Machine Intelligence*, 2(1):4, March 2020.
- [16] Vicky Choi. Adiabatic Quantum Algorithms for the NP-Complete Maximum-Weight Independent Set, Exact Cover and 3SAT Problems. *arXiv e-prints*, page arXiv:1004.2226, April 2010.

- [17] Davide Venturelli and Alexei Kondratyev. Reverse Quantum Annealing Approach to Portfolio Optimization Problems. *Quantum Machine Intelligence*, 1(1-2):17–30, May 2019.
- [18] Tobias Stollenwerk, Elisabeth Lobe, and Martin Jung. Flight gate assignment with a quantum annealer. In Sebastian Feld and Claudia Linnhoff-Popien, editors, *Quantum Technology and Optimization Problems*, pages 99–110, Cham, 2019. Springer International Publishing.
- [19] Minsung Kim, Davide Venturelli, and Kyle Jamieson. Leveraging Quantum Annealing for Large MIMO Processing in Centralized Radio Access Networks. *Proceedings of the ACM Special Interest Group on Data Communication*, pages 241–255, August 2019.
- [20] Daisuke Inoue, Akihisa Okada, Tadayoshi Matsumori, Kazuyuki Aihara, and Hiroaki Yoshida. Traffic signal optimization on a square lattice with quantum annealing. *Scientific Reports*, 11(1):3303, February 2021.
- [21] Koki Kitai, Jiang Guo, Shenghong Ju, Shu Tanaka, Koji Tsuda, Junichiro Shiomi, and Ryo Tamura. Designing metamaterials with quantum annealing and factorization machines. *Physical Review Research*, 2(1):013319, March 2020.
- [22] Florian Neukart, Gabriele Compostella, Christian Seidel, David von Dollen, Sheir Yarkoni, and Bob Parney. Traffic flow optimization using a quantum annealer. *Frontiers in ICT*, 4:29, 2017.
- [23] Alejandro Perdomo-Ortiz, Neil Dickson, Marshall Drew-Brook, Geordie Rose, and Alán Aspuru-Guzik. Finding low-energy conformations of lattice protein models by quantum annealing. *Scientific Reports*, 2(1):571, August 2012.
- [24] Vikram Khipple Mulligan, Hans Melo, Haley Irene Merritt, Stewart Slocum, Brian D. Weitzner, Andrew M. Watkins, et al. Designing peptides on a quantum computer. *bioRxiv*, page 752485, Sep 2019.

- [25] Maritza Hernandez and Maliheh Aramon. Enhancing quantum annealing performance for the molecular similarity problem. *Quantum Information Processing*, 16(5):133, April 2017.
- [26] Jeffrey Marshall, Andrea Di Gioacchino, and Eleanor G. Rieffel. Perils of embedding for sampling problems. *Phys. Rev. Research*, 2:023020, Apr 2020.
- [27] Mario S. Könz, Wolfgang Lechner, Helmut G. Katzgraber, and Matthias Troyer. Embedding overhead scaling of optimization problems in quantum annealing. *PRX Quantum*, 2:040322, Nov 2021.
- [28] Tameem Albash and Daniel A. Lidar. Adiabatic quantum computation. *Rev. Mod. Phys.*, 90:015002, Jan 2018.
- [29] Vasil S. Denchev, Sergio Boixo, Sergei V. Isakov, Nan Ding, Ryan Babbush, Vadim Smelyanskiy, John Martinis, and Hartmut Neven. What is the computational value of finite-range tunneling? *Phys. Rev. X*, 6:031015, Aug 2016.
- [30] Andrew D. King, Jack Raymond, Trevor Lanting, Sergei V. Isakov, Masoud Mohseni, Gabriel Poulin-Lamarre, et al. Scaling advantage over path-integral Monte Carlo in quantum simulation of geometrically frustrated magnets. *Nature Communications*, 12(1):1113, February 2021.
- [31] Rolando D. Somma, Daniel Nagaj, and Mária Kieferová. Quantum speedup by quantum annealing. *Phys. Rev. Lett.*, 109:050501, Jul 2012.
- [32] Matthew B. Hastings. The Power of Adiabatic Quantum Computation with No Sign Problem. *Quantum*, 5:597, December 2021.
- [33] John M. Martinis. Saving superconducting quantum processors from decay and correlated errors generated by gamma and cosmic rays. *npj Quantum Information*, 7(1):1–9, June 2021.

- [34] Antti P. Vepsäläinen, Amir H. Karamlou, John L. Orrell, Akshunna S. Dogra, Ben Loer, Francisca Vasconcelos, et al. Impact of ionizing radiation on superconducting qubit coherence. *Nature*, 584(7822):551–556, August 2020.
- [35] Matt McEwen, Lara Faoro, Kunal Arya, Andrew Dunsworth, Trent Huang, Seon Kim, et al. Resolving catastrophic error bursts from cosmic rays in large arrays of superconducting qubits. *Nature Physics*, 18(1):107–111, January 2022.
- [36] C. D. Wilen, S. Abdullah, N. A. Kurinsky, C. Stanford, L. Cardani, G. D’Imperio, et al. Correlated charge noise and relaxation errors in superconducting qubits. *Nature*, 594(7863):369–373, June 2021.
- [37] N. G. Dickson, M. W. Johnson, M. H. Amin, R. Harris, F. Altomare, A. J. Berkley, et al. Thermally assisted quantum annealing of a 16-qubit problem. *Nature Communications*, 4(1):1903, May 2013.
- [38] Sergio Boixo, Vadim N. Smelyanskiy, Alireza Shabani, Sergei V. Isakov, Mark Dykman, Vasil S. Denchev, Mohammad H. Amin, Anatoly Yu Smirnov, Masoud Mohseni, and Hartmut Neven. Computational multiqubit tunnelling in programmable quantum annealers. *Nature Communications*, 7(1):10327, January 2016.
- [39] Louis Fry-Bouriaux, Daniel T. O’Connor, Natasha Feinstein, and Paul A. Warburton. Locally suppressed transverse-field protocol for diabatic quantum annealing. *Phys. Rev. A*, 104:052616, Nov 2021.
- [40] F. Bloch. Nuclear induction. *Phys. Rev.*, 70:460–474, Oct 1946.
- [41] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [42] W. van Dam, M. Mosca, and U. Vazirani. How powerful is adiabatic quantum computation? In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 279–287, 2001.



- [43] Jérémie Roland and Nicolas J. Cerf. Quantum search by local adiabatic evolution. *Phys. Rev. A*, 65:042308, Mar 2002.
- [44] Zhaohui Wei and Mingsheng Ying. A modified quantum adiabatic evolution for the Deutsch–Jozsa problem. *Physics Letters A*, 354(4):271–273, June 2006.
- [45] Itay Hen. Fourier-transforming with quantum annealers. *Frontiers in Physics*, 2:44, 2014.
- [46] B. Apolloni, C. Carvalho, and D. de Falco. Quantum stochastic optimization. *Stochastic Processes and their Applications*, 33(2):233–244, December 1989.
- [47] Tadashi Kadowaki and Hidetoshi Nishimori. Quantum annealing in the transverse ising model. *Phys. Rev. E*, 58:5355–5363, Nov 1998.
- [48] 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(5):343–348, 1994.
- [49] Ernst Ising. Beitrag zur Theorie des Ferromagnetismus. *Zeitschrift für Physik*, 31(1):253–258, February 1925.
- [50] Sergey Bravyi and Barbara Terhal. Complexity of stoquastic frustration-free hamiltonians. *SIAM Journal on Computing*, 39(4):1462–1485, 2010.
- [51] Milad Marvian, Daniel A. Lidar, and Itay Hen. On the computational complexity of curing non-stoquastic Hamiltonians. *Nature Communications*, 10(1):1571, April 2019.
- [52] S. Bravyi, D.P. DiVincenzo, R. Oliveira, and B.M. Terhal. The complexity of stoquastic local Hamiltonian problems. *Quantum Information and Computation*, 8(5):361–385, May 2008.
- [53] Mohammad H. Amin. Searching for quantum speedup in quasistatic quantum annealers. *Phys. Rev. A*, 92:052323, Nov 2015.

