



Università degli Studi di Napoli *Federico II*

DOTTORATO DI RICERCA IN
FISICA

Ciclo XXXIII

Coordinatore: Prof. Salvatore Capozziello

**Quantum annealing and
advanced optimization strategies
of closed and open quantum systems**

Settore Scientifico Disciplinare FIS/03

Dottorando

Gianluca Passarelli

Tutori

Prof. Vittorio Cataudella

Dott. Procolo Lucignano

Anni 2018/2021

CONTENTS

CONTENTS	II
LIST OF ACRONYMS	IV
LIST OF FIGURES	VI
INTRODUCTION	VIII
1 SUMMARY OF THE MAIN RESULTS	1
1.1 Quantum annealing	1
1.2 D-Wave and the embedding problem	2
1.3 Ferromagnetic p -spin model	3
1.4 Open quantum systems	4
1.5 Dissipative dynamics: p -spin model	6
1.6 Speeding up adiabatic quantum computation	9
1.7 Conclusions	10
2 QUANTUM ANNEALING	11
2.1 Introduction	11
2.2 Adiabatic quantum computation	12
2.3 Adiabatic theorem	13
2.4 D-Wave	17
2.5 Advanced annealing schedules	19
2.6 Conclusions	25
3 FERROMAGNETIC P-SPIN MODEL	26
3.1 Introduction	26
3.2 Hamiltonian	27
3.3 Adiabatic reverse annealing	28
3.4 Spin sectors	30
3.5 Spectral properties and minimal gap	31
3.6 The large- p limit	33
3.7 The embedding problem	34

3.8	Conclusions	41
4	OPEN QUANTUM SYSTEMS	43
4.1	Introduction	43
4.2	Dynamical semigroup	43
4.3	Quantum jumps	49
4.4	p -spin model and dissipation	51
4.5	Beyond the Born-Markov approximation	54
4.6	Conclusions	62
5	RESULTS	63
5.1	Introduction	63
5.2	Quantum annealing	63
5.3	Quantum annealing with pauses	68
5.4	Iterated reverse annealing	74
5.5	Adiabatic reverse annealing	82
5.6	Conclusions	87
6	BEYOND ADIABATIC EVOLUTIONS	89
6.1	Introduction	89
6.2	Shortcuts to adiabaticity	90
6.3	Variational approach	93
6.4	Application: p -spin model, variational approach	97
6.5	Genetic approach	106
6.6	Application: p -spin model, genetic approach	108
6.7	Conclusions	110
	CONCLUSIONS	111
A	MEAN FIELD ANALYSIS	114
A.1	Quantum annealing	114
A.2	Adiabatic reverse annealing	117
B	AME AND MCWF CPU COST	120
	BIBLIOGRAPHY	122

LIST OF ACRONYMS

AME	Adiabatic master equation
AQC	Adiabatic quantum computation
ARA	Adiabatic reverse annealing
CA	Cyclic ansatz
CB	Collective bath
CD	Counterdiabatic
CPT	Classical phase transition
CPU	Central processing unit
CRAB	Chopped random basis
DMRG	Density matrix renormalization group
GPGPU	General-purpose graphics processing unit
GPU	Graphics processing unit
GS	Ground state
HEOM	Hierarchical equation of motion
IB	Independent bath
IRA	Iterated reverse annealing
KMS	Kubo-Martin-Schwinger
LMG	Lipkin-Meshkov-Glick
LS	Lamb shift
LZ	Landau-Zener
MCWF	Monte Carlo wave function
NC	Nested commutators
OCT	Optimal control theory
PIMC	Path-integral Monte Carlo
PUBO	Polynomial unconstrained binary optimization
QA	Quantum annealing

QPT	Quantum phase transition
QPU	Quantum processing unit
QUAPI	Quasiadiabatic propagator
QUBO	Quadratic unconstrained binary optimization
RWA	Rotating wave approximation
SA	Simulated annealing
SBM	Spin-boson model
SIL	Short-iterative Lanczos
SQUID	Superconducting quantum-interference device
STA	Shortcut to adiabaticity
TTS	Time to solution

LIST OF FIGURES

1.1	Annealing schedules used by the D-Wave QPU	2
1.2	Unit cells of the Chimera graph	3
1.3	Physical versus logical qubits for the p -to-2-body decomposition . .	4
1.4	Free energy versus magnetization, p -spin model	5
2.1	Simulated versus quantum annealing	14
2.2	Adiabatic versus nonadiabatic (but favorable) dynamics	16
2.3	Annealing schedules used by the D-Wave QPU	17
2.4	Unit cells of the Chimera graph	18
2.5	Example of minor embedding	19
2.6	Pauses on unitary dynamics, diabatic versus adiabatic	21
2.7	Success probability versus pause length, single qubit	22
2.8	Annealing fraction versus time	23
3.1	Minimal gap, $p = 2$	32
3.2	Minimal gap, $p = 3$	33
3.3	Minimal gap of ARA, $p = 3$	34
3.4	AND embedding	36
3.5	Schematics of a genetic algorithm	39
3.6	Error probability in the effective QUBO	41
3.7	Physical versus logical qubits for the p -to-2-body decomposition . .	42
4.1	Open quantum system	45
4.2	Scaling of the CPU time as a function of N , Lindblad master equation	48
4.3	Convergence to AME and spectrum of thermal processes, m_z	53
4.4	Convergence to AME and spectrum of thermal processes, m_x	54
4.5	Thermal processes versus time	55
4.6	Discretized spectral density of the bath	56
4.7	SIL error in the pure dephasing model	58
4.8	Time evolution of $\langle \sigma_z(t) \rangle$ in the SBM	59
4.9	$\langle \sigma_x \rangle^{\text{eq}}$ versus η , spin-boson model	60
4.10	Error probability, single qubit, AME versus SIL	61
5.1	ϵ_{res} versus t_f , $p = 2$, unitary	64

5.2	ϵ_{res} versus t_f , $p = 5$, unitary	65
5.3	ϵ_{res} versus t_f , $p = 5$, dissipative, finite temperature	66
5.4	ϵ_{res} versus t_f , $p = 5$, dissipative, $T = 0$	67
5.5	ϵ_{res} versus t_f , $p = 5$ and $p = 7$, quantum versus simulated annealing	68
5.6	Eigenvalues of the p -spin and search Hamiltonians	70
5.7	$N = 20$, $p = 19$, unitary and dissipative dynamics, no pauses	71
5.8	Fidelity versus pausing time s_p	72
5.9	Fidelity versus s_p and t_p , density plot	73
5.10	$p_{\text{gs}}(s)$ versus s , optimal versus worst pausing point	74
5.11	IRA, success probability versus inversion time, unitary	76
5.12	Maximum fidelity in IRA versus starting magnetization, unitary	77
5.13	IRA, success probability versus inversion time, dissipative	78
5.14	IRA, fidelity versus s_{inv} , collective and individual bath	79
5.15	Maximum fidelity in IRA versus starting magnetization, dissipative	80
5.16	IRA, fidelity versus s_{inv} , dissipative with a pause	81
5.17	IRA, fidelity versus s_{inv} , collective and individual bath, pausing	82
5.18	ARA, $\langle m_z \rangle$ versus s	83
5.19	TTS versus t_f , ARA, $\Gamma = 1$	84
5.20	TTS versus t_f , ARA, $\Gamma = 2$	85
5.21	TTS versus t_f , ARA, $\Gamma = 1$, grouping by η	86
5.22	TTS versus t_f , ARA, $\Gamma = 2$, grouping by η	87
6.1	Pictorial representation of counterdiabatic driving	90
6.2	Decision tree of OCT methods	92
6.3	Fidelity as a function of N , $p = 1$, diabatic versus CD	99
6.4	Fidelity as a function of N , $p = 2$, diabatic versus CD	100
6.5	Fidelity as a function of N , $p = 3$, diabatic versus CD	101
6.6	$p_{\text{gs}}(s)$ versus s , standard QA versus NC versus CA	102
6.7	$p_{\text{gs}}(s)$ versus s , finite-range p -spin model	104
6.8	Fidelity distributions, random p -spin model, standard QA	105
6.9	Fidelity distributions, random p -spin model, NC and CA	106
6.10	Average fitness versus generation number, genetic OCT optimization	108
6.11	$p_{\text{gs}}(s)$ versus s , genetic OCT operators	109
A.1	Phase diagram, $p = 2$	115
A.2	Phase diagram, $p = 3$	116
A.3	Phase diagram of ARA, $p = 3$	118
A.4	Magnetization in ARA, $p = 3$	119
B.1	CPU time as a function of N , MCWF versus AME	121

INTRODUCTION

Adiabatic quantum computation and quantum annealing are powerful methods designed to solve optimization problems more efficiently than classical computers. The idea is to encode the solution to the optimization problem into the ground state of an Ising Hamiltonian, which can be hard to diagonalize exactly and can involve long-range and multiple-body interactions. The adiabatic theorem of quantum mechanics is exploited to drive a quantum system towards the target ground state.

More precisely, the evolution starts from the ground state of a transverse field Hamiltonian, providing the quantum fluctuations needed for quantum tunneling between trial solution states. The Hamiltonian is slowly changed to target the Ising Hamiltonian of interest. If this evolution is infinitely slow, the system is guaranteed to stay in its ground state. Hence, at the end of the dynamics, the state can be measured, yielding the solution to the problem. In real devices, such as in the D-Wave quantum annealers, the evolution lasts a finite amount of time, which gives rise to Landau-Zener diabatic transitions, and occurs in the presence of an environment, inducing thermal excitations outside the ground state. Both these limitations have to be carefully addressed in order to understand the true potential of these devices. The present thesis aims to find strategies to overcome these limitations.

In the first part of this work, we address the effects of dissipation. We show that a low-temperature Markovian environment can improve quantum annealing, compared with the closed-system case, supporting other previous results known in the literature as thermally-assisted quantum annealing.

In the second part, we combine dissipation with advanced annealing schedules, featuring pauses and iterated or adiabatic reverse annealing, which, in combination with low-temperature environments, can favor relaxation to the ground state and improve quantum annealing compared to the standard algorithm. In general, however, dissipation is detrimental for quantum annealing especially when the annealing time is longer than the typical thermal relaxation and decoherence time scales. For this reason, it is essential to devise shortcuts to adiabaticity so as to reach the adiabatic limit for relatively short times in order to decrease the impact of thermal noise on the performances of QA.

To this end, in the last part of this thesis we study the counterdiabatic driving

approach to quantum annealing. In counterdiabatic driving, a new term is added to the Hamiltonian to suppress Landau-Zener transitions and achieve adiabaticity for any finite sweep rate. Although the counterdiabatic potential is nonlocal and hardly implementable on quantum devices, we can obtain approximate potentials that dramatically enhance the success probability of short-time quantum annealing following a variational formulation.

The themes here outlined are thoroughly discussed throughout this thesis, which is organized as follows.

- IN CHAPTER 1** we give an extended summary of the main results of this project, leaving all technical details to later chapters.
- IN CHAPTER 2** we describe adiabatic quantum computation and some of its variants. Moreover, we briefly discuss the adiabatic theorem and the problem of embedding NP-hard problems into sparsely-connected quantum annealers. Some results concerning adiabatic quantum computation in the unitary limit are presented.
- IN CHAPTER 3** we introduce the ferromagnetic p -spin model, which constitutes the *leitmotiv* of this project. Despite being exactly tractable in the thermodynamic limit, this system is a good toy model of NP-hard optimization tasks and can be used as a test bed for numerical simulations.
- IN CHAPTER 4** we discuss open quantum systems. Physical quantum annealers interact with their surrounding environment, inducing thermal relaxation and decoherence. In order to correctly capture the physics of quantum annealing, it is essential to include the effects of dissipation in the description of the dynamics. In the weak-coupling limit, this is done by using adiabatic master equations that can be either integrated numerically as they are, or by using Monte Carlo wave function. If the coupling with the environment is stronger, other techniques are necessary.
- IN CHAPTER 5** we collect our results concerning quantum annealing of the p -spin model in dissipative environments. We discuss the standard annealing schedule and also variants such as iterated and adiabatic reverse annealing, and quantum annealing with pauses.
- IN CHAPTER 6** we turn our attention to shortcuts to adiabaticity and counterdiabatic driving to speed up adiabatic quantum computation. We revise the variational approach and propose an alternative method to build approximate shortcuts to adiabaticity based on genetic optimization.

The results of this thesis are based on the following original contributions:

- G. Acampora, V. Cataudella, P. R. Hegde, P. Lucignano, **G. Passarelli**, and A. Vitiello. *“An evolutionary strategy for finding effective quantum 2-body Hamiltonians of p -body interacting systems”*. In: Quantum Machine Intelligence 1.3 (2019), pp. 113–122 [1]. Discussed in section 3.7.
- L. M. Cangemi, **G. Passarelli**, V. Cataudella, P. Lucignano, and G. De Filippis. *“Beyond the Born-Markov approximation: Dissipative dynamics of a single qubit”*. In: Phys. Rev. B 98 (18 2018), p. 184306 [2] and L. M. Cangemi, **G. Passarelli**, V. Cataudella, P. Lucignano, and G. De Filippis. *“A Numerically Exact Method for Dissipative Dynamics of Qubits”*. In: Proceedings 12.1 (2019) [3]. Discussed in section 4.5.
- **G. Passarelli**, G. De Filippis, V. Cataudella, and P. Lucignano. *“Dissipative environment may improve the quantum annealing performances of the ferromagnetic p -spin model”*. In: Phys. Rev. A 97 (2 2018), p. 022319 [4] and **G. Passarelli**, G. De Filippis, V. Cataudella, and P. Lucignano. *“May a Dissipative Environment Be Beneficial for Quantum Annealing?”* In: Proceedings 12.1 (2019) [5]. Discussed in section 5.2.
- **G. Passarelli**, V. Cataudella, and P. Lucignano. *“Improving quantum annealing of the ferromagnetic p -spin model through pausing”*. In: Phys. Rev. B 100 (2 2019), p. 024302 [6]. Discussed in section 5.3.
- **G. Passarelli**, K.-W. Yip, D. A. Lidar, H. Nishimori, and P. Lucignano. *“Reverse quantum annealing of the p -spin model with relaxation”*. In: Phys. Rev. A 101 (2 2020), p. 022331 [7]. Discussed in section 5.4.
- **G. Passarelli**, K.-W. Yip, D. A. Lidar, H. Nishimori, and P. Lucignano. *Adiabatic reverse annealing of the p -spin model with dissipation. In preparation [8]. Discussed in section 5.5.*
- **G. Passarelli**, V. Cataudella, R. Fazio, and P. Lucignano. *“Counterdiabatic driving in the quantum annealing of the p -spin model: A variational approach”*. In: Phys. Rev. Research 2 (1 2020), p. 013283 [9]. Discussed in section 6.4.

SUMMARY OF THE MAIN RESULTS

This chapter aims to be an extended summary of the main results of this work. We will outline our motivations and findings so as to support the reader in following the rest of the manuscript. In doing so, we will postpone all technical details to later chapters.

1.1 QUANTUM ANNEALING

Quantum annealing (QA) is the physical realization of *adiabatic quantum computation* (AQC), a quantum computation paradigm polynomially equivalent to the quantum gate model, engineered to solve *quadratic unconstrained binary optimization* (QUBO) problems that classical computers would struggle to solve [10, 11]. In standard QA, a qubit system is prepared in the ground state of a simple transverse field Hamiltonian $H_{\text{TF}} = -\Gamma \sum_i \sigma_i^x$. The target is to read the ground state of a problem Hamiltonian H_p , which encodes the solution to the optimization task at hand. To do so, a time-dependent Hamiltonian is built which reads

$$H_0(s) = A(s)H_{\text{TF}} + B(s)H_p, \quad (1.1)$$

where $A(s)$ and $B(s)$ are the annealing schedules and the annealing fraction $s = s(t)$ is a function of time. In standard annealing, $s(0) = 0$ and $s(t_f) = 1$, where t_f is the annealing time. At $s = 0$, $A(0) \gg B(0) \approx 0$, while, at $s = 1$, $B(1) \gg A(1) \approx 0$, see figure 1.1. Therefore, at the beginning of the evolution the transverse field term induces quantum tunneling between computational basis states, which allows for the exploration of the Hilbert space in the search of the target solution. Near the end of the dynamics, quantum fluctuations vanish and the system is frozen in a local minimum of the target Hamiltonian, hopefully the ground state. This idea is inherited from *simulated annealing* (SA), a classical optimization strategy where the role of quantum fluctuations is played by the temperature, slowly annealed to zero [12].

In order for the algorithm to succeed, the evolution must respect the adiabatic theorem and thus be slow on a time scale of the order of $1/\Delta_{\text{min}}$, where Δ_{min} is the minimal gap between the ground state and the first excited state, otherwise diabatic *Landau-Zener* (LZ) transitions would take place and degrade the ground

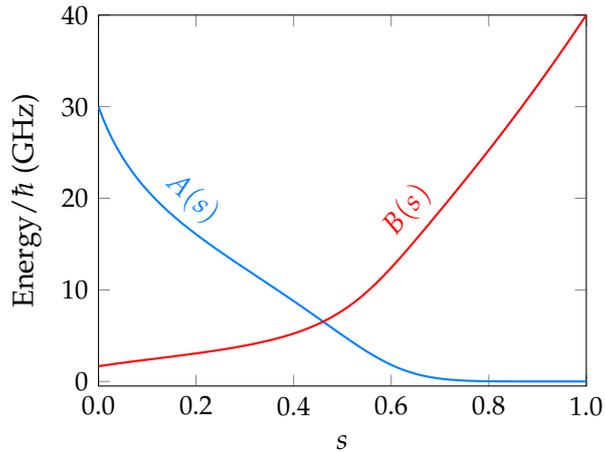


Figure 1.1 | Annealing schedules used by the D-Wave 2000Q QPU.

state (success) probability, or fidelity [13]. Computationally demanding tasks often feature exponentially small spectral gaps in the size of the system, such as when the system undergoes a first-order *quantum phase transition* (QPT), thus the *time to solution* (TTS) usually diverges exponentially as well.

A major downside of having a long annealing time comes from the fact that the quantum processor is an open quantum system and, as such, it interacts with the surrounding environment [14]. Therefore, a longer walltime leaves the qubit system more prone to errors due to thermal noise and can degrade the performance of QA. This is what happens in the D-Wave machines, physical quantum annealers based on superconducting electronics that have been extensively studied in recent years [15–17]. For this reason, an optimal working t_f can usually (but not always [18]) be identified as a compromise between adiabaticity and the necessity of limiting noise [19–21].

1.2 D-WAVE AND THE EMBEDDING PROBLEM

The D-Wave machines are built upon the Chimera graph, a sparse architecture where qubits are arranged in a bipartite lattice and coupled to at most six other qubits [22, 23]. The unit cell is depicted in figure 1.2. In the latest model of D-Wave machines, the D-Wave 2000Q, there are $M = 2048$ physical qubits but only $N_c = 6016$ couplers between them [15]. This poses a severe limitation as many relevant optimization problems are formulated in terms of fully connected Hamiltonians H_p , sometimes even featuring p -body interactions with $p > 2$, such as in satisfiability problems. Therefore, in order to use D-Wave annealers to solve such problems, the preliminary steps are i) decomposing the p -body interactions into 2-body interactions using ancillae and ii) performing minor embedding to represent logical qubits of the original problem as clusters of physical qubits on the machine. D-Wave’s pieces of software PyQUBO [24] and minorminer [25] allow one to perform both these steps with ease.

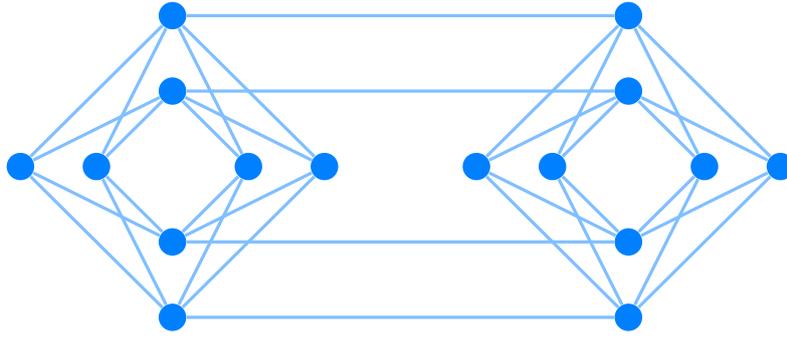


Figure 1.2 | Two unit cells of the Chimera graph.

We contributed to the field by testing the feasibility of a genetic algorithm to perform the p -to-2-body decomposition step, i. e., to find an effective Hamiltonian H'_p containing just 1- and 2-body operators featuring the same spectral properties as the original one [1]. Genetic algorithms are evolutionary strategies inspired by the Darwinian theory of the survival of the fittest. A random set of trial effective Hamiltonians, or chromosomes, undergoes three different genetic operators in succession: mutation, crossing over, and selection of the best candidates according to a problem-specific fitness value. In this case, the fitness function is the mismatch between the spectra of H_p and H'_p . After a suitable number of generations, a close-to-optimal set of trial solutions is returned.

Our results are discussed in section 3.7. The genetic algorithm works well for small system sizes, as summarized in table 3.1, and allows one to replicate eigenvalues and eigenvectors of a 3-body Hamiltonian with 2-body interactions. The main problem is that this decomposition quickly saturates the available physical qubits. As shown in figure 1.3, only small instances of fully connected models with p -body interactions can be embedded in the Chimera graph. Therefore, numerical simulations are of fundamental importance to understand the physics behind these models. In this work, we focus on one of them, the ferromagnetic p -spin model [26, 27]. Its symmetry properties allow us to simulate its dynamics on classical computers. Besides, this system undergoes a first-order QPT when $p > 2$, thus mimicking the hardness of optimization tasks [28].

1.3 FERROMAGNETIC p -SPIN MODEL

The ferromagnetic p -spin model is an Ising spin system where each qubit interacts with $p - 1$ other qubits via infinite-range interactions. In the limit of large and odd p , this model corresponds to the adiabatic search algorithm [29]. Its Hamiltonian reads

$$H_p = -N \left(\frac{1}{N} \sum_{i=1}^N \sigma_i^z \right)^p. \quad (1.2)$$

The ground state is ferromagnetic and nondegenerate (twofold degenerate) for odd (even) values of p . The annealing Hamiltonian of equation (1.1) then commutes

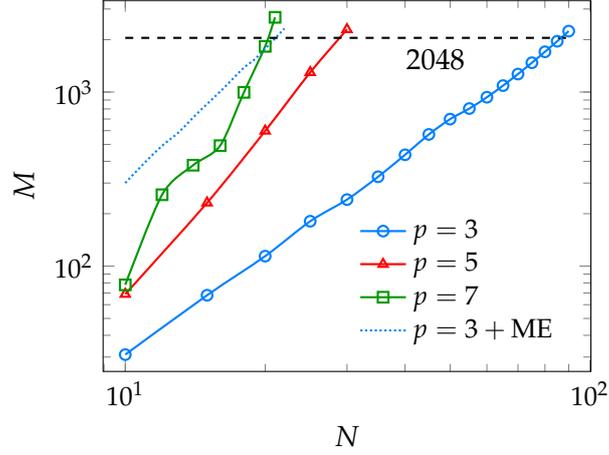


Figure 1.3 | Number of physical qubits M versus logical qubits N needed to decompose a p -body Hamiltonian into an effective 2-body model. The blue dotted line represents the number of physical qubits to perform both the p -to-2-body decomposition and the minor embedding for $p = 3$.

with the total spin operator S^2 at all times: $[H_0(s), S^2] = 0$. Both the starting and the target state belong to the subspace with maximal spin and there cannot be any transitions outside this subspace due to the Hamiltonian dynamics, hence we can always work in the symmetry subspace with maximal spin, substantially reducing the computational effort of simulating this model. This scenario can change in the presence of dissipation, except in some special cases.

The phase diagram of this model can be obtained within a mean-field theory [28]. It is easy to show that for $p = 2$ the p -spin model undergoes a continuous second-order QPT as a function of the transverse field Γ , separating the starting paramagnetic phase and the final ferromagnetic phase. At the quantum critical point, the minimal gap closes as polynomially in the system size N and quantum annealing is efficient. By contrast, for $p > 2$ the QPT is first-order; the minimum of the free energy suddenly jumps from the paramagnetic state $m = 0$ to $m \neq 0$ at a critical value of the transverse field Γ_c . The free energy of the p -spin system is shown in figure 1.4.

1.4 OPEN QUANTUM SYSTEMS

When the coupling with the environment is weak, the dynamics of the quantum annealer can be described via an *adiabatic master equation* (AME) in the Lindblad form [30, 31] for the reduced density matrix of the qubit system [32]. This form ensures the complete positivity of the density operator at all times and is derived within the Born, Markov, and rotating-wave approximation. This entails disregarding correlations, memory effects, and virtual processes that violate energy conservation. The AME reads

$$\partial_t \rho_S(t) = -i[H_0(t) + H_{LS}(t), \rho_S(t)] + \mathcal{D}[\rho_S(t)], \quad (1.3)$$

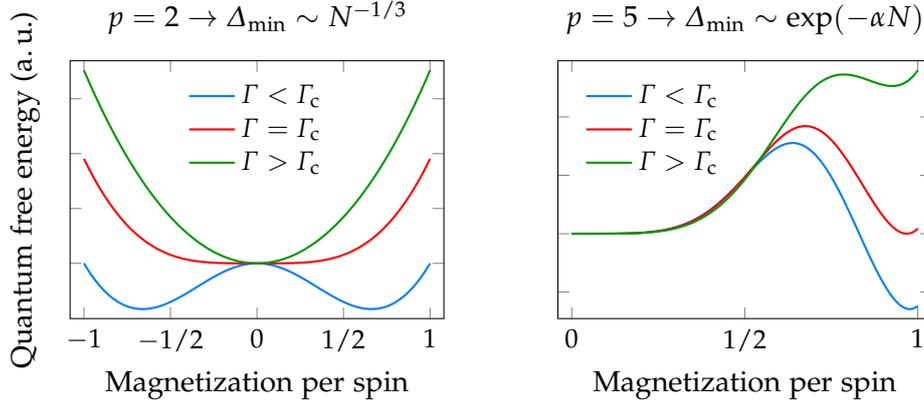


Figure 1.4 | Free energy density as a function of the magnetization m , p -spin model. For $p = 2$, there is a second-order QPT at $\Gamma = \Gamma_c$. The QPT is first-order when $p \geq 3$.

where $H_{\text{LS}}(t)$ is the Lamb shift Hamiltonian and \mathcal{D} is the adiabatic dissipator [see equation (4.13)]. These terms appear in the Lindblad equation due to the interaction with the environment, and are proportional to Lindblad operators, constituting the possible dissipative channels of the qubit system. The system-bath coupling strength η and the channel frequency affect the rate of each dissipative process.

The computational cost of simulating equation (1.3) numerically scales as D^3 , where D is the Hilbert space dimension. A more efficient approach is to recover the adiabatic master equation via the *Monte Carlo wave function* (MCWF) technique, which rewrites the master equation as a Schrödinger equation for a ket state, in which a non-Hermitian Hamiltonian incorporates the dissipative effects of the environment [33]. Many trajectories are run in parallel. In each trajectory, quantum jumps occur at random times and project the reduced system onto its instantaneous eigenstates. Taking averages over all trajectories allows for the reconstruction of the AME dynamics. The computational effort of each trajectory scales as D^2 , therefore this approach is to be preferred when a large number of computer cores are available to run multiple trajectories in parallel. Another advantage of the MCWF approach is that it outputs the spectrum of thermal processes taking place during the dynamics.

To go beyond the Born-Markov approximation, in reference [2] we discussed an alternative technique to account for relaxation and decoherence. In this approach, we combined a discretization scheme of the bath density of states, described as a finite set of harmonic oscillators, a truncation scheme of the bosonic Hilbert space of the bath, and Lanczos propagation of the system + bath ket state. The two main advantages of this technique, known as *short-iterative Lanczos* (SIL), are that we have access to the full state of the combined system and that we can include multiple-phonon processes to improve the accuracy out of the weak-coupling regime, as opposed to the AME, still maintaining a sizable computational effort. The major downside is that the maximum time allowed in simulations is the Poincaré time, determined by the minimal frequency of the discretization, after which the collection of harmonic oscillators ceases to be a good approximation of

an ergodic bath. To compensate, a large number of harmonic modes is needed, increasing the computational cost of the technique.

Short-iterative Lanczos is described extensively in section 4.5 and allows us to recover known results concerning the spin-boson model. We have also applied it to the quantum annealing of a single qubit with dephasing at $T = 0$. Our main result is that the Born-Markov approximation is no longer valid at $T = 0$ except qualitatively at very weak couplings, due to correlations with the bath playing a major role. Nonetheless, we can still ask what the predictions of the adiabatic master equation are, irrespective of its validity.

1.5 DISSIPATIVE DYNAMICS: p -SPIN MODEL

1.5.1 Quantum and simulated annealing

In reference [4], we have studied the dissipative quantum annealing of a p -spin system using the AME approach. Our results are reported in section 5.2. As indicator of the annealing performance, we refer to the residual energy, i. e., the difference between the average energy of the final state and the exact ground state energy at $t = t_f$:

$$\epsilon_{\text{res}} = \frac{\langle H_p \rangle_{t_f} - E_0(t_f)}{N}. \quad (1.4)$$

Our main conclusion is that, if the bath temperature is large enough, the unitary dynamics always yield a smaller residual energy compared to the dissipative setting, whereas, on the other hand, at low temperatures the dissipative residual energy is smaller than that of the closed system at short times if the system-bath coupling strength is not too weak. This indicates that a stronger coupling might be beneficial for quantum annealing.

In addition to that, we have also compared the scaling of the residual energy in quantum and simulated annealing for several values of the exponent p . At low temperature, SA is known to perform better than unitary QA at adiabatic times [34]. However, in the intermediate- t_f regime it is not immediate to see which technique performs best.

Indeed, QA performs better than SA when the system-bath coupling strength is not too weak. We note that the time at which simulated annealing starts to outperform its quantum counterpart is loosely proportional to the exponent p . This could suggest that, for very large values of p , quantum annealing could perform better than simulated annealing, in an accessible time window in the presence of a realistic (i. e., not extremely weak) coupling to the environment, making quantum annealing the method of choice to study the Grover's limit $p \rightarrow \infty$.

1.5.2 Pausing the annealing

Recent literature has shown that pausing the quantum annealing in the presence of dissipation can enhance the performance of the standard algorithm. Indeed, a

pause inserted some time after the minimal gap enhances thermal recombination and increases the ground state probability [35, 36].

This experimental finding is recovered also numerically within the AME approximation, as we show in our recent paper [6], discussed in section 5.3. The aim of that work is to prove that pauses can be beneficial not only for sparse Ising models that can be embedded in the Chimera graph, but also for fully-connected models such as the p -spin system. In addition, we show that even a simple Markovian model of collective dephasing is able to reproduce qualitatively the same features found in the D-Wave devices.

By testing various pause lengths t_p and pausing times s_p during the dynamics of the p -spin system, we found four different regimes resembling the experimental ones.

1. When the pause is inserted well before the time of the minimal gap s_Δ , the fidelity is not affected by pausing because the level spacing is large compared with the temperature energy scale.
2. When $s_p \approx s_\Delta$, thermal processes are more frequent and pauses can have either a positive or a negative impact on the dynamics according to the populations of the first two energy levels: when $s_p \lesssim s_\Delta$, excitations are predominant and the fidelity is decreased, whereas if $s_p \gtrsim s_\Delta$ disexcitations are favored and the fidelity is increased.
3. When the pause is inserted around an optimal pausing point, the fidelity has a peak independently of the pause length. The existence of an optimal pausing point is due to incomplete thermal relaxation and is related to some properties of the thermal relaxation rate [36].
4. When s_p is around $s = 1$, the fidelity is not affected by pausing because the level spacing is large and the dynamics are frozen.

In conclusion, pausing the annealing at the optimal pausing point can substantially enhance the success probability in the case of the p -spin system.

1.5.3 Iterated reverse annealing

Recently, *iterated reverse annealing* (IRA) has been proposed and implemented as a way to refine a trial solution of an optimization problem [37]. As opposed to standard quantum annealing, IRA starts at $s = 1$ from a candidate solution rather than from the ground state of the transverse field Hamiltonian. Then, s is decreased and quantum fluctuations are increased up to an inversion point $s = s_{\text{inv}}$, after which the annealing resumes as usual towards $s = 1$. Wise choices of the turning point and of the starting state can improve the trial solution iteratively, hence its name.

Unitary reverse annealing is not able to improve the annealing performance of the p -spin model, as proven by a recent paper [38]. Yet, experimentation on the D-Wave machines demonstrated that a substantial enhancement in success

probabilities for random spin glass instances under reverse annealing compared to standard (forward) annealing can be achieved [35]. In order to resolve this apparent contradiction, we numerically studied the dynamics of reverse annealing in the presence of Markovian dephasing for the p -spin model [7]. Our results are reported in section 5.4. We show that the associated thermal relaxation results in significant increase in the success probabilities, if s_{inv} is chosen to be close to the avoided crossing point, or before it. Pausing at the inversion point further improves performance. This result is found for two different forms of dissipation: individual dephasing, where each qubit is coupled to its own independent bath, and collective dephasing, where all qubits are coupled to the same bath. In the latter case, we additionally showed that a low-temperature environment can yield a success probability close to one even if the starting state is very far in Hamming distance to the target ferromagnetic solution, as opposed to the unitary case. The thermal effects are further enhanced by inserting a pause at the inversion point.

Simulations using the collective dephasing model have larger success probabilities for almost every s_{inv} , compared to individual dephasing. This is because, in the independent dephasing model, other states not in the subspace of maximum spin become accessible by thermal excitation or diabatic transition during the reverse anneal. Moreover, the maximum success probability achievable is always smaller for the independent dephasing model but still larger than the unitary case. Thus, our conclusion is that thermal relaxation to the ground state (possibly with a pause) improves the performance for both dephasing models compared to unitary IRA.

1.5.4 Adiabatic reverse annealing

Adiabatic reverse annealing (ARA) is a variant of QA similar in spirit to IRA. In this implementation of reverse annealing, another term is added to the Hamiltonian of equation (1.1) so as to enforce a classical starting condition and to make quantum fluctuations nonmonotonic. The aim of the technique is to use this additional degree of freedom to circumvent first-order QPTs in order to speed up AQC, with promising results so far [39]. The ARA Hamiltonian reads

$$H_0(s, \lambda) = sH_p + (1-s)\lambda H_{\text{TF}} + (1-s)(1-\lambda)H_{\text{init}}, \quad (1.5)$$

where the evolution goes from $(s, \lambda) = (0, 0)$ to $(s, \lambda) = (1, 1)$. There are several possible choices for the function $\lambda = \lambda(s)$, the simplest one being $\lambda(s) = s$. The initial Hamiltonian H_{init} enforces the starting condition and is used to prepare the qubit system in a classical configuration supposedly close in Hamming distance to the target solution. In the case of the p -spin system, a parameter $c = N_{\uparrow}/N$ is used to quantify the distance from the target state, having $c = 1$. When c is above a certain threshold, the first-order QPT for $p > 2$ becomes second-order, thus making quantum annealing exponentially more efficient in a unitary setting.

In reference [8], we studied the effects of the environment on the dynamics of adiabatic reverse annealing using the Markovian AME approach. We quantified the efficiency of ARA both in terms of the ground state probability and of the time to solution, representing the effective time it takes to solve the given problem at least

once with a probability greater than 0.99 using runs of duration t_f . Our results are thoroughly discussed in section 5.5 and can be summarized as follows:

1. Comparing unitary and dissipative dynamics for several values of c in the presence of collective dephasing, we see that the low-temperature environment can improve the success probability compared with the unitary case if the system-bath coupling strength is not too weak and c is below the aforementioned threshold. On the other hand, dissipation is detrimental when c is above the threshold.
2. Concerning the time to solution, in the case of collective dephasing we see that generally the dissipative TTS is shorter than the unitary one for short and intermediate annealing times, but is larger than the unitary TTS for longer annealing times. In addition, while in the unitary case ARA yields a shorter time to solution than standard quantum annealing, the presence of dissipation changes this scenario and the performances of the two techniques become comparable for many choices of the starting fraction c , casting doubts on the actual advantages of adiabatic reverse annealing in real-life situations.

1.6 SPEEDING UP ADIABATIC QUANTUM COMPUTATION

Quantum systems with small spectral gaps need impractically long evolution times to have a nonvanishing fidelity at the end of quantum annealing. Not only is this inconvenient for all the practical purposes where the computational time and resources might be limited, but it is also detrimental for realistic QA in the presence of dissipation, where longer annealing times leave the system more prone to thermal noise. This is the reason behind the recent interest of the AQC community in *shortcuts to adiabaticity* (STA), i. e., alternative ways to achieve adiabaticity at short times. In particular, the *counterdiabatic* (CD) driving approach to STA has gained a lot of momentum.

In this approach, the recipe is to add a counterdiabatic operator $H_{cd}(t)$ to the standard annealing Hamiltonian of equation (1.1) so as to completely suppress Landau-Zener diabatic transitions for any finite sweep rate. Demirplak and Rice [40], and later Berry [41] have derived the analytical form of this potential, which is however useless for all practical purposes as it requires the exact spectrum of $H_0(t)$, is highly nonlocal and ill-defined around quantum critical points [42, 43].

Recently, the variational formulation of CD driving, proposed by Sels and Polkovnikov [44], has allowed researchers to build approximate CD operators that can be realized experimentally. To this end, local ansätze such as our proposal, the *cyclic ansatz* (CA), or truncated series expansions of the exact CD operator, such as the *nested commutators* (NC) ansatz [45], have proven useful for many-body quantum systems. Details about these two ansätze can be found in chapter 6.

In our recent work [9], we analyzed the performance of CA and NC in the case of the p -spin model. For a very short annealing time of $t_f = 1$, we compared standard and CD-driven quantum annealing in terms of the scaling of the success

probability with the system size N for $p = 3$. We know that, in the standard case, the success probability F exponentially decreases with N as a consequence of the minimal gap closing exponentially in N . Working in the symmetry subspace of the p -spin system, we were able to study large systems up to $N = 100$ qubits. More details can be found in section 6.4. In short, we found that:

1. Using the NC series expansion, the scaling of the success probability with N remains exponential, but the fidelity is increased of several orders of magnitude compared to standard QA. The efficiency of the NC ansatz quickly saturates with the order l of the truncated series expansion, meaning that an impractically large number of terms (and, thus, variational parameters) would be required to enhance the fidelity of large systems. On the other hand, the cyclic ansatz yields an almost perfect fidelity $F \approx 1$ independently of the system size with just three variational parameters.
2. The efficiency of the cyclic ansatz remains comparable with that of the far more general nested commutators ansatz even when applied to variants of the p -spin model that give up on some of the defining characteristics of the model, such as the full connectivity and/or the infinite-range interactions. This evidence supports the fact that our cyclic ansatz can be successfully applied not only to the p -spin system, but also to more experimentally-realistic Ising models.

1.7 CONCLUSIONS

In this work, we discuss a number of results concerning unitary and dissipative quantum annealing of the ferromagnetic p -spin model. The goal is to analyze all aspects of quantum annealing for an exactly solvable model, in order to have a test bed to validate the technique and gain more insight into the underlying mechanisms. Our main results are as follows:

1. Global dephasing can be beneficial for quantum annealing at low temperature and for nonadiabatic annealing times.
2. The dissipative enhancement can be improved by pausing the system at the optimal pausing point.
3. Other variants of quantum annealing, such as IRA and ARA, can also benefit from dissipation at low temperatures and for nonadiabatic annealing times.
4. Whenever the bath is detrimental for quantum annealing, shortcuts to adiabaticity and CD driving can be used to speed up adiabatic quantum computation so that the impact of the environment is reduced.

These findings have been outlined here and will be discussed with more detail throughout this work.

QUANTUM ANNEALING

2

2.1 INTRODUCTION

Quantum annealing and adiabatic quantum computation [10, 11, 46–48] are physics-inspired heuristic methods whose aim is to solve optimization problems, NP-hard tasks that classical algorithms usually struggle to solve. Using quantum tunneling, quantum annealing can be used to efficiently sample the solution space of optimization problems. The underlying physical idea is to exploit the adiabatic theorem of quantum mechanics [13, 49, 50], that allows one to prepare a quantum system in a trivial ground state and evolve it very slowly to target the ground state of another Hamiltonian, encoding the solution to the optimization problem.

What makes quantum annealing interesting is that the largest quantum processors now on the market, the D-Wave machines, are indeed quantum annealers [16, 17]. They already allowed us to study systems beyond the current capabilities of competitor quantum chips operating in the quantum circuit paradigm. These quantum processors operate at dilution refrigerator temperatures of around 12 mK [15]. The electromagnetic environment and the thermal anchoring contribute to form a dissipative environment. The interaction with the environment tends to destroy quantum coherence, and, even though quantum annealing is found to be robust with respect to decoherence [51, 52], there is no conclusive evidence that quantum annealers are actually quantum machines and not just thermal samplers. In this work, we do not enter this dispute and rather focus on the physical description of quantum annealers, in order to understand their working principles for present and future applications.

This chapter aims to be a brief introduction to adiabatic quantum computation and quantum annealing. In section 2.2, we review the basics of these techniques. In section 2.3, we discuss the adiabatic theorem and the adiabatic time scale for quantum annealing. Section 2.4 provides a number of technical details concerning the D-Wave hardware. In section 2.5, we discuss the latest annealing protocols included in the D-Wave machines, that can sometimes help us reduce the time to solution, compared to the standard algorithm.

2.2 ADIABATIC QUANTUM COMPUTATION

Many interesting problems are expressed as unconstrained optimization tasks, where the goal is to find a vector in a variable space that minimizes a given cost function. Here, in particular, we shall focus on unconstrained combinatorial optimization, where the optimal vector belongs to a finite (even though often impractically large) set of states. Known examples of this sort include minimizing traffic in big cities [53] or finding the most efficient traveling routes for parcel deliveries [54], optimizing supply chains [55] or airline traffic [56]. In computational biology, the study of protein folding [57, 58] and molecular structures for drug discovery are optimization problems raising compelling attention in these days, but similarly interesting problems are also found in chemistry [59], computer science [29, 60], mathematics [61, 62], and physics [11, 63, 64].

