

Durham E-Theses

Suppressing precision errors by connecting copies of Ising models for continuous-time quantum computing

JEMMA ELLEN BENNETT

How to cite:

BENNETT, JEMMA ELLEN (2023) Suppressing precision errors by connecting copies of Ising models for continuous-time quantum computing. Doctoral thesis, Durham University.

Use policy



This work is licensed under a [Creative Commons Attribution 3.0 \(CC BY\)](https://creativecommons.org/licenses/by/3.0/)

Suppressing precision errors by connecting copies of Ising models for continuous-time quantum computing

Jemma Bennett

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



Department of Physics
Durham University
March 2023

Suppressing precision errors by connecting copies of Ising models for continuous-time quantum computing

Jemma Bennett

Abstract

Classical optimization problems can be mapped to Ising models in order to be solved by continuous-time quantum computing. It was recognised in [1–3], that these problems are susceptible to a lack of precision in the fields and couplings of the Ising model. In this thesis we introduce a scheme first described in [4], which aims to suppress errors caused by lack of precision. This scheme was inspired by quantum annealing correction (QAC), first introduced in [1], where physical qubits in multiple copies of a Ising model are linked together. However, we introduce several innovations thereby making our scheme distinct. First, when determining the ground state of the problem, we require only one copy to be correct, because the solution quality can be checked efficiently. Second, using this “one correct copy” setting, we find the optimal strength of links connecting the copies to be anti-ferromagnetic and close to the minimum strength allowed by the precision. Here we find an improvement (on average) above separate copies and copies connected ferromagnetically. Third, we find that configurations of copies that contain frustration (e.g. a loop of three or five copies), provide a further improvement in fraction correct. Numerically testing our innovations on small instances of spin chains and spin glasses, we find improved tolerance to lack of precision equating to around 3 bits of precision improvement at $p = 7$. We develop a link selection protocol which aims to determine in a computationally non-intensive fashion, whether or not to connect corresponding qubits in different copies. Here, we obtain mixed results, with improvement in fraction correct over separate copies only for precisions $p < 4$. Finally, we apply our error suppression scheme when computing with quantum walks. In this setting we find that the improvement from using our technique is lost for all values of precision. We hypothesise this is due to the way our error suppression scheme functions by allowing ‘access’ to excited states, available innately in quantum walks.

Acknowledgements

Firstly, I would like to thank my supervisor Viv Kendon, for all her support and encouragement over the last few years. Your comments and analysis have been vital to keeping me on track during the research process, and your encouragement and ability to answer my questions before I've even asked them has definitely boosted my morale at times. Thanks for the effort you have put in to meet me in person as well as virtually despite us living at opposite ends of the country most of the time!

I have had the privilege to have Nick Chancellor as my co-supervisor. Thank you for all your incredibly useful and often prescient comments, suggestions and contributions to our research discussions. They definitely helped me along the way. Also thanks for teaching me how to program my first quantum walk!

Thank you to Adam Callison, for not only providing many contributions to our research discussions but also saving me from hours of computational strife by contributing the code for both the branch and bound technique and quantum walk dynamics that support an extensive amount of research in this thesis. Thanks for answering my many questions and the time you put into editing the code.

Thanks also to Steph Foulds, for being my go to person for asking about any PhD student admin type questions and for reading an early chapter draft of this thesis. Thanks to Mia West and Tom O'Leary, who both provided disproportionately strong early contributions to the research in this thesis despite the length of time of their projects!

For the portion of my PhD spent in Durham, I would like to thank the members of QLM for their friendly and welcoming atmosphere. Especially the people involved in coffee breaks, fatmol, pubmol and fitmol, which I definitely took advantage of while I had the chance! Specifically thanks to Clare for including me on some of her many parkrun and running outings, and encouraging me to sign up for Edinburgh marathon! Also thanks to the many people involved in the admin side of my PhD:

Simon, Matt, Stuart, Kevin, plus others I'm sure - thanks for keeping the ball rolling! I would also like to acknowledge the funding from a EPSRC funded DTG studentship, project reference 2214392, and Durham university, without which, the research in this thesis would not have been possible.

Outside of the university I would like to thank Durham City Harriers for giving me a great running club to train and compete for and my housemates Diana and Diego for always including me for the short time I was around and for feeding me delicious Mexican food on multiple occasions!

Due certain events in the past few years this PhD has not always followed along the traditional routes, and that has meant spending a lot more time in Newton Abbot than I otherwise might have expected.

For the academic side of my life, I would like to thank the members of the Exeter University quantum theory group, for accepting me as one of their own, and including me in their theory meetings and extra-curricular activities! Outside of the university, I would like to thank the Teignbridge trotters running club and Mid-Devon cycling club for keeping me fit (and more importantly sane!), the folk orchestra of Torbay for welcoming me to play violin, and Ben's friends and the gang at Malting's taphouse for providing me with Friday night entertainment!

For related reasons, I need thank my fiancé Ben's parents Anne and John for their support and especially for allowing me to stay in their house for much longer than anticipated during a national lock-down (announced the day after my arrival).

I also need to thank my family, Mum, Dad, Kerenza and Alex for all their encouragement and support over the years, even before this PhD started, funding my interest in physics by buying me copious amounts of physics and maths related books and even quantum computing T-shirts. Thank you for bearing with me all these years as I tried to explain ever more niche concepts of physics to you in the hope you'd understand - even if you did coin my research as 'Betamax quantum computing' - I'm still in all your debt for helping me get there!

Finally, I couldn't have completed this PhD without the support of my fiancé Ben, who has been there all the way, suffering 7 hour commutes, my frustrations at code and physics, having to share (an at times very warm) home office, and my constant need to drag him outside for some form of exercise. Thanks for always being there for me and for knowing exactly what will cheer me up.

Contents

Declaration	viii
List of Figures	ix
List of Tables	xii
Notation	xiii
1 Introduction	1
1.1 Classical computing	1
1.2 Quantum computing	3
2 Continuous-time quantum computing	7
2.1 Quantum Walks	7
2.1.1 Search problem using a quantum walk	9
2.2 Adiabatic Quantum Computing	12
2.3 Quantum Annealing	13
2.4 Chapter summary	14
3 Optimisation problems	16
3.1 Transverse field Ising Model	17
3.2 Ising spin chains	18
3.3 Sherrington-Kirkpatrick spin glasses	18
3.4 Maximum independent set	19
3.5 Chapter summary	20
4 Error models	22
4.1 Types of Noise/error	22
4.1.1 Bit flip and phase flip errors	22

4.1.2	Quantum decoherence	23
4.1.3	Qubit loss	23
4.1.4	Limited precision	23
4.2	Modelling limited precision	24
4.2.1	Random error model	25
4.2.2	Deterministic mid-point error model	26
4.2.3	Deterministic random error model	27
4.3	Effect of the error model on single copies	27
4.4	Chapter Summary	29
5	Quantum error correction	30
5.1	Quantum error correction for gate based QC	30
5.2	Quantum error correction for continuous QC	32
5.3	QAC	32
5.3.1	Summary of the mean field analysis of QAC	34
5.3.2	Disadvantages to QAC	35
5.4	Our error suppression method	36
5.4.1	An MIS example which retains the original ground state	38
5.5	Chapter summary	40
6	Numerical Methods	41
6.1	Simulation of classical dynamics	42
6.1.1	Branch and bound technique	42
6.2	Simulation of quantum dynamics	43
6.3	Problem instance generation	43
6.4	Fitting and error analysis	44
6.5	Chapter summary	45
7	Link Strength Tests	46
7.1	Optimal link strength	47
7.1.1	Random error model	48
7.1.2	Link strength tests using other error models	50
7.2	Individual instances	52
7.3	The effect of no fields	54
7.4	On the mechanism of error suppression via anti-ferromagnetic links	56
7.4.1	An example of the mechanism of error suppression	57
7.5	Chapter summary	60
8	Error Suppression Results	62
8.1	Two copies of spin glasses	64

8.1.1	Deterministic random error model	64
8.1.2	Random error model	64
8.2	Two copies of spin chains	66
8.3	Three copies	68
8.3.1	Spin glasses in a chain	68
8.3.2	Spin glasses in a loop	70
8.3.3	Spin chains in a loop	71
8.4	Four copies	72
8.4.1	Spin glasses in a loop	72
8.5	Five copies	75
8.5.1	Spin glasses in a loop	75
8.6	Comparison of five and three copies	78
8.7	Chapter summary	79
9	Precision Improvements	82
9.1	Conditional fraction correct	83
9.2	Precision improvements	85
9.2.1	Scaling of fraction correct with n	87
9.3	Quantifying improvements	89
9.3.1	Spin glasses	89
9.3.2	Spin glasses	91
9.3.3	Spin chains	92
9.4	Chapter summary	93
10	Link Selection Protocol	95
10.1	Threshold behaviour	96
10.1.1	Energy levels below and above threshold	98
10.1.2	Inclusion of negative fields	102
10.2	Fraction correct versus J_F for 2-qubit example with positive fields	106
10.2.1	Two copies	106
10.2.2	Three copies	107
10.2.3	2-qubit example with one positive field, one negative field	108
10.3	Larger size Ising models	109
10.4	The link selection protocol	111
10.4.1	Link selection protocol	112
10.5	Error suppression results with the protocol	113
10.5.1	Spin glasses	113
10.5.2	Spin chains	116
10.6	Chapter summary	116

11 Quantum Walk Dynamics	120
11.1 Optimal γ	121
11.2 Average success probability versus precision	123
11.3 Differences in average success probability	125
11.4 Average success probability with the protocol	126
11.5 Average success probability for larger problem sizes	129
11.6 Chapter summary	130
11.7 Future work	132
12 Conclusion and Outlook	133
12.1 Conclusion	133
12.2 Outlook	137
Bibliography	139

Declaration

I confirm that no part of the material offered has previously been submitted by myself for a degree in this or any other university. Where material has been generated through joint work, the contribution of others has been indicated.

Jemma Bennett

Publications

Parts of chapters 7, 8, 9, and 11 contain graphs and analyses that were included in the paper [4] Jemma Bennett, Adam Callison, Tom O’Leary, Mia West, Nicholas Chancellor and Viv Kendon, *Using copies to improve precision in continuous-time quantum computing*, (2022), <https://arxiv.org/abs/2206.02545>. Soon to be published in Quantum Science and Technology.

The copyright of this thesis rests with the author. No quotation from it should be published without the author’s prior written consent and information derived from it should be acknowledged.

List of Figures

1.1	Artist's impression of a dilution refrigerator of a quantum computer . . .	4
2.1	Diagram of 3-dimensional hypercube	9
2.2	Plot of quantum walk on a hypercube, initial state $ 000\rangle$	10
2.3	Plot of quantum walk search on a hypercube	11
2.4	Plot of AQC search on a hypercube	13
2.5	Energy levels of 3 and 5 qubit system undergoing AQC search	14
3.1	Artist's impression of a Rubik's cube in its minimum energy state . . .	16
3.2	Diagram of 5-qubit Ising spin chain	18
3.3	Diagram of 5-qubit SK spin glass	18
3.4	Diagram of 5-qubit MIS example	19
4.1	Diagram of number lines showing precision divisions	24
4.2	Plot comparing error models, single copies, fraction correct vs precision	28
5.1	Diagram of our error suppression scheme	38
7.1	Fraction correct vs J_F , random error model, 5 qubit, spin glass and spin chain	48
7.2	Fraction correct vs J_F , deterministic mid-point error model, 5 qubit, spin glass and spin chain	50
7.3	Fraction correct vs J_F , deterministic random error model, 5 qubit, spin glass and spin chain	51
7.4	Fraction correct vs J_F , $p = 3$, deterministic random error model, 4 single instances, 3 copies of 5 qubit spin glasses.	52
7.5	Fraction correct vs J_F , $p = 3$, random error model, 4 single instances, 3 copies of 5 qubit spin glasses.	53
7.6	Fraction correct vs J_F , 2 and 3 copies, 5 qubit spin glasses, no fields . .	55

7.7	Plot of energy vs ground states, single copy, 5 qubit spin glass, $p = 3$. . .	58
8.1	Fraction correct vs precision, deterministic random error model, 2 copies vs single, 5-qubit, spin glasses	63
8.2	Fraction correct vs precision, random error model, 2 copies vs single, 5-qubit, spin glasses	65
8.3	Fraction correct vs precision, deterministic random error model, 2 copies vs single, 5-qubit, spin chains	67
8.4	Fraction correct vs precision, deterministic random error model, 3 copies chain vs single, 5-qubit, spin glasses	69
8.5	Fraction correct vs precision, deterministic random error model, 3 copies loop vs single, 5-qubit, spin glasses	70
8.6	Fraction correct vs precision, deterministic random error model, 3 copies loop vs single, 5-qubit, spin chains	72
8.7	Fraction correct vs precision, deterministic random error model, 4 copies loop vs single, 5-qubit, spin glasses	73
8.8	Fraction correct vs precision, deterministic random error model, 5 copies loop vs single, 5-qubit, spin glasses	76
8.9	Difference in fraction correct vs precision, deterministic random error model, 5 copies loop vs 3 copies loop, spin glasses	79
8.10	Diagram: 5 copies of 5-qubit spin glasses arranged in “bowtie configuration”	79
9.1	Conditional fraction correct vs precision, 3 copies loop vs single, 5-qubit spin chains	83
9.2	Conditional fraction correct vs precision, 3 copies loop vs single, 5-qubit spin glasses	85
9.3	Conditional fraction correct vs precision, 3 copies loop vs single, 9-qubit spin glasses	86
9.4	Fraction correct vs precision, for 3+1 copies vs single, 5, 7, 8, 9-qubit spin glasses	87
9.5	Average energy gap vs number of qubits, spin glasses	88
9.6	Fraction correct at $p = 3$ vs n , for separate and 3+1cp copies for 5, 6, 7, 8, 9-qubit spin glasses	89
9.7	Fraction broken vs precision, 3+1 copies vs single, 5, 7, 8, 9-qubit spin glasses, log y plot	90
9.8	Precision improvement vs precision, 3+1 copies vs single, 5, 6, 7, 8, 9-qubit spin glasses	91
9.9	Precision improvement vs precision, 3+1 copies vs single, 5, 6, 7, 8, 9-qubit spin chains	92

10.1	Diagram: 2 qubit Ising model with the fields and couplings labelled. . .	96
10.2	Diagram: 2 copies of a 2 qubit Ising model connected with antiferromagnetic links, with fields and couplings labelled.	96
10.3	Diagram: Threshold of 2 qubit ground states when both fields are positive, or both fields are negative.	98
10.4	Diagram: Threshold of 2 qubit ground states when one field positive, one field negative.	99
10.5	Energy levels and energy shifts of two 2 qubit examples, with two positive/negative fields.	100
10.6	Diagram: 2 copies of 2-qubit Ising models above and below threshold (two positive fields)	102
10.7	Energy levels and energy shifts of two 2 qubit examples with one negative, one positive field	103
10.8	Diagram: two copies of 2-qubit Ising above and below threshold (one positive, one negative fields)	105
10.9	Fraction correct vs J_F : Above and below threshold 2 copies 2-qubit chain	105
10.10	Fraction correct vs J_F : Above and below threshold 3 copies 2-qubit chain	107
10.11	Fraction correct vs J_F : Below and above threshold 2 copies, one negative, one positive field, 2-qubit chain	108
10.12	2 and 3 copy 5-qubit spin glass, protocol and no protocol fraction correct/broken vs precision	114
10.13	2 and 3 copy 5-qubit spin glass protocol and no protocol fraction correct/broken vs precision	117
11.1	Average success probability vs γ , 4 single instances, 3 copy vs single, 5-qubit spin glasses	122
11.2	Average success probability vs precision, 10^3 instances, three copies vs single, 5-qubit spin glasses	124
11.3	Difference in average success probability vs precision, three copies vs single, 5-qubit spin glasses	126
11.4	Average success probability vs precision, for 3 copies with link protocol vs single, 5-qubit spin glasses	127
11.5	Normalised average success probability vs precision, single copies, 5, 6, 7, 8, 9-qubit spin glasses	129

List of Tables

5.1	Table comparing correct ground states of three connected and single copies of a 5-qubit MIS problem	39
7.1	Table showing configurations and energy penalties of two anti-ferromagnetically connected Ising models	56
8.1	Table showing configurations of four connected Ising model copies and their energies with penalties	74
8.2	Table showing configurations of five connected Ising model copies and their energies with penalties	77
10.1	Table showing for 2-qubit Ising model, possible start states, end state and transition inequalities	99
10.2	Table showing for 3-qubit Ising model, possible start states, end state and transition inequalities	111
10.3	Table showing for the n -qubit Ising chain, the number and neighbourly-ness of bit flips, transition inequality and diagram	111

Notation

AQC Adiabatic Quantum computing

SK spin glasses Sherrington-Kirkpatrick spin glasses, an fully connected graph of spins (qubits). Where the strength of the fields and couplings lie between -1 and +1. Only two spin couplings are considered (no higher order couplings).

Spin chains A 1D chain of qubits, where each qubit is only coupled to its nearest neighbours, i.e. 0 to 1, 1 to 2, 2 to 3 and so on. The spin chains used in this thesis do not have periodic boundary conditions.

Fraction correct The fraction of times we find the correct ground state across many (10^3 or more) instances. For multiple copies, it adds to the fraction correct if one or more of the copies have the correct ground state.

Random error model Error model where several random samples (repeats) of each value is taken. As the errors here are not fixed, repeats can improve the fraction correct. See section 4.2.1.

Deterministic mid-point error model Error model where the erred value is set to be the mid-point of the interval in which the true value lies. The error here is fixed to this mid-point and therefore repeats do not improve fraction correct. See section 4.2.2.

Deterministic random error model Error model where the erred value is set to be a random value in the interval in which the true value lies. The error here is fixed to this random value and therefore repeats do not improve fraction correct. See section 4.2.3.

Introduction

1.1 Classical computing

The first classical computers evolved as ways of solving practical problems. The Antikythera mechanism, estimated to have been built sometime between 200 [5] and 60 [6] BC, is thought to be the oldest example [7] of an analogue mechanical computer. It was used to predict the location of planets and timings of eclipses decades in advance [8].

Mechanical analogue computers were again developed in between the 3rd and 13th centuries, once again for astronomical purposes [9–11]. Around 1200 AD, an elaborate clock known as Al-Jazari’s castle clock was built. This clock did not only keep time but also showed the zodiac, the path of the Sun and phases of the Moon. Notably, it was also possible to re-program the length of day and night to account for their changes throughout the year [12].

One of the most influential inventions in the history of computing was the Jacquard loom. Patented by Joseph Jacquard in 1804, it was an improvement on previous similar machines [13]. Its main innovation was the use of punched cards which allowed different designs to be ‘programmed’ into the machine [14].

This punch card design for inputting ‘programs’ was seen by 19th century mathematician and inventor Charles Babbage and he planned to incorporate this design into his proposed analytical engine, the successor to the difference engine [15]. The programmability of this machine allowed Ada Lovelace to famously propose a ‘program’ for calculating Bernoulli numbers that could have been run on the machine [16]. However the analytical engine was ultimately never built.

The punch card design was re-used later, in the 1880s, when Herman Hollerith developed a machine which used them for collecting and tabulating data [17]. This

technology was famously used in the 1890 US census. Hollerith founded a company that went on to become part of IBM [18].

While coming up with his solution to the Entscheidungsproblem in his famous ‘On computable numbers’ paper [19], Alan Turing came up with an abstract idea of a mechanical machine which was capable of universal computing. This machine, (later coined a ‘Turing machine’), consisted of a ‘head’ which could read and write a finite number of symbols, a current ‘state’ selected from a finite set of states, and an infinite tape that passed through the machine. The infinite tape is divided up into cells, each cell can hold a single symbol. The ‘head’ moves along the tape reading one cell at a time. Then depending on this reading and the current ‘state’, the machine makes a decision on a symbol to write and then whether to move left, right or stop the computation. As well as being capable of solving the Entscheidungsproblem, Turing realised that this machine would be capable of universal computation, that is to take in any set of instructions (if they could be translated into the right format) and compute a response. This abstract machine became the base for the theory of computation.

Computers up until this point were created for specific purposes, such as collecting and processing data or performing certain mathematical functions and were never fully programmable and therefore not considered Turing complete, that is, capable of universal computation (aside from Babbage’s analytical engine if it had ever been built). Indeed, the term ‘computer’ was mostly used at the time to refer to women employed to carry out complex calculations such as ballistics during world war II [20]. It is estimated that, 75% of the work force involved in codebreaking at Bletchley park were women [21].

The first Turing complete, programmable, electronic, general purpose digital computer, is generally believed to be the Electronic Numerical integrator and computer or ENIAC for short [22], which was announced to the public in 1946 [23]. Out of around 200 women employed as human ‘computers’ for ballistics calculations, six (Kay McNulty, Betty Jennings, Betty Snyder, Marlyn Meltzer, Fran Bilas and Ruth Lichtermann) were chosen to program the new machine. Due to attitudes at the time their contribution to the project was not well acknowledged and they did not receive the same level of recognition as their male counterparts [20].

In 1959 Mohamed Atalla and Dawon Kahng of Bell labs developed the metal-oxide-silicon field-effect transistor (MOSFET) [24, 25]. This was the first compact transistor that could be miniaturised. This development enabled the development of microprocessors which allowed computers to continually shrink as long as the technology to create them was there [26, 27]. As the size of computers decreased, so did the wires within them, making them faster.

In the following years, computers continued to decrease in size, leading Gordon Moore to conjecture in 1965 that the number of transistors on a chip would double every year [28], though he then revised it to every two years in 1975. This became known as Moore's Law. So far Moore's Law has held true, but as we near the length scales of molecules, it is becoming clear that this progress will at some point reach its limit. In 2012 a team from the university of New South Wales developed a single atom transistor [29] and in 2022 IBM announced a 2nm chip, this is smaller than the width of a single strand of human DNA [30]. There is not yet a consensus on whether the limit has been reached. Gordon Moore never got to see the conclusion to this discussion, he died on 24th March 2023, a few days before this thesis was submitted [31].

There is another reason why reaching the length scales of molecules and beyond is interesting. At this length scale particles start following the rules of quantum rather than classical mechanics. This ultimately brings challenges for classical computing as things no longer act as we expect them to. However, if we are able to harness the quantum behaviour of these particles, we may be able to compute in a different way by designing quantum computers [32].

1.2 Quantum computing

In the 1980s, a big question for classical computers, was whether they were able to simulate classical physics. At first it didn't look possible as the first classical gates weren't considered reversible. However, building on the work of Rolf Landauer [33], Charles H. Bennett showed that a Turing machine could be made so that it was both logically and thermodynamically reversible [34].

Paul Benioff then took this one step further and showed that it was possible in principle to build a Turing machine that functioned using quantum mechanical rules [35]. In a 1981 meeting Feynman addressed the question of whether it was possible to simulate quantum physics on classical computer [36]. By repurposing Bell's inequality, he was able to show that this would not be possible [37]. Therefore, in order to simulate quantum physics a new type of *quantum* machine would be needed.

Beyond physics, it was still yet to be seen whether this new type of computer could provide advantages above classical computing. The Deutsch-Jozsa algorithm, published in 1992 and improved upon in 1998 [38], was one of the first examples of an algorithm that could run exponentially faster on a quantum computer [39]. Soon more examples came along. The Bernstein-Vazirani algorithm which solved a re-

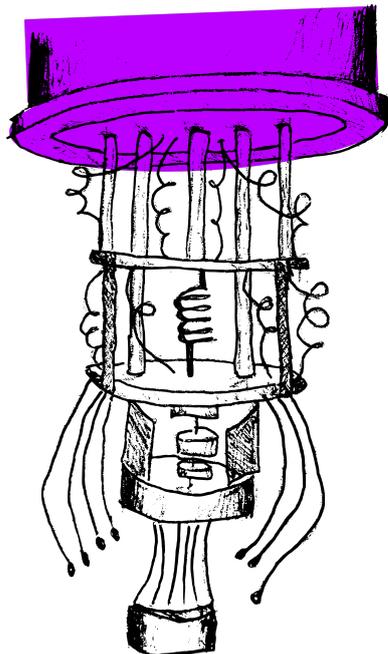


Figure 1.1: Artist's impression of the inside of dilution refrigerator for a quantum computer. The superconducting quantum computer chip would be stuck to the bottom where it's coldest.

stricted version of the Deutsch-Jozsa algorithm was developed in 1993. Then in 1994, Shor developed an algorithm which would allow the prime factors of integers to be found in polylogarithmic time [40]. This algorithm was famously capable of breaking RSA encryption. In 1994, Simon developed an algorithm that could distinguish between a one-to-one or two-to-one function, with an exponential speed-up compared to classical computers [41], and in 1996 Lov Grover devised an algorithm for unstructured search that provided a $O(\sqrt{N})$ speed-up compared the classical brute force approach [42].

Another crucial question was how these quantum machines could be implemented experimentally. The original contenders were techniques such as trapped ions [43, 44], super-conducting qubits[45, 46], quantum dots [47] and nuclear magnetic resonance (NMR)[48–50]. Despite early successes, such as the first demonstration of Shor's algorithm on a quantum computer [51], NMR quantum computing has fallen out of favour due to issues with scalability [52]. Trapped ions, superconducting qubits and quantum dots continue to be active areas of research with a multitude more techniques entering the game, such as neutral atoms in optical lattices [53, 54], linear optical set-ups[55], and diamond NV centers [56]. Despite the diversity of techniques, no single method of quantum computing has yet emerged as dominant and they each face their own set of challenges. For example, to keep

qubits in a coherent regime, superconducting qubits require the cooling of their qubits to extremely low temperatures, by use of a dilution fridge (an artist's impression of which can be seen in figure 1.1). For this reason, the image of the dilution fridge has become synonymous with quantum computing for many people.

Initially quantum computation was thought of in a way that was analogous to classical computing. Qubits were acted upon by logic gates which carried out certain functions. Some quantum gates were equivalent to the gates seen in classical computing, such as the NOT and CNOT (or X and control-X) gates. NOT gates flip a qubit from 0 to 1 or vice versa, and CNOT gates flip from 0 to 1 or vice versa depending on the value of another (the control) qubit. Other gates did not have a classical equivalent, such as the phase flip (Z) gate, which changes the phase of a qubit, or the Hadamard gate which introduces a superposition of states to the qubit. However this was not the only way to compute using quantum computers and soon several other ways of computing were developed, such as: measurement-based quantum computing [57], analogue quantum computing [58] or continuous-time quantum computing. In this thesis we chose to focus on continuous-time quantum computing as it is the basis of our research. An equation free introduction to these other types of quantum computation can be found in [59].

In this thesis, I will first describe the background theory necessary to understand the functioning and results of our error suppression scheme. Then I will present our scheme for error suppression in continuous-time quantum computing, which was first described in [4]. Finally I will present numerical analysis into the functioning of our scheme under classical and quantum dynamics, while speculating on the theory behind its function. The content of the chapters is summarized below.

Chapter 2: In this chapter I introduce the theory behind continuous-time quantum computing, describing in detail quantum walks, adiabatic quantum computing and quantum annealing and the differences between each.

Chapter 3: In this chapter, I first describe the transverse field Ising model, the model on which all optimization models seen in this thesis are mapped. I then describe both Ising spin chains and Sherrington-Kirkpatrick spin glasses, two of the models used in our research. Finally I describe the mapping of the optimisation problem, maximum independent set.

Chapter 4: In this chapter, I first introduce some of the types of possible errors that can occur on quantum computers. Then I present the error model used in our research - limited precision. I then describe several possible implementations of this error model and discuss why we chose the implementation we did.

Chapter 5: In this chapter, I first briefly discuss quantum error correction in

gate-based quantum computers. Then I describe quantum error correction for continuous-time quantum computers, detailing several of the methods. I then describe quantum annealing correction, the technique which inspired our error suppression scheme. Finally I introduce our error suppression scheme, with an illustrative maximum independent set problem example.

Chapter 6: In this chapter I discuss the numerical methods and analysis carried out in the research which makes up this thesis.

Chapter 7: In this chapter I present the numerical analysis into the optimal link strength used to connect copies in our error suppression scheme. I present these results from data collected using several different error models and discuss the differences. I then present the results from a select choice of individual instances and discuss the findings. After this, I present the results from running these tests using Ising models with no fields and discuss the findings. Finally I hypothesise on the mechanism of our error suppression scheme.

Chapter 8: In this chapter I present error suppression results when subjected to two of our error models for both 5-qubit spin glass and spin chain Ising models, for two, three, four and five copies in several configurations.

Chapter 9: In this chapter I present error suppression results in terms of precision improvement for spin glasses and spin chains, of sizes 5 to 10 qubits.

Chapter 10: In this chapter I first describe the threshold behaviour present in our results. I then discuss the theory behind this threshold for 2-qubit Ising models, and extend this to larger models. Building on this theory, I develop a link selection protocol aimed at locally predicting whether corresponding qubits in separate copies should be connected or not. I then present error suppression results on both spin glasses and spin chains after the implementation of this protocol and discuss the findings.

Chapter 11: In this chapter I first present results from analysis of optimal γ when using our error suppression method. I then present results on average success probability versus the strength of our error model. I then present and discuss the differences in average success probability between different instances. I then present average success results, when the link selection protocol has been implemented. Finally I present results of the average success probability on larger problem sizes and individual improvement cases.

Chapter 12: Finally in this chapter, I summarize and discuss the research that has been presented in this thesis.

Continuous-time quantum computing

Continuous-time quantum computing, includes adiabatic quantum computing, quantum annealing, and continuous-time quantum walks. Each of these techniques compute using the evolution of a Hamiltonian. We start with a quantum system in an initial state that is easily prepared, for example, the ground state of a simple Hamiltonian \hat{H}_0 . The system is then driven over time into the ground state of a Hamiltonian \hat{H}_P that encodes the problem to be solved. The Hamiltonian $\hat{H}_c(t)$ that carries out the computation can typically be written,

$$\hat{H}_c(t) = A(t)\hat{H}_0 + B(t)\hat{H}_P, \quad (2.1)$$

where $A(t)$ and $B(t)$ are the time-dependent control functions. These functions differ depending on which type of continuous-time quantum computing you are carrying out.

In section 2.1 we introduce and describe quantum walks, including a description of the solving the search problem in this format 2.1.1. In section 2.2 we introduce and describe adiabatic quantum computing (AQC) and how it differs from quantum walks. Finally, in section 2.3 we briefly introduce quantum annealing.

2.1 Quantum Walks

A quantum walk is the quantum analogue of a classical random walk. There are both continuous-time quantum walks and discrete-time quantum walks, but in this thesis we focus on continuous-time quantum walks, as their behaviour with our error suppression scheme are the subject of analysis in chapter 11.

Continuous-time quantum walks were first considered in [60, 61], but only shown to be capable of universal quantum computing in [62]. Continuous-time quantum walks evolve the system from the ground state of a simple easily-prepared Hamiltonian under a time-independent Hamiltonian with the constant functions $A(t) = \gamma$ and $B(t) = 1$, where γ is a hopping rate between states. γ is a key parameter for quantum walks, as it determines the relative strengths of the driving and problem Hamiltonians. Practically, when performing quantum walks it is important that the success probability is not very sensitive to the exact value of γ , because there is no existing efficient way of finding the optimal γ without solving the problem itself.

An important concept for continuous-time quantum walks, is the graph on which the ‘walk’ takes place. Three examples are: the complete graph, hypercube graph, and d -dimensional periodic lattice [61]. It was found that both the complete graph and hypercube graph achieved \sqrt{N} speed-up for the continuous-time search algorithm [60]. Whereas the d -dimensional periodic lattice achieved speed-up for $d > 4$. Using the complete graph for quantum walks can provide algorithmic benefits but they are hard to achieve in practice due to the high order X terms needed. In this thesis we focus on quantum walks which take place on hypercubes. A hypercube is an n -dimensional cube, otherwise known as a square in 2-dimensions or a cube in 3-dimensions. All possible states of a problem with n qubits can be placed on the vertices of an n -dimensional hypercube. The states are placed such that moving from one vertex to the next along an edge is equivalent to flipping a single qubit. As the number of qubits in the problem increase with n , the dimension of the hypercube also increases with n . The number of vertices in the hypercube (equivalent to the number of possible states) grows as 2^n . An example of a 3-dimensional hypercube can be seen in figure 2.1.

As we saw in the introduction of this chapter, in order to perform a continuous-time quantum computation we must prepare the system into an initial state. An example of a simple easily-prepared state is the ground state of the following Hamiltonian,

$$\hat{H}_0 = (n\hat{\mathbb{I}}_n - \sum_{j=0}^{n-1} \hat{X}_j), \quad (2.2)$$

where n is the number of qubits, $\hat{\mathbb{I}}_n$ is the identity of size n , and $\sum_{j=0}^{n-1} \hat{X}_j$ is (once calculated) equivalent to the adjacency matrix of the hypercube on which the quantum walk is taking place. This is the term which allows ‘the walk’ to pass to adjacent vertices, thereby allowing it to propagate through the hypercube. \hat{X}_j

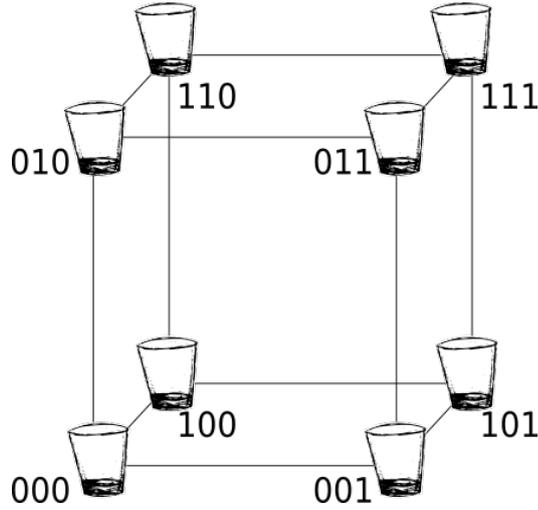


Figure 2.1: Diagram of a 3-dimensional hypercube, where travelling from one vertex to the next along an edge is equivalent to flipping a bit. The liquid in each of the cups represents the probability of being in each of the states. This diagram shows the initial equal superposition state (see equation 2.5) where each cup is 1/8th full.

can be written out as,

$$\hat{X}_j = \left(\bigotimes_{r=0}^{j-1} \mathbb{I}_2 \right) \otimes X_j \otimes \left(\bigotimes_{j+1}^n \mathbb{I}_2 \right), \quad (2.3)$$

where $\bigotimes_{r=0}^{j-1} \mathbb{I}_2$ is the tensor product from 0 to $j - 1$ of the 2×2 identity, and X_j is the Pauli matrix X on qubit j .

In order to perform the quantum walk, equation 2.2 is then substituted into equation 2.1, along with the expressions for $A(t)$ and $B(t)$. 2.1 is then substituted into the time-dependent Schrödinger equation $\hbar \frac{\partial \psi}{\partial t} = \hat{H}_\hbar \psi$ (where $\hbar = 1$) along with the initial state $|\psi(0)\rangle$, and evaluated to find the state at time t ,

$$|\psi(t)\rangle = \exp(-i\hat{H}_c t) |\psi(0)\rangle. \quad (2.4)$$

The evolution of a quantum walk of a 3-qubit system on a 3-dimensional hypercube, with an initial state of $|000\rangle$ is shown in figure 2.2.

2.1.1 Search problem using a quantum walk

The search problem using a quantum walk is equivalent to the gate based version of Grover's algorithm [60]. Here the initial state $|\psi(0)\rangle$ of the system is the equal superposition of all states on the hypercube. The general form for this state on n qubits is,

$$|\psi(0)\rangle = \frac{1}{\sqrt{2^n}} \sum_{j=0}^{2^n-1} |j\rangle. \quad (2.5)$$

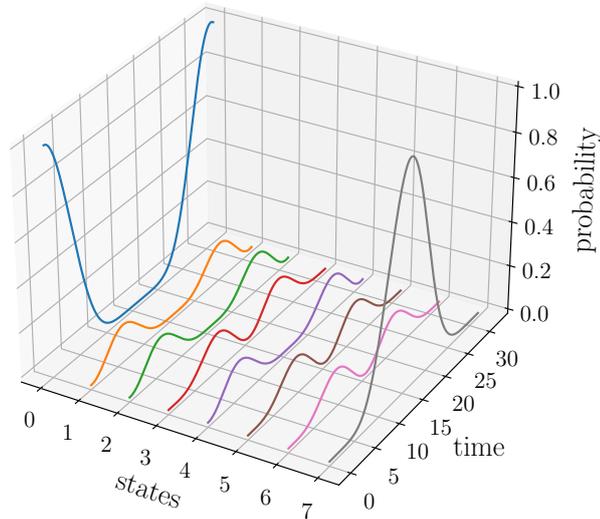


Figure 2.2: 3D plot showing the probability evolution over time of each of the eight 3 qubit states on a hypercube during a quantum walk with initial state $|000\rangle$.

This state happens to be the ground state of the Hamiltonian H_0 , so is a natural state to start in.

One of the states on the hypercube is called the ‘marked state’ and given a lower energy than the rest. This is defined in the Hamiltonian H_p ,

$$\hat{H}_p = \hat{\mathbb{I}} - |m\rangle\langle m|. \quad (2.6)$$

\hat{H}_p is then substituted, along with $A(t)$ and $B(t)$ into the overall Hamiltonian $\hat{H}_c(t)$, and is then evolved according to the Schrödinger equation. $A(t)$ is equal to γ and there is an optimal γ at which the success probability of finding the correct solution to \hat{H}_p is maximised. For quantum search the optimal γ is narrowly defined with a steep fall off [63], meaning it could be hard to extract optimal success rates in experiment. However for other problems such as SK spin glasses, optimal γ is much broader with shallower fall off, meaning is easier to extract close to optimal success rates. In [64] they discuss a way of calculating a suitable γ by using a heuristic.

Figure 2.3, shows the probability vs time for each of the possible states of 3-qubit system undergoing a quantum walk on a 3-dimensional hypercube. The initial state was the equal super-position of all possible states and the marked state was $|000\rangle$. We see here that 1. the probability of the marked state oscillates over

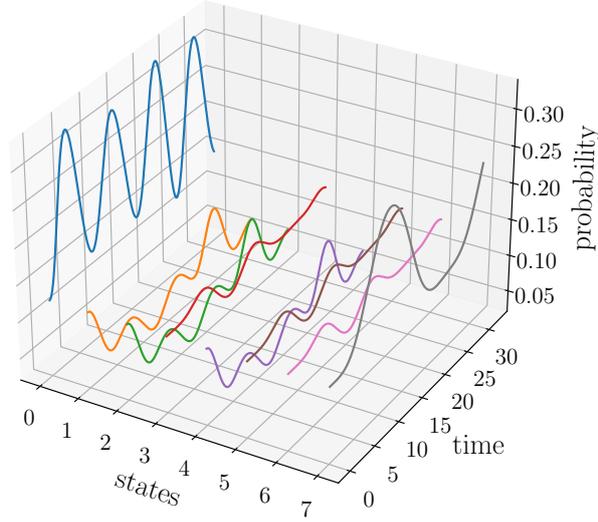


Figure 2.3: 3D plot showing the probability vs time of each of the eight possible states of a 3-qubit system during a continuous quantum walk search. The initial state was an equal superposition all eight possible states and the marked state was $|000\rangle$.

time, and 2. the probabilities of the marked state maxima are well below 1. This means that when computing using quantum walks, we need to make many repeats over many times in order to generate a high probability of successfully finding the correct ground state. As we know, the probabilities are oscillatory so to minimise the chances of always measuring when the probability is abnormally small (or large), we measure at several different times. If we repeat the measurement at many uniformly random values of time t , we can calculate the average success probability as,

$$\bar{P}(t, \Delta t) \equiv \frac{1}{\Delta t} \int_{t+\Delta t}^t dt_f P(t_f), \quad (2.7)$$

where t is the initial time, Δt is the change in time and $P(t_f)$ is the probability at time t_f which is integrated over the interval $[t, t + \Delta t]$. The long time limit of this equation is $P_{\text{inf}} = \bar{P}(0, \text{inf})$. This is a useful value as it can be calculated by numerical diagonalization of the Hamiltonian. It is given by,

$$P_{\text{inf}} = \sum_{a=0}^{N-1} |\langle E_0^{(P)} | E_a(\gamma) \rangle|^2 |\langle E_a(\gamma) | \psi(0) \rangle|^2 \quad (2.8)$$

where $E_0^{(P)}$ is the ground state, $\psi(0)$ is the initial state, and $|E_j(\gamma)\rangle$ is the eigenstate

with eigenvalue $E_j(\gamma)$ [64].

Unfortunately, the quantum search problem for quantum walks is not a realistic problem, as in order to set up the problem, you already need to know the marked state. This means it is useful as a toy problem for understanding how quantum walks work, but has no further applications. Luckily, problems do exist that can be solved by quantum walks that do not have this problem. These tend to be optimization problems, which can be mapped to Ising models a type of model which can be solved using quantum walks and other types of continuous-time quantum computing techniques. Optimization problems and Ising models are introduced and explained in chapter 3.

2.2 Adiabatic Quantum Computing

As with quantum walks, in adiabatic quantum computing (AQC), we start off with an easily prepared initial state, for example the equal superposition state, i.e. the ground state of the Hamiltonian seen in equation 2.2. We then choose our problem Hamiltonian H_p (e.g. equation 2.6), and substitute them both into $\hat{H}_c(t)$ 2.1.

Unlike with quantum walks, in AQC both $A(t)$ and $B(t)$ vary slowly and smoothly with time. $A(t) = (1 - s)$ and varies from 1 to 0. $B(t) = s$ and varies from 0 to 1. s is a time-dependent function called the *schedule* that starts equal to 0 and ends equal to 1 at a final time t_f . According to the adiabatic theorem, if the system starts in the ground state of \hat{H}_0 , if there is an energy gap $g(s) > 0$ and if s is varied slowly enough, the system will remain in the ground state with high probability and therefore finish fully in the ground state of \hat{H}_p , thereby solving the problem encoded into \hat{H}_p [65, 66]. Figure 2.4 shows probability vs time, of each of the possible states in a 3-qubit system, during an AQC quantum search, where $|000\rangle$ is the marked state.

For the AQC quantum search shown in figure 2.4, the schedule s was varied linearly. In figure 2.5, we have plotted the energy levels vs time of 3 (left) and 5 (right) qubit systems, undergoing a quantum search using AQC. Comparing the energy levels of the two, it can be seen that the gap $g(s)$ between the ground state and the next excited state is more narrow in the 5 qubit example than in the 3 qubit example. Also, $g(s)$ is smallest close to the point where the system goes from being mostly in the ground state of the initial Hamiltonian \hat{H}_0 to mostly in the ground state of the problem Hamiltonian \hat{H}_p . As the number of qubits increases this narrowing gap can mean that the time needed to vary the schedule slowly enough can become

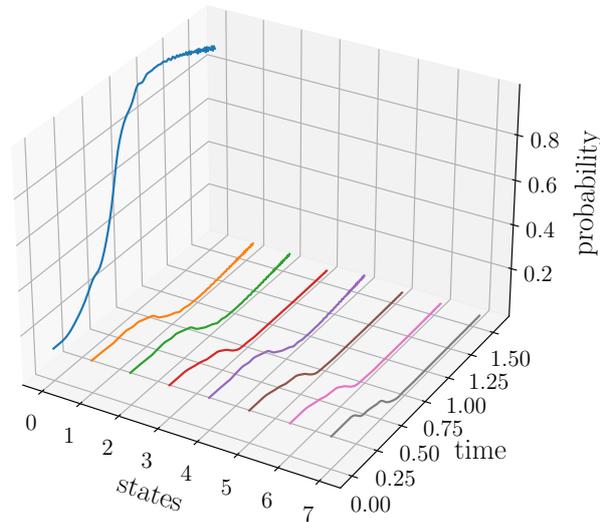


Figure 2.4: Plot of success probability vs time of each of the possible states on the hypercube, during an AQC quantum search, on 3 qubits. The initial state was the equal superposition state of all the states on the hypercube and the marked state was $|000\rangle$.

prohibitively long. (In fact in order for AQC to be possible a small probability of error must be allowed.)

In addition, it was shown in [67] that a linear schedule s does not give a quantum speed-up. In order to achieve the quadratic speed up predicted by Grover's algorithm, s must be chosen so that it varies quickly at very early and late times t , but slowly in between these times, to avoid 'jumping up' to the next excited state (and therefore finding the wrong marked state $|m\rangle$).

2.3 Quantum Annealing

Quantum annealing also evolves the system from \hat{H}_0 to \hat{H}_P , with $A(t)$ and $B(t)$ parametrized the same as for AQC. However, rather than relying on the condition of adiabaticity, in quantum annealing the system is evolved much faster and other effects such as cooling help to find the ground state [68, 69]. Terminology around quantum annealing is not universally agreed upon, but generally it is considered to encompass both cases where environmental dissipation plays a role in the com-

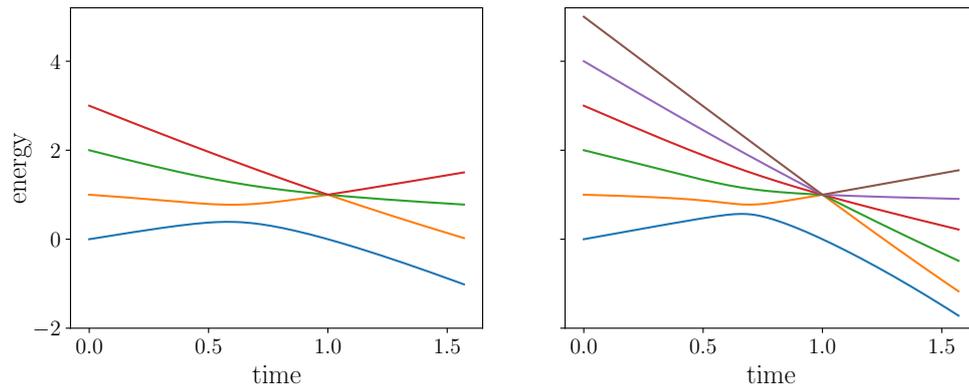


Figure 2.5: The energy levels of a 3 qubit (left) and 5 qubit (right) system on a hypercube undergoing adiabatic evolution to the problem Hamiltonian. In both cases the states have been grouped by Hamming weight.

putation, and where closed system effects dominate but evolution is faster than adiabatic (known as diabatic quantum annealing [70]).