- [54] Marcello Benedetti, John Realpe-Gómez, Rupak Biswas, and Alejandro Perdomo-Ortiz. Estimation of effective temperatures in quantum annealers for sampling applications: A case study with possible applications in deep learning. *Phys. Rev. A*, 94:022308, Aug 2016.
- [55] Brian Hu Zhang, Gene Wagenbreth, Victor Martin-Mayor, and Itay Hen. Advantages of Unfair Quantum Ground-State Sampling. *Scientific Reports*, 7(1):1044, April 2017.
- [56] Mohammad H. Amin, Evgeny Andriyash, Jason Rolfe, Bohdan Kulchytskyy, and Roger Melko. Quantum boltzmann machine. *Phys. Rev. X*, 8:021050, May 2018.
- [57] Jeffrey Marshall, Davide Venturelli, Itay Hen, and Eleanor G. Rieffel. Power of Pausing: Advancing Understanding of Thermalization in Experimental Quantum Annealers. *Physical Review Applied*, 11(4):044083, April 2019.
- [58] Sergei V. Isakov, Guglielmo Mazzola, Vadim N. Smelyanskiy, Zhang Jiang, Sergio Boixo, Hartmut Neven, and Matthias Troyer. Understanding quantum tunneling through quantum monte carlo simulations. *Phys. Rev. Lett.*, 117:180402, Oct 2016.
- [59] Evgeny Andriyash and Mohammad H. Amin. Can quantum Monte Carlo simulate quantum annealing? *arXiv:1703.09277 [quant-ph]*, March 2017.
- [60] Tameem Albash and Daniel A. Lidar. Demonstration of a scaling advantage for a quantum annealer over simulated annealing. *Phys. Rev. X*, 8:031016, Jul 2018.
- [61] Troels F. Rønnow, Zhihui Wang, Joshua Job, Sergio Boixo, Sergei V. Isakov, David Wecker, John M. Martinis, Daniel A. Lidar, and Matthias Troyer. Defining and detecting quantum speedup. *Science*, 345(6195):420–424, 2014.
- [62] Matthew B. Hastings. Obstructions to classically simulating the quantum adiabatic algorithm. *Quantum Inf. Comput.*, 13:1038–1076, 2013.

- [63] Michael Jarret, Stephen P. Jordan, and Brad Lackey. Adiabatic optimization versus diffusion monte carlo methods. *Phys. Rev. A*, 94:042318, Oct 2016.
- [64] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar. Sign problem in the numerical simulation of many-electron systems. *Phys. Rev. B*, 41:9301–9307, May 1990.
- [65] Matthias Troyer and Uwe-Jens Wiese. Computational complexity and fundamental limitations to fermionic quantum monte carlo simulations. *Phys. Rev. Lett.*, 94:170201, May 2005.
- [66] Lalit Gupta, Tameem Albash, and Itay Hen. Permutation matrix representation quantum monte carlo. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(7):073105, jul 2020.
- [67] Layla Hormozi, Ethan W. Brown, Giuseppe Carleo, and Matthias Troyer. Nonstoquastic hamiltonians and quantum annealing of an ising spin glass. *Phys. Rev. B*, 95:184416, May 2017.
- [68] Hidetoshi Nishimori and Kabuki Takada. Exponential enhancement of the efficiency of quantum annealing by non-stoquastic hamiltonians. *Frontiers in ICT*, 4:2, 2017.
- [69] Jacob D. Biamonte and Peter J. Love. Realizable hamiltonians for universal adiabatic quantum computers. *Phys. Rev. A*, 78:012352, Jul 2008.
- [70] Vicky Choi. Essentiality of the Non-stoquastic Hamiltonians and Driver Graph Design in Quantum Optimization Annealing. *arXiv e-prints*, page arXiv:2105.02110, May 2021.
- [71] Itay Hen. Determining quantum monte carlo simulability with geometric phases. *Phys. Rev. Research*, 3:023080, Apr 2021.
- [72] Dominik Hangleiter, Ingo Roth, Daniel Nagaj, and Jens Eisert. Easing the monte carlo sign problem. *Science Advances*, 6(33):8341, 2020.

- [73] Lev D Landau. A theory of energy transfer ii. *Phys. Zs. Sowjet*, 4:46, 1932.
- [74] Clarence Zener. Non-Adiabatic Crossing of Energy Levels. *Proceedings of the Royal Society of London Series A*, 137(833):696–702, September 1932.
- [75] Sabine Jansen, Mary-Beth Ruskai, and Ruedi Seiler. Bounds for the adiabatic approximation with applications to quantum computation. *Journal of Mathematical Physics*, 48(10):102111, 2007.
- [76] Alexander Elgart and George A. Hagedorn. A note on the switching adiabatic theorem. *Journal of Mathematical Physics*, 53(10):102202, 2012.
- [77] Lorenzo Campos Venuti, Tameem Albash, Daniel A. Lidar, and Paolo Zanardi. Adiabaticity in open quantum systems. *Phys. Rev. A*, 93:032118, Mar 2016.
- [78] Thomas Jörg, Florent Krzakala, Guilhem Semerjian, and Francesco Zamponi. First-order transitions and the performance of quantum algorithms in random optimization problems. *Phys. Rev. Lett.*, 104:207206, May 2010.
- [79] Sergey Knysh. Zero-temperature quantum annealing bottlenecks in the spin-glass phase. *Nature Communications*, 7(1):12370, August 2016.
- [80] A. P. Young, S. Knysh, and V. N. Smelyanskiy. First-order phase transition in the quantum adiabatic algorithm. *Phys. Rev. Lett.*, 104:020502, Jan 2010.
- [81] Beatriz Seoane and Hidetoshi Nishimori. Many-body transverse interactions in the quantum annealing of the  $p$ -spin ferromagnet. *Journal of Physics A: Mathematical and Theoretical*, 45(43):435301, oct 2012.
- [82] Thomas Jörg, Florent Krzakala, Jorge Kurchan, and Andrew Colin Maggs. Quantum Annealing of Hard Problems. *Progress of Theoretical Physics Supplement*, 184:290–303, 03 2010.
- [83] Thomas Jörg, Florent Krzakala, Jorge Kurchan, and A. C. Maggs. Simple glass models and their quantum annealing. *Phys. Rev. Lett.*, 101:147204, Oct 2008.

- [84] Boris Altshuler, Hari Krovi, and Jérémie Roland. Anderson localization makes adiabatic quantum optimization fail. *Proceedings of the National Academy of Sciences*, 107(28):12446–12450, 2010.
- [85] C. R. Laumann, R. Moessner, A. Scardicchio, and S. L. Sondhi. Quantum adiabatic algorithm and scaling of gaps at first-order quantum phase transitions. *Phys. Rev. Lett.*, 109:030502, Jul 2012.
- [86] Junichi Tsuda, Yuuki Yamanaka, and Hidetoshi Nishimori. Energy gap at first-order quantum phase transitions: An anomalous case. *Journal of the Physical Society of Japan*, 82(11):114004, 2013.
- [87] Andrew M. Childs, Richard Cleve, Enrico Deotto, Edward Farhi, Sam Gutmann, and Daniel A. Spielman. Exponential algorithmic speedup by a quantum walk. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing, STOC '03*, page 59–68, New York, NY, USA, 2003. Association for Computing Machinery.
- [88] Elizabeth Crosson, Edward Farhi, Cedric Yen-Yu Lin, Han-Hsuan Lin, and Peter Shor. Different Strategies for Optimization Using the Quantum Adiabatic Algorithm. *arXiv:1401.7320 [quant-ph]*, January 2014.
- [89] Siddharth Muthukrishnan, Tameem Albash, and Daniel A. Lidar. Tunneling and speedup in quantum optimization for permutation-symmetric problems. *Phys. Rev. X*, 6:031010, Jul 2016.
- [90] Stephen P. Jordan, David Gosset, and Peter J. Love. Quantum-merlin-arthur-complete problems for stoquastic hamiltonians and markov matrices. *Phys. Rev. A*, 81:032331, Mar 2010.
- [91] Edward Farhi, Jeffrey Goldstone, David Gosset, Sam Gutmann, Harvey B. Meyer, and Peter W. Shor. Quantum adiabatic algorithms, small gaps, and different paths. *Quantum Inf. Comput.*, 11(3&4):181–214, 2011.