There is a strong connection between *quadratic unconstrained binary optimization* (QUBO) problems, where the cost function is a quadratic function of binary variables, and Ising spin lattices with pairwise interactions in statistical mechanics [65]. In fact, the two possible values assumed by classical spins in an Ising model represent the binary variables of the given problem. By tuning the local magnetic fields and the coupling constants between the spins, the Ising lattice can be used to physically embody the given cost function. The Ising energy can then be turned into a quantum model by replacing every binary variable with a qubit described by the Pauli operator σ_z , so that a QUBO can be expressed by a quantum Hamiltonian H_p , diagonal in the computational basis. Therefore, finding the optimal vector solving the optimization problems corresponds to finding the ground state of H_p . A rather larger superclass of QUBO problems is that of *polynomial unconstrained binary optimization* (PUBO) problems, where the cost function is a polynomial of order greater than two. Any PUBO can be turned into an equivalent QUBO, using ancillae and a suitable decomposition scheme [66], as we will discuss in section 3.7.

Optimization problems are NP-hard and digital computers usually struggle to solve them exactly. However, a variety of heuristic methods have been designed to tackle specific optimization tasks in the most efficient possible way. Among global heuristics, searching in the whole configuration space, there is *simulated annealing* (SA) [12], a classical physics-based optimization technique. It exploits thermal energy to perform Monte Carlo spin-flip moves, resulting in a random walk in the configuration space. The temperature is slowly annealed to zero and the system is eventually frozen into a minimum of the cost function. Whether or not this minimum is the global minimum is not assured unless the scaling of the temperature with time is exponentially slow, which is impractical for real-life applications. If the Ising system has a spin glass phase or the temperature is varied too quickly, the technique can easily remain stuck in local minima [34].

Quantum annealing was introduced by Kadowaki and Nishimori [10] as a modification of simulated annealing, in which thermal escape rates are replaced by quantum tunneling. In SA, the temperature is used to provide the energy necessary to overcome the energy barriers separating the local minima of the potential landscape one wants to minimize. Even when the barriers are very thin,

the classical system cannot possibly cross them if it does not have sufficient energy. By contrast, a quantum mechanical system can tunnel through narrow energy barriers even when its energy is below the barrier height. This feature makes quantum annealing appealing for optimization, allowing it to drive a quantum system towards its ground state more efficiently than its classical counterpart. In figure 2.1, we show a visual comparison between the two algorithms. Quantum fluctuations are typically introduced via a transverse field Hamiltonian H_{TF} , whose general form is

$$H_{\text{TF}} = - \sum_{i=1}^N \Gamma_i \sigma_i^x, \quad (2.1)$$

where N is the number of logical qubits and Γ_i is the strength of the local magnetic field acting on the i th qubit in the x -direction. Quantum fluctuations arise from the fact that H_p and H_{TF} do not commute, $[H_p, H_{\text{TF}}] \neq 0$, hence the transverse field term mixes the eigenstates of the problem Hamiltonian. The transverse field intensity is very large at the beginning of the computation, when the whole Hilbert space has to be explored, and vanishes (i. e., is annealed) towards the end of the dynamics, when the algorithm must focus on the low-energy configurations of H_p . This is usually realized by employing the following parametric Hamiltonian,

$$H_0(s) = A(s)H_{\text{TF}} + B(s)H_p, \quad (2.2)$$

where $s \in [0, 1]$ is the annealing fraction, and the functions $A(s)$ and $B(s)$ denote the annealing schedules and satisfy $A(0) \gg B(0) \sim 0$ and $B(1) \gg A(1) \sim 0$. In standard quantum annealing, $s = t/t_f$ where t_f is the annealing time. In theoretical and numerical investigations, it is typical to resort to a linear annealing schedule, where $A(s) = 1 - s$ and $B(s) = s$, even though this form is very different from the one realized in real quantum hardware (see figure 2.3). We will largely use this form throughout this work, even though we will also show some results concerning a more experimentally-realistic annealing schedule in sections 5.3 and 5.4.

In the standard formulation of quantum annealing, the system is prepared at $t = 0$ in the ground state of $H_0(0)$. Then, the qubit state is evolved according to the Schrödinger equation with Hamiltonian (2.2) up to a certain annealing time $t = t_f$, where $s(t_f) = s_f = 1$. Eventually, the system state is measured. If the evolution has been sufficiently slow, then the qubit system will be found in the target ground state of H_p with a large probability, thus providing a trial solution to the optimization problem. How slow we must go in order to find the target with a certain confidence level is quantified by the adiabatic theorem of quantum mechanics.

2.3 ADIABATIC THEOREM

The annealing function $s = s(t)$ determines how rapidly the Hamiltonian (2.2) is varied in time. The adiabatic theorem of quantum mechanics establishes an upper bound for the time derivative of $s(t)$ needed to achieve a certain success probability in quantum annealing.

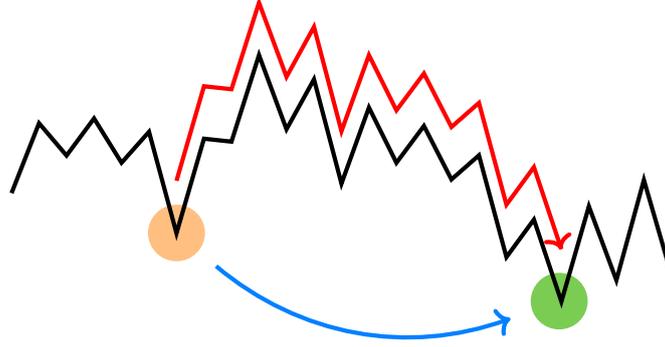


Figure 2.1 | Visual comparison between simulated and quantum annealing during the minimization of a cost function. Simulated annealing (in red) has to overcome local barriers using thermal energy, whereas quantum annealing (in blue) can tunnel through them and reach the global minimum more efficiently.

If we denote by $|\psi(t)\rangle$ the many-qubit system state subjected to the Hamiltonian $H_0(s(t)) \equiv H_0(t)$, the Schrödinger equation reads

$$i \frac{d}{dt} |\psi(t)\rangle = H_0(t) |\psi(t)\rangle, \quad (2.3)$$

with $\hbar = 1$. The Hamiltonian (2.2) can be diagonalized instantaneously, providing the time-dependent eigenbasis $\{|E_n(t)\rangle\}$. Thus, the time-dependent state reads

$$|\psi(t)\rangle = \sum_n a_n(t) e^{-i \int_0^t E_n(t') dt'} |E_n(t)\rangle, \quad (2.4)$$

which can be inserted in the Schrödinger equation to derive a system of differential equations for the coefficients $a_n(t)$. When the external driving of the Hamiltonian is slow and we are in the adiabatic regime, this system can be solved perturbatively and yields

$$a_m(t) \approx a_m(0) e^{i\gamma_m(t)} - i \sum_{n \neq m} a_n^{(0)}(t) \frac{\langle E_m(t) | \partial_t H_0(t) | E_n(t) \rangle}{\Delta_{n,m}^2(t)} e^{i\phi_{n,m}} \Big|_0^t, \quad (2.5)$$

where $a_n^{(0)}(t) = a_n(0) \exp(i\gamma_n(t))$, $\phi_{n,m}$ is a dynamical phase factor, γ_m is the geometric phase

$$\gamma_m(t) = i \int_0^t \langle E_m(t') | \partial_{t'} E_m(t') \rangle dt', \quad (2.6)$$

and $\Delta_{n,m}(t) = E_n(t) - E_m(t)$ is the instantaneous gap between $E_n(t)$ and $E_m(t)$. Requiring that the system stays in an instantaneous eigenstate for the entire dynamics translates into the usually stated adiabatic condition for the m th eigenstate

$$\max_{t, n \neq m} \left| \frac{\langle E_m(t) | \partial_t H_0(t) | E_n(t) \rangle}{\Delta_{n,m}^2(t)} \right| = \epsilon_m \ll 1. \quad (2.7)$$

For instance, for constant annealing velocity $\partial_t s(t) = 1/t_f$ and $m = 0$ (e. g., the ground state), equation (2.7) gives the following estimate for the annealing time in adiabatic quantum computation,

$$t_f \gg \max_{s, n > 0} \left| \frac{\langle E_0(s) | \partial_s H_0(s) | E_n(s) \rangle}{\Delta_{n,0}^2(s)} \right|, \quad (2.8)$$

which in many cases can be expressed in terms of just the ground state and the first excited state:

$$t_f \gg \max_s \left| \frac{\langle E_0(s) | \partial_s H_0(s) | E_1(s) \rangle}{\Delta_{\min}^2} \right|, \quad (2.9)$$

where $\Delta_{\min} = \min_s \Delta_{1,0}(s) = \Delta_{1,0}(s_{\min})$ is the minimal gap between the ground state and the first excited state. In this sense, it is often stated that the minimal gap is the bottleneck of adiabatic quantum computation [63]. The annealing time must scale as the inverse of the square of Δ_{\min} in order to solve the given optimization with a nonvanishing probability [13, 49]. However, many interesting problems feature exponentially small minimal gaps for macroscopic systems, and thus require an exponentially long annealing time to be solved. This happens particularly (but not exclusively) when a quantum system is driven across a first-order quantum phase transition, where the minimal gap at the quantum critical point closes exponentially as a function of N in the thermodynamic limit [67]. This finding is also reminiscent of the *Landau-Zener* (LZ) excitation probability $P_{\text{exp}} = \exp(-2\pi\Delta_{\min}^2 t_f)$, which is exponentially small provided that $t_f \gg \mathcal{O}(\Delta_{\min}^{-2})$ [68, 69]. The LZ model can indeed be used to describe adiabatic quantum computation in a two-level approximation.

It turns out that both adiabatic conditions, equations (2.7) and (2.9), are too strict and do not always provide a good estimate of the error at the end of quantum annealing [70]. First of all, both equations impose the adiabatic condition globally on the entire dynamics, disregarding the fact that the minimal gap varies in time. In fact, the derivative of $s(t)$ could be adjusted locally in order to accommodate for these changes, thus speeding up the dynamics far from the minimal gap and slowing it down next to it, and resulting in an overall faster quantum annealing. Indeed, this is the kind of adaptive schedule used by Roland and Cerf, providing the only possible way to achieve Grover's quadratic speedup in adiabatic search [29]. Secondly, it is important to stress that, for the purpose of quantum optimization, we are not actually interested in keeping the qubit system in its ground state for the entire dynamics, but rather we want it to be in its ground state at $t = t_f$, i. e., when we measure the state. Therefore, even if the adiabatic condition of equation (2.7) is violated locally, we can still have a favorable combination of diabatic excitations and relaxations that eventually brings the system in its ground state at $t = t_f$ with a large probability (see figure 2.2). This statement is also the foundation of modern tools denoted shortcuts to adiabaticity, which allow us to speed up AQC beyond adiabaticity [71]. We will discuss these tools in chapter 6.

In reference [70], the authors take this fact into account and derive an adiabatic condition which is more suitable for AQC, showing that the annealing time scales

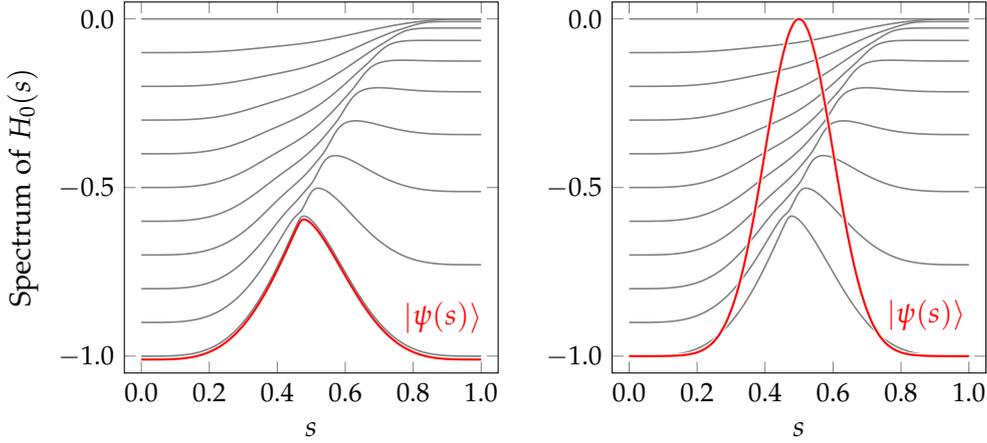


Figure 2.2 | (Left panel) Pictorial example of an adiabatic dynamics. The qubit system, prepared in the ground state of $H_0(s=0)$, evolves following exactly the instantaneous ground state when the annealing time is sufficiently long. (Right panel) The adiabatic condition can be locally violated, but quantum annealing can still manage to yield the solution to the optimization problem at $s=1$.

as Δ_{\min}^{-1} rather than Δ_{\min}^{-2} in many relevant cases. This provides a polynomial speedup compared to the naive adiabatic condition (2.7). In particular, when the annealing time is sufficiently long we can focus on just the ground state and the first excited state. When the instantaneous gap has the form

$$\Delta_{1,0}(s) = \left[(s - s_{\min})^{2a} + \Delta_{\min}^b \right]^{1/b} \quad (2.10)$$

with $b > 0$ and $a \in \mathbb{N}_+$, the inverse of its analytic continuation in the complex plane has a pole at a point \tilde{s} whose real part is close to s_{\min} and whose imaginary part satisfies $|\text{Im } \tilde{s}| \ll 1$. We can express the time derivative of $s(t)$ in terms of the gap and of an auxiliary function $h(s)$ such that the following equation holds:

$$\frac{ds}{dt} = \Delta_{1,0}(s)h(s). \quad (2.11)$$

The authors of reference [70] prove that, as long as

$$h(0) + h(1) \ll 1 \vee \text{Re} \left(i \int_0^{\text{Re } \tilde{s} + i \text{Im } \tilde{s}/2} \frac{ds}{h(s)} \right) \gg 1, \quad (2.12)$$

the upper bound for the annealing time in adiabatic quantum computation is given by

$$t_f = \int_0^1 \frac{ds}{\Delta_{1,0}(s)h(s)} \sim \frac{1}{\Delta_{1,0}(\text{Re } \tilde{s})h(\text{Re } \tilde{s})} |\text{Im } \tilde{s}| \sim \mathcal{O}(\Delta_{\min}^{-1}). \quad (2.13)$$

This is a consequence of the fact that $1/\Delta_{1,0}(s)$ is strongly peaked around $\text{Re } \tilde{s}$ and has a width of $|\text{Im } \tilde{s}|$, whereas h is roughly constant in this interval and satisfies $h \ll |\text{Im } \tilde{s}|$. This estimate of t_f is confirmed by numerical simulations [70]. If global adiabaticity is required, we will still use equations (2.7) and (2.9) as an estimate of the adiabatic time scale.

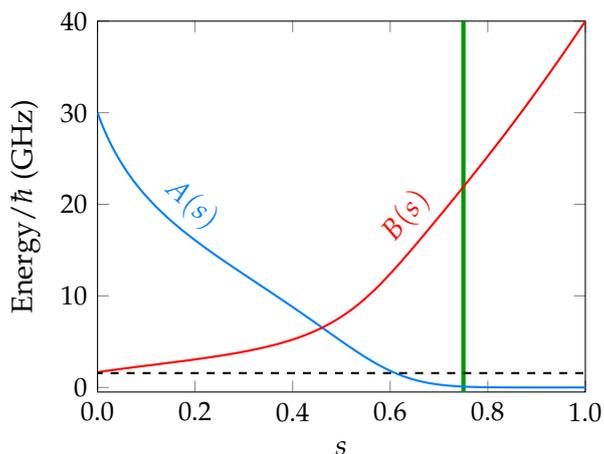


Figure 2.3 | Annealing schedules used by the D-Wave 2000Q QPU. The dashed line represents the thermal energy scale $k_B T$ with $T = 12$ mK. The vertical green line indicates the single-qubit freeze-out time.

2.4 D-WAVE

The D-Wave QPUs realize finite-temperature adiabatic quantum computation on a physical hardware, based on rf-SQUID qubits, that operate at a cryogenic temperature of approximately $T = 12$ mK. The quantum chip is essentially an open quantum system that interacts with a dilution refrigerator acting as a thermal reservoir, which affects the quantum algorithm and provides a source of thermal relaxation and decoherence [15]. Due to technological limitations, the qubits used in superconducting architectures are arranged in a sparse graph known as Chimera graph.

2.4.1 Chimera graph and embedding

The notation CK refers to a Chimera graph of $K \times K$ repeating unit cells arranged on a square two-dimensional lattice [22, 23]. The unit cell of the Chimera graph is shown in figure 2.4 and features eight qubits arranged on a bipartite lattice. The largest quantum annealers available on the market feature a C16 graph, therefore the number of physical qubits is $M = 8 \times 16^2 = 2048$. Each qubit in the Chimera graph is coupled to at most six other qubits, four within the same unit cell and two that belong to adjacent unit cells. The total number of couplers in an ideal C16 is $N_c = 6016$.

The Chimera graph is a sparse architecture. However, many interesting QUBO problems are formulated in terms of fully-connected Ising Hamiltonians, sometimes even featuring k -body interactions with $k > 2$ as in satisfiability problems [65, 72]. This makes it cumbersome to solve these problems using the D-Wave machines. The problem one aims to solve is formulated on a graph $G = (V_G, E_G)$, where V_G are vertices (qubits) and E_G are bonds (couplers). The standard procedure is to map G onto a physical quantum graph $U = (V_U, E_U)$, which in this case is equal to

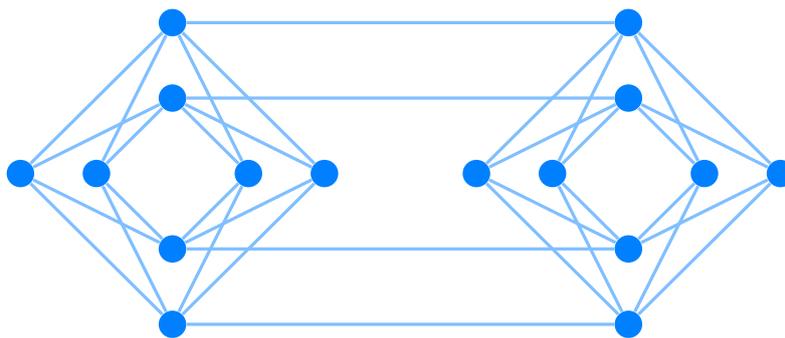


Figure 2.4 | Two unit cells of the Chimera graph. The eight qubits of each cell are arranged in a bipartite lattice.

the Chimera graph C16. This mapping involves two separate components: minor embedding [23] and parameter setting [22]. The goal of minor embedding is to find a subgraph G_{emb} of U such that the logical graph G can be obtained from G_{emb} by contracting edges. In graph theory, we say that G is a minor of U , hence the name. Physically, this means that each logical qubit of the original problem is effectively represented by subtree of the hardware graph, where physical qubits within the subtree are strongly ferromagnetically coupled and behave as a single logical qubit. The parameter setting problem, on the other hand, is to determine the parameters of the embedded model that allow to replicate experimentally the original Hamiltonian, also taking into account the limitations of the analog QPU used for computation. For example, there might be limits on the maximum allowed coupling J between qubits. Thus, in order to have physical qubits within the same subtree act as a single qubit, all other parameters in the Hamiltonian should be rescaled accordingly so that the intra-subtree ferromagnetic coupling is the largest energy scale. This rescaling, though necessary, inevitably impacts the performance of the adiabatic algorithm: firstly, we need to use ancillary qubits to build these trees and represent a smaller number of logical qubits, thus hindering the capability of the QPU of solving large optimization problems; secondly, a smaller energy scale means smaller spectral gaps, which in turn imply a longer run time. A side effect of this fact is that the freeze-out time, i. e., the time after which a qubit state is fixed due to the fact that quantum fluctuations and thermal energy are too small compared with the level spacing, shifts towards $s = 0$ for multi-qubit clusters [15]. The single-qubit freeze-out time is indicated in figure 2.3 by a green vertical line.

More formally, a minor embedding of a graph $G = (V_G, E_G)$ onto a fixed hardware graph $U = (V_U, E_U)$ is an application $\phi: G \rightarrow U$ such that each vertex in V_G is mapped to a connected subtree T_i of U , and such that there is a map $\tau: V_G \times V_G \rightarrow V_U$ such that, for each $(i, j) \in E_G$, there exist corresponding vertices $i_{\tau(i,j)} \in V_{T_i}$ and $j_{\tau(j,i)} \in V_{T_j}$ with $(i_{\tau(i,j)}, j_{\tau(j,i)}) \in E_U$. An example of minor embedding is shown in figure 2.5. Here, we want to embed a logical triangular cell into a square lattice. Using minor embedding, two vertices of the square lattice are strongly bound together and act as a single vertex of the original triangular lattice.

The problem of finding a minor embedding when no assumptions are made

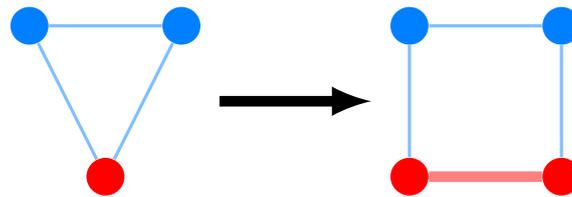


Figure 2.5 | Minor embedding of a triangular graph into a square lattice.

concerning both G and U is NP-hard [73]. However, when the target graph is known (in particular, is the Chimera graph), there exist polynomial-time algorithms that can perform exact minor embedding for graphs with a small number of vertices, up to ~ 10 . The most used codes to perform minor embedding, such as `minorminer` [25], are instead heuristics that trade off accuracy in exchange for a more favorable scaling of the runtime with the graph size. This allows us to perform minor embedding on the C16 graph with a reasonable effort.

2.5 ADVANCED ANNEALING SCHEDULES

The standard schedules commonly used by superconducting quantum annealers are shown in figure 2.3. The annealing fraction $s = s(t)$ is normally defined as $s(t) = t/t_f$. However, it is also possible to study variants of quantum annealing where the annealing schedule is changed so as to improve the performances of the algorithm. In particular, recent theoretical and experimental works have shown that inhomogeneous transverse driving [74, 75], mid-annealing quenches, pauses [6, 35, 76], and reverse quantum annealing [7, 37–39] can be beneficial to improve the success probability of quantum annealing or can be used to gain more insight into the physics of dissipative quantum annealing.

2.5.1 Pauses

Quantum annealing can be temporarily stopped at a certain fraction s_p , and the qubit system can evolve for a certain amount of time t_p with a time-independent Hamiltonian. Experimental and theoretical works have proven that inserting pauses at favorable points during the dynamics can improve the success probability of quantum annealing, under specific conditions regarding the ratio between the thermal relaxation rate and the pause duration [6, 35, 76]. In general, increasing the pause length increases the success probability. In fact, pausing a quantum annealer allows it to relax towards its ground state, and mitigates the effects of diabatic excitations. In order for it to be effective, the pause must be inserted when the level spacing is small enough so that tunneling between eigenstates is not exponentially suppressed, and after the minimal gap, so that a macroscopic fraction of the system is in an excited state and has the chance to relax back by a phonon emission. Moreover, the qubit system must be far from its thermal equilibrium state. Pauses inserted before the minimal gap have generally no effects on the dynamics.

Due to this delicate balance between time and energy scales, pausing affects open quantum systems in a nontrivial way, and only recently we have come to a more in depth understanding of it. In fact, in a recent paper [36], the authors have proven the following necessary conditions for the existence of an optimal pausing point, i. e., a point s_p for which the ground state probability shows a maximum:

1. The time-dependent thermal relaxation rate must be monotonically decreasing after the minimal gap.
2. The relaxation rate must be large right after the minimal gap and small at the end of the annealing (compared to the other energy scales).
3. The ground state population at the end of the annealing must be subthermal, i. e., $p_{\text{gs}}(s = 1) \leq 1/(1 + \exp[-\beta\Delta_{1,0}(s = 1)])$.

This theorem allows to explain all previous numerical and experimental findings concerning pauses in dissipative quantum annealing.

Pauses also affect isolated quantum systems, but the underlying mechanism is far easier to understand and can be related to coherent oscillations and qubit rotations. In particular, an isolated quantum system benefits from pauses inserted specifically around the minimal gap, whereas the effects of pauses are negligible elsewhere. In addition, the beneficial effect of pausing around the minimal gap decreases when the annealing time grows and the dynamics are adiabatic. For these reasons, we can generally disregard pausing effects on unitary dynamics.

Consider for example a single two-level system evolving with the Landau-Zener-like Hamiltonian

$$\frac{H(s)}{\hbar\omega} = -\frac{1-s}{\sqrt{2}}\sigma_x - \frac{s}{\sqrt{2}}\sigma_z, \quad (2.14)$$

where ω fixes the energy and time scales. The minimal gap $\Delta_{\min} = 1$ occurs at $s_{\min} = 1/2$. The evolution starts from the $+1$ eigenstate of σ_x and the target state is $|0\rangle$.

First of all, if the evolution is adiabatic ($t_f \gg 1$), the state vector $|\psi\rangle$ will move along the arc $\phi = 0$, $\theta \in [0, \pi/2]$ on the Bloch sphere [with $\theta(0) = \pi/2$ and $\theta(1) = 0$]. A pause inserted at any point s_p will lock the Hamiltonian to $H(s_p)$. At $s = s_p$, the adiabatically-evolved state would be the $+1$ eigenstate of $H(s_p)$. Therefore, during the pause, the system state will trivially acquire a global phase factor $\Delta\alpha$ proportional to the pause length t_p , which however will not affect the dynamics later on after the pause. This is represented in figure 2.6(a), where the green circle represents the time evolution operator, i. e., a rotation operator around the axis $\mathbf{k} = \mathbf{v}$, with \mathbf{v} representing the system state on the Bloch sphere.

Conversely, if the evolution is nonadiabatic, the effect of the pause will depend on s_p (and, of course, on t_p). If s_p is close to $s = 0$, the system state would still be very close to the adiabatic eigenstate of $H(s_p)$ due to the fact that the gap is large at the beginning of the dynamics and that the time evolution operator is continuous. Therefore, the effect of the pause would be once again to produce

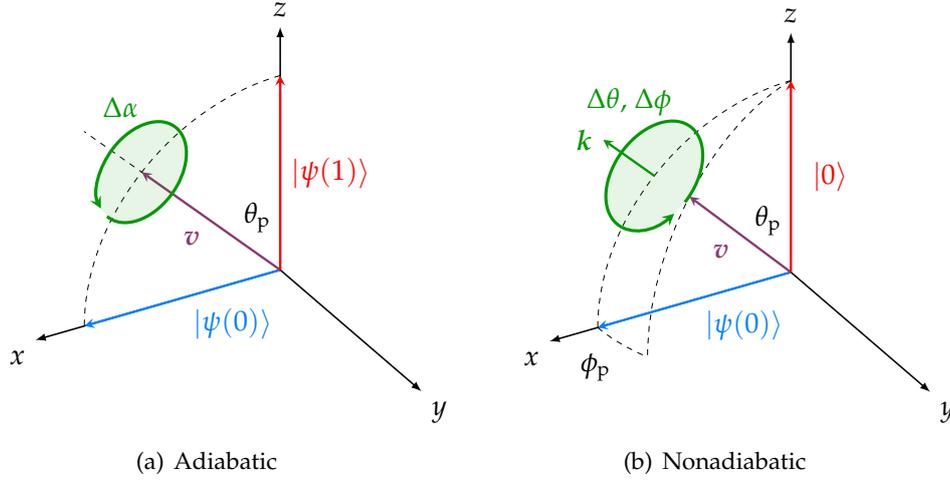


Figure 2.6 | Effect of pausing the unitary dynamics of the single qubit (2.14) at the minimal gap. In panel 2.6(a), the dynamics is adiabatic ($|\psi(s)\rangle \approx |E_0(s)\rangle$) and the pause only adds an irrelevant global phase factor. By contrast, if the evolution is nonadiabatic as in panel 2.6(b), the instantaneous state is not an eigenstate of the time evolution operator at $s = s_p$, and thus rotates around the k -axis during the pause.

a global t_p -dependent phase factor that will not affect the dynamics afterwards. On the other hand, if s_p is close to $s = 1$, the system state before the pause is no longer an adiabatic eigenstate of $H(s_p)$ but a vector on the Bloch sphere that forms an angle θ_p with respect to the z -axis. The Hamiltonian is approximately $H(s_p) \approx -\sigma_z/\sqrt{2}$ and the resulting time evolution operator is the rotation operator around the z -axis. The rotation angle is $\phi(t) = \phi_p + \sqrt{2}t$, where ϕ_p is the state phase at $s = s_p$. This operator acts on the Bloch state by changing its azimuthal angle ϕ but leaving its polar angle θ_p unchanged. However, since the success probability in this case is given by $|\langle 0|\psi(1)\rangle|^2 = \cos^2 \theta \approx \cos^2 \theta_p$, the pause does not affect the outcome of the algorithm.

The only case in which a pause can affect the success probability is when $s_p \approx s_{\min}$. Before the pause, the system state ceases to follow the instantaneous eigenstate of $H(s)$ when the inverse of the instantaneous gap becomes too small, therefore the system state at $s = s_p$ is described by the two angles (ϕ_p, θ_p) , with $\phi_p \neq 0$ in general. The time evolution operator during the pause is the rotation operator around the axis k with angles $\phi = 0$ and $\theta = \pi/4$. The angle that the system state forms with this axis evolves as $\alpha(t) = \alpha_0 + \sqrt{2}t$. This rotation affects both θ_p and ϕ_p in a complicated way that can be worked out using Rodrigues' rotation formula [77]

$$\mathbf{v}_{\text{rot}}(t) = \cos \alpha(t) \mathbf{v} + \sin \alpha(t) \mathbf{k} \times \mathbf{v} + (\mathbf{k} \cdot \mathbf{v}) [1 - \cos \alpha(t)] \mathbf{k}. \quad (2.15)$$

This is represented in figure 2.6(b). Of course, the overall changes $\Delta\theta$ and $\Delta\phi$ periodically depend on t_p (modulus $\sqrt{2}\pi$). Thus, the system state will be generally modified by the pause, which can either be detrimental or beneficial for the success probability according to t_p . For this reason, and unlike the dissipative case, the

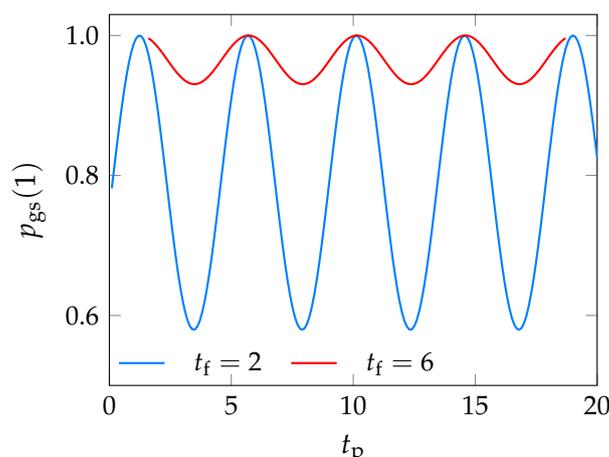


Figure 2.7 | Success probability $p_{gs}(1) = |\langle 0|\psi(1)\rangle|^2$ as a function of the pause length t_p , single qubit, for $s_p = 1/2$. The ground state fidelity oscillates with a period of $\sqrt{2}\pi$. The effect of the pause becomes less evident when the dynamics are more adiabatic, see the main text for further details. The total annealing time is $t'_f = t_f + t_p$.

success probability depends nonmonotonically on t_p in the unitary case, and acquires a periodicity of $\sqrt{2}\pi$, as evident from figure 2.7.

2.5.2 Iterated reverse annealing

In the standard annealing protocol (with or without pauses), $s(t)$ is a nondecreasing function of t that satisfies the boundary conditions $s(0) = 0$ and $s(t_f) = 1$, thus quantum fluctuations monotonically decrease. As for other heuristic methods, in quantum annealing there is no control, *a priori*, of the accuracy of the output vector, which is usually a suboptimal approximation of the correct solution. This limitation can be partially softened by combining multiple heuristics together as steps of a multistage optimization routine, where the output of each step is used as input of the following one. In this way, we can often produce better trial solutions by successive approximation. Reverse quantum annealing (or reverse annealing for short) is a viable tool to perform this kind of multistage optimization.

In reverse annealing, the qubit system evolves according to a different annealing path that starts and ends at $s(0) = s(t_f) = 1$. At $t = 0$, the qubit system is initialized at first in a certain classical state, for example the trial solution coming from another optimization routine. This state is supposedly close to the target ground state of the problem Hamiltonian. The algorithm occurs in two phases. In the first phase, quantum fluctuations are increased by decreasing the annealing fraction from $s(0) = 1$ to $s(t_{inv}) = s_{inv}$, where t_{inv} (s_{inv}) is the inversion time (point). At the inversion point $s = s_{inv}$, a pause can possibly be inserted. In the second phase, the annealing resumes as usual, quantum fluctuations are decreased and the system is measured at $s(t_f) = 1$ [7]. In figure 2.8, we show a comparison between the annealing paths in standard and reverse quantum annealing. In the standard and reverse protocol, $t_f = 500 \mu\text{s}$. In the paused protocol, we have inserted a pause of

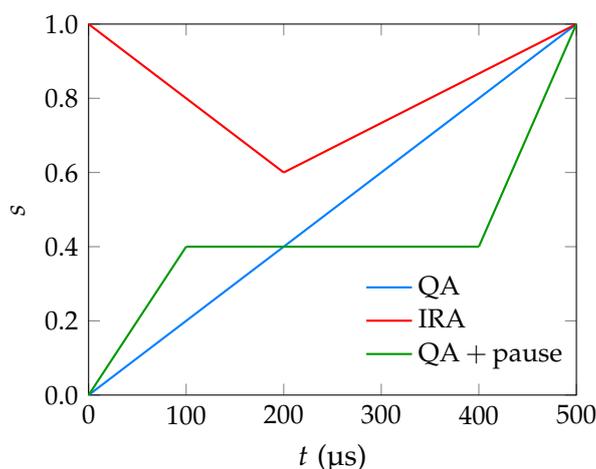


Figure 2.8 | Annealing fraction versus time, in standard quantum annealing (with and without a pause) and in reverse annealing.

length $t_p = 300 \mu\text{s}$ at $s_p = 2/5$, during an annealing of time $t_f = 200 \mu\text{s}$, so that the total evolution time is $t'_f = 500 \mu\text{s}$. The inversion time for reverse annealing is $t_{\text{inv}} = 200 \mu\text{s}$ and the inversion point is $s_{\text{inv}} = 3/5$.

If the inversion point and the starting state are chosen conveniently, the output is an improved trial solution, i. e., a quantum state having a larger superposition with the target state. This state can either be kept as it is or fed again into reverse annealing as a new starting state. In this way, we can systematically improve a trial solution, if at each step the overlap with the target state increases. This is the reason why this kind of reverse annealing is colloquially known as *iterated reverse annealing* (IRA). Reverse annealing can also be employed by itself to study topological systems [78], and for measuring tunneling and relaxation rates [15].

Similarly to pausing, reverse annealing generally fails to improve the success probability of adiabatic quantum computation in a closed system setting. Suppose that the starting state is an (excited) eigenstate of the problem Hamiltonian. Then, if the dynamics are adiabatic, according to the adiabatic theorem the system will remain in the corresponding instantaneous eigenstate for the entire dynamics and reverse annealing will have no effect.

By contrast, if (i) the starting state is close to the target, (ii) the dynamics are nonadiabatic, and (iii) the inversion time is chosen to be $s_{\text{inv}} < s_{\text{min}}$, then reverse annealing can possibly change the state and improve the success probability, although these requirements severely limit its range of application. In fact, in a closed system setting, the only mechanism that can allow the qubit system to move across its energy eigenstates is LZ transitions. Thus, in order for reverse annealing to be effective we have to cross the minimal gap, i. e., $s_{\text{inv}} < s_{\text{min}}$. Unfortunately, however, after t_{inv} the system crosses its minimal gap once again, thus part of the ground state probability recovered after the first crossing is lost after the second passage across the gap. A recent theoretical paper indeed proved that IRA fails to improve trial solutions of a fully-connected model in a closed system settings for

these reasons [38].

Once again, this scenario changes drastically in the presence of dissipation. When the quantum chip is coupled to an environment at low temperature, thermal processes can induce jumps between energy eigenstates and affect the outcome of reverse annealing. This is the reason why there are experimental and theoretical papers claiming that iterated reverse annealing can be useful for quantum annealing [7, 79]. We will discuss one example in section 5.4.

2.5.3 Adiabatic reverse annealing

Finally, *adiabatic reverse annealing* (ARA) is another variant of reverse annealing that has been recently proposed as an alternative to iterated reverse annealing [38, 39]. This technique shares many similarities with IRA:

1. At $t = 0$ the qubit system is prepared in a trial classical state.
2. The qubits are evolved following a path where quantum fluctuations are nonmonotonic.
3. The effectiveness of the technique is directly related to the overlap between the starting state and the target state, the larger the better.

However, as opposed to IRA, the annealing fraction is monotonic in ARA and goes from $s(0) = 0$ to $s(t_f) = 1$ as in standard quantum annealing. At the same time, the original Hamiltonian (2.2) has to be modified in order to include an additional term used to enforce the starting condition without changing the boundary conditions of $s(t)$.

The proposed ARA Hamiltonian can be written as

$$H_0(s, \lambda) = B(s)H_p + A(s)[\lambda H_{\text{TF}} + (1 - \lambda)H_{\text{init}}], \quad (2.16)$$

where λ is an additional time-dependent parameter that ranges in $[0, 1]$. Note that we can recover the standard quantum annealing Hamiltonian (2.2) if we set $\lambda = 1$. Now we adopt a linear annealing schedule. In adiabatic reverse annealing, we start from $(s, \lambda) = (0, 0)$, where the Hamiltonian is $H_0(0, 0) = H_{\text{init}}$. The initial Hamiltonian is diagonal in the computational basis and can be written as

$$H_{\text{init}} = \sum_{i=1}^N \epsilon_i \sigma_i^z. \quad (2.17)$$

The parameters $\epsilon_i = \pm 1$ are used to set the starting state for the annealing, i. e., the ground state of H_{init} . This state is evolved adiabatically following a path (s, λ) in the parameter space that ends at the point $(1, 1)$, where the Hamiltonian is $H_0(1, 1) = H_p$.

The path in the parameter space is parametrized by the function $\lambda(s)$. For example, consider $\lambda(s) = s^\alpha$, with $\alpha > 0^a$. If $0 < \alpha \leq 1$, the coefficient of the

^a $\alpha = 0$ gives standard quantum annealing. Instead, $\alpha < 0$ is excluded because we must have $\lambda(s) \in [0, 1]$ for all $s \in [0, 1]$.

transverse field $s^\alpha(1-s)$ has a maximum at $s^* = \alpha/(1+\alpha)$, hence quantum fluctuations increase until $s = s^*$ and decrease thereafter. For $\alpha > 1$, there is also a minimum at $\bar{s} = 0$, and the coefficient of the transverse field remains nonmonotonic.

The choice of the annealing path (for example, the choice of α) affects the performance of ARA. This additional freedom compared with standard quantum annealing can allow us to build paths in the parameter space that avoid first-order quantum critical points, therefore exponentially speeding up adiabatic quantum computation. More details about this technique will be given in sections 3.3 and 5.5 and appendix A.2.

2.6 CONCLUSIONS

In this chapter, we have introduced quantum annealing and its variants. Quantum annealing is used to perform optimization and solve NP-hard tasks that classical computers have troubles solving. The performance of the quantum algorithm is determined by the minimal gap between the ground state and the first excited state encountered during the annealing. However, in many cases, we are not able to diagonalize the quantum Hamiltonian corresponding to a given optimization problem, hence it can be tricky (sometimes impossible) to estimate its minimal gap and know a priori whether or not quantum annealing will be efficient in solving it. Due to the fact that we do not know the target state in advance, it is also difficult to understand if the algorithm converged to the correct solution. For this reason, it is crucial to test quantum annealing by solving quantum models whose ground states are exactly known. In this way, we can test the theoretical and numerical predictions and gain more insight into the physics of quantum annealers.

There are not many interesting quantum models that are both exactly solvable and hard enough to pose a realistic challenge for quantum annealing. In the next chapter, we are going to introduce the ferromagnetic p -spin model, a quantum system that meets both these requirements. The numerical part of this work will be entirely based on this model, therefore we will introduce it and discuss its main features to set the foundations for our subsequent investigation.

FERROMAGNETIC P-SPIN MODEL

3.1 INTRODUCTION

In this work, we shall focus on the ferromagnetic p -spin model, featuring infinite-range all-to-all p -body interactions [26, 27]. In the large N and p limit, this model becomes the adiabatic version of Grover’s search problem [29, 60]. The ferromagnetic p -spin model is simple enough to simulate its dynamics using classical hardware, and even evaluate its exact ground state in the thermodynamic limit using a mean field theory [28]. Yet, on the other hand, finding this ground state by means of an adiabatic evolution poses some serious threats that closely resemble the complexity of solving NP-hard optimization problems. The fact that we have access to a reliable analytical solution means that we can use this model as a benchmark of the performance of quantum annealing. This is the reason why it is so popular in the adiabatic quantum computation community [1, 4–7, 28, 34, 38, 39, 63, 74, 80–83].

The simplicity of numerically simulating this system comes from its Hamiltonian being permutationally invariant: any permutation of qubit variables leaves all properties of the system unchanged. Physically, this corresponds to the fact that the Hamiltonian commutes with the square of the total angular momentum operator S^2 , hence the total Hilbert space is decomposed as a direct sum of eigenspaces of S^2 . This symmetry allows us to simulate the unitary adiabatic evolution of this system with resources scaling as $\text{poly}(N)$ rather than exponentially in N [84]. However, when dissipation is present, this symmetry can be broken due to the environment causing transitions between otherwise uncoupled symmetry eigenspaces. For instance, this happens when each qubit is coupled to its own *independent bath* (IB). There are cases, however, where the characteristic wavelength of the interaction is much larger than the system size, so that all qubits are collectively coupled to the same *collective bath* (CB). Numerically, this is a very favorable case that preserves the permutational symmetry and allows us to study the dissipative dynamics of the p -spin model for large N . Later in this work, we will show our results concerning both kinds of dissipation.