The D-Wave machines are probably the most famous example of quantum annealers. These machines act in the de-coherent regime as open quantum systems subject to noise, which are attempting to evolve the Hamiltonian from its initial state to the problem state. This noise means many excited states (not generally the solutions we are looking for) are populated along the way, meaning the success probability at the end of a quantum anneal is unlikely to be close to one. Therefore repeats are used to boost the success probability of a computation.

Recently it was found that a coherent regime can be accessed by these machines in order to simulate closed system dynamics, if the anneal is done on extremely short time-scales [71]. The dissipation effects found in D-Wave machines can also have some beneficial effects [72], and are a key part of the popular reverse annealing technique [73]. Quantum annealing has found applications across a variety of subjects, such as scheduling and job optimization [74–76], biology [77, 78], chemistry [79, 80] and finance [81, 82], plus many more. For a more comprehensive reviews of these applications see [83, 84]

2.4 Chapter summary

In this chapter, we have introduced and described the the three forms of continuous-time quantum computing, starting with quantum walks, with the example of solving the search problem using this technique. Then we introduced and described adiabatic quantum computing and its differences to quantum walks. Finally we briefly described quantum annealing. The three continuous-time tech-

niques, quantum walks, AQC and quantum annealing, that we have described in this section, can be thought of as extremal points in a space of hybrid methods that can be interpolated between [63], to find methods suited to a particular problem and hardware.

Optimisation problems

In continuous-time quantum computing we like to solve classical optimisation problems. These are problems where the goal is to find solutions which minimise (or maximise) a certain function (called the cost function). We call the solution state which minimises (or maximises) the cost function, the minimum energy state or ground state. Pop culture examples of this type of problem are Rubik's cubes and sudoku puzzles, but they can also be found in less frivolous problems such as route planning and finance. A Rubik's cube is an optimisation problem because the end goal is to try and maximise the number of cubies (the smaller cubes that make up a Rubik's cube) that have the same colour on each of the faces of the cube, so that we end up with the six faces only having one colour of cubie (see figure 3.1). Therefore a function parametrizing this problem (its cost function) would penalise configurations of the Rubik's cube which contain faces with more than one colour of cubie. Likewise in sudoku, we are aiming to have only the numbers 1 to 9 once in each of the 3×3 boxes, rows and columns, therefore its cost function would

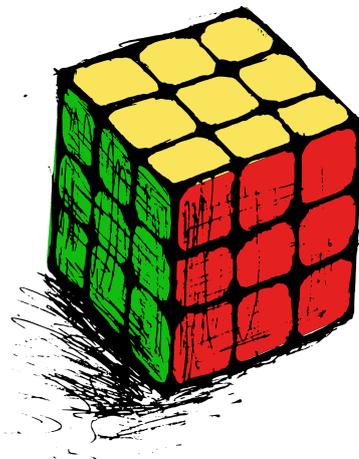


Figure 3.1: Artist's impression of a Rubik's cube in its minimum energy state.

penalise configurations of sudoku puzzle which violate these rules.

In section 3.1, we introduce the transverse field Ising model, the model to which optimization models can be mapped so that they can be solved by continuous-time quantum computers. In section 3.2, we introduce Ising spin chains, a simple (trivially solvable) version of the Ising model which has been used in many simulations described in this thesis. In section 3.3, we introduce Sherrington-Kirkpatrick (SK) spin glasses and their mapping to the transverse field Ising model. SK spin glasses are a model whose ground state is both NP-hard to find and also ‘uniformly hard’ to find. It was used for the majority of simulations and numerical analysis in this thesis. In section 3.4, we describe maximum independent set (MIS), a simple example of an optimization problem that can be mapped to the transverse field Ising model. A 5-qubit example of this model is used to introduce and demonstrate our error suppression scheme in this thesis. In section 3.5 we summarize the chapter.

3.1 Transverse field Ising Model

Classical optimization problems can be efficiently encoded [85] into an n -qubit Ising Hamiltonian \hat{H}_P of the form

$$\hat{H}_P = \sum_{j=0}^{n-1} h_j \hat{Z}_j + \sum_{j \neq k=0}^{n-1} J_{jk} \hat{Z}_j \hat{Z}_k, \quad (3.1)$$

where the fields with strengths h_j act on the j th qubit and the couplings with strengths J_{jk} act between the qubits j and k . The \hat{Z}_j operators are tensor products

$$\hat{Z}_j = \left(\bigotimes_{r=0}^{j-1} \mathbb{I}_2 \right) \otimes \hat{Z} \otimes \left(\bigotimes_{j+1}^n \mathbb{I}_2 \right), \quad (3.2)$$

which act non-trivially with a Pauli-Z operator \hat{Z} on only the j th qubit. This Ising Hamiltonian can be applied to the qubits on a quantum computer but the strengths of the fields h_j and couplings J_{jk} will be limited in range by the hardware capabilities. To model this limitation, we used a range restricted to the interval $[-1,1]$ for both the h_j and the J_{jk} of the Ising models in the numerical simulations in this thesis.

As explained in the previous section 2, we would find the ground state of this transverse field Ising model using a continuous-time technique, thereby introducing a transverse field Hamiltonian \hat{H}_0 , (which can be of the form we saw in 2.2) to the system.

Another limitation of the transverse field Ising model, is specification errors, we numerically modelled these by introducing a lack of precision in the h_j ’s and the

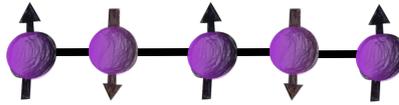


Figure 3.2: Diagram of an 5-qubit Ising spin chain, without periodic boundaries. Qubits are shown in purple. Couplings are shown in black.

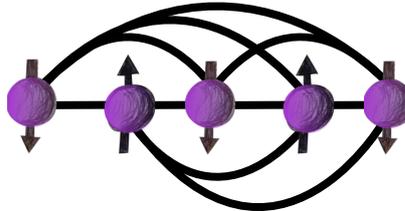


Figure 3.3: Diagram of a 5-qubit Ising spin glass. Qubits are shown in purple. Couplings are shown in black.

J_{jk} 's. We aimed to reduce their impact by the introduction of our error suppression scheme. The models we used to apply our specification errors are described in section 4. We ran numerical simulations on several different types of Ising model. Each of these models are described below.

3.2 Ising spin chains

For our initial numerical simulations and later as a tool for comparison, we used Ising spin chains. These are chains of qubits linked with couplings $J_{jk} \neq 0$ when $k = j + 1$, with a field h_j applied to each qubit. The chains used in this thesis project did not have periodic boundaries, see figure 3.2. This type of model is similar to the multi-qubit variable mappings used on D-Wave machines for minor embedding [86–88]. However, finding the ground state of these spin chains can be solved efficiently classically, and they have been shown to exhibit counter-intuitive statistical effects [89], so they are not suitable for gaining insight into the behaviour of more general types of problem Hamiltonians. Nonetheless, they are important for real hardware architectures (e.g., minor embedding), and their simplicity is useful in illustrative diagrams.

3.3 Sherrington-Kirkpatrick spin glasses

The majority of numerical simulations in this thesis were based on a data set of Sherrington-Kirkpatrick (SK) spin glass instances which were generated for research reported in [64] and made available here [90]. Finding the ground state of

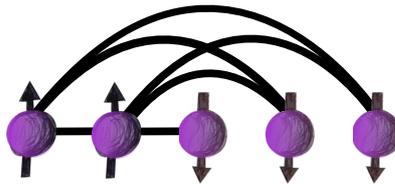


Figure 3.4: Diagram of a 5-qubit MIS example. Qubits are shown in purple. Couplings are shown in black.

SK spin glasses is NP-hard, and also ‘uniformly hard’, meaning almost all problem instances are hard at large sizes. Sherrington-Kirkpatrick spin glasses [91] are defined by,

$$H_{SK} = -\frac{1}{2} \sum_{(j \neq k)=0}^{n-1} J_{jk} S_j S_k, \quad (3.3)$$

where S_j are the classical spins ($S_j \in \{-1, 1\}$) and the couplings J_{jk} are drawn independently from a normal (Gaussian) distribution with mean zero and standard deviation σ . The qubits are all to all connected (unless one of the couplings happens to equal zero), see figure 3.3. Single-body field terms, $\sum_{j=0}^{n-1} h_j S_j$ are added to break the spin inversion symmetry, where h_j are the field strengths and are also drawn independently from the same normal distribution as the couplings. In order to map this to the quantum Ising model, the classical spin variables S_j are simply mapped to Pauli- Z operators and the problem Hamiltonian becomes,

$$\hat{H}_{SK} = -\frac{1}{2} \sum_{j \neq k=0}^{n-1} J_{jk} \hat{Z}_j \hat{Z}_k - \sum_{j=0}^{n-1} h_j \hat{Z}_j. \quad (3.4)$$

3.4 Maximum independent set

As it is an example of an optimisation problem that can be solved by a continuous-time quantum computer, we next looked at the maximum independent set (MIS) problem. This is an NP-hard problem that is well-studied in computer science/graph theory. Here, an independent set is a set of vertices in a graph G that are not adjacent (i.e. not connected by an edge). The maximum independent set (MIS) is the maximum possible size of this set for a given graph G .

To set up this problem we start with a graph G . This graph G has adjacency matrix M , where each entry $M_{kl} = 1$ if vertices k, l are connected by an edge. To maximise the number of sets, as many as possible adjacent vertices must be of different sets. In order to study this problem on a continuous-time quantum computer, we must encode this problem onto an Ising model. We may represent

vertices of different sets by the opposing qubit values $|0\rangle, |1\rangle$. Using $+J_{ij}\hat{Z}_i\hat{Z}_j$ terms this makes it energetically favourable for connected qubits to have opposing spins (if J_{ij} is positive). So to map the MIS problem to the Ising model, we set each coupling strength $J_{ij} = M_{kl}$. The Ising model will have the same structure as the MIS graph G .

In order to ensure the number of independent sets is maximised we must ensure highly connected vertices are of a different set to as many of their adjacent vertices as possible. The $+h_i\hat{Z}_i$ terms in the Ising model, energetically favour $|1\rangle$ for each qubit (if h_i is positive). (This effect opposes that of the coupling terms.) In order to maximise the chance that a highly connected qubit is in a different set to those qubits adjacent to it, we set the fields as,

$$h_i = - \sum_{j=0}^{N-1} J_{ij} + \kappa, \quad (3.5)$$

where N is the number of vertices. This penalises highly connected qubits which remain in $|1\rangle$. (Although the penalty strength is reduced by κ . κ is generally less than 1.)

We can see that each additional vertex in an MIS problem will add another qubit (and therefore coupling and field) to the Ising model. The energy scale of the MIS problem will therefore increase. In any experimental hardware there will be limitations to the strength at which fields/couplings can be set. In order to cope with the increasing energy scale, the strengths of the fields and/or couplings will have to be re-scaled down. This means that the energy levels between states will become far closer as the problem size increases and small differences between the theoretical and experimental precision are more likely to lead to errors in the ground state of the Ising model. We use a 5 vertex problem (see figure 3.4) as an example to show the effect of our error suppression method in section 5.4.1.

3.5 Chapter summary

In this chapter, we begin by introducing and describing the transverse field Ising model, the model which optimization problems are mapped to in order to be solved by continuous-time quantum computing. Then we introduced Ising spin chains, the first (trivially solvable) model on which we carried out our error suppression simulations. We next introduced Sherrington-Kirkpatrick (SK) spin glasses, on which most of the error suppression simulations in this thesis were done. This is a model which has ground states that are NP-hard and ‘uniformly hard’ to solve. Finally, we introduced the maximum independent set (MIS) problem, a

simply described NP-hard to solve optimization problem which can be mapped to a transverse Ising model. In section 5.4.1, we use this model to introduce and demonstrate our error suppression scheme.

Error models

In quantum computing there are various sources of error. In this thesis we focus on errors which result from a lack of precision in the fields and couplings of the Ising model. This type of error (if strong enough) causes a change in the ground state, resulting in us getting the wrong answer to our optimisation problem.

In section 4.1, we first look at some of the types of error that can occur on quantum computers, before introducing the limited precision errors that this thesis focuses on. There are multiple ways of modelling precision errors. Section 4.2 introduces and describes several different methods, along with the advantages and disadvantages of each. Each of the methods described were trialled at different points of this research, before the introduction of the deterministic random error model, which was used in simulations for the majority of this research. Finally section 4.3 shows a comparison of the effect of each of the error models on the measurement of fraction correct of single disconnected copies of 5-qubit SK spin glasses versus precision p .

4.1 Types of Noise/error

4.1.1 Bit flip and phase flip errors

In classical computing, bits are susceptible to bit flip errors. This is when a bit is flipped to its opposite value e.g. from 0 to 1. To ensure the information stored on this bit is protected, redundancy is created by simply copying the value of a single bit onto multiple bits. This means that even if bit flip errors do occur, the true value of the original bit can still be recovered by performing a majority vote.

Like classical bits, qubits are also susceptible to bit flip (X) errors. However unlike classical bits, they are additionally susceptible to phase flip (Z) errors, which are

a vulnerability of the extra information stored by the qubit's superposition. Unfortunately, due to the no cloning theorem, we are prohibited for directly copying qubits. Further, we cannot directly measure a qubit on which we think an error has occurred as this would cause a collapse of its wave function, thereby losing any information stored in its superposition. These issues were overcome for gate-based quantum computing with the development of quantum error correction, which is described in more detail in section 5.1. For continuous-time, quantum error correction is less well-developed, though many approaches exist, which are described in section 5.2.

4.1.2 Quantum decoherence

In order to perform the quantum mechanical operations involved in quantum computing, the quantum system must remain coherent. However when the system comes into contact with the external environment (such as when a measurement is performed) the system gradually loses its coherence. This is known as quantum decoherence, which was first proposed by H. D. Zeh in [92]. For a review on the subject see [93]. For a description of the challenges decoherence provides in quantum computing, and techniques which may be used to avoid these challenges see [94] and [95].

4.1.3 Qubit loss

Qubit loss occurs when a qubit is lost from the system. This can be due to its actual loss from the system, e.g. an ion escapes from the ion trap, or it might be excited to a state which is outside the state space which is used for computations.

There are several platform specific techniques for remedying this situation. For gate-based quantum computing [96–100], there are error correcting code based techniques which can be used against these losses [101–104]. For example for lattice based error correcting codes such as the surface code and color code, their lattices may be redrawn in order to avoid using the nodes where a qubit has been lost [105–107].

4.1.4 Limited precision

In any real hardware, there will be a limit to how precisely parameters can be set. In continuous-time quantum computing where we map optimization problems to Ising models, this can be seen as a limit to how precisely we may set the fields

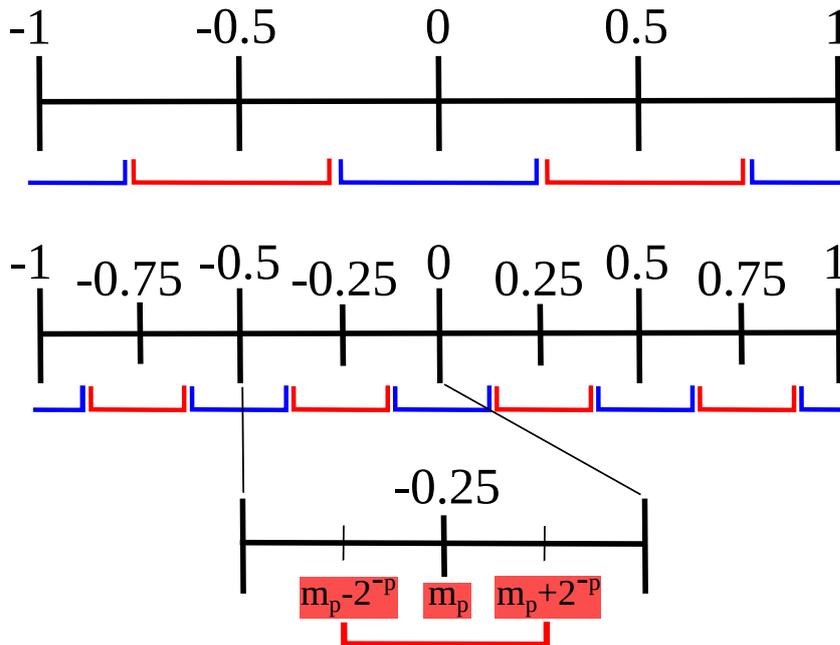


Figure 4.1: Two number lines showing the resolution available between -1 and 1, from $p = 2$ bits ($2^2 = 4$ values) and $p = 3$ bits ($2^3 = 8$ values). The uncertainty of each of the values is shown below the line in alternating red and blue delineations of the divisions. Underneath the number line for $p = 3$ bits, the area between 0 and -0.5 is expanded to show the location of the midpoint m_p and the upper and lower limits: $m_p \pm 2^{-p}$.

h_j and couplings J_{jk} . The theoretical Ising model equation (3.1), introduced in section 3.1 allows h_j and J_{jk} to be real numbers, but in practice we only have a fixed number of possible values available. Thus, we cannot in general represent the problem Hamiltonian exactly, the limited resolution of the hardware will be a potential source of error, even before we carry out the computation. This issue was first recognised in the AQC setting by Young et al. [1] and later in [2, 3].

In this thesis we choose to focus only on errors due to limited precision. This is a simple and easy to understand model that is of practical importance. Additionally a deeper understanding of one type of error is useful for designing error mitigation strategies that can handle realistic situations with multiple sources of error.

4.2 Modelling limited precision

There are multiple ways of modelling limited precision. In this section we describe a few possible modelling methods, each of which were used at some point in this research. We also detail the advantages and disadvantages to each technique.

In our model of limited precision, we quantify the number of possible values in a range $[-1, 1]$ (the resolution) by dividing this range into $2^p + 1$ values, where p is the precision. The natural way to do this for periodic boundary conditions is as shown in figure 4.1. The divisions shown below the line are labelled by their midpoints m_p , and zero is one of the midpoints. The endpoints of the division with the midpoint of zero are $\pm 2^{-p}$, so are equal to ± 0.25 with $p = 2$ for the upper diagram of figure 4.1 and ± 0.125 with $p = 3$ for the lower diagram. The precision p is thus the number of classical bits needed to represent the possible values in the range $[-1, 1]$.

The reason this is the natural setting for periodic boundary conditions is because in situations where the measurements ‘wrap’ back on themselves (like a dial) we still get a full division for each value represented. However for non-periodic boundary conditions we are left with “half divisions” at either end of the scale, which are harder to deal with.

In real hardware situations, we would not have periodic boundary conditions for the fields and coupling strengths h_j and J_{jk} . The natural way to divide a range with non-periodic endpoints is to shift the divisions in figure 4.1 so that the midpoints are $\pm 2^{-p}$, etc. However, this means there are two divisions that could represent approximately zero, and for reasons which will become clear in chapter 7, we prefer to have only one division representing approximately zero. Hence, we use the model in figure 4.1, but with two “half divisions” at each end.

4.2.1 Random error model

The first error model for simulating a lack of precision is the random error model. In this model we apply a random error to each of the fields and couplings in the Ising model and then measure its ground state. We then repeat this process n times, taking n error samples, where n is usually 10 unless otherwise specified. In general, the ground state will only be incorrect for some of these error samples, so the fraction correct we calculate is an average over the repeats and the instances. Equivalently, we can look at the fraction correct per instance here as the probability of the single copy being correct or incorrect. Looking at it this way, we can see that as long as there is some probability that the ground state is correct, using extra disconnected copies to do repeats can help to find the correct ground state. If we assume completely uncorrelated errors, we can find the fraction correct for three disconnected copies $P(3)$ (the same as three repeats of a single copy) using,

$$P(3) = 1 - (1 - P(1))^3, \tag{4.1}$$

where $P(1)$ is the fraction correct for the single copy.

When using hardware with limited precision, if we were given a parameter too precise to set, we would have to instead approximate it by setting the parameter to m_p , the midpoint of the division that its value lies within. This means that the actual value obtained in the hardware could be anywhere between $m_p \pm 2^{-p}$, with some probability distribution. More generally, if we set a value x such that $m_p - 2^{-p} \leq x \leq m_p + 2^{-p}$, we instead obtain y such that $m_p - 2^{-p} \leq y \leq m_p + 2^{-p}$ with probability $P(y|x)$. If $P(y|x)$ is the uniform distribution on the interval $m_p \pm 2^{-p}$, the average error in x scales quadratically in x , as $\varepsilon_u = 2^{p-1}\{(x - m_p)^2 + 2^{-2p}\}$, ranging from $1/2 2^{-p}$ for $x = m_p$ to 2^{-p} for x at the endpoints of the division, with the average over x being $\langle \varepsilon_u \rangle_x = 2/3 2^{-p}$.

The advantage of using this random error model is that it closely follows the process that would occur on real hardware. However, it is unlikely that the error distribution $P(y|x)$ would be uniform, so this would introduce small errors here. Also using this model, we need to take several error samples at each parameter value in order to approximate the error distribution and then average over it. This makes numerical simulation more intensive and means we are able to analyse fewer instances in total.

4.2.2 Deterministic mid-point error model

A way of reducing the intensity of numerical simulations is to use a deterministic error model. Instead of averaging over many randomly chosen y for each x we aim to set, we chose a single y for each x at precision p . This also means that using repetitions will not increase the probability of a successful computation. If the chosen y gives the wrong ground state the first time round, it will continue to be wrong even if repeats are carried out. This allows us to separate the effects of repeat runs from the effects of our error mitigation strategies.

A natural model is to set the chosen y to be the midpoint m_p of each division. We call this model the deterministic mid-point error model. The error in the deterministic mid-point model scales linearly, $\varepsilon_m = |x - m_p|$, ranging from zero at the midpoint, to 2^{-p} at the endpoints, with the average over x being $\langle \varepsilon_m \rangle_x = 1/2 2^{-p}$. This underestimates the error compared to the random error model, especially when the exact x is close to m_p , but gives a similar error rate overall at each precision (see figure 4.2).

The advantage of this error model compared to the random error model is that it decreases the intensity of the numerical simulation for each instance, meaning more instances can be analysed. However, especially when the precision p is low, using the mid-point model means there is a high likelihood of terms within the

Hamiltonian cancelling. This leads to a high level of degeneracy in the ground state of the Ising model instance with error. This high level of degeneracy is not realistic for experimental hardware, where for example in flux qubit quantum annealers, the programmable coupling values are not evenly spaced [108]. The high level of degeneracy also causes problems in numerical simulation, if the degenerate ground states include the correct solution.

4.2.3 Deterministic random error model

In order to avoid the high levels of degeneracy found in the deterministic mid-point error model, we instead developed the deterministic random error model and used this in the majority of the analyses in this project. In the deterministic random error model, we set the chosen y to be a random value y_r in each division. Here the error also scales linearly, $\varepsilon_r = |x - y_r|$, ranging from zero when $x = y_r$ to $y_r - (m_p - 2^{-p})$ at the lower endpoint and $(m_p + 2^{-p}) - y_r$ at the upper endpoint of the division. The average over x becomes $\langle \varepsilon_r \rangle_x = 2^{p-1}(y_r - m_p)^2 + 1/2 2^{-p}$. When $y_r = m_p$ this reduces to $1/2 2^{-p}$, meaning that as long as y_r is not less than or equal to m_p , the average error in x will be larger than $1/2 2^{-p}$. This suggests that this error model will not underestimate the error as strongly as the deterministic mid-point error model and therefore have a more similar error rate (see figure 4.2).

The advantage of this error model is that we retain the reduction in numerical simulation intensity gained from using the deterministic mid-point error model, whilst lowering the chance of degeneracy in the ground state with limited precision. Hence, we use the deterministic random error model, for most of the simulations in this project. Real hardware is likely to have a non-uniform distribution for $P(y|x)$, so neither the deterministic random nor the deterministic mid-point error model is preferred on realistic grounds.

4.3 Effect of the error model on single copies

We illustrate the difference between the error models described in this chapter in figure 4.2, by calculating the fraction of instances with correct ground states (known as the **fraction correct**) when rounded to a given precision p .

The data plotted have been computed using 5 qubit SK spin glass instances from instances from dataset [90]. Each of the 10^4 instances was subjected to each of the error models, which generated an approximated Hamiltonian in each case. The ground state of each approximated Hamiltonian was then compared to the ground

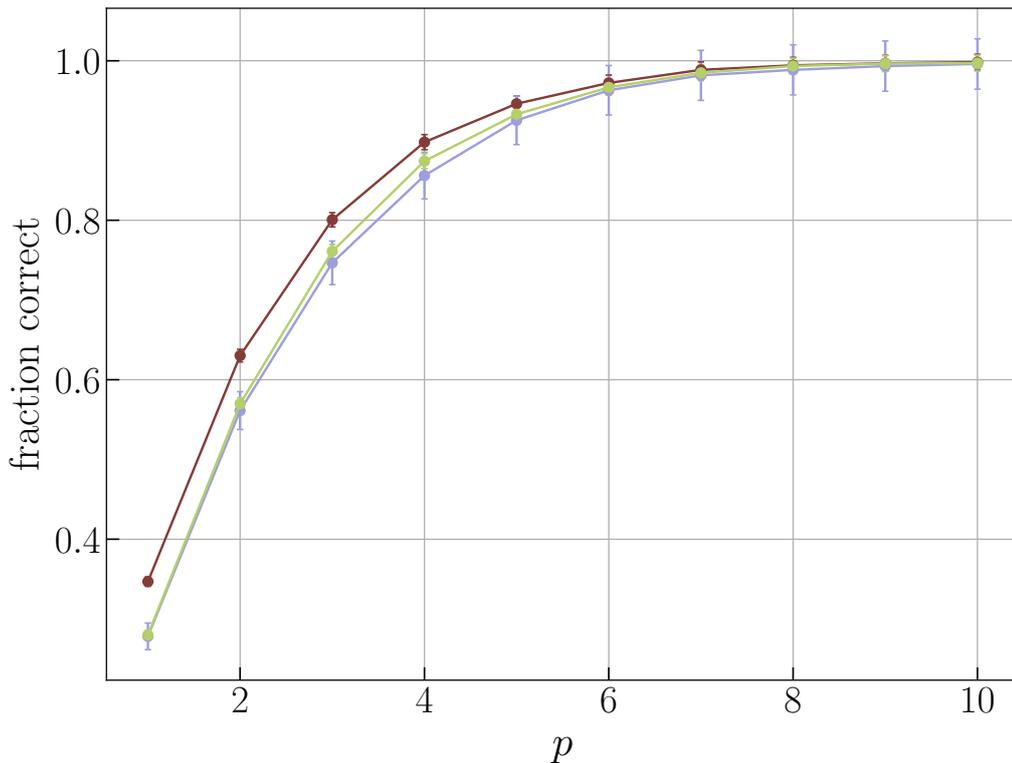


Figure 4.2: Fraction of instances with the correct ground state vs precision p for 10^4 instances of 5 qubit SK spin glasses. Precision was implemented first using the random error model (blue), then the deterministic mid-point error model (red) and finally using the deterministic random error model (green).

state of the exact Hamiltonian. The fraction of instances which retained the ground state of the exact Hamiltonian (the ‘correct’ ground state) were measured, and plotted in figure 4.2. We called this fraction the **fraction correct**. This process was repeated using all three of the error models: the random error model (blue), the deterministic mid-point error model (red) and the deterministic random error model (green) across each of the 10^4 instances, for each integer value of precision from $1 \leq p \leq 10$.

As predicted, the deterministic mid-point error model found the correct ground state slightly more often than the random error model and the deterministic random error model at low values of p . However, the differences between the deterministic mid-point error model, deterministic random error model and random method are small, and the qualitative behaviour is the same for all three models.

Figure 4.2 also shows the typical size of the errors introduced by reducing the precision. For the small sizes we study (up to 9 qubits, limited by computational capabilities), to break more than 5% of the instances requires reducing the precision

to around $p \lesssim n$ where n is the number of qubits. This makes sense intuitively, because n qubits can only represent 2^n different outcomes, meaning that the average gap between adjacent energy eigenstates will be of order 2^{-n} , and errors smaller than this are unlikely to change the ground state.

4.4 Chapter Summary

In this chapter, we first briefly described some of the noise and error types that occur in quantum computing as well as introducing limited precision errors, the error type on which this thesis focuses. We then described three ways of modelling limited precision: the random error model, the deterministic mid-point model and the deterministic random error model. We explained how these models would be implemented, analysed theoretically their error distribution and average error, and finally described their advantages and disadvantages. Next, we measured fraction correct versus precision in the range $1 \leq p \leq 10$, for 10^4 instances of single copies of 5-qubit SK spin glasses, implementing precision with each of the three error models. We found the results to be consistent with the theoretical error distributions of each described in the previous section. Finally, we commented on the reduction in fraction correct caused by the precision errors and discussed whether this reduction was consistent with theory.

Quantum error correction

As we saw in the previous chapter, quantum computers are susceptible to errors. In order to be able to perform computations on quantum computers and be confident of our answers, we are going to need error correction. For gate-based quantum computing a framework for how this could be done, has already been established, with the development of fault-tolerant quantum error correction in [109, 110] and its scalability in [111, 112]. For other types of quantum computing such as continuous-time quantum computing, error correction is not so well-developed, and it is subject to several caveats [113–115]. Despite these challenges, a variety of schemes have been established, which may be generally referred to as Hamiltonian error suppression [70]. Our error suppression method is originally inspired by one of these continuous-time error correction schemes known as quantum annealing correction (QAC). However, our method has several key differences and innovations compared to QAC, making it a distinct method.

In section 5.1 we first briefly look at quantum error correction for gate based quantum computing. Then in section 5.2, we look at quantum error correction and its continuing development for continuous-time quantum computing. In section 5.3, we focus on the QAC technique, describing its set-up and functioning. In section 5.4 we describe the set-up and functioning of our error suppression regime and in section 5.5 we summarize the chapter.

5.1 Quantum error correction for gate based QC

In order to get an idea of how we protect the information on a quantum computer, it first makes sense to look at how information is protected on classical computers. Information on a classical computer is stored as strings of bits. These bits are susceptible to bit flip errors, which is when a bit is accidentally flipped to its opposite

value, e.g. from 0 to 1. If these bits are unprotected, this means the information that this string of bits was encoding is now lost. To protect the information from this type of error, redundancy is created by copying the value of one bit onto multiple bits. This means that if a bit flip occurs on one of these bits, the value of the original bit can still be recovered by majority vote. By increasing the redundancy we can protect a bit string from a larger number of bit flip errors. This type of classical error correction is known as an error correction code (ECC) and was pioneered by Hamming [116] in the early 1950s, while working at Bell labs. These types of codes are still used today in areas such as data transmission and storage.

In quantum computing, instead of bits, we use qubits, these can be in $|0\rangle$, $|1\rangle$ or a superposition of the two $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. Similarly to bits, qubits can also suffer from bit flip errors, however they are also susceptible to phase flip errors. Due to the no cloning theorem, we can't simply copy the value of one qubit on to one or more extra qubits (as you would do in classical computing). Also, we can't directly measure the qubits we'd like to protect as by doing this the qubit must 'decide' whether it is in a $|0\rangle$, or $|1\rangle$ state and we thereby lose the information stored in its superposition.

For gate based quantum computing, a framework for overcoming these obstacles was established with the development of fault-tolerant quantum error correction in [109, 110] and its scalability in [111, 112]. Instead of copying, redundancy is created by entangling multiple qubits together. To detect errors, rather than directly measuring the qubits on which the information is stored, ancilla qubits are connected to these qubits in a particular way. Measuring these ancilla qubits tells us about the parity between qubits. We call this the 'error syndrome'. By looking at this error syndrome, the qubit (or qubits) on which an error has occurred may be found and then corrected.

Fault-tolerant error correction methods have also been adapted to measurement-based quantum computing [117] using techniques such as lattice surgery [118]. Many refinements and improvements continue to be developed in the gate model setting (for a review, see [119]). In particular, the large number of physical qubits required per logical qubits is a large obstacle to the future development of fault-tolerant gate-based quantum computers. Hence, much research is devoted to the reduction of physical to logical qubit number.

5.2 Quantum error correction for continuous QC

Scalable fault-tolerant QEC for continuous-time quantum computing has yet to be established and is subject to several caveats [113–115], in particular, that two-local commuting Hamiltonians are not sufficient for constructing ground subspace encodings [115]. Despite the challenges, a variety of schemes for error correction and suppression in continuous-time quantum computing have been developed, generally referred to as Hamiltonian error suppression [70]. Most can be grouped into categories: energy penalty Hamiltonians [120–123]; dynamical decoupling [124–126]; subsystem codes [127–131]; continuous-in-time techniques [113, 132, 133]; via qubit ensembles [134]; the Zeno effect [135]; and QAC [1, 136–145]. In addition to these techniques for explicit error suppression and/or correction, quantum annealing may be carried out in some circumstances without error correction, as long as sufficiently many repetitions are implemented [146–150].

5.3 QAC

Quantum annealing correction (QAC) is the error correction technique that inspired the error suppression scheme that is the subject of this thesis. QAC was first introduced in [1, 137] and is closely related to the energy penalty form of HES. QAC focuses on the implementation of quadratic unconstrained binary optimization (QUBO) problems on to generalised Ising model Hamiltonians on a graph. The Hamiltonians are of the form,

$$H_Q = \sum_{i \in \nu} h_j Z_j + \sum_{\langle j, k \rangle \in \varepsilon} J_{jk} Z_j Z_k, \quad (5.1)$$

where ε was the set of edges in the graph and ν the set of vertices in the graph.

For the error model, like for our error suppression scheme, the Ising models are subject to a lack of precision in the fields h_j and couplings J_{jk} . Here, the lack of precision is modelled as an additional error Hamiltonian of the form,

$$\Delta H = \sum_{\langle i, j \rangle \in \varepsilon} \epsilon_{ij}^{(J)} Z_i Z_j + \sum_{i \in \nu} \epsilon_i^{(h)} Z_i. \quad (5.2)$$

In both QAC and our error suppression method, in order to protect the ‘correct’ ground state of the exact Ising model problem Hamiltonian, several copies are used at the same time. The corresponding qubits in each of the copies are linked to each other with a coupling strength J_F .

The Hamiltonian at precision p for the system of C Ising model copies, each with a set of corresponding qubits connected according to a graph G can be written as

$$\hat{H}_p^{(C)} = \sum_{c=0}^{C-1} \left\{ \sum_{j=0}^{n-1} h_j^{(p)} \hat{Z}_j^{(c)} + \sum_{j \neq k=0}^{n-1} J_{jk}^{(p)} \hat{Z}_j^{(c)} \hat{Z}_k^{(c)} \right\} - \sum_{j=0}^{n-1} \sum_{c, c' \in G} J_F^{(p)} \hat{Z}_j^{(c)} \hat{Z}_j^{(c')}, \quad (5.3)$$

where $c, c' \in G$ means copies c and c' , and corresponds to endpoints of an edge in graph G .

In QAC, the C copies are connected together in chain or grid structures via strong ferromagnetic links $J_F^{(p)} > 0$. The final term in the Hamiltonian 5.3, makes it energetically favourable for the qubits in each of the copies to align. It is therefore energetically unfavourable for qubits that are linked together between copies to have opposite spins. This scheme looks similar to the repetition codes used in gate-based quantum computing or classical error correction.

For the origin exact Hamiltonian, all of the copies in the system will thus have the correct ground state, i.e. the links between copies will not introduce any extra errors. Ferromagnetic links provide an energy barrier which suppresses bit-flip errors. However, if a bit-flip error does occur in one or more copies, the ferromagnetic links tend to propagate this error to all the other copies, and to neighbouring qubits within copies, potentially making them all incorrect.

Similarly to the repetition codes in gate-based quantum computing, in QAC, after the initial computation has taken place, there is a further decoding step, which aims to detect and correct any error that may have occurred during the computation. Correction of errors is most often carried out via majority vote [2, 136, 138–144], but ‘coin tossing’ or energy minimization may also be used [137].

QAC may be split into three further subtly different types: penalty-QAC, simple-QAC and nested-QAC. Penalty-QAC defines specific qubits as penalty qubits. These qubits provide an energy penalty to those physical qubits that suffer an bit flip error, thereby making it energetically unfavourable [2, 136–138, 140–142, 151].

The other two types of QAC do not define penalty qubits but instead define their energy penalty through the connection of the physical qubits within the logical qubit. Simple-QAC uses a chain or grid-like structure of two or more qubits to create a logical qubit [1, 137]. Whereas nested-QAC uses a complete graph of ferromagnetic links between physical qubits to define a logical qubit [2, 139, 143–145, 151].

Unlike penalty-QAC and simple-QAC, nested-QAC is capable of encoding arbitrary Ising models and can suppress effective temperature in a scalable manner [151]. On the other hand, penalty-QAC is the only type of QAC which improves the scaling of time-to-solution for random Ising-models [138]. All three types of QAC have been shown to improve performance across the four generations of DWave. [2, 136–139, 141, 143, 145].

Anti-ferromagnetic chains were studied in [136] and found to be particularly challenging for QAC, due to the domination of domain wall errors in this type of problem. Minor embedding employs chain like structures and hence reduces the performance of QAC.

5.3.1 Summary of the mean field analysis of QAC

In [140] a mean field analysis was carried out in order to determine the free energy of penalty-QAC. The analysis showed that, in the thermodynamic limit, depending on the strength of the penalty term, for simple Ising or p -spin ferro- or anti-ferromagnetic problem Hamiltonians, penalty-QAC can either soften (when the temperature T is greater than zero) or prevent (when T tends to zero) closing of the minimum energy gap. For nested-QAC, further analysis in [143, 151], showed this type of QAC can reduce the effective temperature by a factor of C^2 , where C is the number of copies. Specifically they found $\lambda \rightarrow \lambda$, $\beta \rightarrow \beta C^2$, $\Gamma \rightarrow \Gamma/C$. This is equivalent to increasing the effective energy scale, both of which mitigate against thermal and other environment-induced errors. The steps of the mean field analysis are outlined in the following section.

They start by defining the partition function of QAC, and separating the Ising model and transverse field terms using the Trotter-Suzuki formula,

$$Z = \exp(-\beta(H^x + H^z)) = \lim_{M \rightarrow \infty} (e^{-\frac{\beta}{M}H^x} e^{-\frac{\beta}{M}H^z})^M = \lim_{M \rightarrow \infty} Z_M, \quad (5.4)$$

where H^x is the transverse field Hamiltonian and H_z is the Ising model Hamiltonian. They then write the partition function as,

$$Z_M = \sum_{\{\sigma^z\}} \langle \{\sigma^z\} | (e^{-\frac{\beta}{M}H^x} e^{-\frac{\beta}{M}H^z})^M | \{\sigma^z\} \rangle, \quad (5.5)$$

where $\sum_{\{\sigma^z\}}$ is the sum over all possible $2^{(C+1)N}$ spin configurations in the z basis and $|\{\sigma^z\}\rangle = \otimes_{i=1}^N \otimes_{k=0}^K |\sigma_{ik}^z\rangle$. They then introduce M copies of the identity operator closure relations, each labeled by the Trotter time α and introduce this into the partition function Z_M , so that they can write H^x and H^z as c-numbers,

$H^z(\alpha)$ and $H^x(\alpha)$. Where,

$$H^z(\alpha) = - \sum_{k=1}^K [N(S_k^z(\alpha))^2 + \gamma H_k^P(\alpha)]. \quad (5.6)$$

They next introduce auxiliary Hubbard-Stratonovich fields $m_{k\alpha}$ and $m'_{k\alpha}$, in order to rewrite the p -body (we only consider $p = 2$ in this thesis) interactions in terms of one-body interactions, by successively using the elementary δ function identities,

$$f(a) = \int_{-\infty}^{\infty} f(m_{k\alpha}) \delta(m_{k\alpha} - a) dm_{k\alpha}, \quad \delta(m_{k\alpha} - a) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{i(m_{k\alpha} - a)m'_{k\alpha}} dm'_{k\alpha}, \quad (5.7)$$

and subbing them into Z_M , with $a = S_k^z(\alpha)^2$.

They then use the static approximation, which allows them to drop the α 's, i.e. $m_{k\alpha} \mapsto m_k$ and $m'_{k\alpha} \mapsto m'_k$, and make a change of variables to $m'_k = \frac{N}{M} \tilde{m}_k$ and then find another expression for Z_M . They then take $M \rightarrow \infty$ and go from the classical system back to the quantum system. They then find the saddle points of Z . They inserted the saddle point equations back into the expression for Z and find the free energy F , by using that it is defined by,

$$F = -\frac{1}{\beta N} \ln(Z). \quad (5.8)$$

By taking the thermodynamic limit ($N \rightarrow \infty$) on the expression they find, they read off the following scaling behaviour,

$$\lambda \rightarrow \lambda, \quad \beta \rightarrow a^2 N^{2b}, \quad \Gamma \rightarrow \Gamma / (aN^b), \quad (5.9)$$

where $C = aN^b$, where $a, b > 0$ are constants. These are the same expressions as stated in the first paragraph of this section. The full derivations can be seen in the supplementary material of [140] and the appendices of [142, 151].

As this derivation is based on the Ising model of QAC, rather than the decoding mechanism, it can be argued that a similar derivation could be applied to our method, as long as the anti-ferromagnetic links rather than ferromagnetic links connecting copies are accounted for.

5.3.2 Disadvantages to QAC

One of the main disadvantages to QAC, is that it requires strong ferromagnetic links between copies. This means that when scaling to larger problem sizes, constraints are put on the precisions of variables in the problem Hamiltonian due to dynamic range limitations as was recognised in [144]. These scaling requirements

were improved in a series of papers which developed the mean field analysis which is described above [139, 140, 142–144, 151] and derived an effective temperature reduction of T/C for penalty-QAC and T/C^p for nested-QAC [151]. However, these results still rely on adjusting for the optimal ferromagnetic strength, which is still not optimal for problems that use large amounts of the dynamic range of the hardware.

5.4 Our error suppression method

For our error suppression scheme, similarly to QAC, we focus on protecting the original ‘correct’ ground state of an exact Ising model problem Hamiltonian which has been subjected to a lack of precision in its fields h_j and couplings J_{jk} . We focus on three error models: the random error model, the deterministic mid-point model and the deterministic random error model. All three of these techniques were introduced and described in section 4.2. Since Ising Hamiltonians and the problems encoded in them are classical (solutions are represented by computational basis states), only bit-flip errors need to be considered to protect the problem Hamiltonian. Phase errors may become important in the dynamics when the driver Hamiltonian \hat{H}_0 , which drives the quantum walk or adiabatic evolution (see section 2), is used, but are not relevant to the error models in the setting we considered during this research.

As was shown in figure 4.2, in section 4.3, across all the error models, for precision $p \lesssim n$, the ground state changes from the original correct state to a different incorrect ground state in a significant fraction of the spin glass instances. This means that, similarly to the error model used in QAC research [1, 137], a computation process that successfully finds the ground state does not necessarily always solve the original problem. Using our error suppression scheme, we aim to find a way to reduce the number of times we don’t solve the original problem. We aim to do this by using several copies at the same time, like a repetition code.

When using the random error model, by repeating the computation, we improve our chances of finding the original correct ground state. This is because sometimes the erred h_j ’s and J_{jk} ’s will be close enough to their original values as to prevent an error in the ground state, whereas other times they will not be. However, in the deterministic mid-point and deterministic random error models, the error is deterministic. Repeating the computation will not change the value of the erred h_j ’s and J_{jk} ’s and how close they are to their original values. Therefore stand-alone repeats of the computation will not help, unless there are degenerate ground states that include the correct solution. This means that using the deterministic

mid-point or deterministic random error model provides a tougher test for our error correction methods: none of the improvement we see can be due to repetition alone.

Although originally inspired by QAC, we introduce several innovations in our error suppression scheme. First, instead of connecting copies with strong ferromagnetic links $J_F^{(p)} > 0$, we instead connect copies with weak anti-ferromagnetic links $J_F^{(p)} < 0$. This changes the final term in the Hamiltonian 5.3, so that it makes it energetically favourable for connected corresponding qubits in different copies to be anti-aligned. This means that, even for the exact Hamiltonian (no precision errors on h_j 's or J_{jk} 's) one or more of the connected copies is likely to have the incorrect ground state. However, these qubits, if anti-aligned, still provide an energy barrier which suppresses bit-flip errors. Also, if a bit-flip error occurs in a single copy, the anti-ferromagnetic links in the linked copies can prevent the error occurring on all copies, making it more likely a correct copy persists despite the bit-flip errors.

We note that a consequence of our first innovation is that it is no longer likely that all connected copies in a system will have the original correct ground state even when a precision error has not yet occurred. However, we also note that in the setting where we are solving hard classical optimization problems, we only need one of the copies to provide a correct answer. This is because it is efficient to compare the quality of the candidate solutions, and select the best one (i.e. the one with the lowest energy). Therefore unlike QAC we do not perform a decoding step after computation. Post computation, we would compare candidate solutions and select the best one, as this is an efficient process. This is the strategy employed in practical implementations using multiple runs, regardless of any error correction techniques used. This fundamentally changes the criteria for evaluating the performance of the error correction strategies, and thus identifies different strategies that work well. This is also why unlike QAC we do not describe our scheme as error correction (there is no way of confirming we have the correct answer after each computation).