- [92] Tameem Albash. Role of nonstoquastic catalysts in quantum adiabatic optimization. *Phys. Rev. A*, 99:042334, Apr 2019.
- [93] Tameem Albash and Matthew Kowalsky. Diagonal catalysts in quantum adiabatic optimization. *Phys. Rev. A*, 103:022608, Feb 2021.
- [94] Vrinda Mehta, Fengping Jin, Hans De Raedt, and Kristel Michielsens. Quantum annealing with trigger hamiltonians: Application to 2-satisfiability and nonstoquastic problems. *Phys. Rev. A*, 104:032421, Sep 2021.
- [95] Natasha Feinstein, Louis Fry-Bouriaux, Sougato Bose, and P. A. Warburton. Effects of XX-catalysts on quantum annealing spectra with perturbative crossings. *arXiv e-prints*, page arXiv:2203.06779, March 2022.
- [96] Mustafa Demirplak and Stuart A. Rice. Adiabatic population transfer with control fields. *The Journal of Physical Chemistry A*, 107(46):9937–9945, 2003.
- [97] Adolfo del Campo. Shortcuts to adiabaticity by counterdiabatic driving. *Phys. Rev. Lett.*, 111:100502, Sep 2013.
- [98] Christopher Jarzynski. Generating shortcuts to adiabaticity in quantum and classical dynamics. *Phys. Rev. A*, 88:040101, Oct 2013.
- [99] Bogdan Damski. Counterdiabatic driving of the quantum ising model. *Journal of Statistical Mechanics: Theory and Experiment*, 2014(12):P12019, dec 2014.
- [100] Dries Sels and Anatoli Polkovnikov. Minimizing irreversible losses in quantum systems by local counterdiabatic driving. *Proceedings of the National Academy of Sciences*, 114(20):E3909–E3916, 2017.
- [101] Pieter W. Claeys, Mohit Pandey, Dries Sels, and Anatoli Polkovnikov. Floquet-engineering counterdiabatic protocols in quantum many-body systems. *Phys. Rev. Lett.*, 123:090602, Aug 2019.

- [102] Andreas Hartmann and Wolfgang Lechner. Rapid counter-diabatic sweeps in lattice gauge adiabatic quantum computing. *New Journal of Physics*, 21(4):043025, apr 2019.
- [103] Luise Prielinger, Andreas Hartmann, Yu Yamashiro, Kohji Nishimura, Wolfgang Lechner, and Hidetoshi Nishimori. Two-parameter counter-diabatic driving in quantum annealing. *Phys. Rev. Res.*, 3:013227, Mar 2021.
- [104] Neil G. Dickson and M. H. S. Amin. Does adiabatic quantum optimization fail for np-complete problems? *Phys. Rev. Lett.*, 106:050502, Feb 2011.
- [105] Marek M Rams, Masoud Mohseni, and Adolfo del Campo. Inhomogeneous quasi-adiabatic driving of quantum critical dynamics in weakly disordered spin chains. *New Journal of Physics*, 18(12):123034, dec 2016.
- [106] Yuki Susa, Yu Yamashiro, Masayuki Yamamoto, and Hidetoshi Nishimori. Exponential speedup of quantum annealing by inhomogeneous driving of the transverse field. *Journal of the Physical Society of Japan*, 87(2):023002, 2018.
- [107] Juan I Adame and Peter L McMahon. Inhomogeneous driving in quantum annealers can result in orders-of-magnitude improvements in performance. *Quantum Science and Technology*, 5(3):035011, jun 2020.
- [108] Eliot Kapit and Vadim Oganessian. Noise-tolerant quantum speedups in quantum annealing without fine tuning. *Quantum Science and Technology*, 6(2):025013, feb 2021.
- [109] Zhijie Tang and Eliot Kapit. Unconventional quantum annealing methods for difficult trial problems. *Phys. Rev. A*, 103:032612, Mar 2021.
- [110] A. T. Rezakhani, W.-J. Kuo, A. Hamma, D. A. Lidar, and P. Zanardi. Quantum adiabatic brachistochrone. *Phys. Rev. Lett.*, 103:080502, Aug 2009.

- [111] A. T. Rezakhani, D. F. Abasto, D. A. Lidar, and P. Zanardi. Intrinsic geometry of quantum adiabatic evolution and quantum phase transitions. *Phys. Rev. A*, 82:012321, Jul 2010.
- [112] Patrick R. Zulkowski and Michael R. DeWeese. Optimal protocols for slowly driven quantum systems. *Phys. Rev. E*, 92:032113, Sep 2015.
- [113] Walter Vinci and Daniel A. Lidar. Non-Stoquastic Interactions in Quantum Annealing via the Aharonov-Anandan Phase. *npj Quantum Information*, 3(1):38, December 2017.
- [114] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.
- [115] J. Brooke, D. Bitko, T. F. Rosenbaum, and G. Aeppli. Quantum annealing of a disordered magnet. *Science*, 284(5415):779–781, 1999.
- [116] Andrew Lucas. Ising formulations of many np problems. *Frontiers in Physics*, 2:5, 2014.
- [117] D. T. O’Connor, L. Fry-Bouriaux, and P. A. Warburton. Perturbed ferromagnetic chain: Tunable test of hardness in the transverse-field ising model. *Phys. Rev. A*, 105:022410, Feb 2022.
- [118] Huo Chen and Daniel A. Lidar. Why and when pausing is beneficial in quantum annealing. *Phys. Rev. Applied*, 14:014100, Jul 2020.
- [119] Tameem Albash and Jeffrey Marshall. Comparing relaxation mechanisms in quantum and classical transverse-field annealing. *Physical Review Applied*, 15(1):014029, January 2021.
- [120] Nicolas Gisin, Grégoire Ribordy, Wolfgang Tittel, and Hugo Zbinden. Quantum cryptography. *Rev. Mod. Phys.*, 74:145–195, Mar 2002.
- [121] A. J. Berkley, A. J. Przybysz, T. Lanting, R. Harris, N. Dickson, F. Altomare, et al. Tunneling spectroscopy using a probe qubit. *Phys. Rev. B*, 87:020502, Jan 2013.



- [122] T. Lanting, A. J. Przybysz, A. Yu. Smirnov, F. M. Spedalieri, M. H. Amin, A. J. Berkley, et al. Entanglement in a quantum annealing processor. *Phys. Rev. X*, 4:021041, May 2014.
- [123] Tameem Albash, Itay Hen, Federico M. Spedalieri, and Daniel A. Lidar. Reexamination of the evidence for entanglement in a quantum annealer. *Phys. Rev. A*, 92:062328, Dec 2015.
- [124] G. Vidal and R. F. Werner. Computable measure of entanglement. *Phys. Rev. A*, 65:032314, Feb 2002.
- [125] Peter J. Love, Alec Maassen van den Brink, A. Yu. Smirnov, M. H. S. Amin, M. Grajcar, E. Il'ichev, A. Izmailkov, and A. M. Zagoskin. A Characterization of Global Entanglement. *Quantum Information Processing*, 6(3):187–195, June 2007.
- [126] Barbara M. Terhal. Bell inequalities and the separability criterion. *Physics Letters A*, 271(5):319–326, July 2000.
- [127] Otfried Gühne and Géza Tóth. Entanglement detection. *Physics Reports*, 474(1):1–75, April 2009.
- [128] Federico M. Spedalieri. Detecting entanglement with partial state information. *Phys. Rev. A*, 86:062311, Dec 2012.
- [129] Philipp Hauke, Lars Bonnes, Markus Heyl, and Wolfgang Lechner. Probing entanglement in adiabatic quantum optimization with trapped ions. *Frontiers in Physics*, 3:21, 2015.
- [130] J. Batle, C. H. Raymond Ooi, Ahmed Farouk, M. Abutalib, and S. Abdalla. Do multipartite correlations speed up adiabatic quantum computation or quantum annealing? *Quantum Information Processing*, 15(8):3081–3099, August 2016.