This chapter aims to be a collection of known results concerning the p -spin model, which will help us discuss our personal contributions to the field in the next chapters. It is mainly based on Refs. [1, 28, 38, 39], where the interested reader

will find additional information and more thorough discussion. This chapter is organized as follows. In section 3.2, we present the p -spin Hamiltonian. In section 3.4, we discuss the decomposition of the Hilbert space in spin sectors and motivate our choice of focusing on the maximal spin subspace exclusively. In section 3.5, we discuss the spectral properties of the p -spin model and the scaling of the minimal gap in the thermodynamic limit. In addition, we recover the adiabatic Grover's search as the $p \rightarrow \infty$ limit in section 3.6. Finally, in section 3.7, we discuss the two interlaced problems of transforming the p -spin model into a QUBO and its embedding in the Chimera graph of the D-Wave chip. More details regarding the mean field analysis of the p -spin model can be found in appendix A.

3.2 HAMILTONIAN

We consider a system of N qubits and define the transverse and local magnetization operators in terms of single-qubit Pauli operators as

$$m^{x,z} = \frac{1}{N} \sum_{i=1}^N \sigma_i^{x,z} = \frac{2}{N} S^{x,z}. \quad (3.1)$$

The Hamiltonian of the infinite-range ferromagnetic p -spin model reads

$$H_p = -N(m^z)^p = -\frac{1}{N^{p-1}} \left(\sum_{i=1}^N \sigma_i^z \right)^p = -\frac{1}{N^{p-1}} \sum_{i_1, \dots, i_p} \sigma_{i_1}^z \cdots \sigma_{i_p}^z, \quad (3.2)$$

with $p \geq 1$. For odd values of p , its ground state is the product state with all qubits in $|0\rangle$. For even values of p , the ground state is twofold degenerate as the system is Z_2 -invariant. We pair Hamiltonian (3.2) with a homogeneous transverse field Hamiltonian [see equation (2.1) with $\Gamma_i = \Gamma$ for all i], which reads

$$H_{\text{TF}} = -N\Gamma m^x. \quad (3.3)$$

In this section, we consider a linear annealing schedule $A(s) = (1-s)$ and $B(s) = s$ [see equation (2.2)], hence the quantum annealing Hamiltonian of this model reads

$$H_0(s) = sH_p + (1-s)H_{\text{TF}} = -Ns(m^z)^p - N\Gamma(1-s)m^x. \quad (3.4)$$

This Hamiltonian only features collective spin operators, hence it commutes with the square of the total spin operator: $[H_0(s), S^2] = 0$. The case $p = 1$ represents a single $(N+1)$ -level system and is trivial, thus we will not discuss it here. The case $p = 2$ is equivalent to the anisotropic *Lipkin-Meshkov-Glick* (LMG) model with no interactions along the y direction [85–87]. We will show later that this system undergoes a second-order QPT when the parameter s is swept from $s = 0$ to $s = 1$, separating the starting disordered phase (where $\langle m^z \rangle = 0$) and the final magnetic phase (where $\langle m^z \rangle \neq 0$). The transition occurs at a quantum critical point s_c . For finite sizes, the minimal gap Δ_{\min} closes as $1/\text{poly}(N)$ in the thermodynamic limit [28, 88, 89], hence, according to the adiabatic criterion [equation (2.7)], the

annealing time for adiabaticity scales as $\text{poly}(N)$ and quantum annealing is efficient in finding the ground state [90].

Instead, for $p > 2$ the p -spin model is subjected to a first-order QPT in the thermodynamic limit. For finite-size systems, the signature of the first-order QPT is usually the fact that Δ_{\min} scales as $\exp(-\alpha N)$, even though there are tailored examples in which the gap closes nonexponentially at a first-order phase transition [91, 92]. Usually, however, the annealing time grows exponentially in N , thus finding the ferromagnetic ground state by means of an adiabatic evolution is inefficient [93], and mimics the complexity of NP-hard problems, where the computational time usually scales exponentially with the input size.

The appearance of this first-order quantum phase transition for $p > 2$ can be avoided by adding nonstoquastic catalyst terms (such as XX interactions) [81, 94] to the annealing Hamiltonian of equation (3.4), or by considering inhomogeneous transverse fields [74]. The effect of nonstoquastic Hamiltonians has been demonstrated in the thermodynamic limit, but there is still an open debate concerning their actual effectiveness on finite-size systems, where, in many cases, nonstoquastic Hamiltonians actually reduce the success probability of quantum annealing, as proven by a recent paper [95].

An inhomogeneous transverse field allows each qubit to cross the quantum critical point s_c at different times. Hence, the quantum phase transition, a collective phenomenon, is bound not to occur [67]. However, such a technique inevitably destroys the permutational symmetry of the Hamiltonian (3.4), making the numerical analysis of the system exponentially more demanding. For these reasons, we will not discuss these two techniques in this work.

3.3 ADIABATIC REVERSE ANNEALING

As mentioned in section 2.5.3, adiabatic reverse annealing [38, 39] is a type of quantum annealing where the system starts from a given classical configuration rather than from the ground state of the transverse field Hamiltonian, i. e., the uniform quantum superposition of all computational basis states. The starting state can be a trial solution of the optimization problem, such as the output of another annealing routine, and is supposedly close to the target solution according to some metrics [37, 96]. At first, quantum fluctuations are increased, and then decreased again. This algorithm can improve the trial solution, if the starting state is already sufficiently close to the target. The meaning of this statement will be clarified in the following.

In section 3.2, we introduced the p -spin Hamiltonian H_p [equation (3.2)] and the corresponding transverse field [equation (3.3)]. In the case of ARA, the annealing Hamiltonian is defined as

$$H_0(s, \lambda) = sH_p + (1-s)\lambda H_{\text{TF}} + (1-s)(1-\lambda)H_{\text{init}}, \quad (3.5)$$

where $s, \lambda \in [0, 1]$. Standard quantum annealing is recovered when $\lambda = 1$, in which case the additional term drops out and we are left with Hamiltonian (3.4). In ARA,

on the other hand, the quantum evolution starts from $(s, \lambda) = (0, 0)$ at $t = 0$ and goes to $(s, \lambda) = (1, 1)$ at $t = t_f$. Thus, the Hamiltonian at $t = 0$ is H_{init} , which reads

$$H_{\text{init}} = - \sum_{i=1}^N \epsilon_i \sigma_i^z, \quad (3.6)$$

where $\epsilon_i = \pm 1$ is used to set the initial condition. In fact, the starting state is the ground state of H_{init} and, in the computational basis, is equal to $|\psi_0\rangle = |(1 - \epsilon_1)/2, (1 - \epsilon_2)/2, \dots, (1 - \epsilon_N)/2\rangle$, which can be chosen to be close to the target ferromagnetic solution $|\psi_{\text{fm}}\rangle = |0, 0, \dots, 0\rangle$ (for odd p). How close the starting state is to $|\psi_{\text{fm}}\rangle$ is expressed by the fraction of spins that are in state $|0\rangle$ and thus are already ferromagnetically aligned:

$$c = \frac{1}{N} \sum_{i=1}^N \delta_{\epsilon_i, 1}. \quad (3.7)$$

The ferromagnetic state has $c = 1$. Therefore, the larger is c in the starting state, the closer the starting state is to the ferromagnetic solution at $s = \lambda = 1$. The fraction c is given by the mean value over the starting state of the operator

$$C = \frac{1}{N} \sum_{i=1}^N \frac{\sigma_i^z + 1}{2} = \frac{1}{2}(m^z + 1). \quad (3.8)$$

Due to the presence of H_{init} , the Hamiltonian of equation (3.5) no longer commutes with the total angular momentum: $[H_0(s, \lambda), S^2] \neq 0$. However, the qubit system is partitioned into two groups according to $\epsilon = \pm 1$ in Hamiltonian (3.6), thus we can build total spin operators related to each of the two groups:

$$m_1^{x,z} = \frac{1}{N} \sum_{i=1}^N \sigma_i^{x,z} = \frac{2}{N} S_1^{x,z}, \quad m_2^{x,z} = \frac{1}{N} \sum_{i=1}^N \sigma_i^{x,z} = \frac{2}{N} S_2^{x,z}. \quad (3.9)$$

It is easy to show that the Hamiltonian of equation (3.5) commutes with the set of commuting operators $\{S_1^2, S_1^z, S_2^2, S_2^z\}$. Hence, each term in the Hamiltonian (3.5) can be expressed as

$$H_p = -N(m_1^z + m_2^z)^p, \quad (3.10)$$

$$H_{\text{TF}} = -N(m_1^x + m_2^x), \quad (3.11)$$

$$H_{\text{init}} = -N(m_1^z - m_2^z); \quad (3.12)$$

we can also rewrite equation (3.8) as $C = m_1^z$. In appendix A, we show that when c is above a certain threshold (i. e., the starting state is already partially magnetized), adiabatic reverse annealing can allow us to avoid the first-order quantum phase transition for $p > 2$.

3.4 SPIN SECTORS

The Hamiltonian of equation (3.4) is invariant with respect to any permutations of spin indices, due to its highly-symmetric structure. Given any permutation operator $P_{\zeta_1, \dots, \zeta_N}$ which constitutes a bijection $\{1, \dots, N\} \leftrightarrow \{\zeta_1, \dots, \zeta_N\}$, we have that $[H_0(s), P_{\zeta_1, \dots, \zeta_N}] = 0$. Accordingly, the Hilbert space of the system can be decomposed as a direct sum of vector subspaces, identified by their symmetry properties. After fixing a basis in each subspace, using the theory of representations and the Young diagrams, the Hamiltonian will then have a block-diagonal structure, which provides an important simplification in the numerical analysis if, during the dynamics, no external operators break this symmetry.

From a more practical viewpoint, we can simplify the construction of these subspaces by noting that the permutational symmetry of the p -spin Hamiltonian is equivalent to rotational invariance, i. e., the Hamiltonian of equation (3.4) commutes with the total angular momentum at each s : $[H_0(s), S^2] = 0$. The eigenvalues of S^2 can be written as $S(S+1)$, where S can be either integer or half-integer. The maximum allowed angular momentum is obtained when all N qubits are up-aligned, in which case the eigenvalue of the z -component of the angular momentum operator is also equal to $S = N/2$. Therefore, the possible values of S are given by $S(K) = N/2 - K$, with $K = 0, 1, \dots, \lfloor N/2 \rfloor$. From the theory of addition of angular momenta, we know that there are

$$\mathcal{N}_K^N = \binom{N}{K} - \binom{N}{K-1} \quad (3.13)$$

equivalent subspaces for each K , each having degeneracy $D_K = 2S(K) + 1 = N + 1 - 2K$. The total dimension of the Hilbert space is 2^N ,

$$\sum_{K=0}^{\lfloor N/2 \rfloor} \mathcal{N}_K^N (N + 1 - 2K) = 2^N, \quad (3.14)$$

but each subspace has a dimension that scales linearly with N . In particular, the subspace with $K = 0$, i. e., the totally symmetric space with $S = N/2$, has dimension $D_0 = N + 1$ and is unique ($\mathcal{N}_0^N = 1$).

Let us denote the restriction of $H_0(s)$ to one of the subspaces of spin $S = N/2 - K$ as $H_0^{(K)}(s)$. These subspaces are spanned by the Davies basis, i. e., the basis of common eigenvectors of $\{S^2, m^z\}$, which can be indicated as $|m; K\rangle^a$, with $m \in \mathcal{M}_K^N = \{-1 + 2K/N, -1 + 2K/N + 2/N, \dots, 1 - 2K/N\}$. In this basis, the matrix representation of $H_0^{(K)}(s)$ is tridiagonal due to the fact that m^x only couples states differing by $\Delta m = \pm 2/N$; its matrix elements are

$$\left[H_0^{(K)}(s) \right]_{m,m} = -Nsm^p, \quad (3.15)$$

$$\left[H_0^{(K)}(s) \right]_{m,m'} = -N \frac{1-s}{2} \sqrt{[1-2k + \max(m, m')][1-2k - \min(m, m')]}, \quad (3.16)$$

^aWe can equivalently indicate these states in terms of the z -component of the total spin S , i. e., by $|(N/2 - K)m; K\rangle$.

with $k = K/N$ and $m - m' = \pm 2/N$. We will extensively use this representation in our numerical simulations.

In particular, as discussed in chapter 2, we can perform several kinds of quantum annealing, which differ by the starting state and/or the form of the Hamiltonian used during the anneal. In standard QA, the system is originally prepared in the ground state of the transverse field Hamiltonian. This state is the quantum eigenstate of m^x with eigenvalue $m^{(x)} = 1$ and belongs to the subspace with $K = 0$. In iterated reverse annealing, we start from a trial solution at $s = 1$, perform the dynamics in reverse up to $s = s_{\text{inv}}$, and go back to $s = 1$. As the starting state, we can always choose an eigenstate of $H_0^{(K=0)}(s = 1)$. In both cases, the target state is the ferromagnetic state with $m = 1$, which belongs to the $K = 0$ subspace as well. Therefore, both these two kinds of quantum annealing can be worked out in the subspace with $K = 0$. For this reason, from now on, we will drop the superscript (K) in the representation of the Hamiltonian $H_0^{(K)}(s)$, and we will always work in the fully symmetric subspace $K = 0$ unless explicitly mentioned. The dimension of the relevant space is $D = N + 1$ and the total spin is $S = N/2$.

The case of adiabatic reverse annealing is only apparently more complicated. As discussed in section 3.3, the Hamiltonian (3.5) can be expressed in terms of collective operators related to the two subspaces, up and down, in which the ARA system is partitioned. The dynamics entirely occurs in the subspaces with $K = 0$, for both S_1^2 and S_2^2 . This is because at $t = 0$ both subsystems are prepared in the eigenstates of $m_{1,2}^x$ having the maximum allowed eigenvalues $m_1^{(x)} = c$ and $m_2^{(x)} = 1 - c$, respectively, which belong to the $K = 0$ subspaces, and the target state for the two subsystems are the fully symmetric states with $m_1 = c$ and $m_2 = 1 - c$ (so that $m = m_1 + m_2 = 1$). The Hamiltonian commutes with S_1^2 and S_2^2 and the dynamics cannot leave the subspaces with $K = 0$. The total angular momenta of the two subspaces read $S_1 = \lfloor Nc \rfloor / 2$ and $S_2 = (N - \lfloor Nc \rfloor) / 2$. Thus, the relevant Hilbert space for the ARA dynamics has dimension $D = (\lfloor Nc \rfloor + 1)(N - \lfloor Nc \rfloor + 1)$, quadratic in N . The maximum of this function of c is $D = (N/2 + 1)^2 = N^2/4 + N + 1$ for $c = 1/2$.

In conclusion, we can always work in the subspace(s) with $K = 0$. Compared to the total Hilbert space dimension $D = 2^N$, this provides an exponential reduction of walltime and resources in the numerical simulation of the p -spin model, which allows us to study very large systems up to $N = \mathcal{O}(10^3)$ with reasonable effort.

3.5 SPECTRAL PROPERTIES AND MINIMAL GAP

In this section, we will numerically discuss the relevant spectral properties of the p -spin Hamiltonian at $T = 0$, and compare them with some analytical results shown in reference [28]. We will not repeat analytical calculations here as they are beyond the scope of this section.

We start from the case $p = 2$, where the mean-field analysis predicts a second-order QPT at $s = s_c = \Gamma / (2 + \Gamma)$. For simplicity, we consider the case $\Gamma = 1$, so that the critical annealing fraction is $s_c = 1/3$. It is possible to analytically compute

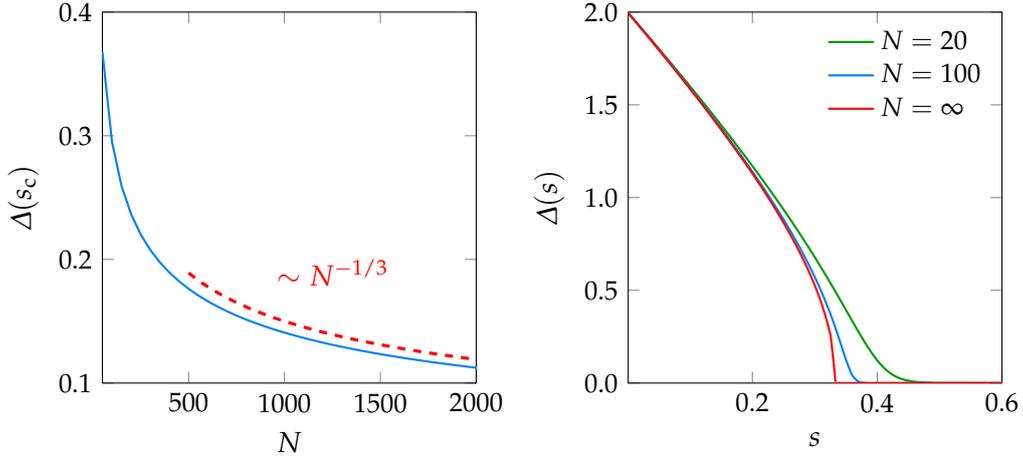


Figure 3.1 | Spectral gap of the p -spin model with $\Gamma = 1$ and $p = 2$. Left panel: scaling of the gap at the critical point $s_c = 1/3$ as a function of the system size N . The scaling of the gap is compatible with a power law $\Delta_{\min} \sim N^{-1/3}$ (red dashed line). Right panel: instantaneous gap for finite size systems, compared to the thermodynamic limit [equation (3.17)].

the instantaneous spectral gap between the ground state and the first excited state in the thermodynamic limit. The final result is

$$\lim_{N \rightarrow \infty} [E_1(s) - E_0(s)] = \begin{cases} \sqrt{12(1-s)(s_c - s)} & s \leq s_c \\ 0 & s > s_c. \end{cases} \quad (3.17)$$

For $s > s_c$, i. e., after the phase transition, the ground state and the first excited state become degenerate in the thermodynamic limit, as a consequence of the Z_2 symmetry, and exponentially close to each other for finite-size systems.

In the left panel of figure 3.1, we plot the scaling of the gap between the ground state and the first excited state at $s = 1/3$, as a function of the system size. The red dashed line represents the asymptotic behavior $\Delta_{\min} \sim N^{-1/3}$, which is in good agreement with numerical diagonalization for finite-size systems already for $N \sim \mathcal{O}(10^2)$. In the right panel, we show the instantaneous gap of equation (3.17), and compare it with the instantaneous gap of finite-size systems. The latter approaches the square root behavior of equation (3.17) for large N .

For $p = 3$, the quantum phase transition is first-order and the minimal gap closes exponentially with N . The instantaneous gap remains of order one until s is very close to the critical point. At the critical point, the slope of the ground state energy is discontinuous as a consequence of the first-order transition. The paramagnetic state survives as a metastable state for all s . By contrast, the ferromagnetic state only exists beyond the spinodal point. In reference [28], the authors have proven that the exponential closing rate α_p of the gap, i. e., the coefficient appearing in the scaling relation $\Delta_{\min}(N) \sim \exp(-\alpha_p N)$, can be written as

$$\alpha_p = \frac{1}{2} \int_0^{m_c} \operatorname{arccosh} \frac{e_c + s_c m^p}{(s_c - 1) \sqrt{1 - m^2}} dm \quad (3.18)$$

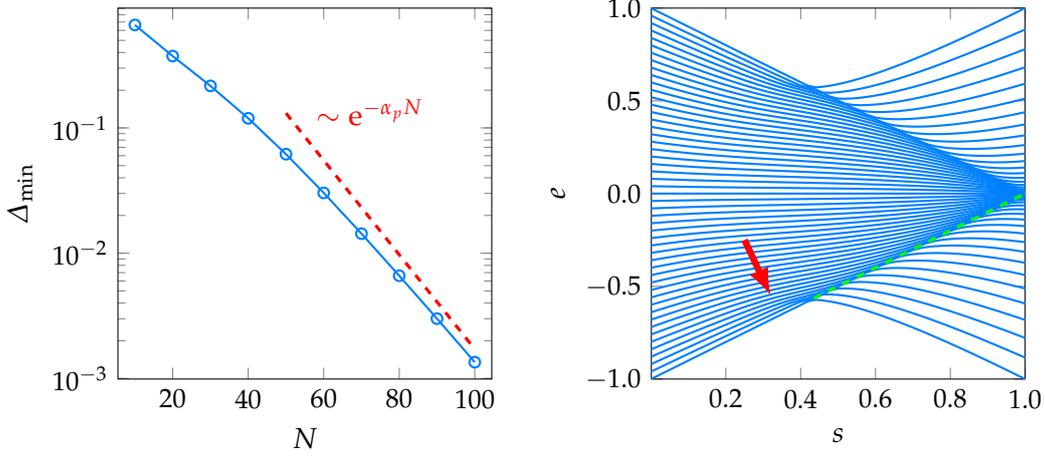


Figure 3.2 | (Left panel) Minimal gap of the p -spin model with $\Gamma = 1$ and $p = 3$, as a function of N . The red dashed line indicates the asymptotic exponential scaling, with $\alpha_p = 8.677 \times 10^{-2}$. (Right panel) Instantaneous spectrum (over N) for $N = 50$. Visual inspection confirms that the paramagnetic state survives as a metastable state beyond the critical point (see the green dashed line), whereas the ferromagnetic state only exists after the spinodal point, indicated by the red arrow.

for all $p \geq 3$. For $p = 3$, this rate is $\alpha_{p=3} = 8.677 \times 10^{-2}$. In figure 3.2 (left panel), we show the scaling of the minimal gap obtained from the numerical diagonalization of Hamiltonian (3.4), compared with the asymptotic exponential scaling with rate α_p . Numerical results are in excellent agreement with the predictions of equation (3.18). We also show in the right-hand panel of figure 3.2 the instantaneous spectrum of the p -spin model with $p = 3$, $\Gamma = 1$ and $N = 50$, in the symmetric sector $K = 0$. The ferromagnetic state, existing as a stable state past the critical point, can be analytically continued back as a metastable state up to the spinodal point, marked by the red arrow. The metastable paramagnetic state after the critical point is indicated by the green dashed line.

Finally, in figure 3.3 we show the scaling of the minimal gap in ARA as a function of the system size, for $p = 3$ and $\Gamma = 1$, for several values of the starting fraction c . The annealing path is $\lambda(s) = s$. Similarly to standard quantum annealing, the scaling of the gap stays exponential for $c = 0.7$ and $c = 0.8$, although in the latter case Δ_{\min} decreases more slowly. This is expected, as for these values of c the chosen annealing path crosses the first-order QPT. By contrast, for $c = 0.9$ the minimal gap is almost constant in the displayed window of system sizes and generally scales polynomially as a function of N [38].

3.6 THE LARGE- p LIMIT

We can represent the p -spin Hamiltonian of equation (3.2) on the Davies basis introduced in section 3.4, in the spin sector with $K = 0$. We obtain

$$H_p = -N \sum_m m^p |m; 0\rangle \langle m; 0|, \quad (3.19)$$

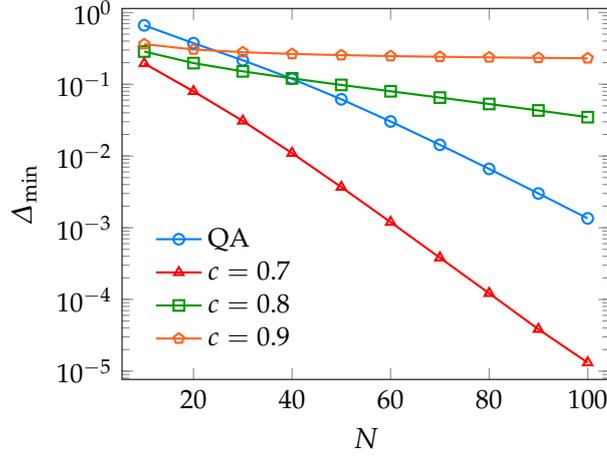


Figure 3.3 | Minimal gap of the ARA Hamiltonian (3.5) with $\Gamma = 1$ and $p = 3$, as a function of N .

where $m \in \{-1, -1 + 2/N, \dots, 1\}$. When p is large, m^p tends to zero for all $m \neq \pm 1$. Let us consider odd values of p , so that $m^p = -1$ when $m = -1$. Thus, the large- p limit (odd) of Hamiltonian (3.19) is

$$H_{p=\infty} = -N |1;0\rangle\langle 1;0| + N |-1;0\rangle\langle -1;0|. \quad (3.20)$$

The eigenvalues of $H_{p=\infty}$ are all zero except for the ground state ($e_0 = E_0/N = -1$) and the highest excited state ($e_N = E_N/N = 1$). The ground state of this simplified model can be interpreted as the target state of a quantum search. Following reference [29], we can then use the following oracular Hamiltonian to perform the search,

$$H_g = -|1;0\rangle\langle 1;0|, \quad (3.21)$$

which is very similar to Hamiltonian (3.20) except for the fact that the energy does not scale extensively with the number of entries ($\|H_g\| = 1$), and for the presence of the highly excited state in the large- p Hamiltonian (3.20). In fact, in Grover's search, all excited states have energy $E = 0$ and the target ground state has energy $E = -1$, so that only this state is selected by Grover's filter. However, it is highly unlikely that an adiabatic dynamics will involve the excited state with $E_N = N$ due to the large separation with the low-energy spectrum, meaning that the occupation probability of that state will be close to zero during the entire dynamics. Therefore, from the adiabatic quantum computation standpoint, the large- p p -spin model and the adiabatic Grover's search are equivalent for all practical purposes. Additional numerical proofs of this equality can be found in appendix B of reference [6] and in section 5.3.

3.7 THE EMBEDDING PROBLEM

We have mentioned in section 2.4 that the D-Wave machines are built upon a sparse architecture known as Chimera graph, that makes it difficult to study

fully-connected models such as the p -spin Hamiltonian (3.4) on the quantum chip due to the necessity to perform minor embedding [22, 23]. On top of that, the p -spin Hamiltonian also features infinite-range p -body interactions (e. g., it is a PUBO problem), whereas the physical hardware only allows for 2-body interactions between logical qubits (e. g., QUBO problems). Therefore, a preliminary step for mapping the p -spin system onto the Chimera graph is the decomposition of the p -body Hamiltonian (3.4) into an auxiliary Hamiltonian only featuring 2-body interactions, which typically still corresponds to a dense graph. Then, after this step, we can perform minor embedding and map the dense graph onto the sparse Chimera graph. There are proposals for improving the latter step by making use of denser architectures [97–101].

Concerning the D-Wave, both these steps require the introduction of ancillary degrees of freedom, that substantially reduce the number of logical qubits that can be embedded in the quantum chip. In fact, if M is the number of physical qubits, the scaling relation between M and N is polynomial and takes the form $M(p) \sim \alpha(p)N^{\beta(p)}$, where α and β both depend on the exponent p . Finding the actual embedding is an NP-hard problem [73] and can only be tackled heuristically.

In this section, we focus on the p -to-2-body decomposition. Our goal is therefore to rewrite the p -spin Hamiltonian (3.2) as

$$H'_p = K + \sum_{i=1}^M h_i \sigma_i^z + \sum_{i,j>i}^M J_{i,j} \sigma_i^z \sigma_j^z, \quad (3.22)$$

where N_a is the number of ancillae we need to add to perform the embedding ($M = N + N_a$). The parameters $\{K, h_i, J_{i,j}\}$ are the unknowns of this problem. They have to be chosen in such a way that the low-energy spectrum of Hamiltonian (3.22) is equal (up to a certain tolerance) to the low-energy spectrum of Hamiltonian (3.2). Moreover, all spurious energy levels introduced by the ancillae must be well-separated from the relevant spectrum so that they do not affect significantly the dynamics of the low-energy states. In this way, we can perform adiabatic quantum computation using the effective Hamiltonian (3.22) and still obtain a solution of the optimization problem we want to solve. We stress that even though Hamiltonian (3.22) mimics the spectrum of the starting Hamiltonian, the quantum dynamics produced by using it in the annealing Hamiltonian (3.4) can be totally different from that of the original model, due to the additional ancillary qubits we add. However, if we are only interested in reading the final state and disregard the instantaneous dynamics, which is often the case in optimization tasks, this approach constitute a viable method.

In the purely adiabatic limit, only the ground state and the first excited state of Hamiltonian (3.22) are relevant. Even for faster dynamics, only few of the eigenstates of the p -spin Hamiltonian might be involved in the quantum evolution. For this reason, in general, we do not need to match the whole spectrum and we can match just the first L eigenvalues of the two Hamiltonians. For very simple instances, we can always consider $L = 2^N$ and match the entire spectrum of the original Hamiltonian (3.2) without much effort.

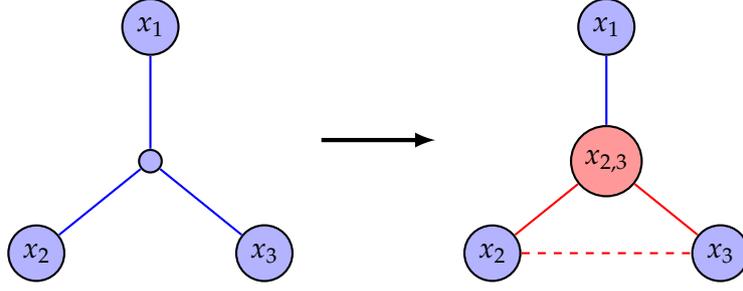


Figure 3.4 | Pictorial representation of the AND embedding of equation (3.24). On the left, we show the graph corresponding to the 3-body term $Jx_1x_2x_3$. On the right, we show the resultant 2-body Hamiltonian. The blue circles are the original qubits and the red circle is the ancilla. The blue lines represent the coupling J , the solid red lines represent -2δ and the dashed red line corresponds to a coupling $-\delta$.

In order to perform the p -to-2-body decomposition, we can use the AND embedding [66]. This algorithm is more straightforward if we convert the binary representation $\sigma \in \{-1, 1\}$ to the equivalent $x \in \{0, 1\}$ by using the linear transformation $\sigma = 1 - 2x$. Given two binary variables x_1 and x_2 , we want to associate their product x_1x_2 (corresponding to the logical AND) to an ancilla $x_{1,2}$, in such a way that $x_{1,2} = x_1 \wedge x_2$. When either x_1 or x_2 is zero (or both), $x_{1,2} = 0$ as well. If both x_1 and x_2 are equal to one, also $x_{1,2} = 1$. When we use a physical qubit to represent the new logical variable $x_{1,2}$, however, it may happen that during the dynamics the logical AND could be broken as the three variables evolve independently, therefore at some point we can have $x_{1,2} \neq x_1 \wedge x_2$. Naturally, these configurations are not allowed due to the physical interpretation of $x_{1,2}$ as the product of x_1 and x_2 , and must be disregarded. We can do that by applying a large penalty $E_{\text{pen}}^a(x_1, x_2, x_{1,2})$ every time the AND clause is violated. For example, we can use the following function:

$$E_{\text{pen}}^a(x_1, x_2, x_{1,2}) = \delta(3x_{1,2} + x_1x_2 - 2x_{1,2}x_1 - 2x_{1,2}x_2), \quad (3.23)$$

where $\delta > 0$ is a large value compared to all other energies. The penalty term is $E_{\text{pen}}^a = 0$ when the AND clause is satisfied, and is $E_{\text{pen}}^a \geq \delta$ if $x_{1,2} \neq x_1 \wedge x_2$.

As an example, consider the 3-body term $Jx_1x_2x_3$. If we define $x_{2,3} = x_2 \wedge x_3$ and use equation (3.23), the resulting Hamiltonian can be written as

$$Jx_1x_2x_3 \equiv Jx_1x_{2,3} + \delta(3x_{2,3} + x_2x_3 - 2x_{2,3}x_2 - 2x_{2,3}x_3), \quad (3.24)$$

where the equivalence is intended as an equality between the $L = 8$ lowest eigenvalues of the two Hamiltonians. A schematic representation of this procedure is shown in figure 3.4.

In terms of the binary representation $x_i \in \{0, 1\}$, the 2-body Hamiltonian (3.22) can be rewritten as

$$H'_p = c_0 + \sum_{i=1}^M c_i x_i + \sum_{i,j>i}^M d_{i,j} x_i x_j, \quad (3.25)$$

where the correspondence between the free parameters in the σ and x representation is given by

$$K = c_0 + \frac{1}{2} \sum_{i=1}^M c_i + \frac{1}{4} \sum_{i,j>i}^M d_{i,j}, \quad (3.26a)$$

$$h_i = -\frac{1}{2} c_i - \frac{1}{4} \sum_{j=1}^M (d_{j,i} + d_{i,j}) \quad (3.26b)$$

$$J_{i,j} = \frac{1}{4} d_{i,j}. \quad (3.26c)$$

If we denote by $\{E_{p,i}\}$ ($\{E'_{p,i}\}$) the eigenvalues of the original (mapped) Hamiltonian, our goal is to minimize the objective fitness function

$$F = \frac{1}{L} \sum_{i=1}^L (E_{p,i} - E'_{p,i})^2 + E_{\text{pen}}^a + E_{\text{pen}}^v, \quad (3.27)$$

where the first term matches eigenvalues, the penalty term E_{pen}^a ensures that all AND clauses are satisfied and the final term E_{pen}^v ensures that eigenvectors are sorted correctly. If there are l unsorted eigenvectors, this term can be written as $E_{\text{pen}}^v = l\delta$. Once the fitness function is defined, the optimal solution is $\{c_0, c_i^*, d_{i,j}^*\} = \arg \min F$.

There are many heuristic algorithms designed to face this problem. For example, the piece of software PyQUBO [24] represents the starting Hamiltonian as a tree, folds it into a polynomial of order p and then simplifies it to a second-order polynomial by a convenient application of equation (3.24). Instead of using this iterative scheme, we decided to test a genetic global optimization algorithm to achieve the same goal.

3.7.1 Genetic algorithms

Genetic algorithms are population-based stochastic metaheuristics used to solve optimization tasks, inspired by the Darwinian theory of evolution [102, 103]. Similarly to natural selection, which leads to the survival of only the fittest individuals, genetic algorithms perform an evolution process that favors the best solutions to the optimization problem. More precisely, genetic algorithms operate on encoded representations of the trial solution denoted as chromosomes, where each free parameter is a gene. In our case, the chromosomes are

$$\mathbf{v} \equiv (c_0, c_1, \dots, c_M, d_{1,2}, \dots, d_{M-1,M}). \quad (3.28)$$

The chromosome's length is the number of free parameters in Hamiltonian (3.22), i. e., $W = (M^2 + M + 2)/2$. The criterion used to rank different chromosomes is the fitness function of equation (3.27): the fittest individuals are those who minimize the fitness function or, alternatively, maximize $-F$. It is implicit that choosing the best fitness function is usually the most crucial step of genetic algorithms.

Once the fitness function is specified, the algorithm goes through several phases. Initially, a population of random individuals is generated. Then, the chromosomes are propagated through a set of generations until a termination criterion is satisfied, such as the achievement of a specific tolerance or a maximum number of generations. In each generation, the population is subjected to three different operators whose goal is to increase the genetic variability of the trial solutions and to select the fittest individuals that will survive to next generations. These are the following:

MUTATION In nature, individuals can spontaneously develop specific traits to adapt to changes in their surroundings. This ability is mimicked in genetic algorithms by assigning a mutation probability p_m , so that chromosomes can randomly mutate their genes to improve their fitness. Typically, the mutation operator extracts a random number from a Gaussian distribution with zero mean and variance σ^2 , and then adds it to a random gene.

CROSSOVER During the crossover phase, portions of two parent chromosomes are combined together to produce two children. The crossover occurs with a certain probability p_c and can happen in multiple ways. For example, in one-point crossover, a random gene is selected, the two parents are split at this point and their tails are exchanged. Instead, in two-point crossover, two random genes are extracted and the parent chromosomes are split into three pieces; the offspring is built by combining alternating parts of the parent chromosomes.

SELECTION In the final step, the genetic algorithm selects the best individuals that will become parents in the next generation. This is usually done by setting up a tournament between N_T randomly selected individuals at a time (with repetitions allowed), and ranking them according to their fitness value.

In figure 3.5, we show a schematic representation of these steps. The net result is the improvement of the average fitness of the individuals across the generations.

In reference [1], we tested the feasibility of using the genetic algorithm to solve the p -to-2-body mapping. As a proof of principle, we focused on two very simple instances whose QUBO decomposition can be worked out analytically using equation (3.24), i. e., $N = 3$ and $p = 3$ (for which the number of physical qubits after the decomposition is $M = 4$), and $N = 4$ with $p = 3$ ($M = 6$). We aimed to match all $L = 2^N$ eigenvalues of the starting Hamiltonian. We used the Python library DEAP [104] to perform the genetic evolution.

A preliminary step in every kind of heuristic optimization is the setting of hyperparameters, which can be done by analyzing the performance of the algorithm for several choices and by selecting the best configuration [105, 106]. In our case, we varied the variance σ^2 of the Gaussian mutation operator, the number of individuals involved in the tournament selection N_T , and the type of crossover, one- or two-point. By contrast, the mutation and crossover probability have been fixed to $p_c = 0.4$ and $p_m = 0.7$ as we verified that the performance of the algorithm was

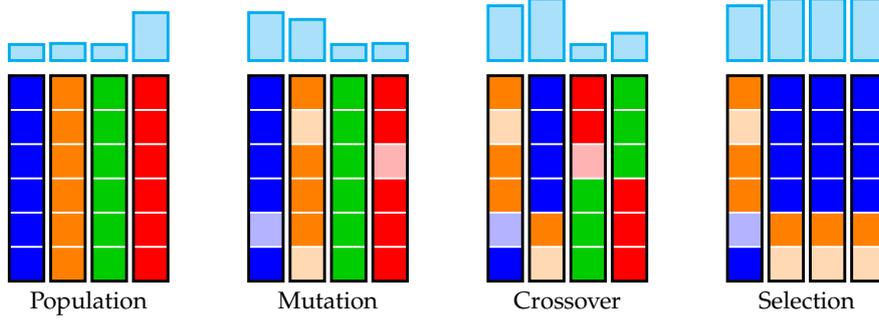


Figure 3.5 | Schematic representation of the elementary step of a genetic algorithm. The genes of a starting population of chromosomes are randomly mutated; then, pairs of individuals are crossed-over and the best children are selected according to their fitness value (represented by the light blue bars above).

almost independent on the choice of these parameters. We performed $N_r = 100$ runs for each configuration of hyperparameters, where in each run we propagated a population of $N_{\text{pop}} = 20$ random individuals for $N_g = 1000$ generations. We then chose the best configuration of hyperparameters according to the minimal median fitness. In order to ensure a wide separation between the physical and unphysical parts of the spectrum, we fixed $\delta = 50$ in the penalty terms.

As an example, we report here our results concerning the case $N = 3$ and $p = 3$, and point the readers to reference [1] for further details concerning $N = 4$. For $N = 3$, the best configuration of hyperparameters involves a two-point crossover, standard deviation of the mutation operator $\sigma = 0.2$, and a tournament pool size $N_T = 5$. Using this information, we evolved a population of $N_{\text{pop}} = 20$ chromosomes for $N_g = 25\,000$ generations. At the end of the algorithm, we selected the best individual in the final population, i. e., the one having the smallest fitness value ($F = 9.88 \times 10^{-8}$). In table 3.1, we report our results for $N = 3$, compared to the analytic solution. We note that the both the free parameters and the eigenvalues are very close to their theoretical value. Also the eigenvectors are correctly replicated. In the genetic eigenvector column, the first qubit is the ancilla and is defined as $x_{1,2} = x_1 \wedge x_2$. The AND clause is always satisfied in this lower part of the spectrum. Eigenvectors that violate the AND clause have higher energy and are not included in the comparison.

With this solution at hand, we can perform an adiabatic quantum computation and compare it with the original p -spin Hamiltonian (3.4), with a caveat. In fact, the global adiabatic criterion of equation (2.7) states that, in order for the evolution to be adiabatic, the annealing time must satisfy the condition

$$t_f \gg \max_{s,a \neq b} \frac{|\langle E_a(s) | \partial_s H_0(s) | E_b(s) \rangle|}{\Delta_{\min}^2}. \quad (3.29)$$

Numerically, we can check that the minimal gap of the effective 2-body model is very similar to the original model ($\Delta_{\min} \sim 1.3$ versus $\Delta_{\min} \sim 1.2$). On the other hand, due to the presence of energy penalties, the rate of variation of the effective

Table 3.1 | Comparison between the best chromosome obtained by the genetic algorithm and the analytically computed solution for $N = 3$ and $p = 3$ ($M = 4$). In the genetic eigenvectors, the first qubit is the ancilla and is defined as $x_{1,2} = x_1 \wedge x_2$.

Free parameters		Eigenvectors		Eigenvalues	
Analytic	Genetic	Analytic	Genetic	Analytic	Genetic
-3	-2.99919	000⟩	0⟩ 000⟩	-3	-2.99919
-150	-150.853	001⟩	0⟩ 100⟩	-1/9	-0.11138
26/9	2.88781	010⟩	0⟩ 001⟩	-1/9	-0.11129
26/9	2.88795	100⟩	0⟩ 010⟩	-1/9	-0.11124
26/9	2.88790	011⟩	1⟩ 110⟩	1/9	0.11111
100	100.720	101⟩	0⟩ 101⟩	1/9	0.11120
100	101.118	110⟩	0⟩ 011⟩	1/9	0.11120
16/3	5.33174	111⟩	1⟩ 111⟩	3	2.99999
-158/3	-53.6496				
-8/3	-2.66531				
-8/3	-2.66545				

Hamiltonian is typically much larger than that of the original p -body model, in particular it scales linearly with δ . Therefore, a larger δ implies a wider separation from the spurious part of the spectrum but a longer adiabatic time scale, and vice versa, a smaller δ implies a shorter adiabatic time but closer unphysical energy states that can potentially affect the dynamics. Naturally, this trade-off is present in all kinds of QUBO mappings involving ancillae.