The work in this thesis investigates how and when it is effective to use anti-ferromagnetic links between copies to reduce bit-flip errors, and estimates the level of error reduction achieved. Due to computational limitations, we focus on $C = 2$ and $C = 3$, but also briefly look at $C = 4$ and $C = 5$. For $C = 2$, unless we are allowing double links, there is just one way of connecting the two copies. Just one $J_F^{(p)}$ link connects each pair of qubits in the two copies, see the top left diagram of figure 5.1. For $C = 3$, we may either connect linearly (as for two copies), see bottom left diagram of figure 5.1 or in a loop (triangle) such that the first copy connects to the last, see bottom right diagram of figure 5.1. For $C = 4$ and $C = 5$, the space of possible configurations rapidly increases. In both cases we consider only the configuration where the copies are connected in a loop (square, pentagon

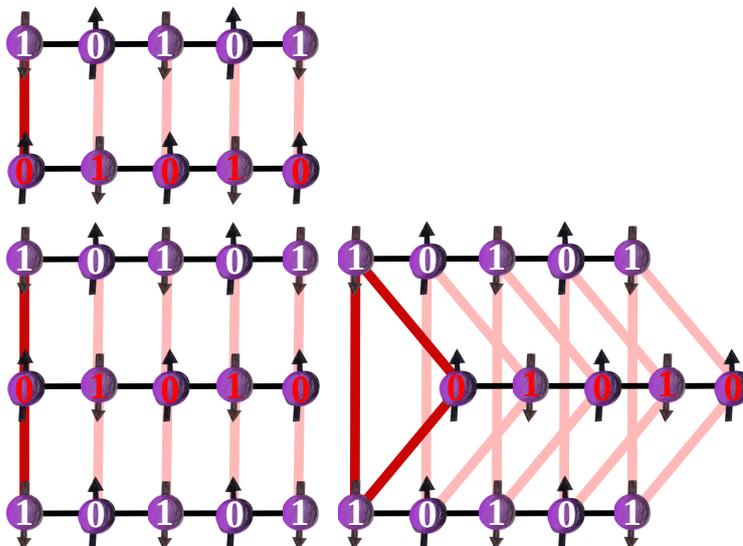


Figure 5.1: Diagrams of two and three copies of a 5 qubit Ising chain (thin black links) connected with anti-ferromagnetic links (red/pink links). Two copies in a chain (left) and three copies in a loop (triangle) (right). The red links identify the corresponding copy qubits on the left end of the chain; the pattern continues with the pink links for the other four qubits.

respectively) in this work.

An important thing to note is that when copies are connected in a linear configuration, the corresponding copy qubits are able to simply anti-align, in order to satisfy the anti-ferromagnetic links. This is also true for an even number of copies in a loop. However for an odd number of copies in a loop (e.g. $C = 3$), one of the three links connecting the corresponding copy qubits must always be unsatisfied. This thereby introduces frustration into the system, which can allow bit-flip errors to be shared more evenly across the copies, potentially increasing the probability of the presence of the original correct ground state. From our results in the following chapters we find this frustration improves the suppression of errors further.

5.4.1 An MIS example which retains the original ground state

To illustrate the error suppressing effect of connecting copies with anti-ferromagnetic links according to our scheme, we present table 5.1, which shows the ground state versus precision (from 1-4) of a 5-qubit maximum independent set problem, which is first simulated as a single disconnected copy and then three anti-ferromagnetically connected copies. This error suppressing MIS instance was first analysed and brought to my attention by Tom O’Leary. We used the deterministic mid-point error model in this example.

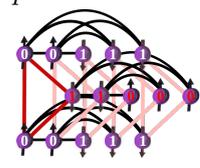
	$p \geq 4$		$p = 2$	
p		# ✓		# ✓
1	00000	0	00010, 11101, 00010	0
2	10100	0	00111, 11000, 00111	2
3	00000	0	00111, 10000, 00111	2
4	00111	1	00111, 10000, 00111	2

Table 5.1: Table comparing the numbers of correct ground states, for a single copy of a 5 qubit MIS Ising problem versus three connected copies, at precisions 1, 2, 3, and 4. For each ground state, correct qubits are shown in black and incorrect qubits in red.

The graph of the MIS problem is shown in the top row, second column, in its single copy form and in the top row fourth column, in its three connected copies form. Couplings are represented by the black connections between the qubits and a qubit in state one (zero) corresponds to a vertex which is (is not) in the independent set. The single copy is shown in its $p \geq 4$ ground state and the three copies are shown in their $p = 2$ ground state. The description of the MIS problem and its encoding onto an Ising model can be found in section 3.4.

The ground states of the single disconnected copy for $p = 1$ to 4 are shown in the second column of table 5.1. Correctly valued qubits are coloured black, whereas incorrectly valued qubits are coloured red. The number of correct copies present at each value of p is indicated in the third column. Here, we see that the ground state is incorrect for $p = 4$.

The fourth column of table 5.1 shows the ground states of three copies which have been connected with $-J_F^{\min}$ strength anti-ferromagnetic links, where J_F^{\min} is the smallest possible non-zero value of J_F allowed at that precision (see chapter 7, for more explanation). The data is presented in the same way as for the single copy, in the fourth and fifth columns. Here we see that except for $p = 1$, where all copies have the incorrect ground state, for $p > 1$, two of the three connected copies have the correct ground state. This shows that connecting the three copies anti-ferromagnetically allows for a lowering of two bits of precision, whilst still being able to find the correct ground state, for this example. This indicates an improvement in the robustness of the ground state to a lack of precision by using our scheme.

This example shows that our error suppression method works in principle. Even with some incorrect copies, it is easy to check which of the three bit strings provides

the best candidate solution to the MIS problem, by calculating (classically) the energy with the exact Hamiltonian.

5.5 Chapter summary

In this chapter, we first described how the ideas from classical error correction were re-applied in a quantum setting in the development of quantum error correction (QEC) for gate-based quantum computers as well as its current challenges. In the next section, we introduced the currently developed options for error correction and suppression in continuous-time quantum computing. Next, we introduced and described QAC. Then we briefly mentioned its disadvantages, which our error suppression scheme aims to overcome. After this, we introduced our error suppression scheme. We described its differences to QAC and how we expected our scheme to be able to suppress precision errors. Finally we presented an example of an MIS problem, which had an error in its ground state successfully suppressed by our scheme.

Numerical Methods

All of the code for the simulations and numerical analysis in this thesis that was written by me, was written in Python3 [152] from the anaconda distribution [153] and conda virtual environments were utilised. The majority of the code was initially written in the Pycharm IDE [154], however in the latter stages of research, code was written and edited using the VS code IDE [155], which has very useful remote explorer capabilities. Source control was done using Git [156], combined with a remote repository of all code written by me on Github [157].

Much of the numerical analysis was done using the NumPy [158] and SciPy [159] packages and much of the data was stored using the dataframes of the Pandas [160] package. All graphs plotted for the research in this thesis were done so using the Matplotlib [161] package. Numerical simulations were run on desktop workstations and on HPC facilities based at Durham University, and at Oxford University via the UK Quantum Technology Quantum Computing and Simulation Hub.

The research in this thesis focuses on the suppression of errors caused by a lack of precision in the fields h_j and couplings J_{jk} of an Ising model which has an optimization problem encoded onto it. These errors essentially change the ground state of the Ising model leading to us finding the incorrect answer to our optimization problem. The simulation of precision errors was a key part of the numerical methods in this research. How these errors were simulated is described in section 4.2.

In section 6.1, we describe how we simulated classical dynamics. In section 6.2 we describe how we carry out the quantum dynamics described in section 11. We next look at problem instance generation in section 6.3. Finally in section 6.4, we look into how we did fitting in a key part of our research and how we carried out error analysis across this research.

6.1 Simulation of classical dynamics

The majority of simulations for the research carried out in this thesis, used classical techniques to calculate the ground states of Hamiltonians. In the initial stages of research, ground states were found using a brute force approach with code in Python3 written by me. Here, where needed, `numpy.linalg.eig()` was used to diagonalise the Hamiltonian matrix, finding the eigenvalues and eigenvectors. Then the Python in-built `min()` function was used to find the ground states. When the increasing size of Ising models made this approach too computationally intensive, my initial code for generating a lack of precision on several copies of an Ising model was integrated with code written by Adam Callison (first used in [64]) which used a branch and bound technique to find the ground states.

The results of the branch code at low sizes were compared to the results from my initial ground state finding code, in order to check the integration was successful and both codes were functioning correctly. Subsequently the branch and bound code was updated and improved by Adam Callison following my feedback in order that the two codes functioned correctly together.

6.1.1 Branch and bound technique

The branch and bound technique was first proposed in [162], and actually consists of a family of algorithms that find the solution to an optimization problem by using the steps of ‘branching’ and ‘bounding’.

A branch and bound algorithm starts by storing several partial solutions (called subproblems) and complete possible solutions in a tree structure. Each partial solution branch leads to a less and less partial solution until it gets to a complete possible solution at the end of a branch. We first go fully down a single branch in order to get a ball-park complete possible solution. We then explore down each of the other branches in turn, in order to make a bound on the energy of the complete possible solution at the end of the branch. If we get a bound which is higher than the energy of the ball-park solution we already have then we can stop exploring this current branch (as we’re not going to get any better solutions) and can move on to the next branch. This way we avoid having to look at all the possible solutions along the way, and gain our speed-up. If we end up getting to the complete possible solution at the end of any branch, this solution now becomes the new ball-park solution and we move on to the next branch from here. Eventually you will have looked at all of the tree (other than the bits you avoided by the

bounding) and therefore the current ball-park solution is the actual solution to your optimization problem.

Several reviews of branch and bound techniques can be found in [163–166] . A step-by-step description of the branch and bound technique, an updated version of which (following collaboration between me and Adam) was used to compute the ground states for the research in this thesis, can be found in [167].

6.2 Simulation of quantum dynamics

For the initial simulations of quantum walks and adiabatic quantum computing seen in chapter 2, the code was written in Python3 by me. Here, functions such as `numpy.kron()` were used for the tensoring, such as when computing equation (2.3), `scipy.linalg.expm()` were used for exponentiating matrices such as when computing $\exp(-i\hat{H}_c t)$ in equation (2.4), and `numpy.matmul()` were used to multiply matrices and vectors such as also in equation (2.4).

For the quantum walk dynamics in chapter 11, the dynamics simulations were done by integrating my code (for generating the Hamiltonians of instances and subecting them to a lack of precision) with code written by Adam Callison (first used in [64]) using the quimb library [168] for simulation of quantum walks. Quimb was used as a interface for the function `scipy.integrate.complex_ode()`, which adapts times-steps using an 8th order Runge-Kutta method [167, 169].

6.3 Problem instance generation

In section 5.4.1, we demonstrate our error suppression scheme using a 5-qubit maximum independent set (MIS) example. This MIS example was selected from 13 examples at $n = 5$ and its behaviour under our scheme was first analysed by Tom O’Leary. These 13 example instances were generated randomly from code written by Nicholas Chancellor. MIS problems are described in more detail in section 3.4.

Throughout this thesis (specifically in chapters 4, 7, 8, 9, 10, 11) we use a data set of Ising spin chains and SK spin glasses.

The Ising spin chains dataset used in this thesis were generated by code written by me for the purposes of the research included in this thesis. The dataset consists of 10^4 instances of 2, 5, 6, 7, 8, and 9 qubit Ising spin chains (which are described in section 3.2). Each instance is identified by a universally unique identifier (UUID), which were generated by the function `uuid.uuid4()` from the Python

library uuid. Each instance consists of two numpy arrays containing the strengths of the fields h_j and strengths of the coupling strengths J_{jk} respectively. For the fields array, the index of the array corresponds to the qubit number (vertex of the graph) that it acts upon. For the coupling strengths array, the strengths are ordered so that it corresponds to the counting of the combinations $nC2$ (e.g. using `itertools.combinations(range(n), 2)`) where n is the number of qubits in the model. Each of the fields h_j and coupling strengths J_{jk} were randomly generated from a uniform distribution between -1 and +1, unless otherwise stated, using the function `numpy.random.uniform()`.

The SK spin glasses dataset used in this thesis were first generated for use in [64] and are stored at [90]. The dataset consists of 10^4 instances of $5 \leq n \leq 20$ n -qubit Sherrington-Kirkpatrick (SK) spin glasses (which are described in section 3.3). Each instance is identified by a unique ID (UID). Each instance is stored as two numpy arrays containing the strengths of the fields h_j and the coupling strengths J_{jk} respectively. The fields strengths are stored in a 1D length n array, where the index of the array corresponds to the qubit (vertex) on which it acts. The coupling strengths are stored in a 2D $n \times n$ array, where the index of the array corresponds to the qubit (indices) that it is acting between. This looks like the Laplacian matrix in graph theory. Each of the fields h_j and coupling strengths J_{jk} were randomly generated from a normal distribution with mean 0 and standard deviation 1, using the function `numpy.random.normal()`.

6.4 Fitting and error analysis

In chapter 11, we applied a fit of the form $f(p) = A \exp(-bp)$ (where fraction correct was equal to $(1 - f(p))$) to the single disconnected copies data, using the Python function `scipy.optimize.curve()`. This fit was then used to estimate the value of precision p needed for a single copy to achieve a certain value of fraction correct. This estimated precision for the single copy was then subtracted from the value of precision for three connected copies at the same value of fraction correct in order to estimate the precision improvement gained by using the three connected copies. The error in the estimated precision (and therefore precision improvement) was calculated by estimating the errors in the fraction correct of the three copies and the error in the A and b parameters of the fit, using a functional approach (method described in [170]) and combining them in quadrature. Elsewhere, we estimated errors in our average based on the number of repeats /number of samples we took.

6.5 Chapter summary

In this chapter, we first described the software and tools which allowed us to carry out this research. We then mentioned the simulation of precision errors, a key part of this research, which is described in detail in section 4.2. We then described how the simulation of classical dynamics was carried out. First in simplistic cases and then when scaling to larger sized Ising models by using branch and bound. After this, we described how the quantum dynamics were carried out. First in simplistic cases (e.g. in chapter 2) using code written by me and then in more complex cases (e.g. in chapter 11) using code written by Adam Callison. Next we described how we generated our MIS and spin chain instances and how we sourced our SK spin glass instances, all of which we used to carry out our numerical simulations. Finally we described how we used a fitting function in chapter 11, in order to make an estimate of the bits of precision improvement we gained through using our error suppression technique and how we carried out our error analysis across the research in this thesis.

Link Strength Tests

For QAC it was found in [1, 136–138] that, the optimal improvement comes when using strong ferromagnetic links between the Ising model copies. These ferromagnetic links connecting the qubits in different copies must be stronger than the typical field and coupling strengths in the problem Hamiltonian, so that there is a smaller chance of the connected copy qubits being accidentally flipped. Corresponding qubits in connected copies need to be restrained from flipping as much as possible, as in order to correctly recover the original state when using a majority vote technique, over half of the copy qubits need to be in the same correct state.

However, using these strong ferromagnetic links between copies is not favourable when scaling to larger sizes or more copies as, because the coupling strengths between copies must always be larger than the coupling strengths within the copies, further constraints are put on the precision of the variables in the problem Hamiltonian due to the dynamic range limitations of the hardware. These scaling requirements were improved in subsequent work on QAC in [139, 140, 142–144, 151], but the results still rely on optimising the ferromagnetic link strength, and are therefore still sub-optimal for problem Hamiltonians that use the full dynamic range of the hardware [143].

For our error suppression technique, we mandate neither the ferromagnetic couplings or majority votes method. It therefore follows that the optimal link strength between our copy qubits may be different. Hence, we set about determining the optimal link strength between corresponding qubits in copies. In section 7.1, we start by measuring the fraction correct versus link strength J_F between two and three connected copies of 10^3 instances of 5-qubit SK spin glasses and spin chains, at precisions $p = 3$ and $p = 4$. The two copies were connected in a chain and the three copies in a loop. We varied the strength of J_F from -1 to +1. In section 7.1.1 we applied (lack of) precision using the random error model. In section 7.1.2

we repeated the same measurement of fraction correct versus link strength J_F first using the deterministic mid-point error model and then the deterministic random error model. Noting that the improvement effect seen in figures 7.1, 7.2, and 7.3 is an average effect across 10^3 instances, we next looked at individual instances, to see if and how this improvement carries over and if any patterns can be found. In section 7.2, we measured the fraction correct versus link strength J_F between four singular instances of three connected (in a loop) copies of 5-qubit SK spin glasses, for precisions $p = 3$ and $p = 4$. We first applied precision using the deterministic random error model. Then we repeated the measurement using the random error model with 10^3 error samples per value of J_F . We next investigated the role of fields in the Ising problem models, and their effect on our error suppression scheme. We did this by measuring the fraction correct versus J_F , when there were no fields present, on two and three copies of 10^3 instances of 5-qubit spin glasses, for precisions $p = 3$ and $p = 4$. The precision was applied using the deterministic random error model. In section 7.4, based on our results in the previous sections, we hypothesised on the mechanism of error suppression due connecting copies with minimum strength anti-ferromagnetic links. We further discuss this hypothesis by the description of an example in section 7.4.1. Finally in section 7.5, we summarize the chapter.

This chapters contributions are as follows. Adam Callison wrote the code for the classical branch and bound method, which was used to calculate the ground states which were used to calculate the fraction correct in the numerical analysis of this chapter. Mia West carried out the original research which discovered the optimal link strength and error suppression benefit available from connecting copies with anti-ferromagnetic links. My contributions were: further analysis of the optimal link strength under different error models, the discovery of no improvement for no field strengths and the development of some of the theoretical understanding of how our scheme works.

7.1 Optimal link strength

In order to find the optimal strength of couplings between corresponding qubits in copies, we measured the fraction correct (using our “one correct copy” criteria see section 5.4), at precisions $p = 3$ and $p = 4$, of 10^3 instances of two (in a chain) and three (in a loop) connected copies, of 5-qubit SK spin glasses and spin chains, whilst varying the strength of the links between copies J_F from -1 to +1. We first implemented errors due to lack precision using the random error model, then

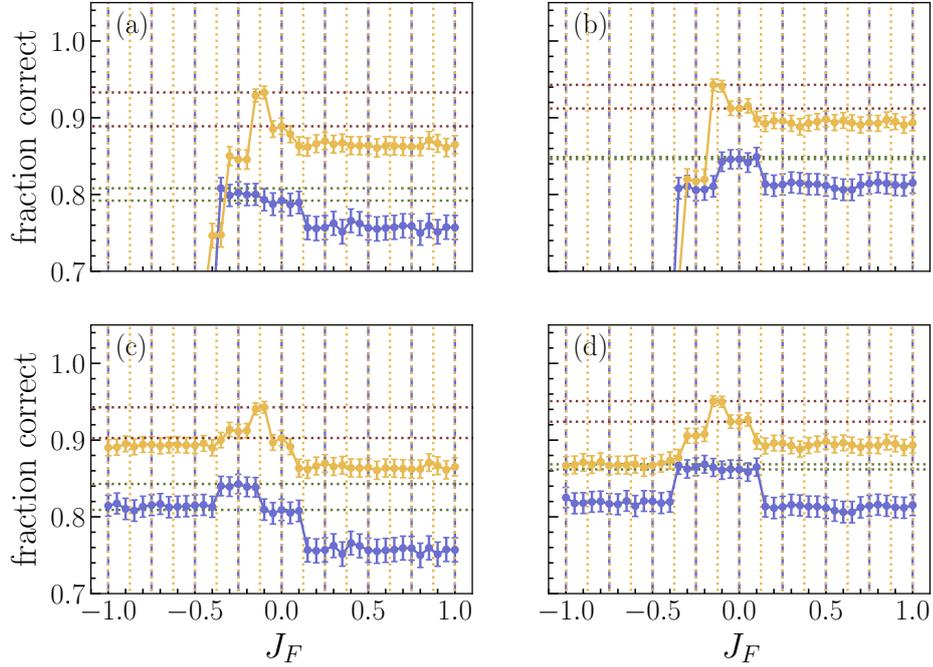


Figure 7.1: Fraction correct against J_F for $p = 3$ (blue) and $p = 4$ (yellow) of 10^3 instances of 5-qubit problems for 2 copies (chain, top row) and 3 copies (loop, bottom row) of spin glasses (left column) and spin chains (right column). Where precision was applied using the random error model. Vertical dotted lines indicate the mid-points of each of the allowed divisions at that precision for $p = 3$ (blue) and $p = 4$ (yellow). Horizontal dotted lines indicate the fraction correct for $J_F^{(p)} = 0$ and the maximum fraction correct for any $J_F^{(p)}$, in green for $p = 3$ and red for $p = 4$.

the deterministic mid-point error model and finally the deterministic random error model.

7.1.1 Random error model

Figure 7.1, shows the fraction correct against link strength $J_F^{(p)}$ for precisions $p = 3$ (blue) and $p = 4$ (yellow), for two copies (chain, top row) and three copies (loop, bottom row) of 10^3 instances of 5-qubit SK spin glasses (left column) and 5-qubit spin chains (right column). Here the errors due to precision were implemented using the random error model (with 10 error samples) (see 4.2.1). Each of the four graphs shown show results from both precision $p = 3$ (blue) and $p = 4$ (yellow), with dotted vertical lines (of the same colour) indicating the resolution values (see section 4.2)

available at that precision. Horizontal dotted lines indicate the fraction correct for $J_F^{(p)} = 0$, i.e., the baseline for any improvement, and the maximum fraction correct over all $J_F^{(p)}$ values, in green for $p = 3$ and red for $p = 4$. The results are similar for both spin chains and SK spin glasses, showing the robustness of the effect.

We see that in all four cases, at both precisions (except from $p = 3$ for 2 and 3 copies of 5-qubit spin chains), there is an maximum in the fraction correct (indicating optimal improvement) when J_F is close to $J_F^{(p)} = -2^{-p+1}$, i.e. at $J_F = 0.25$ for $p = 3$ and at $J_F = 0.125$ for $p = 4$. For consistency, we refer to this value (equal to the minimum allowed strength by the precision) as J_F^{\min} regardless of the error model used. Note that we also tested intermediate values of $J_F^{(p)}$, to check that the value of J_F^{\min} is robust to variation within a division. The choice of divisions in figure 4.2 with only one representing approximately zero is now justified. The other possible choice, with two divisions representing approximately zero, one with ferromagnetic character, the other with anti-ferromagnetic character, would have made it more difficult to extract the optimal value of J_F^{\min} .

In all cases we see that the fraction correct remains relatively stable for all values of $J_F > 0$ shown, before increasing in all cases at $J_F = 0$. This means that in all cases tested here, leaving copies disconnected provides a better fraction correct than connecting copies with ferromagnetic links. We note that this is a contrary result to that of QAC, however it is important to remember that unlike QAC, we do not use the majority voting method to find our fraction correct and should therefore not expect the same conclusions. We also expect that at values of $J_F \gg 0$ the fraction correct will tend towards the value at $J_F = 0$.

For precision $p = 3$ and $p = 4$ for cases (a) and (c) and $p = 4$ for cases (b) and (d), there is a further increase in fraction correct at J_F^{\min} . For $p = 4$ in all cases the optimal fraction correct is around 0.94, increased from around 0.9. For $p = 3$, case (a), shows an improvement from just below 0.8 to just above, case (c) shows an improvement from just above 0.8 to around 0.84. For cases (a) and (b), which are two copies (in a chain) of SK spin glasses and spin chains respectively, for both precisions, there is a sharp decrease in fraction correct for values of $J_F < J_F^{\min}$. For cases (c) and (d), which are three copies (in a loop) of SK spin glasses and spin chains respectively, for both precisions, there is a slight decrease in fraction correct when $J_F < J_F^{\min}$, before it then stabilises. For case (c), the fraction correct stabilises at value close to its value when $J_F = 0$. For case (d), the fraction correct stabilises at a value close to its value when $J_F > 0$.

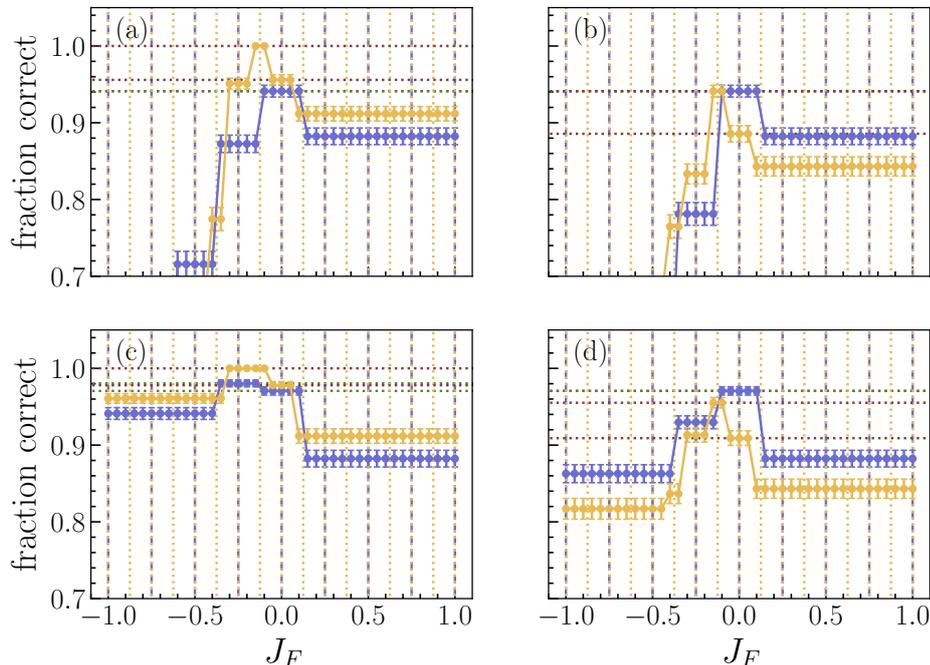


Figure 7.2: As figure 7.1, where precision was applied using the deterministic mid-point error model.

7.1.2 Link strength tests using other error models

Due to limited computational resources, we were unable to continue to use the random error model for Ising models with higher numbers of qubits. We therefore developed the more computationally efficient deterministic mid-point error model (see 4.2.2) and deterministic random error model (see 4.2.2). In order to test whether we obtained similar results, we repeated the fraction correct versus links strength $J_F^{(p)}$ simulations twice using these two models.

Figure 7.2 shows the fraction correct versus link strength $J_F^{(p)}$ results, where precision was applied using the deterministic mid-point error model. We see that some but not all of the qualitative features of figure 7.1 are retained. We see that in all cases at both precisions, the fraction correct at $J_F > 0$ is stable and lower than the fraction correct at $J_F = 0$. We still see that the fraction correct decreases sharply when $J_F < J_F^{\min}$ for both precisions when there are two copies for both models. We also retain the improvement in fraction correct at J_F^{\min} for all cases for $p = 4$. However, we no longer see improvement here for any of the cases at $p = 3$. We also see that this model tends to overestimate the fraction correct, with the

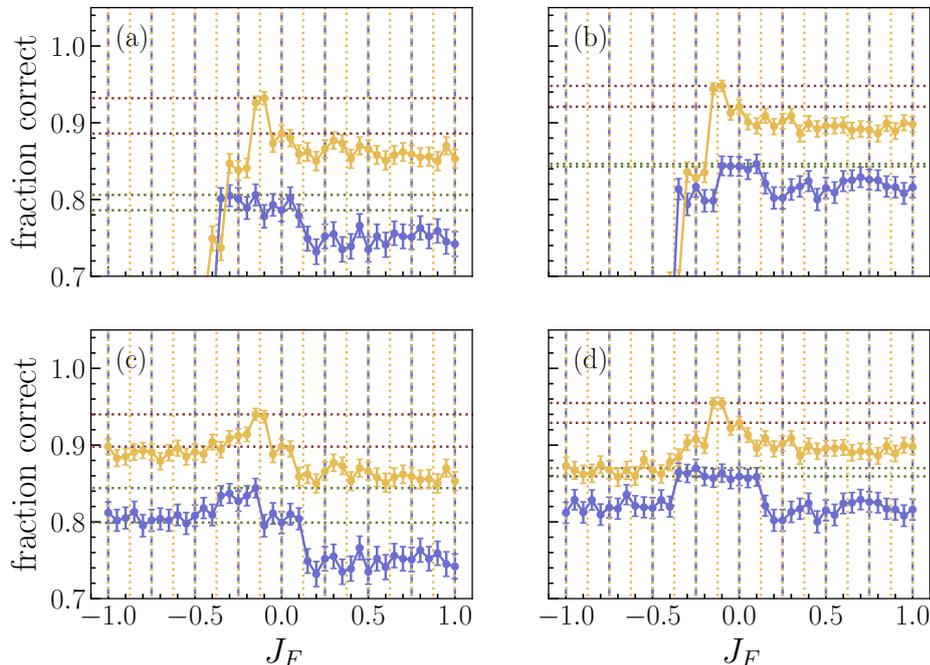


Figure 7.3: As figure 7.1, where precision was applied using the deterministic random error model.

maximal values of fraction correct at around 1 for the SK spin glass cases and at around 0.95 for the spin chain cases. However, this was not an unexpected result as an overestimate in fraction correct when using this error model was predicted theoretically in our error analysis (see 4.2.2) and shown numerically to be the case for a single copy in figure 4.2.

We next repeated the fraction correct versus link strength $J_F^{(p)}$ simulation for the same 10^3 SK spin glasses and spin chains, this time applying the precision using the deterministic random error model. The results can be seen in figure 7.3. These results are even closer to those seen in figure 7.1. In all cases at both precisions, we see the stability of fraction correct at $J_F > 0$ and the increase in fraction correct at $J_F = 0$. Like figure 7.1, there is an optimal value of the fraction correct at J_F^{\min} for all cases at precision $p = 4$ and there is also an optimal fraction correct for the SK spin glasses at $p = 3$. Also the values of fraction correct at these maxima are similar to those in the random error model simulations. (From around 0.9 to 0.94 at $p = 4$ in each case. From just below 0.8 to just above for two copies of SK spin glasses and from around 0.8 to around 0.84 for 3 copies.) Due to the close similarity to the results using the random error model but its relative computational

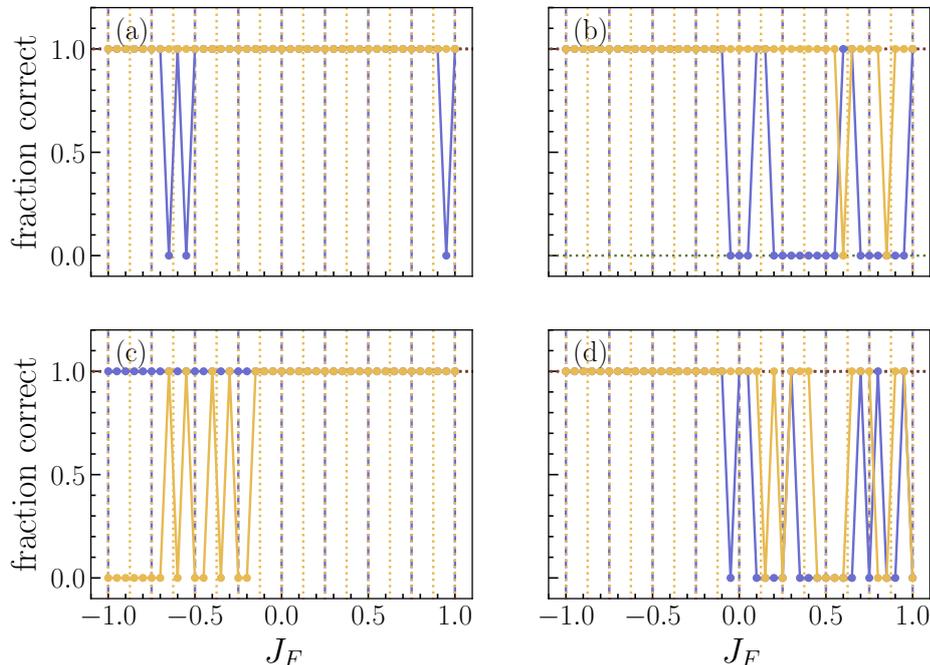


Figure 7.4: Fraction correct against J_F for $p = 3$ (blue) and $p = 4$ (yellow) for 4 single instances of 3 copies of 5 qubit spin glasses. Here the deterministic random error model was used. Vertical dotted lines indicate the mid-points of each of the allowed divisions at that precision, for $p = 3$ (blue) and $p = 4$ (yellow). The four plotted instances (a), (b), (c), and (d) have the unique i.d.'s of: ‘aaavmaiqliolnplcovmzxjazkyvyayz’, ‘aacrpcjsbugeteaageltzcpnpovkcm’, ‘aagtzdpchgtnrzilrnhpvxqtvqiql’, and ‘aakxejqunlcpqhmnftnrckailrczyp’ respectively. The ground states of each of the examples at true precision are $[1\ 1\ 1\ 1\ 1]$, $[0\ 0\ 1\ 0\ 1]$, $[0\ 0\ 1\ 1\ 1]$, and $[1\ 1\ 1\ 0\ 1]$, respectively.

efficiency, we chose to use the deterministic random error model for the majority of our simulations.

7.2 Individual instances

The fraction correct results from figures 7.1, 7.2, 7.3, are an average over 10^3 instances of SK spin glasses and spin chains. To further understand the working of our error suppression method, we measured the effect of our technique on singular 5-qubit SK spin glass instances. In figure 7.4, we plot fraction correct against $J_F^{(p)}$, where precision was applied using the deterministic random error model, for precisions $p = 3$ (blue) and $p = 4$ (yellow) of four singular instances of three copies (con-

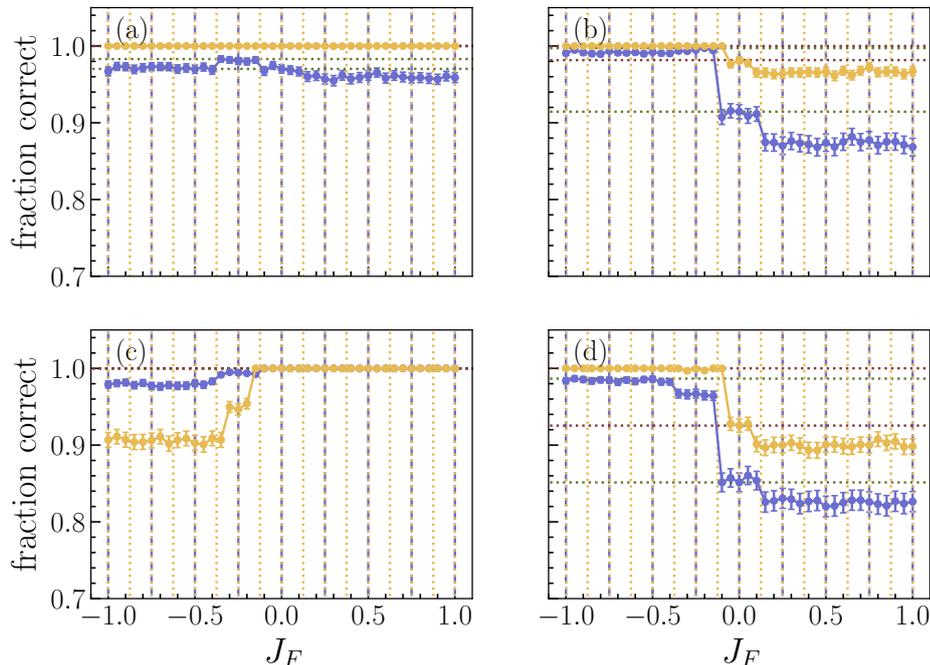


Figure 7.5: Fraction correct against J_F for $p = 3$ (blue) and $p = 4$ (yellow) for 4 single instances of 3 copies of 5 qubit spin glasses. Here the random error model was used with 1000 error samples per point. Vertical dotted lines indicate the mid-points of each of the allowed divisions at that precision for $p = 3$ (blue) and $p = 4$ (yellow). The four plotted instances (a), (b), (c), and (d) have the unique i.d.'s of: ‘aaavmaiqliolnplcovmzxjazkyvyayz’, ‘aacrcpjsbugeteaageltzcpnpovkcm’, ‘aagtzdpchgtknrzihrnhpvxqtvtqql’, and ‘aakxejqunlcpqhmnftrckailrczyp’ respectively. The true ground states remain the same as the as figure 7.4.

nected in a loop) of 5-qubit SK spin glasses. The unique i.d.'s of the four instances are: ‘aakxejqunlcpqhmnftrckailrczyp’ (a), ‘aagtzdpchgtknrzihrnhpvxqtvtqql’ (b), ‘aacrcpjsbugeteaageltzcpnpovkcm’ (c), and ‘aaavmaiqliolnplcovmzxjazkyvyayz’ (d). The (correct) ground states of a single copy of each of the instances with true precision are $[1\ 1\ 1\ 1\ 1]$, $[0\ 0\ 1\ 0\ 1]$, $[0\ 0\ 1\ 1\ 1]$, and $[1\ 1\ 1\ 0\ 1]$ respectively.

When using a deterministic error model, the ground state of a single instance can only either be correct (same as the true precision ground state) or incorrect, so we see in figure 7.4, the fraction correct is either 0 or 1. We also see that the value of fraction correct can flip from 0 to 1 or vice versa, for only a small change in $J_F^{(p)}$. The reason for these ‘flips’ is randomness introduced by the error model. As $J_F^{(p)}$ is varied, although the true values of the fields and couplings in the SK spin glass remain the same, their exact value is chosen randomly (within the resolution

division) each time $J_F^{(p)}$ is varied. This means they change slightly as $J_F^{(p)}$ is varied. When the SK spin glass model's ground state is sensitive to these changes it can cause the 'flipping' we see.

For this reason we study the effect of our error suppression scheme on singular instances using the random error model. As we take multiple error samples (10^3 in this case), instead of getting a fraction correct of either 0 or 1 at each value of $J_F^{(p)}$, we instead calculate a probability between 0 and 1 of fraction correct. As the fraction correct is now a probability calculated over many random samples, we no longer see the 'flipping' and instead see the average trend in probability. This way we get an idea what $J_F^{(p)}$ strength works on average for each instance.

Looking at figure 7.5, we see there is a different behaviour in fraction correct versus $J_F^{(p)}$ for each of the four instances ((b) and (d)), we can see that connecting three copies anti-ferromagnetically ($J_F < 0$) in a loop, provides a benefit in fraction correct. However we see for instance (a) connecting the copies with anti-ferromagnetic links provides no benefit at all and for instance (c) we see that connecting with this strength of link is actually disadvantageous. This means that the benefit at $J_F = J_F^{\min}$ we see in figures 7.1, 7.2, 7.3, is only an overall average effect. Whilst connecting copies with $J_F = J_F^{\min}$ helps improve fraction correct for some instances, it also reduces fraction correct in others.

Whilst our technique of connecting copies is beneficial overall, it is clear that for some instances it would be better to leave the copies disconnected. A method which could predict whether (or not) corresponding qubits in different copies should be connected or not, would clearly be advantageous to extract the maximum benefit. This prediction method should preferably not be computationally expensive, and should therefore rely on variables local to the qubit which you are deciding whether to connect. Introduction and analysis of a protocol which attempts to achieve this criteria can be seen in chapter 10.

7.3 The effect of no fields

In order to gain further insight into the functioning of our error suppression scheme, we next studied the effect of setting all the fields in the SK spin glass to zero. Figure 7.6 shows the fraction correct vs $J_F^{(p)}$, where precision was implemented using the deterministic random error model, for same 10^3 instances of 5-qubit SK spin glasses, with the fields set to zero, for three copies (loop; left column) and two copies (chain; right column). We see that in contrast to the results in figures 7.1, 7.2, 7.3, we no longer see any improvement in fraction correct at $J_F = J_F^{(\min)}$. However we do see

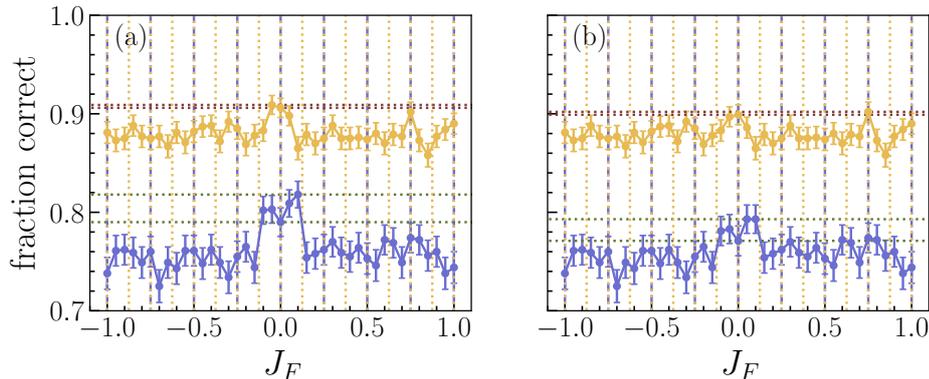


Figure 7.6: Fraction correct against J_F for $p = 3$ (blue) and $p = 4$ (yellow) for 10^3 instances of 5 qubit spin glasses with no fields applied to each qubit, for 2 copies (chain; right) and 3 copies (loop; left). Vertical dotted lines indicate the mid-points of each of the allowed divisions at that precision for $p = 3$ (blue) and $p = 4$ (yellow). Horizontal dotted lines indicate the fraction correct for $J_F^{(p)} = 0$ and the maximum fraction correct for any $J_F^{(p)}$, in green for $p = 3$ and red for $p = 4$.

a very small improvement at $p = 3$, for two and three copies, at around $J_F = 0$. This indicates that in this regime there is no advantage to connecting copies of SK spin glasses at all, either anti-ferromagnetically or ferromagnetically.

We suggest the following reasoning for this result. In this ‘no fields’ setting there are always two degenerate ground states, which are two ‘opposite’ states. For example, if the ground state was 01, then 10 would also be the ground state. The lack of advantage seen in connecting copies together can be understood qualitatively in the following way: If the copies are connected ferromagnetically, then both copies ‘want’ to be in the same state so the probability of either of them being in the ‘correct’ state is simply the probability of the singular state being in the ‘correct’ state at that precision minus the probability that the precision or the connecting links causes the copies to become incorrect (hence the slightly higher fraction correct at $J_F^{(p)} = 0$). Conversely, if the copies are connected anti-ferromagnetically, then both copies ‘want’ to be in opposite states. Whilst this would normally force one of the copies to be in the ‘incorrect’ state, in this setting the opposite state is also the ground state, so here there is no change in probability of being in the ‘correct’ state between the ferromagnetically connected copies and the anti-ferromagnetically connected copies.

States	Energy penalty
	$+5J_F^{(p)\min}$
	$+3J_F^{(p)\min}$
	$+J_F^{(p)\min}$
	$-J_F^{(p)\min}$
	$-3J_F^{(p)\min}$
	$-5J_F^{(p)\min}$

Table 7.1: Table showing the states of each possible spin flip configuration of two anti-ferromagnetically connected Ising models (left column), with their associated energy penalty (right column)

7.4 On the mechanism of error suppression via anti-ferromagnetic links

We can use the insight gained from understanding why connecting SK spin glass copies with no fields does not work, to gain further insight into why it does work when fields are present. In the setting where the SK spin glass does have fields, the ground state is no longer degenerate with its ‘opposite’ state, e.g. the system ‘prefers’ 01 over 10. If the copies are connected ferromagnetically, both copies ‘want’ to be in the same state. (This is the situation that minimises the energy penalty.) This is ideal when the model has the correct ground state and the ferromagnetic links even provide an energy barrier to incorrect states. However when there is enough energy to surpass this barrier, it is likely that both states will become incorrect.

If the copies are connected anti-ferromagnetically, it is less likely that the overall ground state of the connected system will be two (or more) ‘correct’ ground states of the singular model. This is because two of the same states connected together anti-ferromagnetically will have an energy penalty of $+nJ_F^{(p)\min}$, see row one of table 7.1.

If one of the anti-ferromagnetic links are satisfied i.e. one of the qubits in either copy flips, this penalty reduces to $+(n-2)J_F^{(p)\min}$ and then $+(n-4)J_F^{(p)\min}$ and so on until there are a greater number of satisfied anti-ferromagnetic links, than unsatisfied. At this point the anti-ferromagnetic links instead give an energy benefit. For two fully opposite states this energy benefit will be maximised at $-nJ_F^{(p)}$. However, due to the bias caused by the fields on the singular model, the fully opposite state, is unlikely to be the ground state of the singular model and therefore this state is also unlikely to minimise the energy of the system. The ground state of the connected system of copies is therefore likely to be a compromise between the lowest energy states of the SK spin glass and satisfying the most anti-ferromagnetic links between the copies. We next present an example which illustrates this reasoning.

7.4.1 An example of the mechanism of error suppression

Figure 7.7 shows the energy versus states of a singular 5 qubit spin glass with unique i.d.: affbxcwhjowsiiqwvnavfuagfralz at $p = 3$. The numerical states 0-31 each represent a length 5 bit string each representing each of the 32 possible 5-qubit ground states, ordered in the following way 00000, 00001, 00010, 00011, 00100,..., 11111. The ground state of the single copy at $p = 3$ is highlighted in green and the ground state of the single copy with true precision is highlighted in orange. The two 5-qubit states that make up the ground state of the two anti-ferromagnetically connected (with J_F^{\min}) copies are marked with red crosses. The black arrows to dotted horizontal lines indicate energy benefits or penalties from the anti-ferromagnetic links by connecting various states (explained in the proceeding paragraphs).