- [131] Roman Martoňák, Giuseppe E. Santoro, and Erio Tosatti. Quantum annealing by the path-integral monte carlo method: The two-dimensional random ising model. *Phys. Rev. B*, 66:094203, Sep 2002.
- [132] H. Rieger and N. Kawashima. Application of a continuous time cluster algorithm to the two-dimensional random quantum Ising ferromagnet. *The European Physical Journal B - Condensed Matter and Complex Systems*, 9(2):233–236, May 1999.
- [133] Giuseppe E. Santoro, Roman Martoňák, Erio Tosatti, and Roberto Car. Theory of quantum annealing of an ising spin glass. *Science*, 295(5564):2427–2430, 2002.
- [134] Lucas T. Brady and Wim van Dam. Quantum monte carlo simulations of tunneling in quantum adiabatic optimization. *Phys. Rev. A*, 93:032304, Mar 2016.
- [135] Elizabeth Crosson and Mingkai Deng. Tunneling through high energy barriers in simulated quantum annealing. *arXiv e-prints*, page arXiv:1410.8484, October 2014.
- [136] Andrew D. King, Juan Carrasquilla, Jack Raymond, Isil Ozfidan, Evgeny Andriyash, Andrew Berkley, et al. Observation of topological phenomena in a programmable lattice of 1,800 qubits. *Nature*, 560(7719):456–460, August 2018.
- [137] Sergio Boixo, Tameem Albash, Federico M. Spedalieri, Nicholas Chancellor, and Daniel A. Lidar. Experimental signature of programmable quantum annealing. *Nature Communications*, 4(1):2067, June 2013.
- [138] Sergio Boixo, Troels F. Rønnow, Sergei V. Isakov, Zhihui Wang, David Wecker, Daniel A. Lidar, John M. Martinis, and Matthias Troyer. Evidence for quantum annealing with more than one hundred qubits. *Nature Physics*, 10(3):218–224, March 2014.

- [139] Erik Luijten and Henk W.J. Blöte. Monte carlo method for spin models with long-range interactions. *International Journal of Modern Physics C*, 06(03):359–370, 1995.
- [140] Philipp Werner, Klaus Völker, Matthias Troyer, and Sudip Chakravarty. Phase diagram and critical exponents of a dissipative ising spin chain in a transverse magnetic field. *Phys. Rev. Lett.*, 94:047201, Jan 2005.
- [141] Philipp Werner and Matthias Troyer. Cluster Monte Carlo Algorithms for Dissipative Quantum Systems. *Progress of Theoretical Physics Supplement*, 160:395–417, 06 2005.
- [142] Tameem Albash, Sergio Boixo, Daniel A. Lidar, and Paolo Zanardi. Quantum adiabatic Markovian master equations. *New Journal of Physics*, 14(12):123016, December 2012.
- [143] Tameem Albash and Daniel A. Lidar. Decoherence in adiabatic quantum computation. *Physical Review A*, 91(6):062320, June 2015.
- [144] Dazhi Xu and Jianshu Cao. Non-canonical distribution and non-equilibrium transport beyond weak system-bath coupling regime: A polaron transformation approach. *Frontiers of Physics*, 11(4):110308, May 2016.
- [145] Huo Chen and Daniel A. Lidar. HOQST: Hamiltonian Open Quantum System Toolkit. *arXiv:2011.14046 [quant-ph]*, November 2020.
- [146] Helmut G Katzgraber. Viewing vanilla quantum annealing through spin glasses. *Quantum Science and Technology*, 3(3):030505, jun 2018.
- [147] Philipp Hauke, Helmut G Katzgraber, Wolfgang Lechner, Hidetoshi Nishimori, and William D Oliver. Perspectives of quantum annealing: methods and implementations. *Reports on Progress in Physics*, 83(5):054401, may 2020.
- [148] Alejandro Perdomo-Ortiz, Alexander Feldman, Asier Ozaeta, Sergei V. Isakov, Zheng Zhu, Bryan O’Gorman, et al. Readiness of quantum optimization machines for industrial applications. *Phys. Rev. Applied*, 12:014004, Jul 2019.

- [149] John Preskill. Reliable quantum computers. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 454(1969):385–410, 1998.
- [150] Kevin C. Young, Mohan Sarovar, and Robin Blume-Kohout. Error suppression and error correction in adiabatic quantum computation: Techniques and challenges. *Phys. Rev. X*, 3:041013, Nov 2013.
- [151] Milad Marvian and Daniel A. Lidar. Error suppression for hamiltonian-based quantum computation using subsystem codes. *Phys. Rev. Lett.*, 118:030504, Jan 2017.
- [152] Kristen L. Pudenz, Tameem Albash, and Daniel A. Lidar. Error-corrected quantum annealing with hundreds of qubits. *Nature Communications*, 5(1):3243, February 2014.
- [153] Walter Vinci, Tameem Albash, and Daniel A. Lidar. Nested quantum annealing correction. *npj Quantum Information*, 2(1):1–6, August 2016.
- [154] Kristen L. Pudenz, Tameem Albash, and Daniel A. Lidar. Quantum annealing correction for random ising problems. *Phys. Rev. A*, 91:042302, Apr 2015.
- [155] Yoshiki Matsuda, Hidetoshi Nishimori, and Helmut G Katzgraber. Ground-state statistics from annealing algorithms: quantum versus classical approaches. *New Journal of Physics*, 11(7):073021, jul 2009.
- [156] Salvatore Mandrà, Zheng Zhu, and Helmut G. Katzgraber. Exponentially biased ground-state sampling of quantum annealing machines with transverse-field driving hamiltonians. *Phys. Rev. Lett.*, 118:070502, Feb 2017.
- [157] Mario S. Könz, Guglielmo Mazzola, Andrew J. Ochoa, Helmut G. Katzgraber, and Matthias Troyer. Uncertain fate of fair sampling in quantum annealing. *Phys. Rev. A*, 100:030303, Sep 2019.

- [158] Andrew D. King, Emile Hoskinson, Trevor Lanting, Evgeny Andriyash, and Mohammad H. Amin. Degeneracy, degree, and heavy tails in quantum annealing. *Phys. Rev. A*, 93:052320, May 2016.
- [159] Tameem Albash, Walter Vinci, Anurag Mishra, Paul A. Warburton, and Daniel A. Lidar. Consistency Tests of Classical and Quantum Models for a Quantum Annealer. *Physical Review A*, 91(4):042314, April 2015.
- [160] Tobias Stollenwerk, Bryan O’Gorman, Davide Venturelli, Salvatore Mandrà, Olga Rodionova, Hokkwan Ng, et al. Quantum annealing applied to de-conflicting optimal trajectories for air traffic management. *IEEE Transactions on Intelligent Transportation Systems*, 21(1):285–297, 2020.
- [161] Masayuki Ohzeki, Akira Miki, Masamichi J. Miyama, and Masayoshi Terabe. Control of automated guided vehicles without collision by quantum annealer and digital devices. *Frontiers in Computer Science*, 1:9, 2019.
- [162] Román Orús, Samuel Mugel, and Enrique Lizaso. Forecasting financial crashes with quantum computing. *Phys. Rev. A*, 99:060301, Jun 2019.
- [163] Michael Marzec. *Portfolio Optimization: Applications in Quantum Computing*, chapter 4, pages 73–106. John Wiley & Sons, Ltd, 2016.
- [164] Gili Rosenberg, Poya Haghnegahdar, Phil Goddard, Peter Carr, Kesheng Wu, and Marcos López de Prado. Solving the optimal trading trajectory problem using a quantum annealer. *IEEE Journal of Selected Topics in Signal Processing*, 10(6):1053–1060, 2016.
- [165] Zhaokai Li, Nikesh S. Dattani, Xi Chen, Xiaomei Liu, Hengyan Wang, Richard Tanburn, et al. High-fidelity adiabatic quantum computation using the intrinsic Hamiltonian of a spin system: Application to the experimental factorization of 291311. *arXiv e-prints*, page arXiv:1706.08061, June 2017.
- [166] Olawale Titiloye and Alan Crispin. Quantum annealing of the graph coloring problem. *Discrete Optimization*, 8(2):376–384, 2011.