With that said, it makes little sense to compare the dynamics of the original p -body model with those of the effective model for a fixed annealing time, as the same numerical value has entirely different meanings in the two cases, in relation to their own adiabatic time scale. The important thing to stress is that with the p -to-2-body mapping we have access to a physically realizable Hamiltonian whose ground state provides the solution to the starting problem, at the expense of consuming more qubits and/or having to run the program for a longer annealing time. Therefore, in the left panel of figure 3.6, we show the scaling of the error probability $p_e(t_f) = 1 - |\langle E_0(t_f) | \psi(t_f) \rangle|^2 \equiv 1 - p_{gs}(t_f)$ with the annealing time t_f in three cases: the original 3-body Hamiltonian, the mapped 2-body model with $\delta = 15$, and the mapped model with $\delta = 50$. The plot confirms our analysis, as the error probability decreases more slowly with t_f for larger δ , meaning that the adiabatic regime of the effective model is reached at longer annealing times. For example, if our goal were to find the ferromagnetic ground state with 95% probability, numerical simulations would yield $t_f(p = 3) = 5.4$, $t_f(\delta = 15) = 12.6$, and $t_f(\delta = 50) = 31.5$ in appropriate units. Indeed, as evident from the right panel of figure 3.6, although the dynamics are different, the final ground state probability is $p_{gs} = 0.95$ in all cases if we use the aforementioned annealing times.

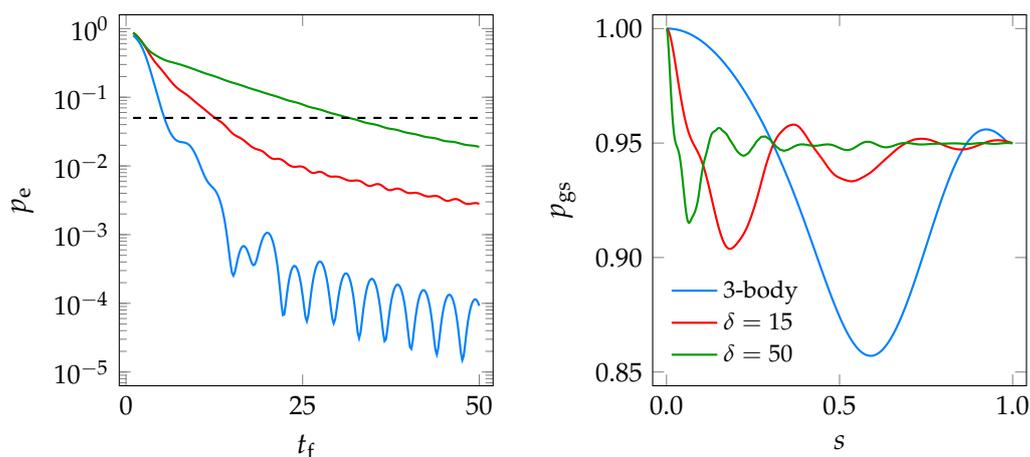


Figure 3.6 | (Left panel) Error probability as a function of the annealing time for the p -body model and the corresponding QUBO mapping. (Right panel) Instantaneous ground state probability as a function of s . The annealing time has been chosen to yield a confidence level of 95% of finding the ferromagnetic ground state (see main text).

3.8 CONCLUSIONS

In this chapter, we have introduced the ferromagnetic p -spin model and discussed its features both in the thermodynamic limit and for finite-size systems. Due to its mean-field character and the fact that it resembles the adiabatic Grover's search, this model is heavily studied in the field of AQC as a template of NP-hard problems. As it is a fully-connected system, many ancillae are necessary in order to perform quantum annealing on the D-Wave machine. The largest chip now available features 2048 physical qubits. As an example, for a p -spin system with $p = 3$ we saturate $M = 2048$ physical qubits just by decomposing the p -body Hamiltonian of a system of $N = 86$ logical qubits into a 2-body one, but the maximum number of logical qubits quickly decreases with increasing p . For $p = 7$, the maximum embeddable size is $N = 20$ logical qubits (see figure 3.7). These numbers do not even factor in the additional ancillae needed for minor embedding, which are usually the bottleneck for resources. If we take into account the Chimera graph, the number of logical qubits that we can embed sharply drops. For $p = 3$, we can only embed approximately $N = 20$ qubits, as shown by the blue dotted line in figure 3.7. Therefore, numerical simulations are fundamental in order to describe the behavior of larger instances of the p -spin system during a quantum annealing, especially in the large- p limit.

D-Wave chips are physical objects, and as such they are embedded in an environment and interact with it. Thus, any numerical analysis cannot leave aside the environmental effects, which affect the quantum evolution in many ways [4, 19, 82, 83, 107–109]. Simulating the dissipative dynamics of a quantum annealer is a very demanding problem involving several different approximations. However, it appears that many qualitative features of dissipative quantum annealing can be captured by just resorting to a Markovian approximation. Therefore, in the next

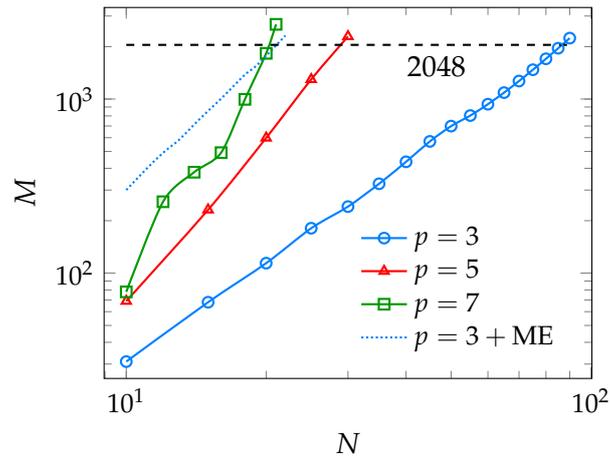


Figure 3.7 | Number of physical qubits M versus logical qubits N needed to decompose the p -body Hamiltonian (3.2) into the effective 2-body model of equation (3.22). The blue dotted line represents the number of physical qubits to perform both the p -to-2-body decomposition and the minor embedding for $p = 3$. These scalings have been computed using PyQUBO [24] and minorminer [25].

chapter we will revisit the microscopic derivation of a Markovian quantum master equation in the Lindblad form and we will discuss its numerical implementation. We will then heavily use this tool in the following chapters.

OPEN QUANTUM SYSTEMS

4.1 INTRODUCTION

In real-world devices, the ensemble of qubits constituting a QPU is an open quantum system, whose dynamics cannot be described by a Schrödinger equation, due to the interaction with the environment. In this chapter, we are going to review the mathematical description of open quantum systems within the Born, Markov, and rotating wave approximations. In the weak-coupling regime, the dynamics of an open quantum system can be described by a quantum dynamical semigroup leading to a time-local quantum master equation in the Lindblad form for the reduced density matrix of the qubits. This equation can be rephrased as a stochastic Schrödinger equation for a ket state, which simplifies numerical analysis for large systems and allows for an easy parallelization. These tools are invaluable to provide a qualitatively accurate description of physical quantum annealers.

This chapter is organized as follows. In section 4.2, we introduce the concept of quantum dynamical semigroup and derive the adiabatic quantum master equation in the Lindblad form that we will use in the remainder of this thesis. In section 4.3, we show that the adiabatic master equation has an analogue formulation in terms of a stochastic wave function approach, in which the environment induces quantum jumps between pairs of energy eigenstates at random times. The cost comparison between these two formulations as a function of the system size is discussed in appendix B. As an example of an application, we discuss the ferromagnetic p -spin model with collective dephasing in section 4.4. Finally, in section 4.5, we discuss a way to include effects beyond the Born-Markov approximation using exact diagonalization and a technique known as short-iterative Lanczos, and we show our results for a single open qubit as a test bed for this approach.

4.2 DYNAMICAL SEMIGROUP

In quantum mechanics, the state of a system is a ket $|\psi(t)\rangle$ that evolves in time according to the Schrödinger equation, or, equivalently, a density operator $\rho(t)$ that evolves according to the Liouville-Von Neumann equation

$$\partial_t \rho(t) = -i [H(t), \rho(t)] \equiv \mathcal{L}(t) \rho(t) \quad (4.1)$$

where $H(t)$ is the system Hamiltonian. The Liouvillian superoperator $\mathcal{L}(t)$ is introduced so that we can formally write the dynamics of the density operator as

$$\rho(t) = \mathcal{T} e^{\int_0^t \mathcal{L}(t') dt'} \rho(0), \quad (4.2)$$

where \mathcal{T} is the time ordering operator.

Suppose we have a quantum system whose dynamics we want to unveil, such as the QPU of a quantum annealer. This system is never fully isolated as it interacts with the surrounding environment, which acts as a thermal bath. If \mathcal{H}_S is the Hilbert space of the qubit system S and \mathcal{H}_B is the Hilbert space of the environment B , the total Hilbert space is $\mathcal{H} = \mathcal{H}_S \otimes \mathcal{H}_B$. The system $S + B$ is a closed quantum system (see figure 4.1) and its dynamics is described by equation (4.2). On the other hand, S is an open quantum system exchanging energy with B , therefore the dynamics of the reduced density operator $\rho_S(t) = \text{tr}_B \rho(t)$, obtained by taking the partial trace over the degrees of freedom of the bath, cannot be expressed by the right-hand side of equation (4.2). Generally, we are not interested in following the dynamics of the bath and we want a closed-form equation that allows us to just follow the reduced system S . In other words, we want to find the superoperator $V(t)$, denoted dynamical map, that gives us

$$\rho_S(t) = \text{tr}_B \left[\mathcal{T} e^{\int_0^t \mathcal{L}(t') dt'} \rho(0) \right] \equiv V(t) \rho_S(0). \quad (4.3)$$

Under certain assumptions, the dynamical map $V(t)$ has the structure of a semigroup, i. e., it satisfies

$$V(t_1)V(t_2) = V(t_1 + t_2), \quad t_1, t_2 > 0. \quad (4.4)$$

In particular, this equation (Chapman-Kolmogorov) is valid when the bath correlation functions decay more rapidly than the characteristic time scales of the system evolution, which corresponds to a Markovian approximation where we can disregard any memory effects. In turn, this corresponds to the fact that the reduced system is only weakly coupled to its environment. The fact that we are limited to $t > 0$ is a consequence of the irreversibility of dissipative dynamics [14].

When the reduced dynamics are expressed by the dynamical semigroup (4.4), it is possible to find a generator $\mathcal{L}(t)$ so that $V(t) = \mathcal{T} \exp\left(\int_0^t \mathcal{L}(t') dt'\right)$ and the dynamical equation for $\rho_S(t)$ takes the form of a time-local master equation in the Gorini-Kossakowsky-Sudarshan-Lindblad form [30, 31] (Lindblad for brevity) formally similar to equation (4.1). $\mathcal{L}(t)$ is the Lindbladian of the system. It contains both the unitary part of the dynamics and a dissipative term that incorporates the effects of the environment on the reduced dynamics. The Lindblad equation reads

$$\partial_t \rho_S(t) = \mathcal{L}(t) \rho_S(t) = -i [\tilde{H}(t), \rho_S(t)] + \mathcal{D}[\rho_S(t)], \quad (4.5)$$

where $\tilde{H}(t)$ may contain additional terms due to the coupling with the environment and \mathcal{D} is the dissipator superoperator, whose general form is

$$\mathcal{D}[\rho_S(t)] = \sum_k \gamma_k \left[L_k(t) \rho_S(t) L_k^\dagger(t) - \frac{1}{2} \{ L_k^\dagger(t) L_k(t), \rho_S(t) \} \right]. \quad (4.6)$$

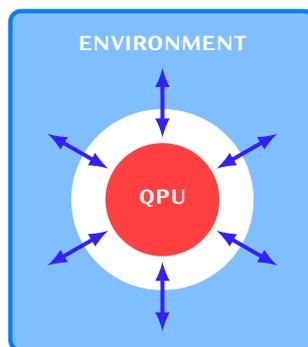


Figure 4.1 | Schematic representation of an open quantum system.

The Lindblad operators $L_k(t)$ represent the different dissipation channels, each occurring with a rate γ_k . The Markovian approximation requires $\gamma_k \geq 0$. The Lindblad equation (4.5) ensures the complete positivity of the density operator at all times and thus preserves the probabilistic interpretation of its diagonal elements when represented on a basis [14].

4.2.1 Microscopic Lindblad equation

The Lindblad equation (4.5) can be derived from first principles in the adiabatic limit when the system-bath coupling strength is weak compared to other energy scales [32]. The corresponding master equation is known as *adiabatic master equation* (AME) and is often used in the literature to model the dissipative dynamics of quantum annealers [4, 6, 7, 110]. In this section, we want to briefly discuss this derivation as the AME will be the starting point of our numerical analysis.

The $S + B$ Hamiltonian reads

$$H(t) = H_0(t) + H_B + H_{SB}, \quad (4.7)$$

where H_B is the bath Hamiltonian and H_{SB} couples the two subsystems. In particular, we model the environment as a collection of noninteracting harmonic oscillators with Hamiltonian

$$H_B = \sum_k \omega_k a_k^\dagger a_k, \quad (4.8)$$

in equilibrium at inverse temperature $\beta = 1/T$. The interaction Hamiltonian reads

$$H_{SB} = \sum_\alpha g_\alpha A_\alpha \otimes B_\alpha, \quad (4.9)$$

where A_α (B_α) are dimensionless qubit (bath) operators and g_α are coupling energies. The operators $B_\alpha(t) = \exp(i H_B t) B_\alpha \exp(-i H_B t)$ enter the definition of the bath correlation function $C_{\alpha\beta}(t) = \langle B_\alpha(t) B_\beta(0) \rangle_B$, where angular brackets denote thermal averages. We expect the modulus of $C_{\alpha\beta}(t)$ to decay as $|C_{\alpha\beta}(t)| \sim e^{-t/\tau_B}$, where τ_B is a typical time scale of bath correlations [32].

At $t = 0$, we can always prepare the full system in the uncorrelated state $\rho(0) = \rho_S(0) \otimes \rho_B$, where ρ_B is the Gibbs density operator at temperature T . However, at $t > 0$ the system state will not be separable in general, which makes it cumbersome to evaluate the partial traces over B . The AME is recovered when the Born, Markov, and rotating wave approximation are valid:

BORN APPROXIMATION At all times, we can neglect correlations among qubits and bath, i. e., we can always write $\rho(t) = \rho_S(t) \otimes \rho_B + \chi(t) \approx \rho_S(t) \otimes \rho_B$. This is valid in the *weak coupling* regime, where $g\tau_B \ll 1$.

MARKOV APPROXIMATION The system has no memory of itself at preceding times. This happens because the correlation time scale τ_B is much shorter than the relaxation time scale $1/g$. Thus, if $g\tau_B \ll 1$, we can use the Markov approximation and write the adiabatic master equation as a *time-local equation*.

ROTATING WAVE APPROXIMATION In the interaction picture, *we disregard terms in the Hamiltonian that contain fast oscillating terms* and only keep counterrotating ones. This enforces the energy conservation and guarantees the complete positivity of the reduced density matrix during the evolution.

More precisely, it is shown in reference [32] that these approximations are valid when the following inequalities are satisfied,

$$\frac{v\tau_B}{t_f\Delta_{\min}} \ll \min\left(\tau_B\Delta_{\min}, \frac{1}{\tau_B\Delta_{\min}}\right), \quad g\tau_B \ll \min(1, \Delta_{\min}/g), \quad (4.10)$$

where $v = \max_{s \in [0,1], a \neq b} \langle E_a(s) | \partial_s H_0(s) | E_b(s) \rangle$. The first inequality enforces the (global) adiabatic condition of equation (2.9) and ensures that the instantaneous changes in the energy eigenbasis are small in the time scale τ_B . The second inequality ensures the validity of the Born-Markov approximation and of the perturbative expansion of the Lindbladian.

In conclusion, the AME has the form

$$\partial_t \rho_S(t) = -i[H_0(t) + H_{\text{LS}}(t), \rho_S(t)] + \mathcal{D}[\rho_S(t)], \quad (4.11)$$

where $H_{\text{LS}}(t)$ is the Lamb shift Hamiltonian and \mathcal{D} is the adiabatic dissipator. They are both expressed in terms of Lindblad jump operators, which are

$$L_{\alpha,ab}(t) = \langle E_a(t) | A_\alpha | E_b(t) \rangle | E_a(t) \rangle \langle E_b(t) |, \quad (4.12)$$

where $\{|E_a(t)\rangle\}$ are the instantaneous eigenstates of $H_0(t)$. Each Lindblad operator $L_{\alpha,ab}(t)$ describes a jump of frequency $\omega_{ba}(t) = \Delta_{ba}(t)$ ($\hbar = 1$), from $|E_b(t)\rangle$ to $|E_a(t)\rangle$ for $a \neq b$. The Lindblad operators with $a = b$ provide dephasing channels. There are D^2 independent Lindblad operators, where D is the dimension of the Hilbert space \mathcal{H}_S . Omitting the time-dependence of all quantities and operators for

brevity, the expressions of H_{LS} and \mathcal{D} are

$$H_{\text{LS}} = \sum_{\alpha\beta, a \neq b} S_{\alpha\beta}(\omega_{ba}) L_{\alpha,ab}^\dagger L_{\beta,ab} + \sum_{\alpha\beta ab} S_{\alpha\beta}(0) L_{\alpha,aa}^\dagger L_{\beta,bb}, \quad (4.13a)$$

$$\begin{aligned} \mathcal{D}[\rho_S] = & \sum_{\alpha\beta, a \neq b} \gamma_{\alpha\beta}(\omega_{ba}) \left(L_{\beta,ab} \rho_S L_{\alpha,ab}^\dagger - \frac{1}{2} \left\{ L_{\alpha,ab}^\dagger L_{\beta,ab}, \rho_S \right\} \right) \\ & + \sum_{\alpha\beta ab} \gamma_{\alpha\beta}(0) \left(L_{\beta,aa} \rho_S L_{\alpha,bb}^\dagger - \frac{1}{2} \left\{ L_{\alpha,aa}^\dagger L_{\beta,bb}, \rho_S \right\} \right), \end{aligned} \quad (4.13b)$$

where the spectral functions $\gamma_{\alpha\beta}(\omega)$ and $S_{\alpha\beta}(\omega)$ are the real and imaginary part of the Fourier transform of the bath correlation functions:

$$\Gamma_{\alpha\beta}(\omega) = g_\alpha g_\beta \int_0^\infty e^{i\omega t} C_{\alpha\beta}(t) dt = \frac{1}{2} \gamma_{\alpha\beta}(\omega) + i S_{\alpha\beta}(\omega). \quad (4.14)$$

The two functions $\gamma_{\alpha\beta}$ and $S_{\alpha\beta}$ are connected to each other via the Kramers-Kronig relations. It can be proven that, if the bath is in thermal equilibrium, then $\gamma_{\alpha\beta}(\omega)$ satisfies the detailed balance [or *Kubo-Martin-Schwinger* (KMS)] condition $\gamma_{\alpha\beta}(-\omega) = e^{-\beta\omega} \gamma_{\alpha\beta}(\omega)$.

The function $\gamma_{\alpha\beta}(\omega)$ is related to the bath spectral density $J(\omega)$. In the continuous limit, it is customary to model the low-frequency behavior of $J(\omega)$ as a power law. At high frequencies, a rigid or exponential cutoff is usually employed to keep a finite number of states, up to a frequency ω_c :

$$J(\omega) = \sum_k \delta(\omega - \omega_k) = \tilde{\eta} \omega^\nu \omega_c^{1-\nu} e^{-\omega/\omega_c}. \quad (4.15)$$

The parameter $\tilde{\eta}$ has the dimensions of time squared. The exponent $\nu > 0$ distinguishes different kinds of dissipation: sub-Ohmic $0 < \nu < 1$, Ohmic ($\nu = 1$), or super-Ohmic ($\nu > 1$). In terms of $J(\omega)$, we can write

$$\gamma_{\alpha\beta}(\omega) = 2\pi \eta_{\alpha\beta} \frac{\omega^\nu \omega_c^{1-\nu} e^{-\omega/\omega_c}}{1 - e^{-\beta\omega}}, \quad (4.16)$$

where $\eta_{\alpha\beta} = \tilde{\eta} g_\alpha g_\beta$ is a dimensionless coupling strength. In the following, we will always consider Ohmic baths ($\nu = 1$) and equal couplings between system and bath operators, i. e., $g_\alpha = g$. Moreover, we assume that $\gamma_{\alpha\beta}$ is diagonal, thus $\eta_{\alpha\beta} = \eta \delta_{\alpha\beta} = \tilde{\eta} g^2 \delta_{\alpha\beta}$ ^a. We will use the parameter η to express the system-bath coupling strength in a dimensionless way.

Since the spectral matrices are diagonal, equations (4.13) assume the form

$$H_{\text{LS}} = \sum_{\alpha,\omega} S_\alpha(\omega) L_{\alpha,\omega}^\dagger L_{\alpha,\omega}, \quad (4.17a)$$

$$\mathcal{D}[\rho_S] = \sum_{\alpha,\omega} \gamma_\alpha(\omega) \left(L_{\alpha,\omega} \rho_S L_{\alpha,\omega}^\dagger - \frac{1}{2} \left\{ L_{\alpha,\omega}^\dagger L_{\alpha,\omega}, \rho_S \right\} \right) \quad (4.17b)$$

^aIt is always possible to diagonalize the rate matrix $\gamma_{\alpha\beta}$ by a unitary transformation u such that $\gamma'_i \delta_{ij} = \sum_{\alpha\beta} u_{i\alpha} \gamma_{\alpha\beta} u_{j\beta}^\dagger$. This transformation also redefines the Lindblad operators $L_{\alpha,\omega} \rightarrow \sum_{\alpha'} u_{i\alpha'}^\dagger L_{\alpha',\omega}$. Here, we are assuming that the spectral matrix is already diagonal, which corresponds to having uncorrelated baths. Of course, the matrix $S_{\alpha\beta}$ is diagonal as well.

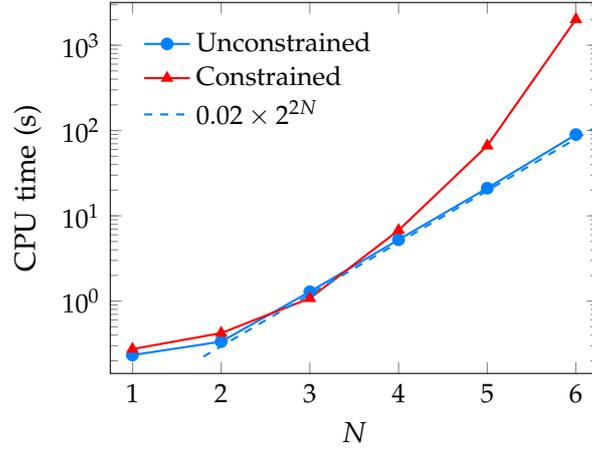


Figure 4.2 | Scaling of the CPU time as a function of N when simulating the Lindblad master equation. Notice the logarithmic y axis.

where we also rewrote the sums over eigenvalues as a constrained sum over Bohr frequencies, e. g., $\omega(t) = E_b(t) - E_a(t)$, in such a way that

$$L_{\alpha,\omega}(t) = \sum_{E_b(t)-E_a(t)=\omega} \langle E_a(t) | A_\alpha | E_b(t) \rangle | E_a(t) \rangle \langle E_b(t) |. \quad (4.18)$$

Equations (4.17) and (4.18) are useful to write the Lindblad equation in a more compact form. However, from the numerical standpoint, using the unconstrained form of equation (4.13) is more efficient as it eliminates the additional step of grouping equal frequencies together. In addition, we can easily work out the form of the matrix elements of both H_{LS} and \mathcal{D} in the energy eigenbasis:

$$(H_{LS})_{ii} = \sum_{\alpha a} S_\alpha(\omega_{ia}) |A_\alpha|_{ia}^2; \quad (4.19a)$$

$$(\mathcal{D})_{ii} = \sum_{\alpha a} |A_\alpha|_{ia}^2 [\gamma_\alpha(\omega_{ai}) \rho_{aa} - \gamma_\alpha(\omega_{ia}) \rho_{ii}]; \quad (4.19b)$$

$$(\mathcal{D})_{ij} = -\frac{\rho_{ij}}{2} \sum_{\alpha} \left[\gamma_\alpha(0) \left((A_\alpha)_{ii} - (A_\alpha)_{jj} \right)^2 + \sum_{a \neq i} \gamma_\alpha(\omega_{ia}) |A_\alpha|_{ia}^2 + \sum_{a \neq j} \gamma_\alpha(\omega_{ja}) |A_\alpha|_{ja}^2 \right]. \quad (4.19c)$$

The last equation holds for $i \neq j$.

In figure 4.2, we show the scaling of the computational time for a typical master equation simulation as a function of the number of qubits N . The Hilbert space dimension is $D = 2^N$. If we use the constrained sums of equation (4.18), the computational time is more than exponential in N . By contrast, if we use the unconstrained sums and the analytic matrix elements of equation (4.19), the scaling of the computational time is $\approx 0.02 \times 2^{2N}$, which is far better than the previous case. Even though this is the fastest scaling we can attain using a quantum master equation, it is still not ideal for computation and only small systems up to $\mathcal{O}(10)$ qubits can be actually simulated in a reasonable time.

4.3 QUANTUM JUMPS

In order to improve the scaling of the CPU time with N , the quantum trajectories method, or *Monte Carlo wave function* (MCWF), can be used for the unraveling of the adiabatic master equation [14]. In this approach, the AME dynamics is reconstructed as an ensemble average over a large number N_T of quantum trajectories. In each trajectory, a ket state is evolved, rather than a density matrix, subjected to a non-Hermitian Hamiltonian and under the action of a stochastic force causing instantaneous jumps between energy eigenstates. The resulting trajectory is the result of a continuous indirect measurement of the environment in the energy eigenbasis. In this way, the computational cost of the simulation is reduced from D^2 to (at least) D , at the cost of having to run multiple trajectories to recover the correct dynamics. This tradeoff can be readily compensated by running the quantum trajectories in parallel on high-performance computers using MPI/openMP, as MCWF is particularly well-suited for parallelization.

The quantum jumps method is a common tool in quantum optics [111–113] and statistical physics [114], and can also be extended to describe non-Markovian quantum master equations [115]. Recently, it has been extended to systems with time-dependent Hamiltonian [33], therefore paving the way to using the MCWF method in quantum annealing.

It is useful to isolate the dephasing operator:

$$L_{\alpha,0}(t) = \sum_a \langle E_a(t) | A_\alpha | E_a(t) \rangle | E_a(t) \rangle \langle E_a(t) |. \quad (4.20)$$

For each bath index α , the remaining $D(D-1)$ Lindblad operators can be labeled using an index $h = 1, \dots, D(D-1)$, so that there are $D(D-1) + 1$ Lindblad operators per bath. Excluding accidental degeneracies, each of them is uniquely accompanied by its frequency, which we shall indicate as ω_h . Once the frequency is specified, $\gamma_\alpha(\omega_h)$ and $S_\alpha(\omega_h)$ are real numbers, which we will indicate as $\gamma_{\alpha,h}$ and $S_{\alpha,h}$. If we define $C_{\alpha,h} = \sqrt{\gamma_{\alpha,h}} L_{\alpha,h}$, then the dissipator can be rewritten compactly as

$$\mathcal{D}[\rho_S] = \sum_{\alpha,h} \left(C_{\alpha,h} \rho_S C_{\alpha,h}^\dagger - \frac{1}{2} \{ C_{\alpha,h}^\dagger C_{\alpha,h}, \rho_S \} \right). \quad (4.21)$$

We can rewrite the Lindblad equation as

$$\frac{d\rho_S(t)}{dt} = -i \left(H_{\text{eff}}(t) \rho_S(t) - \rho_S(t) H_{\text{eff}}^\dagger(t) \right) + \sum_{\alpha,h} C_{\alpha,h}(t) \rho_S(t) C_{\alpha,h}^\dagger(t), \quad (4.22)$$

where we defined the non-Hermitian Hamiltonian $H_{\text{eff}}(t)$ as

$$H_{\text{eff}}(t) = H_0(t) + H_{\text{LS}}(t) - \frac{i}{2} \sum_{\alpha,h} C_{\alpha,h}^\dagger(t) C_{\alpha,h}(t), \quad (4.23)$$

diagonal in the energy eigenbasis:

$$(H_{\text{eff}})_{aa} = E_a + \sum_{\alpha,b} \left[S_\alpha(\omega_{ba}) - \frac{i}{2} \gamma_\alpha(\omega_{ba}) \right] |A_\alpha|_{ab}^2. \quad (4.24)$$

Let us focus on the two terms on the right-hand side of equation (4.22). The first term is deterministic and tells us that the reduced system effectively evolves non-unitarily in time in presence of dissipation. The environment shifts the energy levels of the spectrum of the Hamiltonian $H_0(t)$ and introduces a characteristic decay time of the norm of the state of the reduced system. The power of this mapping is that this part of the dynamics can be rephrased as a Schrödinger equation for a ket state $|\psi(t)\rangle$:

$$i \frac{d}{dt} |\psi(t)\rangle = H_{\text{eff}}(t) |\psi(t)\rangle. \quad (4.25)$$

Of course, this first term alone would violate the probability conservation. The second term accounts of equation (4.22) accounts for this violation and restores the correct probability, instantaneously projecting the ket state onto the energy eigenbasis by applying a Lindblad operator on $|\psi(t)\rangle$. The probability of a jump (α, h) is given by $dp_{\alpha,h}(t) = dt \|C_{\alpha,h}(t) |\psi(t)\rangle\|^2$, therefore this term is stochastic. Summarizing, at each time step dt , the ket states evolve non-unitarily with H_{eff} ; however, there is a finite probability that the system may jump to an energy eigenstate, and this projective measurement restores the norm of the states. The jump rates for each Lindblad operator is $\sqrt{\gamma_{\alpha,h}}$. We can distinguish three kinds of processes: (i) decay from a higher state to a lower state, when $\omega_{ba} > 0$; (ii) pump to a higher state, when $\omega_{ba} < 0$; (iii) dephasing, when $\omega_{ba} = 0$.

The update rule at first order in the time step dt for the unnormalized ket state $|\tilde{\psi}(t)\rangle$ reads

$$d|\tilde{\psi}(t)\rangle = -i H_{\text{eff}}(t) dt |\tilde{\psi}(t)\rangle + \sum_{\alpha,h} dN_{\alpha,h}(t) [A_{\alpha,h}(t) - \mathbb{1}] |\tilde{\psi}(t)\rangle. \quad (4.26)$$

The stochastic variable $dN_{\alpha,h}(t)$ is the number of jumps of type (α, h) in the time interval dt . The probability of having more than one jump per time step goes to zero with $dt \rightarrow 0$, therefore

$$dN_{\alpha,h}(t) = \begin{cases} 1 & \text{with probability } dp_{\alpha,h}(t) \\ 0 & \text{with probability } 1 - dp_{\alpha,h}(t). \end{cases} \quad (4.27)$$

Simulating equation (4.26) can be computationally inefficient. This is due to the fact that, at each step, we have to extract a random number to check whether the evolution is deterministic or stochastic. However, in general, $dp(t) = \sum_{\alpha,h} dp_{\alpha,h}(t) \ll 1$, thus this computational time is largely wasted as the evolution is mostly deterministic. In the rare cases where a jump occurs, we then have to extract another random number to determine the index of the Lindblad operator to be applied.

A more efficient algorithm employs the waiting-time distribution [14], exploiting the fact that the squared norm of the unnormalized state at t is equal to the probability of not having jumps until t . We report here the pseudo code for the single quantum trajectory:

1. Draw a random number r uniformly distributed in $[0, 1]$.

2. Starting from a normalized state $|\psi(0)\rangle$, evolve non-unitarily using $H_{\text{eff}}(t)$ until $\langle \tilde{\psi}(t^*) | \tilde{\psi}(t^*) \rangle = r$. At $t = t^*$, a quantum jump occurs.
3. Draw another random number μ uniformly in $[0, 1]$ and select the quantum jump according to the probabilities $dp_{\alpha,h}(t)$. The probability of a jump $k = (\alpha, h)$ is given by $P_k = dp_k(t) / dp(t)$. The index of the occurring jump is the smallest non-negative integer m satisfying $\sum_{k=0}^m P_k \geq \mu$.
4. Update the state as $|\tilde{\psi}(t^* + 0^+)\rangle = C_m |\tilde{\psi}(t^*)\rangle$ and renormalize. Draw another random number r and use $|\psi(t^* + 0^+)\rangle$ as the new normalized starting state.
5. Repeat steps 1–4 until the wanted annealing time.

Once the simulation is complete, averaging over the N_T trajectories returns the correct observables. The Monte Carlo standard error affecting a given observable \mathcal{O} reads

$$\sigma_t^2(\mathcal{O}) = \frac{\lambda_{\mathcal{O}}(D)}{N_T} \sim \frac{1}{N_T}, \quad (4.28)$$

where $\lambda_{\mathcal{O}}(D)$ is the variance

$$\lambda_{\mathcal{O}}(D) = \frac{1}{N_T - 1} \sum_{r=1}^{N_T} (\langle \psi^r(t) | \mathcal{O} | \psi^r(t) \rangle - M_t)^2, \quad (4.29)$$

and $M_t = \sum_{r=1}^{N_T} \langle \psi^r(t) | \mathcal{O} | \psi^r(t) \rangle / N_T$. It can be proven that $\lambda_{\mathcal{O}}(D)$ is a nonincreasing function of D , thus the number of trajectories needed to achieve a certain standard error does not need to be increased with increasing N . In particular, it scales as $\lambda_{\mathcal{O}}(D) = \Lambda_{\mathcal{O}} D^{-x}$, with $0 \leq x \leq 1$, depending on the observable [14].

In order for the first-order update rule of equation (4.26) to be valid, the time step dt must satisfy the following conditions:

$$dt \ll \min_t \left\{ \frac{2 \|H_{\text{eff}}(t)\|}{\|\dot{H}_{\text{eff}}(t)\|}, \frac{1}{\|H_{\text{eff}}(t)\|}, \left| \frac{\lambda(t)}{\lambda^2(t) - \dot{\lambda}(t)} \right| \right\}, \quad (4.30)$$

where $\|\cdot\|$ is the operator norm (e. g., the largest singular value), and $\lambda(t) = \frac{dp(t)}{dt}$. In practice, it is easier to proceed heuristically by decreasing dt until convergence is obtained within the desired tolerance. In particular, the second term on the right-hand side of equation (4.30) implies that the time step must be decreased for increasing N to achieve convergence.

4.4 p -SPIN MODEL AND DISSIPATION

In chapter 3, we introduced the p -spin model as a useful test bed to benchmark quantum annealing. Its main advantage is that its spin symmetry allows us to work in a smaller subspace with dimension $D = N + 1$, simplifying numerical simulations. This simplification is unfortunately lost in the presence of dissipation. In fact, the coupling to an environment generally involves local (single-qubit) operators that break the permutational symmetry of the p -spin model, due to the

fact that each qubit is coupled to its own bath. This model is known as *independent bath* (IB). In this case, numerical simulations are very demanding as we have to represent states and operators in the whole Hilbert space, thus we are limited to systems of ~ 10 qubits, even factoring in the speedup provided by MCWF.

Under certain conditions, however, a different, collective form of dissipation can arise. For example, if the qubit system is coupled to a low-temperature environment with only low-energy modes, then this ensures that only long-wavelength bosonic modes are occupied. If the qubits are spatially close to each other, the bosons will not distinguish them and the permutational symmetry will be preserved. This kind of dissipation is known as *collective bath* (CB) and allows us to perform dissipative simulations for large systems sizes [6, 7].

As a side note, some forms of collective dissipation support decoherence-free subspaces, i. e., subspaces that evolve unitarily despite the coupling to the bath and are one of the reasons why QA is believed to be more robust to decoherence compared to the quantum circuit model [51, 52]. For example, if the p -spin system is subjected to collective dephasing, where the coupling to the environment occurs via the m_z operator, the $S = 0$ subspace is a decoherence-free subspace (for even N). This subspace is unsuitable for quantum computation as the Hamiltonian of the p -spin model preserves it, but could be used to perform quantum computation if we added Heisenberg exchange terms to the p -spin Hamiltonian [7]. We will not explore this possibility in this work.

In chapter 5, we will show many results concerning the dissipative QA of the p -spin model subjected to both forms of dissipation. In the remainder of this section, we will discuss some very simple tests showcasing convergence with N_T . In addition, we will highlight the fact that MCWF can provide information concerning the spectrum of thermal processes affecting quantum dynamics, thus giving an insight into the microscopic effects of the environment on the dynamics of the ground state [33]. In turn, this knowledge can be useful to engineer error correction codes based on control paths that push the system back to its ground state right after thermal excitations, thus counteracting the effect of the bath.

Here, we consider a p -spin system of $N = 8$ qubits with $p = 3$ subjected to a collective bath. The QPU is coupled to the environment via $A = m_z$. We consider standard forward annealing with no pauses and with a linear schedule, i. e., $s(t) = t/t_f$, $A(s) = 1 - s$, and $B(s) = s$. The transverse field is $\Gamma = 1$ GHz. The minimal gap at $s_{\min} = 0.42$ is $\Delta_{\min} = 0.76$. The annealing time is $t_f = 100$ ns. The bath is Ohmic and in equilibrium at $T = 12$ mK, with $\omega_c = 25$ GHz. The dimensionless system-bath coupling strength is $\eta = 1 \times 10^{-4}$. This choice of parameters is consistent with equation (4.10) and ensures that the hypotheses underlying the AME are valid. In particular, the unitary evolution is adiabatic and the ground state fidelity at $t = t_f$ is $p_{\text{gs}}(t_f) = 0.99998$.

In figure 4.3(a), we show a comparison between the dynamics of the ground state probability as described by the adiabatic master equation and those resulting from the Monte Carlo wave function approach. Convergence with respect to dt is achieved already for $dt = 1 \times 10^{-2}$ ns. Instead, even though $N_T = 100$ is already

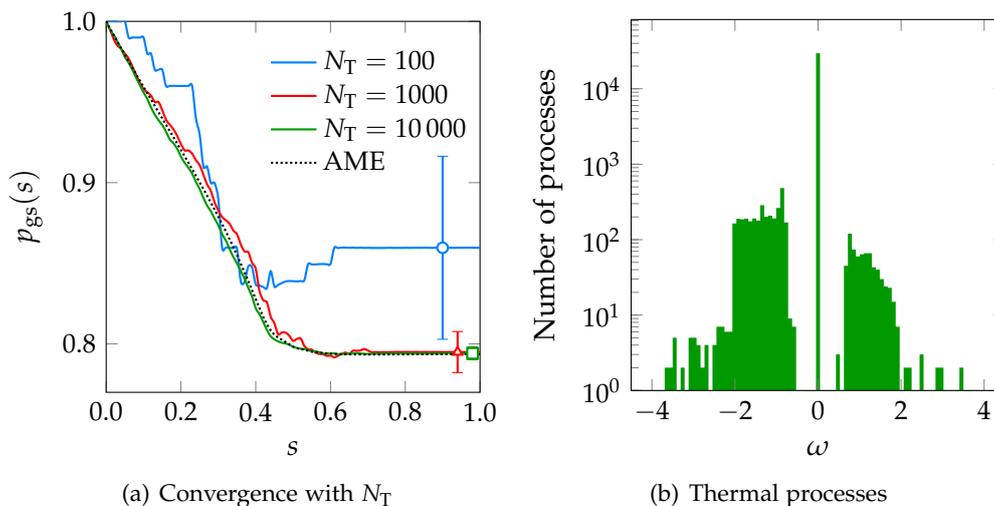


Figure 4.3 | Convergence of MCWF to AME as a function of N_T [panel 4.3(a)] and spectrum of thermal processes for $N_T = 10000$ [panel 4.3(b)]. The coupling operator is $A = m_z$.

loosely in agreement with the master equation solution within Monte Carlo errors (shown only at one point per curve for visual clarity), we need $N_T \approx 1000$ to 10000 to recover a behavior visually similar to the AME solution.

For $N_T = 10000$, we also plot in figure 4.3(b) the spectrum of thermal processes affecting the qubit system, i. e., the number of processes per frequency after binning the frequency interval into $N_b = 1000$ bins. Here, we can see that the m_z -coupling mostly induces dephasing, whereas the total number of thermal excitations and disexcitations with $\omega \neq 0$ is one order of magnitude smaller. This happens because, although the coupling operator is nondiagonal in the instantaneous energy eigenbasis (e. g., where dissipation occurs in the weak coupling limit), its off-diagonal elements are much smaller than its diagonal ones, therefore dephasing is more likely to occur.

In figure 4.4, we repeat the same analysis when the QPU is coupled to the environment via $A = m_x$. Some differences arise due to the different coupling direction. Indeed, the two couplings are complementary to each other as shown in figure 4.5, where we plot the number of processes as a function of time. This is easily understood as we can perform a global rotation around the y -axis in the spin space so that $m_x \rightarrow m_z$ and $m_z \rightarrow -m_x$, where the minus sign does not affect the dynamics since the spectrum is symmetric, thus coupling via $A = m_x$ is equivalent to using $A = m_z$ when performing an annealing in reverse from $s = 1$ to $s = 0$.