For example for the instance with unique i.d.: affbxcwhjowsiiqwvnavfuagfralz, in figure 7.7, we see that state (7) 00111 (highlighted in green), is the ground state for a single copy at $p = 3$. This is different to the ground state at true precision which is (15) 01111 (highlighted in orange), indicating that the single copy of this spin glass is ‘broken’ at $p = 3$. For two disconnected copies of this spin glass, the ground state of one copy at $p = 3$ would be the ground state for both copies i.e. simply two copies of (7): 00111 00111, two incorrect copies. However, by connecting corresponding qubits in the two copies together with anti-ferromagnetic links, we either introduce an energy penalty or benefit (depending on whether/how many of these links are satisfied). This essentially either pushes the energy of one of the connected states up (penalty) or down (benefit). (In figure 7.7, we’ve represented this potential change in energy, by labelling several of the key states with black arrows (a) - (d).)

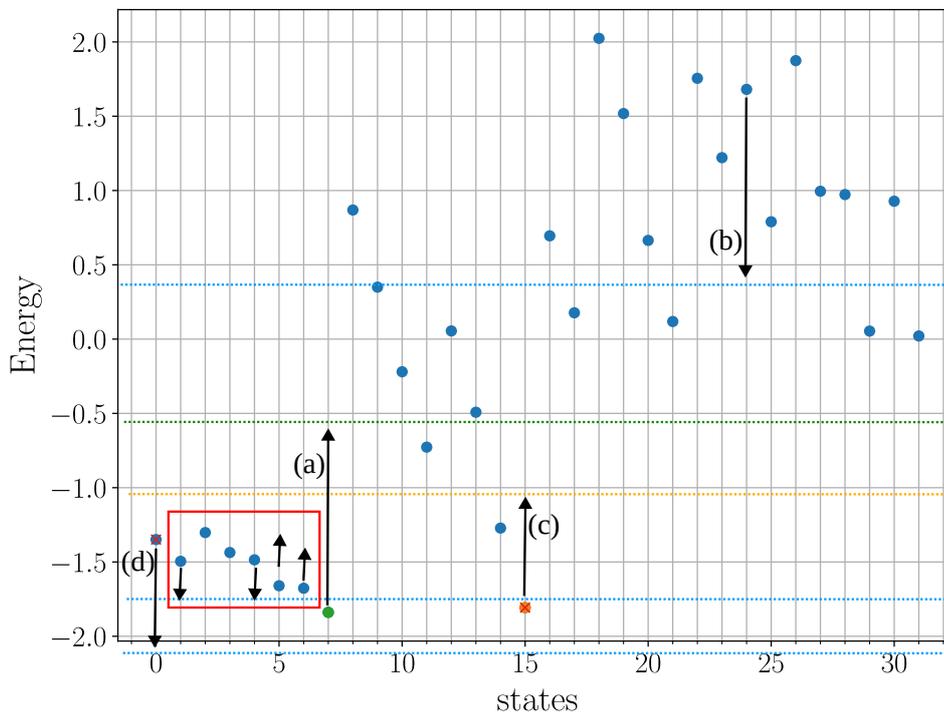


Figure 7.7: Graph showing the Energy vs ground states, of a single copy of a 5 qubit Ising spin glass with $p = 3$ with unique i.d.: affxbxcwhjowsiiqwvnavfuagfralz. The states represent length 5 bit strings representing qubits, ordered as following: 00000, 00001, 00010, 00011, 00100,... The ground state of the single copy at $p = 3$ is highlighted in green and the ground state of the single copy with true precision is highlighted in orange. The two states that make up the ground state of the two connected copies at $p = 3$ are marked with red crosses. Black arrows (a)-d)) to dotted horizontal lines indicate energy benefits or penalties from the anti-ferromagnetic links by connecting the ground state at $p = 3$ (green) with the state on which the arrow is (as explained fully in the text).

The end effect of this pushing up and down of energy is that two copies of (7): 00111 00111 is no longer the ground state of the two connected copy system. The ground state becomes (15) and (0): 01111 00000. We see that we have also recovered in one of the two copies the ground state at true precision (a success for our scheme!). We explain in more detail how this occurred below.

First we look at the situation if we simply connected two of the $p = 3$ ground states: (7) 00111, with anti-ferromagnetic links of strength $J_F^{(p)\min}$. This corresponds to row 1 in table 7.1. In this example $p = 3$, so $J_F^{(p)\min} = 0.25$ and as there are 5-qubits, the energy penalty is $+1.25$. This is represented by arrow (a). We can see this is a large penalty which takes us above many low lying states, so this is

not favourable for a ground state.

On the other hand, the situation which maximises the benefit from the anti-ferromagnetic links, is to connect (7) 00111 with the completely opposite state (24) 11000, to give: 00111 11000. This corresponds to the bottom row in table 7.1 and in this case the maximum energy benefit is -1.25. We represent this by arrow (b) on figure 7.7. However we can see in this case, (24) 11000 is a very high lying state, so its large energy benefit is not enough to overtake the low lying states.

So next we focus on the low-lying states as these are the most likely to form part of the two copy system ground state. Looking at the ground state at true precision (15) 01111, we see that it is also low lying, very close to the energy of the ground state at $p = 3$ (7) 00111. However if we anti-ferromagnetically connect (7) and (15) to give 01111 00111, we only satisfy one anti-ferromagnetic link, meaning the energy penalty corresponds to the second row in table 7.1 and in this case has an energy penalty of +0.75. We represent this by arrow (c) on figure 7.7, and can see that this penalty still takes us above many low lying states.

Most of these low lying states occur between 0-6 and each start with two leading zeros, meaning if connected with (7) 00111, we can at maximum only satisfy three links. For example, if we connect with (0) 00000, giving 00111 00000, this corresponds to the fourth row of table 7.1 and in this case an energy benefit of -0.25. However, we can also see that if we instead connect (0) 00000 with the ground state at true precision (15) 01111, giving 01111 00000, we are able to satisfy one more link. This corresponds to the fifth row of table 7.1 and in this case gives an energy benefit of -0.75. We represent this by arrow (d) on figure 7.7. We can see that this larger energy penalty is bigger than the energy difference between (7) and (15) allowing the state containing (0) and (15) to overtake the state containing (0) and (7) to become the overall ground state of the two copy system. In this way we have also recovered the ground state at true precision as one of the two connected copies, meaning a success for our scheme!

Therefore the mechanism of our error suppression technique can be described as follows: Due to the energy penalty introduced by the anti-ferromagnetic links, the ground state of the two copy system will likely become one ‘correct’ state and one ‘incorrect’ state, that is low lying in energy and has at least one or more of the connected spins flipped. Having only one ‘correct’ state is not an issue for our method as we only require one or more of the states to be ‘correct’. When an energy shift imparted by a precision error is enough that it causes the ‘correct’ single state to become ‘incorrect’, by the same mechanism the ground state of the two copy system is more likely to contain one state that is the ground state at true precision as long as said state is low lying in energy and has at least one or more

of the connected spins flipped.

It can be seen that the two copy system will not *always* contain the ‘correct’ ground state. But that is *not* the claim of this scheme, we simply claim that the two copy system contains the ‘correct’ ground state *more often* than two disconnected copies.

7.5 Chapter summary

In this chapter, we first looked at the fraction correct versus the strength of the links connecting the corresponding qubits in two connected (in a chain) and three connected (in a loop) copies of 5-qubit SK spin glasses and spin chains averaged over 10^3 instances for precisions $p = 3$ and $p = 4$. We first applied precision using the random error model, second the deterministic mid-point model and third the deterministic random error model. We found in most cases tested, that there was an optimum in fraction correct close to $J_F^{(p)} = -2^{-p+1}$, e.g. near $J_F = 0.25$ for $p = 3$ and $J_F = 0.125$ for $p = 4$. To avoid ambiguity across error models we refer to this value of J_F as J_F^{\min} . This was an improvement above the value of fraction correct at $J_F = 0$ or $J_F > 0$, indicating there was a benefit to connecting copies anti-ferromagnetically with this strength of link. We also found that for two copies (in a chain) of both SK spin glasses and spin chains, when $J_F < 0$, as the J_F decreased, the fraction correct rapidly decreased, to be well below the fraction correct at $J_F = 0$ and $J_F > 0$. However, for three copies (in a loop) of both SK spin glasses and spin chains, when $J_F < 0$, as J_F decreased, the fraction correct decreased, however not as rapidly or as far as for two copies, and for both types of model levelled out. This levelling out was at a similar value of fraction correct to $J_F = 0$ for SK spin glasses, and for spin chains this levelling out was around the value of fraction correct for $J_F > 0$. This behaviour was consistent across all three error models tested here, despite variations in the fraction correct measurements for each.

As the fraction correct results of the preceding section were an average across 10^3 instances, we next looked at the the fraction correct versus link strength J_F for four example instances of three copies of 5-qubit SK spin glasses connected in a loop for precisions $p = 3$ and $p = 4$. We first applied the precision using the deterministic random error model. However, due to the binary nature of this error model on single instances, the results were unclear. Therefore, we repeated the measurement, applying precision using the random error model, with 10^3 error samples at each value of J_F . In each of the instances the behaviour for $p = 3$ was similar to the behaviour for $p = 4$. We found that in two of the four instances there was a benefit to connecting the three copies with anti-ferromagnetic links and this

improvement was not affected by the magnitude of $J_F < 0$. In one of the instances there was no visible advantage or disadvantage from connecting the three copies with anti-ferromagnetic or ferromagnetic links. In the other instance we found that there was a disadvantage in fraction correct from connecting the three copies with anti-ferromagnetic links. This indicated the difference in behaviours of fraction correct across different instances. We hypothesised that a numerically efficient technique for identifying whether or not a specific instance should be connected anti-ferromagnetically or left disconnected, could enhance the error suppression capabilities of our scheme further.

We next looked into the effect the fields h_j present in the problem Ising Hamiltonian, had on our error suppression scheme. We measured the fraction correct versus link strength J_F for 10^3 instances of two (connected in a chain) and three copies (connected in a loop) of 5-qubit spin glasses which had no fields (fields set to zero) for precisions $p = 3$ and $p = 4$. We applied precision using the deterministic random error model. We found that for both the two and three copies of spin glasses, the optimum in fraction correct when J_F was close to J_F^{\min} , was no longer present. We hypothesised that this may be down to the symmetry of the ground state (0 or 1 equivalent) when no fields were present. This may indicate that the efficacy of our error suppression scheme is dependent on the asymmetry of the ground state of the problem Ising model.

Using the insight gained from the preceding two sections, we next hypothesised on the mechanism of our error suppression scheme. We hypothesised that the mechanism of our error suppression scheme, was due to the balance of needing to find the ground state of the problem Ising Hamiltonian and also needing to minimise the energy penalty given by the anti-ferromagnetic links. This hypothesis was backed up by the analysis of a single instance of a 5-qubit spin glass which has lost its original correct ground state by lack of precision at $p = 3$. We looked at the potential penalties given by the anti-ferromagnetic connection of a second copy. In this example these energy penalties, meant that connecting this instance anti-ferromagnetically to a second copy would recover the original correct ground state on one of the connected copies. We use this example to suggest that this is the mechanism for all error suppression by our technique, with the caveat that this does not work for all instances, but it finds the correct ground state *more often* than separate copies.

Error Suppression Results

In chapter 7, we saw in figures 7.1, 7.2, 7.3, that there is, on average, an improvement in fraction correct gained by connecting copies of 5-qubit SK spin glasses and spin chains with anti-ferromagnetic links, at precisions $p = 3$ and $p = 4$, with optimal improvement when the strength of the links connecting the copies was at or close to $J_F^{(p)} \simeq J_F^{\min}$ (the minimum allowed value by the precision). In order to see whether this improvement was consistent across a larger range of precisions and copy configurations, we next measured the fraction correct versus precision of two connected (in a chain), three connected (in a chain and loop), four (connected in a loop) and five (connected in a loop) copies of 5-qubit SK spin glass and spin chain Ising models, for precision in the range $1 \leq p \leq 10$. We connected the copies with links of strength $J_F^{(p)} \simeq J_F^{\min}$. For the two copies connected in a chain, we ran simulations first with precision implemented using the deterministic error model and then using the random error model for comparison. For all other copy configurations we implemented the precision using the deterministic random error model only. Adam Callison wrote the code for the classical branch and bound method, which calculated the ground states used to calculate the fraction correct in this chapter.

In section 8.1, we measured the fraction correct versus precision of two copies of 5-qubit SK spin glasses connected using our error suppression scheme compared to two repeats of single disconnected copies. First in subsection 8.1.1, we apply precision using the deterministic random error model. Then in subsection 8.1.2 we apply precision using the random error model and compare and discuss the differences between the results. In section 8.2, we found the fraction correct versus precision of two copies of 5-qubit Ising spin chains connected (in a chain) using our error suppression scheme compared to two repeats of single disconnected copies. Here, the precision was applied using the deterministic random error model. All simulations in the preceding sections were done using the deterministic ran-

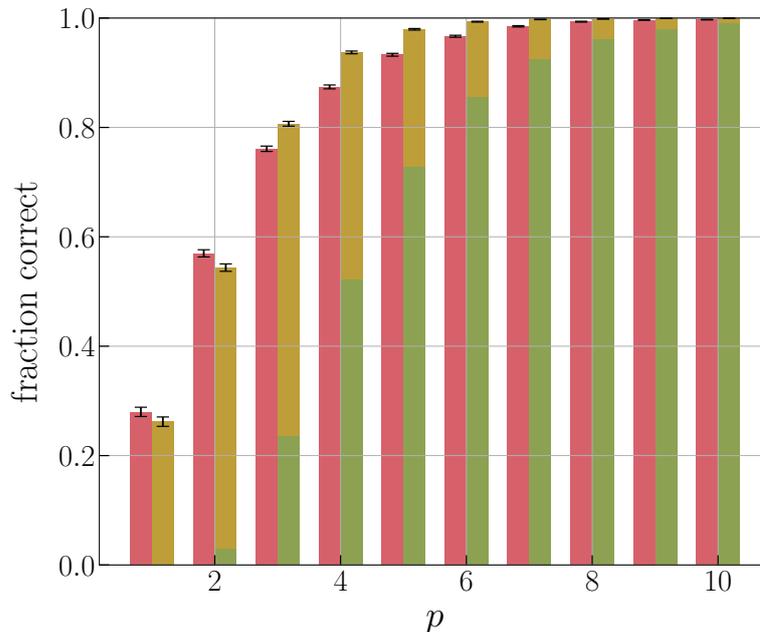


Figure 8.1: Fraction correct versus precision p for 10^4 instances of 5 qubit SK spin glasses. Left bars show fraction correct for single copies (red). Right bars show fraction correct for 2 copies connected in a chain with links $J_F^{(p)} = J_F^{(\min)}$ split into 2 correct (green), 1 correct (yellow). Precision was applied using the deterministic random error model.

dom error model. In section 8.3, we looked at the fraction correct versus precision for three copies in two configurations (chain and loop) of 5-qubit SK spin glasses and spin chains. We first looked at three copies of spin glasses connected anti-ferromagnetically in a chain 8.3.1. We then looked at three copies of spin glasses connected anti-ferromagnetically in a loop 8.3.2. Finally, we looked at three copies of spin chains connected in a loop 8.3.3. In section 8.4, we measured the fraction correct versus precision of four copies of 5-qubit SK spin glasses connected with anti-ferromagnetic links in a loop and discuss the results. In section 8.5, we measured the fraction correct versus precision of five copies of 5-qubit SK spin glasses connected with anti-ferromagnetic links in a loop and discuss the results. Finally, in section 8.7, we summarize the chapter.

8.1 Two copies of spin glasses

8.1.1 Deterministic random error model

We first looked at two copies of SK spin glasses connected in a chain with minimum strength antiferromagnetic links. Precision was applied using the deterministic random error model. Figure 8.1 shows the fraction correct versus precision p in the range $1 \leq p \leq 10$ for 10^4 instances of two copies of 5-qubit SK spin glasses. The total height of the right hand composite bars (green and yellow) show the the fraction correct of two copies connected together in a chain with anti-ferromagnetic links with strength $J_F^{(p)} \simeq J_F^{\min}$. The green portion of the bars is the fraction of instances where both copies were correct and the yellow portion of the bars is the fraction where just one of the two copies was correct. Note that we consider both these cases a success for our scheme and therefore they both add to the total fraction correct. We compare these results to the fraction correct of two repeats of a single disconnected copy of the same 10^4 instances of 5-qubit SK spin glasses (this uses the same resources as two connected copies). These results are shown by left hand bars (red) of figure 8.1. We note that in the deterministic random error model, the improvement gained from repeats is negligible. However, in each case we calculated, we compare the fraction correct of the connected copies to the fraction correct of the equivalent number of disconnected copies.

We can see from figure 8.1, that at $p = 1$ and $p = 2$ there is no improvement in fraction correct for the two anti-ferromagnetically connected copies. At both these precisions it would be better to leave the two copies disconnected. However for $p \geq 3$, we see an advantage in fraction correct for the two connected copies, above the two disconnected copies, indicating an improvement gained from using our scheme at these precisions. The error bars (black) at the top of each bar show that the behaviour we see is much larger than statistical effects. These results clearly show how the anti-ferromagnetic links frequently ‘break’ one of the two connected copies, the bar for two correct connected copies is always lower, than the bar for two disconnected single copies. However, it also shows that at $p \geq 3$, the antiferromagnetic links are able to keep at least one of the two copies correct more often, even when the single disconnected copy is broken by reduced precision.

8.1.2 Random error model

In the deterministic random error model, we apply a single fixed error value for each field and coupling of each instance. Once the errors for an instance have

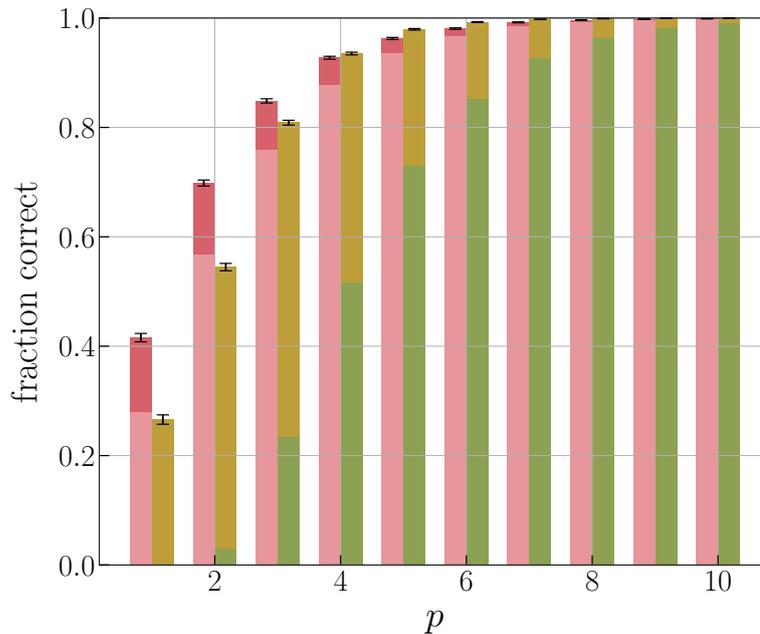


Figure 8.2: Fraction correct versus precision p for 10^4 instances of 5 qubit SK spin glasses. Left bars show fraction correct for single copies (red). Right bars show fraction correct for 2 copies connected in a chain with links $J_F^{(p)} = J_F^{(\min)}$ split into 2 correct (green), 1 correct (yellow). Here precision was applied using the random error model.

been applied, the ground state we then measure cannot be changed. It will always either be correct or incorrect, i.e. have a success probability of either 0 or 1. This outcome will remain the same even if we repeat the measurement. This means that doing repeats of single copies when using this error model has negligible impact on the fraction correct. However this is not necessarily a true reflection of what would happen in experiment. On the other hand, for the random error model, the random error that is applied to each field and coupling in the Ising model instance, will change for each each repeat (or error sample) you take. This means that the success probability of each instance will likely be a value between 0 and 1. This means that when using this error model, we can increase the likelihood of finding the correct ground state and thereby improve the fraction correct by carrying out repeats.

Figure 8.2 shows the fraction correct versus precision p , of 10^4 instances of 5-qubit SK spin glasses, where precision was this time applied using the random error model. The height of the right hand composite bars (green and yellow) show the fraction correct of two copies anti-ferromagnetically connected with minimum strength links. In this figure, we note, the left hand bars are now also composite.

The total height of these left hand bars (pink and red) show the fraction correct of two repeats of a single copy of each of the 10^4 spin glass instances. The light pink portion of the bar represents the fraction correct of a single repeat of a single disconnected copy (very close to the height of the red bars in figure 8.1) and the red portion of the bar shows the improvement gained from the second repeat of the single disconnected copy. Compared to disconnected copy fraction correct when using the deterministic random error model, for precision implemented using the random error model, in figure 8.2, we see that the disadvantage from connecting the two copies of the SK spin glasses anti-ferromagnetically with minimum strength links, at $p = 1$ and $p = 2$ is increased and extended to $p = 3$. However for $p \geq 4$ we still see a moderate improvement from connecting the two copies with minimum strength anti-ferromagnetic links. This indicates that the connected copies still provide an additional effect beyond independent repetition, albeit less than in the deterministic error situation.

As we assume completely uncorrelated errors when applying this error model, we use equation 4.1 in section 4.2.1 to calculate the fraction correct for the two repeats of single disconnected copies. However, in hardware it is likely that there would be a mixture of both uncorrelated and correlated errors present, meaning the true improvement from a repeat would be smaller than that seen in figure 8.2. The best balance between repeats of single copies and connecting multiple copies thus depends on the specific hardware implementation. We expect the effectiveness of repeating single copy runs to reduce for larger sizes, because there is more chance that while one random error is removed, another is introduced, but this was not yet tested.

8.2 Two copies of spin chains

We next measured the effect on fraction correct of connecting two copies (in a chain) of Ising spin chains, with minimum strength anti-ferromagnetic links. Here, the precision was applied using the deterministic random error model. Figure 8.3, shows the fraction correct against the precision p in the range $1 \leq p \leq 10$, for 10^4 instances of 5-qubit Ising spin chains. Like the previous figures, the height of the right hand composite bars (green and yellow) shows the fraction correct of the two anti-ferromagnetically connected copies, whereas the height of the right hand bar (red), shows the fraction correct of two repeats of a single disconnected copy. Likewise, the green portion of the right hand bar shows the fraction of the connected copies instances which both had the correct ground state, and the yellow portion of the right hand bar shows the fraction of connected copies instances which

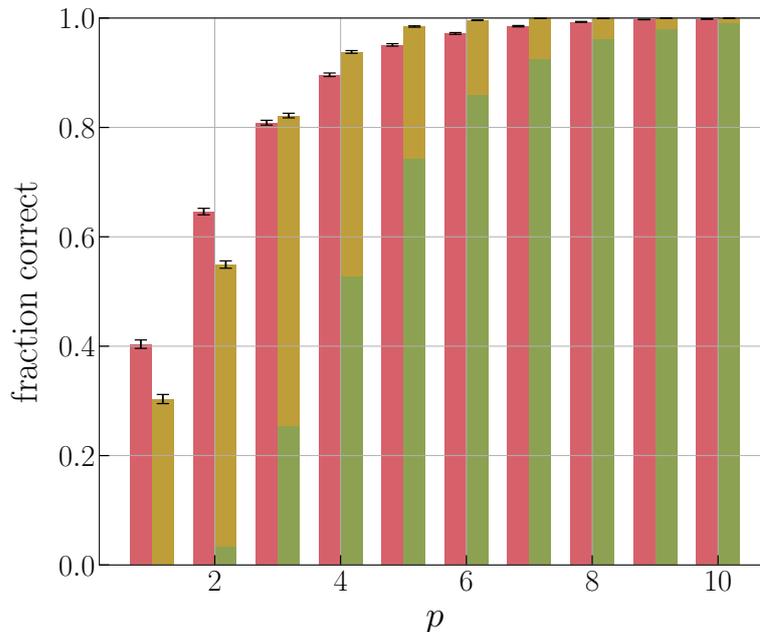


Figure 8.3: Fraction correct versus precision p for 10^4 instances of 5 qubit spin chains. Left bars show fraction correct for single copies (red). Right bars show fraction correct for 2 copies connected in a chain with links $J_F^{(p)} = J_F^{(\min)}$ split into 2 correct (green), 1 correct (yellow). Here precision was applied using the deterministic random error model.

only had one of the two copies correct. The error bars (black) at the top of each of the bars, again show that the effect is much larger than statistical effects.

We see in figure 8.3 that the disadvantage between the fraction correct of two disconnected copies and two anti-ferromagnetic connected copies of 5-qubit spin chains for $p = 1$ and $p = 2$, is much larger than the disadvantage we see for the same configurations of 5-qubit SK spin glasses. We estimate that the disadvantage between two disconnected and two connected 5-qubit spin chains at $p = 1$ and $p = 2$ is around 10%. Whereas for two disconnected and two connected 5-qubit spin glasses at $p = 1$ and $p = 2$, we estimate the disadvantage to be around 2–5%. However like SK spin glasses we saw an advantage from connecting the two copies for $p \geq 3$ and at higher precision the difference in fraction correct between the two types of Ising model becomes negligible.

These results clearly show that the Ising spin chains behave in a similar way to the SK spin glasses, showing again that the anti-ferromagnetic links, tend to break one of the copies whilst keeping one of the copies correct more often than disconnected copies. This indicates our error suppression method is effective at a range of

precisions for at least more than one type of Ising model.

8.3 Three copies

In chapter 7, we saw that for both two and three anti-ferromagnetically connected copies of SK spin glasses and spin chains, at precisions $p = 3$ and $p = 4$, there was an optimal value of fraction correct, when the strength of the anti-ferromagnetic links was equal or close to J_F^{\min} . However we saw that for two copies that, as the strength of these anti-ferromagnetic links increased, the fraction correct falls sharply. We hypothesise that this is due to the anti-ferromagnetic links favouring the completely opposite state (less likely to be the correct ground state) more strongly as their strength increases.

On the other hand, for three copies connected in a loop, as the strength of the anti-ferromagnetic links increased, we saw a much smaller reduction in fraction correct before it levelled off. We hypothesise that this is because, in order to maximise the energy benefit from the anti-ferromagnetic links, unlike the two copies, the three copies do not need to be in completely opposite states, as long as at least one flip is maintained within each loop of three qubits. This means that more combinations of states are likely to be part of the total ground state of the three copies, and it is therefore more likely that one of these states will be the ‘correct’ ground state. Therefore, we wanted to measure the fraction correct against the precision $1 \leq p \leq 10$ for several configurations of three copies of Ising models. In all proceeding data presented in this chapter the precision was applied using the deterministic random error model.

8.3.1 Spin glasses in a chain

We measured the fraction correct against precision $1 \leq p \leq 10$, for 10^4 instances of three copies of 5-qubit SK spin glasses connected in a chain, and compared this to three repeats of a single disconnected copy of the same instances. Figure 8.4 shows the results. The left hand bars (red) shows the fraction correct of three repeats of single disconnected copies. The total height of the right hand composite bars show the fraction correct of three anti-ferromagnetically connected copies. The blue portion of the right hand bar is the fraction of instances that had all three copies correct. The green portion is the fraction where two copies were correct and the yellow portion is the fraction where just a single copy was correct.

We can see from figure 8.4, that there was disadvantage in fraction correct for connecting three copies in a chain at precisions $p = 1$, $p = 2$ and $p = 3$ compared

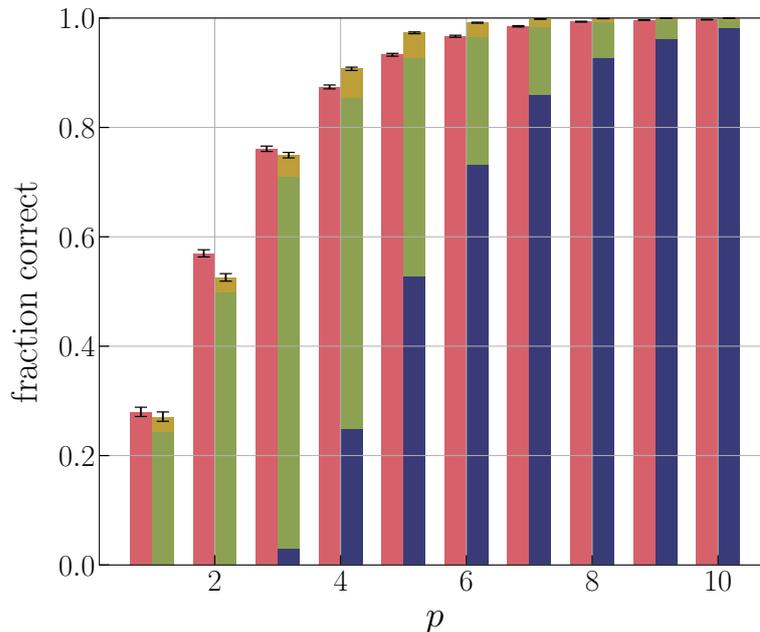


Figure 8.4: Fraction correct versus precision p for 10^4 instances of 5 qubit SK spin glasses. Left bars show fraction correct for 3 repeats of single disconnected copies (red). Right bars show fraction correct for 3 copies connected in a chain with links $J_F^{(p)} = J_F^{(\min)}$ split into 3 copies correct (blue), 2 correct (green), and 1 correct (yellow). Here precision was applied using the deterministic random error model.

to three disconnected single copies. However, for $p \geq 4$, there was a small advantage to connecting the three copies in a chain. This configuration of three copies in a chain therefore provides a similar but less strong effect as connecting just two copies in a chain. (For 2 copies in a chain we saw an improvement in fraction correct at $p = 3$.) We see that in the majority of the three connected (in a chain) copies cases, there is either all three, or two copies correct. Only a very small fraction cases have just one copy correct. We hypothesise that this is due to the following. For three copies connected in a chain, to maximise the energy benefit of the anti-ferromagnetic links, the central copy would have to be the opposite of the two states it is connected to. To also maximise the energy benefit of the optimisation problem, the two outer states would have the lower energy state compared to the central state. If we assume normally these two states will be the ‘correct’ ground state, we can see that there would normally be two correct states and one incorrect state. On the other, hand if these two states are incorrect due to the lack of precision, there is a small chance that the flipped central state will now be the only state that has the ‘correct’ ground state, and this accounts for the section of yellow bar in figure 8.4.

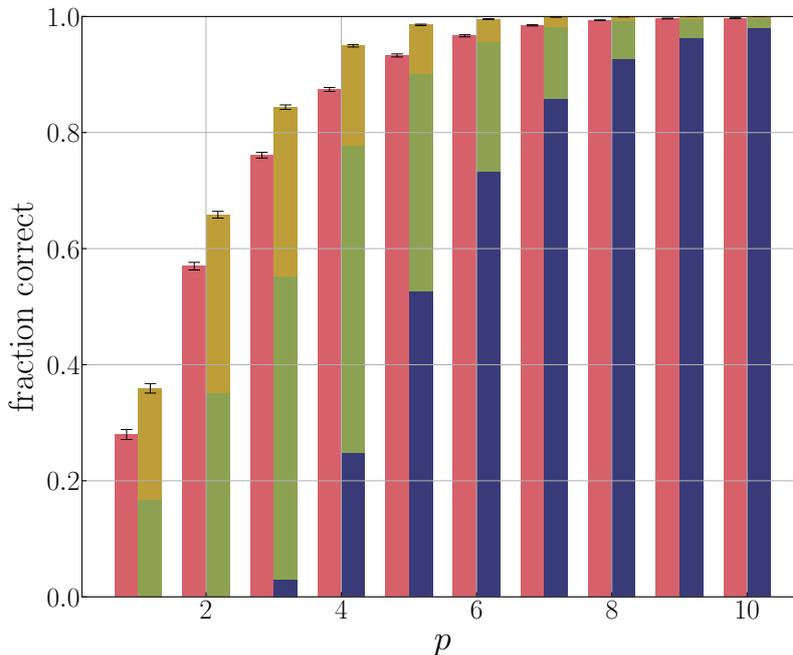


Figure 8.5: Fraction correct versus precision p for 10^4 instances of 5 qubit SK spin glasses. Left bars show fraction correct for 3 repeats of single disconnected copies (red). Right bars show fraction correct for 3 copies connected in a loop with links $J_F^{(p)} = J_F^{(\min)}$ split into all 3 copies correct (dark blue), 2 correct (green), 1 correct (yellow). The precision was applied using the deterministic random error model.

8.3.2 Spin glasses in a loop

After seeing that there was no improvement at low precisions and only a small advantage at higher precisions from connecting three copies in a chain, we next connected three copies of 5-qubit SK spin glasses in a loop. This introduces frustration into the system, meaning that there would now be more configurations of states that could maximise the energy benefit from the anti-ferromagnetic links i.e. gain more access to low-lying states which potentially include the ‘correct’ ground state.

In figure 8.5, the left hand bars (red) show the fraction correct against precision p of three repeats of a single disconnected copy of a 5-qubit SK spin glass averaged over 10^4 instances, for precisions in the range $1 \leq p \leq 10$. The total height of the right hand composite bars (blue, green, yellow), shows the fraction correct of a single repeat of three anti-ferromagnetically connected (in a loop, with link strength $J_F^{(p)} = J_F^{(\min)}$) 5-qubit SK spin glasses averaged over the same 10^4 instances, for the same precision range. (Note: the improvement gained from repeats in the deterministic random error model is negligible.) The composite bars (right) are

broken down into the three possible cases: all three copies correct (blue, lowest), two of three copies correct (green, middle) and one of three copies correct (yellow, top). The error bars (black) at the top of each bar show that the effect is much larger than statistical effects.

We can see that for precision $1 \leq p \leq 10$, there is always an improvement (on average) in fraction correct from connecting the three copies together in a loop rather than leaving them disconnected. These results clearly show how the anti-ferromagnetic links, tend to break one or more copies; the bar for all three connected copies correct (blue) is always lower than the three repeats of single disconnected copies fraction correct (red). Additionally this blue section of the bar is of a similar height to the same section in figure 8.4, indicating this is an effect also present for three copies connected in a chain. Therefore this effect is likely due to the anti-ferromagnetic links only, rather than the configuration of the three copies. Conversely, compared to figure 8.4, the section of bar for two correct copies (green) is slightly reduced, whereas the the section of bar for one copy correct (yellow) is largely increased. This change results in the stronger improvement in fraction correct we see by connecting the three copies in a loop rather than chain. This shows that the frustration introduced by connecting the three copies in a loop (instead of a chain) tends to keep at least one of the three connected copies correct, even when the single disconnected copy is broken by reduced precision, and more often than anti-ferromagnetic links alone are able to do so.

8.3.3 Spin chains in a loop

Still using the deterministic random error model, we next connected three copies of 5-qubit Ising spin chains in a loop using minimum strength anti-ferromagnetic links, in order to see how the introduction of frustration into the system affected the fraction correct in this setting.

In figure 8.6, the left hand bars (red) show the fraction correct of three repeats of single disconnected copies of 5-qubit spin chains averaged over 10^4 instances between precision $1 \leq p \leq 10$. The total height of the right hand composite bars (blue, green, yellow), shows the fraction correct of one repeat of three anti-ferromagnetically connected copies (in a loop, $J_F^{(p)} = J_F^{(\min)}$), of 5-qubit spin chains, averaged over 10^4 instances. The composite bars (right) are broken down into three possible cases: all three copies correct (blue, lowest), two of three copies correct (green, middle) and one of three copies correct (yellow, top). The error bars (black) at the top of each bar show that the effect is much larger than statistical effects.

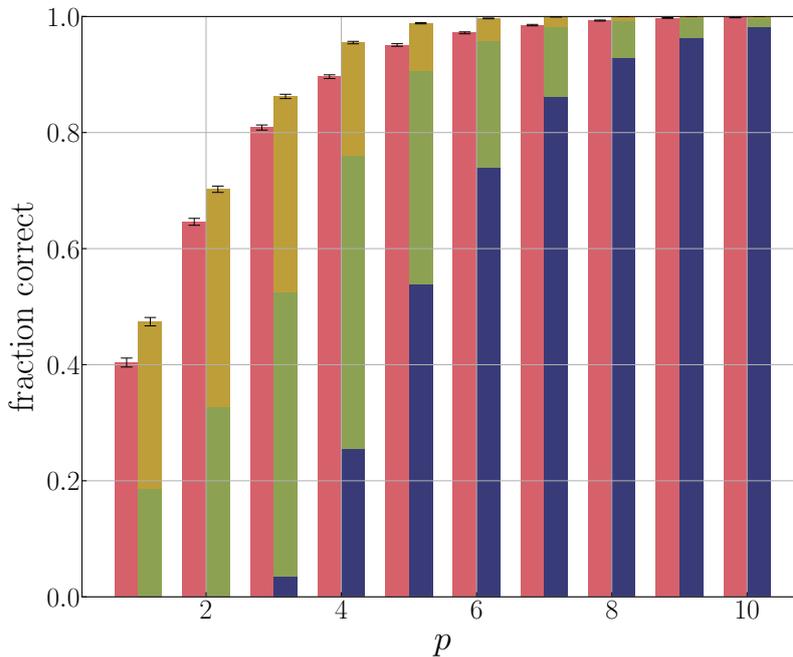


Figure 8.6: As figure 8.5 for 10^4 instances of 5-qubit Ising spin chains.

As with SK spin glasses, we see that for $1 \leq p \leq 10$, there is always an improvement in fraction correct for the three connected copies compared to the disconnected copies. Therefore we can make similar conclusions for spin chains as we did for SK spin glasses. Compared to SK spin glasses, for both disconnected copies and three connected copies, for low values of precision $1 \lesssim p \lesssim 4$, the fraction correct for spin chains is larger. At higher precision this difference becomes negligible. These results clearly show that the Ising spin chains behave in a similar way to the SK spin glasses, i.e. the anti-ferromagnetic links cause one or more copies to break and the frustration allows more low lying states to be accessed. This indicates that our error suppression method is effective for a range of precisions for at least two types of Ising model.

8.4 Four copies

8.4.1 Spin glasses in a loop

In order to see whether further improvement might be gained by increasing the number of copies connected in a loop, we next measured the effect on fraction correct of connecting four copies of 5-qubit SK spin glasses in a loop. Here precision was again applied using the deterministic random error model. In figure 8.7 the left

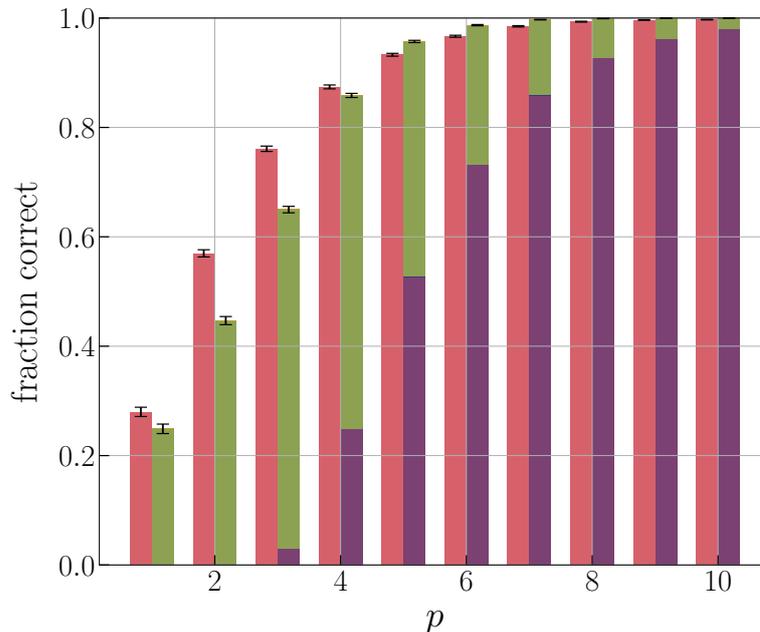


Figure 8.7: Fraction correct versus precision p for 10^4 instances of 5-qubit SK spin glasses. Left bars show fraction correct for four repeats of single copies (red). Right bars show fraction correct for three copies connected with links $J_F^{(p)} = J_F^{(\min)}$ split into all four correct (purple), three correct (dark blue), two correct (green), one correct (yellow).

hand bars (red) show the fraction correct of four repeats of a single disconnected copy of a 5-qubit spin glass, averaged over 10^4 instances, between the precisions $1 \leq p \leq 10$. The total height of the right hand composite bars (purple, blue, green, yellow), show the fraction correct of a single repeat of four copies of a 5-qubit SK spin glass connected in a loop with anti-ferromagnetic links $J_F^{(p)} = J_F^{(\min)}$, for the same precision range, averaged over the same 10^4 instances. The right hand composite bar is broken down into the four possible cases which count as success for our error suppression scheme: all four copies correct (purple, lowest), three copies correct (blue), two copies correct (green) and one copy correct (yellow, highest). The black error bars at the top of each of the bars show that this effect is much larger than statistical effects.

We see in figure 8.7, that similar to three copies in a chain and two copies in a chain, we only see improvement in fraction correct here for precisions $p \geq 5$. The disadvantage in connecting 4 copies in a loop here is stronger than for connecting two or three copies in a chain and the advantage when $p \geq 5$ is smaller. We can see in figure 8.7 that the majority of the fraction correct of connecting four copies together in a loop comes from having all four copies correct and from having two of

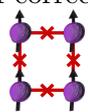
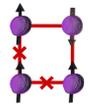
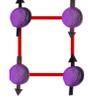
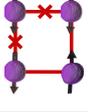
Copies correct	Energy
4 correct 	$E_4 + 4nJ_F^{(p)\min}$
3 correct 	E_3
2 correct 	$E_2 - 4nJ_F^{(p)\min}$
1 correct 	E_1

Table 8.1: Table showing configurations of four copies with possible frustrated (represented by crosses on the link) and unfrustrated links according to how many copies are correct (left column) and the energy of each configuration and its energy penalty/benefit (if present) according to the anti-ferromagnetic links. Here, each qubit represents a single copy.

four copies correct. There are no visible contributions to the fraction correct from cases with three copies correct or just a single copy correct.

We present the following argument for the proportions of contributions in figure 8.7. For simplicity, we think in terms of minimizing the energy penalty of the anti-ferromagnetic links between copies, which means states have to be opposite to each other. However, it is important to note that in the real simulations, it is always a compromise between minimizing the energy penalty and minimizing the optimisation problem. Therefore we are more likely to get intermediate states with a few bits flipped compared to the original. Nevertheless, the minimizing energy penalty argument is a useful way to try and understand what is occurring these systems.

In table 8.1, in the left hand column, we have drawn configurations of four copies (each copy is represented by a single qubit here), that have four, three, two and one copy correct. In the right hand column we have written the energy of each of the configurations, including the energy penalty/benefit given by the anti-ferromagnetic links (if present). We have described the energy of four, three, two and one copy correct (with the others incorrect) as E_4 , E_3 , E_2 and E_1 respectively. This is an oversimplification as the incorrect states could vary in energy, but is

useful for this explanation.

In terms of the energy penalty introduced by the anti-ferromagnetic links: this penalty is maximised when all four copies have the same state. This can be seen in the first row of table 8.1. When all four copies have the ‘correct’ ground state, this is represented by the purple section of the right hand bars in figure 8.7 . On the other hand, this penalty is minimized (acts as a maximum benefit) when there are two and two alternating opposite states, see third row of table 8.1. When two of these same copies have the ‘correct’ ground state, this is represented by the green section of the bar in figure 8.7 . If we try and introduce a third correct copy to the two correct copies, the energy penalty is increased by $4nJ_F$ (see second row of table 8.1), so if the energy difference gained by having the third correct copy is less than this energy penalty, then the three copy correct configuration is not energetically viable. We face the same penalty in reverse if we try to reduce the number of correct copies by one, to have just one correct copy. If the reduction in energy between the one correct copy configuration and the two correct copies configuration is not greater than $4nJ_F$, then the one correct copy configuration is also not energetically viable.

We further note that the frustration present in a loop of three copies is not present in a loop of four. This may further explain, why we no longer see such a strong benefit to connecting copies in a loop of four.

8.5 Five copies

8.5.1 Spin glasses in a loop

We hypothesised that the main reason for such a strong improvement in fraction correct for three copies connected in a loop, compared to three copies connected in a chain and four copies connected in a loop, was the presence of frustration in the loop of three, which may allow more low lying states to be accessed and therefore increase the chance of recovering the ‘correct’ ground state. In order to test our theory that frustration is a key element of our error suppression scheme, we next measured the effect on fraction correct of connecting five copies of 5-qubit SK spin glasses in a loop. For five copies connected in a loop we recover the frustration that was present in a loop of three but lost in a loop of four, so we might expect that we recover a strong improvement in fraction correct if our frustration theory holds. Here precision was applied using the deterministic random error model.