- [167] Daniele Ottaviani and Alfonso Amendola. Low Rank Non-Negative Matrix Factorization with D-Wave 2000Q. *arXiv:1808.08721 [quant-ph]*, August 2018.
- [168] Daniel O’Malley, Velimir V. Vesselinov, Boian S. Alexandrov, and Ludmil B. Alexandrov. Nonnegative/Binary matrix factorization with a D-Wave quantum annealer. *PLOS ONE*, 13(12):e0206653, December 2018.
- [169] Davide Venturelli, Dominic J. J. Marchand, and Galo Rojo. Quantum Annealing Implementation of Job-Shop Scheduling. *arXiv e-prints*, page arXiv:1506.08479, June 2015.
- [170] Alan Crispin and Alex Syrichas. Quantum annealing algorithm for vehicle scheduling. In *2013 IEEE International Conference on Systems, Man, and Cybernetics*, pages 3523–3528, 2013.
- [171] Tony T. Tran, Minh Binh Do, Eleanor Gilbert Rieffel, Jeremy D. Frank, Zhihui Wang, Bryan O’Gorman, Davide Venturelli, and J. Christopher Beck. A hybrid quantum-classical approach to solving scheduling problems. In *Symposium on Combinatorial Search*, 2016.
- [172] A. Perdomo-Ortiz, J. Fluegemann, S. Narasimhan, R. Biswas, and V.N. Smelyanskiy. A quantum annealing approach for fault detection and diagnosis of graph-based systems. *The European Physical Journal Special Topics*, 224(1):131–148, February 2015.
- [173] Hartmut Neven, Vasil S. Denchev, Geordie Rose, and William G. Macready. Training a Large Scale Classifier with the Quantum Adiabatic Algorithm. *arXiv e-prints*, page arXiv:0912.0779, December 2009.
- [174] Kristen L. Pudenz and Daniel A. Lidar. Quantum adiabatic machine learning. *Quantum Information Processing*, 12(5):2027–2070, May 2013.

- [175] Alexander Zlokapa, Alex Mott, Joshua Job, Jean-Roch Vlimant, Daniel Lidar, and Maria Spiropulu. Quantum adiabatic machine learning by zooming into a region of the energy surface. *Phys. Rev. A*, 102:062405, Dec 2020.
- [176] Alex Mott, Joshua Job, Jean-Roch Vlimant, Daniel Lidar, and Maria Spiropulu. Solving a Higgs optimization problem with quantum annealing for machine learning. *Nature*, 550(7676):375–379, October 2017.
- [177] Walter Winci, Lorenzo Buffoni, Hossein Sadeghi, Amir Khoshaman, Evgeny Andriyash, and Mohammad H Amin. A path towards quantum advantage in training deep generative models with quantum annealers. *Machine Learning: Science and Technology*, 1(4):045028, nov 2020.
- [178] David H. Ackley, Geoffrey E. Hinton, and Terrence J. Sejnowski. A Learning Algorithm for Boltzmann Machines. *Cognitive Science*, 9(1):147–169, 1985.
- [179] Geoffrey E. Hinton, Simon Osindero, and Yee-Whye Teh. A Fast Learning Algorithm for Deep Belief Nets. *Neural Computation*, 18(7):1527–1554, May 2006.
- [180] Tijmen Tieleman. Training restricted Boltzmann machines using approximations to the likelihood gradient. In *Proceedings of the 25th International Conference on Machine Learning, ICML '08*, pages 1064–1071, New York, NY, USA, July 2008. Association for Computing Machinery.
- [181] Steven H. Adachi and Maxwell P. Henderson. Application of Quantum Annealing to Training of Deep Neural Networks. *arXiv e-prints*, page arXiv:1510.06356, October 2015.
- [182] Marcello Benedetti, John Realpe-Gómez, Rupak Biswas, and Alejandro Perdomo-Ortiz. Quantum-assisted learning of hardware-embedded probabilistic graphical models. *Phys. Rev. X*, 7:041052, Nov 2017.

- [183] Daniel Crawford, Anna Levit, Navid Ghadermarzy, Jaspreet S. Oberoi, and Pooya Ronagh. Reinforcement learning using quantum boltzmann machines. *Quantum Info. Comput.*, 18(1–2):51–74, feb 2018.
- [184] Florian Neukart, David Von Dollen, Christian Seidel, and Gabriele Compostella. Quantum-enhanced reinforcement learning for finite-episode games with discrete state spaces. *Frontiers in Physics*, 5:71, 2018.
- [185] Amir Khoshaman, Walter Vinci, Brandon Denis, Evgeny Andriyash, Hossein Sadeghi, and Mohammad H Amin. Quantum variational autoencoder. *Quantum Science and Technology*, 4(1):014001, sep 2018.
- [186] Max Wilson, Thomas Vandal, Tad Hogg, and Eleanor G. Rieffel. Quantum-assisted associative adversarial network: Applying quantum annealing in deep learning. *Quantum Machine Intelligence*, 3(1):19, June 2021.
- [187] Daniel O’Connor and Walter Vinci. RBM-Flow and D-Flow: Invertible Flows with Discrete Energy Base Spaces. *arXiv:2012.13196 [cs, stat]*, July 2021.
- [188] Alejandro Perdomo-Ortiz, Bryan O’Gorman, Joseph Fluegemann, Rupak Biswas, and Vadim N. Smelyanskiy. Determination and correction of persistent biases in quantum annealers. *Scientific Reports*, 6, 2016.
- [189] Andrew D. King, Sei Suzuki, Jack Raymond, Alex Zucca, Trevor Lanting, Fabio Altomare, et al. Coherent quantum annealing in a programmable 2,000 qubit Ising chain. *Nature Physics*, 18(11):1324–1328, November 2022.
- [190] Yuki Bando, Yuki Susa, Hiroki Oshiyama, Naokazu Shibata, Masayuki Ohzeki, Fernando Javier Gómez-Ruiz, et al. Probing the universality of topological defect formation in a quantum annealer: Kibble-Zurek mechanism and beyond. *Physical Review Research*, 2(3):033369, September 2020.
- [191] Phillip Weinberg, Marek Tylutki, Jami M. Rönkkö, Jan Westerholm, Jan A. Åström, Pekka Manninen, Päivi Törmä, and Anders W. Sandvik. Scaling and