When the coupling is along z , the coupling with the environment would induce pumps and decays between energy levels when $s \ll s_{\min}$ because m_z is nondiagonal in the energy eigenbasis. The p -spin system eventually freezes out right after the minimal gap. After the gap, many dephasing processes occur but the only effect is to mix the relative phases of the components over the different eigenstates, thus the ground state probability is left unchanged. On the other hand, when the coupling

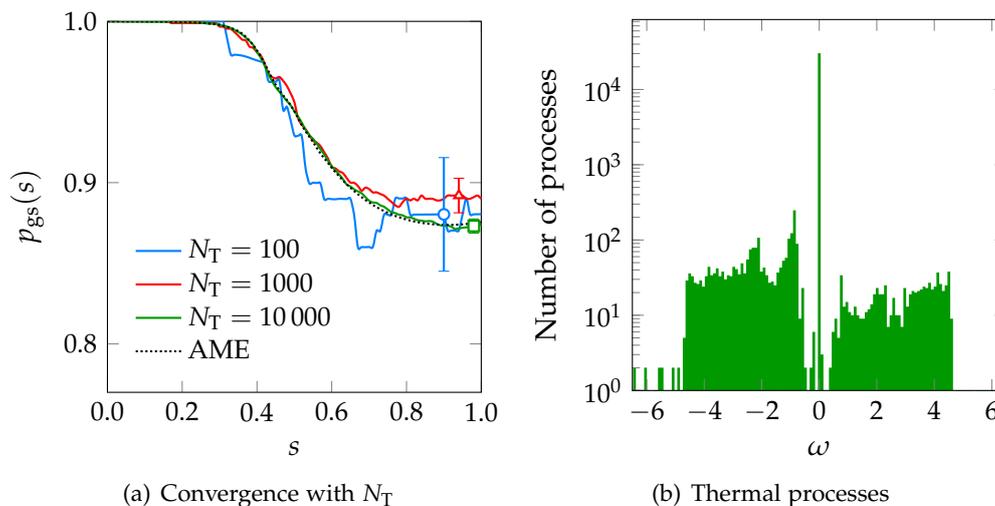


Figure 4.4 | Convergence of MCWF to AME as a function of N_T [panel 4.4(a)] and spectrum of thermal processes for $N_T = 10000$ [panel 4.4(b)]. The coupling operator is $A = m_x$.

is via m_x , the system is unable to leave the starting ground state at the beginning of the dynamics, hence there is a plateau in p_{gs} for s close to zero. The reason is that, for one, the starting state is approximately an eigenstate of the Hamiltonian as $s \ll s_{\min}$, and secondly, the system-bath coupling operator is almost diagonal in the energy eigenbasis as well, thus it can only induce dephasing which, however, has no effect on p_{gs} as only the ground state is populated. After the minimal gap, m_x is no longer diagonal in the energy eigenbasis and the rate of dephasing is substantially reduced. We can still have pumps and decays between energy levels. As shown in figure 4.4(b), the two bands around $\omega = \pm 1$ become larger, even though the number of processes with $\omega \neq 0$ is still one order of magnitude smaller than those at $\omega = 0$. In the following, we will focus on the case $A = m_z$ as it is more likely to occur in experiments [116]. Moreover, we will always consider $N_T = 5000$ as a compromise between CPU time and accuracy. This leads to a relative Monte Carlo error on the fidelity of about 1% [6].

4.5 BEYOND THE BORN-MARKOV APPROXIMATION

The adiabatic master equation and the variant of MCWF presented in section 4.3 are based on the Born, Markov and rotating wave approximation, and are designed to include the effects of the environment on the dynamics of the QPU in a simplified way that can be treated numerically. There are many other approaches similar to those presented in this work that allow one to go beyond the Born-Markov approximation, extending this formalism to non-Markovian systems in the non-weak-coupling limit. In recent years, some of these tools have gained a lot of momentum, such as the *density matrix renormalization group* (DMRG) [117, 118], *quasiadiabatic propagator* (QUAPI) [119], *real-time path-integral Monte Carlo*

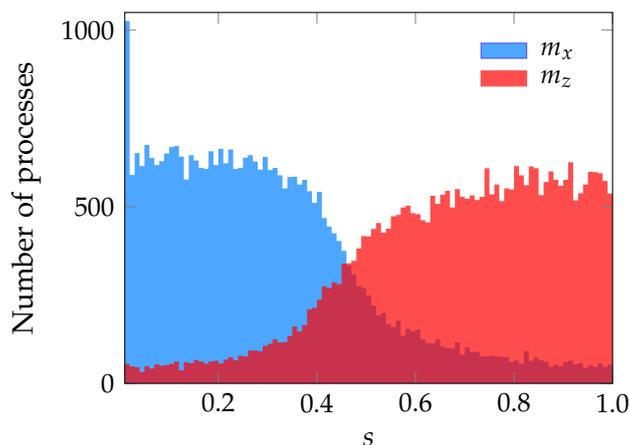


Figure 4.5 | Number of thermal processes as a function of s , for $N_T = 10\,000$.

(PIMC) [120], non-Markovian master equations and quantum jumps [115, 121–123], *hierarchical equations of motion* (HEOM) [124, 125], and many others.

All these methods disregard the back-action of the quantum processor on the environment by tracing away the bath degrees of freedom, so that the environment only acts as a static container that embed the qubit dynamics. This is what happens in thermodynamics, where the thermal reservoir is so large that any recurrence phenomena occur over time scales that are longer than any realistic experiment. However, it may be also worth considering the feedback of the QPU on the environment, in which case the aforementioned techniques are insufficient.

In reference [2], we discussed an alternative technique to account for relaxation and decoherence of open quantum systems beyond the Born-Markov approximation. In this approach, we combined a discretization scheme of the bath density of states, now described as a finite set of harmonic oscillators, a truncation scheme of the bosonic Hilbert space of the bath, and Lanczos propagation of the system + bath ket state. The two main advantages of this technique [*short-iterative Lanczos* (SIL)] are that we have access to the full state of the combined system and that we can include multiple-phonon processes in the description to improve the accuracy in the non-weak-coupling regime, as opposed to the AME, still maintaining a sizable computational effort.

In order to calculate the time evolution operator of the full $S + B$ Hamiltonian of equation (4.9), we start by discretizing the spectrum of bosonic frequencies by considering M equally-spaced modes

$$\omega_k = k \delta\omega = k \frac{\omega_c}{M}, \quad k = 1, 2, \dots, M. \quad (4.31)$$

We consider a hard (Heaviside theta) cutoff in the bath spectral density $J(\omega)$ and also include the coupling energy squared g^2 in its definition. Moreover, we assume that each mode is coupled differently to the qubit system by setting $g \rightarrow g_k$. The couplings g_k that replicate the correct spectral density can be obtained by

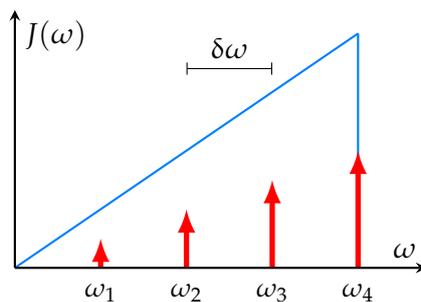


Figure 4.6 | Discretized spectral density of the bath. The heights of the arrows increase with k to represent the strength of the effective couplings [see equation (4.32)].

integrating equation (4.15) around $\omega = \omega_k$, obtaining

$$g_k^2 \approx \eta \omega_k^\nu \omega_c^{1-\nu} \frac{k^\nu}{M^{\nu+1}}. \quad (4.32)$$

This uniform sampling is the simplest one, and different alternatives have already been proposed in the literature [126–129]. A schematic representation is shown in figure 4.6.

After this discretization step, the bath space is spanned by the ket states $\{|n_1, n_2, \dots, n_M\rangle\}$, where $n_k = 0, 1, \dots, N_{\max}$ is the occupation number of the k -th mode and N_{\max} is a cutoff on the excitation number. The creation or annihilation of a phonon leads to a variation Δn_k with respect to the equilibrium value. In order to only keep relevant processes for the dynamics of the qubit system, we propose to truncate the bath Hilbert spaces to those states with $\sum_k |\Delta n_k| = N_{\text{ph}}$, where N_{ph} is the maximum number of allowed excitations. In the weak-coupling regime, it is sufficient to consider $N_{\text{ph}} = 1$. Stronger couplings require $N_{\text{ph}} > 1$, but $N_{\text{ph}} = 3$ is typically enough to describe coupling strengths of the order of $\eta = 1 \times 10^{-2}$ and computations remain affordable in this intermediate-coupling regime. Finite-temperature simulations require taking averages over different possible extractions of the bath starting state according to the Boltzmann distribution.

Once the bath basis is fixed, we can represent the full Hamiltonian of equation (4.9) and diagonalize it using Lanczos in order to evaluate the time evolution operator similarly to equation (B.4). The main limitation of this approach is that the minimum frequency ω_1 is related to the Poincaré recurrence time $T_P = 2\pi/\omega_1$, which in turns determines an upper bound on the annealing time that we can simulate with this method. After T_P , in fact, the discretized collection of harmonic oscillators is no longer a good approximation of the original, ergodic bath. Therefore, this method is unfeasible to study long, adiabatic dynamics, except in the weak-coupling limit, where the dimension of the bath Hilbert space with $N_{\text{ph}} = 1$ scales linearly with M , allowing us to easily simulate a large number of modes and reasonably long annealing times.

4.5.1 Case study 1: single qubit, pure dephasing

We consider a single qubit with time-independent Hamiltonian $H_S = -\epsilon\sigma_z$ linearly coupled to the environment via $A = \sigma_z$, so that the full $S + B$ Hamiltonian reads

$$H = -\epsilon\sigma_z + \sum_k \omega_k b_k^\dagger b_k + \sigma_z \sum_k g_k (b_k + b_k^\dagger). \quad (4.33)$$

This is the pure dephasing model. We use it as a test bed as its Hamiltonian can be diagonalized exactly using the Lang-Firsov transformation [130]

$$V = e^{-S}, \quad S = \sigma_z \sum_k \frac{g_k}{\omega_k} (b_k^\dagger - b_k), \quad (4.34)$$

which shifts the center of oscillations of the harmonic oscillators according to the qubit state. The transformed Hamiltonian reads

$$\tilde{H} = V^\dagger H V = -\epsilon\sigma_z - \sum_k \frac{g_k^2}{\omega_k} + \sum_k \omega_k b_k^\dagger b_k. \quad (4.35)$$

The eigenstates in the original basis are coherent states in the bath degrees of freedom. This Hamiltonian commutes with σ_z , therefore the occupation probabilities of the qubit states are constant.

Conversely, we prepare the system in the tensor state $\rho(0) = |\hat{x}; +\rangle\langle\hat{x}; +| \otimes \rho_B = \rho_S(0) \otimes \rho_B$, where $\sigma_x |\hat{x}; +\rangle = |\hat{x}; +\rangle$ and ρ_B is the thermal equilibrium density operator of the environment:

$$\rho_B = \frac{e^{-\beta H_B}}{\mathcal{Z}_B}, \quad \mathcal{Z}_B = \prod_k \frac{1}{1 - e^{-\beta\omega_k}}. \quad (4.36)$$

We evaluate here the mean value $\langle\sigma_x(t)\rangle$, which, in this simple case, is proportional to the off-diagonal elements of the reduced density matrix in the energy eigenbasis, hence it is a decoherence indicator. The final result is

$$\langle\sigma_x(t)\rangle = \cos(2\epsilon t) e^{-8 \sum_k \frac{g_k^2}{\omega_k} \sin^2\left(\frac{\omega_k t}{2}\right) \coth\left(\frac{\beta\omega_k}{2}\right)} \equiv \cos(2\epsilon t) e^{-\eta K(t,\beta)}, \quad (4.37)$$

where $K(t, \beta)$ is the decoherence function [2].

On the other hand, solving the AME yields the following result [110],

$$\langle\sigma_x(t)\rangle_{\text{AME}} = \cos(2\epsilon t) e^{-2\gamma(0)t}. \quad (4.38)$$

Comparing equations (4.37) and (4.38), we notice that the coherent part of the dynamics is well-captured by the Lindblad approximation. However, the damping factor of the AME is only correct in the adiabatic limit. In fact, using the Ohmic spectral density $J(\omega)$, it is easy to show that the following limit in the space of distributions holds ($\beta < \infty$)

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{K(t, \beta)}{t} &= 8 \lim_{t \rightarrow \infty} \int_0^{\omega_c} \frac{1}{\omega t} \sin^2\left(\frac{\omega t}{2}\right) \coth\left(\frac{\beta\omega}{2}\right) d\omega \\ &= 2 \int_0^{\omega_c} \pi \omega \delta(\omega) \coth\left(\frac{\beta\omega}{2}\right) d\omega \\ &= \frac{4\pi}{\beta} = 2\gamma(0). \end{aligned} \quad (4.39)$$

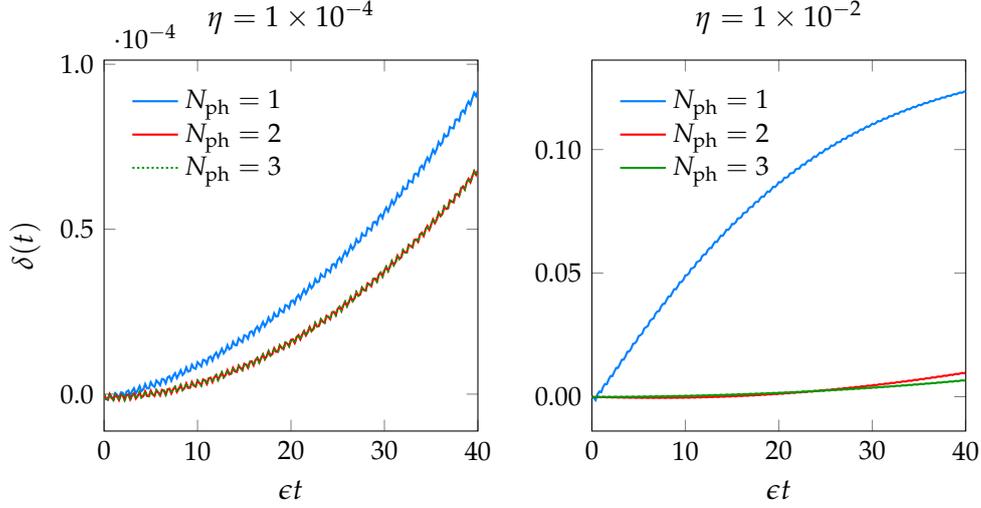


Figure 4.7 | SIL error $\delta(t)$ in the pure dephasing model, for $\eta = 1 \times 10^{-4}$ and $\eta = 1 \times 10^{-2}$.

Besides, the AME result is always incorrect in the zero-temperature limit $\beta \rightarrow \infty$, where the decoherence function can be computed exactly:

$$K(t, \infty) = 4[\gamma - \text{Ci}(\omega_c t) + \log(\omega_c t)], \quad (4.40)$$

where γ is the Euler-Mascheroni constant and $\text{Ci}(x)$ is the cosine integral. This term grows logarithmically in t , whereas any finite-temperature contributions grow linearly in t . Therefore, $K(t, \infty)$ is the relevant part of the decoherence function when the annealing time is shorter than the bath self-correlation time scale $\tau_B = \beta/2\pi$ [32], while the finite-temperature correction is relevant at longer times. The Lindblad approximation does not contain the zero-temperature contribution. However, the bath time scale τ_B diverges when $T \rightarrow 0$, meaning that the Lindblad equation is always inadequate in the zero-temperature limit.

In reference [2], we compared our SIL with the exact result of equation (4.37). We used $\epsilon = 1$ GHz as energy scale and $1/\epsilon$ as unit time. The cutoff frequency has been fixed to $\omega_c = 10$ in these units and we have simulated $M = 200$ bosonic modes. The Poincaré recurrence time is $T_P = 40\pi$, therefore we fixed $t_f = 40 < T_P$. We considered a maximum of $N_{\text{ph}} = 3$ bosonic excitations starting from the equilibrium state at $T = 0$. The Hilbert space dimensions are $D(N_{\text{ph}} = 1) = 402$, $D(N_{\text{ph}} = 2) = 40\,602$, and $D(N_{\text{ph}} = 3) = 2\,747\,402$. We used 30 Lanczos iterations. We simulated $\eta = 1 \times 10^{-4}$ (weak coupling) and $\eta = 1 \times 10^{-2}$ (intermediate coupling), and computed the relative error with respect to the exact solution, e. g.,

$$\delta(t) = \frac{\langle \sigma_x(t) \rangle_{\text{SIL}} - \langle \sigma_x(t) \rangle}{\langle \sigma_x(t) \rangle}. \quad (4.41)$$

In figure 4.7, we plot $\delta(t)$ as a function of t in the analyzed cases. For $\eta = 1 \times 10^{-4}$, setting $N_{\text{ph}} = 1$ is already sufficient to obtain a small error, and $\delta(t)$ saturates when $N_{\text{ph}} = 2$. The increasing trend is solely due to the fact that we are

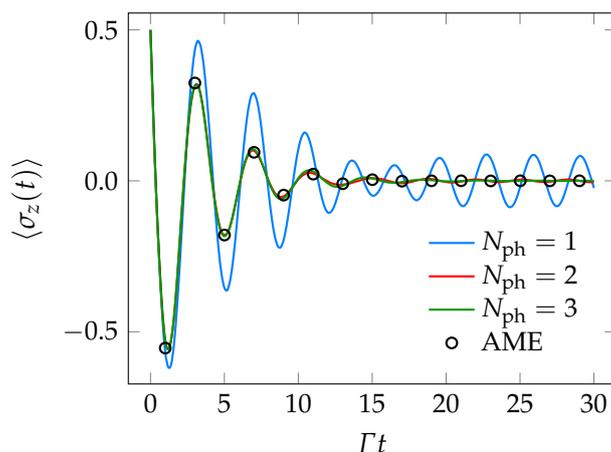


Figure 4.8 | Time evolution of $\langle \sigma_z(t) \rangle$ in the SBM, with $\eta = 5 \times 10^{-2}$, $\omega_c = 10\Gamma$ and $T = 0$. We fixed $M = \{1000, 500, 300\}$, for $N_{\text{ph}} = \{1, 2, 3\}$, respectively. SIL results are compared with the AME (black circles).

approaching T_P and can be softened by increasing the number of modes. On the other hand, for $\eta = 1 \times 10^{-2}$ just one phonon is not enough to recover the correct solution. There is a very small difference between $N_{\text{ph}} = 2$ and $N_{\text{ph}} = 3$, therefore we can conclude that two excitations are enough to obtain the correct solution with a reasonable accuracy.

4.5.2 Case study 2: spin-boson model

The spin-boson model [116, 131] is described by the following Hamiltonian:

$$H = -\Gamma \sigma_x + \sum_k \omega_k b_k^\dagger b_k + \sigma_z \sum_k g_k (b_k + b_k^\dagger). \quad (4.42)$$

The energy scale is $\Gamma = 1$ GHz. No analytic solutions are known for the non-trivial dynamics of $\langle \sigma_z(t) \rangle$ and $\langle \sigma_x(t) \rangle$, however several approximation schemes have been proposed over the last four decades in order to tackle this problem in various windows of system parameters, based on a standard path-integral formulation [116].

We analyze the dynamics of the SBM using SIL at $T = 0$. We consider an Ohmic bath and choose a cutoff frequency $\omega_c = 10\Gamma$. The coupling parameter η is in the range $\eta = 5 \times 10^{-4}$ to 1×10^{-1} . In addition, following this choice of parameters, we perform the basis truncation including up to $N_{\text{ph}} = 3$ excitations per mode. We prepare the system at time $t = 0$ in a linear combination of the basis states at fixed starting values $\langle \sigma_x(0) \rangle = \langle \sigma_z(0) \rangle = 1/2$, i. e., $|\psi(0)\rangle = \cos(\xi/2) |0\rangle + \sin(\xi/2) \exp(i\phi) |1\rangle$, with $\xi = \pi/3$ and $\phi = \text{acos}(1/\sqrt{3})$. Then, we calculate the time-evolved mean values $\langle \sigma_x(t) \rangle$ and $\langle \sigma_z(t) \rangle$ and compare them with their analytic closed-form counterparts obtained from the AME.

In figure 4.8, we show the results for $\langle \sigma_z(t) \rangle$ for $\eta = 5 \times 10^{-2}$, compared with the AME, for $t_f = 30$. Choosing a minimum value of $N_{\text{ph}} = 2$, the time evolution of $\langle \sigma_z(t) \rangle$ converges to the exact physical behavior, featuring underdamped

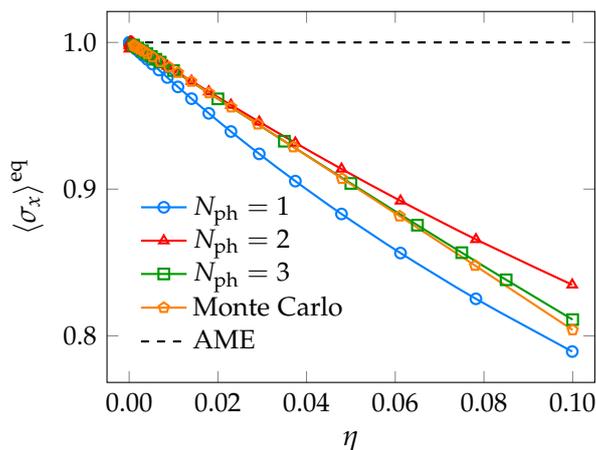


Figure 4.9 | Equilibrium value $\langle \sigma_x \rangle^{\text{eq}}$, extrapolated with exponential fits, as a function of the dimensionless coupling η , at $T = 0$, for an Ohmic bath. The simulated data of the numerical diagonalization, up to $N_{\text{ph}} = 3$ bosonic excitations from the vacuum state, are compared to the AME and Monte Carlo predictions at equilibrium.

oscillations due to decoherence. At long times, the equilibrium value $\langle \sigma_z \rangle^{\text{eq}} = 0$ is reached and the system completely loses its coherence. The Lindblad AME is able to capture the behavior of quantum coherence accurately in this setting.

On the other hand, $\langle \sigma_x(t) \rangle$ is a measure of thermal relaxation. The AME predicts that at long annealing times the system reaches equilibrium at temperature $T = 0$, therefore $\langle \sigma_x \rangle_{\text{AME}}^{\text{eq}} \rightarrow 1$, which is correct only in the weak-coupling regime [110]. Instead, our calculations show that the stationary value is related to the ground state of the qubit-bath system: at long times, qubit and bath remain entangled, as expected at equilibrium. Therefore, we can plot the equilibrium values $\langle \sigma_x \rangle^{\text{eq}}$ as a function of the coupling parameter η compared with the Lindblad result. These values have been obtained by performing an exponential fit of $\langle \sigma_x(t) \rangle$ to extract the saturation value. In figure 4.9, we show the fitted equilibrium values $\langle \sigma_x \rangle^{\text{eq}}$ as a function of the coupling strength η compared with the Lindblad result. In order to further test the reliability of our calculations, we also plot the equilibrium values calculated using a Monte Carlo approach at thermal equilibrium, whose errors are smaller than the point size. We note that, as the coupling factor becomes larger than $\eta = 1 \times 10^{-3}$, the Born-Markov approach misses the correct physical behavior. It follows that, at long times, the unavoidable system-bath entanglement effects start to play a role, noticeably reducing the value of $\langle \sigma_x \rangle^{\text{eq}}$. A good physical description up to $\eta = 1 \times 10^{-1}$ can be achieved by truncating the phonon basis to three excitations per mode.

4.5.3 Case study 3: single qubit, quantum annealing

Finally, we consider the QA Hamiltonian with a linear schedule

$$H_S(s) = -(1-s)\Gamma\sigma_x - s\epsilon\sigma_z. \quad (4.43)$$

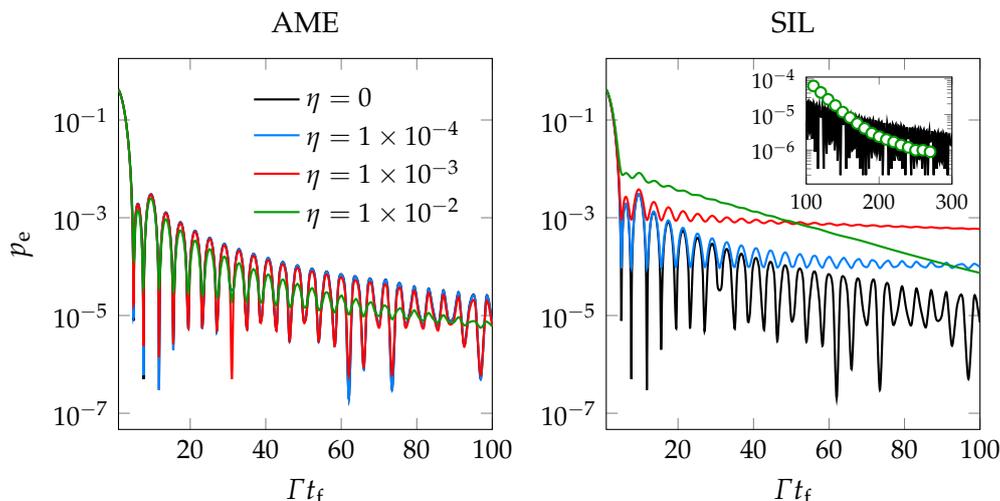


Figure 4.10 | Error probability p_e as a function of the annealing time t_f for a single qubit with the Hamiltonian of equation (4.43), coupled via σ_z to a zero-temperature environment with several coupling constants η . The left panel is AME and the right panel is SIL. In the inset, we plot the unitary curve and some points relative to $\eta = 1 \times 10^{-2}$ at longer annealing times.

The transverse field strength is the energy scale ($\Gamma = 1$ GHz) and we consider $\epsilon = \Gamma$. We prepare the system in the $+1$ eigenstate of σ_x and the environment in the vacuum state at $T = 0$. The coupling operator is $A = \sigma_z$. The AME for the Hamiltonian of equation (4.43) can be solved exactly in the adiabatic limit (see for instance reference [110]). In particular, the Lindblad equation predicts that, when $t_f \gg \Delta_{\min}^{-1}$, the ground state fidelity at $T = 0$ saturates at $p_{gs}(1) = 1$, hence the only effect of the environment is to drive the qubit towards its zero-temperature equilibrium state. However, as already discussed in section 4.5.2, this picture is short-sighted as it does not consider effects beyond the Born-Markov approximation, especially arising at low temperatures.

Here, we consider $M = 200$ bosonic modes and a maximum of $N_{\text{ph}} = 3$ excitations per mode, with $\omega_c = 10$. In figure 4.10, we plot the error probability $p_e = 1 - p_{gs}$ as a function of the annealing time, for several values of the system-bath coupling strength η . The left panel shows the results obtained by using the AME, whereas the right panel is the SIL method. The two panels suggest once again that the AME is inadequate at $T = 0$, where the system-bath correlations play a role and modify the error probability of quantum annealing. In particular, in the AME approach, the sole effect of the environment is the damping of the nonadiabatic oscillations seen in the unitary case. By contrast, SIL is in agreement with the Lindblad approach at very short times, but at intermediate times the system-bath interaction tends to increase the error probability relative to the isolated case. The dissipative curves show a transient plateau at intermediate times, that anticipates a decrease of the error probability towards the unitary (and AME) curve at adiabatic times (see the inset in the right panel, obtained by considering $M = 450$ bosonic modes to increase the recurrence time T_P). The plateau exhaustion time inversely

depends on η , thus we can observe a nonmonotonic behavior of the error probability with η in the time window where this decrease takes place.

For these reasons, in the following, we will always consider $T > 0$ when using the AME of MCWF, being sure to choose every other parameter within a range where the Born, Markov, and rotating wave approximation are valid. However, as we saw that at short times the AME dynamics is mostly in agreement with the results of the exact diagonalization, in the following we will push the limit of applicability of the AME by extending it also to nonadiabatic times, asking ourselves what the predictions of the Lindblad equation are, irrespective of its derivation.

4.6 CONCLUSIONS

In this chapter, we have discussed a number of approximations to the dynamics of open quantum systems, which will allow us to study numerically the dynamics of quantum annealers. In the next chapter, we will use these numerical tools to study different kinds of quantum annealing: standard, iterated reverse, and adiabatic reverse annealing. Our aim is to understand how the performance of the different paradigms of adiabatic quantum computation is modified in the presence of dissipation, and whether or not there are cases in which the environment can improve the success probability of the algorithm, yielding higher fidelities compared to the isolated case. We shall focus on the ferromagnetic p -spin system as a toy model, easy to simulate numerically, that might give us an insight into the mechanisms regulating dissipative quantum annealing.

RESULTS

5

5.1 INTRODUCTION

In this chapter, we are going to summarize some of our personal contributions to the field of AQC with dissipation. To this end, we are going to use the numerical tools described in the previous section, focusing on the ferromagnetic p -spin model. This discussion is based on references [4, 6–8].

In particular, this chapter is organized as follows. In section 5.2, we will discuss the standard quantum annealing of the p -spin model with and without dissipation, and compare its performances with simulated annealing. We will show that, although thermal annealing is more efficient than QA in the unitary limit, the situation can change thanks to the coupling to a low-temperature environment. In section 5.3, we will add pauses to the annealing and show that fully-connected models can benefit from pauses inserted around an optimal pausing point, as previously observed experimentally for sparse Ising models. In section 5.4, we will move to reverse annealing and demonstrate the feasibility of this advanced schedule to refine an already available trial solution. Finally, in section 5.5, we will discuss adiabatic reverse annealing.

5.2 QUANTUM ANNEALING

5.2.1 The unitary limit

In this section, we describe the QA of the p -spin system, described by the Hamiltonian of equation (3.4) and reported here for convenience:

$$H_0(s) = -Ns(m^z)^p - N\Gamma(1-s)m^x. \quad (5.1)$$

At $t = 0$, the system is prepared in the paramagnetic ground state of the transverse field Hamiltonian and then is slowly driven towards the ground state of the p -spin Hamiltonian.

As discussed in appendix A.1 and section 3.5, for $p = 2$ the p -spin model undergoes a second-order quantum phase transition where the minimal gap closes as $\Delta_{\min} \sim N^{-1/3}$ in the thermodynamic limit, whereas for $p > 2$ the QPT is of first-

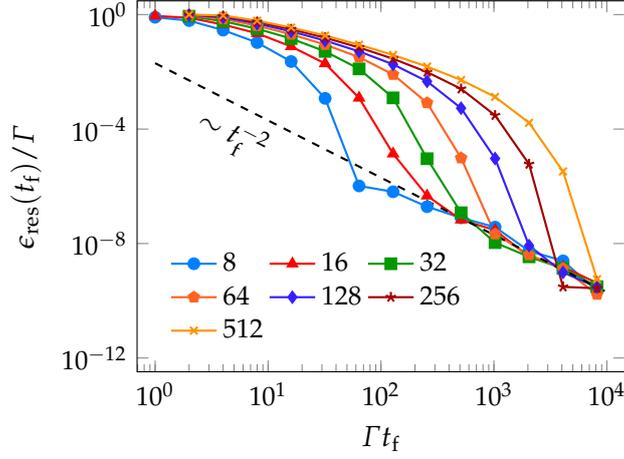


Figure 5.1 | Residual energy as a function of the annealing time, for several values of the system size N , in the unitary case of $p = 2$. The dashed line denotes the asymptotic behavior $\epsilon_{\text{res}} \sim t_f^{-2}$.

order and the minimal gap vanishes exponentially with N , with a rate depending on p [see equation (3.18)].

When p is even, the ground state is doubly degenerate and the fidelity, i. e., the success probability of the adiabatic algorithm, must be defined as the sum of the populations of the two degenerate ground states. Alternatively, a measure of the accuracy of QA is given by the excess or residual energy, which is the difference between the exact ground-state energy and the mean value of the Hamiltonian H_p at $t = t_f$. Energy is an extensive quantity, therefore we normalize it to the number of qubits N so as to have a way to reliably compare results relative to different system sizes:

$$\epsilon_{\text{res}} = \frac{\langle \psi(t_f) | H_p | \psi(t_f) \rangle - E_0(t_f)}{N}. \quad (5.2)$$

The residual energy depends on both t_f and N . In this section, we show the scaling of the residual energy with t_f for several values of N , which reflects the dependence of the gap with the system size. Units are chosen so that the transverse field is $\Gamma = 1$ GHz. In figure 5.1, we show the scaling of ϵ_{res} as a function of t_f for $p = 2$.

We observe three different regimes at short, intermediate, and long annealing times. In the first one, the system is trapped in the starting paramagnetic state as the annealing time is too short to follow the ground state across the critical point, therefore the residual energy is almost constant and close to one. In the long-time region, ϵ_{res} scales as t_f^{-2} as a consequence of the adiabatic theorem. The intermediate regime is governed by Landau-Zener transitions. Here, the residual energy scales as

$$\epsilon_{\text{res}}(t_f) = \frac{C}{N} e^{-t_f/\tau_N}, \quad (5.3)$$

where C is a dimensional constant and τ_N is proportional to Δ_{min}^{-2} and thus to $N^{2/3}$. Therefore, the larger is N , the longer is the adiabatic timescale [34].

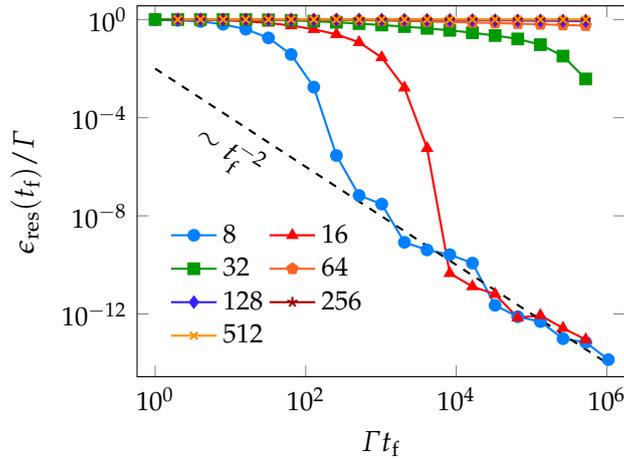


Figure 5.2 | Residual energy as a function of the annealing time, for several values of the system size N , in the unitary case of $p = 5$.

In figure 5.2, we show the residual energy as a function of t_f in the case of $p = 5$. The first and third regimes and relative explanations are similar to those of the case of $p = 2$. Conversely, in the intermediate-time regime, the characteristic time of LZ transitions increases exponentially in N following the scaling law of the gap as a function of the system size. Thus, we can only observe the adiabatic regime for $N = 8$ and $N = 16$ in the analyzed time window.

5.2.2 Quantum annealing with dissipation

We now use of the tools discussed in chapter 4 to study QA with dissipation. The dynamics of the open quantum p -spin system are derived within the Born, Markov, and rotating-wave approximation, and unveiled either by solving the AME directly or via MCWF. Here, we focus on a system of $N = 8$ qubits with $p = 5$, collectively coupled to the an environment of harmonic oscillators via the total spin operator $A = m_z$. The energy scale is the transverse field $\Gamma = 1$ GHz. The dimensionless coupling constant is varied between $\eta = 1 \times 10^{-4}$ and $\eta = 1 \times 10^{-1}$. The cutoff frequency in the Ohmic spectral function is $\omega_c = 25$ and we study two different temperatures: $\beta = 2$, corresponding to $T = 4$ mK in our units, and $\beta = 10$ ($T = 0.75$ mK). Despite the fact that these temperatures are currently beyond the reach of the available quantum devices, at these low temperatures some very interesting effects can arise. For both these temperatures, the validity conditions for the dynamical semigroup approximation [see equation (4.10)] are satisfied for $t_f \gtrsim 50/\Gamma$. Nonetheless, we will still use the adiabatic master equation at shorter annealing times, where we have shown in section 4.5.3 that the Lindblad equation gives results in agreement with our numerically exact short-iterative Lanczos method.

In figure 5.3, we show the scaling of the residual energy as a function of the annealing time, compared to the isolated case. We notice that the effect of the bath is generally negligible at small t_f and becomes relevant at longer t_f . Except for very

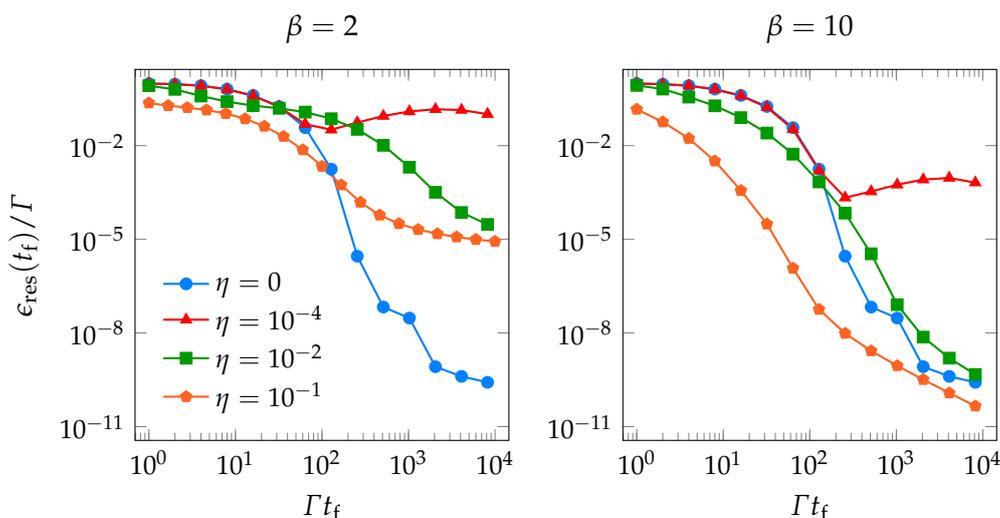


Figure 5.3 | Residual energy as a function of t_f in the case of $p = 5$, comparison between unitary and dissipative dynamics. The left-hand panel is for $\beta = 2$ and the right-hand panel is for $\beta = 10$.

short times, where the accuracy of the AME is questionable, for $\beta = 2$ the unitary dynamics always yield a smaller residual energy compared to the dissipative curves. On the other hand, at the lowest temperature $\beta = 10$ the scenario becomes more interesting. From the right-hand panel of figure 5.3, we indeed see that in the weak-coupling regime $\eta = 1 \times 10^{-4}$ the bath has always a detrimental effect on the annealing performance, but, unexpectedly enough, for larger values of η this picture changes radically. In particular, for $\eta = 1 \times 10^{-2}$, the residual energy is smaller than that of the closed system at short times, whereas for $\eta = 1 \times 10^{-1}$ it is always smaller independently of t_f . This scenario is very similar to the partial thermal speed up discussed in section 4.5.3 for the single qubit case.

These results show that the velocity of convergence to the target ground state at low temperatures is strongly influenced by the system-bath coupling strength, in particular a stronger coupling is responsible for a faster decrease of the residual energy as a function of the annealing time. This speed up can either be quantum or thermal in nature, or a combination of both. We claim that this is indeed a quantum effect because it occurs also at the zero-temperature limit $\beta \rightarrow \infty$, as shown in figure 5.4, and is probably due to the formation of an entangled system-bath state that favors the driving of the reduced system towards its ground state.

Finally, we compare quantum annealing with simulated annealing. In order to perform SA, we linearly decrease the system temperature from the initial temperature $T_0 = 2$ GHz (in units in which $k_B = 1$) to a final temperature T_f in a time t_f . In order to make a fair comparison between SA and QA, we set the final temperature of SA as the equilibrium temperature of the bosonic bath in QA, i. e., $T_f = 1/\beta$. The classical p -spin system is subjected to a *classical phase transition* (CPT) as a function of the temperature, which can either be of second order (for $p = 2$) or first order (for $p > 2$). The initial temperature is chosen to be larger than the critical

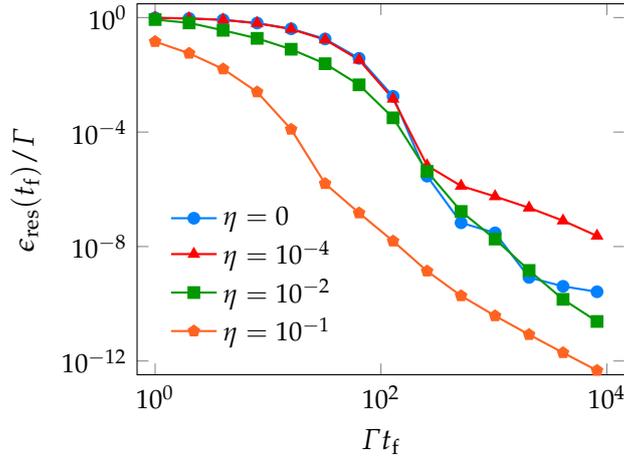


Figure 5.4 | Residual energy as a function of t_f in the case of $p = 5$, comparison between unitary and dissipative dynamics at $T = 0$.

temperature T_c , so as to start from the paramagnetic phase and then move towards the ferromagnetic phase, similarly to the QA procedure.

Thermal annealing simulations are carried out using Glauber's master equation for the probability $P(m, t)$ of observing a magnetization m at time t :

$$\begin{aligned} \frac{\partial P(m, t)}{\partial t} = & \frac{N}{2} \sum_{\alpha=\pm} \left(1 + \alpha m + \frac{2}{N}\right) W_{m, m+2\alpha/N} P\left(m - \alpha \frac{2}{N}\right) \\ & - \frac{N}{2} \sum_{\alpha=\pm} (1 + \alpha m) W_{m-2\alpha/N, m} P(m, t). \end{aligned} \quad (5.4)$$

The elements $W_{m, m\pm 2/N}$ are the rates for a single spin-flip, which we choose in the heat bath form:

$$W_{a,b} = \frac{e^{-\beta \Delta E_{ab}/2}}{e^{-\beta \Delta E_{ab}/2} + e^{\beta \Delta E_{ab}/2}}. \quad (5.5)$$

There are four terms in the right-hand side of equation (5.4): the first two increase the probability $P(m, t)$ because of transitions from the states with a magnetization that differs of $\pm 2/N$ from m ; the remaining two terms represent the inverse processes.

The authors of reference [34] have shown that for $p = 2$ SA outperforms quantum annealing due to a combination of two factors, i. e., the residual energy in SA scaling exponentially in t_f and it being independent of the system size N . On the other hand, for $p > 2$ the residual energy in simulated annealing is no longer size-independent. Therefore, we here focus on the two cases $p = 5$ and $p = 7$, where there can be a competition between quantum and simulated annealing in terms of performance. We consider $\beta = 10$. Our results are summarized in figure 5.5.