In figure 8.8, the left hand bars (red) show the fraction correct of five repeats of a

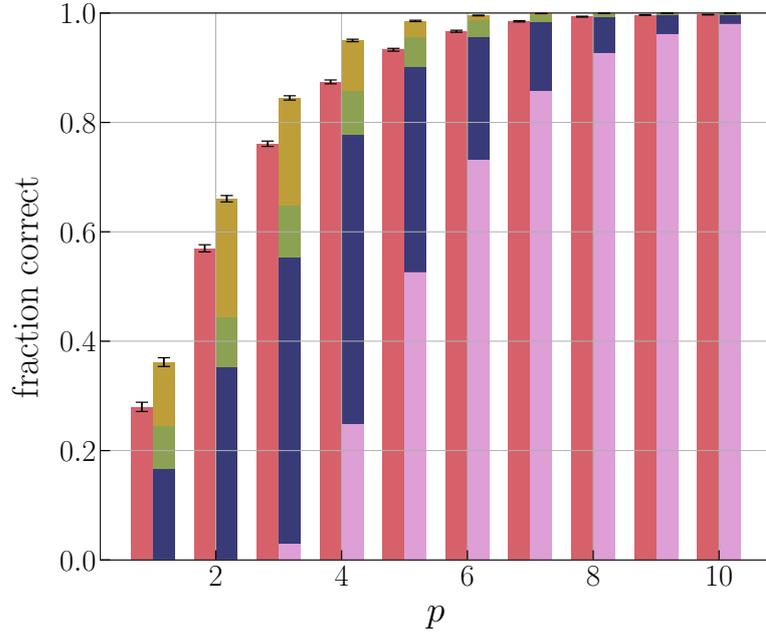


Figure 8.8: Fraction correct versus precision p for 10^4 instances of 5 qubit SK spin glasses. Left bars show fraction correct for 5 repeats of single copies (red). Right bars show fraction correct for 3 copies connected with links $J_F^{(p)} = J_F^{(\min)}$ split into all 5 correct (light pink), 4 correct (purple), 3 correct (dark blue), 2 correct (green), 1 correct (yellow).

single disconnected copy of a 5-qubit SK spin glass, averaged over 10^4 instances, between the precisions $1 \leq p \leq 10$. The total height of the right hand composite bars (pink, purple, blue, green, yellow), shows the fraction correct of a single repeat of five copies of a 5-qubit SK spin glass connected in a loop with anti-ferromagnetic links of strength $J_F^{(p)} = J_F^{(\min)}$, for the same precision range, averaged over the same 10^4 instances. The right hand composite bar is broken down into the five possible cases which count as success for our error suppression scheme: all five copies correct (pink, lowest), four copies correct (purple), three copies correct (blue), two copies correct (green) and one copy correct (yellow, highest). The black error bars at the top of each of the bars indicate that this effect is much larger than statistical effects.

We see in figure 8.8, that we again recover the improvement in fraction correct from connecting the copies anti-ferromagnetically across all the precisions $1 \leq p \leq 10$. The height of the left and right hand bars in figure 8.8 are similar in value to those in figure 8.5 indicating a similar improvement is gained. In terms of the proportions of number of correct copies contributing to the total fraction correct, we see here that the fraction where all five copies are correct is similar in value to the fraction

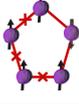
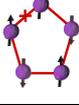
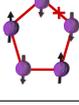
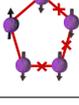
Copies correct	Energy
5 correct 	$E_5 + 5nJ_F^{(p)\min}$
4 correct 	$E_4 + nJ_F^{(p)\min}$
3 correct 	$E_3 - 3nJ_F^{(p)\min}$
2 correct 	$E_2 - 3nJ_F^{(p)\min}$
1 correct 	$E_1 + nJ_F^{(p)\min}$

Table 8.2: Table showing configurations of five copies with possible frustrated (represented by red crosses on the link) and unfrustrated links according to how many copies are correct (left column) and the energy of each configuration and its energy penalty/benefit (if present) according to the anti-ferromagnetic links. Here, each qubit represents a single copy.

where all three are correct, for three copies are connected in a loop. We see that the second largest contribution comes from the fraction where three of the five copies are correct. The third largest contribution comes from when there is just a single copy is correct, and there is a small contribution from when there are two copies correct. There is no visible contribution from when four copies are correct.

We present the following argument for the proportions of the contributions in figure 8.8. For simplicity, we think again in terms of minimizing the energy penalty of the anti-ferromagnetic links between copies.

In table 8.2, in the left hand column, we have drawn configurations of five copies (each copy is represented by a single qubit), that have five, four, three, two, one copy correct. In the right hand column we have written the energy of each of the configurations, including the energy penalty/benefit given by the anti-ferromagnetic links (if present). Despite the oversimplification (as with four copies), we have described the energy of five, four, three, two, and one copy correct (with the others incorrect) as E_5, E_4, E_3, E_2, E_1 respectively.

In terms of the energy penalty introduced by the anti-ferromagnetic links: this penalty is maximised when all five copies have the same state. This configuration can be seen in the first row of table 8.2. When all five copies have the ‘correct’ ground state, this is the situation which is represented by the pink section of the right hand bars in figure 8.8.

On the other hand the energy penalty is minimized (with the energy benefit maximized) when there are three and two of the same state (and one of the three is alternated with another state). This can be seen in row three and four of table 8.2. When three of five copies have the ‘correct’ ground state, this is the situation represented by the blue section of the right hand bars in figure 8.8. If we try to introduce a fourth correct copy to the three correct copy configuration, the energy penalty is increased by $4nJ_F$, so if the difference in energy by gaining one more correct copy is less than $4nJ_F$, the four copy correct situation is not energetically viable. Hence we do not see a contribution from four copies correct in figure 8.8.

On the other hand if we try to reduce the number of correct states from three correct to two correct, we do not receive an energy penalty from the anti-ferromagnetic links, (but we do from the optimisation problem), so we should only see two correct copies when the ‘incorrect’ state has a lower energy than the ‘correct’ state. This situation is represented by the green section of the right hand bars in figure 8.8. If we try and reduce the number of correct states from two to one, we could also face $4nJ_F$ penalty, if the four incorrect states were the same (i.e. completely flipped compared to the correct state). However, for this number of incorrect copies it is likely the copies will flip bits between each other, reducing the energy penalty in this way. This is why we still see contribution from one copy correct which is represented by the yellow section of the right hand bars in figure 8.8.

8.6 Comparison of five and three copies

In order to see how the improvement from five connected copies in a loop compares to the improvement from three connected copies in a loop of 5-qubit spin glasses, we measured the difference in fraction correct between the two.

Figure 8.9, shows the difference in fraction correct versus precision between five copies in a loop and three copies in a loop of 10^4 instances of 5-qubit spin glasses. Here we see that that for precisions $p = 1$ to $p = 4$, the mean of the difference is slightly above zero at around 0.0025 for $p = 1$ however this mean decreases closer to zero as p increases. Across the precisions tested here $[1, 10]$, the error bars of each data point overlap with zero, with maximum uncertainty of around ± 0.010 .

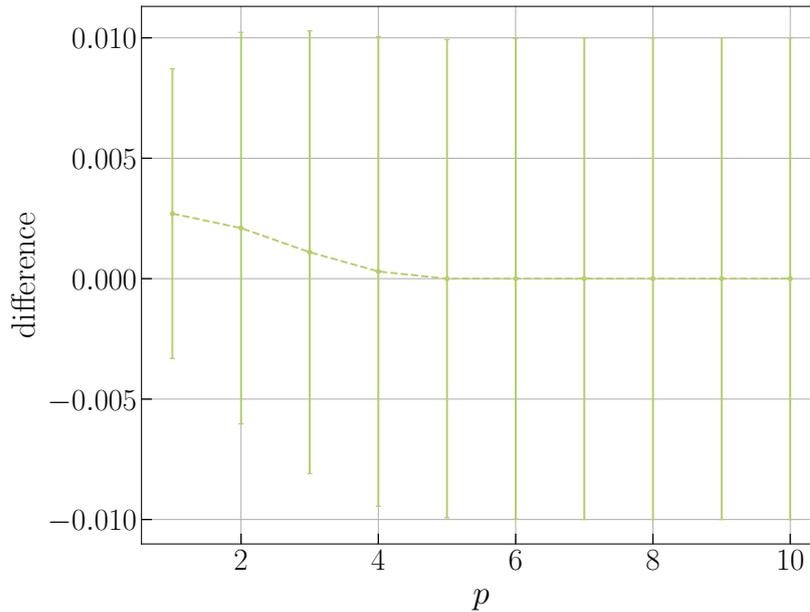


Figure 8.9: Difference in fraction correct versus precision p between 5 copies in a loop and three copies in a loop of 10^4 instances of 5 qubit SK spin glasses.

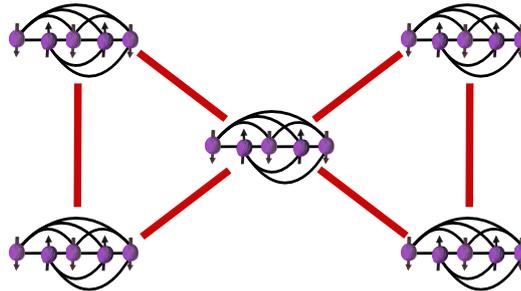


Figure 8.10: Diagram of a “bowtie” configuration of five copies of 5-qubit spin glasses. For simplicity only one link between models is shown (in practice there would be links between each corresponding qubit in each copy).

This indicates that the difference in fraction correct between five and three copies connected in a loop is consistent with zero. Hence we see no improvement from using five copies (though it would require more resources), however we also see no numerical disadvantage from using five copies.

8.7 Chapter summary

In this chapter, we measured the fraction correct versus the precision p in the range $1 \leq p \leq 10$, for two, three, four and five copy configurations of 5-qubit SK spin glasses and two and three copy configurations of 5-qubit spin chains. Copies were

connected by anti-ferromagnetic links of strength $J_F^{(p)} \simeq J_F^{\min}$ and we compared the fraction correct of the connected copies to the fraction correct of the equivalent number of repeats of single disconnected copies. We have found the following.

For two SK spin glass copies connected in a chain, applying precision according to the deterministic random error model, there was no advantage to connecting copies below $p = 3$, but there was a small advantage in connecting at higher precision. We hypothesised that this small advantage was due to the anti-ferromagnetic links between the copies, forcing the second copy to become correct when a precision error occurred on the first copy. When using the random error model, for two SK spin glass copies connected in a chain, the advantage we saw for the deterministic random error model was reduced. We saw no improvement in fraction correct from the connected copies compared to the two repeats of single disconnected copies below $p = 4$.

We next looked at the fraction correct versus precision, for two copies of spin glasses connected in a chain, applying precision according to the deterministic random error model. We saw similar results to the SK spin glasses. Compared to the fraction correct of two repeats of single disconnected copies, we saw no improvement below $p = 3$, for two copies connected anti-ferromagnetically in a chain. In all proceeding simulations, precision was applied using the deterministic random error model.

We next looked the fraction correct versus precision for three copies of SK spin glasses, connected anti-ferromagnetically in a **chain**. Here we saw a reduction in the improvement gained, compared to two copies of SK spin glasses connected in a chain. There was no advantage in fraction correct of three copies connected anti-ferromagnetically in a chain compared to three repeats of single disconnected copies below $p = 4$.

Next, we measured the fraction correct versus precision for three copies of SK spin glasses, connected anti-ferromagnetically in a **loop**. Here we saw an advantage in fraction correct for the three copies connected (in a loop) over the three repeats of single disconnected copies, across the range $1 \leq p \leq 10$ of precision p . We hypothesised that this large gain in improvement from three copies connected in a loop compared to three copies connected in a chain was due to the inclusion of frustration in the loop of three copies.

We next investigated whether the use of more copies could prove/disprove this hypothesis. We next connected four copies in a loop and found that there was no gain in fraction correct over four repeats of single disconnected copies for $p < 5$. For $p \geq 5$, we still saw a small improvement of four copies connected anti-

ferromagnetically in a loop, compared to four repeats of single disconnected copies. We note that there is no frustration present in a loop of four copies.

Next, we measured the fraction correct versus precision, of five copies connected with anti-ferromagnetic links in a loop. This five copy loop now contained frustration and we recovered a similar level of improvement in fraction correct (over five repeats of single disconnected copies) as three SK spin glass copies connected in a loop did (for three repeats of single disconnected copies). This result indicated frustration plays a key role in enhancing the fraction correct in our error suppression scheme. Finally we measured the difference in fraction correct versus precision between five copies in a loop and three copies in a loop. We found this difference to be consistent with zero, indicating no improvement from using five copies in a loop over three.

In future, the effect on fraction correct of different frustration containing configurations of five copies (such as a “bow-tie configuration” see figure 8.10), or more could be investigated. The improvement from these different configurations could be compared to see if particular configurations or number of copies gain more advantage than others.

Precision Improvements

In order to gain an insight into how our error suppression scheme affected individual instances on average, we next measured the conditional fraction correct against precision p for three anti-ferromagnetically connected (in a loop) copies of 5 and 9-qubit SK spin glasses and 5-qubit spin chains. In order to see if the improvement gained by our scheme continued for larger sizes of Ising model we then measured the fraction correct against precision p of three anti-ferromagnetically connected (in a loop) copies of 5, 6, 7, 8 and 9-qubit SK spin glasses and spin chains. Using this result we attempted to quantify in terms of number of bits of precision the improvement gained by connecting the three copies over three disconnected copies. For all simulations in this chapter the precision was implemented by the deterministic random error model.

In section 9.1, we look at the breakdown of conditional fraction correct versus precision given the outcome of the single disconnected copy. We do this for 5 and 9-qubit SK spin glasses and 5-qubit spin chains. In section 9.2, we measure fraction correct versus precision p , for three connected copies and three repeats of a single disconnected copy using the deterministic random error model, for SK spin glasses of size 5, 6, 7, 8, and 9-qubits. In section 9.3 we calculate a quantitative estimate in terms of bits of precision of the improvement gained from our scheme versus precision, in the deterministic random error model, for both spin glasses 9.3.1 and spin chains 9.3.3. In section 9.4 we summarize this chapter.

The contributions for this chapter were as follows. Nick Chancellor and Adam Callison provided discussion on the explanation as to the levelling off of precision improvement at higher values of precision. The branch and bound code used to classically calculate ground states and solution states in this chapter was written by Adam Callison.

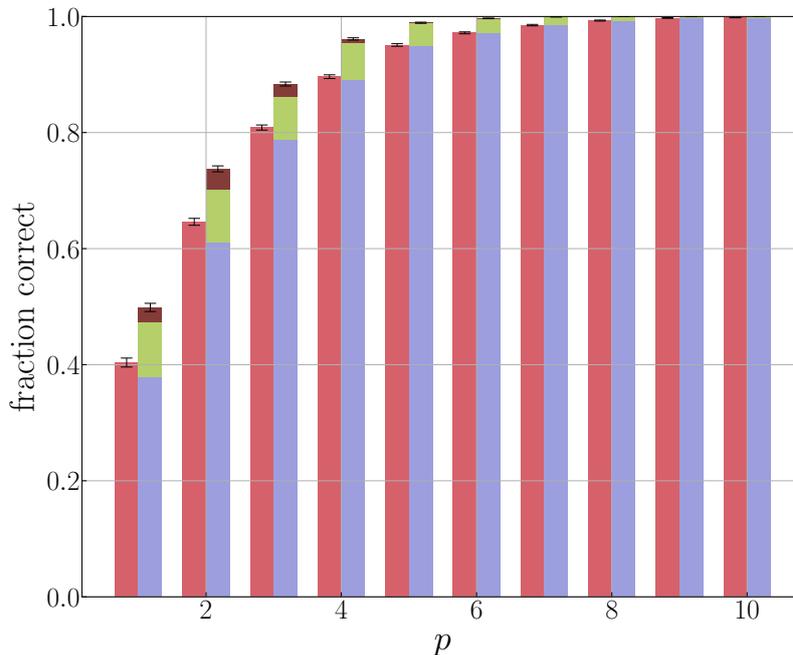


Figure 9.1: Fraction correct versus precision p for 10^4 instances of 5-qubit spin chains. Left bars show fraction correct for a single copy (red). Right bars show fraction correct broken down into cases where: both single copy and three copies are correct (bottom, light blue); the single copy is incorrect but the three copies have at least one correct (middle, light green); the three connected copies are incorrect but the single copy is correct (top, dark red).

9.1 Conditional fraction correct

In section 8.3.2 and 8.3.3, we saw that three copies of 5-qubit SK spin glasses and spin chains, connected anti-ferromagnetically with $J_F^{(p)} = J_F^{(\min)}$ in a loop gained an improvement in fraction correct compared to three repeats of single disconnected copies, for precisions $1 \leq p \leq 10$, where precision was applied using the deterministic random error model. We note that this improvement in fraction correct was an average effect with some instances seeing a strong benefit from being connected anti-ferromagnetically in a loop, but other instances seeing no improvement or disadvantageous effects. Hence, another useful measure for us would be, how often connecting the three copies, caused us to lose the correct ground state i.e. the fraction of times all three copies had the incorrect ground state, given the ground state of the single copy prior to connection was correct. We calculated this conditional fraction correct using a different analysis on the same data, as was collected for the plots in sections 8.3.2 and 8.3.3. The results are shown in figures 9.1 and 9.2.

First we look at the results from spin chains. In figure 9.1, the left hand bars (light red) show the fraction correct of three repeats of single disconnected copies of 5-qubit spin chains averaged over 10^4 instances between precision $1 \leq p \leq 10$. These bars (and the data they present) are exactly the same as the left hand bars in figure 8.6 and are provided as a comparison. The height of the right hand composite bars (light blue, light green, dark red) show the conditional fraction correct of three anti-ferromagnetically connected copies (in a loop, $J_F^{(p)} = J_F^{(\min)}$) of 5-qubit spin chains, averaged over 10^4 instances. The composite bars are broken down into the three following cases: a) the single copy was *correct* and the three connected copies had at least one correct copy *correct* (light blue, lowest), b) the single copy was *incorrect* but the three connected copies had at least one copy *correct* *correct* (light green, middle) and c) the single copy was *correct* but all three connected copies were *incorrect* (dark red, top). We consider a) – acceptable, b) – good and c) – bad, in terms of our suppression scheme. The error bars (black) at the top of each of the bars show that the effect is much larger than statistical effects. The sum of the light blue and light green bars is the same as the total height of the right hand bars in figure 8.6, i.e. all the cases where the three copies give the correct result. The sum of the dark red and light blue bars is the same as the height of the light red left hand bars (in the same figure), i.e., all the cases where the single copy is correct.

In figure 9.1 we see that, in the vast majority of cases our error suppression scheme has no effect on the outcome of finding the correct ground state (light blue section), however crucially, for a small but significant number of instances, connecting the three copies means that we are able to recover the correct ground state despite it being lost for the single copy (light green section). We also see that there is only a very small fraction of instances, where connecting the three copies causes the correct ground state to be lost. Though the increase in fraction correct is small, we note that by running a single (disconnected) copy as well as the 3 connected copies (on a machine one third the size and therefore cheaper) and comparing the candidate solutions to choose the best, allows us to include the dark red section of the right hand bars into our total fraction correct. Therefore in the proceeding figures and data, we now include the fraction of instances represented by the dark red section (where the single disconnected copies were correct but the three connected copies were incorrect) into our measurement of total fraction correct.

The equivalent figure for 5-qubit SK spin glasses is shown in figure 9.2. Its results are similar to the 5-qubit spin chain results (in figure 9.1), indicating that the small but significant number of instances where our error suppression scheme provides a benefit are also crucial for spin glasses. As before, we see that for low precisions

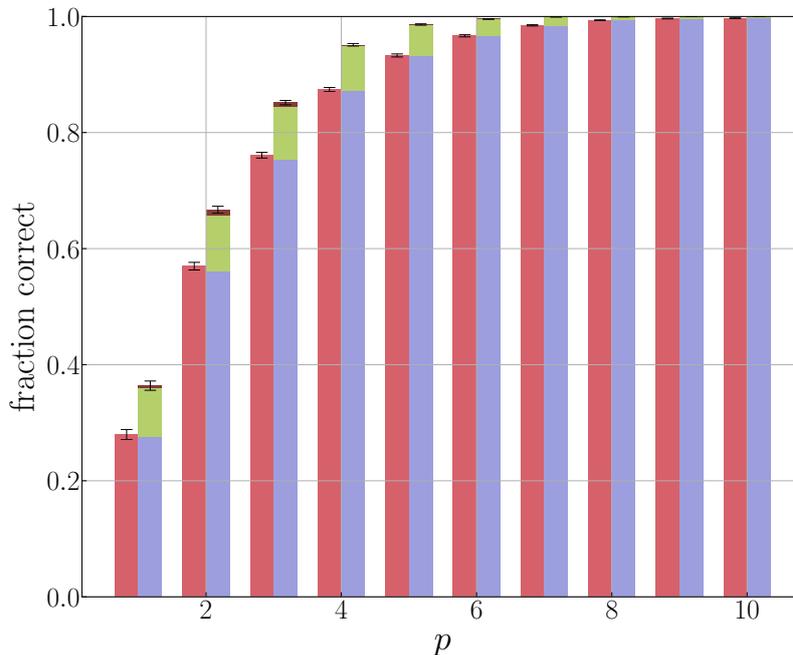


Figure 9.2: As figure 9.2 for 10^4 instances of 5-qubit SK spin glasses.

$1 \lesssim p \lesssim 4$, the fraction correct for the SK spin glasses is lower than for the spin chains. However, we also see at this low precision, that the fraction of instances represented by the dark red section (where connecting three copies causes the ground state to be lost) is smaller for SK spin glasses. For both spin glasses and spin chains, at higher precisions $p \geq 5$ the size of the dark red section reduces so that it becomes negligible.

The equivalent figure for 9-qubit SK spin glasses is shown in figure 9.3. It is similar to the 5-qubit results in figure 9.2, indicating the improvement in fraction correct gained from using our three connected copies technique still exists at this larger size, but shows how larger system size problems are broken more easily by low precision, as would be expected.

9.2 Precision improvements

We next wanted to get an estimate of the improvement gained from connecting three copies anti-ferromagnetically in a loop, opposed to three repeats of a single disconnected copy, in terms of bits of precision. We also wanted to see how this improvement varied over different sizes of SK spin glasses. For this estimation we used the fraction correct against precision p data collected using the deterministic

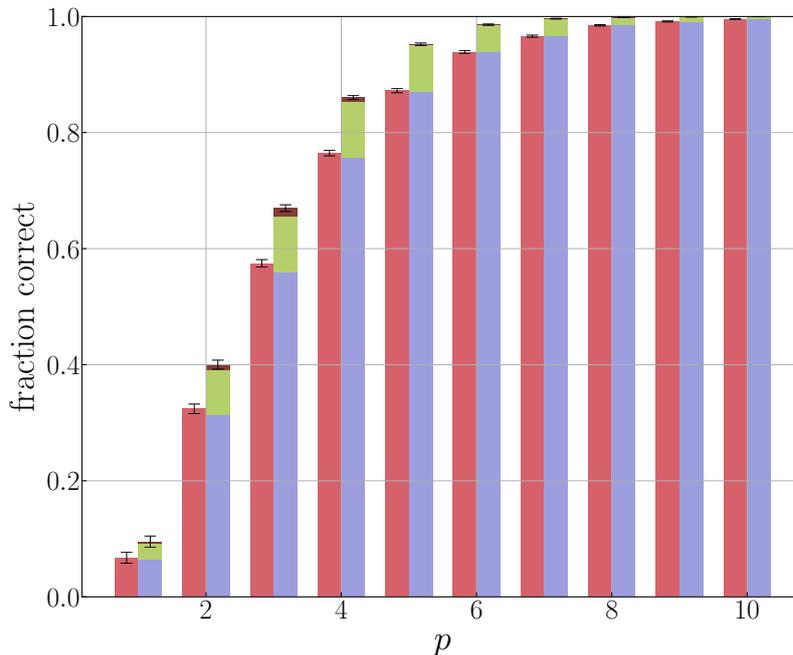


Figure 9.3: As figure 9.2 for 10^4 instances of 9-qubit Ising spin glasses.

error model for 5, 6, 7, 8 and 9-qubit SK spin glasses.

Figure 9.4 shows the total fraction correct (the peaks of the bars in figures 9.2 and 9.3) against precision p for three repeats of single disconnected copies (solid lines) and three connected copies plus a separate single copy (dotted lines) for 10^4 instances of SK spin glasses for $n = 5$ (top left), 7 (top right), 8 (bottom left) and 9 (bottom right). Not shown, $n = 6$ is very similar. The limit we can compute using reasonable computational resources is $n = 9$, for which three copies requires 27 qubits.

Looking at 9.4 we can see that for each of the four SK spin glass sizes, the improvement gained from using the three connected copies (plus a separate single copy), rather than an equivalent number of repeats of a single disconnected copy, is still present. We also see that in each case, single copy fraction correct approaches one around $p = n + 1$. The intuition for this is that there are only 2^n different possible states for n -qubits, therefore when increasing precision, we eventually reach a point where it will no longer improve the fraction correct of the ground state. This is because on average the gap between the energy levels reduces as 2^{-n} as the number of qubits n in the Ising model increases. When the errors in the fields and couplings due to the precision become smaller than the average gap, they are unlikely to cause the ground state to change.

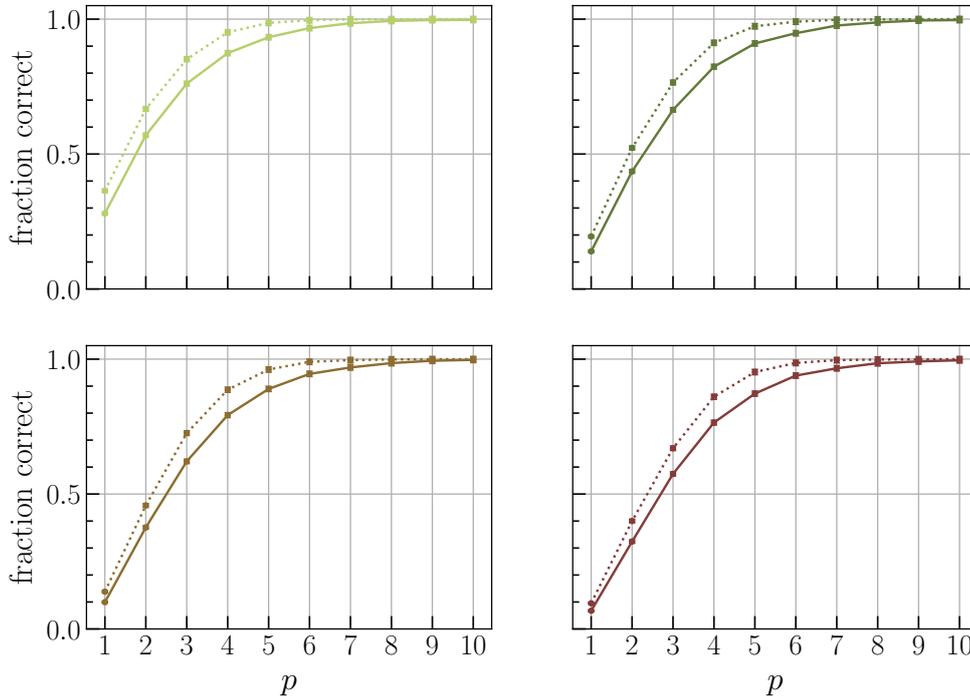


Figure 9.4: Fraction correct vs precision p for 10^4 instances of 5 (light green), 7 (dark green), 8 (brown) and 9 (red) qubit spin glasses with four repeats of a single copy (solid) and three connected copies plus 1 separate single copy (dotted).

Figure 9.5 shows the average energy gap size (blue points) calculated from the single copy Hamiltonians over the 10^4 instances of SK spin glasses of sizes 5, 6, 7, 8, and 9 qubits. Error bars for each data point are plotted but too small to see. A fit to this data of the form $a \exp(-bx)$ is plotted in red, where $a = 1.09 \pm 0.00$ and $b = 0.71 \pm 0.01$. As $2^{-n} = e^{-n \ln(2)}$, we can see this fit closely agrees with the exponential reduction in gap between energy levels as the number of qubits n increases.

For the three connected copies (plus a separate single copy), the fraction correct (dotted lines) tends to approach one sooner. In other words, the three connected copies still find the correct solution at lower precision than for a single copy. This is the effect we are looking for, to protect against lack of precision in the hardware.

9.2.1 Scaling of fraction correct with n

Having seen that the improvement gain from connecting three copies plus a separate single copy still continues to exist for SK spin glasses up to 9 qubits, we next

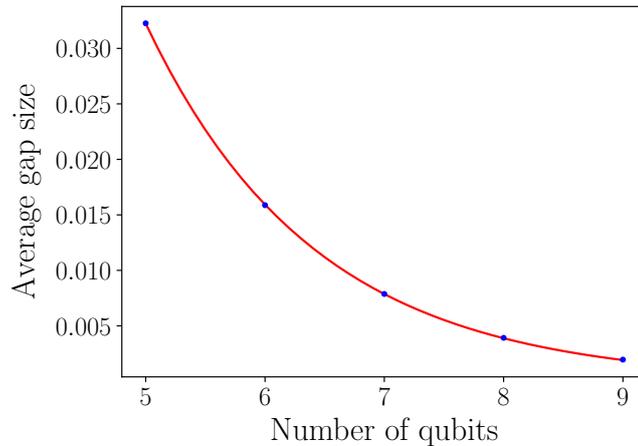


Figure 9.5: Average energy gap size (normalised) against the number of qubits in a SK spin glass. Average gaps size data (blue) were calculated using the single copy Hamiltonian of 10^4 instances at each spin glass size, 5 to 9. An exponential fit to this data of the form $a \exp(-bx)$ is plotted in red, where $a = 1.09 \pm 0.00$ and $b = 0.71 \pm 0.01$.

wanted to see how this improvement scaled with n . We can see from figure 9.4, that the fraction correct of both disconnected and connected copies scales down with increasing problem size. However, we wanted to see the extent of this scaling and whether there was any scaling reduction in the size of the improvement as the SK spin glasses increased in size.

Figure 9.6, shows the fraction correct at $p = 3$, versus n for four repeats of disconnected single copies (blue dots) and three connected plus one separate copies (green dots) of 5, 6, 7, 8, and 9 qubit SK spin glasses. We can see that for both the disconnected copies and the three connected plus one disconnected copies, the fraction correct scales down with increasing problem size in a linear fashion.

For both datasets we have applied a linear fit of the form $y = mx + c$, which are plotted in figure 9.6. The four repeats of single disconnected copies fit is a blue solid line and the three connected plus one disconnected copies fit is a dotted green line. The gradient value of the single disconnected copies fit is -0.047 ± 0.001 and its intercept value is 0.994 ± 0.005 . The gradient value of the three connected plus one separate copy fit is -0.0045 ± 0.002 and its intercept value is 1.076 ± 0.011 . The uncertainties of the both the gradients overlap indicating their similarity. This indicates that despite the reduction in fraction correct in both the disconnected copies and the three connected copies plus one separate copy as n increases, the gap between the two configurations appears to remain constant. If this scaling continues to larger sizes, it indicates that the improvement from our error suppression scheme

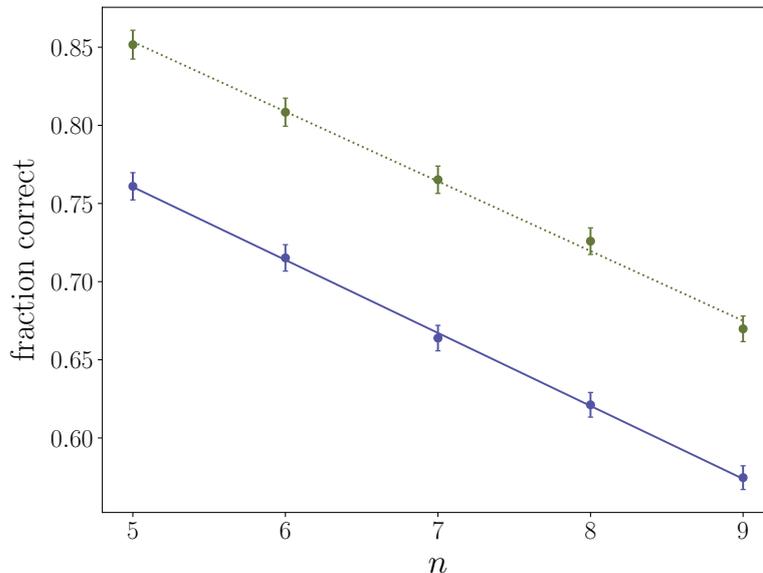


Figure 9.6: Fraction correct at $p = 3$ vs problem size n , for four repeats of single disconnected copies (blue dots) and three connected copies plus one separate copy (green dots), for 5, 6, 7, 8, and 9 qubit SK spin glasses. A linear fit of the form $y = mx + c$ is also fitted to both datasets. The fit for single copies (blue dotted line) has $m = -0.047 \pm 0.001$ and $c = 0.994 \pm 0.005$. The fit for three connected copies plus one separate copy (green solid line) has $m = -0.045 \pm 0.002$ and $c = 1.076 \pm 0.011$.

should continue to larger sizes.

9.3 Quantifying improvements

9.3.1 Spin glasses

The gap between the solid lines and dotted lines in figures 9.4 and 9.7 at each of the sizes indicates that the improvement gained by connecting three copies with anti-ferromagnetic links in a loop is still present for each of these SK spin glass sizes. Next we wanted to quantify this improvement in terms of the bits of precision gained. We did this by looking at the fraction correct of three-connected-copies-plus-a-single-disconnected-copy for the same 10^4 instances of SK spin glasses and comparing it to the fraction correct of four repeats of a single disconnected copy. Here we used the random deterministic error model, which meant there was negligible difference in fraction correct between one repeat of a single disconnected copy and four repeats of a single disconnected copy.

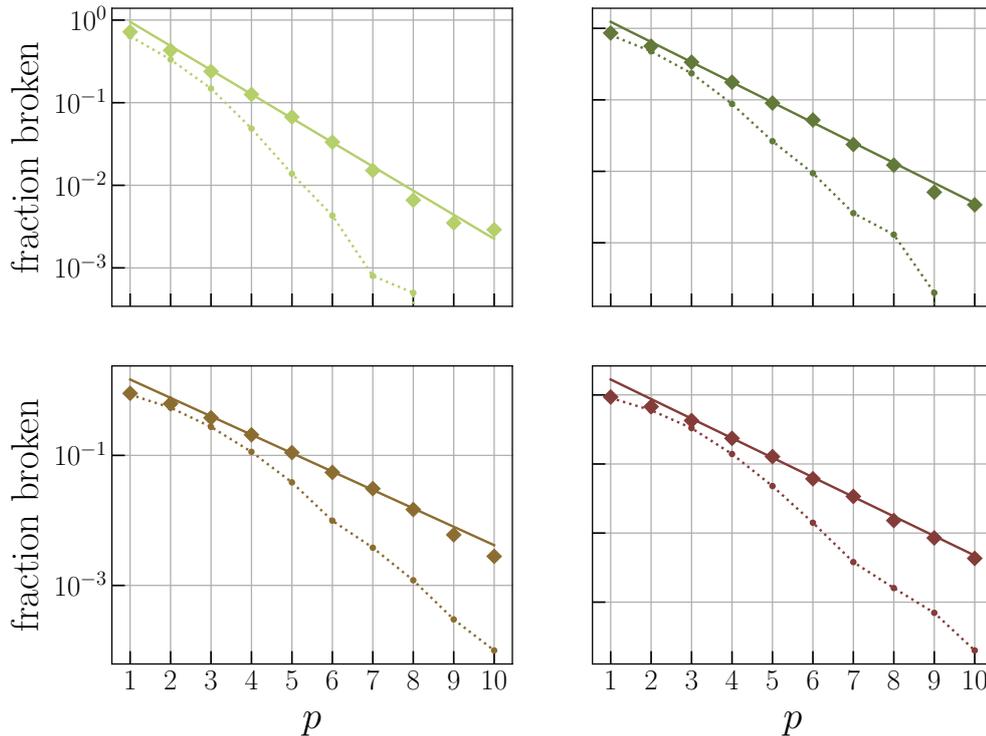


Figure 9.7: Same data as figure 9.4, but plotted as fraction broken = $1 -$ fraction correct on a log scale, with a linear fit to the single copy data (solid lines) excluding the first two points ($p = 1, 2$) from the fit.

We defined the *precision improvement* as the difference in bits of precision p (the horizontal displacement) between the single-disconnected-copies and three-connected-copies-plus-single-copy at the same value of fraction correct. In figure 9.4, this is the horizontal distance between the dotted and solid lines. However, as the data points on the dotted line (three-connected-copies-plus-single-copy) do not in general have corresponding data points on the solid line (single-disconnected-copies) at the same fraction correct, we calculated *precision improvement* using the data points for three-connected-copies-plus-single-copy data and estimating the effective precision for the single-disconnected-copies using an exponential fit to the single-disconnected-copies data. The exponential fit was of the form $f(p) = A \exp(-bp)$ where fraction correct = $(1 - f(p))$. All fits exclude the single-disconnected-copies data points at $p = 1$ and $p = 2$, which exhibit non-exponential effects due to the low precision. These fits are plotted in figure 9.7, alongside the single-disconnected-copies and three-connected-copies-plus-single-copy data on a log fraction broken = $(1 -$ fraction correct) vs precision p plot.

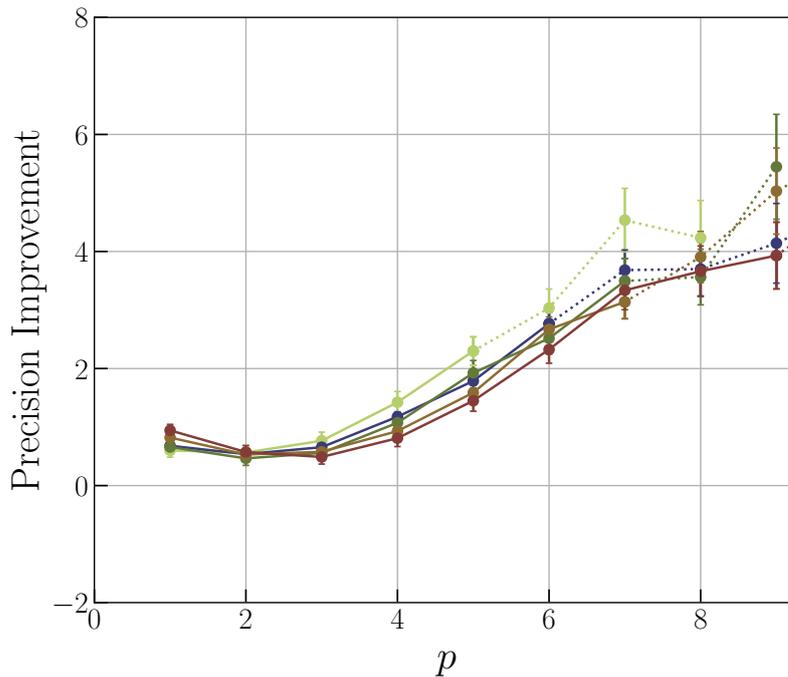


Figure 9.8: Precision improvement in bits versus precision, for 10^4 instances of 5 (light green) 6 (blue), 7 (dark green), 8 (brown) and 9 (red) qubit spin glasses. Precision improvement was calculated between a linear fit to the 4 repeats of a single copy data and the 3 connected copies plus 1 single copy data at the same value fraction correct. Data below $p = n$ is plotted as solid lines. Data above $p = n$ is plotted as dotted lines.

9.3.2 Spin glasses

The precision improvement vs p for 10^4 instances of 5, 6, 7, 8, and 9 qubit SK spin glasses is plotted in figure 9.8. The error in the variables A and b (arising from the fit) and the error in fraction correct (from the three-connected-copies-plus-single-copy data points), were combined via a functional approach [170] in order to calculate the error in the single-copy precision. This is the error shown by the error bars in figure 9.8. The precision improvement data for 5, 6, 7, 8 and 9 qubit SK spin glasses is shown in light green, blue, dark green, brown and red respectively. Due to the single copy fraction correct approaching one around $p = n + 1$, for the reasons explained in section 9.2, the precision improvement data $p > n$ is plotted as dotted lines whereas the data $p \leq n$ is plotted in solid lines.

In figure 9.8, it can be seen that, for each size of SK spin glass $n = 5$ to 9, between precision $p = 2$ to around $p = 7$, as precision increases, the precision improvement also increases, from less than 1 at $p = 2$ to around 3 at around $p = 7$. At $p = 1$ the precision improvement appears to be larger than at $p = 2$ in some cases, but

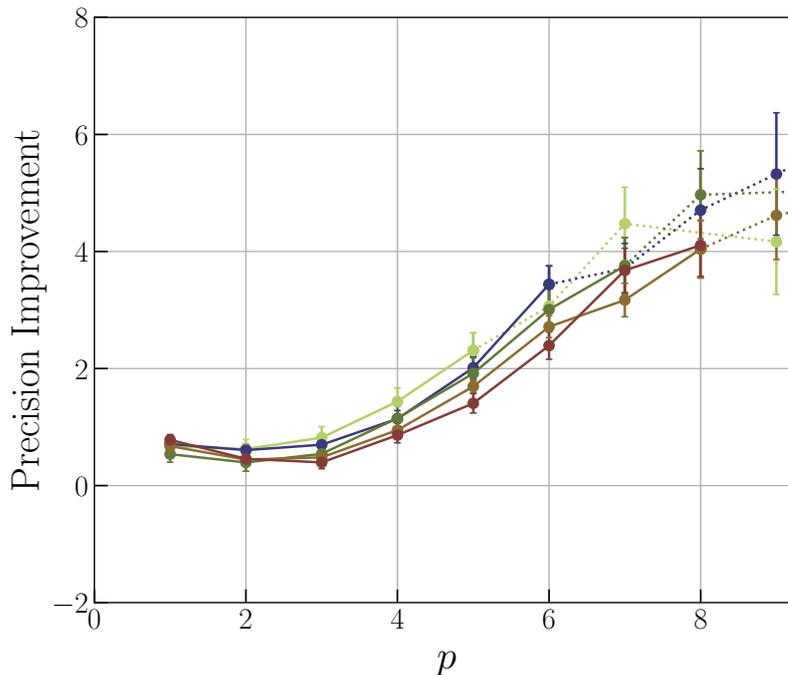


Figure 9.9: As figure 9.8, for 10^4 instances of 5-qubit Ising spin chains.

this is due to the exponential fit not matching the data at this very low precision. Data points for $p \gtrsim 6$ have larger error bars due to the small number of incorrect single copy counts at high precision.

9.3.3 Spin chains

Using the same method, we also calculate the *precision improvement*, for 10^4 instances of spin chains of sizes 5, 6, 7, 8, and 9 qubits. We also see that between precision $p = 2$ to around $p = 7$, as precision increases, the precision improvement also increases. We also see that the precision improvement starts off close to 1 at $p = 2$, for the spin chains. Interestingly we see that for the spin chains the precision improvement approaches 3, slightly sooner than for the spin glasses at $p = 6$. As with the spin glasses the apparent higher precision improvement at $p = 1$ is due to the exponential fit not matching the data at this precision. We also see larger error bars for data points $p \gtrsim 6$ due to the lower number of incorrect single copy counts at this higher precision.

In both figures 9.8 and 9.9, we see a suggestion in the data of a levelling off in the improvement, for higher p . This is expected for the finite resources of three copies. As this analysis was done only for small problem sizes, we cannot reliably extract trends for larger problem sizes. However, these results offer proof-of-concept that

our error suppression method can provide several extra bits of precision by using three anti-ferromagnetically linked copies compared to a single disconnected copy.

9.4 Chapter summary

In this chapter we first looked at the conditional fraction correct, i.e. the fraction correct of the three connected copies given the outcome of the single disconnected copies. This was split into three cases: 1) the fraction when both three connected and single disconnected copies were correct, 2) the fraction when the three connected copies were correct but the single disconnected copies were incorrect and finally 3) when the three connected copies were incorrect but the single disconnected copies were correct. We found that for 5 and 9-qubit spin glasses and 5-qubit spin chains, the majority of cases were unaffected by our error suppression scheme, but a small significant proportion fell into case 2) where our error suppression scheme was having a beneficial effect on them. For all models tested here, there was only a very small proportion of case 3), meaning that there was only a small negative effect from our error suppression scheme. We also noted that we could run a single copy of each instance alongside our three connected copies in order to avoid this small negative effect proportion. The small negative effect was larger at low precision for spin chains than SK spin glasses (5-qubits) but at higher precision this proportion became negligible for both types of Ising model.

We next repeated the measurement of fraction correct versus precision for SK spin glasses of larger sizes in order to see if the benefit gained from connecting three copies anti-ferromagnetically in a loop was still present. We measured the fraction correct of three anti-ferromagnetically connected copies plus a single disconnected copy of 5, 6, 7, 8 and 9-qubit SK spin glasses versus precision and compared this to the fraction correct of four repeats of a single disconnected copy at each size. We found that the benefit gained from using our error suppression scheme was still present for SK spin glasses of up to size 9-qubits. Due to computational cost spin glasses with more qubits were not tested. We noted that at each size the fraction correct reached one at about $p = n + 1$. we hypothesised that this was due to the reduction in average gap between the energy level by 2^{-n} .

We next aimed to quantify the improvement gained by using our error suppression scheme in terms of number of bits of precision gained between the fraction correct of the four repeats of single disconnected copies and the fraction correct of the three connected copies plus single disconnected copy. This could be calculated by measuring the horizontal distance between the three-connected-copies-plus-single-copy data points and a exponential fit to the single-disconnected-copy data. We

did this for both SK spin glasses and spin chains and got similar results. For SK glasses the *precision improvement* reaches roughly 3 bits at $p = 7$ before showing signs of levelling off. For spin chains the *precision improvement* was slightly higher, reaching roughly 4 bits at $p = 7$ before showing signs of levelling off.

Link Selection Protocol

In chapter 7.2, we saw that connecting copies of SK spin glasses with anti-ferromagnetic links was beneficial for some instances but not others. In this chapter, to evaluate further why this might be the case, we analysed single and multiple copies of 2-qubit Ising models, like the one in figure 10.1, gaining knowledge of their behaviour. Using this knowledge, we developed a link selection protocol which suggests whether or not to connect qubits between copies. We then tested the effectiveness of this protocol on 5-qubit SK spin glasses and spin chains.