- diabatic effects in quantum annealing with a d-wave device. *Phys. Rev. Lett.*, 124:090502, Mar 2020.
- [192] Bartłomiej Gardas, Jacek Dziarmaga, Wojciech H. Zurek, and Michael Zwolak. Defects in Quantum Computers. *Scientific Reports*, 8(1):4539, March 2018.
- [193] Kohji Nishimura, Hidetoshi Nishimori, and Helmut G. Katzgraber. Griffiths-mccoy singularity on the diluted chimera graph: Monte carlo simulations and experiments on quantum hardware. *Phys. Rev. A*, 102:042403, Oct 2020.
- [194] Paul Kairys, Andrew D. King, Isil Ozfidan, Kelly Boothby, Jack Raymond, Arnab Banerjee, and Travis S. Humble. Simulating the Shastry-Sutherland Ising Model using Quantum Annealing. *PRX Quantum*, 1(2):020320, December 2020.
- [195] R. Harris, Y. Sato, A. J. Berkley, M. Reis, F. Altomare, M. H. Amin, et al. Phase transitions in a programmable quantum spin glass simulator. *Science*, 361(6398):162–165, 2018.
- [196] Helmut G. Katzgraber, Firas Hamze, Zheng Zhu, Andrew J. Ochoa, and H. Munoz-Bauza. Seeking quantum speedup through spin glasses: The good, the bad, and the ugly. *Phys. Rev. X*, 5:031026, Sep 2015.
- [197] Salvatore Mandrà, Zheng Zhu, Wenlong Wang, Alejandro Perdomo-Ortiz, and Helmut G. Katzgraber. Strengths and weaknesses of weak-strong cluster problems: A detailed overview of state-of-the-art classical heuristics versus quantum approaches. *Phys. Rev. A*, 94:022337, Aug 2016.
- [198] Yuki Bando and Hidetoshi Nishimori. Simulated quantum annealing as a simulator of nonequilibrium quantum dynamics. *Phys. Rev. A*, 104:022607, Aug 2021.

- [199] Alejandro Perdomo-Ortiz, Salvador E. Venegas-Andraca, and Alán Aspuru-Guzik. A study of heuristic guesses for adiabatic quantum computation. *Quantum Information Processing*, 10(1):33–52, February 2011.
- [200] Masaki Ohkuwa, Hidetoshi Nishimori, and Daniel A. Lidar. Reverse annealing for the fully connected  $p$ -spin model. *Phys. Rev. A*, 98:022314, Aug 2018.
- [201] Yu Yamashiro, Masaki Ohkuwa, Hidetoshi Nishimori, and Daniel A. Lidar. Dynamics of reverse annealing for the fully connected  $p$ -spin model. *Phys. Rev. A*, 100:052321, Nov 2019.
- [202] Nicholas Chancellor. Modernizing quantum annealing using local searches. *New Journal of Physics*, 19(2):023024, feb 2017.
- [203] Nicholas Chancellor and Viv Kendon. Experimental test of search range in quantum annealing. *Physical Review A*, 104(1):012604, July 2021.
- [204] Hannes Bernien, Sylvain Schwartz, Alexander Keesling, Harry Levine, Ahmed Omran, Hannes Pichler, et al. Probing many-body dynamics on a 51-atom quantum simulator. *Nature*, 551(7682):579–584, November 2017.
- [205] Steven J. Weber, Gabriel O. Samach, David Hover, Simon Gustavsson, David K. Kim, Alexander Melville, et al. Coherent coupled qubits for quantum annealing. *Phys. Rev. Applied*, 8:014004, Jul 2017.
- [206] Paul I. Bunyk, Emile M. Hoskinson, Mark W. Johnson, Elena Tolkacheva, Fabio Altomare, Andrew J. Berkley, et al. Architectural considerations in the design of a superconducting quantum annealing processor. *IEEE Transactions on Applied Superconductivity*, 24(4):1–10, 2014.
- [207] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, et al. Quantum annealing with manufactured spins. *Nature*, 473(7346):194–198, May 2011.

- [208] D. Rosenberg, D. Kim, R. Das, D. Yost, S. Gustavsson, D. Hover, et al. 3D integrated superconducting qubits. *npj Quantum Information*, 3(1):1–5, October 2017.
- [209] S. Poletto, F. Chiarello, M. G. Castellano, J. Lisenfeld, A. Lukashenko, C. Cosmelli, G. Torrioli, P. Carelli, and A. V. Ustinov. Coherent oscillations in a superconducting tunable flux qubit manipulated without microwaves. *New Journal of Physics*, 11(1):013009, January 2009.
- [210] L. DiCarlo, J. M. Chow, J. M. Gambetta, Lev S. Bishop, B. R. Johnson, D. I. Schuster, et al. Demonstration of two-qubit algorithms with a superconducting quantum processor. *Nature*, 460(7252):240–244, July 2009.
- [211] John M. Martinis, S. Nam, J. Aumentado, K. M. Lang, and C. Urbina. Decoherence of a superconducting qubit due to bias noise. *Physical Review B*, 67(9):094510, March 2003.
- [212] D. Vion, A. Aassime, A. Cottet, P. Joyez, H. Pothier, C. Urbina, D. Esteve, and M. H. Devoret. Manipulating the quantum state of an electrical circuit. *Science*, 296(5569):886–889, 2002.
- [213] A. Lupaşcu, S. Saito, T. Picot, P. C. de Groot, C. J. P. M. Harmans, and J. E. Mooij. Quantum non-demolition measurement of a superconducting two-level system. *Nature Physics*, 3(2):119–123, February 2007.
- [214] David M. Berns, Mark S. Rudner, Sergio O. Valenzuela, Karl K. Berggren, William D. Oliver, Leonid S. Levitov, and Terry P. Orlando. Amplitude spectroscopy of a solid-state artificial atom. *Nature*, 455(7209):51–57, September 2008.
- [215] Douglas A. Bennett, Luigi Longobardi, Vijay Patel, Wei Chen, Dmitri V. Averin, and James E. Lukens. Decoherence in rf SQUID qubits. *Quantum Information Processing*, 8(2):217–243, June 2009.

- [216] F. Yoshihara, Y. Nakamura, and J. S. Tsai. Correlated flux noise and decoherence in two inductively coupled flux qubits. *Phys. Rev. B*, 81:132502, Apr 2010.
- [217] E. Il'ichev, S. N. Shevchenko, S. H. W. van der Ploeg, M. Grajcar, E. A. Temchenko, A. N. Omelyanchouk, and H.-G. Meyer. Multiphoton excitations and inverse population in a system of two flux qubits. *Phys. Rev. B*, 81:012506, Jan 2010.
- [218] Guido Burkard, Roger H. Koch, and David P. DiVincenzo. Multilevel quantum description of decoherence in superconducting qubits. *Phys. Rev. B*, 69:064503, Feb 2004.
- [219] Philip Krantz, Morten Kjaergaard, Fei Yan, Terry P. Orlando, Simon Gustavsson, and William D. Oliver. A Quantum Engineer's Guide to Superconducting Qubits. *Applied Physics Reviews*, 6(2):021318, June 2019.
- [220] R. Harris, J. Johansson, A. J. Berkley, M. W. Johnson, T. Lanting, Siyuan Han, et al. Experimental demonstration of a robust and scalable flux qubit. *Phys. Rev. B*, 81:134510, Apr 2010.
- [221] R. Harris, M. W. Johnson, S. Han, A. J. Berkley, J. Johansson, P. Bunyk, et al. Probing Noise in Flux Qubits via Macroscopic Resonant Tunneling. *Physical Review Letters*, 101(11):117003, September 2008.
- [222] Catherine McGeoch and Pau Farre. The D-Wave Advantage System: An Overview. [https://www.dwavesys.com/media/s3qbjp3s/14-1049a-a\\_the\\_d-wave\\_advantage\\_system\\_an\\_overview.pdf](https://www.dwavesys.com/media/s3qbjp3s/14-1049a-a_the_d-wave_advantage_system_an_overview.pdf), 2020. Accessed: 2022-05-03.
- [223] Probing Mid-Band and Broad-Band Noise in Lower-Noise D-Wave 2000Q Fabrication Stacks. [https://www.dwavesys.com/media/4j4gg1gf/14-1034a-a\\_probling\\_noise\\_in\\_ln\\_2000q\\_fabrication\\_stacks.pdf](https://www.dwavesys.com/media/4j4gg1gf/14-1034a-a_probling_noise_in_ln_2000q_fabrication_stacks.pdf), 2019. Accessed: 2022-05-03.