At such a low temperature, SA is expected to perform better than QA, as the adiabatic theorem predicts a t_f^{-2} -scaling of ϵ_{res} in the quantum case, as opposed to the t_f^{-1} asymptotic behavior in simulated annealing [34]. However, the time at

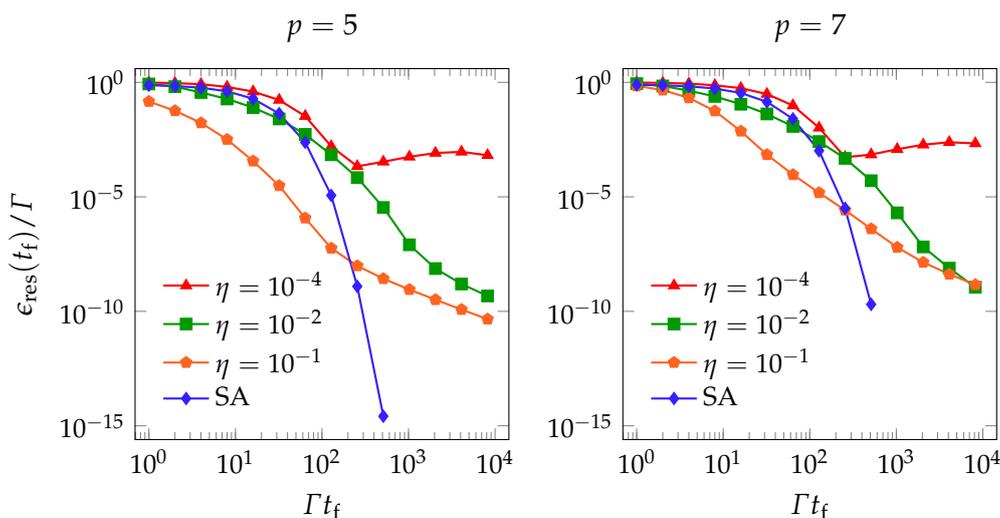


Figure 5.5 | Residual energy as a function of t_f , comparison between dissipative quantum annealing and simulated annealing. The left-hand panel is for $p = 5$ and the right-hand panel is for $p = 7$. The final temperature of SA is $T_f = 1/\beta$ with $\beta = 10$.

which the system moves from the Landau-Zener regime to the adiabatic regime is an exponential function of N , which for macroscopic systems is reached only at impractically long annealing times. Therefore, the performance of the two techniques must be compared in the intermediate- t_f regime, where it is not immediate to see which technique performs best.

In particular, what we see is that, in the intermediate-time regime, QA performs better than SA when the system-bath coupling strength is not too weak. The time at which simulated annealing starts to outperform its quantum counterpart seems to be proportional to the exponent p , as inferred from the comparison between the left- and right-hand panels of figure 5.5 and with the case of $p = 3$ (not shown).

This could suggest that, for very large values of p , quantum annealing could perform better than simulated annealing, in an accessible time window in the presence of a realistic (i. e., not extremely weak) coupling to the environment. This effect is probably not due to thermal fluctuations, but rather arises because of a renormalization of the quantum p -spin Hamiltonian for the effect of the bath. For $p \rightarrow \infty$, this may indicate that quantum annealing is more efficient than thermal annealing when studying the Grover's limit (see section 3.6).

5.3 QUANTUM ANNEALING WITH PAUSES

As anticipated in section 2.5.1, the success probability of quantum annealing can be improved significantly by pausing the dynamics, thanks to the environment favoring thermal relaxation towards the target ground state. The advanced schedule mechanics of D-Wave machines have allowed the scientific community to test this claim experimentally in various settings, all involving short-range sparse Ising models that can be embedded in the Chimera architecture of current-generation

quantum annealers. Our personal contribution to the field was to demonstrate that this effect, rather than being problem-dependent, is quite general and found also in fully-connected models such as the p -spin system. Not only do we observe a similar effect in a model that is radically different from the ones that can be studied experimentally, but we are also able to recover the fidelity enhancement due to pauses even with the simplest model of dissipation at hand, i.e., the Markovian AME. From another point of view, these results indicate once more that a Markovian approximation is sufficient to describe, at least qualitatively, the behavior of quantum annealers [132].

Here, we focus on a system of $N = 20$ qubits with $p = 19$. All energies are expressed as multiples of $\Gamma = 1$ GHz and all times are multiples of Γ^{-1} . In order to make our simulations more similar to what we would observe should we use the D-Wave device, we rescale the annealing Hamiltonian of equation (2.2) by a factor 2 and consider the experimentally-realistic annealing schedules of the D-Wave 2000Q depicted in figure 2.3, with $A(0) = 30$ and $B(1) = 40$.

There are two main reasons why we decided to focus on $p = 19$. First of all, we have shown in the previous section that, if a thermal speed up exists, we should be able to detect it at short and intermediate annealing times compared with the adiabatic time scale. Given the energy scales of the annealing schedules and of the p -spin model Hamiltonian, this choice of p results in a minimal gap of $\Delta_{\min} = 0.14$, at $s_{\Delta} = 0.334$, corresponding to an adiabatic time scale $t_{\text{f}}^{\text{ad}} \sim 1 \times 10^4$ according to equation (2.7). In the following, we will choose $t_{\text{f}} = 100$ as annealing time, so that we are far from the adiabatic regime and, at the same time, computations remain affordable. With this choice of t_{f} , the fidelity $F = p_{\text{gs}}(t = t_{\text{f}})^a$ in the unitary limit is small ($F \approx 5.51 \times 10^{-3}$).

The second reason is more conceptual and has to do with the large- p limit of the p -spin model. In fact, we can easily see that for $p = 19$ the spectrum of the p -spin system indeed closely resembles the one of the adiabatic Grover's search, as opposed to smaller values of p . For instance, in figure 5.6 we plot the eigenvalues of the p -spin model compared to those of the adiabatic search Hamiltonian of equation (3.21) rescaled by a factor $N/2$ to match energy scales, as a function of the Hamming distance from the target ground state. The Hamming distance d_{H} between two states is defined as the number of spin flips connecting them. In the case of the p -spin or search Hamiltonian, where the target state is ferromagnetic, the Hamming distance counts the number of spins misaligned with respect to the true ground state and is related to the magnetization via the rule $d_{\text{H}} = N(1 - m)/2$. By analyzing figure 5.6 we realize that the p -spin model with $p = 3$ and $p = 7$ has quite a different spectrum than that of the search Hamiltonian, whereas, in the case of $p = 19$, the eigenvalues are nonzero only close to the target state and are similar to the search ones. As opposed to the search Hamiltonian, the p -spin model features a highly excited eigenstate with energy $E_{N+1} = -E_0$. However, this state does not participate in the dynamics as it is unlikely to be coupled to the low-energy subspace.

^a $F = p_{\text{gs}}(t = t'_{\text{f}})$ if there is a pause, with $t'_{\text{f}} = t_{\text{f}} + t_{\text{p}}$.

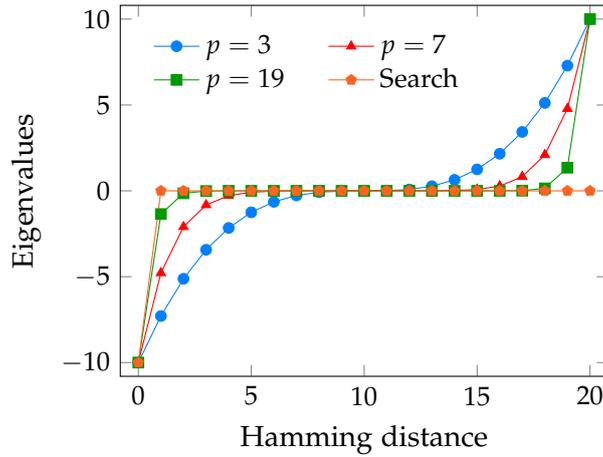


Figure 5.6 | Eigenvalues of H_p , dimensionless, as a function of the Hamming weight of the corresponding eigenstates, for $p = 3$, $p = 7$, and $p = 19$. They are compared with the eigenvalues of the adiabatic search Hamiltonian of equation (3.21). The lines serve to guide the eye.

We collectively couple the p -spin system via the operator $A = m_z$ to a bath of harmonic oscillators at $T = 12 \text{ mK} = 1.57 \text{ GHz}$. This is one order of magnitude larger than the minimal gap, therefore we expect thermal effects to play an important role for this instance. The coupling constant is $\eta = 1 \times 10^{-3}$ and we employ a cutoff frequency of $\omega_c = 1 \text{ THz}$. These parameters allow us to adopt the Born-Markov approximation and use the adiabatic master equation to study dynamics. For convenience, we use the MCWF unraveling with $K = 5000$ trajectories as explained in section 4.3.

First of all, in figure 5.7 we show the dynamics of the ground state probability $p_{\text{gs}}(s)$ for a standard quantum annealing with no pauses. The blue line is the unitary case $\eta = 0$, while the red line corresponds to the dissipative simulation with $\eta = 1 \times 10^{-3}$. In the unitary case, Landau-Zener transitions due to the fast dynamics compared with the adiabatic timescale are responsible for the abrupt decrease of the ground state probability at the avoided crossing. After $s = s_\Delta$, the energy levels quickly separate and the population remains constant, yielding $F = 5.51 \times 10^{-3}$. In the dissipative case, on the other hand, the bath acts around the gap and partly compensates the diabatic excitations by favoring thermal relaxation towards the ground state. The net result is an increase in the ground state probability and hence the fidelity is substantially larger ($F = 0.799$). Once again, far from $s = s_\Delta$, the level spacing is too large for the bath to have an effect on the dynamics.

Right before $s = s_\Delta$, we observe a small decrease in $p_{\text{gs}}(s)$ in the dissipative case that can be easily explained using a two-level picture. Let us consider the ground state and the first excited state, and focus on the region $s \in [s_T, s_\Delta]$, where s_T is the time before s_Δ where the instantaneous gap is equal to the temperature T . At shorter times, the thermal processes have scarce influence on the dynamics. At $s = s_T$, the system is almost completely in its ground state. As our simplifying hypothesis, we suppose that for $s \in [s_T, s_\Delta]$ the gap is constant and equal to Δ_{min} . In

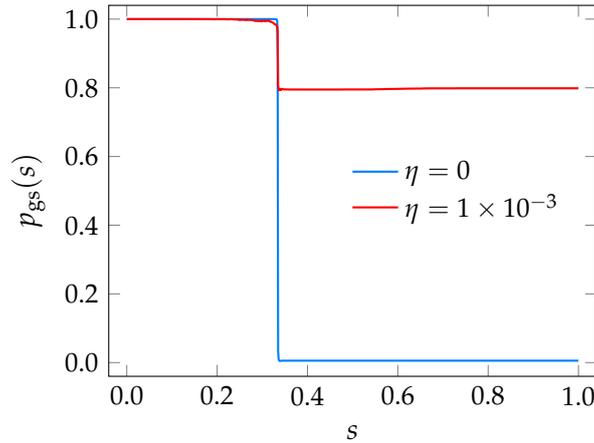


Figure 5.7 | Dynamics of the ground state probability without any pauses, unitary and dissipative evolution. See the main text for details about the parameters of the simulation.

the energy eigenbasis, the density matrix of the qubit system at $s = s_T$ is diagonal as $B(s \leq s_T) \ll A(s \leq s_T)$, and can be written as

$$\rho(s_T) = \begin{pmatrix} \rho_{11}(s_T) & 0 \\ 0 & \rho_{22}(s_T) \end{pmatrix}, \quad (5.6)$$

where $\rho_{11}(s_T) \approx 1$ and $\rho_{22}(s_T) = 1 - \rho_{11}(s_T) \approx 0$. The population transfer due to thermal processes can be effectively modeled by a classical master equation of the form

$$\frac{1}{t_f} \frac{d\rho_{11}(s)}{ds} = \Gamma_{2 \rightarrow 1} \rho_{22}(s) - \Gamma_{1 \rightarrow 2} \rho_{11}(s), \quad (5.7)$$

where $\Gamma_{1 \rightarrow 2} = \gamma(-\Delta_{\min}) = e^{-\beta \Delta_{\min}} \gamma(\Delta_{\min})$ and $\Gamma_{2 \rightarrow 1} = \gamma(\Delta_{\min})$. The solution to this master equation reads

$$\rho_{11}(s) = \rho_{11}(s_T) - C[1 - e^{-(s-s_T)/s_1}], \quad (5.8)$$

where $C = \rho_{11}(s_T) - t_f s_1 \gamma(\Delta_{\min})$ and $s_1 = 1/t_f \sqrt{\gamma(\Delta_{\min})} (1 + e^{-\beta \Delta_{\min}})$. The ground state population at $s \rightarrow s_{\Delta}^-$ is then $\rho_{11}(s_{\Delta}^-) \approx 0.975$, in agreement with numerical simulations.

Thermal effects can be enhanced by pausing the quantum annealing. We tested various pause lengths t_p , from 100 to 900. The maximum pause length t_p gives a total annealing time $t'_f = t_f + t_p = 10t_f$, whose corresponding fidelity, in the absence of pauses, is $F = 5.39 \times 10^{-2}$ for the unitary dynamics and $F = 0.664$ for the dissipative one. The effects of pausing substantially increase this fidelity. These results are t_p -independent, although there are small quantitative differences. In figure 5.8, we show the fidelity F as a function of the pause time $s_p \in [0, 1]$, for different values of the pause duration, $t_p = 100, 400$ and 900 . The right-hand panel is a zoom in the region around the optimal pausing time $s_p^{\text{opt}} \approx 0.55$. Four different regions can be distinguished for all t_p .

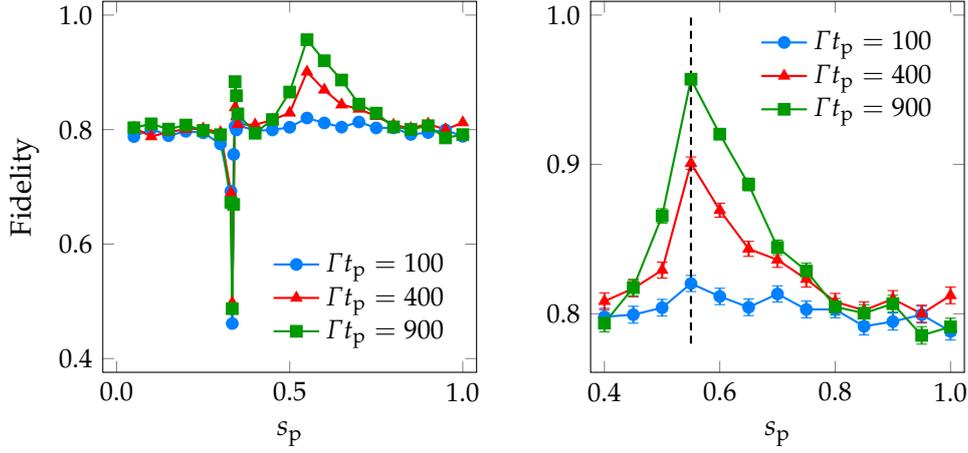


Figure 5.8 | Fidelity as a function of the pausing time s_p , for several values of the pause length t_p . The right-hand panel zooms in on the region around the optimal pausing point $s_p^{\text{opt}} = 0.55$, highlighted by the vertical dashed line. The optimal pausing point is t_p -independent. Error bars represent MCWF errors [see equation (4.28)].

1. When $s_p < s_\Delta$, the fidelity is not affected by pausing. Here, the tunneling amplitude is large compared with the thermal relaxation rate, and the system evolves quantum mechanically with negligible influence from the environment.
2. When $s_p \approx s_\Delta$, thermal processes are more frequent and, correspondingly, the relaxation rate is maximum. In this region, $\exp(-\beta\Delta_{\text{min}}) \sim 1$, hence the excitation rate $\Gamma_{1 \rightarrow 2}$ is comparable with the decay rate $\Gamma_{2 \rightarrow 1}$. When $s_p \lesssim s_\Delta$, most of the population is in the adiabatic ground state. Thus, transitions from the ground state to the first excited state are more probable than reverse processes. This imbalance causes a decrease in the observed fidelity after the pause, and the worst pausing time is sharply localized right before s_Δ . A recent paper proves that this effect is likely to be an artifact of the Markovian approximation and is not typically found in experimental data [132]. The situation is reversed for $s_p \gtrsim s_\Delta$, and here the fidelity is slightly enhanced. This effect is more pronounced for longer t_p .
3. For $s_p = s_p^{\text{opt}}$, we observe a peak in the success probability for any $t_p > 100$. The peak height increases with increasing t_p , following a saturation law of the form

$$F(t_p) = F_{\text{sat}} \left[1 - \alpha e^{-(t_p - l_0)/T_r} \right], \quad (5.9)$$

with fitted parameters $F_{\text{sat}} = 0.976 \pm 0.007$, $\alpha = 0.160 \pm 0.005$, and $T_r = 4.1 \text{ ns} \pm 0.4 \text{ ns}$, and l_0 fixed to $l_0 = 100 \text{ ns}$. In particular, T_r is related to the thermal relaxation time of the many-body system, and F_{sat} is an estimate of the maximum fidelity that can be achieved by pausing the dynamics at the optimal point. The fidelity shows a peak almost at the time $s_p = s_p^{\text{opt}}$ independently of the pause duration t_p . By contrast, t_p influences the time at

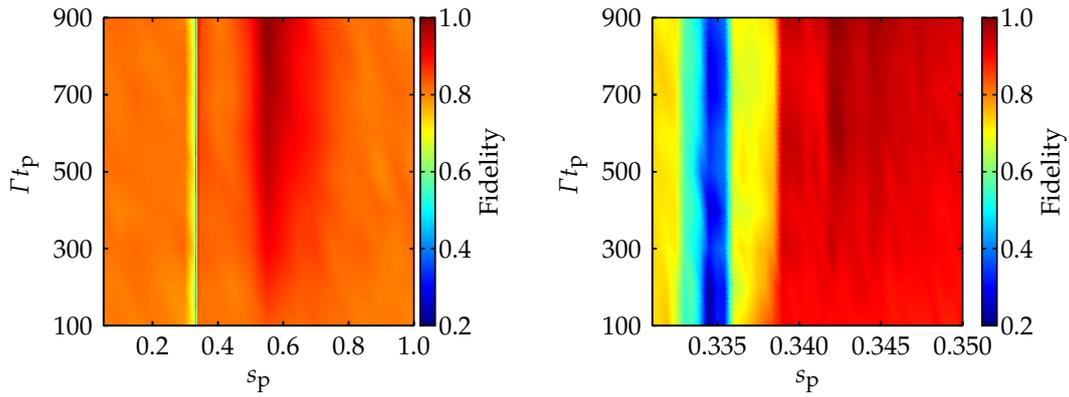


Figure 5.9 | Fidelity as a function of pausing time s_p and duration t_p (in units of $1/\Gamma$). The right-hand panel is a zoom around $s_\Delta = 0.334$.

which the fidelity goes back to its baseline value. In particular, for $t_p = 900$, we register a 20% increase of the fidelity with respect to the dissipative dynamics with no pause and total annealing time t_f , and a 45% increase with respect to a dissipative dynamics of total annealing time t'_f . The optimal pausing point s_p^{opt} is independent of the pause length.

4. When $s_p \gg s_\Delta$, the fidelity is not influenced much by pausing. The ground state is well-separated in energy from the other levels. The eigenstates of the Hamiltonian are almost diagonal in the σ_z basis and the qubit-bath coupling operator has exponentially small off-diagonal matrix elements. The perturbation theory predicts that thermal processes are exponentially suppressed, and the dynamics are frozen.

All these results are summarized in figure 5.9, where we show the fidelity as a function of both the pause time and pause duration. The right-hand panel is a zoom in the region around s_Δ . The heat maps show even more evidently that around an optimal pausing time s_p^{opt} the fidelity abruptly increases, almost independently of t_p . This is evident in figure 5.9, where the dark shadow shows up at $s_p \approx s_p^{\text{opt}}$ and $t_p > 100$. At shorter pausing lengths $t_p < 100$, this phenomenon is no longer visible. The sharp decrease of the fidelity around $s = s_\Delta$ can be better visualized in the right-hand panel. For the fully-connected p -spin model, the largest fidelity enhancement occurs for $s_p^{\text{opt}} \approx 1.65 s_\Delta$, or, equivalently, $s_p^{\text{opt}} \approx s_\Delta + 0.22$, while the worst pausing point is $s_p^{\text{wor}} \approx s_\Delta$ in the AME approximation.

Figure 5.10 shows the differences in the dynamics of $p_{\text{gs}}(s)$ when a pause of duration $t_p = 500$ is inserted at s_p^{opt} (blue line) or s_p^{wor} (red line). When $s_p = s_p^{\text{opt}}$, the ground state population grows monotonically during the pause because of thermal relaxation. In fact, thermal excitations out of the ground state are suppressed as the spectral gap is large compared to the temperature T , e. g., $E_1(s_p^{\text{opt}}) - E_0(s_p^{\text{opt}}) \approx 100$. On the other hand, for $s_p = s_\Delta$, the system can be excited due to the fact that $\Delta_{\text{min}} \ll T$ and $p_{\text{gs}}(s)$ is reduced compared to the previous case. Excitations and

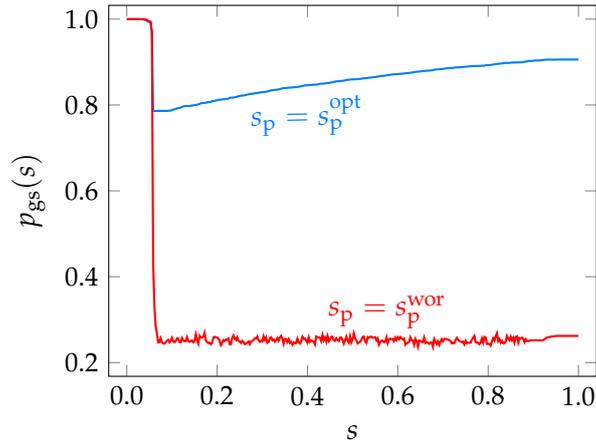


Figure 5.10 | Ground state probability as a function of time, when a pause of length $t_p = 500$ is inserted during a quantum annealing of time $t_f = 100$, either at the optimal or at the worst pausing point.

decays are almost equally probable, and produce evident noisy oscillations of $p_{\text{gs}}(s)$ around a stationary value $\bar{p}_{\text{gs}} \approx 0.25$ during the pause.

5.4 ITERATED REVERSE ANNEALING

In this section, we shall apply iterated reverse annealing (see section 2.5.2) to the ferromagnetic p -spin model with $p = 3$. We remind to the reader that in IRA the system is prepared in a classical state or in an excited eigenstate of the final Hamiltonian, rather than in the uniform quantum superposition. The annealing starts from $s(0) = 1$, where quantum fluctuations are zero, and goes back up to $s(t_{\text{inv}}) = s_{\text{inv}}$. At $s = s_{\text{inv}}$, a pause can possibly inserted, and then the annealing resumes towards $s(t_f) = 1$ (or $s(t'_f) = 1$ if there is a pause) once again. See figure 2.8 for a sketch of the IRA annealing schedule. This procedure can improve an already available solution in an open system setting. Indeed, we will show in this section using both collective and individual dephasing as a model of dissipation that open system dynamics substantially enhance the performance of reverse annealing, compared to the unitary case.

In general, s_{inv} and t_{inv} can be chosen independently of each other. However, in this section we choose the following linear relation, in order to have only one free parameter:

$$t_{\text{inv}} = t_f(1 - s_{\text{inv}}), \quad s_{\text{inv}} \neq 0, 1. \quad (5.10)$$

In this way, we have that

$$s(t) = \begin{cases} 1 - t/t_f & \text{for } t \leq t_{\text{inv}}, \\ \frac{1-s_{\text{inv}}}{t_f s_{\text{inv}}} + \frac{2s_{\text{inv}}-1}{s_{\text{inv}}} & \text{for } t > t_{\text{inv}}. \end{cases} \quad (5.11)$$

In addition, we use the D-Wave 2000Q annealing schedules of figure 2.3 and rescale the p -spin Hamiltonian by a factor 2 as done in the previous section.

In the following, we will start reverse annealing in each of the N excited states of H_p in the symmetric sector with maximal spin, i. e., $S = N/2$. The similarity to the ferromagnetic ground state is quantified by the corresponding starting eigenvalue of m_z , denoted m_0 . Note that the w th excited state differs from the ferromagnetic ground state by w spin flips. Therefore, the initial state and the target solution differ by a fraction $c = N_\uparrow/N = 1 - w/N$ of up-aligned qubits. These parameters are also related to the Hamming distance d_H , via $d_H = N - N_\uparrow = N(1 - c)$.

5.4.1 Unitary reverse annealing

In this section, we study the closed system case of a system of $N = 20$ qubits, with $p = 3$. For our choice of parameters, the p -spin system has a minimal gap $\Delta_{\min} \approx 2.45$ at $s_\Delta \approx 0.309$. We consider an annealing time of $t_f = 100$.

In figure 5.11, we report the fidelity $F = p_{\text{gs}}(t_f)$ as a function of the inversion point s_{inv} , for several initial states: $m_0 = 0.9, 0.8, 0$ and -1 .

The rightmost part of figure 5.11 corresponds to cases in which the anneal is reversed too early, i. e., for $s_{\text{inv}} > s_\Delta$. The system does not cross the avoided crossing, and the success probability is zero. Therefore, there are no visible effects on the outcome of the procedure, as the dynamics is slow compared with the minimal inverse level spacing and diabatic transitions are exponentially suppressed. Thus, the system is forced to stay in its initial state, or is excited to other high-energy states. In fact, avoided crossings between pairs of excited eigenstates occur at $s > s_\Delta$ for this model, and Landau-Zener processes can further excite the p -spin system.

On the other hand, if $s_{\text{inv}} < s_\Delta$ the system crosses the minimal gap twice. Here, the success probability benefits from Landau-Zener processes, inducing transitions towards the ground state. In this region, we also note some non-adiabatic oscillations of the success probability, due to the finite annealing time. These oscillations are more evident for large m_0 . As expected from the adiabatic theorem, they disappear by increasing the annealing time. The sharp rise of the success probability for $m_0 = 0.9$ occurs exactly at $s_{\text{inv}} = s_\Delta$. For smaller values of m_0 , the success probability rises more smoothly, as the ground state is reached after a preliminary sequence of Landau-Zener transitions between pairs of excited states, whose corresponding avoided crossings occur at $s > s_\Delta$. For $m_0 = 0.8$, a very small rise of the success probability can still be observed around $s_{\text{inv}} = s_\Delta$. This is due to the fact that during the reverse annealing, the system, prepared in the second excited state, firstly encounters an avoided crossing with the first excited state, where part of the population is transferred to the latter, and then the avoided crossing with the ground state, where the system populates its ground state. After reversing the dynamics, the two avoided crossings are encountered again (in the reverse order) and part of the population gets excited, thus reducing the success probability F .

As expected, reverse annealing is more effective when the initial state is close to the correct ground state. Moreover, as is also clear from figure 5.11, the inversion time s_{inv} must be increasingly close to $s_{\text{inv}} = 0$ for decreasing m_0 , in order to obtain

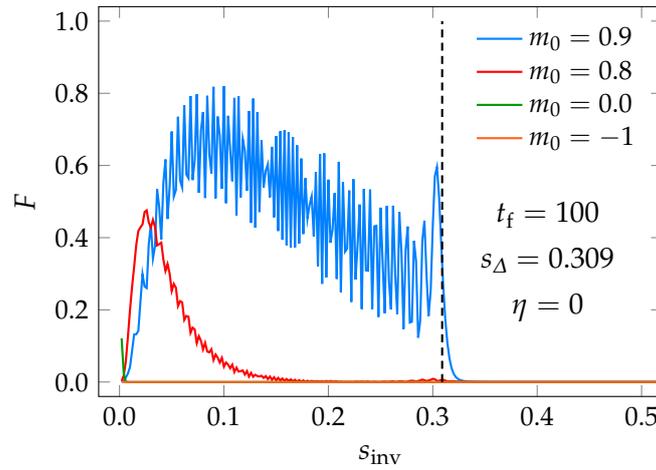


Figure 5.11 | Success probability in unitary reverse annealing as a function of the inversion point s_{inv} , for several values of the magnetization of the initial state. The dashed vertical line indicates $s_{\text{inv}} = s_{\Delta} = 0.309$. The annealing time is $t_f = 100$. We sampled the interval $s_{\text{inv}} \in (0, 1)$ using a step size of $\Delta s_{\text{inv}} = 0.002$ and repeated the dynamics for each choice of s_{inv} .

a nonzero success probability at $t = t_f$. This means that almost the entire dynamics is spent in the reverse part of the annealing, and the system is eventually quenched towards $s = 1$ for $t \approx t_f$. Even so, if the initial state is too far in energy from the correct solution, the success probability of reverse annealing is always close to zero, as evident from the curves for $m_0 = 0$ and $m_0 = -1$ in figure 5.11.

The maximum success probability decreases rapidly as a function of m_0 . This is clearly seen in figure 5.12, where we report the maximum attainable success probability as a function of m_0 , for annealing times $t_f = 100$ and 1000. Increasing the annealing time reduces nonadiabaticity and results in a lower success probability, compared with that at the end of a faster reverse anneal. As shown in the inset of figure 5.12, which zooms in on the region $m_0 \in [-1, 0]$, this decrease can be of several orders of magnitude for poorly chosen trial solutions. The influence of the annealing time is less pronounced close to $m_0 = 1$, and more evident for intermediate and lower values of m_0 . This is consistent with the adiabatic theorem, since a longer anneal time guarantees that the system will have a higher probability of remaining close to the initial eigenstate it has the largest overlap with (not necessarily the ground state) [133].

Adopting a conventional quantum annealing procedure, the success probability in the case of $t_f = 100$ would be $P_0 = 0.96$. This value is larger than any F achievable using reverse annealing. However, this argument cannot be used to discredit reverse annealing for two main reasons. First of all, it is clear that, in the analyzed case, we are very close to adiabaticity, where conventional annealing is efficient. Secondly, as clarified in the next sections, the role of dissipation may strongly affect this scenario.

The results of this section are in agreement with those reported in reference [38]. Namely, as is clear from figure 5.12, upon iteration the IRA protocol will only

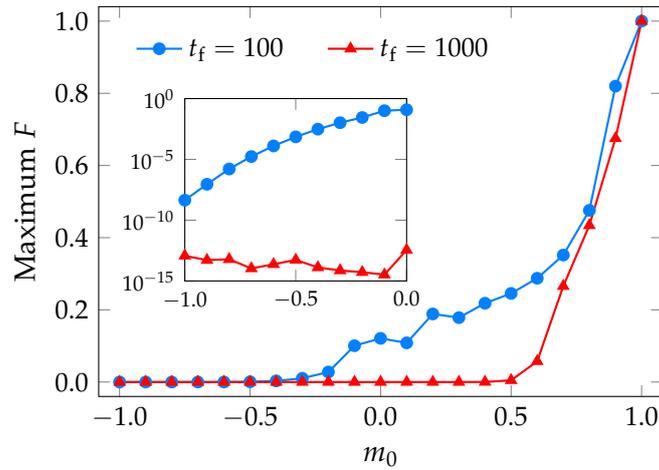


Figure 5.12 | Maximum success probability achievable with unitary reverse annealing, as a function of the magnetization of the initial state, for two annealing times: $t_f = 100$ and $t_f = 1000$. The inset zooms in on the region $m \in [-1, 0]$.

decrease the success probability under unitary, closed system dynamics, unless the initial state was already chosen as the solution of the optimization problem.

5.4.2 Reverse annealing with dephasing

In this section, we couple the p -spin system to an environment of harmonic oscillators and simulate its dynamics using the adiabatic master equation. We consider two different models of dissipation: a collective bath where the qubit system is coupled to the environment via the total magnetization, i. e., $A = m_z$, and an independent bath model where each qubit is coupled to its own bath. We consider identical baths and all equal coupling constants for simplicity. Of course, in the latter case, numerical simulations are more demanding as the interaction with the environment destroys the spin symmetry. In fact, we have to include N times as many Lindblad operators in our description, where each Lindblad operator is

$$L_{ab,i}(t) = \langle E_a(t) | \sigma_i^z | E_b(t) \rangle | E_a(t) \rangle \langle E_b(t) |. \quad (5.12)$$

For this reason, in this case we will only investigate reverse annealing starting from the first excited state in the symmetric subspace with maximum spin, i. e., $|w = 1\rangle$, for $N \in \{3, \dots, 8\}$. Moreover, for the particular cases of $N = 7$ and $N = 8$, we truncate our system to the lowest 29 and 37 eigenstates, respectively, to speed up the numerics. This choice is made since the first three levels of the maximum spin subspace at $s = 1$ are spanned by $\sum_{i=0}^2 \binom{7}{i} = 29$ (for $N = 7$) and $\sum_{i=0}^2 \binom{8}{i} = 37$ (for $N = 8$) energy eigenstates. We confirm that this is a good approximation by checking that the total population among these levels is close to one during the reverse annealing when additional levels are included in the simulation. The parameters of the dissipative simulation are $\omega_c = 1$ THz, $\eta = 1 \times 10^{-3}$, and $T = 12.1$ mK = 1.57 GHz. As always the energy scale is $\Gamma = 1$ GHz.

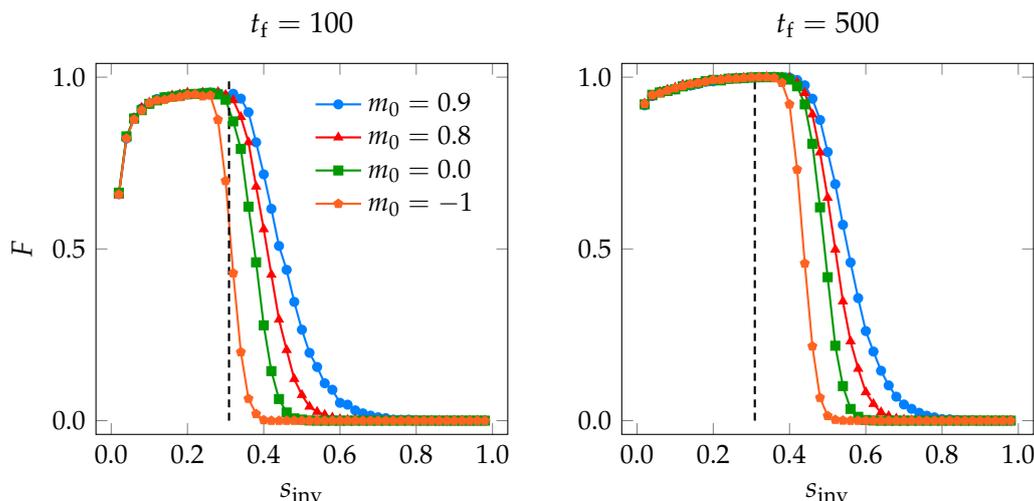


Figure 5.13 | Success probability in IRA as a function of the inversion point s_{inv} , for several values of the magnetization of the initial state. The p -spin system is coupled to a collective dephasing bosonic environment with a coupling strength of $\eta = 1 \times 10^{-3}$. The dashed vertical line denotes the time s_{Δ} of the avoided crossing between the ground state and the first excited state.

First of all, we repeat the simulations we reported in section 5.4.1 for $N = 20$, $p = 3$ and $t_f = 100$, but now include the role of the collective environment.

In the left-hand panel of figure 5.13, we show the success probability F as a function of the inversion time s_{inv} , for the four initial magnetizations $m_0 = 0.9, 0.8, 0$ and -1 . Monte Carlo errors are of the order of the point size in all cases.

As in the unitary case, if the inversion occurs too early (i. e., for $s_{\text{inv}} \gg s_{\Delta}$), the reverse annealing protocol fails to find the ferromagnetic ground state. In fact, thermal excitations are suppressed, as well as Landau-Zener transitions, due to the large level spacing, compared with the temperature and the inverse of the annealing time. For $s_{\text{inv}} \approx s_{\Delta}$, however, the scenario is drastically different from the unitary case of figure 5.11.

The first difference is that the success probability can be nonzero even if the inversion occurs for $s_{\text{inv}} \gtrsim s_{\Delta}$, especially for $m_0 = 0.9$, where the tail of the curve protrudes to $s_{\text{inv}} \approx 0.75$. When the instantaneous gap is of the same order of magnitude as the temperature, thermal processes influence reverse annealing even before crossing the minimal gap. Secondly, for all m_0 we observe a sudden increase in the success probability around $s_{\text{inv}} \approx s_{\Delta}$, that eventually brings all curves to an almost flat region at $s_{\text{inv}} < s_{\Delta}$, where the success probability reaches the large value $F \approx 0.957$. The value of the maximum success probability at the plateau is m_0 -independent within Monte Carlo errors. The time at which the success probability starts to increase with respect to the baseline depends on m_0 . Moreover, the flat region is wider for larger m_0 , although it has a finite width for all m_0 .

These results show that even trial solutions far in Hamming distance from the ferromagnetic ground state can result in a large success probability at the end of a reverse anneal. Moreover, the time window in which inverting the annealing favors

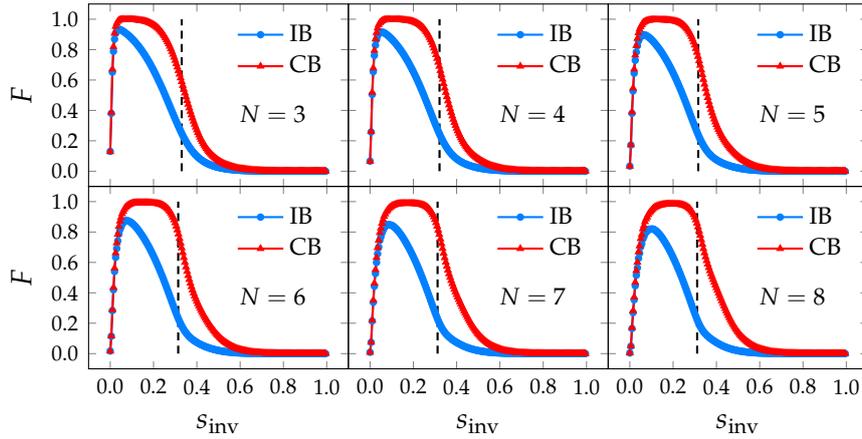


Figure 5.14 | Success probability in reverse annealing as a function of the inversion time s_{inv} , for $N \in \{3, \dots, 8\}$. The initial state is the first excited state of the maximum spin subspace ($m_0 = 1 - 2/N$). The dashed vertical line denotes the time s_{Δ} of the avoided crossing between the ground state and the first excited state. The annealing time is $t_f = 100$.

the ferromagnetic ordering is relatively large.

We also studied a longer annealing time, $t_f = 500$, as shown in the right-hand panel of figure 5.13. Here, we note that the onset of the success probability plateau shifts towards longer values of s_{inv} , compared with the case of $t_f = 100$. Therefore, the plateau is wider, and the maximum success probability at the plateau is $F \approx 1$ within Monte Carlo errors for all considered m_0 . This is in contrast with the unitary case of figure 5.12, where increasing the annealing time had detrimental effects on the algorithm. This evidence supports the idea that the success probability enhancement is due to thermal effects, rather than due to purely quantum effects [4]. Moreover, the adiabatic theorem for open quantum system guarantees convergence to the steady state of the superoperator generator of the dynamics in the large t_f limit [134, 135]. This helps to explain our observations: the steady state of the Lindblad generator of the open system dynamics considered here is the Gibbs distribution of the final Hamiltonian, which, at sufficiently low temperature relative to the gap, is the ferromagnetic ground state. Recall that in our case $\Delta_{\text{min}} \approx 2.45$ GHz (at $s_{\Delta} \approx 0.309$) and $T = 1.57$ GHz.

Comparing these results with conventional forward annealing in the presence of dissipation, we note that the success probability at the plateau is similar to that of standard quantum annealing for $\eta = 1 \times 10^{-3}$ ($F = 0.98$). The reason is that collective dephasing favors the ferromagnetic ordering in the p -spin model and its induced relaxation increases the success probability compared to the isolated case, in agreement with previous findings [4]. We also observe that when the annealing is far from adiabaticity (e.g. $t_f = 1$), reverse annealing becomes a more favorable approach to forward annealing for the p -spin model.

We also compare the collective and independent dephasing models of dissipation. Figure 5.14 shows the simulation results for the two models using the

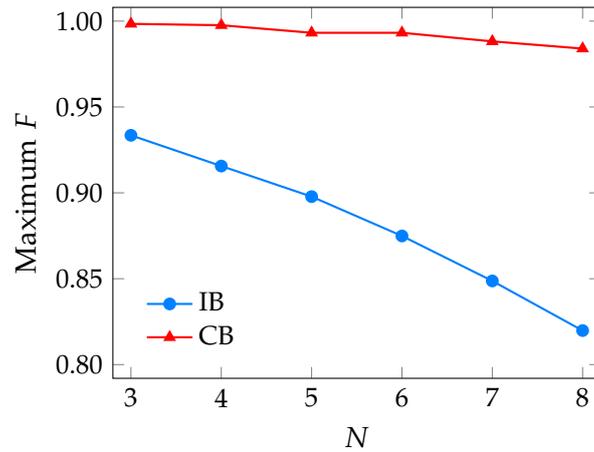


Figure 5.15 | Maximum success probability achievable with reverse annealing as a function of the number of qubits N , using the collective and the independent dephasing models ($t_f = 100$).

adiabatic master equation for $N \in \{3, \dots, 8\}$, with $t_f = 100$ and starting from the first excited state, having magnetization $m_0 = 1 - 2/N$. As shown in the figure, simulations using the collective dephasing model have larger success probabilities for almost every s_{inv} . This is because, in the independent dephasing model, other states not in the subspace of maximum spin become accessible by thermal excitation or diabatic transition during the reverse anneal. For all of the system sizes we simulated, we had to reverse anneal to a smaller inversion point s_{inv} for the independent dephasing model to achieve the same success probability as the collective dephasing model. Moreover, the maximum success probability achievable is always smaller for the independent dephasing model. The success probabilities from both models, however, are very similar as $s_{\text{inv}} \rightarrow 0$, i. e., in the quench limit of the direct part of the evolution.

Figure 5.15 shows how the maximum success probability of both bath models depends on the number of qubits. As N increases, the maximum success probability of the independent dephasing model decreases more rapidly than that of the collective dephasing model. While we can infer that if we modeled independent dephasing for $N = 20$ we would not observe as large success probabilities as in figure 5.13, we stress that reverse annealing in the independent dephasing model still yields a significantly larger success probability (for the same N values) than the unitary dynamics case described in section 5.4.1.