In section 10.1, we presented analysis of the location of the ground states of a 2-qubit Ising models with either: two positive or negative fields, or one positive and one negative field. Looking at the strength of the fields h_1 and h_2 versus the coupling J_{12} , we locate thresholds in J_{12} where ground states change from one to another. In section 10.1.1, we looked at the energy levels of specific 2-qubit examples, designed to be above and below the predicted threshold. Alongside the energy levels, we also looked at the predicted energy shifts, for each of the states when connected anti-ferromagnetically to one copy of the model in its ground state. We used these findings to hypothesise on whether this affects the effectiveness of our scheme in figure 10.6 and 10.8. In section 10.2, we measured fraction correct vs J_F , for two and three copy examples, above and below the predicted J_{12} thresholds. In section 10.4, we present a link selection protocol that aims to predict whether corresponding qubits in separate copies should be connected or left disconnected. In section 10.5, we tested the effectiveness of our link selection protocol, by measuring the fraction correct vs precision for two and three copies of 5-qubit SK spin glasses and spin chains. In section 10.6, we summarize the chapter.

10.1 Threshold behaviour

In chapter 7.2, we found that connecting copies anti-ferromagnetically provided an improvement in fraction correct above separate copies for some Ising model instances but not others. To try and understand why this was the case, we first look at how the strength of the fields h_1 , h_2 versus the coupling J_{12} affects the ground state of a 2-qubit Ising model, like the example in figure 10.1. The Hamiltonian for this Ising model can be written,

$$\hat{H}_I = h_1 \hat{Z}_1 + h_2 \hat{Z}_2 + J_{12} \hat{Z}_1 \hat{Z}_2. \quad (10.1)$$

For a single copy of this 2-qubit Ising model, there are four possible ground states 11, 00, 01, and 10. The ground state of this model is dependent on the interaction of fields h_1 and h_2 and the coupling J_{12} between them.

When we connect a second copy of this Ising model with anti-ferromagnetic links, the Hamiltonian becomes,

$$\begin{aligned} \hat{H}_{2I} = & h_1 \hat{Z}_{1,1} + h_2 \hat{Z}_{2,1} + J_{12} \hat{Z}_{1,1} \hat{Z}_{2,1} \\ & + h_1 \hat{Z}_{1,2} + h_2 \hat{Z}_{2,2} + J_{12} \hat{Z}_{1,2} \hat{Z}_{2,2} + J_F \hat{Z}_{1,1} \hat{Z}_{1,2} + J_F \hat{Z}_{2,1} \hat{Z}_{2,2}, \end{aligned} \quad (10.2)$$

where $J_F < 0$ and equal to the minimum value allowed by the precision. For this system of two copies, there are 16 possible ground states. We can see that for this two copy system, with the addition of the two J_F terms, the ground state of this system is no longer just the solution to the 2-qubits Ising model, but is now

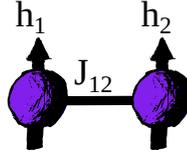


Figure 10.1: Single copy of a 2-qubit Ising model with fields h_1, h_2 and coupling J_{12} labelled.

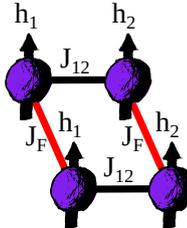


Figure 10.2: Two copies connected with anti-ferromagnetic links labelled J_F of a 2-qubit Ising model with the fields h_1, h_2 and coupling J_{12} labelled.

a compromise between the ground state of the 2-qubit Ising model and the anti-ferromagnetic links connecting the two copies, which want the states of the qubits in different copies to oppose.

If we first focus on a single copy, we can see that by varying the strengths of the fields and couplings in this 2-qubit model and analysing the effect on its ground state, we can find regimes in which each of the four possible states become the ground state.

We first consider the scenario where both fields h_1, h_2 are either both in the range $[0, 1]$ (positive) or $[-1, 0]$ (negative). We allow J_{12} to vary in the range $[-1, 1]$. We find the following regime which can be seen in figure 10.3. It can be seen that if the strength of J_{12} is small compared to the strength of the fields h_1, h_2 , the ground state will always be either 00 (when both fields are negative) or 11 (when both fields are positive). We can also see that there is a threshold value in the coupling strength J_{th} . Below this value the ground state of the model remains in 00 or 11. However, above this threshold value, the ground state becomes 10 (if both fields are negative and $h_1 < h_2$, or if both fields are positive and $h_2 < h_1$) or 01 (if both fields are negative and $h_2 < h_1$ or if both fields are positive and $h_1 < h_2$).

Next, we consider the scenario where one field h_i is in the interval $[0, 1]$ (positive, h_{+ve}) and the other field h_j is in the interval $[-1, 0]$ (negative, h_{-ve}). We allow the coupling strength J_{12} to vary in the interval $[-1, 1]$. We find the following regime which can be seen in figure 10.4. In this scenario, it can be seen that if the strength of J_{12} is small compared to the strength of the fields h_1, h_2 , the ground state will always be either 10 (when h_1 is positive) or 01 (when h_2 is positive). We can also see that there is a threshold value for the in the coupling strength J_{th} . Above this value, the ground state of the model remains 01 or 10. However, below this threshold value the ground state becomes 00 (if $|h_{-ve}| > |h_{+ve}|$) or 11 (if $|h_{-ve}| < |h_{+ve}|$).

We note that in both cases shown in figures 10.4 and 10.3, the threshold value of J_{12} indicates the point at which the coupling strength begins to dominate over the field strengths in the Ising model. We saw in figure 7.6 in section 7.3, that connecting 2 or 3 copies (anti-ferromagnetically or ferromagnetically) of 5-qubit Ising spin glasses with no fields (field strengths set to zero) is never beneficial compared to separate copies. Therefore we theorize, that this threshold value is a key value in determining whether or not our scheme will be successful for a particular instance. However, we see that as the number of qubits in the Ising model increases, determining this threshold value becomes more computationally intensive. An algorithm or heuristic which could estimate this threshold value would be extremely valuable. This motivates the further investigation and establishment of a link selection

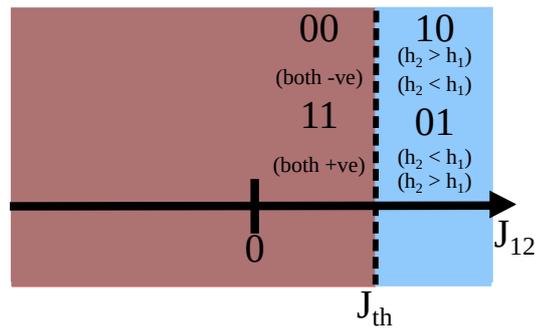


Figure 10.3: This figure shows the ground state of the 2-qubit Ising model when both fields are positive, or both fields are negative, and J_{12} is allowed to vary in the range $[-1, +1]$. We see at a positive value of J_{12} there is a threshold value J_{th} , where the ground state changes from 00 or 11 to 10 or 01.

protocol which proceeds in the rest of this chapter.

In table 10.1, we have written all the possible start states (if we start with J_{12} close to zero) and end states, depending on the relative strengths of h_1 and h_2 . Alongside each of these, we have written the inequality (i.e. the threshold value of J_{12}) that must be satisfied for each of the transitions to take place. Each of these inequalities has been derived because for the start state to change to the end state, the energy of the end state E_{end} must be less than the energy of the start state E_{start} . For example, for the state to change from 11 to 01, then $E_{01} < E_{11}$, this gives,

$$+h_1 - h_2 - J_{12} < -h_1 - h_2 + J_{12}. \quad (10.3)$$

We see that the terms h_2 cancel, as the field on this qubit is unaffected by the bit flip of the first. Hence we see the expression simplifies to $J_{12} > h_1$ which we see in the first row of table 10.1.

We see that in each of the cases in table 10.1, the transition relies in each case on J_{12} being larger in magnitude than just a single field h_1 or h_2 . This means in the two qubit Ising model the location of the threshold is easy to locate. However, we might expect that as the problem size increases, due to the addition of more fields and couplings the location of the threshold becomes harder to locate.

10.1.1 Energy levels below and above threshold

In order to establish how important this threshold value J_{th} was, in the positive fields only regime, we next predicted a value of J_{th} and created two 2-qubit Ising model examples: one which was below J_{th} (both $h_1 < J_{th}$ and $h_2 < J_{th}$) and another which was above J_{th} (either $h_1 > J_{th}$ or $h_2 > J_{th}$). We then analysed the energy levels of both these examples and looked at the shifts in the energy

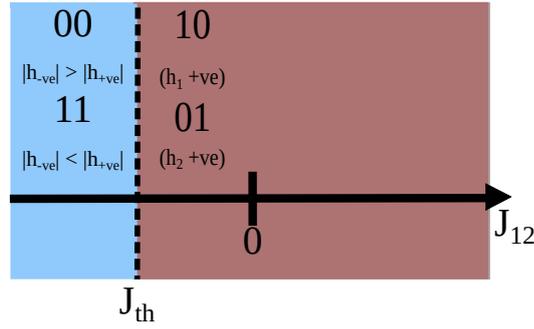


Figure 10.4: This figure shows the ground state of the 2-qubit Ising model when one field is positive and one field is negative, and J_{12} is allowed to vary in the range $[-1, +1]$. We see at a negative value of J_{12} there is a threshold value J_{th} , where the ground state changes from 10 or 01 to 00 or 11.

Start state	End state	Inequality
11	01	$J_{12} > h_1$
$(h_1, h_2, +ve)$	10	$J_{12} > h_2$
00	01	$J_{12} > h_2 $
$(h_1, h_2, -ve)$	10	$J_{12} > h_1 $
01	11 ($ h_1 < h_2 $)	$ J_{12} > h_2 $
$(h_1 -ve, h_2 +ve)$	00 ($ h_1 > h_2 $)	$ J_{12} > h_1 $
10	11 ($ h_2 < h_1 $)	$ J_{12} > h_1 $
$(h_1 +ve, h_2 -ve)$	00 ($ h_2 > h_1 $)	$ J_{12} > h_2 $

Table 10.1: Table showing for the 2-qubit Ising model, the starting state and end state, and the inequality (threshold) that must be satisfied for the transition to take place. In the top half the fields must be both positive or negative. In the bottom half one field is negative, the other positive.

levels that could be caused by connecting two of the same example together with anti-ferromagnetic links. Then we hypothesised on the effect of a lack of precision and how errors might propagate.

Our first 2-qubit example had fields of $h_1 = 0.4$ and $h_2 = 0.7$. We set the coupling strength to $J_{12} = 0.5$. As both h_1 and h_2 were positive, according to figure 10.3, we predicted the threshold value of J_{12} to also be positive. As the threshold value occurs when $J_{12} > h_1$ or $J_{12} > h_2$, we expected the threshold value for this example to be $J_{th} = 0.4$. Therefore, our first 2-qubit example was above threshold. We first plotted the energies of this example's four possible states (00, 01, 10, 11) in figure 10.5 (left). Our second 2-qubit example had the same value fields of $h_1 = 0.4$ and $h_2 = 0.7$, but we set the coupling strength to $J_{12} = 0.3$. Hence this example was

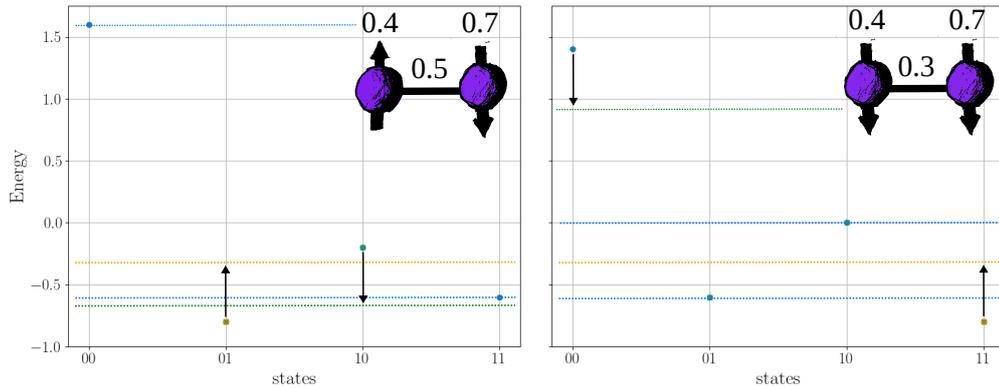


Figure 10.5: Left figure shows the energy levels (blue dots) of a 2 qubit example with a ground state of 01. Right figure shows the energy levels of a 2 qubit example with a ground state of 11 (blue dots). In both figures, the ground state of the single copy is highlighted in orange and the states that make the two connected copies ground state are shown with green crosses. The fields and coupling of both Ising models is shown in the top right corner of their respective plot. Arrows and dotted lines ((a)-(d)) represent the potential energy penalty/benefit gained by connecting the respective state with the ground state of the single copy.

below threshold. We plotted the energies of this example's four possible states in figure 10.5 (right).

In the top right corner of both of the graphs, there is a diagram of the two qubit Ising model it corresponds to. In both graphs, the blue dots represent the energy of each of the possible states of the 2-qubit model, with the energy of the ground state highlighted in orange. The two ground states which make up the total ground state of two anti-ferromagnetically connected copies of each of the examples are marked by green crosses. The black arrows (if present) and dotted lines represent the energy shifts which would occur if that particular state was connected anti-ferromagnetically (with $J_F^{\min} = 0.25$) to the ground state of the singular copy (01, left graph or 11, right graph).

In the left hand graph, we can see that 01 is the ground state of the singular copy (orange dot), whereas the ground state of the two copy system is 01 10 (green crosses). Looking at each of the four possible states we can see that connecting them each with 01 results in one of three situations: an energy shift up, no energy shift or an energy shift down. If states 00 or 11 connect with 01, the total ground state becomes 01 00 or 01 11 respectively. Neither of these states receive an energy shift from their connection using anti-ferromagnetic links. This is because one link is satisfied (0 to 1 or 1 to 0) whereas the other is frustrated (0 to 0 or 1 to 1), which results in the two J_F terms in the Hamiltonian (see 10.2) cancelling. On the other

hand when we connect 01 to 01, we receive a positive energy shift as both links are frustrated (0 to 0 and 1 to 1). This positive energy shift is likely to move this state above others. In this case we see that this positive energy shift moves the state 01 01 above both 01 11 and 01 10. Finally, we can see that if we connect 01 to 10, we receive a negative energy shift as both links are satisfied (0 to 1 and 1 to 0). This negative energy shift is likely to move this state below others. We see in this case that this negative energy shift moves the state 01 10, below both 01 11 and 01 01, making it the ground state of the two copy system (as represented by the green crosses on each of the participating states).

On the other hand, in the right hand graph we can see the ground state of the single copy is 11 (orange dot), whereas the ground state of the two copy system is 11 01 (green crosses). Looking at the arrows and dotted lines drawn on figure 10.5 (right), we can see that the state that receives the negative energy shift is now 11 00. However, as the energy of 00 is so high, the negative energy shift is not large enough for 11 00 to become the ground state. On the other hand, we see that 11 11 now receives a positive energy shift. This energy shift is large enough such that the state 11 01 (which does not receive an energy shift) becomes the ground state instead.

We see in these examples that when J_{12} is above threshold (left hand graph), the ground state of the two copy system will be 01 10 whereas when J_{12} is below threshold (right hand graph), the ground state will be 11 01. We hypothesise that the difference in ground states above and below threshold seen here may be the underlying cause of the difference in fraction correct behaviour between models.

This hypothesis may be reasoned in the following way. In most cases an error caused by the lack of precision is likely to cause a bit flip which changes the original ground state into the next lowest lying state. Continuing with our example from figure 10.5, if the ground state of the single copy 2-qubit model is 01, then an error due to a lack of precision is likely to cause the ground state to change to 11. This is shown in the top row (in green box) of the left hand diagram of figure 10.6. If this copy is connected to a second by minimum strength anti-ferromagnetic links the ground state of the system before the precision error is likely to be 01 10. After the precision error, the ground state of the system becomes 11 01. We see that despite the precision error on the first copy causing its original ground state to be lost, we find that due to the connection by anti-ferromagnetic links, the original ground state is instead retained on the second copy. A success for our scheme! This is described diagrammatically in the left hand diagram of figure 10.6.

On the other hand, if the ground state of the single copy 2-qubit model is 11, then an error due to a lack of precision is likely to cause the ground state to change

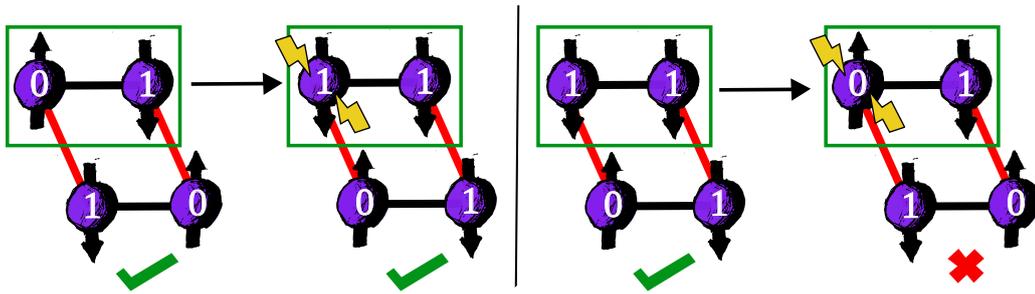


Figure 10.6: Left hand diagram shows the ground state of a two copies of a 2-qubit Ising model with two positive fields, that has J_{12} above the threshold, before a precision error has occurred (left) and after a precision error has caused the left qubit to flip (right). In this case the true precision single copy ground state is retained. Right hand diagram shows the 2-qubit Ising model with the same two positive fields, this time with J_{12} below the threshold, before a precision error has occurred (left) and after a precision error has caused the left qubit to flip (right). In this case the true precision single copy ground state is lost. In both cases ground state of the single copy only (before and after the precision error) is shown within the green boxes. The qubit on which the precision error (resulting in a bit flip) has occurred is indicated by yellow lightning. Under each two copy system a green tick (red cross) indicates that the system has retained (lost) the true precision single copy ground state.

to 01. This is shown in the top row (in green box) of the right hand diagram of figure 10.6. If this copy is connected to the a second by minimum strength anti-ferromagnetic links, the ground state of the system before the precision error is likely to be 11 01. After the precision error, the ground state of the system becomes 01 10. Unfortunately we see in this case that for both the original copy and second connected copy, the original ground state is lost. A failure for our scheme! This is described diagrammatically in the right hand diagram of figure 10.6.

This means the success of our scheme depends strongly on the ground state of the singular copy. In this scenario, we saw that the system with J_{12} above threshold *was* able to recover the ground state, whereas the system with J_{12} below threshold *was not*. This means that if we are able to predict the threshold of an Ising model, we may be able to predict whether or not it is beneficial to connect copies anti-ferromagnetically.

10.1.2 Inclusion of negative fields

We note that in the previous example, both fields in the 2-qubit Ising model were positive. This meant that according to figure 10.3, the threshold value of the coupling strength would be positive. Now we look at a 2-qubit example where one

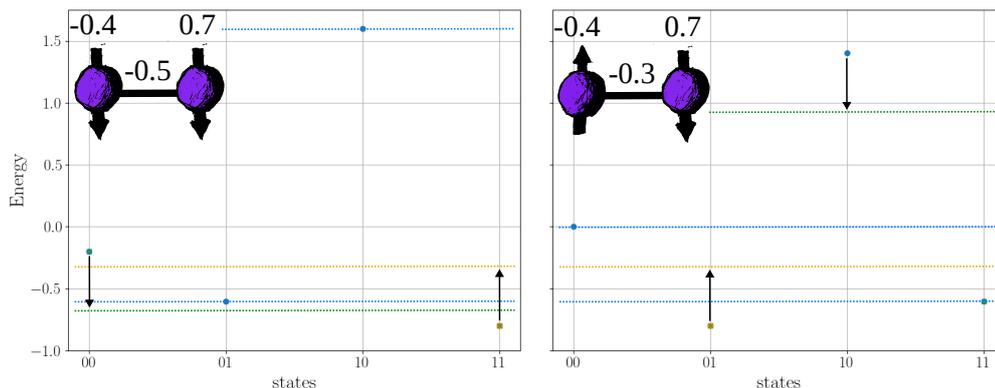


Figure 10.7: Left figure shows the energy levels (blue dots) of a 2 qubit example with a ground state of 01. Right figure shows the energy levels of a 2 qubit example with a ground state of 11 (blue dots). In both figures, the ground state of the single copy is highlighted in orange and the states that make the two connected copies ground state are shown with green crosses. The fields and coupling of both Ising models is shown in the top right corner of their respective plot. Arrows and dotted lines ((a)-(d)) represent the potential energy penalty/benefit gained by connecting the respective state with the ground state of the single copy.

of the fields is negative. In this situation our first 2-qubit example had fields of $h_1 = -0.4$ and $h_2 = 0.7$. We set the coupling strength to be $J_{12} = -0.5$. As one of the fields is negative, according to figure 10.4, we expect the threshold value of J_{12} to also be negative. As the threshold value occurs when either $J_{12} < h_1$ or $J_{12} < h_2$, we expect the threshold value to be $J_{th} = -0.4$. Therefore, our first 2-qubit example was below threshold. We first plotted the energies of this example's four possible states (00, 01, 10, 11) in figure 10.7 (left). Our second 2-qubit example had the same value fields of $h_1 = -0.4$ and $h_2 = 0.7$, but we set the coupling strength to $J_{12} = -0.3$. Hence this example was above threshold. We plotted the energies of this example's four possible states in figure 10.7 (right).

We have put the diagram of the corresponding 2-qubit Ising model in the top left hand corner of both graphs. In both graphs the blue dots represent the energy of each of the possible states of the 2-qubit model, with the ground state of the singular 2-qubit model highlighted in orange. The two ground states which make up the total ground state of the two anti-ferromagnetically connected copies are marked by green crosses. The black arrows (if present) and dotted lines represent the energy shifts which would occur if that particular state was connected anti-ferromagnetically (with $J_F^{min} = 0.25$) to the ground state of the singular copy (11 left graph, 01 right graph).

In the left hand graph of figure 10.7, we see that 11 is the ground state of the

singular copy (orange dot), whereas the ground state of the two copy system is 11 00 (green crosses). If we look at all four of the possible states in figure 10.7, we can see that connecting 11 with any of them anti-ferromagnetically will cause one of three options: a positive energy shift, a negative energy shift or no energy shift at all. We can see that the states that receive no energy shift by connecting with 11 are 01 and 10. This is because in both these cases there is one satisfied J_F and one frustrated J_F link. This means that the energy shift from these two links cancel and we see no overall energy shift. On the other hand we see that connecting 11 with 11 receives a positive energy shift, whereas connecting 11 with 00 receives a negative energy shift. We see in this example that the positive energy shift causes 11 11 to move above 11 01, however the negative energy shift for 11 00 causes it to move below 11 01, making 11 00 the ground state of the two copy system.

In the right hand graph of figure 10.7, we see that 01 is the ground state of the singular copy (orange dot), whereas the ground state of the two copy system is 01 11 (green crosses). Connecting each of four of the possible energy states to the singular ground state 01, causes a positive energy shift, a negative energy shift or no energy shift. Connecting 01 with 11 or 00 receives no energy shift, as positive and negative couplings cancel. On the other hand, 01 01 receives a positive energy shift, whereas 01 10 receives a negative energy shift. In this example, the positive energy shift of 01 01 causes it to move above 01 11. However, the negative energy shift of 01 10 is not enough to move it below any states as the energy of 10 is too high. Therefore we see that the ground state of the two copy system is 01 11.

In this example, when J_{12} is below threshold (left graph), the ground state of the two copy system will be 11 00, whereas when J_{12} is above threshold (right hand graph), the ground state of the two copy system will be 01 11. Like for the previous examples, we hypothesise that the difference in fraction correct behaviour between Ising model instances may be due to differences in the two copy ground state caused by whether the coupling of an instance is above or below a threshold value.

The hypothesis reasoned in previous section 10.1.1 again applies here for an explanation of why this behaviour may be occurring. A diagrammatic explanation of this behaviour may be seen in figure 10.8. However, we note that we must be careful to treat the thresholds in the right manner, as in this case the rules have been somewhat reversed. For two positive (and with the same reasoning two negative) fields we expect the threshold in the coupling strength J_{12} to be *positive*. When the threshold is positive we expect to recover the correct ground state when we are *above* threshold. For the case we are considering here, (one positive and one negative field) we expect the threshold in J_{12} to be *negative*. When this is the case, we expect to recover the correct ground state when the coupling strength J_{12}

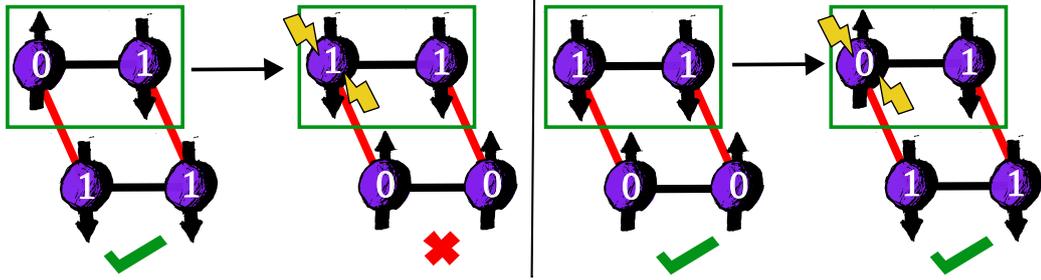


Figure 10.8: Same as figure 10.6 but with one positive and one negative field. Left hand diagram has J_{12} above the threshold, before a precision error has occurred (left) and after a precision error has caused the left qubit to flip (right). In this case the true precision single copy ground state is lost. Right hand diagram has J_{12} below the threshold, before a precision error has occurred (left) and after a precision error has caused the left qubit to flip (right). In this case the true precision single copy ground state is retained.

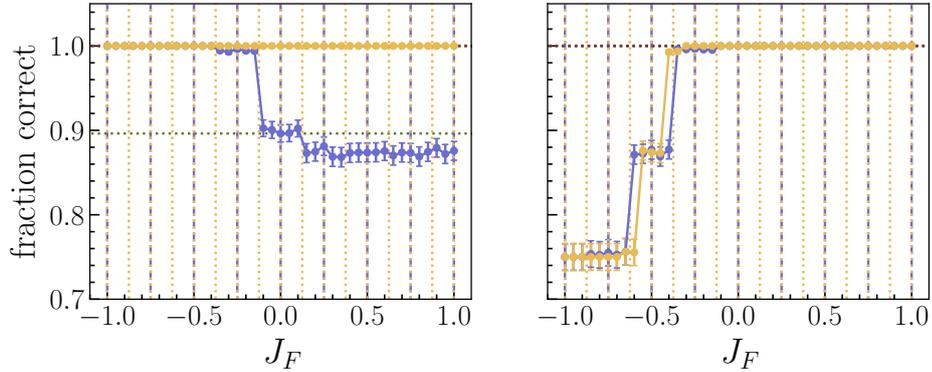


Figure 10.9: Left figure shows the fraction correct vs $J_F^{(p)}$ of 2 copies connected in a chain of a 2 qubit Ising spin chain example with fields $h_1 = 0.4$ and $h_2 = 0.7$ and coupling strength $J_{12} = 0.5$, for $p = 3$ (blue) and $p = 4$ (yellow). The right figure shows the same with $h_1 = 0.4$ and $h_2 = 0.7$ but the coupling strength is $J_{12} = 0.3$. Vertical dotted lines indicate the mid-points of each of the allowed divisions at that precision for $p = 3$ (blue) and $p = 4$ (yellow). Horizontal dotted lines indicate the fraction correct for $J_F^{(p)} = 0$ and the maximum fraction correct for any $J_F^{(p)}$, in green for $p = 3$ and red for $p = 4$.

is *below* threshold.

10.2 Fraction correct versus J_F for 2-qubit example with positive fields

10.2.1 Two copies

In order to establish the validity of the hypothesis (at least for 2-qubit Ising models) reasoned in section 10.1.1, we next measured the fraction correct vs $J_F^{(p)}$ for two 2-qubit Ising chain examples with positive fields: one with J_{12} above threshold and the other with J_{12} below threshold. First we looked at an instance where $h_1 = 0.4$ and $h_2 = 0.7$. Here we estimated the threshold to be at $J_{\text{th}} = 0.4$. In the first example we set the coupling above threshold at $J_{12} = 0.5$ and in the second example we set the coupling below threshold at $J_{12} = 0.3$. We did this measurement for two copies connected anti-ferromagnetically in a chain, and then three copies connected anti-ferromagnetically in a loop. Like the data in figure 7.5, as there was only one instance on which this measurement was carried out, we used the random error model, taking 10^3 error samples of the instance, to ensure a reliable value fraction correct was found.

Figure 10.9 shows the fraction correct vs $J_F^{(p)}$ for two copies connected in a chain, of a single 2-qubit Ising chain instance with fields $h_1 = 0.4$ and $h_2 = 0.7$, for precisions $p = 3$ (blue) and $p = 4$ (yellow), where the precision was applied using the random error model with 10^3 error samples. The left graph shows the fraction correct results from the 2-qubit Ising chain instance where the coupling strength is above the threshold value at $J_{12} = 0.5$, whereas the right hand graph shows the fraction correct results when the coupling strength is below the threshold value at $J_{12} = 0.3$. According to our hypothesis, as both of the fields are positive, we expect the threshold J_{th} to be positive. We estimate the threshold to be at $J_{\text{th}} = 0.4$. Therefore the Ising model with $J_{12} = 0.5$, i.e. the left hand graph, is above threshold and we expect that it should gain more benefit from being connected by antiferromagnetic links, than the Ising model with $J_{12} = 0.3$.

We see in figure 10.9 that, for the Ising model above threshold (left hand graph), at $p = 4$, the fraction correct is always ~ 1 . This indicates that neither the precision of the fields and couplings nor the strength of the links connecting the copies break the ground state of the Ising model. Whereas at $p = 3$, at $J_F = 0$ (two disconnected copies), the fraction correct is ~ 0.9 . This is further reduced when the two copies are connected with ferromagnetic links $J_F > 0$. However, when the two copies are connected using anti-ferromagnetic links, the fraction correct increases to ~ 1 , indicating a benefit to connecting copies anti-ferromagnetically here for this instance.

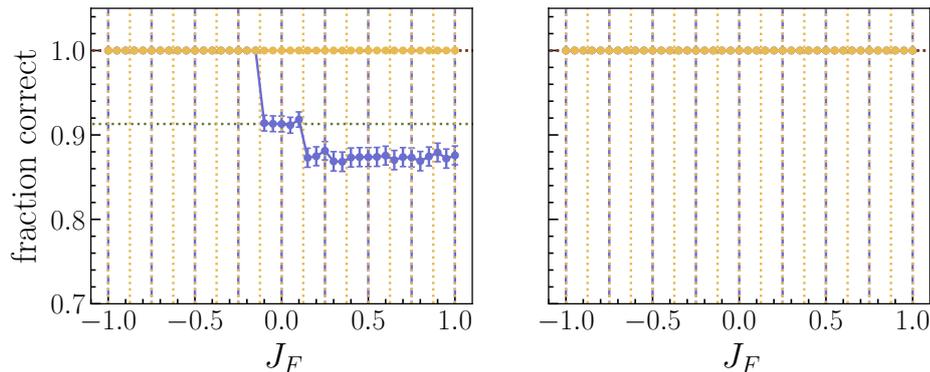


Figure 10.10: As figure 10.9, for 3 copies of a 2-qubit chain

On the other hand, for the Ising model below threshold (right hand graph), for both $p = 3$ and $p = 4$, we see a distinct behaviour from that of the model above threshold. For $J_F \geq 0$, the fraction correct is ~ 1 , indicating that neither the precision nor the strength of the links connecting the copies break the ground state of the Ising model. However, when the two copies are connected with anti-ferromagnetic links the fraction correct begins to decrease, first to ~ 0.88 at $J_F \sim -0.5$ and then further to ~ 0.74 at $J_F \lesssim -0.5$. This indicates there is a negative effect on the fraction correct for this instance by connecting the two copies using anti-ferromagnetic links.

Therefore, we only see a benefit to connecting copies with anti-ferromagnetic links for the instance above threshold when $p = 3$. Conversely, we only see a negative effect to connecting the copies with anti-ferromagnetic links for the instance below threshold. These results are in agreement with the hypothesis that connecting copies with anti-ferromagnetic links is more beneficial for instances with J_{12} above threshold, than with J_{12} below threshold.

10.2.2 Three copies

To see whether the validity of our hypothesis extended to three connected copies, we next repeated the fraction correct vs $J_F^{(p)}$ measurement for three copies connected in a loop of the same 2-qubit Ising chain instances, with fields $h_1 = 0.4$ and $h_2 = 0.7$, and two strengths of coupling: $J_{12} = 0.5$ above (left hand graph) and $J_{12} = 0.3$ below (right hand graph) the threshold J_{th} . Precision in the range $[1, 10]$ was applied using the random error model with 10^3 error samples.

In figure 10.10 we see that for the Ising model above threshold (left hand graph), at $p = 4$ the fraction correct is always ~ 1 , indicating that neither the strength of the links connecting the copies nor the precision break the ground state of the Ising model. This is the same behaviour as the left hand graph in figure 10.9.

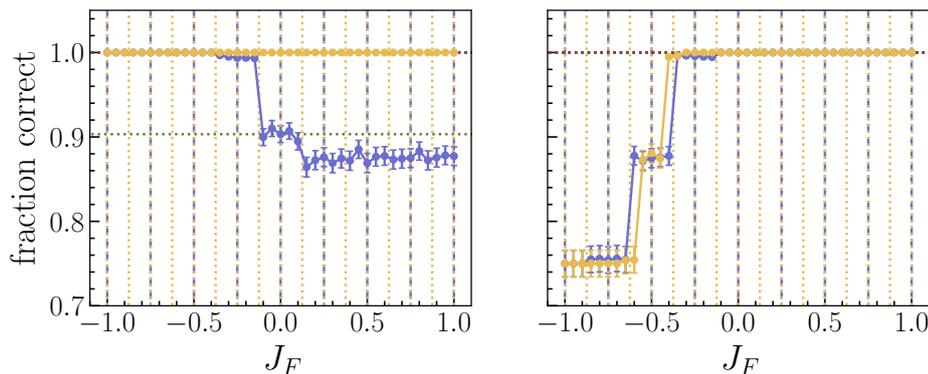


Figure 10.11: Left figure shows the fraction correct vs $J_F^{(p)}$ of 2 copies connected in a chain of a 2 qubit Ising spin chain example, with fields $h_1 = -0.4$ and $h_2 = 0.7$ and coupling strength $J_{12} = -0.5$, for precisions $p = 3$ (blue) and $p = 4$ (yellow). The right figure shows the same for fields $h_1 = -0.4$ and $h_2 = 0.7$ and coupling strength $J_{12} = -0.3$ for the same precisions p . Vertical dotted lines indicate the same as figure 10.9.

Similarly again to figure 10.9, we also see that for $p = 3$, at $J_F = 0$, the fraction correct is also ~ 0.9 , and that connecting the three copies with ferromagnetic links reduces the fraction correct. This is a slightly larger decrease than we saw for the two connected copies. We also see that we retain the increase of fraction correct to ~ 1 , when the three copies are connected anti-ferromagnetically $J_F < 0$.

On the other hand, for the Ising model below threshold (right hand graph), we again see a distinct effect compared to the instance above threshold, for both $p = 3$ and $p = 4$. However this behaviour is different to that of two copies. In this case, we see that the fraction correct is always ~ 1 , indicating that neither the precision nor the links connecting the copies are causing the ground state to break. These results also in agreement with our hypothesis, as we see the only benefit from connecting three copies anti-ferromagnetically, for the instance above threshold.

10.2.3 2-qubit example with one positive field, one negative field

We next repeated the measurement of fraction correct vs $J_F^{(p)}$, returning to two copies connected in a chain for a similar 2-qubit example instance, this time with one positive and one negative field: $h_1 = -0.4$ and $h_2 = 0.7$, with coupling strengths: $J_{12} = -0.5$ below (left hand graph) and $J_{12} = -0.3$ above (right hand graph) the threshold J_{th} . As explained in section 10.1, here we expect the threshold value to be $J_{\text{th}} = -0.4$. Precision $p = 3$ (blue), $p = 4$ (yellow) was again applied using the random error model with 10^3 samples.

In the left hand graph of figure 10.11, we see, for $p = 3$, a benefit in fraction correct for two copies connected with anti-ferromagnetic links and a negative effect for two copies connected with ferromagnetic links. In contrast to the results in the left hand graph of figure 10.9, this behaviour is seen when the coupling strength $J_{12} = -0.5$ is below (*rather than above*) threshold (and negative). In the right hand graph of figure 10.11, for both $p = 3$ and $p = 4$, we see a negative effect for copies connected using anti-ferromagnetic links. Conversely again to figure 10.9, this behaviour is seen when the coupling strength $J_{12} = -0.3$ is above threshold (*rather than below*).

These results agree with our prediction, that we expect the threshold value of the coupling strength J_{th} to be negative, when one field is negative and one field is positive. They also agree with the prediction described in section 10.1.2, that we expect connecting copies with anti-ferromagnetic links to provide more benefit when the coupling strength J_{12} is below threshold when one field is positive and the other is negative.

10.3 Larger size Ising models

So far in this section we have looked at the threshold behaviour of the 2-qubit Ising model. These models are useful for proof of principle, but trivial to solve and therefore not necessarily indicative of the behaviour of a larger sized model. We next take a quick look at the threshold behaviour of a 3-qubit Ising spin chain model, to see how it compares. We will then try and generalise the threshold behaviour of the model up to larger sizes.

To avoid confusion, we first look at the 3-qubit Ising model, limiting the strength of the fields and couplings to be in the range $[0, 1]$, i.e. positive only. This is equivalent to just looking at the top row of table 10.1, but the complexity of this table will increase as the number of qubits in the model increase. Also, we can see from the 2-qubit example, that the threshold behaviour of the model allowing for negative fields and couplings was similar (if somehow ‘reflected’) to the behaviour of the positive only case, so we could argue that if we derive the positive only case, the cases including negative fields and couplings should follow.

We see in table 10.2, the possible start state 111 (if the fields are positive and the couplings are positive and small) and the possible end states given one or two bit flips. Alongside each of these, we have written the inequalities that must be satisfied for each of these transitions to take place. Each of these inequalities has been derived the same way as for 2-qubits, by setting the energy of the end E_{end}

state as lower than the start state E_{start} . For example, for the state to change from 111 to 011, then $E_{011} < E_{111}$, this gives,

$$+h_1 - h_2 - h_3 - J_{12} + J_{23} < -h_1 - h_2 - h_3 + J_{12} + J_{23}. \quad (10.4)$$

We see the terms h_2 , h_3 and J_{23} cancel as the fields on the two other qubits are not affected by the bit flip on the first, and the coupling not connected to the first qubit also is unaffected. This leaves the inequality $J_{12} > h_1$ which you can see in the first row of table 10.2.

We can see in table 10.2, that the inequalities of the first two rows are similar to the first two rows of table 10.1, however the inequality in the third row is different, featuring two coupling strengths. The reason the inequalities of the first two rows are similar to that of the 2-qubit Ising model, is because of their configuration. Like the qubits in the 2-qubit model the qubits where the bit flip occurs in row one and two, (h_1 and h_3) they each only have one coupling link connected to them. On the other hand, for the qubit that undergoes the bit flip in the third row (h_2), there are two links connected to it. This causes the additional coupling term in the inequality of the third row. So we see that the inequality will depend on the field that is acting on the qubit and the couplings that are linked to the qubit.

We also see in table 10.2, rather than just one bit flip, we can get two bit flips. For the three qubit model, these can be either neighbouring e.g. 001 and 100 or separated by one qubit e.g. 010. In the neighbouring cases, we see that we only have dependence from one coupling as the coupling internal to the neighbours cancelled. (We note that if there were four qubits or a link back to qubit one, there would still be two couplings (they just have to be external to the neighbours). On the other hand, for the separate case, we get dependence of two couplings and two fields. We note that if there were four qubits or the extra one to three link, there would be two couplings per field. This result here is similar to just combining row one and row two (although we note that J_{12} and J_{23} can now work together to overcome h_1 and h_2), i.e. just having two single bit flips. We also note that in the spin chain setting the maximum number of links per qubit is two, meaning we would never get more than two couplings per qubit. Whereas for the spin glass setting, we would get $n - 1$ couplings per qubit (where n is the problem size).

Having looked at the transition inequalities of 2 and 3-qubit Ising chain models, we now try to generalise the transition inequalities to larger problem sizes. We first note that if there is only one bit flip in a transition or the bit flips are spaced one or more apart, we can treat the inequalities (almost) like separate models. That means we only need to look at the field on the qubit and strength of the couplings connected to the qubit to get a good idea of whether the inequality would be

Start state	End state	Inequality
111	011	$J_{12} > h_1$
	110	$J_{23} > h_3$
	010	$J_{12} + J_{23} > h_2$
111	001	$J_{23} > h_1 + h_2$
	100	$J_{12} > h_2 + h_3$
	010	$J_{12} + J_{23} > h_1 + h_3$

Table 10.2: Table showing for the 3-qubit Ising spin chain model, the starting state and possible end state, and the inequality (threshold) that must be satisfied for the transition to take place. In this table we set the fields and couplings to only be positive.

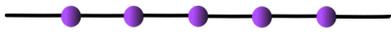
State	Transition	
One x	$h_i \leq J_{(i-1)i} + J_{i(i+1)}$	
> One x	$\sum_i^X h_i \leq \sum_i^X J_{(i-1)i} + J_{i(i+1)}$	
Two n.b.s	$h_i + h_j \leq J_{(i-1)i} + J_{j(j+1)}$	
Three n.b.s	$h_i + h_j + h_k \leq J_{(i-1)i} + J_{k(k+1)}$	

Table 10.3: Table showing for the n-qubit Ising spin chain model, the number and neighbourly-ness (n.b.) of bit flips (x), the transition inequality (threshold) that must be satisfied for the transition to take place and a diagram of the spin chain (qubits in purple, links in black) indicating the location of the bit-flip/s and affected links (in green). In this table we are in the setting where the fields and couplings are only positive.

satisfied. This picture becomes more complicated when neighbouring bits flip as only the couplings external to the neighbours are involved in the inequality. Table 10.3 gives an idea of how this might work in length n Ising spin chains. The picture is yet more complicated if Ising spin glasses are used.

10.4 The link selection protocol

We hypothesise that we might expect the threshold in coupling strength that occurs for 2-qubit models, may continue to occur for larger size Ising models. The reasoning behind this is as follows. If we had an Ising model with of size n-qubits, if each of the fields h_i applied to the qubits was much stronger than the couplings

J_{ij} , if the model had a Hamiltonian of the following form,

$$H = \sum_{i=0}^{n-1} h_i \hat{Z}_i + \sum_{i \neq j=0}^{n-1} J_{ij} \hat{Z}_i \hat{Z}_j, \quad (10.5)$$

and we forced the fields to be in the range $[0, 1]$ (i.e. positive), then we would expect the ground state of the model to be the all 1's state. As we increase the strength of the couplings connecting the qubits, at some point we reach a threshold above which at least one of the qubits will flip from 1 to 0 and the ground state of the model will change. For the two qubit model, when the strength of the coupling J_{12} was above this threshold (for two positive fields or two negative fields), we were more likely to recover the correct ground state when this state was subjected to an error caused by a lack of precision. We therefore hypothesise, that for a larger n-qubit Ising model, it becomes more likely that we recover the correct ground state (once subjected to precision) if original ground state contained more 0's.

If we make a large assumption and treat the n-qubit Ising model as multiple 2-qubit Ising models connected together, we can predict whether each 2-qubit model is above or below threshold, and therefore make a prediction on whether connecting each 2-qubit model to its copy could be beneficial. Herein this assumption and prediction lies the basis for the link selection protocol.

We note that once we allow negative fields to be applied to the qubits in an Ising model, the possibility of negative thresholds occurs and the accompanying mirrored threshold situation. However this rule can be consistently applied within a link selection protocol.

We also note that the validity of splitting the n-qubit Ising model into multiple 2-qubit models is most legitimate for models with fewer connecting links, such as a spin chain, rather than those with more connecting links such as spin glasses. So we might expect that this link selection protocol would work more effectively for spin chains than spin glasses.

The link selection protocol is as follows:

10.4.1 Link selection protocol

1. For each qubit i in Ising model (h_i is field on qubit i):
2. Calculate sum of the coupling strengths J_{ij} which are connected to qubit $i \rightarrow$ call this sum K .

3. Determine the sign/s compared to qubit i of the fields h_j on the neighbouring (connected) qubits. Same, opposite or both and go to the respective step a), b) or c).
 - a) If the sign of all the neighbouring qubits are the same as qubit i , use step 4 only.
 - b) If the sign of all the neighbouring qubits are opposite to qubit i , use step 5 only.
 - c) If the signs of the neighbouring qubits are some the same, some the opposite, use both step 4 and 5.
4. If $K > 0$, if $|K| > |h_i|$ connect qubit i to its copy qubit with minimum strength anti-ferromagnetic link. Else leave qubit i disconnected.
5. If $K < 0$, if $|K| > |h_i|$ connect qubit i to its copy qubit with minimum strength anti-ferromagnetic link. Else leave qubit i disconnected.