- [224] Bibek Pokharel, Zoe Gonzalez Izquierdo, P. Aaron Lott, Elena Strbac, Krzysztof Osiewalski, Emmanuel Papathanasiou, Alexei Kondratyev, Davide Venturelli, and Eleanor Rieffel. Inter-generational comparison of quantum annealers in solving hard scheduling problems. *arXiv e-prints*, page arXiv:2112.00727, December 2021.
- [225] Tristan Zaborniak and Rogério de Sousa. Benchmarking hamiltonian noise in the d-wave quantum annealer. *IEEE Transactions on Quantum Engineering*, 2:1–6, 2021.
- [226] P. J. D. Crowley, T. Đurić, W. Vinci, P. A. Warburton, and A. G. Green. Quantum and classical dynamics in adiabatic computation. *Phys. Rev. A*, 90:042317, Oct 2014.
- [227] T. Lanting, M. H. S. Amin, M. W. Johnson, F. Altomare, A. J. Berkley, S. Gildert, et al. Probing high-frequency noise with macroscopic resonant tunneling. *Physical Review B*, 83(18):180502(R), May 2011.
- [228] Kelly Boothby, Andrew D King, and Jack Raymond. Zephyr Topology of D-Wave Quantum Processors. [https://www.dwavesys.com/media/2uznec4s/14-1056a-a\\_zephyr\\_topology\\_of\\_d-wave\\_quantum\\_processors.pdf](https://www.dwavesys.com/media/2uznec4s/14-1056a-a_zephyr_topology_of_d-wave_quantum_processors.pdf), 2021. Accessed: 2022-05-03.
- [229] QPU-Specific Characteristics — D-Wave System Documentation documentation. [https://docs.dwavesys.com/docs/latest/doc\\_physical\\_properties.html](https://docs.dwavesys.com/docs/latest/doc_physical_properties.html). Accessed: 2022-05-03.
- [230] Jun Cai, William G. Macready, and Aidan Roy. A practical heuristic for finding graph minors. *arXiv e-prints*, page arXiv:1406.2741, June 2014.
- [231] Wolfgang Lechner, Philipp Hauke, and Peter Zoller. A quantum annealing architecture with all-to-all connectivity from local interactions. *Science Advances*, 1(9):e1500838, 2015.

- [232] Andrea Rocchetto, Simon C. Benjamin, and Ying Li. Stabilizers as a design tool for new forms of the Lechner-Hauke-Zoller annealer. *Science Advances*, 2(10):e1601246, 2016.
- [233] Yan-Long Fang and P. A. Warburton. Minimizing minor embedding energy: An application in quantum annealing. *Quantum Information Processing*, 19(7):191, May 2020.
- [234] Vicky Choi. Minor-embedding in adiabatic quantum computation: I. The parameter setting problem. *Quantum Information Processing*, 7(5):193–209, October 2008.
- [235] Vicky Choi. Minor-embedding in adiabatic quantum computation: II. Minor-universal graph design. *Quantum Information Processing*, 10(3):343–353, June 2011.
- [236] Kristen L. Pudenz. Parameter setting for quantum annealers. *2016 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6, 2016.
- [237] Davide Venturelli, Salvatore Mandrà, Sergey Knysh, Bryan O’Gorman, Rupak Biswas, and Vadim Smelyanskiy. Quantum optimization of fully connected spin glasses. *Phys. Rev. X*, 5:031040, Sep 2015.
- [238] Nicholas Chancellor. Domain wall encoding of discrete variables for quantum annealing and QAOA. *Quantum Science and Technology*, 4(4):045004, aug 2019.
- [239] Jie Chen, Tobias Stollenwerk, and Nicholas Chancellor. Performance of domain-wall encoding for quantum annealing. *IEEE Transactions on Quantum Engineering*, 2:1–14, 2021.
- [240] Masayuki Ohzeki. Breaking limitation of quantum annealer in solving optimization problems under constraints. *Scientific Reports*, 10(1):3126, February 2020.

- [241] Charles J. Geyer and Elizabeth A. Thompson. Constrained monte carlo maximum likelihood for dependent data. *Journal of the Royal Statistical Society. Series B (Methodological)*, 54(3):657–699, 1992.
- [242] Koji Hukushima and Koji Nemoto. Exchange monte carlo method and application to spin glass simulations. *Journal of the Physical Society of Japan*, 65(6):1604–1608, 1996.
- [243] J. Houdayer. A cluster Monte Carlo algorithm for 2-dimensional spin glasses. *The European Physical Journal B - Condensed Matter and Complex Systems*, 22(4):479–484, August 2001.
- [244] J. Houdayer. The wormhole move: A new algorithm for polymer simulations. *The Journal of Chemical Physics*, 116(5):1783–1787, February 2002.
- [245] Zheng Zhu, Andrew J. Ochoa, and Helmut G. Katzgraber. Efficient cluster algorithm for spin glasses in any space dimension. *Phys. Rev. Lett.*, 115:077201, Aug 2015.
- [246] Firas Hamze and Nando de Freitas. From fields to trees. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, UAI '04, page 243–250, Arlington, Virginia, USA, 2004. AUAI Press.
- [247] Alex Selby. Efficient subgraph-based sampling of Ising-type models with frustration. *arXiv e-prints*, page arXiv:1409.3934, September 2014.
- [248] K. Hukushima and Y. Iba. Population annealing and its application to a spin glass. *AIP Conference Proceedings*, 690(1):200–206, 2003.
- [249] J. Machta. Population annealing with weighted averages: A monte carlo method for rough free-energy landscapes. *Phys. Rev. E*, 82:026704, Aug 2010.
- [250] E. Schrödinger. An undulatory theory of the mechanics of atoms and molecules. *Phys. Rev.*, 28:1049–1070, Dec 1926.

- [251] A.G. Redfield. The theory of relaxation processes. In John S. Waugh, editor, *Advances in Magnetic Resonance*, volume 1 of *Advances in Magnetic and Optical Resonance*, pages 1–32. Academic Press, 1965.
- [252] A. J. Leggett, S. Chakravarty, A. T. Dorsey, Matthew P. A. Fisher, Anupam Garg, and W. Zwerger. Dynamics of the dissipative two-state system. *Rev. Mod. Phys.*, 59:1–85, Jan 1987.
- [253] E. Paladino, Y. M. Galperin, G. Falci, and B. L. Altshuler.  $1/f$  noise: Implications for solid-state quantum information. *Rev. Mod. Phys.*, 86:361–418, Apr 2014.
- [254] Evgeny Mozgunov and Daniel Lidar. Completely positive master equation for arbitrary driving and small level spacing. *Quantum*, 4:227, February 2020.
- [255] E. B. Davies and H. Spohn. Open quantum systems with time-dependent Hamiltonians and their linear response. *Journal of Statistical Physics*, 19(5):511–523, November 1978.
- [256] G. Lindblad. On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics*, 48(2):119–130, June 1976.
- [257] Vittorio Gorini, Andrzej Kossakowski, and E. C. G. Sudarshan. Completely positive dynamical semigroups of  $n$ -level systems. *Journal of Mathematical Physics*, 17(5):821–825, 1976.
- [258] Frederik Nathan and Mark S. Rudner. Universal lindblad equation for open quantum systems. *Phys. Rev. B*, 102:115109, Sep 2020.
- [259] H. Dekker. Noninteracting-blip approximation for a two-level system coupled to a heat bath. *Phys. Rev. A*, 35:1436–1437, Feb 1987.
- [260] Anurag Mishra, Tameem Albash, and Daniel A. Lidar. Finite temperature quantum annealing solving exponentially small gap problem with non-monotonic success probability. *Nature Communications*, 9(1):2917, July 2018.