5.4.3 Pausing reverse annealing

We repeat the simulations inserting a pause of duration t_p at $t = t_{\text{inv}}$, so that the total annealing time, including the pause, is $t'_f = t_f + t_p$. In the case of collective dissipation, we consider $t_p = 100$ and $t_f = 400$, so that the total annealing times are $t'_f = 200$ and $t'_f = 500$, respectively.

In figure 5.16, we report the success probability as a function of the inversion

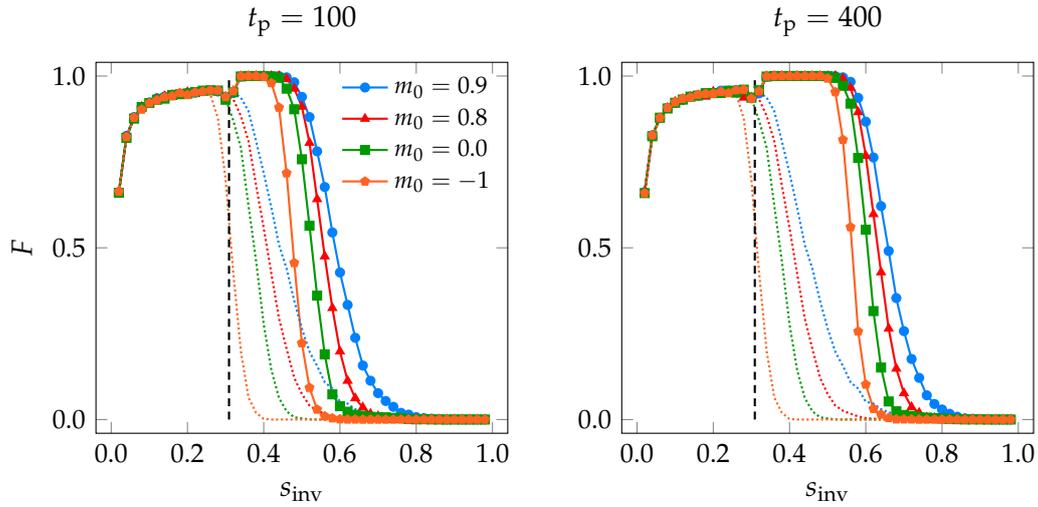


Figure 5.16 | Success probability in IRA as a function of the inversion point s_{inv} , for several values of the magnetization of the initial state. A pause of length $t_p = 100$ (left) or $t_p = 400$ (right) is added at the inversion point. The p -spin system is coupled to a collective dephasing bosonic environment with a coupling strength of $\eta = 1 \times 10^{-3}$. The other bath parameters are described in the main text. The dashed vertical line denotes the time s_Δ of the avoided crossing between the ground state and the first excited state.

point, for starting magnetizations $m_0 = 0.9, 0.8, 0$ and -1 . The left-hand panel is for $t_p = 100$ and the right-hand panel is for $t_p = 400$. We compare the paused case with the “unpaused” case. Here again, if the dynamics is reversed too early ($s_{\text{inv}} \gg s_\Delta$), the success probability at the end of the anneal vanishes. The level spacing is large compared with the temperature. The relaxation rate is small and the pause is too short to have impact on the dynamics.

However, the presence of a pause significantly changes the outcome of the annealing around $s_{\text{inv}} \approx s_\Delta$. In fact, when a pause is inserted at $s_{\text{inv}} \gtrsim s_\Delta$, the success probability reaches $F \approx 1$ for a wide range of inversion points and for all m_0 , within Monte Carlo errors. Here, the ground state is completely repopulated by thermal relaxation. This is in contrast with conventional quantum annealing, where the success probability exhibits a peak as a function of the pausing time, when the pause is inserted about 20% later than s_Δ , and then rapidly returns to its baseline value (see section 5.3). Conversely, for $s_{\text{inv}} < s_\Delta$, the effect of the pause is negligible; the solid (with pause) and dotted (no pause) lines in figure 5.16 overlap in this region.

As shown in the right-hand panel of figure 5.16, the longer pause duration affects the results only marginally. Comparing with the left-hand panel, we note that the qualitative behavior of the curves is the same in the two cases. The pause duration affects mostly the region $s_{\text{inv}} \gtrsim s_\Delta$. A longer pause enhances thermal relaxation, thus the success probability starts to increase from its baseline earlier than for shorter t_p . This results in a wider plateau where the success probability is large, compared with the case of $t_p = 100$.

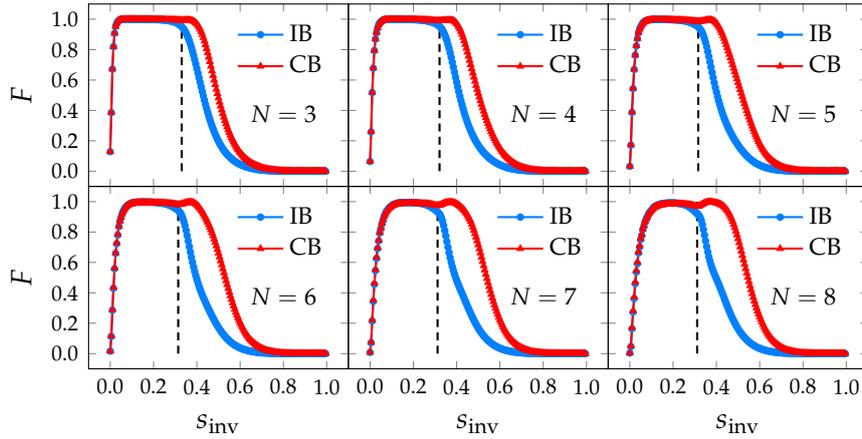


Figure 5.17 | Success probability in reverse annealing as a function of the inversion time s_{inv} , for $N \in \{3, \dots, 8\}$. The initial state is the first excited state of the maximum spin subspace ($m_0 = 1 - 2/N$). The dashed vertical line denotes the time s_{Δ} of the avoided crossing between the ground state and the first excited state. The annealing time is $t_f = 100$ and a pause of $t_p = 100$ is inserted at the inversion point.

Finally, we compare the collective and independent dephasing models while including pausing, starting from the first excited state of the maximal spin sector and for $N = 3, 4, 5, 6, 7$ and 8 . The results are shown in figure 5.17, for a pause of duration $t_p = 100$ ns inserted at the inversion point. The collective dephasing model continues to exhibit higher success probabilities than the independent dephasing model, as in the case discussed in the previous section, but the results of the two models coincide when $s_{\text{inv}} < s_{\Delta}$. Thus, relaxation to the ground state during the pause improves performance for both dephasing models. Note that, as N increases, the maximum success probability of the collective dephasing model is achieved at $s_{\text{inv}} > s_{\Delta}$, while it is achieved at $s_{\text{inv}} < s_{\Delta}$ in the independent dephasing model. This is in agreement with the $N = 20$ result shown in figure 5.13.

5.5 ADIABATIC REVERSE ANNEALING

In this last section, our aim is to study adiabatic reverse annealing applied to the p -spin model in the presence of both collective and individual dissipation. We work in the setting described in section 3.3. In particular, i) we adopt a linear annealing schedule $A(s) = 1 - s$, $B(s) = s$; ii) we use $\lambda(s) = s$ as auxiliary function in the ARA Hamiltonian of equation (2.16); iii) we use the p -spin and homogeneous transverse field Hamiltonians of equations (2.1) and (3.2). We consider the case of $p = 3$. The (Ohmic) environment is at equilibrium at $T = 12$ mK and the cutoff frequency is $\omega_c = 8\pi$. This time we are also interested in varying the strength of the transverse field relative to other parameters, therefore energies will be expressed in terms of a fixed energy scale $J = 1$ GHz. The proximity of the starting state to the target ferromagnetic state is quantified by the parameter c of equation (3.7), expressing the fraction of up-aligned spins in the starting state. The target solution

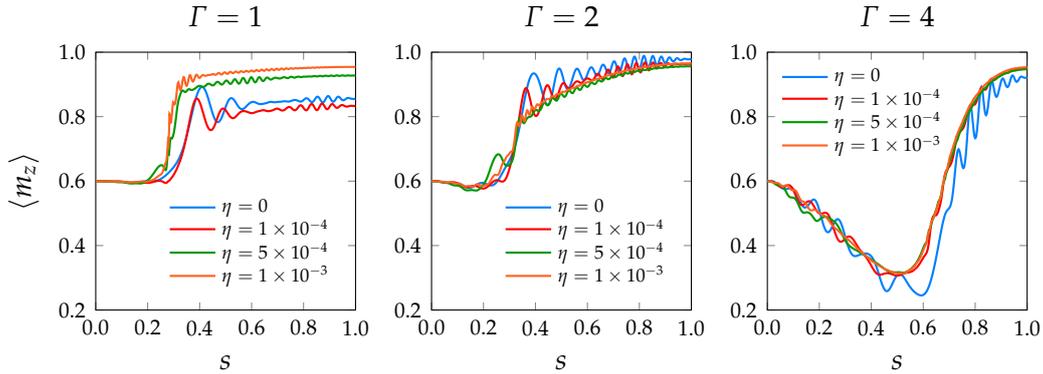


Figure 5.18 | Mean magnetization $\langle m_z \rangle$ as a function of s , for $\Gamma = 1, 2$ and 4 , and several values of the system-bath coupling strength ($\eta = 1 \times 10^{-4}, 5 \times 10^{-4}$ and 1×10^{-3}). Other parameters are $N = 50, c = 0.8, t_f = 40$.

has $c = 1$.

5.5.1 Dynamics

First of all, we study the dynamics of a system of $N = 50$ qubits and fix $t_f = 40$. We fix $c = 0.8$ so that $N_\uparrow = cN = 40$. We vary the strength of the transverse field ($\Gamma = 1, 2$ and 4) and the coupling to a collective dephasing bath ($\eta = 1 \times 10^{-4}, 5 \times 10^{-4}$ and 1×10^{-3}). The time evolution of the mean magnetization $\langle m_z \rangle$ is plotted in figure 5.18, in which we also report the unitary curves for comparison.

For $\Gamma = 1$, we see that if the system-bath coupling is weak ($\eta = 1 \times 10^{-4}$), then dissipation negatively affects the mean magnetization compared to the unitary case. However, the general features of the curve are still clearly discernible and similar to the unitary case. By contrast, for stronger interactions the environment favors the ferromagnetic alignment and the mean magnetization increases, confirming previous findings concerning dissipation-assisted quantum annealing [4, 6, 7]. The coherent oscillations are damped, compared to weaker and zero coupling strengths.

For $\Gamma = 2$, on the other hand, dissipation always reduces the mean magnetization compared to the unitary case, probably due to the fact that the final magnetization is already very large for $\eta = 0$ ($F > 0.95$), therefore excitations from the ground state are favored compared to the reverse processes. Overall, all curves are very similar to each other and the impact of dissipation appears to be marginal. Of course, this may change for longer annealing times, which we discuss later on.

Similarly to $\Gamma = 2$, for $\Gamma = 4$ dissipative dynamics are all very similar to each other and to the unitary case. As opposed to $\Gamma = 2$, in this case the final mean magnetization is slightly increased compared to $\eta = 0$.

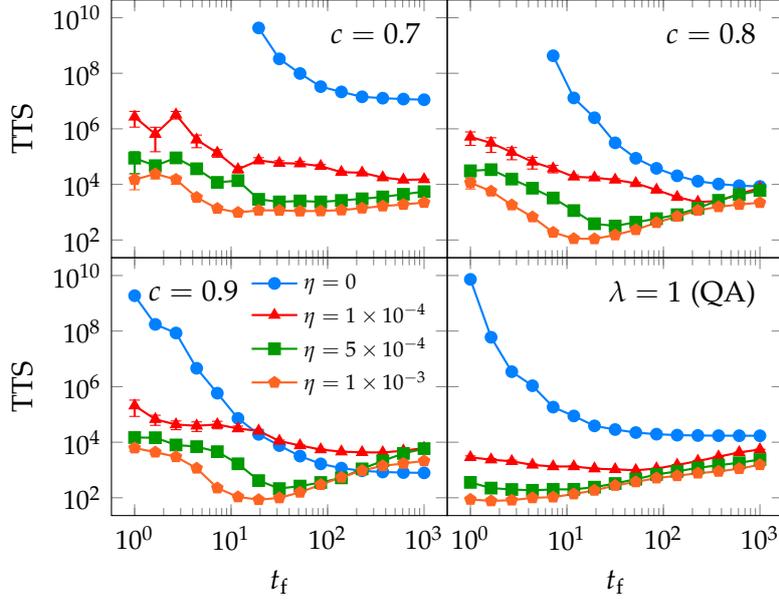


Figure 5.19 | Time to solution as a function of the annealing time t_f for $\Gamma = 1$. Top-left: $c = 0.7$; top-right: $c = 0.8$; bottom-left: $c = 0.9$; bottom-right: standard QA ($\lambda = 1$). The unitary results ($\eta = 0$) are compared with the dissipative ones in the presence of collective dephasing, for several values of the system-bath coupling strength $\eta = 1 \times 10^{-4}$, 5×10^{-4} and 1×10^{-3} .

5.5.2 Time to solution

The *time to solution* (TTS) is a performance indicator of quantum annealing and is defined as

$$\text{TTS}(t_f, p_d) = t_f \frac{\log(1 - p_d)}{\log p_e(t_f)}, \quad (5.13)$$

where $p_e(t_f) = 1 - p_{gs}(t_f)$ is the error probability and p_d is a threshold probability. The TTS represents the effective time it takes to solve the given problem at least once with a probability greater than p_d using runs of duration t_f . We fix $p_d = 0.99$ as usual in the recent literature [38, 132]. Note that the typical behavior of the error probability $p_e(t_f)$ with t_f for closed systems is that of an exponential decay at short and intermediate annealing times reflecting the Landau-Zener formula [68, 69], followed by a power-law tail for adiabatic annealing times, where the error probability scales as t_f^{-2} . As a consequence, the time to solution is expected to have a plateau in the Landau-Zener region, followed by a $t_f / \log t_f$ growth in the adiabatic regime. In the following, we will focus on the Landau-Zener region.

In this section, we study a system of $N = 45$ qubits. For annealing times $t_f \in [1, 1000]$, we have computed the TTS for several values of the system-bath coupling strength η and of the initial fraction of up-aligned spins c . In figure 5.19, we report our results for a transverse field strength of $\Gamma = 1$. For $\eta \neq 0$, the error bars represent MCWF errors [see equation (4.28)] and are smaller than the point size in many cases. We can distinguish three different behaviors at short, intermediate, and long annealing times.

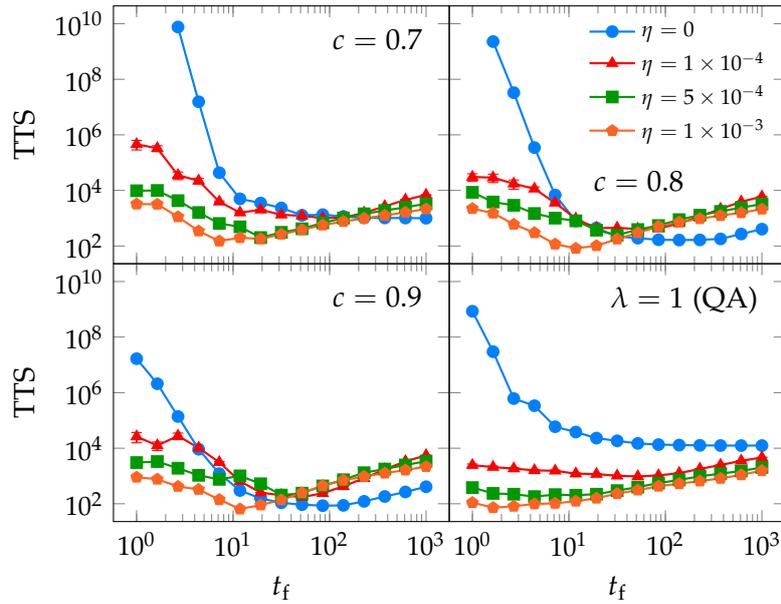


Figure 5.20 | Time to solution as a function of the annealing time t_f for $\Gamma = 2$. Top-left: $c = 0.7$; top-right: $c = 0.8$; bottom-left: $c = 0.9$; bottom-right: standard QA ($\lambda = 1$). The unitary results ($\eta = 0$) are compared with the dissipative ones in the presence of collective dephasing, for several values of the system-bath coupling strength $\eta = 1 \times 10^{-4}$, 5×10^{-4} and 1×10^{-3} .

For $\eta = 0$, the time to solution is a monotonically decreasing function of t_f . At short t_f , the fidelity of both standard QA and ARA is small, hence the TTS is very large due to vanishingly small denominators in equation (5.13). For long annealing times, the fidelity grows and the time to solution decreases, until the adiabatic regime, where the TTS eventually saturates to a plateau. These results are thoroughly discussed in reference [38].

Conversely, for $\eta \neq 0$ some of the curves show a nonmonotonic behavior, more evident for larger values of η . At short times, the environment is beneficial in all analyzed cases and the TTS is reduced compared to the isolated result, in agreement with many previous findings reporting the enhancement of the success probability of (several kinds of) quantum annealing of the p -spin model subjected to collective dephasing [4, 6, 7]. At intermediate times, the effect of the environment is still generally beneficial for the TTS except for the case $c = 0.9$, $\eta = 1 \times 10^{-4}$. For $c = 0.8$ and 0.9 the curves for larger values of η show a minimum, corresponding to the optimal working point of dissipative quantum annealing [19]. The behavior at long times depends on the value of c . If $c \lesssim 0.8$, the TTS in the dissipative case is reduced compared to the isolated case. Instead, if $c > 0.8$, the dissipative TTS is larger than the unitary one. In the case of $c = 0.9$, the starting state is already close to the target ground state, hence the TTS is already very short in the unitary case and is harmed by dissipation. Concerning standard QA, in the analyzed time window the dissipative TTS is always shorter than the unitary TTS.

By analyzing curves relative to the same coupling strength, we report that

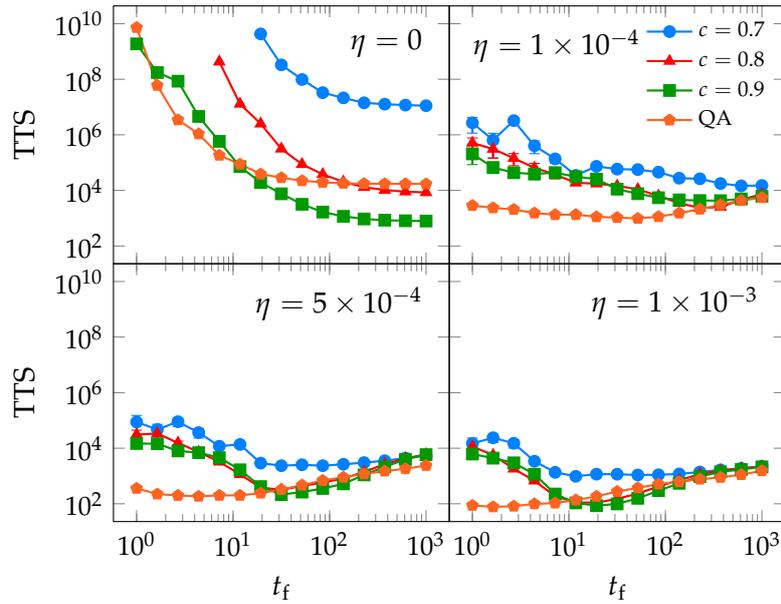


Figure 5.21 | Time to solution as a function of the annealing time t_f for $\Gamma = 1$. Top-left: $\eta = 0$; top-right: $\eta = 1 \times 10^{-4}$; bottom-left: $\eta = 5 \times 10^{-4}$; bottom-right: $\eta = 1 \times 10^{-3}$. The unitary results ($\eta = 0$) are compared with the dissipative ones in the presence of collective dephasing.

both at short and long annealing times the TTS for standard quantum annealing is shorter than that of ARA. By contrast, for $\eta > 1 \times 10^{-4}$, at intermediate times the TTS at the optimal working point in adiabatic reverse annealing for $c \gtrsim 0.8$ is smaller than the time to solution of standard quantum annealing.

In figure 5.20, we report our results concerning the TTS as a function of the annealing time t_f for a transverse field strength of $\Gamma = 2$. In the unitary case of ARA, previous findings show that increasing the transverse field strength from $\Gamma = 1$ to $\Gamma = 2$ causes the appearance of a minimum TTS at intermediate annealing times [38]. This feature survives also in the presence of dissipation, with the minimum being lower as the system-bath coupling strength increases. For short and intermediate annealing times (up to $t_f \sim 50$ to 100), we see that dissipation is beneficial for ARA. On the other hand, for longer annealing times dissipation increases the TTS compared to the isolated case. For $t_f = 10^3$, the time to solution decreases as η increases. In the case of standard quantum annealing, dissipation in the presence of collective dephasing always reduces the time to solution, compared to the isolated case. All curves are very similar to the case $\Gamma = 1$ of figure 5.19 (bottom-right panel). In the isolated case, this feature has already been reported in reference [38]. In the dissipative case, we are not surprised to observe these similarities as changing the transverse field mostly affects the beginning part of the dynamics (before the avoided crossing), where the effect of dissipation is negligible.

By grouping our curves according to the system-bath strength η , we clearly see that, while in the unitary case ARA outperforms standard QA at intermediate and long annealing times, dissipation changes this feature. Our results for $\Gamma = 1$ and

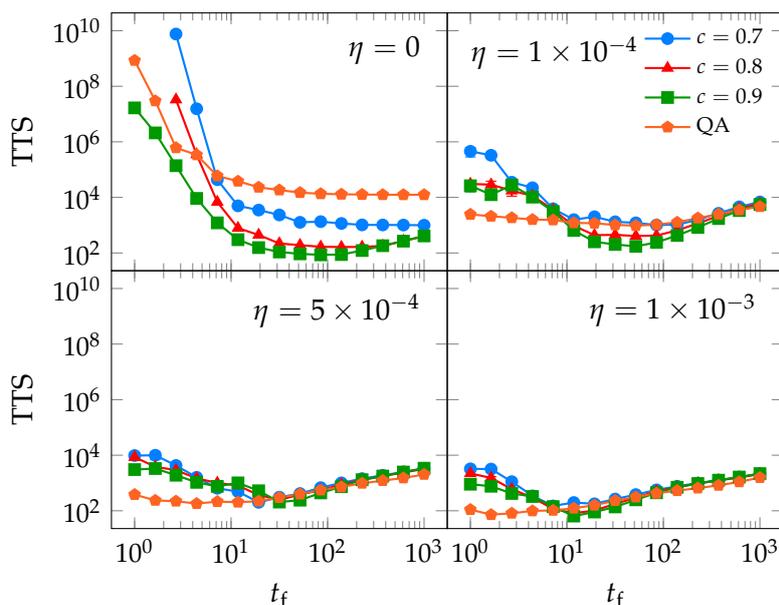


Figure 5.22 | Time to solution as a function of the annealing time t_f for $\Gamma = 2$. Top-left: $\eta = 0$; top-right: $\eta = 1 \times 10^{-4}$; bottom-left: $\eta = 5 \times 10^{-4}$; bottom-right: $\eta = 1 \times 10^{-3}$. The unitary results ($\eta = 0$) are compared with the dissipative ones in the presence of collective dephasing.

$\Gamma = 2$ are summarized in figures 5.21 and 5.22, respectively. In the unitary limit (top left panel), we see that the time to solution of ARA is indeed shorter compared to that of standard QA, in line with the mean field analysis discussed in appendix A.2 and with the literature [38]. This means that, for finite-size closed quantum systems, it is convenient to choose a starting state that is already sufficiently magnetized ($c \gtrsim 0.8$) so as to be closer to the target ferromagnetic ground state of the p -spin model and avoid exponentially vanishing spectral gaps.

On the other hand, the remaining panels tell a different story. By comparing the curves of ARA and QA for a given coupling strength η , we notice that the advantage that ARA has in the unitary case is lost in the presence of dissipation. In particular, the time to solution of the two approaches is similar and, moreover, changing the starting state of ARA does not yield any improvement. Therefore, at this level, there is no evidence in favor of using adiabatic reverse annealing in real-world devices to improve on standard quantum annealing, as dissipation would likely negate any possible advantages of ARA. These findings concern global dephasing dynamics. A possible future line of work would be to test these results in the presence of individual dephasing.

5.6 CONCLUSIONS

In this chapter, we analyzed the quantum annealing of the ferromagnetic p -spin model in the presence of dissipation. Contrary to common belief, we showed that the presence of a low temperature environment can be beneficial for quantum

annealing in some cases. In addition, the enhancing effects of the bath can be improved by exploiting advanced annealing protocols, such as pauses and iterated reverse annealing. Conversely, we have shown that novel proposals of quantum annealing such as adiabatic reverse annealing, which promises to improve QA in a unitary setting, are as effective as the standard algorithm when dissipation is taken into account.

Except for some counterexamples that are mostly model-dependent, the fact remains that the environment is often detrimental for the performance of quantum computation. This is particularly relevant for the case of adiabatic quantum computation, where the need for adiabatically long annealing times leaves the quantum system more prone to thermal effects. Therefore, the ultimate goal in AQC is to design alternative algorithms resulting in high-fidelity quantum optimization even at antiadiabatic time scales. On the one hand, this would make the system less sensitive to thermal effects. On the other hand, moving too fast across the minimal gap reduces the success probability because of Landau-Zener transitions. Promising results in this direction come from the counterdiabatic approach to adiabatic quantum computation, allowing us to circumvent LZ transitions while still achieving a large success probability. We will discuss this approach in the next chapter.

BEYOND ADIABATIC EVOLUTIONS

6.1 INTRODUCTION

Shortcuts to adiabaticity (STA) are a family of methods whose aim is to drive a quantum system adiabatically, without the prerequisite of the evolution being slow. A quantum evolution is adiabatic when the value of certain observables, denoted adiabatic invariants, are left unchanged. The adiabatic theorem establishes a bridge between these geometric quantities and the evolution time, which must be longer than the inverse of the minimal gap. Nonetheless, it is possible for a dynamics to be adiabatic even if the evolution time is shorter than the inverse of the level spacing. This result is achieved by appropriately tuning the control parameters of the Hamiltonian, or, alternatively, by including additional terms in the Hamiltonian that suppress the diabatic transitions resulting from finite sweep rates.

In quantum annealing, STA often follows the latter approach. By adding a so-called CD potential to the conventional annealing Hamiltonian, the quantum system can be driven unscathed across vanishing spectral gaps, resulting in high-fidelity adiabatic quantum computation even for short annealing times. Reducing the walltime additionally results in less decoherence and thermal noise. Figure 6.1 informally represents this concept.

A formal expression for the CD operator has been derived by Demirplak and Rice [40] and, later, by Berry [41]. However, this operator can hardly be implemented in experiments, as it involves infinite-range multiple-body interactions and, in addition, it diverges around quantum critical points in the thermodynamic limit. Moreover, it requires the exact spectrum of the original Hamiltonian, which is unknown in principle.

A lot of effort has been devoted, in recent years, in developing simpler strategies to derive approximate CD operators that can be more readily implemented on the available platforms and/or are well-defined also for many-body quantum systems and do not require the exact spectrum of the original Hamiltonian to be computed. One of the most successful attempts involves a variational approximation, and expands the exact counterdiabatic operator in terms of local Pauli operators or nested commutators between the Hamiltonian and its time derivative.

In this chapter, we are going to review some results concerning CD driving in

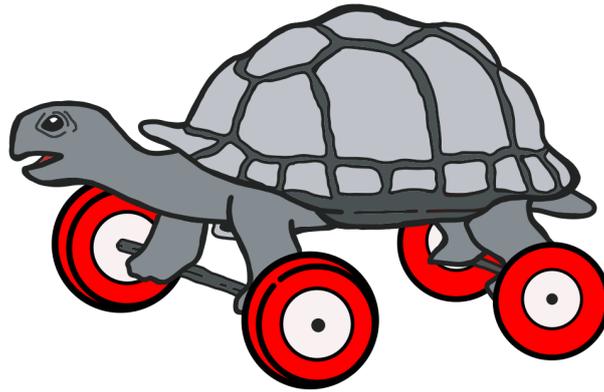


Figure 6.1 | Pictorial representation of the CD driving approach to adiabatic quantum computation. Image taken from reference [71] and inspired by the artistic work of Andree Richmond [136].

quantum annealing, with a reference to the p -spin system. This chapter is organized as follows. In section 6.2, we will introduce the several paradigms of STA, with an in-depth focus on counterdiabatic driving. In section 6.3, we will review the variational approach to CD driving as proposed by Sels and Polkovnikov [44], and we will apply it to the ferromagnetic p -spin model in section 6.4. In section 6.5, we will discuss our alternative proposal to derive an approximate CD operator by means of a genetic algorithm. We apply it to the p -spin model in section 6.6.

6.2 SHORTCUTS TO ADIABATICITY

In figure 2.2, we showed that the global adiabatic condition is too strict for adiabatic quantum computation. Indeed, in AQC the goal is to find the ground state of the target Hamiltonian H_p . In order to do so, one usually prepares the qubit system in the ground state of a simple Hamiltonian H_{TF} and let it evolve slowly, changing its Hamiltonian towards H_p , thus rigorously tracking the instantaneous ground state at all times. In this approach, most of the QPU time is spent to track the ground state for $0 < s < 1$. This portion of the dynamics is of scarce interest for optimization problems, whose goal is to read the ground state at $s = 1$. Therefore, even if the QPU does not always follow the ground state, it can still end up in the target GS at $s = 1$ thanks to a favorable combination of diabatic transitions. This fact can help us engineer faster protocols that drive the system back to its ground state despite the fact that it gets excited at intermediate times.

A faster schedule is well-suited to improve both adiabatic quantum computation, because it reduces the effect of relaxation and decoherence by speeding up the evolution, and the circuit model, where STA can be used to speed up gate operations or state preparation. In addition, STA can be fruitfully applied to all platforms currently used for quantum computation, such as superconducting qubits, nitrogen vacancies, trapped ions, or cold atoms.

When STA are used to tune the control fields of a Hamiltonian so as to realize a

specific quantum process as quickly as possible, they belong to the class of *optimal control theories* (OCT) [71, 137].

Suppose we have a Hamiltonian H_0 depending on a set of time-dependent driving fields $W(t)$ ^a. The goal in OCT is to find the optimal form of $W(t)$ to drive a starting state $|\psi_0\rangle$ into a target one $|\psi_f\rangle$ in a time t_f , according to some cost functionals we want to minimize. Typical examples of cost functionals f include:

1. The infidelity in the preparation of the target state, i. e., $f_1(|\psi(t_f)\rangle) = 1 - |\langle\psi(t_f)|\psi_f\rangle|^2$.
2. The energy of the final state. This is usually the case in AQC, where the target state is the ground state of the final Hamiltonian H_p : $f_2(|\psi(t_f)\rangle) = \langle\psi(t_f)|H_p|\psi(t_f)\rangle$.
3. More generally, a specific property of the final state; for instance, if we know that the final state is ferromagnetic, we can use the functional $f_3(|\psi(t_f)\rangle) = -\langle\psi(t_f)|m_z|\psi(t_f)\rangle$.

There may also be some constraints concerning the control fields. In particular, fields that are too strong or too rapidly varying can be impossible to realize experimentally, therefore sometimes it is useful to include constraints on the power of the driving fields, so that the quantities

$$C_i = \int_0^{t_f} |W_i(t)|^2 dt \quad (6.1)$$

are minimized. Optimal control theory aims to minimize one (or a linear combination weighted by Lagrange multipliers) of the functionals here described, while solving the Schrödinger equation with Hamiltonian $H_0[W(t)]$ and starting condition $|\psi(0)\rangle = |\psi_0\rangle$.

Many algorithms have been developed in order to face this problem. The choice of the control method is mostly driven by the number of control parameters and by whether or not the control fields are analytic. Additionally, we can distinguish between gradient-free methods, which do not require the gradient of the functional to minimize, and methods that instead use it to achieve a faster convergence. When the number of control parameters is small, gradient-free methods such as the *chopped random basis* (CRAB) optimization routine, which expands the control fields as a Fourier series on a suitably-chosen basis of special functions, are usually the best choice. For more complex problems, where the user can supply the derivative of f with respect to basis states, Krotov's algorithm gives excellent convergence and, more importantly, naturally provides an optimization scheme that improves the solution at each iteration. Some of these methods are already present in the QuTiP library [138, 139], whereas, for some others, external packages have been developed. In figure 6.2, we report a decision tree for the choice of several optimization algorithms as described in reference [140].

^aIn the QA Hamiltonian of equation (2.2), for example, these are the two schedules $A(t)$ and $B(t)$.

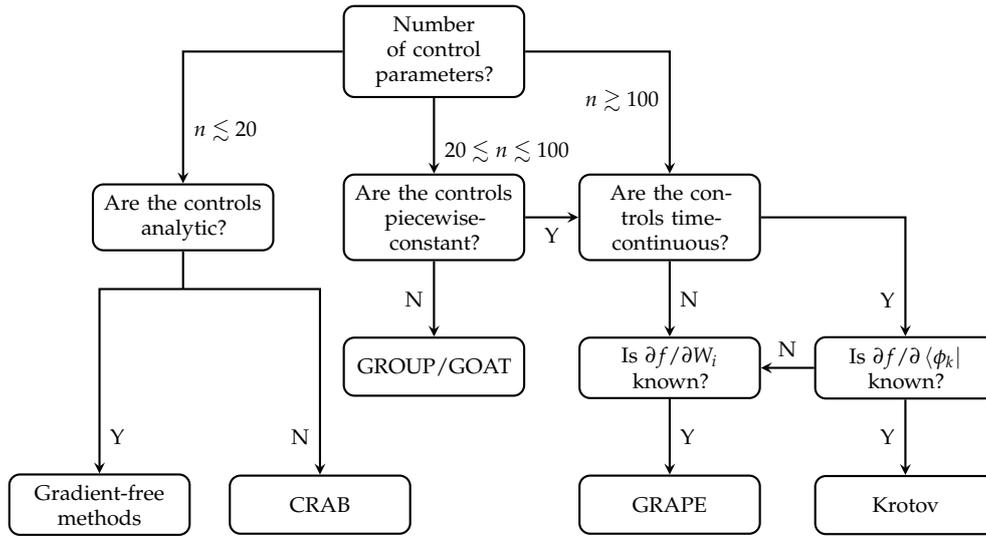


Figure 6.2 | Decision tree of OCT methods, from reference [140]. See references [141] and [142] for details about GROUP and GOAT.

The major downside of optimal control theory is that control fields must be extremely precise. Even a small deviation from the optimum can degrade performance, therefore it is paramount to have a way to properly implement the control pulses experimentally. So far, this level of accuracy is not available in D-Wave quantum annealers, where the form of the annealing schedules is dictated by the architecture and can only be changed rigidly as discussed in section 2.5. This is the reason why the community of AQC has shifted towards the complementary idea of counterdiabatic driving, that, especially in its variational formulation, can offer more viable ways to implement approximate shortcuts to adiabaticity in quantum annealers.

6.2.1 Counterdiabatic driving

A time-dependent Hamiltonian $H_0(t)$ generally induces transitions between instantaneous energy eigenstates for any finite sweep rates. The slower the change in the control fields, the smaller is the probability of Landau-Zener diabatic transitions. Counterdiabatic driving was originally proposed in references [40, 143, 144] and then rediscovered by Berry [41], who showed that it is always possible, given a Hamiltonian $H_0(t)$, to build a family of auxiliary counterdiabatic Hamiltonians $H_{cd}(t)$ so that $H(t) = H_0(t) + H_{cd}(t)$ would result in a transitionless driving of the quantum system for any annealing times.

The eigenvalue problem for $H_0(t)$ reads

$$H_0(t) |E_n(t)\rangle = E_n(t) |E_n(t)\rangle. \quad (6.2)$$

If the system is prepared in the eigenstate $|E_n(0)\rangle$, then in the adiabatic limit (i. e., for very slow evolutions) the quantum state at time t will be

$$|\psi_n(t)\rangle = e^{i\xi_n(t)} |E_n(t)\rangle, \quad (6.3)$$

where ξ_n includes both the dynamical and the geometric phases:

$$\xi_n(t) = -i \int_0^t E_n(t') dt' - \int_0^t \langle E_n(t') | \partial_{t'} E_n(t') \rangle dt'. \quad (6.4)$$

The key idea is to reverse-engineer a Hamiltonian $H(t)$ that is able to replicate the time evolution of equation (6.3) for any sweep velocity, thus suppressing diabatic transitions. It is fairly easy to build $H(t)$ by starting from the time evolution operator $U(t)$ and then using the equation $H(t) = i[\partial_t U(t)]U^\dagger(t)$. The time evolution operator is readily available and reads

$$U(t) = \sum_n e^{i\xi_n(t)} |E_n(t)\rangle \langle E_n(0)|. \quad (6.5)$$

It leads to the following Hamiltonian,

$$H(t) = \sum_n E_n |E_n\rangle \langle E_n| + i \sum_n (|\partial_t E_n\rangle \langle E_n| - \langle E_n | \partial_t E_n \rangle |E_n\rangle \langle E_n|), \quad (6.6)$$

where we omitted the explicit time dependence for brevity. The first term on the right-hand side is $H_0(t)$. The additional term is responsible for the suppression of LZ transitions and therefore is the CD Hamiltonian $H_{\text{cd}}(t)$. It is nondiagonal in the energy eigenbasis, as can be seen by rewriting it in the following form:

$$H_{\text{cd}}(t) = i \sum_{m \neq n} \frac{|E_m(t)\rangle \langle E_m(t) | \partial_t H_0(t) | E_n(t)\rangle \langle E_n(t)|}{E_n(t) - E_m(t)}. \quad (6.7)$$

This expression, albeit exact, can hardly be used in real-life applications to speed up adiabatic quantum computation, for several reasons. First of all, the Hamiltonian of equation (6.7) is ill-defined around quantum critical points, ubiquitous in adiabatic quantum computation, due to the vanishingly small denominators in the thermodynamic limit. Secondly, there are only a few (trivial) cases where we can actually compute the CD operator. In addition, in all these cases the resultant operator is highly nonlocal and impossible to implement on the available hardware. Finally, the most compelling issue is that, in order to evaluate equation (6.7), we would need access to the exact spectrum of the Hamiltonian $H_0(t)$ for all t . However, should this information be available, there would be no reason to use AQC to solve the corresponding optimization problem in the first place.

More realistically, we would need a way to build an approximate counterdiabatic operator that does not suffer from these limitations and does not require the exact spectrum of $H_0(t)$. Much effort has been devoted to this end in recent times [42, 145–150]. A substantial leap forward in this regard comes with the variational approach to counterdiabatic driving, proposed in references [44, 146, 151].

6.3 VARIATIONAL APPROACH

We here consider a linear annealing schedule in the annealing Hamiltonian H_0 , i. e., $A(s) = 1 - s$, $B(s) = s$. According to equation (6.7), the Hamiltonian is

modified by a term proportional to $\partial_t H_0(t)$. However, in standard forward AQC, the Hamiltonian at $t = 0$ and $t = t_f$ has very precise meanings. At $t = 0$, we want it to be the transverse field Hamiltonian of equation (2.1) so that the starting condition is the uniform superposition state over all computational basis states. In addition, at $t = t_f$, the Hamiltonian must be H_p , i. e., the problem Hamiltonian whose ground state we want to read. Therefore, the CD potential must vanish at $t = 0$ and $t = t_f$, thus in order to apply the CD approach in AQC we have to use an annealing function $s = s(t)$ such that $\partial_t s(t)|_{t=0,t_f} = 0$. In addition, we must have $s(0) = 0$ and $s(1) = 1$. There are many functions that satisfy these constraints. For instance, we can use the polynomial function

$$s(t) = 10 \left(\frac{t}{t_f} \right)^3 - 15 \left(\frac{t}{t_f} \right)^4 + 6 \left(\frac{t}{t_f} \right)^5, \quad (6.8)$$

or the sinusoidal function

$$s(t) = \sin^2 \left(\frac{\pi}{2} \sin^2 \left(\frac{\pi t}{2t_f} \right) \right). \quad (6.9)$$

In the following, we will prefer equation (6.9) for consistency with references [9, 44, 45]. The choice of the annealing function does not significantly affect the discussion.

The CD Hamiltonian of equation (6.7) now reads

$$H_{\text{cd}}(s) = i \partial_t s \sum_{m \neq n} \frac{|E_m(s)\rangle \langle E_m(s)| \partial_s H_0(s) |E_n(s)\rangle \langle E_n(s)|}{E_n(s) - E_m(s)} \equiv \dot{s} A_s, \quad (6.10)$$

where A_s is the counterdiabatic gauge potential and the dot denotes time derivation. From here we also see that, when $\dot{s} \rightarrow 0$, like in the adiabatic limit, the CD potential goes to zero as well.

Starting from the Schrödinger equation, we switch to the adiabatic frame, where the Hamiltonian $H_0(s)$ is diagonal at all times, using an s -dependent unitary operator $U(s)$ such that $\tilde{H}_0(s) = U^\dagger(s) H_0(s) U(s) = \sum_n E_n(s) |E_n(s)\rangle \langle E_n(s)|$. In this frame, the rotated state $|\tilde{\psi}\rangle = U(s) |\psi\rangle$ satisfies the equation

$$i \partial_t |\tilde{\psi}\rangle = [\tilde{H}_0(s(t)) - \dot{s} \tilde{A}_s] |\tilde{\psi}\rangle, \quad (6.11)$$

where quantities with the tilde are in the adiabatic frame^b. In particular,

$$\tilde{H}_0(s) = U^\dagger(s) H_0(s) U(s) = \sum_n E_n(s) |E_n(s)\rangle \langle E_n(s)|, \quad (6.12)$$

$$\tilde{A}_s = i U(s) \partial_s U^\dagger(s) = -i \partial_s U(s) U^\dagger(s). \quad (6.13)$$

By differentiating equation (6.12), one arrives to

$$i [\partial_s H_0(s) + F_{\text{ad}}(s)] = [A_s, H_0(s)], \quad (6.14)$$

^bMore precisely, in equation (6.11) there should also be an additional diagonal term proportional to the geometric term $\langle E_n(s) | \partial_s E_n(s) \rangle$, but it is inessential in the derivation.

where $F_{\text{ad}}(s) = -U^\dagger \partial_s \tilde{H}_0 U = -\sum_n \partial_s E_n(s) |E_n(s)\rangle \langle E_n(s)|$ and $A_s = i \partial_s U^\dagger(s) U(s)$. Therefore, A_s satisfies the following equation:

$$[i \partial_s H_0 - [A_s, H_0], H_0] = 0. \quad (6.15)$$

Despite its apparent simplicity, solving equation (6.15) can be cumbersome. Sels and Polkovnikov [44] proved that the problem of finding A_s can be rephrased as the minimization of the Hilbert-Schmidt norm of the operator

$$G_s(A_s) = \partial_s H_0 + i [A_s, H_0], \quad (6.16)$$

that is, it is equivalent to minimizing the action $S(A_s) = \text{tr}[G_s^2(A_s)]$,

$$\frac{\delta S(A_s)}{\delta A_s} = 0, \quad (6.17)$$

in the operator space acting on the Hilbert space of the problem.