10.5 Error suppression results with the protocol

In order to measure the effectiveness of our link selection protocol described above, we next used it to prescribe whether or not to connect two and three copies of 10^4 instances of 5-qubit SK spin glasses and spin chains, using minimum strength anti-ferromagnetic links. We then subjected these instances to a lack of precision and measured their fraction correct. We compared the results of the copies connected using the protocol, to the results of the copies connected without using the protocol, and to the results of an equivalent number of repeats of single disconnected copies. We note that all of the results shown here connect copies using minimum strength anti-ferromagnetic links. The 5-qubit SK spin glass connected using the protocol measurements were repeated using $J_F = -1$, however no improvement above that seen from using minimum strength anti-ferromagnetic links was gained. The precision was applied using the deterministic random error model. The SK spin glass and spin chain instances were the same as those used in the previous chapters. The ground state calculation, used in our calculation of fraction correct, was again done using the branch and bound code written by Adam Callison.

10.5.1 Spin glasses

The top graph in figure 10.12 shows the fraction correct vs precision, for two and three copies (in various configurations), of 10^4 instances of 5-qubit SK Ising spin

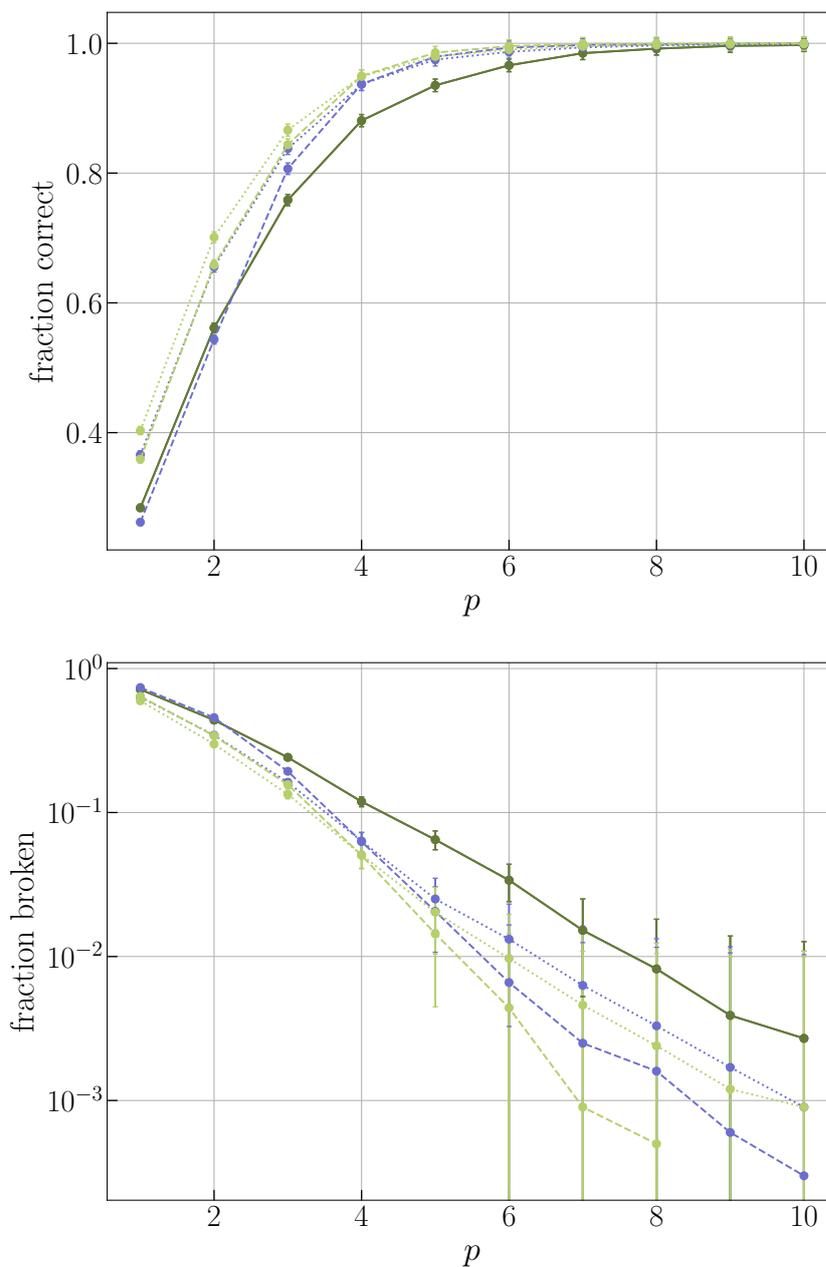


Figure 10.12: Top figure shows the fraction correct vs precision of 10^4 instances of 5-qubit spin glasses which have been subjected to precision by the deterministic random error model and have been configured in the following ways: three repeats of single disconnected copies (dark green, solid), two copies connected in a chain not using the protocol (blue dashed, dotted), three copies connected in a loop not using the protocol (light green, dotted), two copies connected in a chain using the protocol (blue, dashed) and three copies connected in a loop using the protocol. Bottom figure shows the log y plot of $1 - \text{fraction correct} = \text{fraction broken}$ vs precision for the same data.

glasses. The bottom graph shows the same data as in the top graph, presented differently. Instead of fraction correct, we measured the fraction broken ($1 - \text{fraction correct}$) and plotted this on a log y plot, vs precision. Therefore we can see the results more clearly when the fraction broken gets close to zero (same as fraction correct getting closer to 1).

The blue dotted lines on the top (bottom) graphs indicate the fraction correct (fraction broken) results from two copies connected in a chain where the protocol was used in order to determine whether or not to connect each qubit in the Ising model to its copy qubit. The light green dotted lines on the top (bottom) graph show the fraction correct (fraction broken) for three copies connected in a loop, where whether to connect or leave disconnected each qubit was also determined by the protocol. The blue dashed lines on the top (bottom) graphs indicate the fraction correct (fraction broken) results for two copies connected in a chain where every qubit in the Ising model was connected to its copy qubit (i.e. the protocol was not used). The light green dashed lines on the top (bottom) show the fraction correct (fraction broken) for three copies connected in a loop also not using the protocol. Finally, the dark green solid lines top (bottom) indicate the fraction correct (fraction broken) results for three disconnected copies (the same as three repeats of a single copy). Because we used the deterministic random error model, these results are identical to the results from two and one disconnected copies.

In the top graph of figure 10.12, for both two and three copies at $p < 4$, copies connected using the protocol have a higher fraction correct than copies connected not using the protocol and three disconnected copies. Still looking at the left hand graph, for both two and three copies, at $p \geq 4$ we see little or no difference between copies connected using the protocol and copies connected not using the protocol. We continue to see some improvement for both two and three copies connected using the protocol compared to three disconnected copies until $p \sim 8$.

In the bottom graph of figure 10.12, we see similarly for $p < 4$, two and three copies connected using the protocol outperform the same number of copies connected without the protocol and three disconnected copies. However we see in this graph for $p > 4$ the performance of the two and three copies connected without the protocol outperforms the copies connected using the protocol.

Comparing the results from two and three copies, we see that three copies connected with the protocol outperforms two copies connected with the protocol for all p in $[1, 10]$. We also see that three copies connected without the protocol outperforms two copies connected without the protocol for p in range $[1, 10]$. However we see that for $p > 5$ two copies connected without the protocol outperform three copies connected using the protocol.

10.5.2 Spin chains

We noted in section 10.4 that we expected the protocol to perform better for spin chains, due to their fewer couplings making the approximation to multiple two qubit Ising models (made whilst developing the protocol) more valid. We therefore now test the performance of the protocol on two and three copies of spin chains.

In the same way as the spin glasses, the top graph of figure 10.13 shows the fraction correct vs precision, for 10^4 instances of 5-qubit spin chains. Further, the bottom graph of the figure shows the fraction broken vs precision on a log y plot of the same data as the top plot.

In the top graph of figure 10.13, we see that two and three copies connected using the protocol outperform copies connected without using the protocol for roughly $p \leq 5$ (note that this is an improvement compared to the spin glasses). However, in the bottom graph of figure 10.13, we see that for $p > 5$ the two and three copies connected without using the protocol outperform those connected using the protocol.

Comparing the results from two and three copies, we see that the three copies connected using the protocol outperform the two copies connected in the same way for all p in $[1, 10]$. We also see that three copies connected without the protocol outperforms two copies connected without the protocol. Like for the spin glasses, we also see that for $p > 5$, the two copies connected without using the protocol outperform three copies connected using the protocol.

In conclusion, we see only a limited benefit from using the link selection protocol to connect copies and then only for $p < 4$ for spin glasses, and $p < 5$ for spin chains. This indicates we do not yet have a full understanding of the threshold behaviour in larger Ising models and more research is needed to understand whether a more effective protocol could be developed. We note that using a higher strength of our anti-ferromagnetic links when applying our protocol, also does not result in a greater improvement in overall fraction correct. This indicates that the protocol does not benefit every instance but may improve the benefits of some.

10.6 Chapter summary

In this chapter we first analysed 2-qubit spin chain Ising models, finding a threshold between the states that depends on the relative strengths of its fields and coupling. The location of the threshold can be explained in the following way. If both fields are positive or negative, there is a positive threshold, whereas if one field is positive

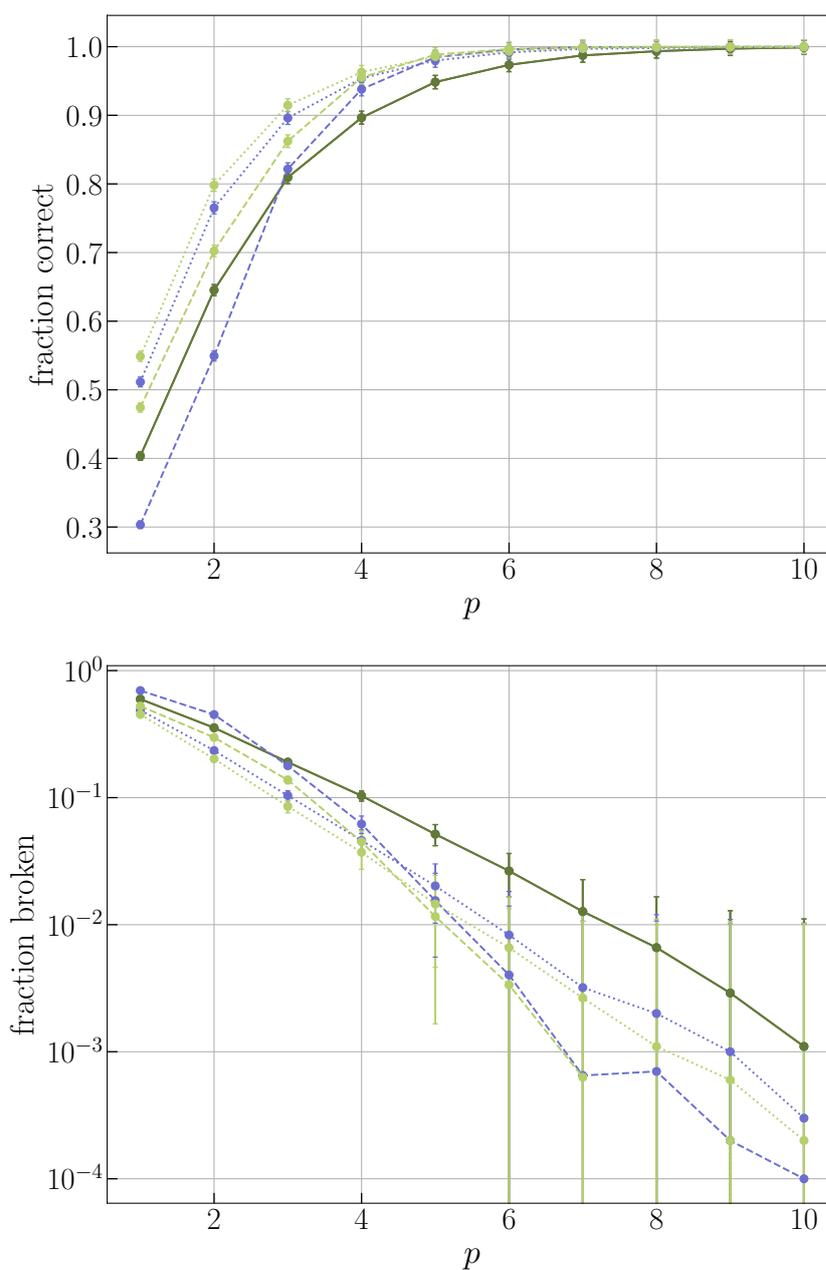


Figure 10.13: Top figure shows the fraction correct vs precision of 10^4 instances of 5-qubit spin chains which have been subjected to precision by the deterministic random error model and have been configured in the following ways: three repeats of single disconnected copies (dark green, solid), two copies connected in chain not using the protocol (blue, dashed), three copies connect in a loop not using the protocol (light green, dashed), two copies connected in a chain using the protocol (blue, dotted) and three copies connected in a loop using the protocol (light green, dotted). Bottom figure shows the shows log y plot of $1 - \text{fraction correct} = \text{fraction broken}$ vs precision for the data.

and the other negative, there is a negative threshold. In either case, the threshold occurs when the magnitude (positive or negative) of the coupling becomes larger than the magnitude of the smallest field.

We next hypothesised that the composition of the ground state of the 2-qubit model and whether it was above or below this threshold might have an influence on whether our error suppression method was effective or not.

We then analysed the energy levels of the 2-qubit model above and below threshold, and the possible energy shifts available when connecting each of the states with the original ground state. This showed which overall ground state (made up of two single copy ground states) would be chosen depending on whether the model was above or below the predicted threshold. We hypothesised that the composition of the single copy ground states (which are contained in the two copy system), influence whether or not our error suppression scheme is effective.

We next looked at the fraction correct vs J_F for four examples of 2-qubit Ising models, with two copies (in a chain) and three copies (in a loop). In the first two examples, for both two and three copies, we set the fields to be both positive. For the first example we set the coupling strength J_{12} to be above the expected threshold value J_{th} . In the second example we set J_{12} to be below expected J_{th} . For both two and three copies, in these first two examples we found that the example with the coupling above the expected J_{th} , performed better when the copies were connected with anti-ferromagnetic links than the example with the coupling strength below the threshold.

In the next two examples, for both two and three copies, we set one field to be positive and the other field to be negative. For the first of these two examples we set the coupling strength J_{12} to be below the expected threshold value J_{th} . In the second example, we set J_{12} to be above the expected J_{th} . For both two and three copies, we found that the example with the coupling strength below the expected J_{th} , performed better when the copies were connected with anti-ferromagnetic links, than the example with J_{12} above the threshold.

The above results agreed with our hypothesis that the effectiveness of our scheme depends on the ground state of the model. Using our knowledge gained from our fraction correct vs J_F measurements, we could now predict which 2-qubit models would benefit from being connected by anti-ferromagnetic links and which would prefer to be left disconnected. We then extended our hypotheses to 3-qubit Ising chain models, and tried to generalise to an n -qubit Ising chain.

By following this knowledge about 2-qubit models and approximating larger Ising models as multiple 2-qubit models, we developed a protocol which allowed us to

prescribe whether or not a particular qubit within an n -qubit Ising model should be connected to its corresponding copy qubit/s.

We next tested the effectiveness of this protocol on two and three copies of 5-qubit SK spin glasses and Ising spin chains, by measuring the fraction correct vs precision. We found for spin glasses that connecting two or three copies (in a chain or loop respectively) increased the fraction correct compared to those connected without the protocol for $p < 4$. However for $p > 4$ we found that the copies connected using the protocol performed worse than those connected without using the protocol. For spin chains, the precisions for which copies connected using the protocol outperformed those connected without, increased to $p < 5$. However beyond this, once again copies connected without using the protocol outperformed those connected using it.

Therefore, in this chapter we have found evidence to suggest that the ground state of the Ising model affects whether connecting it to its copies using anti-ferromagnetic links is beneficial or not. However the development of an initial protocol designed to suggest whether or not to connect certain qubits to their copies has lead to only very limited improvements in fraction correct at low precision for 5-qubit Ising models. This indicates that for larger Ising models the mechanism which determines whether connecting copies with anti-ferromagnetic links will be beneficial or not is more complicated. Therefore more research is required to establish whether an effective link selection protocol for larger models could be developed.

Quantum Walk Dynamics

To calculate the fraction correct in the previous chapters, we considered the problem Hamiltonian only and found the ground state using a classical branch and bound technique (see section 6.1.1). This provides good estimates for the success rate when computing using adiabatic processes as these remain in the ground state throughout the computation. As more dynamic continuous-time techniques (such as quantum walks and quantum annealing) populate excited states during their process, it is important for us to check that this does not introduce adverse effects that negate the precision improvement gained through using our scheme. In this section Adam Callison contributed the Python code which was used to simulate the quantum walks. The optimal γ values were calculated for [64] and can be found in a data archive [90].

As its time-independent Hamiltonian is easier to simulate, we opted to first test our scheme on quantum walks. As quantum walks tend to populate more excited states than the approximately adiabatic dynamics found in quantum annealing, it can also be seen as a harsher test of whether the populated states seen in these techniques negate the improvement gained from our scheme.

In section 11.1 we analyse whether the optimal hopping rates γ (a key quantity for quantum walks) which were determined for singular disconnected SK spin glasses in [64], were strongly affected by our error suppression scheme. In section 11.2 we measure the average success probability vs precision for three copies of 5-qubit SK spin glasses connected using our error suppression scheme and compared this to the average success probability of three repeats of a single disconnected copy. In section 11.3, we assess and compare the differences in success probability between instances where the average success probability for three copies connected using our scheme was *higher* and the differences when the average success probability for three copies connected using our scheme was *lower*. In section 11.4, we measure

the effect on average success probability of anti-ferromagnetically connected copies when connected according to the link selection protocol (described in section 10.4.1). In section 11.5, we establish the behaviour of normalised average success probability as we increase the problem size n of the SK spin glasses from 5 to 9 qubits. Finally in section 11.6 and 11.7, we summarize the chapter and describe potential future research.

11.1 Optimal γ

As explained in 2.1, a key parameter for quantum walks is the hopping rate γ , that determines the relative strengths of the driving and problem Hamiltonians. For practical applications, the success probability must not be very sensitive to the exact value of γ , because there is in general no efficient way to determine the optimal value without solving the problem itself. For SK spin glasses, [64] discusses in detail how to estimate a suitable heuristic value of γ . In order to check whether our error suppression scheme had an effect on the optimal gamma of the problem Hamiltonian, we measured the average success probability (as defined in equation 2.7) for $t = 30$ and $\Delta t = 70$, versus γ , for several comparative examples.

Figure 11.1 shows the average success probability $\bar{P}(t, \Delta t)$ for $t = 30$ and $\Delta t = 70$ versus the hopping rate γ of four 5-qubit SK spin glass instances, labelled $a) - d)$. Each of the four instances were chosen so that, at $p = 3$, (precision applied using the deterministic random error model) when using the classical branch and bound technique, their ground state was found to be incorrect when three disconnected copies were used, but they had at least one copy with a correct ground state, when three copies connected using our scheme were used. As was found by Callison et al in [64], the SK spin glasses have a fairly broad range of suitable γ values. A red dotted vertical line indicates the optimal value of γ for each instance, as determined in [64].

The green curves on figure 11.1 show the average success probability of three repeats of a single disconnected 5-qubit SK spin glass, with true precision. For this case, the probability of obtaining a correct solution from three repeats can be calculated as follows, where $P(y)$ is the probability of succeeding at least once in y runs. We have $P(3) = 1 - (1 - P(1))^3$, i.e., subtract the probability of failing all three runs from one. This is the correct quantity to compare with three connected copies, as it uses the same number of qubits in total, and can be run in the same time (by just setting the anti-ferromagnetic links to zero).

The orange curves show the average success probability of the same disconnected

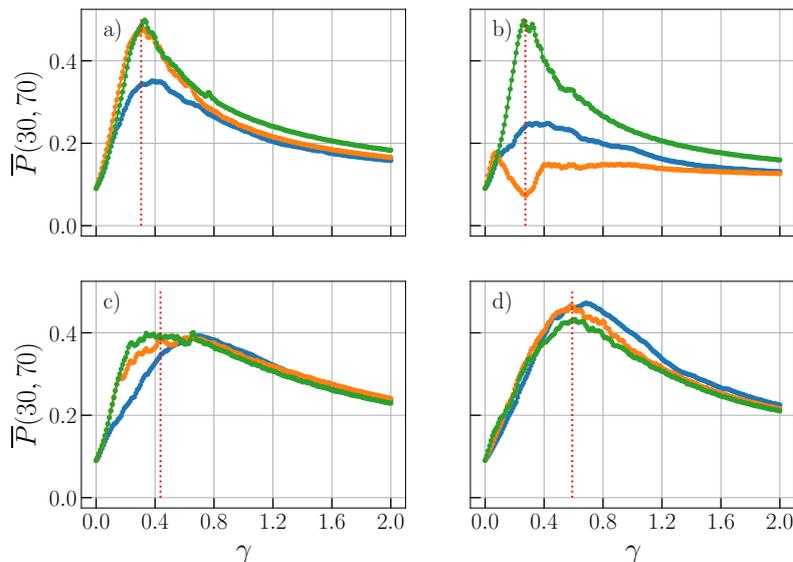


Figure 11.1: Average success probability, $P(t, \Delta t)$ for $t = 30$, $\Delta t = 70$ versus γ , for four instances a) to d) of a 5-qubit spin glass for: no links ($J_F = 0$) and no precision reduction (green); with $J_F = 0$ and precision set to $p = 3$ (orange); with $J_F \simeq J_F^{\min}$ and $p = 3$ (blue). The precision reduction used the deterministic random error model. Each instance a) to d) has an incorrect ground state for a single copy but is correct for at least one of three connected copies for $p = 3$. The optimal value of γ for the case with no links and no precision reduction, (peak of the green line) is shown as a vertical red dotted line. For reproducibility, the instances a) to d) were taken from data archive: [90], and have the unique I.D.'s 'acyenjvndejjyh-bcfmkjefgzjtjqjt', 'aakxejqunlcpqhmnftnrckailrczyp', 'aclwrzmpznazfjkjktzcswfdxj-prfth', 'acjdimxkqejkngndlykgxntdtxrgij', respectively.

three copies of 5-qubit SK spin glass instances, subjected to a precision of $p = 3$. In three of the four cases, the lack of precision causes a subtle difference in the curve, however the location of γ_{opt} is largely unaffected. However in case b), the lack of precision causes a large change to the curve, with a large decrease in the success probability at γ_{opt} .

The blue curves show the average success probability of the same 5-qubit SK spin glass instances, subjected to precision $p = 3$, however, this time the three copies were connected with minimum strength anti-ferromagnetic links in a loop.

There are a few subtleties to correctly comparing the three-connected-copy performance with three-repeats-of-single-copy performance in a quantum walk setting. In 2.1, we defined the average success probability as the overlap of the quantum state with the solution state, time averaged (2.7). Usually the solution state corresponds to the ground state of the problem Hamiltonian and indeed this is the case

for just a single copy. However, for three copies, due to the way we have defined ‘success’ (i.e. only one or more copies need to be correct), there are many different states that are considered a solution and should therefore add to the success probability. Not all of them are ground states. We define the success probability to be the sum of the probabilities of obtaining any of these solution states (since what we are interested in is correct solutions, not ground states) and then we time averaged in the same way.

In two of the four examples (*c*) and *d*)), the curves of the three copies connected using our scheme remain relatively similar compared to those of both the true precision and $p = 3$ curves. However in example *a*) the three connected copies result in a reduction of success probability at γ_{opt} , whereas in example *b*) it results in an increase. These four instance examples plotted in figure 11.1 therefore illustrate the range of behaviours, when measuring the success probability for a 5-qubit SK spin glass using a quantum walk. In these examples, the probability of obtaining the correct solution when using our three connected copies scheme is sometimes higher, sometimes lower and sometimes similar that of a single copy, but in all cases, the broad peak remains, and the optimal value of γ would provide good performance.

This shows that the dynamics and heuristic parameter estimation are not significantly impacted by introduction of a lack of precision or by the anti-ferromagnetic links introduced by our scheme. Moreover, given that the success probability is in any case not sensitive to the exact value of γ , it is not necessary to use any error mitigation techniques to increase the precision for the driver Hamiltonian settings, when solving this type of problem.

11.2 Average success probability versus precision

In order to determine whether using our scheme still provided a benefit when computing with quantum walks, we found the average long time success probability (2.7), for precision in the range $[1, 10]$, over 10^3 instances of three repeats of single disconnected copies and three anti-ferromagnetically connected copies of 5-qubit SK spin glasses, using quantum walks.

Figure 11.2 shows the long time average success probabilities, $P(t, \Delta t)$ for $t = 30$, $\Delta t = 70$ of three repeats of single disconnected copies (left bars, orange) and three anti-ferromagnetically connected copies (right bars, composite blue and purple) of 10^3 instances of 5-qubit SK spin glasses [90], versus four values of precision, $p = 3, 5, 7, 10$. The lower fraction of the composite right hand bars (blue) indicate the fraction of instances where the average success probability is lower for the

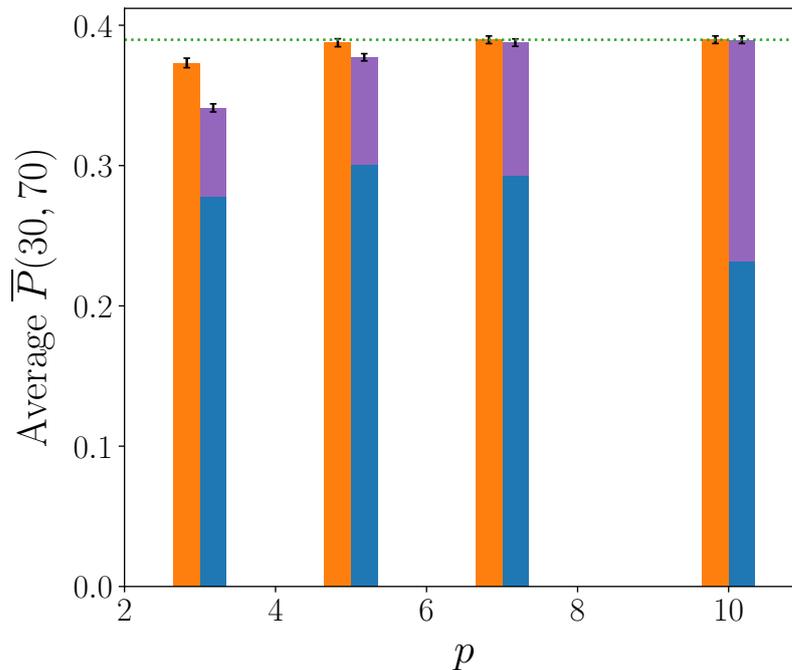


Figure 11.2: Average success probability (equation (2.7)), $\bar{P}(t, \Delta t)$ for $t = 30$, $\Delta t = 70$ at γ_{opt} , averaged over 1000 5-qubit spin glass instances versus precision. Left bars show the $\bar{P}(t, \Delta t)$ of three repeats of single disconnected copies subject to precision (orange). Right bars show the same for three anti-ferromagnetically connected copies subject to precision (composite blue and purple). $\bar{P}(t, \Delta t)$ for three disconnected copies with perfect precision is shown by a horizontal dotted green line, (with a constant average success probability of ~ 0.39) as a comparison for the configurations subjected to precision.

three anti-ferromagnetically connected copies than the three disconnected copies, whereas the upper fraction (purple) shows the reverse. The horizontal dotted green line shows the average success probability of the three disconnected copies with true precision (constant ~ 0.39). Note that this value (despite being less than one) is the ‘best’ we can do when computing this way. Therefore this line indicates the maximum success probability possible for the instances and serves as a benchmark for comparison.

Interestingly we see in figure 11.2 that despite the reduction in average success probability for the three repeats of single copies due to lack of precision, their average success probability is still always higher than that of the three anti-ferromagnetically connected copies. This is a completely opposite result compared to when we found the solution states using the classical branch and bound technique. This indicates that our error suppression scheme provides no benefit when finding the ground state using quantum walks.

Despite this, looking at the right hand composite bars of figure 11.2, we see that, at each precision there is a small fraction of instances where using our scheme does provide an improvement in fraction correct. We also see this fraction increases as precision increases, indicating that, at a higher value of precision, we may be able to recover an improvement from using our scheme. The lower average success probability found here does not necessarily nullify the effectiveness of our scheme when using quantum walks (as there are still a proportion of instances for which our scheme is beneficial), but indicates the interplay of excited states found in quantum walks is important in the calculation of average success probability.

11.3 Differences in average success probability

We next looked at the difference in average success probability between three repeats of single disconnected copies and three copies connected using our scheme. We split the difference into two cases. Case 1: when the average success probability was *higher* for the three copies connected using our scheme. Case 2: when the average success probability was *lower* for the three copies connected using our scheme.

Figure 11.3 shows the difference in average success probability $P(t, \Delta t)$ for $t = 30$, $\Delta t = 70$ at γ_{opt} on a log scale, between three repeats of single disconnected copies and three copies connected anti-ferromagnetically using our scheme, versus precisions $p = 3, 5, 7, 10$. The orange points show the differences when the average success probability is *higher* for the three anti-ferromagnetically connected copies, whereas the blue points show the differences when the average success probability is *lower* for the three anti-ferromagnetically connected copies.

We see in both cases that as the precision increases, the difference between the average success probability decreases. We also see that at all precisions (apart from $p = 3$) the difference, when the average success for three copies connected using our scheme was *higher* (orange), was larger than the difference, when the average success for three copies connected using our scheme was *lower* (blue). This indicates that when our scheme ‘helps’ an instance, it ‘helps’ these instances by on average more than it ‘hurts’ other instances by. However the fraction of instances it ‘helps’ is far lower than the instances ‘hurts’. This means on average there is no advantage (and some disadvantage) to using our scheme when computing with quantum walks.

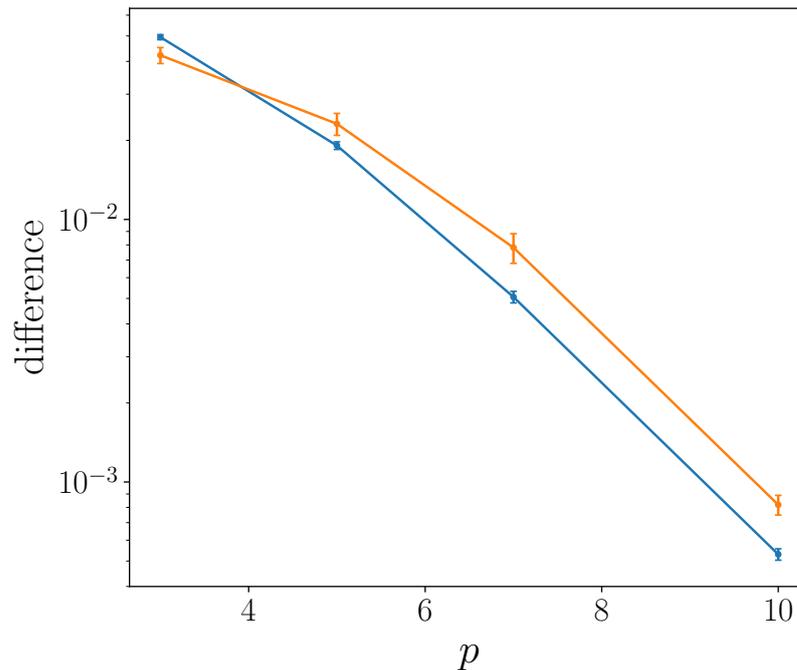


Figure 11.3: Difference in average success probability, $\bar{P}(t, \Delta t)$ for $t = 30$, $\Delta t = 70$ at γ_{opt} on a log scale, between three repeats of single disconnected copies and three anti-ferromagnetically connected copies, averaged over 10^3 5-qubit spin glass instances versus precision. The orange points (and line) show the difference in average success probability when the average success probability is *higher* for three anti-ferromagnetically connected copies than three disconnected copies and the blue points (and line) show when the reverse is true. Error bars are plotted for both cases.

11.4 Average success probability with the protocol

We next investigated the effect the average success probability $\bar{P}(t, \Delta t)$ of connecting the three copies with anti-ferromagnetic links using the link selection protocol developed in section 10.4.1. We measured this on 10^3 instances of 5-qubit spin glasses, for values of precision $p = 3, 5, 7, 10$. Precision was applied using the deterministic random error model.

Figure 11.4 shows the average success probability $\bar{P}(t, \Delta t)$ for $t = 30$, $\Delta t = 70$ of several configurations of three copies of 5-qubit spin glasses. Blue points (and line) show $\bar{P}(t, \Delta t)$ for three copies connected anti-ferromagnetically **without** using the link selection protocol. Red points (and line) show $\bar{P}(t, \Delta t)$ for three copies connected with anti-ferromagnetic links using the link selection protocol. Orange points (and line) show $\bar{P}(t, \Delta t)$ for three repeats of single disconnected copies. All the preceding configurations were subjected to precision $p = 3, 5, 7, 10$, applied

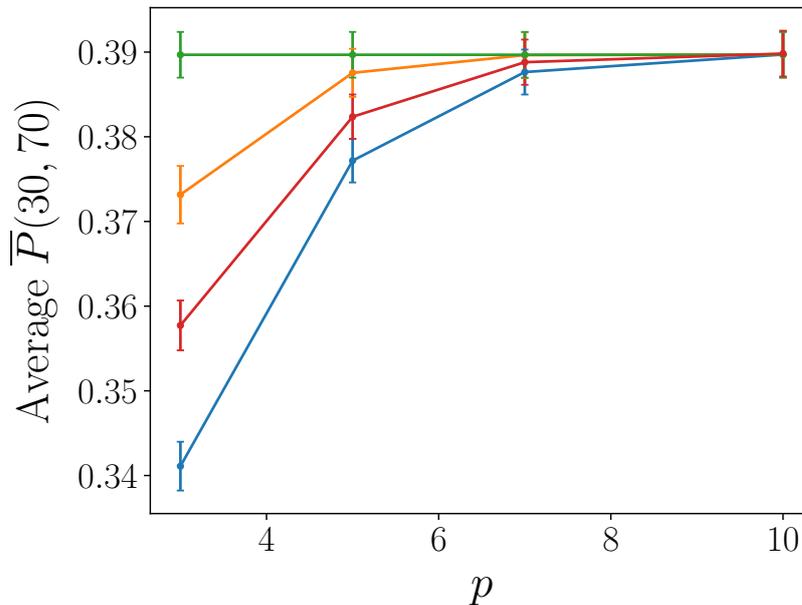


Figure 11.4: Average success probability (equation (2.7)), $\bar{P}(t, \Delta t)$ for $t = 30$, $\Delta t = 70$ at γ_{opt} , averaged over 1000 5-qubit spin glass instances versus precision p . Orange points (and line) shows the $\bar{P}(t, \Delta t)$ for three repeats of single disconnected copies, while subject to precision. Blue points (and line) show the $\bar{P}(t, \Delta t)$, for three copies connected with anti-ferromagnetic links, while subject to precision. Red points (and line) show the $\bar{P}(t, \Delta t)$ for three copies connected with anti-ferromagnetic links, selected using the link selection protocol 10.4.1, while subject to precision. For comparison, green points (and line) show the $\bar{P}(t, \Delta t)$, for three repeats of single copies with true precision.

using the deterministic random error model. As a benchmark of the maximum possible success rate, $\bar{P}(t, \Delta t)$ for three repeats of single disconnected copies with true precision is plotted (green points and line). Error bars are shown in all cases.

We note that the $\bar{P}(t, \Delta t)$ data for three copies connected without the protocol subjected to precision (blue), three repeats of single disconnected copies subjected to precision (orange) and three repeats of single disconnected copies with true precision (green) is the same as the data shown in figure 11.2 (without the fractioning of three connected copies). However, we felt it necessary to replot in order to provide context for the behaviour of $\bar{P}(t, \Delta t)$ for three copies connected using the link selection protocol.

In figure 11.4, we see that the copies connected using the protocol outperform the copies connected without using the protocol, at $p = 3$, $p = 5$ and $p = 7$ (although the error bars do overlap at $p = 7$). This contrasts to figures 10.13 and 10.12 in section 10.5, where we found that the copies connected using the protocol performed worse than the copies connected without for $p > 5$ for spin chains and for $p > 4$

for SK spin glasses. This indicates that using the protocol is more beneficial when computing using quantum walks.

In figure 11.4, we also see that the copies connected using the protocol never outperform three repeats of single disconnected copies for all values of p tested (although the error bars overlap at both $p = 7$ and $p = 10$). This also contrasts with figures 10.13 and 10.12 in section 10.5, where we saw for spin chains that, three copies connected with and without using the protocol and two copies connected with the protocol out-perform three repeats of single disconnected copies for all p and two copies connected without using the protocol out-perform three repeats of single copies for $p > 3$. For SK spin glasses we saw the same, apart from two copies connected without the protocol which outperforms three repeats of single copies for $p > 2$. This result indicates that whilst using the link selection protocol to connect copies when computing using quantum walks provides more of a benefit than it did for computing by classical techniques, we still see that simply using three repeats of single disconnected copies is still more effective for error suppression due to lack of precision than either.

These results indicate that quantum walks have an inbuilt robustness to lack of precision, that is more effective than our error suppression scheme. We hypothesise this may be the case for the following reasons. Our error suppression scheme has the net effect of allowing at least one other low lying state (most likely with one or two bit flips compared to the original correct ground state) to be included in the overall ground state. This means that if an error occurs on the original ground state changing it to an incorrect but low-lying state, then it is likely that the one of the copies it is connected to will be the original correct ground state. This works well when using classical or adiabatic techniques to compute because we have no other ways of accessing these low lying states. However, when computing with quantum walks, we will already have some non-zero probability of being in these low-lying excited states. If these low-lying states include the original correct ground state then they will also add to the success probability. Therefore when we are connecting copies using our error suppression scheme when computing using quantum walks, we are not really adding anything that will make us more likely to find the original correct ground state, but we are potentially adding a disadvantage by breaking copies due to the anti-ferromagnetic links.

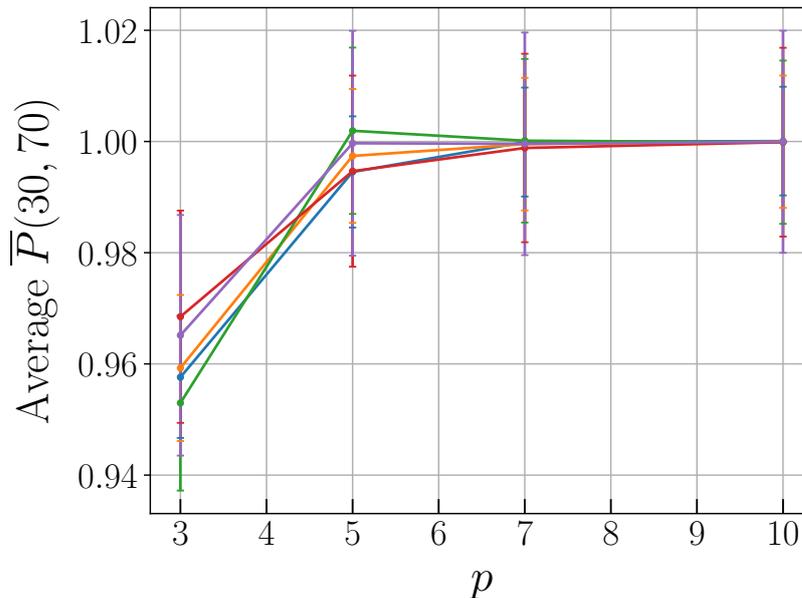


Figure 11.5: Normalised average success probability $\bar{P}(t, \Delta t)$, for $t = 30$, $\Delta t = 70$, versus precision p , for 10^3 instances of three repeats of single disconnected copies of 5 (blue), 6 (orange), 7 (green), 8 (red) and 9 (purple) qubit SK spin glasses. Normalisation was done by dividing each of the data points by the maximum average success probability of three repeats of single copies with true precision.

11.5 Average success probability for larger problem sizes

Next we wanted to see how the average success probability $\bar{P}(t, \Delta t)$ for three repeats of single disconnected copies varied, for problem sizes larger than 5-qubits. Figure 11.5 shows the normalised average success probability $\bar{P}(t, \Delta t)$ for $t = 30$ and $\Delta t = 70$ for 10^3 instances of three repeats of single disconnected copies of 5 (blue), 6 (orange), 7 (green) and 8 (purple) qubit SK spin glasses versus precision of $p = 3, 5, 7, 10$ which was applied using the deterministic random error model.

Normalisation was done by taking the dataset of one problem size and dividing the value of each of the data points by the maximum average success probability of three repeats of a single disconnected copies at true precision. The value of the maximum average success probability at each size were: 5 qubits: ~ 0.39 , 6 qubits: ~ 0.29 , 7 qubits: ~ 0.19 , 8 qubits: ~ 0.13 , 9 qubits: ~ 0.09 . As can be seen the maximum average success probability reduces with increasing problem size (although the rate at which it decreases by also reduces). This is as expected as with increasing problem size, there are more possible states for that the system to be in and therefore more opportunities to be incorrect.

In figure 11.5, we see that there is little difference in normalised $\bar{P}(t, \Delta t)$ versus precision p , between sizes, and all error bars overlap. All sizes start with a normalised $\bar{P}(t, \Delta t)$ of around 0.96 at $p = 3$ which then increases to around 1.00 at $p = 5$. As p increases further to $p = 7$ and $p = 10$ the normalised $\bar{P}(t, \Delta t)$ remains around 1.00, however the variance of the data around this value reduces. This means at $p = 3$, when computing with quantum walks, we are only reducing the average success probability due to a lack of precision by around 4%. This percentage also doesn't seem to vary with increasing problem size. This contrasts to what we found when computing using the classical branch and bound technique (see figure 9.4) to find the SK spin glass ground states. There we saw that when qubit number $n = 5$, the fraction correct (equivalent to average success probability here) was reduced by around 25% for three repeats of single disconnected copies. This grew to over 40% when $n = 9$. This finding further indicates that quantum walks have an intrinsic robustness to precision errors of this type, perhaps for the reasons hypothesised in the previous section.

11.6 Chapter summary

In this chapter, we assessed whether our error suppression scheme affected the optimal hopping rate γ of quantum walks. In order to do this, we measured the average success probability $\bar{P}(t, \Delta t)$ for $t = 30$ and $\Delta t = 70$ versus the hopping rate γ , of four individual 5-qubit SK instances, all subjected to precision of $p = 3$ (applied using the deterministic random error model) in the following configurations: three repeats of single disconnected copies and three copies connected anti-ferromagnetically. We compared this to the $\bar{P}(t, \Delta t)$ versus γ of three repeats of single disconnected copies of each instance at true precision and their optimal value of γ , which was calculated in [64] and stored at [90]. We found that there was no great difference in the optimal γ for copies subjected to lack of precision and connected using our scheme meaning that the optimal γ 's calculated in [64] and the heuristic used to calculate them, could be continued to be used when using our error suppression technique.

Next we looked at the average success probability $\bar{P}(t, \Delta t)$ vs precision p of 10^3 instances of 5-qubit SK spin glasses. We looked at the following configurations: three repeats of single disconnected copies and three copies connected with anti-ferromagnetic links according to our error suppression scheme. Both were subjected to precision using the deterministic random error model. For the three copies connected using our error suppression scheme, we split the total average success probability into the average success probability when the average success probabil-

ity of three repeats of single copies was *higher* and the average success probability when the average success probability of three repeats of single copies was *lower* than the average success probability of the three anti-ferromagnetically connected copies. As when computing with quantum walks the average success probability is always below one, we also measured the average success probability of three repeats of single disconnected copies not subjected to precision and used this as the benchmark for the maximum achievable average success probability. We found that for all values of precision tested ($p = 3, 5, 7, 10$), three repeats of single copies outperformed three copies connected using our error suppression scheme. This was a completely opposite result to when the ground states were computed using the classical branch and bound technique in section 8, indicating our error suppression scheme was largely ineffective when computing with quantum walks. However we did see that the fraction of average success probability of the three connected copies where the average success probability of the three repeats of single copies was *lower* did increase with increasing p . Therefore there may be a potential benefit from using our error suppression scheme when precision p is high enough.

We next looked at differences in success probability for the three copies connected using our error suppression scheme, for two cases. Case 1: when the average success probability was *higher* for three copies connected using our error suppression scheme, than three repeats of single disconnected copies. Case 2: when average success probability was *lower* for three copies connected using our error suppression scheme, than three repeats of single disconnected copies. We found in both cases the difference decreases with increasing precision p . We also found that other than for $p = 3$, the difference for case 1 was always larger than for case 2. This indicated that when our error suppression benefited a particular instance, the amount the particular instance benefited was on average larger than the amount the instances were disadvantaged by in case 2.

We next measured the effect on average success probability versus precision p of introducing our link selection protocol when connecting the three copies anti-ferromagnetically. We found that the copies connected with anti-ferromagnetic links selected by our link selection protocol performed better than the copies connected anti-ferromagnetically without using our protocol. However this improvement, was not enough to outperform the three repeats of single disconnected copies. This result further indicated the intrinsic robustness of quantum walks to the precision errors we had introduced and we hypothesised why this might be the case.

Finally, we measured the normalised average success probability versus precision p , for three repeats of disconnected single copies for sizes 5, 6, 7, 8, and 9-qubit SK spin glasses. We normalised the average success probability of each size, by

dividing each of the results by the average success probability at true precision for each size. We found that the average success probability for each size was only reduced by $\sim 4\%$ at $p = 3$ before tending toward ~ 1 at $p = 5, 7, 10$. This was a much smaller decrease in success probability than what was found when the solution state was computed using the classical branch and bound technique. Also the percentage by which the average success probability was reduced at $p = 3$ was fairly constant for all n tested here, whereas for the solution states computed using the classical technique, this reduced with increasing n . This further indicates the intrinsic robustness of quantum walks to these type of precision errors.