- [261] Yuki Bando, Ka-Wa Yip, Huo Chen, Daniel A. Lidar, and Hidetoshi Nishimori. Breakdown of the weak coupling limit in experimental quantum annealing. *arXiv e-prints*, page arXiv:2111.07560, November 2021.
- [262] M. H. S. Amin and Dmitri V. Averin. Macroscopic resonant tunneling in the presence of low frequency noise. *Phys. Rev. Lett.*, 100:197001, May 2008.
- [263] Anatoly Yu Smirnov and Mohammad H Amin. Theory of open quantum dynamics with hybrid noise. *New Journal of Physics*, 20(10):103037, oct 2018.
- [264] Masuo Suzuki. Relationship between d-Dimensional Quantal Spin Systems and (d+1)-Dimensional Ising Systems: Equivalence, Critical Exponents and Systematic Approximants of the Partition Function and Spin Correlations. *Progress of Theoretical Physics*, 56(5):1454–1469, 11 1976.
- [265] Hidetoshi Nishimori and Gerardo Ortiz. *Elements of Phase Transitions and Critical Phenomena*, volume 9780199577224. Oxford University Press, India, January 2011.
- [266] Bettina Heim, Troels F. Rønnow, Sergei V. Isakov, and Matthias Troyer. Quantum versus classical annealing of ising spin glasses. *Science*, 348(6231):215–217, 2015.
- [267] Glen Bigan Mbeng, Lorenzo Privitera, Luca Arceci, and Giuseppe E. Santoro. Dynamics of simulated quantum annealing in random ising chains. *Phys. Rev. B*, 99:064201, Feb 2019.
- [268] Robert H. Swendsen and Jian-Sheng Wang. Nonuniversal critical dynamics in monte carlo simulations. *Phys. Rev. Lett.*, 58:86–88, Jan 1987.
- [269] Ulli Wolff. Collective monte carlo updating for spin systems. *Phys. Rev. Lett.*, 62:361–364, Jan 1989.
- [270] Tameem Albash, Gene Wagenbreth, and Itay Hen. Off-diagonal expansion quantum monte carlo. *Phys. Rev. E*, 96:063309, Dec 2017.

- [271] Itay Hen. Off-diagonal series expansion for quantum partition functions. *Journal of Statistical Mechanics: Theory and Experiment*, 2018(5):053102, may 2018.
- [272] D. C. Handscomb. A monte carlo method applied to the heisenberg ferromagnet. *Mathematical Proceedings of the Cambridge Philosophical Society*, 60(1):115–122, 1964.
- [273] Anders W. Sandvik. Stochastic series expansion method for quantum ising models with arbitrary interactions. *Physical Review E*, 68(5), Nov 2003.
- [274] Lalit Gupta and Itay Hen. Elucidating the interplay between non-stoquasticity and the sign problem. *Advanced Quantum Technologies*, 3(1):1900108, 2020.
- [275] John R. Klauder. Path integrals and stationary-phase approximations. *Physical Review D*, 19(8):2349–2356, April 1979.
- [276] F. T. Arecchi, Eric Courtens, Robert Gilmore, and Harry Thomas. Atomic coherent states in quantum optics. *Phys. Rev. A*, 6:2211–2237, Dec 1972.
- [277] Seung Woo Shin, Graeme Smith, John A. Smolin, and Umesh Vazirani. How "Quantum" is the D-Wave Machine? *arXiv:1401.7087 [quant-ph]*, May 2014.
- [278] John Smolin and Graeme Smith. Classical signature of quantum annealing. *Frontiers in Physics*, 2:52, September 2014.
- [279] Ryogo Kubo and Natsuki Hashitsume. Brownian Motion of Spins. *Progress of Theoretical Physics Supplement*, 46:210–220, 06 1970.
- [280] A. M. Jayannavar. Brownian motion of spins; generalized spin Langevin equation. *Zeitschrift für Physik B Condensed Matter*, 82(1):153–156, February 1991.
- [281] T.L. Gilbert. A phenomenological theory of damping in ferromagnetic materials. *IEEE Transactions on Magnetics*, 40(6):3443–3449, 2004.

- [282] Lev Davidovich Landau and E Lifshitz. On the theory of the dispersion of magnetic permeability in ferromagnetic bodies. *Phys. Z. Sowjet.*, 8:153, 1935.
- [283] Richard M. Karp. *Reducibility among Combinatorial Problems*, pages 85–103. Springer US, Boston, MA, 1972.
- [284] Hui Kang, Ya-nan Zhao, and Fang Mei. A Graph Coloring Based TDMA Scheduling Algorithm for Wireless Sensor Networks. *Wireless Personal Communications*, 72(2):1005–1022, September 2013.
- [285] I. Chlamtac, A. Ganz, and G. Karmi. Lightpath communications: an approach to high bandwidth optical wan's. *IEEE Transactions on Communications*, 40(7):1171–1182, 1992.
- [286] Julia Kwok and Kristen Pudenz. Graph Coloring with Quantum Annealing. *arXiv e-prints*, page arXiv:2012.04470, December 2020.
- [287] Paul Turán. On an extremal problem in graph theory. *Matematikai*, 48:436–452, 1941.
- [288] K. Appel and W. Haken. Every planar map is four colorable. Part I: Discharging. *Illinois Journal of Mathematics*, 21(3):429 – 490, 1977.
- [289] K. Appel, W. Haken, and J. Koch. Every planar map is four colorable. Part II: Reducibility. *Illinois Journal of Mathematics*, 21(3):491 – 567, 1977.
- [290] M. H. S. Amin and V. Choi. First-order quantum phase transition in adiabatic quantum computation. *Phys. Rev. A*, 80:062326, Dec 2009.
- [291] William D. Oliver, Yang Yu, Janice C. Lee, Karl K. Berggren, Leonid S. Levitov, and Terry P. Orlando. Mach-zehnder interferometry in a strongly driven superconducting qubit. *Science*, 310(5754):1653–1657, 2005.
- [292] Salvatore Mandrà and Helmut G. Katzgraber. A deceptive step towards quantum speedup detection. *Quantum Science and Technology*, 3(4):04LT01, July 2018.

- [293] H. A. Kramers and G. H. Wannier. Statistics of the Two-Dimensional Ferromagnet. Part I. *Physical Review*, 60(3):252–262, August 1941.
- [294] Leo Zhou, Sheng-Tao Wang, Soonwon Choi, Hannes Pichler, and Mikhail D. Lukin. Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices. *Phys. Rev. X*, 10:021067, Jun 2020.
- [295] Mika Sillanpää, Teijo Lehtinen, Antti Paila, Yuriy Makhlin, and Pertti Hakonen. Continuous-time monitoring of landau-zener interference in a cooper-pair box. *Phys. Rev. Lett.*, 96:187002, May 2006.
- [296] Humberto Munoz-Bauza, Huo Chen, and Daniel Lidar. A double-slit proposal for quantum annealing. *npj Quantum Information*, 5(1):1–11, May 2019.
- [297] D. O’Connor. Quantum annealing monte carlo tools. <https://github.com/dtoconnor/QAMCT>, June 2023. Accessed: 2023-06-04.
- [298] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007.
- [299] Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. Exploring network structure, dynamics, and function using networkx. In Gaël Varoquaux, Travis Vaught, and Jarrod Millman, editors, *Proceedings of the 7th Python in Science Conference*, pages 11 – 15, Pasadena, CA USA, 2008.
- [300] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *Society for Industrial and Applied Mathematics review*, 59(1):65–98, 2017.
- [301] S. Behnel, R. Bradshaw, C. Citro, L. Dalcin, D.S. Seljebotn, and K. Smith. Cython: The best of both worlds. *Computing in Science Engineering*, 13(2):31–39, April 2011.

- [302] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
- [303] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, et al. Array programming with NumPy. *Nature*, 585(7825):357–362, September 2020.