11.7 Future work

Characterizing the effects of limited precision and anti-ferromagnetically linked copies on the outcomes of quantum walk computation requires significant work to unravel the contributions from excited states that also provide correct solutions. The above work showing that quantum walks have an intrinsic robustness to precision errors which renders our error suppression scheme ineffective, has only been carried out for using the deterministic random error model. Therefore further analysis using the deterministic mid-point error model and random error model would be highly useful. Furthermore we saw in figure 11.2, that our error suppression scheme does provide a benefit for some fraction of instances (and this fraction grows with increasing precision). Therefore it would be interesting to conduct a further analysis into which instances are benefiting from our scheme and if these instances share any similarities that could be exploited.

Conclusion and Outlook

12.1 Conclusion

In this thesis I have introduced and provided numerical analysis for small n , of a scheme that suppresses errors in the ground state of Ising models due to lack of precision in its fields and couplings. This scheme was inspired by, but has several key differences to QAC an error correction scheme first introduced in [1]. After introducing the necessary theory in chapters 1, 2, 3, 4, and 5, in chapters 7, 8, 9, 10 and 11, I presented the numerical analysis and research of our error suppression scheme. I provide a summary of each of the numerical analysis and research chapters below.

Chapter 7: Link Strength Tests. In this chapter, we first looked how the fraction correct varied with the strength of the links J_F between copies, in order to determine if there was a benefit to connecting copies and if so which was the optimal value of J_F . We applied a lack of precision using first the random error model, then the deterministic mid-point error model and finally the deterministic random error model. For each model tested we simulated precisions of $p = 3$ (blue) and $p = 4$ (yellow). We measured both 5-qubit SK spin glasses and spin chains, in two copies in a chain and three copies in a loop configurations. We found that there was a benefit to connecting copies with close to minimum (allowed by precision) strength anti-ferromagnetic links, $J_F^{(p)} \approx -2^{-p+1}$, which we called J_F^{\min} . We also found that for three connected (in a loop) copies, fraction correct remained high, close to that of separate copies, as the magnitude of the anti-ferromagnetic links increased, whereas for two connected (in a chain) copies, the fraction correct rapidly decreased to well below that of separate copies. We hypothesised this was due to differences in the way the three copies suppressed errors compared to the two copies. We noted that the error suppression benefit we saw from connecting

copies with links $J_F^{(p)} = J_F^{\min}$, was an average effect over 10^3 instances of SK spin glasses and spin chains. We next looked at the fraction correct versus J_F for four randomly selected singular instances. We applied precision using the random error model with 10^3 error samples, so that fraction correct statistics could be built up in this way. We found that for two of the instances, connecting the copies with anti-ferromagnetic links provided a beneficial effect, for one of the instances no effect, and for the other instance a negative effect. This showed that our scheme was not helping every instance. We hypothesised that a non-computationally intensive technique that could identify whether our scheme would help or not would allow us to extract a stronger advantage. We next repeated the measurement of fraction versus J_F across 10^3 instances of 5-qubit spin glasses with no fields. We found that the benefit in fraction correct when $J_F^{(p)} \approx -2^{-p+1}$, was completely lost in this scenario for both two and three copies. We then hypothesised on the mechanism of our error suppression scheme.

Chapter 8: Error Suppression Results. In this chapter we measured the fraction correct vs precision p in the range $[1, 10]$, of multiple configurations of two, three, four and five copy 5-qubit SK spin glasses and two and three copy 5-qubit spin chains. All copies were connected with anti-ferromagnetic links of strength J_F^{\min} . We first looked at two copies connected in a chain of 5-qubit SK spin glasses. We applied precision first using the deterministic error model and then using the random error model. We found an improvement in fraction correct over two separate copies above $p = 3$ and $p = 4$ respectively. We then repeated this measurement for two connected copies of 5-qubit spin chains, using the same two error models to apply precision. In both of these cases we saw an improvement in fraction correct above $p = 3$. In all the following simulations in this chapter, only the deterministic error model was used. We next looked at the fraction correct versus precision for three copies of 5-qubit SK spin glasses connected anti-ferromagnetically in a chain. Here we found an improvement in fraction correct above three separate copies for values of precision greater than $p = 4$. We repeated this measurement for three copies of 5-qubit SK spin glasses connected in a loop. Here we saw an improvement in fraction correct across the whole range of precision $[1, 10]$. After this we measured the fraction correct versus precision for four copies in a loop of 5-qubit SK spin glasses. Here we only saw an improvement in fraction correct above four separate copies for precisions $p \geq 5$. Finally we measured the fraction correct versus precision of five copies in a loop of 5-qubit spin glasses. Here we recovered the improvement in fraction correct over five separate copies for the whole precision range $[1, 10]$. Due to gaining the best performance improvement using loops of three and five copies, we hypothesised that the frustration present in these configurations was key to maximising the performance of our error suppression scheme.

Chapter 9: Precision Improvements. In this chapter we aimed to find the precision improvement in terms of bits of precision gained by connecting three copies of SK spin glasses and spin chains of size 5 to 9 qubits, with anti-ferromagnetic links of strength J_F^{\min} in a loop. We first looked at the conditional fraction correct versus precision of three copies in a loop of 5 and 9-qubit SK spin glasses and 5-qubit spin chains. The conditional fraction correct was the fraction correct when: 1) the single copy was correct and the three copies were correct, 2) the single copy was incorrect and the three copies were correct and 3) the single copy was correct and the three copies were incorrect. Here we found that most of the instances fell into category 1), only very few fell into category 3) and a small but important quantity fell into category 2), indicating the source of our improvement. By running a single separate copy alongside our three connected copies we could also include category 3) into our measure of total fraction correct. We next measured the total fraction correct versus precision for sizes 5 to 9 qubits of three copies (connected with anti-ferromagnetic links in a loop) plus one disconnected copy of SK spin glasses and spin chains and compared this to the fraction correct of the four repeats of separate copies. We found that the improvement in fraction correct was retained across these sizes. We next quantified this improvement in fraction correct by measuring the horizontal distance between the same value of fraction correct for separate copies and three connected copies plus one disconnected copy of the same size. For SK spin glasses we found that the improvement in precision from three connected copies plus one disconnected copy was about 3 at $p = 7$ before showing signs of levelling off. For spin chains we found the improvement in precision from three connected copies plus one disconnected copy was about 4 at $p = 7$, before also showing signs of levelling off.

Chapter 10: Link Selection Protocol. In this chapter, based on an analysis of 2-qubit Ising models, we developed a link selection protocol which aimed to enhance the improvements from our error suppression scheme, by selecting which corresponding qubits should be connected and which qubits should be left disconnected, in a computationally non-intensive way. We first analysed the location of a threshold between a ferromagnetic and anti-ferromagnetic ground state in 2-qubit Ising models, which depends on the strength of link connecting the two qubits. We next analysed the energy levels of the 2-qubit model, above and below this threshold and predicted that our error suppression scheme should perform better on one side of this threshold. We next confirmed this prediction by measuring the fraction correct versus J_F for two and three copies of 2-qubit models. By making a large assumption and approximating an n -qubit Ising model as n 2-qubit Ising models, we developed a link selection protocol, that decided whether or not corresponding qubits in separate copies should be connected or left disconnected, in a

non-computationally intensive way. We tested the performance of this link selection protocol on two (in a chain) and three (in a loop) copies of 5-qubit SK spin glasses and spin chains. For SK spin glasses, we found two and three copies connected using the link selection protocol gained an improvement over two and three copies connected without the protocol for $p < 4$. For spin chains, we found two and three copies connected using the link selection protocol gained an improvement over two and three copies connected without the protocol for $p < 5$. This indicated our link selection protocol was successful at low precision but not at high. This indicated the need for more research in this area to gain a further improvement.

Chapter 11: Quantum Walk Dynamics. In this chapter, we looked at the effect of our error suppression scheme when computing using quantum walks. We first looked at the effect of our scheme on the location of optimal γ . Looking at four random individual 5-qubit SK spin glass instances, we found that in all four cases the value of optimal γ was not significantly impacted by our scheme. We also found that in two of the cases, the success probability at optimal γ was not strongly impacted by our scheme. In one case the success probability was reduced by our scheme and in the other case the success probability was improved by our scheme. In order to determine whether we still saw an improvement on average from using our scheme with quantum walks, we next measured the fraction correct versus precision for three copies connected anti-ferromagnetically in a loop at $p = 3, 5, 7, 10$. We found that there was no improvement from the three connected copies over the three separate copies at any of the precisions tested. Despite the lack of overall improvement, we saw that a small fraction of instances did gain an improvement from our scheme and this fraction increased with precision. We also saw that the difference in success probability between three separate copies and three connected copies was larger for instances where an improvement was gained with three connected copies for $p = 5, 7, 10$. Next, we tested the effect of using the link selection protocol to connect three copies when computing using quantum walks. We found that for all precisions tested, the three copies connected using the link selection protocol outperformed those connected without using the link selection protocol. Despite this improvement, the three copies connected using the link selection protocol still did not outperform three separate copies at any value of precision. We hypothesised that the lack of improvement may be due to access to excited states. This is gained from using our scheme in a classical or adiabatic setting but is already present when computing with quantum walks. We then measured the normalised success probability versus precision for sizes 5- to 9-qubit SK spin glasses. We found that compared to models computed classically, for quantum walks, the reduction in success probability due to lack of precision was much lower and there was also little to no scaling with problem size. This

indicated an innate robustness to lack of precision present when computing using quantum walks.

12.2 Outlook

In chapter 8, we saw that configurations of copies that contained frustration (loops of three and five copies) provided the strongest improvement in fraction correct compared to separate copies. An interesting avenue of research, would be to analyse the performance of several different frustration containing configurations in order to determine if there is an optimum configuration. Investigating the viability of concatenating several frustrated loops, as a way of improving performance further could also be interesting. Our work clearly indicates the importance of providing frustration in a system when designing a continuous-time quantum computer/annealer.

We have tested our scheme only on sizes of up to 9 qubits. If the improvement gained by our scheme extends to higher numbers of qubits, it could provide a way of scaling beyond precision limits in continuous-time quantum computing. In order to determine if our technique was helpful on a useful size of quantum computer, testing of larger size models would be required. On a classical computer the maximum number of qubits that can be simulated is around 50 [171], though for certain tasks more can be simulated [172, 173]. Another way of determining helpfulness would be to show an advantage analytically. The authors in [140, 142, 143] (described in section 5.3.1) use a mean field approach to show that QAC reduces the effective temperature. Potentially this technique could be applied to our scheme, thereby showing favourable scaling.

An improvement could also be potentially shown experimentally if a suitable model was possible to run on the hardware. On machines such as D-wave, minor embedding used to map some models to the machine, is expected to have a negative effect on our scheme, which could lead to inconclusive results. However, it could still be tested if a suitable direct embedding onto the native graph could be found. Experimental results from QAC [137, 139] indicate testing like this could be possible on current D-Wave systems.

In chapter 10, we saw that copies connected with link selection protocol we devised from approximating an n -qubit Ising model as multiple 2-qubit Ising model only outperformed copies connected without the protocol at low precision (for classically computed ground states). This indicated the unsuitability of this approximation on larger size Ising models. An interesting avenue of research would be to investigate

a more rigorous link selection protocol, to see if a balance could be struck between optimal performance and a non-computationally intensive protocol.

In chapter 11, we found that when computing with quantum walks, the improvement gained from connecting copies anti-ferromagnetically according to our scheme was completely lost. We hypothesised this was due to the access to excited states that is available when computing with quantum walks, indicating that they have an innate robustness to this type of precision error. This contrasts to AQC, where we expect performance like we found using the classical branch and bound technique, due to AQC not having access to excited states. It would therefore be an interesting avenue of research to look into the performance of our scheme in quantum annealing or hybrid approaches (interpolation between quantum walks and AQC first described in [63]) as a more mixed result might be expected.

Finally, another interesting avenue to explore, would be the application of our scheme in a gate-based setting, particularly given the similarities between QAOA and quantum annealing [174]. Whether our scheme could be extended beyond classical optimization problems and classical repetition codes to fully universal quantum error correcting codes, e.g., for quantum simulators, where quantum Hamiltonians are evolved continuously in time, is also an interesting open question. The weak anti-ferromagnetic links used in our scheme effectively rewards finding good quality (low energy) solutions which differ by large Hamming distances (i.e. maximising the number of bit-flips). This may have application beyond error mitigation in hybrid algorithms, or for problems where a diverse set of solutions are desired.

Bibliography

- [1] K. C. Young, R. Blume-Kohout and D. A. Lidar, “Adiabatic quantum optimization with the wrong Hamiltonian”, *Phys. Rev. A* **88**, 062314 (2013).
- [2] A. Pearson, A. Mishra, I. Hen and D. Lidar, “Analog errors in quantum annealing : doom and hope”, *npj Quantum Information*, 107 (2019).
- [3] T. Albash, V. Martin-Mayor and I. Hen, “Analog errors in ising machines”, *Quantum Science and Technology* **4**, 02LT03 (2019).
- [4] J. Bennett, A. Callison, T. O’Leary, M. West, N. Chancellor and V. Kendon, “Using copies to improve precision in continuous-time quantum computing”, *arXiv:2206.02545* (2022).
- [5] C. Carman and J. Evans, “On the epoch of the Antikythera mechanism and its eclipse predictor”, *Arch. Hist. Exact Sci.* **68**, 693 (2014).
- [6] P. A. Iversen, “The Calendar on the Antikythera Mechanism and the Corinthian Family of Calendars”, *Hesperia: The Journal of the American School of Classical Studies at Athens* **86**, 129 (2017).
- [7] K. Efstathiou and M. Efstathiou, “Celestial gearbox”, *Mechanical Engineering* **140**, 31 (2018).
- [8] T. Freeth, Y. Bitsakis, X. Moussas, J. H. Seiradakis, A. Tselikas, E. Magkou, M. Zafeiropoulou, R. Hadland, D. Bate, A. Ramsey, M. Allen, A. Crawley, P. Hockley, T. Malzbender, D. Gelb, W. Ambrisco and M. G. Edmunds, “Decoding the Antikythera Mechanism: Investigation of an Ancient Astronomical Calculator”, *Nature* **444**, 587 (2006).
- [9] U. of Southern California, *Islam, knowledge and science*, (2008) https://web.archive.org/web/20080119010548/http://www.usc.edu/dept/MSA/introduction/woi_knowledge.html (visited on 19/01/2008).

- [10] J. Al-Khalili, *Pathfinders, the golden age of arabic science* (Allen Lane, London, 2010), pp. 41, 198.
- [11] R. P. Lorch, “The Astronomical Instruments of Jabir ibn Aflah and the Torquetum”, *Centaurus* **20**, 11 (1976).
- [12] S. T. S. Al-Hassani, *How it works: mechanism of the castle clock*, (2008) <https://web.archive.org/web/20090810052611/http://muslimheritage.com/topics/default.cfm?ArticleID=901> (visited on 02/03/2023).
- [13] C. Razy, *Étude analytique des petits modèles de métiers exposés au musée des tissus* (Musée historique des tissus, Lyon, France, 1913).
- [14] E. A. Posselt, *The Jacquard machine analyzed and explained: with an appendix on the preparation of Jacquard cards* (Pennsylvania museum and school of industrial art, Philadelphia, 1887).
- [15] C. H. Museum, *The engines*, (2023) <https://www.computerhistory.org/babbage/engines/> (visited on 02/03/2023).
- [16] L. F. Menabrea and A. Lovelace, *Sketch of the Analytical Engine invented by Charles Babbage... with notes by the translator. Translated by Ada Lovelace* (Richard Taylor (ed.). Scientific Memoirs. Vol. 3, London, 1843), pp. 666–731.
- [17] F. da Cruz, *Herman Hollerith*, (2022) <http://www.columbia.edu/cu/computinghistory/hollerith.html> (visited on 02/03/2023).
- [18] W. Rodgers, *THINK: A Biography of the Watsons and IBM* (Stein and Day, New York, 1969), p. 83.
- [19] A. M. Turing, “On computable numbers, with an application to the entscheidungsproblem”, *Proceedings of the London Mathematical Society*, 2nd ser. **42** (1937).
- [20] J. S. Light, “When computers were women”, *Technology and Culture* **40**, 455 (1999).
- [21] B. Park, *Who were the codebreakers*, (2021) <https://bletchleypark.org.uk/our-story/who-were-the-codebreakers/> (visited on 02/03/2023).
- [22] H. H. Goldstine and A. Goldstine, “The Electronic Numerical Integrator and Computer (ENIAC)”, *Math. Comp.* **2**, 97 (1947).
- [23] T. R. Kennedy Jr, “Electronic computer flashes answers may speed engineering”, *New York Times* (1946).
- [24] D. Kahng, “Silicon-silicon dioxide field induced surface devices”, *S.M. Semiconductor Devices: Pioneering Papers.*, 583 (1961).

- [25] M. M. Atalla, “Stabilization of silicon surfaces by thermally grown oxides”, *Bell System Technical Journal* **38**, 749 (1959).
- [26] F. Faggin and M. E. Hoff, “Standard parts and custom design merge in four-chip processor kit”, *Electronics*, 112 (1972).
- [27] F. Faggin, M. E. Hoff, H. Feeny and S. Mazor, “The MCS4 - An LSI micro-computer system”, *IEEE Region Six Conference*, 8 (1972).
- [28] G. E. Moore, “Cramming more components onto integrated circuits”, *Electronics* **38** (1965).
- [29] M. Fuechsle, J. Miwa, S. Mahapatra and et al., “A single-atom transistor”, *Nature Nanotech* **7**, 242 (2012).
- [30] J. Frougier and D. Guo, *Introducing the world’s first 2nm node chip*, (2021) <https://research.ibm.com/blog/2-nm-chip> (visited on 03/03/2023).
- [31] B. News, *Gordon Moore, Intel co-founder and creator of Moore’s Law, dies aged 94*, (2023) <https://www.bbc.co.uk/news/world-us-canada-65073812> (visited on 26/03/2023).
- [32] J. Gribbin, *Computing with quantum cats* (Penguin, London, 2015).
- [33] R. Landauer, “Irreversibility and heat generation in the computing process”, *IBM Journal of Research and Development* **5**, 183 (1961).
- [34] C. H. Bennett, “Logical reversibility of computation”, *IBM Journal of Research and Development* **17**, 525 (1973).
- [35] P. 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**, 563 (1980).
- [36] R. Feynman, “Simulating physics with computers”, *Int. J. Theor. Phys.* **21**, 467 (1982).
- [37] J. Bell, “On the Einstein Podolsky Rosen Paradox”, *Physics Physique Fizika* **1**, 195 (1964).
- [38] R. Cleve, A. Ekert, C. Macchiavello and M. Mosca, “Quantum algorithms revisited”, *Proceedings of the Royal Society A* **454** (1998).
- [39] D. Deutsch and R. Jozsa, “Rapid solution of problems by quantum computation”, *Proceedings of the Royal Society A* **439** (1992).
- [40] P. Shor, “Algorithms for quantum computation: discrete logarithms and factoring”, in *Proceedings 35th annual symposium on foundations of computer science* (1994), pp. 124–134.

- [41] D. R. Simon, “On the Power of Quantum Computation”, in Proceedings 35th annual symposium on foundations of computer science (1994), pp. 116–123.
- [42] L. K. Grover, “A fast quantum mechanical algorithm for database search”, in Stoc '96: proceedings of the twenty-eighth annual acm symposium on theory of computing (1996), pp. 212–219.
- [43] J. I. Cirac and P. Zoller, “Quantum computations with cold trapped ions”, *Phys. Rev. Lett.* **74**, 4091 (1995).
- [44] I. Pogorelov, T. Feldker, C. D. Marciniak, L. Postler, G. Jacob, O. Kriegelsteiner, V. Podlesnic, M. Meth, V. Negnevitsky, M. Stadler, B. Höfer, C. Wächter, K. Lakhmanskiy, R. Blatt, P. Schindler and T. Monz, “Compact ion-trap quantum computing demonstrator”, *PRX Quantum* **2**, 020343 (2021).
- [45] J. Clarke and F. Wilhelm, “Superconducting quantum bits”, *Nature* **453**, 1031 (2008).
- [46] W. M. Kaminsky, S. Lloyd and T. P. Orlando, “Scalable superconducting architecture for adiabatic quantum computation”, arXiv:040390 (2004).
- [47] D. Loss and D. P. DiVincenzo, “Quantum computation with quantum dots”, *Phys. Rev. A* **57**, 120 (1998).
- [48] S. Lloyd, “A potentially realizable quantum computer”, *Science* **261**, 1569 (1993).
- [49] D. P. DiVincenzo, “Two-bit gates are universal for quantum computation”, *Phys. Rev. A* **51**, 1015 (1995).
- [50] D. P. DiVincenzo, “Quantum computation”, *Science* **270**, 255 (1995).
- [51] L. Vandersypen, M. Steffen, G. Breyta and et al, “Experimental realization of Shor’s quantum factoring algorithm using nuclear magnetic resonance”, *Nature* **414**, 883 (2001).
- [52] W. S. Warren, “The Usefulness of NMR Quantum Computing”, *Science* **277**, 1688 (1997).
- [53] M. Khazali and K. Mølmer, “Fast multiqubit gates by adiabatic evolution in interacting excited-state manifolds of Rydberg atoms and superconducting circuits”, *Phys. Rev. X* **10**, 021054 (2020).
- [54] L. Henriët, L. Beguin, A. Signoles, T. Lahaye, A. Browaeys, G.-O. Raymond and C. Jurczak, “Quantum computing with neutral atoms”, *Quantum* **4**, 327 (2020).
- [55] E. Knill, R. Laflamme and G. Milburn, “A scheme for efficient quantum computation with linear optics”, *Nature* **409**, 46 (2001).

- [56] A. Nizovtsev, S. Kilin, F. Jelezko and et al., “A quantum computer based on NV centers in diamond: Optically detected nutations of single electron and nuclear spins”, *Opt. Spectrosc.* **99**, 233 (2005).
- [57] R. Raussendorf and H. J. Briegel, “A one-way quantum computer”, *Phys. Rev. Lett.* **86**, 5188 (2001).
- [58] V. M. Kendon, K. Nemoto and W. J. Munro, “Quantum analogue computing”, *Phil. Trans. R. Soc. A* **368** (2010).
- [59] V. Kendon, “Quantum computing using continuous-time evolution”, *Interface Focus* **10** (2020).
- [60] E. Farhi and S. Gutmann, “Quantum computation and decision trees”, *Phys. Rev. A* **58**, 915 (1998).
- [61] A. M. Childs and J. Goldstone, “Spatial search by quantum walk”, *Physical Review A - Atomic, Molecular, and Optical Physics* **70**, 022314 (2004).
- [62] A. M. Childs, “Universal computation by quantum walk”, *Phys. Rev. Lett.* **102**, 180501 (2009).
- [63] J. G. Morley, N. Chancellor, S. Bose and V. Kendon, “Quantum search with hybrid adiabatic–quantum-walk algorithms and realistic noise”, *Phys. Rev. A* **99**, 022339 (2019).
- [64] A. Callison, N. Chancellor, F. Mintert and V. Kendon, “Finding spin glass ground states using quantum walks”, *New Journal of Physics* **21**, 123022 (2019).
- [65] E. Farhi, J. Goldstone, S. Gutmann and M. Sipser, “Quantum Computation by Adiabatic Evolution”, arXiv:0001106 (2000).
- [66] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren and D. Preda, “A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem”, *Science* **292**, 472 (2001).
- [67] N. J. Cerf and J. Roland, “Quantum search by local adiabatic evolution”, *Phys. Rev. A* **65**, 1 (2002).
- [68] A. Finnila, M. Gomez, C. Sebenik, C. Stenson and J. Doll, “Quantum annealing: a new method for minimizing multidimensional functions”, *Chemical Physics Letters* **219**, 343 (1994).
- [69] T. Kadowaki and H. Nishimori, “Quantum annealing in the transverse Ising model”, *Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics* **58**, 5355 (1998).
- [70] E. J. Crosson and D. A. Lidar, “Prospects for Quantum Enhancement with Diabatic Quantum Annealing”, *Nature Reviews Physics* **3**, 466 (2021).

- [71] A. D. King, S. Suzuki, J. Raymond, A. Zucca, T. Lanting, F. Altomare, A. J. Berkley, S. Ejtemaee, E. Hoskinson, S. Huang, E. Ladizinsky, A. J. MacDonald, G. Marsden, T. Oh, G. Poulin-Lamarre, M. Reis, C. Rich, Y. Sato, J. D. Whittaker, J. Yao, R. Harris, D. A. Lidar, H. Nishimori and M. H. Amin, “Coherent quantum annealing in a programmable 2,000 qubit Ising chain”, *Nature Physics* **18**, 1324 (2022).
- [72] N. G. Dickson, M. W. Johnson, M. H. Amin, R. Harris, F. Altomare, A. J. Berkley, P. Bunyk, J. Cai, E. M. Chapple, P. Chavez, F. Cioata, T. Cirip, P. deBuen, M. Drew-Brook, C. Enderud, S. Gildert, F. Hamze, J. P. Hilton, E. Hoskinson, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Lanting, T. Mahon, R. Neufeld, T. Oh, I. Perminov, C. Petroff, A. Przybysz, C. Rich, P. Spear, A. Teaciuc, M. C. Thom, E. Tolkacheva, S. Uchaikin, J. Wang, A. B. Wilson, Z. Merali and G. Rose, “Thermally assisted quantum annealing of a 16-qubit problem”, *Nature Communications* **4** (2013).
- [73] N. Chancellor, “Modernizing quantum annealing using local searches”, *New J. Phys.* **19** (2017).
- [74] E. Rieffel, D. Venturelli, B. O’Gorman, B. D. Minh, E. M. Prystay and V. N. Smelyanskiy, “A case study in programming a quantum annealer for hard operational planning problems.”, *Quantum Inf Process* **14**, 1 (2015).
- [75] D. Venturelli, D. J. J. Marchand and G. Rojo, “Quantum annealing implementation of job-shop scheduling”, arXiv:1506.08479 (2016).
- [76] K. Ikeda, Y. Nakamura and T. Humble, “Application of quantum annealing to nurse scheduling problem”, *Sci Rep* **9** (2019).
- [77] A. Perdomo-Ortiz, N. Dickson, M. Drew-Brook, G. Rose and A. Aspuru-Gzik, “Finding low-energy conformations of lattice protein models by quantum annealing”, *Sci Rep* **2** (2012).
- [78] V. K. Mulligan, H. Melo, H. I. Merritt, S. Slocum, B. D. Weitzner, A. M. Watkins, P. D. Renfrew, C. Pelissier, P. S. Arora and R. Bonneau, “Designing peptides on a quantum computer”, *bioRxiv*, 10.1101/752485 (2020).
- [79] M. Streif, F. Neukart and M. Leib, “Solving Quantum Chemistry Problems with a D-Wave Quantum Annealer”, arXiv:1811.05256 (2019).
- [80] R. Harris, Y. Sato, A. J. Berkley, M. Reis, F. Altomare, M. H. Amin, K. Boothby, P. Bunyk, C. Deng, C. Enderud, S. Huang, E. Hoskinson, M. W. Johnson, E. Ladizinsky, N. Ladizinsky, T. Lanting, R. Li, T. Medina, R. Molavi, R. Neufeld, T. Oh, I. Pavlov, I. Perminov, G. Poulin-Lamarre, C. Rich, A. Smirnov, L. Swenson, N. Tsai, M. Volkmann, J. Whittaker and J.

- Yao, “Phase transitions in a programmable quantum spin glass simulator”, *Science* **361**, 162 (2018).
- [81] G. Rosenberg, P. Haghnegahdar, P. Goddard, P. Carr, K. Wu and M. L. de Prado, “Solving the optimal trading trajectory problem using a quantum annealer”, *IEEE Journal of Selected Topics in Signal Processing* **10**, 1053 (2016).
- [82] D. Venturelli and A. Kondratyev, “Reverse quantum annealing approach to portfolio optimization problems”, *Quantum Mach. Intell.* **1**, 17 (2019).
- [83] S. Yarkoni, E. Raponi, T. Bäck and S. Schmitt, “Quantum annealing for industry applications: introduction and review”, *Reports on Progress in Physics* **85** (2022).
- [84] P. Hauke, H. G. Katzgraber, W. Lechner, H. Nishimori and W. D. Oliver, “Perspectives of quantum annealing: methods and implementations”, *Reports on Progress in Physics* **83** (2020).
- [85] V. Choi, “Adiabatic Quantum Algorithms for the NP-Complete Maximum-Weight Independent Set, Exact Cover and 3SAT Problems”, arXiv:1004.2226, 1 (2010).
- [86] V. Choi, “Minor-embedding in adiabatic quantum computation: II. minor-universal graph design”, *Quantum Information Processing* **10**, 343 (2011).
- [87] V. Choi, “Minor-embedding in adiabatic quantum computation: I. the parameter setting problem”, *Quantum Information Processing*, 193 (2008).
- [88] *Minorminer*, (2018) <https://github.com/dwavesystems/minorminer>.
- [89] N. Chancellor, P. J. Crowley, T. Đurić, W. Vinci, M. H. Amin, A. G. Green, P. A. Warburton and G. Aeppli, “Error measurements for a quantum annealer using the one-dimensional Ising model with twisted boundaries”, *npj Quantum Information* **8** (2022).
- [90] N. Chancellor, A. Callison, V. Kendon and F. Mintert, *Finding spin-glass ground states using quantum walks, data archive at Durham University, UK*, for spin glass instances used in [64], <https://doi.org/10.15128/r21544bp097>, 2019.
- [91] D. Sherrington and S. Kirkpatrick, “Solvable model of a spin-glass”, *Physical Review Letters* **35**, 1792 (1975).
- [92] H. D. Zeh, “On the interpretation of measurement in quantum theory”, *Foundations of Physics* **1**, 69 (1970).
- [93] M. Schlosshauer, “Decoherence, the measurement problem, and interpretations of quantum mechanics”, *Reviews of Modern Physics* **76**, 1267 (2004).

- [94] D. Bacon, “Decoherence, Control, and Symmetry in Quantum Computers”, arXiv:0305025 (2003).
- [95] D. A. Lidar and K. Birgitta Whaley, “Decoherence-Free Subspaces and Subsystems”, *Irreversible Quantum Dynamics*, Springer **622** (2003).
- [96] J. A. Sherman, M. J. Curtis, D. J. Szwer, D. T. Allcock, G. Imreh, D. M. Lucas and A. M. Steane, “Experimental recovery of a qubit from partial collapse”, *Physical Review Letters* **111**, 1 (2013).
- [97] J. Ghosh, A. G. Fowler, J. M. Martinis and M. R. Geller, “Understanding the effects of leakage in superconducting quantum-error-detection circuits”, *Physical Review A - Atomic, Molecular, and Optical Physics* **88**, 1 (2013).
- [98] A. Galiutdinov, “Quantum Zeno effect: A possible resolution to the leakage problem in superconducting quantum computing architectures”, arXiv:1805.06877 (2018).
- [99] A. Strikis, A. Datta and G. C. Knee, “Quantum leakage detection using a model-independent dimension witness”, *Physical Review A* **99** (2019).
- [100] M. A. Rol, F. Battistel, F. K. Malinowski, C. C. Bultink, B. M. Tarasinski, R. Vollmer, N. Haider, N. Muthusubramanian, A. Bruno, B. M. Terhal and L. Dicarlo, “Fast, High-Fidelity Conditional-Phase Gate Exploiting Leakage Interference in Weakly Anharmonic Superconducting Qubits”, *Physical Review Letters* **123**, 120502 (2019).
- [101] M. Grassl, T. Beth and T. Pellizzari, “Codes for the quantum erasure channel”, *Physical Review A - Atomic, Molecular, and Optical Physics* **56**, 33 (1997).
- [102] M. Suchara, A. W. Cross and J. M. Gambetta, “Leakage suppression in the toric code”, *IEEE International Symposium on Information Theory - Proceedings*, 1119 (2015).
- [103] N. Delfosse and G. Zémor, “Linear-time maximum likelihood decoding of surface codes over the quantum erasure channel”, *Physical Review Research* **2**, 1 (2020).
- [104] N. Delfosse and N. H. Nickerson, “Almost-linear time decoding algorithm for topological codes”, *Quantum* **5** (2021).
- [105] T. M. Stace, S. D. Barrett and A. C. Doherty, “Thresholds for topological codes in the presence of loss”, *Physical Review Letters* **102**, 1 (2009).
- [106] T. M. Stace and S. D. Barrett, “Error correction and degeneracy in surface codes suffering loss”, *Physical Review A - Atomic, Molecular, and Optical Physics* **81**, 1 (2010).

- [107] D. Vodola, D. Amaro, M. A. Martin-Delgado and M. Müller, “Twins Percolation for Qubit Losses in Topological Color Codes”, *Physical Review Letters* **121**, 1 (2018).
- [108] S. H. Van Der Ploeg, A. Izmailkov, A. M. Van Den Brink, U. Hübner, M. Grajcar, E. Il’ichev, H. G. Meyer and A. M. Zagoskin, “Controllable coupling of superconducting flux qubits”, *Physical Review Letters* **98**, 2 (2007).
- [109] P. W. Shor, “Fault-Tolerant Quantum Computation”, in 37th symposium on foundations of computing, IEEE computer society press (1996), pp. 56–65.
- [110] E. Knill and R. Laflamme, “Concatenated Quantum Codes”, arXiv:9608012, 1 (1996).
- [111] E. Knill, R. Laflamme and W. H. Zurek, “Resilient quantum computation”, *Science* **279**, 342 (1998).
- [112] D. Aharonov, A. Kitaev and J. Preskill, “Fault-tolerant quantum computation with long-range correlated noise”, *Physical Review Letters* **96**, 050504 (2006).
- [113] M. Sarovar and K. C. Young, “Error suppression and error correction in adiabatic quantum computation: Non-equilibrium dynamics”, *New Journal of Physics* **15** (2013).
- [114] K. C. Young, M. Sarovar and R. Blume-Kohout, “Error suppression and error correction in adiabatic quantum computation: Techniques and challenges”, *Physical Review X* **3**, 041013 (2013).
- [115] I. Marvian and D. A. Lidar, “Quantum error suppression with commuting Hamiltonians: Two local is too local”, *Physical Review Letters* **113**, 260504 (2014).
- [116] R. W. Hamming, “Error Detecting and Error Correcting Codes”, *The Bell System Technical Journal* **196**, 519 (1923).
- [117] R. Raussendorf and H. J. Briegel, “A one-way quantum computer”, *Physical Review Letters* **86**, 5188 (2001).
- [118] C. Horsman, A. G. Fowler, S. Devitt and R. V. Meter, “Surface code quantum computing by lattice surgery”, *New J. Phys* **14**, 123011 (2012).
- [119] S. M. Girvin, “Introduction to quantum error correction and fault tolerance”, arXiv:2111.08894 (2021).
- [120] S. P. Jordan, E. Farhi and P. W. Shor, “Error-correcting codes for adiabatic quantum computation”, *Phys. Rev. A* **74**, 052322 (2006).

- [121] A. D. Bookatz, E. Farhi and L. Zhou, “Error suppression in Hamiltonian-based quantum computation using energy penalties”, *Physical Review A - Atomic, Molecular, and Optical Physics* **92**, 022317 (2015).
- [122] I. Marvian, “Exponential suppression of decoherence and relaxation of quantum systems using energy penalty”, arXiv:1602.03251, 1 (2016).
- [123] M. Marvian and D. A. Lidar, “Error suppression for Hamiltonian quantum computing in Markovian environments”, *Physical Review A* **95**, 032302 (2017).
- [124] D. A. Lidar, “Towards fault tolerant adiabatic quantum computation”, *Physical Review Letters* **100**, 160506 (2008).
- [125] G. Quiroz and D. A. Lidar, “High-fidelity adiabatic quantum computation via dynamical decoupling”, *Physical Review A - Atomic, Molecular, and Optical Physics* **86**, 042333 (2012).
- [126] A. Ganti, U. Onunkwo and K. Young, “Family of $[[6k, 2k, 2]]$ codes for practical and scalable adiabatic quantum computation”, *Physical Review A - Atomic, Molecular, and Optical Physics* **89**, 042313 (2014).
- [127] Z. Jiang and E. G. Rieffel, “Non-commuting two-local Hamiltonians for quantum error suppression”, *Quantum Information Processing* **16**, 89 (2017).
- [128] M. Marvian and D. A. Lidar, “Error Suppression for Hamiltonian-Based Quantum Computation Using Subsystem Codes”, *Physical Review Letters* **118**, 030504 (2017).
- [129] M. Marvian and S. Lloyd, “Robust universal Hamiltonian quantum computing using two-body interactions”, arXiv:1911.01354, 1 (2019).
- [130] D. Bacon, K. R. Brown and K. B. Whaley, “Coherence-Preserving Quantum Bits”, *Physical Review Letters* **87**, 247902–247902 (2001).
- [131] D. A. Lidar, “Arbitrary-time error suppression for Markovian adiabatic quantum computing using stabilizer subspace codes”, *Physical Review A* **100**, 022326 (2019).
- [132] M. Sarovar and G. J. Milburn, “Continuous quantum error correction by cooling”, *Physical Review A - Atomic, Molecular, and Optical Physics* **72**, 012306 (2005).
- [133] J. Atalaya, S. Zhang, M. Y. Niu, A. Babakhani, H. C. H. Chan, J. M. Epstein and K. B. Whaley, “Continuous quantum error correction for evolution under time-dependent Hamiltonians”, *Phys. Rev. A* **103**, 042406 (2021).

- [134] N. Mohseni, M. Narozniak, A. N. Pyrkov, V. Ivannikov, J. P. Dowling and T. Byrnes, “Error suppression in adiabatic quantum computing with qubit ensembles”, *npj Quantum Inf* **7**, 71 (2021).
- [135] G. A. Paz-Silva, A. T. Rezakhani, J. M. Dominy and D. A. Lidar, “Zeno Effect for Quantum Computation and Control”, *Physical Review Letters* **108**, 080501 (2012).
- [136] K. L. Pudenz, T. Albash and D. A. Lidar, “Error-corrected quantum annealing with hundreds of qubits”, *Nature Communications* **5**, 3243 (2014).
- [137] W. Vinci, T. Albash, G. Paz-Silva, I. Hen and D. A. Lidar, “Quantum annealing correction with minor embedding”, *Physical Review A - Atomic, Molecular, and Optical Physics* **92**, 042310 (2015).
- [138] K. L. Pudenz, T. Albash and D. A. Lidar, “Quantum annealing correction for random Ising problems”, *Physical Review A - Atomic, Molecular, and Optical Physics* **91**, 042302 (2015).
- [139] W. Vinci, T. Albash and D. A. Lidar, “Nested quantum annealing correction”, *npj Quantum Information* **2**, 16017 (2016).
- [140] S. Matsuura, H. Nishimori, T. Albash and D. A. Lidar, “Mean Field Analysis of Quantum Annealing Correction”, *Physical Review Letters* **116**, 220501 (2016).
- [141] A. Mishra, T. Albash and D. A. Lidar, “Performance of two different quantum annealing correction codes”, *Quantum Information Processing* **15**, 609 (2016).
- [142] S. Matsuura, H. Nishimori, W. Vinci, T. Albash and D. A. Lidar, “Quantum-annealing correction at finite temperature: Ferromagnetic p -spin models”, *Physical Review A* **95**, 022308 (2017).
- [143] W. Vinci and D. A. Lidar, “Scalable effective-temperature reduction for quantum annealers via nested quantum annealing correction”, *Physical Review A* **97**, 022308 (2018).
- [144] S. Matsuura, H. Nishimori, W. Vinci and D. A. Lidar, “Nested quantum annealing correction at finite temperature : p -spin models”, *Physical Review A* **99**, 062307 (2019).
- [145] R. Y. Li, T. Albash and D. A. Lidar, “Limitations of error corrected quantum annealing in improving the performance of Boltzmann machines”, *Quantum Science and Technology* **5**, 1 (2020).
- [146] J. Brooke, D. Bitko, T. F. Rosenbaum and G. Aeppli, “Quantum annealing of a disordered magnet”, *Science* **284**, 779 (1999).

- [147] T. Lanting, A. J. Przybysz, A. Y. Smirnov, F. M. Spedalieri, M. H. Amin, A. J. Berkley, R. Harris, F. Altomare, S. Boixo, P. Bunyk, N. Dickson, C. Enderud, J. P. Hilton, E. Hoskinson, M. W. Johnson, E. Ladizinsky, N. Ladizinsky, R. Neufeld, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, S. Uchaikin, A. B. Wilson and G. Rose, “Entanglement in a quantum annealing processor”, *Physical Review X* **4**, 021041 (2014).
- [148] V. S. Denchev, S. Boixo, S. V. Isakov, N. Ding, R. Babbush, V. Smelyanskiy, J. Martinis and H. Neven, “What is the computational value of finite-range tunneling?”, *Physical Review X* **6**, 031015 (2016).
- [149] M. W. Johnson, M. H. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, C. J. Truncik, S. Uchaikin, J. Wang, B. Wilson and G. Rose, “Quantum annealing with manufactured spins”, *Nature* **473**, 194 (2011).
- [150] S. Boixo, V. N. Smelyanskiy, A. Shabani, S. V. Isakov, M. Dykman, V. S. Denchev, M. H. Amin, A. Y. Smirnov, M. Mohseni and H. Neven, “Computational multiqubit tunnelling in programmable quantum annealers”, *Nature Communications* **7**, 10327 (2016).
- [151] S. Matsuura, “Mean Field Quantum Annealing Correction”, *Journal of the Physical Society of Japan* **88**, 1 (2019).
- [152] G. Van Rossum and F. L. Drake, *Python language reference manual*, Network theory (United Kingdom, 2003).
- [153] Anaconda, *Anaconda software distribution*, version 2-2.4.0. 2016.
- [154] JetBrains, *Pycharm IDE - community edition*, version 2021.2.3. 2021.
- [155] Microsoft, *Visual studio code*, (2023) <https://github.com/microsoft/vscode> (visited on 16/03/2023).
- [156] L. Torvalds, *Git*, (2023) <https://git-scm.com/> (visited on 16/03/2023).
- [157] Github, *Github*, (2023) <https://github.com/github> (visited on 16/03/2023).
- [158] T. E. Oliphant, *A guide to numpy* (Tregol Publishing USA, 2006).
- [159] E. Jones, T. Oliphant and P. Peterson, *Scipy: open source scientific tools for python*, 2001-.
- [160] W. McKinney, “Data Structures for Statistical Computing in Python”, *Proceedings of the 9th Python in Science Conference* **1**, 56 (2010).
- [161] J. D. Hunter, “Matplotlib: a 2d graphics environment”, *Computing in Science and Engineering* **9**, 90 (2007).

- [162] A. H. Land and A. G. Doig, “An automatic method of solving discrete programming problems”, *Econometrica* **28**, 497 (1960).
- [163] D. R. Morrison, S. H. Jacobson, J. J. Sauppe and E. C. Sewell, “Branch-and-bound algorithms: a survey of recent advances in searching, branching, and pruning”, *Discrete Optimization* **19**, 79 (2016).
- [164] E. L. Lawler and D. E. Wood, “Branch-and-bound methods: a survey”, *Operations Research* **14**, 699 (1966).
- [165] D. Bertsimas and J. N. Tsitsiklis, *Introduction to linear optimization* (Athena Scientific, Nashua, NH, USA, 1997).
- [166] C. H. Papadimitrou and K. Steiglitz, *Combinatorial optimization: algorithms and complexity* (Dover Publications Inc., Mineola, NY, USA, 2000).
- [167] A. Callison, “Continuous-time quantum computing”, PhD thesis (Imperial College London, 2021).
- [168] J. Gray, *Journal of Open Source Software* **3**, 819 (2018).
- [169] SciPy, *Scipy.integrate.complex_ode*, (2023) https://docs.scipy.org/doc/scipy/reference/generated/scipy.integrate.complex%5C_ode.html (visited on 17/03/2023).
- [170] I. Hughes and T. Hase, *Measurements and their uncertainties : a practical guide to modern error analysis* (Oxford University Press, 2010).
- [171] H. De Raedt, F. Jin, D. Willsch, M. Willsch, N. Yoshioka, N. Ito, S. Yuan and K. Michielsen, “Massively parallel quantum computer simulator, eleven years later”, *Computer Physics Communications* **237**, 47 (2019).
- [172] E. Pednault, J. A. Gunnels, G. Nannicini, L. Horesh, T. Magerlein, E. Solomonik, E. W. Draeger, E. T. Holland and R. Wisnieff, “Pareto-Efficient Quantum Circuit Simulation Using Tensor Contraction Deferral”, arXiv:1710.05867 (2020).
- [173] Z.-Y. Chen, Q. Zhou, C. Xue, X. Yang, G.-C. Guo and G.-P. Guo, “64-qubit quantum circuit simulation”, *Science Bulletin* **63**, 964 (2018).
- [174] L. T. Brady, L. Kocia, P. Bienias, A. Bapat, Y. Kharkov and A. V. Gorshkov, “Behavior of analog quantum algorithms”, arXiv:2107.01218, arXiv:2107.01218 (2021).

Colophon

This thesis is based on a template developed by Matthew Townson and Andrew Reeves. It was typeset with L^AT_EX 2_ε. It was created using the *memoir* package, maintained by Lars Madsen, with the *madsen* chapter style. The font used is Latin Modern, derived from fonts designed by Donald E. Knuth.