The exact counterdiabatic operator would allow one to track the whole eigenspectrum of the original Hamiltonian H_0 . For the purpose of adiabatic quantum computation, however, this is often asking too much as we are only interested in the low-energy spectrum. We can adapt this variational formulation to focus on the low-energy spectrum by simply noting that traces in the whole Hilbert space are thermal averages at infinite temperature ($\beta \rightarrow 0$). By introducing a fictitious temperature $T < \infty$, we can redefine the action $S(A_s)$ as

$$S(A_s, \beta) = \text{tr} \left[\frac{e^{-\beta H_0}}{\mathcal{Z}} G_s^2(A_s) \right] \quad (6.18)$$

and tune $\beta = 1/T$ so that the procedure only focuses on the interesting part of the spectrum. The downside of this scheme is that we would need the spectrum of H_0 or at least an efficient way to compute the matrix exponential of equation (6.18). Therefore, in the following we will focus on the infinite-temperature case.

The exact solution to equation (6.17) would yield the counterdiabatic gauge field of equation (6.10). However, this variational formulation is particularly well-suited to deriving approximate CD potentials. In particular, one can restrict the minimization to a subspace of the operator space, containing only physically relevant quantum operators. A trial CD gauge potential A_s^* can be built as a linear combination of a basis of this subspace and the coefficients of the superposition can be chosen so as to minimize the action $S(A_s^*)$.

To date, two different approaches have been proposed to this end.

LOCAL APPROXIMATION The CD gauge potential is a linear combination of operators acting on a small number of qubits. For example, we can expand A_s^* as a linear combination of 1- and 2-body operators so that we can readily implement it on a quantum machine. There are also proposals in which the technique is applied to other architectures that naturally host multiple-body interactions [152]. Given an Ising chain, the relevant subspace is then spanned by the following l operators, $\mathcal{O} = \{ \sigma_i^y, \sigma_i^y \sigma_{i+1}^x + \text{h. c.}, \sigma_i^y \sigma_{i+1}^z + \text{h. c.} \}$.

The exact CD operator breaks time reversal, therefore only operators containing an odd number of σ^y matrices are allowed in this expansion, yielding $A_s^{(l)} = i \sum_k \alpha_k(s) \mathcal{O}_k$. In this approach, only the coefficients of the linear combination are time-dependent.

NESTED COMMUTATORS The CD gauge potential is a linear combination of *nested commutators* (NC) between the Hamiltonian H_0 and its time derivative $\partial_s H_0$. Only odd commutators, i. e., commutators involving an odd number of H_0 operators, contribute to this expansion. The exact potential can be represented in this form when the number l of nested commutators goes to infinity, or in the case of closed algebras. More practically, we can truncate the series to a finite (small) number l and obtain an approximation $A_s^{(l)}$ of A_s :

$$A_s^{(l)} = i \sum_{k=1}^l \alpha_k(s) \underbrace{[H_0, [H_0, \dots [H_0, \partial_s H_0]]]}_{2k-1} = i \sum_{k=1}^l \alpha_k(s) \mathcal{O}_k(s). \quad (6.19)$$

Matrix elements of equation (6.19) in the energy eigenbasis give

$$\langle E_m | A_s^{(l)} | E_n \rangle = i \sum_{k=1}^l \alpha_k (E_m - E_n)^{2k-1} \langle E_m | \partial_s H | E_n \rangle; \quad (6.20)$$

these are very similar to the expression obtained by taking matrix elements of the exact counterdiabatic potential of equation (6.10), except for the fact that the factor $1/\omega_{mn}$ appearing in the exact case is replaced by a “series expansion” $\sum_k \alpha_k \omega_{mn}^{2k-1}$. This expansion obviously fails around $\omega_{mn} \sim 0$, but can be used to approximate the CD potential elsewhere. In addition, this form is particularly well-suited to be realized experimentally via Floquet engineering [45].

Since A_s is a linear function of α_k , the action $S(A_s)$ is a quadratic form whose minimum is easy to find. In fact, given the expansion $A_s^{(l)} = \sum_k \alpha_k \mathcal{O}_k$, we have that

$$S(A_s^{(l)}) = \text{tr}(T_0^2) + 2 \sum_k \alpha_k \text{tr}(T_0 T_k) + \sum_{jk} \alpha_j \alpha_k \text{tr}(T_j T_k), \quad (6.21)$$

where $T_0 = \partial_s H_0$ and $T_k = [H_0, \mathcal{O}_k]$. We can define $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_l)$ and recast S as a quadratic polynomial

$$S(\boldsymbol{\alpha}) = A + 2\mathbf{B} \cdot \boldsymbol{\alpha} + \boldsymbol{\alpha}^T \cdot \underline{\mathbf{C}} \cdot \boldsymbol{\alpha}, \quad (6.22)$$

$$A = \text{tr}(T_0^2), \quad (6.23)$$

$$B_i = \text{tr}(T_0 T_i), \quad (6.24)$$

$$C_{ij} = \text{tr}(T_i T_j). \quad (6.25)$$

We can introduce the matrix U that diagonalizes $\underline{\mathbf{C}}$, i. e., $\underline{\mathbf{D}} = U^T \underline{\mathbf{C}} U$. Then, S is rewritten as

$$S(\boldsymbol{\alpha}') = A + \sum_{k=1}^l \left(2B'_k \alpha'_k + D_{kk} \alpha'^2_k \right), \quad (6.26)$$

with $\alpha' = U^T \alpha$ and $B' = U^T B$. The stationary point is

$$\frac{\partial S}{\partial \alpha'_k} = 2B'_k + 2D_{kk}\alpha'_k = 0 \implies \alpha'_k = -\frac{B'_k}{D_{kk}}. \quad (6.27)$$

This allows us to avoid the usage of minimization routines in favor of more efficient linear solvers. This algebraic discussion has been recently rediscovered in reference [153].

6.4 APPLICATION: p -SPIN MODEL, VARIATIONAL APPROACH

We here apply the variational approach to the ferromagnetic p -spin model. The energy scale is $\Gamma = 1$ GHz and we consider $t_f = 1/\Gamma$. We consider three typical values of the exponent p .

- $p = 1$: In this case, the qubit system is a single spin $S = N/2$. Using the expansion into NC, it can be easily seen that all operators \mathcal{O}_k are proportional to the operator $S_y = \sum_{i=1}^N \sigma_i^y$. Thus, one variational parameter is sufficient to recover the limit $l \rightarrow \infty$ of the nested commutators ansatz and yields the exact counterdiabatic operator up to numerical errors.
- $p = 2$: The system exhibits a second-order quantum phase transition, where the minimal gap Δ_{\min} scales as $\Delta_{\min} \sim N^{-1/3}$. In this case, the CD operator derived within the nested commutator ansatz improves the success probability of quantum annealing as a function of l . We show that the number of NC required grows with the number of qubits N .
- $p = 3$: The system exhibits a nondegenerate ground state and shows a first-order QPT. The exponent $p = 3$ is the smallest odd integer for which the p -spin model has this property. A large number l of NC is needed in order to have a significant improvement in the success probability of quantum annealing. Moreover, this number increases with the system size. However, in this case, we found an ansatz, denoted *cyclic ansatz* (CA), yielding an almost perfectly-efficient and size-independent counterdiabatic driving in the symmetry subspace with maximum spin, with just three variational parameters. The corresponding gauge potential reads

$$A_s^{\text{CA}} = \alpha_1 S_y + \alpha_2 S_y^3 + \alpha_3 (S_x S_y S_z + \text{h. c.})^{\text{c}}. \quad (6.28)$$

As a side note, the variational procedure described in the previous section requires the traces in the action $S(A_s)$ to be evaluated over the whole Hilbert space. However, in the case of the p -spin model, where we want to exploit the permutational invariance and work in the maximum spin subspace, we can define another functional $\bar{S}(A_s)$ where the traces are restricted to the $D = N + 1$ eigenstates of S^2 with maximum spin, which satisfies the same variational equations as the original $S(A_s)$. Therefore, in the following we will drop the overbar and work in the symmetry subspace, unless explicitly mentioned.

^cIt is easy to prove that $S_x S_y S_z + \text{h. c.} = S_y S_z S_x + \text{h. c.} = S_z S_x S_y + \text{h. c.}$

6.4.1 $p = 1$

Although the case $p = 1$ is trivial, it is however useful so as to set the stage for further calculations. The Hamiltonian of the p -spin model for $p = 1$ is

$$H_0(s) = -2(1-s)S_x - 2sS_z. \quad (6.29)$$

It is easy to see that $\mathcal{O}_0 = 2(S_x - S_z)$, while $\mathcal{O}_k \propto S_y$ for all $k > 0$. Thus, the NC ansatz only contains one variational parameter ($l = 1$):

$$A_s^{(1)} = \alpha S_y. \quad (6.30)$$

For $l = 1$, the quadratic action $S_1(\alpha)$ is trivially minimized by

$$\alpha_{\min} = -\frac{\text{tr}[\mathcal{O}_0\mathcal{O}_1]}{\text{tr}[\mathcal{O}_1^2]}. \quad (6.31)$$

For the p -spin system of equation (6.29), $\mathcal{O}_1 = -8(1-s)S_z + 8sS_x$, thus, in the symmetric sector,

$$\begin{aligned} \text{tr}[\mathcal{O}_0\mathcal{O}_1] &= 16 \text{tr}[S_z^2 + s(S_x^2 - S_z^2) - S_xS_z] \\ &= 16 \sum_{i=0}^N \left(\frac{N}{2} - i\right)^2 = \frac{16N(N+1)(N+2)}{12}. \end{aligned} \quad (6.32)$$

Similarly, it is possible to prove that

$$\text{tr}[\mathcal{O}_1^2] = \frac{64N(N+1)(N+2)}{12} (1 - 2s + 2s^2), \quad (6.33)$$

so that

$$\alpha_{\min} = -\frac{1}{4 - 8s + 8s^2}, \quad (6.34)$$

independently of the system size.

We numerically simulate the dynamics of the p -spin system in the symmetric spin sector, for $t_f = 1$ in units $1/\Gamma$, for system sizes ranging from $N = 10$ to 100. The probability of being in the instantaneous ground state is

$$p_{\text{gs}}(t) = |\langle E_0(t)|U(t)|\psi(t=0)\rangle|^2, \quad (6.35)$$

where $U(t)$ is the time evolution operator, and the fidelity F is $F = p_{\text{gs}}(t = t_f)$, i. e., the probability of being in the ground state at the annealing time $t = t_f$. In the absence of the CD term, the fidelity F is very small in the analyzed cases for this choice of t_f ($F \approx 1 \times 10^{-3}$ for $N = 10$, $F < 1 \times 10^{-15}$ for $N = 50$ and above). The scaling of the fidelity as a function of the system size is summarized in figure 6.3, up to $N = 100$. This clearly shows that the ansatz of equation (6.30) indeed, in this simple case, gives the exact counterdiabatic potential of equation (6.15).

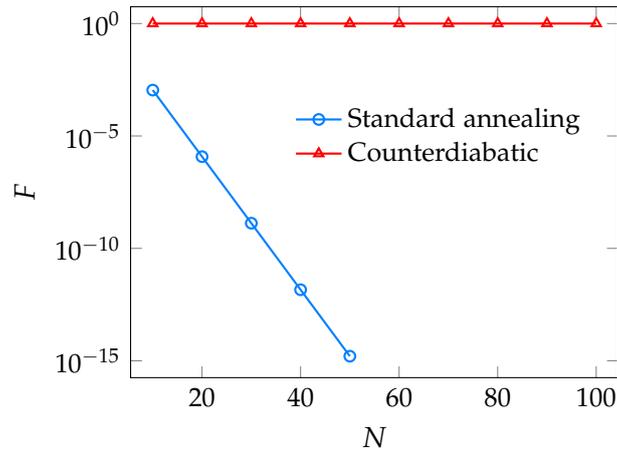


Figure 6.3 | Fidelity F as a function of the system size, for the bare annealing (blue line with circles) and for the CD driving (red line with triangles), for $p = 1$ and $t_f = 1/\Gamma$. The ansatz of equation (6.30) yields a size-independent fidelity, up to truncation errors of the numerical integrator.

6.4.2 $p = 2$

For $p = 2$, there are two degenerate ground states at $t = t_f$. They both belong to the subspace with maximum spin, i. e., they are the two fully spin polarized states $|N/2\rangle = |\uparrow\uparrow \cdots \uparrow\rangle$ and $|-N/2\rangle = |\downarrow\downarrow \cdots \downarrow\rangle$. In this case, the fidelity is given by $F = p_{N/2}(t_f) + p_{-N/2}(t_f)$.

For $p = 2$, we simulate the quantum annealing up to a final time $t_f = 1$ and study the scaling of F as a function of the system size in the maximum spin subspace and for different orders of approximation of the counterdiabatic operator $A_s^{(l)}$ with the NC ansatz.

In particular, in figure 6.4 we report the scaling of F as a function of N , for $l = 1, 3$, and 8 , compared to standard quantum annealing with no CD terms. These results are obtained using a time step of $dt = 1 \times 10^{-3}$. In the standard case, the fidelity rapidly goes to zero as we increase the number of qubits N . This is easily understood, as the energy levels become more and more dense for increasing N and, for fixed t_f , the dynamics quickly leaves the adiabatic regime. The starting paramagnetic state is metastable for all t and the p -spin system occupies this state for the whole dynamics. At the annealing time $t = t_f$, the system state would be almost identical to the starting state,

$$|\psi(t_f)\rangle \approx \frac{1}{2^{N/2}} \sum_{w=0}^N \binom{N}{w}^{1/2} \left| \frac{N}{2} - w \right\rangle, \quad (6.36)$$

where $|N/2 - w\rangle$ are Davies states in the $K = 0$ spin sector (see section 3.4). Only two terms of this sum contribute to the fidelity:

$$F = |\langle -N/2 | \psi(t_f) \rangle|^2 + |\langle N/2 | \psi(t_f) \rangle|^2 \sim 2^{-N}. \quad (6.37)$$

A single variational parameter ($l = 1$) yields a good improvement for small systems, but eventually the fidelity goes to zero for large N . By increasing the

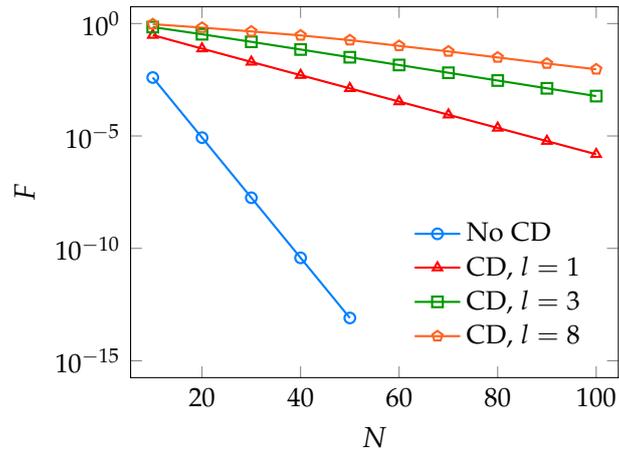


Figure 6.4 | Fidelity F as a function of the system size N , for standard annealing and for the CD ansatz at different orders ($p = 2$). Increasing the order of the approximation yields a higher fidelity F .

number of variational parameters, the improvement in F can even be of several orders of magnitude for small systems. However, the general trend is that, for large N , this improvement still gives a fidelity very close to zero and the proposed variational ansatz seems to be inefficient in the thermodynamic limit.

6.4.3 $p = 3$

For $p = 3$, the ground state is nondegenerate. We discovered an alternative ansatz tailored for this case, which we named cyclic ansatz, yielding strikingly large fidelities in the symmetric sector, almost independently of the system size, having only three variational parameters [see equation (6.28)].

We start this section by showing the scaling of the fidelity as a function of the system size. Due to the first-order QPT, we expect the fidelity to scale as

$$F = \phi e^{-\gamma N}. \quad (6.38)$$

We ask whether the CD driving can change this scaling law or not.

In figure 6.5, we show the fidelity F as a function of the system size for $t_f = 1/\Gamma$ and we compare standard quantum annealing with the nested commutator ansatz ($l = 1, 3$, and 8) and the CA. We perform a fit of the results conjecturing the exponential behavior of equation (6.38) even in the presence of CD driving. The coefficients ϕ and γ are summarized in table 6.1, comparing standard quantum annealing, the NC ansatz (for several orders l), and the cyclic ansatz.

Note that the exponent γ in the CA is three orders of magnitude smaller than both the unitary case and the NC ansatz. Moreover, in the latter case, we observe that the fidelity grows with increasing order l only for small system sizes, whereas for larger systems the fidelity shows a maximum as a function of l and then decreases. In fact, the exponent γ for $l = 8$ is larger than that for $l = 3$.

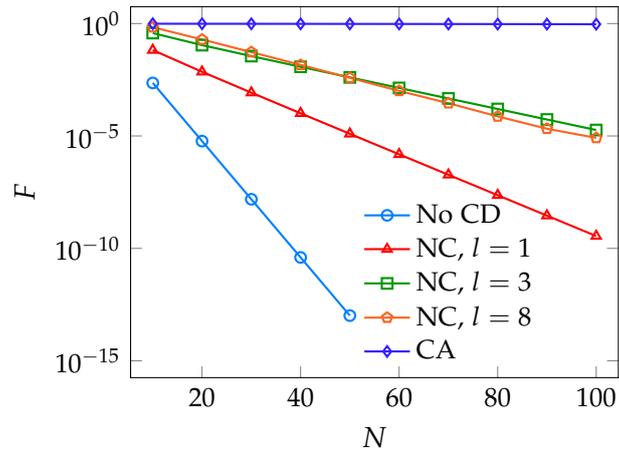


Figure 6.5 | Fidelity F as a function of the system size, for standard annealing and for the CD driving with the nested commutators and cyclic ansätze ($p = 3$).

To summarize, in the presence of a CD driving the scaling of F with N remains exponential, but the coefficient γ is reduced with respect to the bare case. Moreover, the CA yields an almost constant fidelity up to $N = 100$, and a large one ($F > 1/2$) up to $N = 1000$, providing a robust mechanism to counteract the exponentially vanishing spectral gap for macroscopic systems.

In figure 6.6, we show the time evolution of the ground state probability $p_{\text{gs}}(s)$ for $t_f = 1/\Gamma$, for $N = 10$ (left) and $N = 20$ (right) in the maximum spin subspace. The blue line indicates the fidelity of standard quantum annealing with no CD driving. The lightest green line corresponds to a CD dynamics with $l = 8$. Darker lines are for all orders starting from $l = 1$ (the lighter is the line, the larger is l). The red line represents the CA. Figure 6.6 clearly shows that the NC ansatz can improve the fidelity of quantum annealing for $p = 3$. However, the CA yields an even larger fidelity. In comparison, a similar fidelity could be reached only going beyond order $l = 8$. However, increasing the number of NC the improvement in the fidelity is gradually smaller, and we guess that in order to achieve results similar to the one obtained using the CA an unpractical large number of nested commutator would be required. Moreover, as N grows, more and more variational parameters are required to achieve a similar level of fidelity, whereas the CA requires only

Table 6.1 | Table of coefficients for the exponential fit in equation (6.38).

	ϕ	γ
No CD	0.903	5.96×10^{-1}
NC, $l = 1$	0.492	2.11×10^{-1}
NC, $l = 3$	0.999	1.09×10^{-1}
NC, $l = 8$	2.438	1.29×10^{-1}
CA	0.990	5.54×10^{-4}

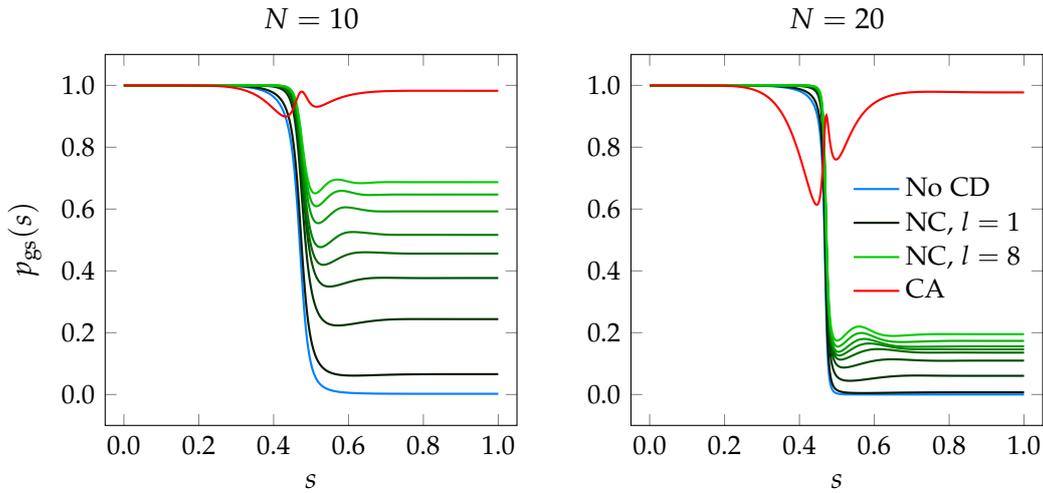


Figure 6.6 | Dynamics of the ground state probability $p_{gs}(s)$, for $p = 3$, $t_f = 1$, and $N = 10$ (left) and $N = 20$ (right). The blue line represents the fidelity of the bare evolution without CD driving. The darkest line is the ground state probability relative to $l = 1$ (NC). Analogously, the lightest line refers to $l = 8$ and all lines in between correspond to increasing orders of the NC ansatz. The red line refers to the CA.

three variational parameters. Hence, in this particular case, the cyclic ansatz is extremely efficient and outperforms the other known approximation schemes. This is even more evident increasing the number of qubits. For instance, the right-hand panel of figure 6.6 shows the same results for $N = 20$. Here, the fidelity for $l = 8$ is significantly smaller than the previous case ($F \approx 0.20$ versus $F \approx 0.68$). By contrast, the fidelity of the CA is almost unchanged with respect to $N = 10$, while the ground state probability at intermediate times is still affected by the system size.

6.4.4 Finite range p -spin model

The efficiency of the CA could depend on the peculiarities of the p -spin model, i. e., spin symmetry and infinite-range interactions. However, it is difficult to prove this statement theoretically, therefore in the following we will try to limit the range of the interactions and to break spin symmetry to gain some insights into this problem.

In the absence of spin symmetry, we need to extend the analysis presented in section 6.4.3 to the whole Hilbert space. Of course, this process is exponentially more demanding, since now we have to consider all the 2^N computational basis states. Therefore, in this section, we will limit our analysis to small systems, i. e., $N = 3$ to 8.

The total Hilbert space can be decomposed as the direct sum of the eigenspaces of the maximum total spin operator, corresponding to $K = 0$ (see section 3.4), and of its orthogonal:

$$\mathcal{H} = \mathcal{H}_{K=0} \oplus \mathcal{H}_{K=0}^\perp. \quad (6.39)$$

In principle, all traces in equation (6.22) have to be evaluated over \mathcal{H} . However, we have already discussed that the variational procedure can be easily restricted to the interesting subspace. Using a parameter $0 \leq \xi \leq 1$, we can choose whether to evaluate the traces in the whole Hilbert space or rather in the symmetric subspace. For any operator \mathcal{O} , we replace

$$\mathrm{tr}(\mathcal{O}) \longrightarrow (1 - \xi) \mathrm{tr}(\mathcal{O})_{K=0} + \xi \mathrm{tr}(\mathcal{O})_{K=0}^\perp. \quad (6.40)$$

The case $\xi = 0$ is analogous to section 6.4.3. Minimizing in the whole Hilbert space corresponds to choosing $\xi = 1/2$. Finally, $\xi = 1$ is the minimization in the orthogonal subspace $\mathcal{H}_{K=0}^\perp$. Here, focus our attention on the cases in which $\xi = 0$ and $\xi = 1/2$.

To highlight the infinite-range nature of the p -spin model, it is more convenient to rewrite its Hamiltonian [equation (3.2)] in the following form:

$$H_p = -\frac{1}{N^{p-1}} \sum_{i_1, \dots, i_p} J \sigma_{i_1}^z \dots \sigma_{i_p}^z, \quad (6.41)$$

where $i_j = 1, \dots, N$ for all j and we have introduced the unit energy J ($J = \Gamma$). As $\{\sigma_i^z, \sigma_j^z\} = 2\delta_{i,j}$, the Hamiltonian in equation (6.41) is a polynomial function of order p of Pauli operators, containing terms of orders $P = p, p-2, \dots, 1$ (odd p). Each term represents an infinite-range P -body interaction between qubits, with uniform coupling constant J .

A possible way for turning this infinite-range p -spin model into a finite-range model is by weighting J with the “distance” between the qubits involved in the p -body term. This can be easily understood by considering the simple case $p = 2$, where the Hamiltonian would be

$$H_{p=2} = \mathrm{const.} - \frac{2}{N} \sum_{i,j} J \sigma_i^z \sigma_j^z. \quad (6.42)$$

In this case, we can replace $J \rightarrow J_{i,j} = J/|i-j|^\nu$ and build a finite-range version of the p -spin model, where the exponent ν determines how punctual the interactions between the qubits are: $\nu = 0$ is the infinite-range model and $\nu \rightarrow \infty$ represents nearest-neighbor interacting qubits.

The same reasoning holds for any value of p . In particular, we can always replace J by $J_{i_1, \dots, i_p} = J/\mathrm{dist}(i_1, \dots, i_p)^\nu$. Here, we propose to consider the following form for the distance function:

$$\mathrm{dist}(i_1, \dots, i_p) = \begin{cases} \sum_{j,k>j} |i_k - i_j|/Z & \text{if } i_j\text{'s are all distinct;} \\ 1 & \text{otherwise.} \end{cases} \quad (6.43)$$

The parameter $Z = (p^3 - p)/6$ is chosen so that $\mathrm{dist}(1, 2, \dots, p) = 1$. For $p = 3$, that normalization factor is $Z = 4$. We also note that this choice of the distance function does not allow us to have finite-range models for $N = 3$ with $p = 3$, as in that case $J_{i_1, i_2, i_3} = 1$ for all combinations of indices. Of course, for all other values

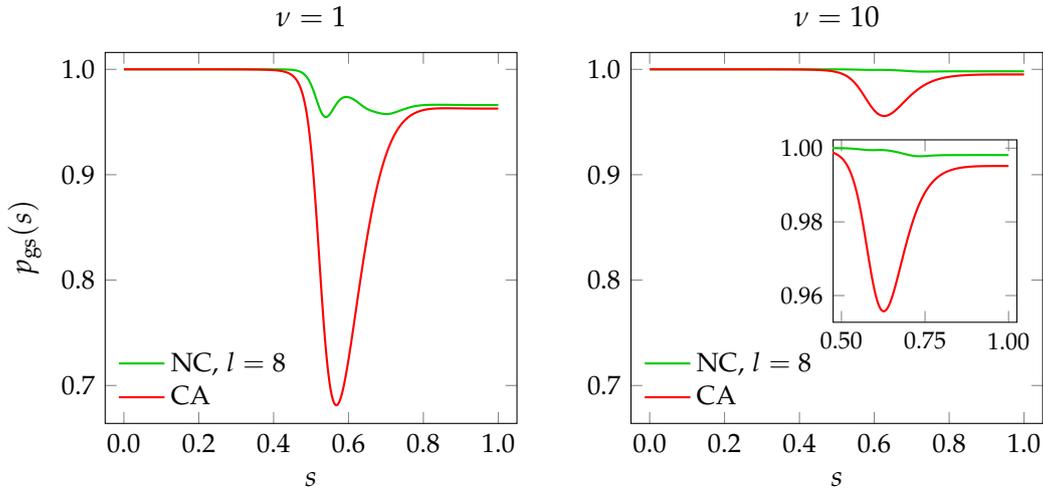


Figure 6.7 | Ground state probability $p_{\text{gs}}(s)$ as a function of time, for $N = 8$ and $p = 3$. The annealing time is $t_f = 1/\Gamma$. The minimization occurs in the whole Hilbert space ($\zeta = 1/2$). The left-hand panel is for $\nu = 1$ and right-hand panel is for $\nu = 10$ in equation (6.43). In both panels, the green line refers to the NC ansatz ($l = 8$), while the red line refers to the CA. The inset zooms in on the region $s \in [0.5, 1]$.

of N , this procedure breaks the spin symmetry of the p -spin model, therefore we will work in the whole Hilbert space and consider $\zeta = 1/2$.

In figure 6.7, we show the ground state probability $p_{\text{gs}}(s)$ as a function of time, for $N = 8$ and $p = 3$. The green line corresponds to NC while the red line corresponds to the CA. Moving from the infinite-range model to the finite-range one, we note that the efficiency of both ansätze, NC and CA, is improved, as both curves are pushed upwards. In particular, for $\nu = 10$, the fidelity F in the CA case is $F \approx 0.995$, comparable with the nested commutators ansatz with $l = 4$. Can we conclude that the reason why the CA works so well cannot be the fact that the model is infinite-range as it works even better without this feature? It is difficult to compare to the p -spin model as (i) we are not working in the symmetric subspace and (ii) the two models have different spectra, however the results of this section motivated us to go even deeper and to analyze different (random) instances and check our ansatz also in that case.

6.4.5 Random instances

Starting from the Hamiltonian of equation (6.41), here we randomly suppress some of the coupling constants J with a certain probability. In this way, we can build a family of infinite-range models, where the full-connectivity of the p -spin model is progressively lost.

The resulting Hamiltonian is identical to that in equation (6.41), but the coupling

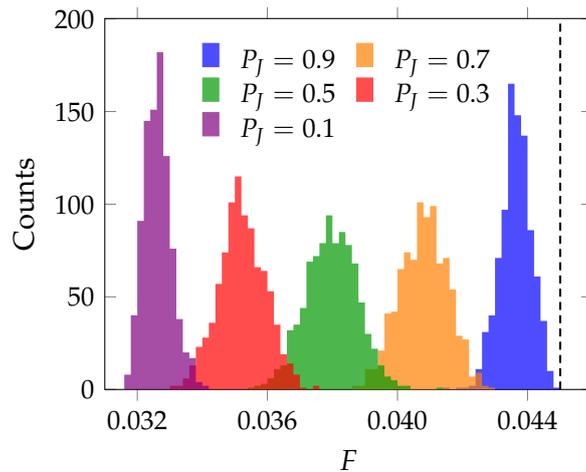


Figure 6.8 | Fidelity distributions for the random p -spin model for several values of the probability P_J of having nonzero coupling constants, for $N = 5$ and $p = 3$. The annealing time is $t_f = 1$ and there are no counterdiabatic terms. The vertical dashed line indicates the fidelity for $P_J = 1$.

constant J is replaced by a random variable \mathcal{J} satisfying

$$\mathcal{J} = \begin{cases} J & \text{with probability } P_J; \\ 0 & \text{with probability } 1 - P_J. \end{cases} \quad (6.44)$$

This model is the usual infinite-range ferromagnetic p -spin model when $P_J = 1$. For any $P_J \neq 0, 1$, this model breaks the spin symmetry and we have to work in the whole Hilbert space, with $\xi = 1/2$.

For several choices of P_J , we performed dynamical simulations for M randomly generated instances of this infinite-range random p -spin model and measured the fidelity, both with and without CD terms. We focus here on the case $N = 5$ and $p = 3$, however we obtained qualitatively similar results also for larger system sizes up to $N = 8$.

In figure 6.8, we show the fidelity distributions for $P_J = 0.1, 0.3, 0.5, 0.7$ and 0.9 . We divided the fidelity interval $[0.30, 0.45]$ into $N_b = 100$ bins and counted the occurrences over $M = 900$ repetitions of the dynamics.

The peaks of the distributions are equally spaced, which implies that the mean fidelity $\langle F \rangle$ linearly depends on P_J . According to the Landau-Zener formula, in a two-level approximation around the avoided crossing the mean fidelity would be

$$\langle F \rangle \approx 1 - e^{-2\pi \langle \Delta_{\min} \rangle^2 t_f} \approx 2\pi \langle \Delta_{\min} \rangle^2 t_f. \quad (6.45)$$

As a consequence, $\langle F \rangle \approx \langle \Delta_{\min} \rangle^2 \propto P_J$ which implies that $\langle \Delta_{\min} \rangle \sim P_J^{1/2}$.

As $P_J < 1$, all randomly generated instances have smaller gaps than the infinite-range ferromagnetic p -spin model with $P_J = 1$, therefore the corresponding fidelity is always smaller than that of the original model in the absence of CD terms. This is shown in figure 6.8, using a black dashed line to highlight the fidelity for $P_J = 1$.

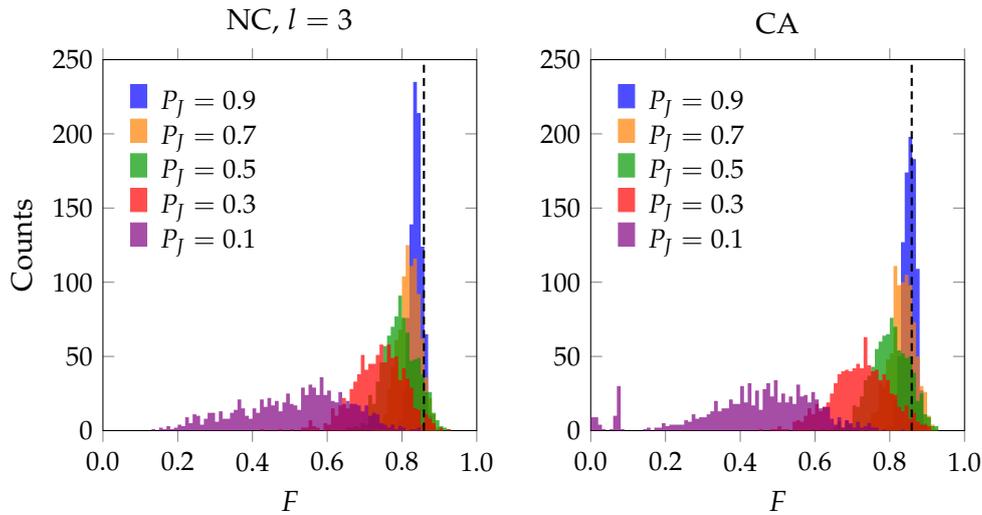


Figure 6.9 | Fidelity distributions for the random p -spin model, in the presence of counterdiabatic driving. The left-hand panel refers to the NC ansatz with $l = 3$. In the right-hand panel, we show data for the cyclic ansatz. In both panels, the black dashed line indicates the fidelity for $P_j = 1$.

In figure 6.9, we show the same fidelity distributions in the presence of CD driving. The left-hand panel is for the NC ansatz of order $l = 3$. The right-hand panel is for the cyclic ansatz. Here, we consider $N_b = 100$ bins for the fidelity interval $[0, 1]$. Except for a few minor differences, such as the presence of outliers around $F = 0$ for $P_j = 0.1$, the two plots look similar. However, we note that the mean fidelity in the cyclic case is smaller than that in the NC case for $P_j \leq 1/2$, while the opposite is true for $P_j > 1/2$. In both cases, the presence of counterdiabatic driving allows for a significantly larger fidelity (~ 15 – 20 times), compared to the case with no CD driving. Increasing the order l of the NC ansatz, the mean values of all distributions move towards $F = 1$. Moreover, all distributions become narrower and acquire a nonzero skewness.

As opposed to the case with no CD driving, here there are some instances showing larger fidelity than that for $P_j = 1$. This is highlighted in figure 6.9, where the black dashed line indicates the fidelity of the case $P_j = 1$. This evidence confirms that the efficiency of counterdiabatic driving does not entirely depend on the spectral properties of the analyzed model. In fact, even if the average minimal gap is smaller than that for $P_j = 1$, there are instances of the $P_j < 1$ case where quantum annealing with CD driving is more efficient than the case $P_j = 1$ with CD.

6.5 GENETIC APPROACH

The aim of all kinds of STA discussed above is to determine the control fields of the Hamiltonian or the coefficient of the counterdiabatic operator by minimizing one or more of the functionals described in section 6.2. Here, we want to introduce

a complementary approach that treats the time-dependence of the counterdiabatic term as fixed and looks for the best time-independent operator so as to maximize the fidelity with the target ground state. This idea is motivated by the fact that, due to technological limitations, the functional form of the control fields on a quantum machine cannot be varied within the range and accuracy required by OCT, and, likewise, time-dependent operators such as the nested commutators described in section 6.3 are difficult to implement. We stress that calling this approach “counterdiabatic driving” is incorrect because we are not suppressing diabatic transitions in any way. Therefore, we will refer to it as genetic optimal control theory.

We propose to adopt the following form for the OCT potential, i. e.,

$$H_{\text{oct}}(s) = C(s)M, \quad (6.46)$$

where $C(0) = C(1) = 0$ and M is a time-independent Hermitian operator. In the following, we will choose $C(s) = A(s)B(s)$ with $A(s) = 1 - s$ and $B(s) = s$. In this approach, M is the unknown optimal control theory operator that we aim to find.

In particular, our aim is to minimize the average energy at $s = 1$, i. e., the target functional is f_2 described in section 6.3. The minimization procedure follows the genetic algorithm outlined in section 3.7.1.

The preliminary step here is the choice of the free parameters to optimize, i. e., of the chromosomes, which reflects the structure of M . We propose two different approaches.

In the first one, the complex matrix elements of the upper triangular of M in the computational basis are the genes of a real-valued chromosome, so that each chromosome has a length of $L = D^2$. In this way, we are looking for the OCT operator in the whole Liouville space and the returned solution is generally close to the optimum, however the obvious disadvantage of this strategy is that the resulting operator is nonlocal. A variant of this approach is to assume that the OCT operator has a specific structure, for example, it is represented by a banded matrix. In this way, we can consider a family of matrices M_k , with $k = 0, 1, \dots, D - 1$, where M_0 is diagonal and $M_{k>0}$ has k nonzero subdiagonals and a nonzero diagonal. The corresponding chromosome has length $L = D + (2D - 1)k - k^2$.

In the second approach we propose, we consider $M = \sum_k \alpha_k \mathcal{O}_k$, where \mathcal{O}_k are 1- and 2-body operators and α_k are rearranged to form the chromosome. In this way, the resulting operator can easily be embedded in the quantum hardware, but it is not guaranteed to be efficient as a OCT operator.

The genetic routine works as follows. We initialize a population of N_{pop} random individuals. We prepare the qubit system in the eigenstate of $H_0(0)$ as in standard quantum annealing. For each individual, we propagate this state up to $s = 1$ using the time-dependent Hamiltonian $H(s) = H_0(s) + H_{\text{oct}}(s; M)$, and eventually measure the mean energy f_2 of the final state. We define the fitness value assigned to each chromosome as $f = -f_2$.

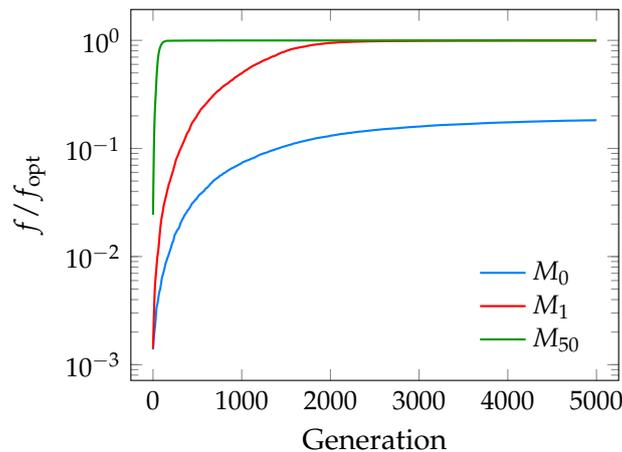


Figure 6.10 | Behavior of the mean fitness, normalized to the optimal value $f_{\text{opt}} = 50$, as a function of the generations during the genetic optimization for forms M_0 (red), M_1 (blue), and M_{D-1} (green), for a typical run of the algorithm.

6.6 APPLICATION: p -SPIN MODEL, GENETIC APPROACH

Here, we apply the genetic approach to a p -spin system of $N = 50$ qubits with $p = 3$. The ground state energy at $t = t_f$ is $E_0(t_f) = -N$ and the optimal fitness therefore is $f_{\text{opt}} = 50$. The annealing time is $t_f = 1$ (in units $1/\Gamma$) and would lead to a very small fidelity in the absence of STA mechanisms ($F \approx 6.16 \times 10^{-11}$). The minimal gap is $\Delta_{\text{min}} = 0.079$ in these units.

Following the genetic procedure, we repeatedly apply the three genetic operators, mutation with variance σ^2 and a probability of p_m , two-point crossover with a probability of p_c , and selection by tournament among N_T individuals at a time, until the termination criterion established by the maximum number of generations N_g is reached. We tested the same operators as those of section 3.7.1. In this case, we verified that the performance of the algorithm was scarcely dependent on the numerical value of the hyperparameters. Therefore, we fixed $N_{\text{pop}} = 20$, $N_g = 5000$, $p_c = 0.7$, $p_m = 0.3$, $\sigma = 1$, and $N_T = 3$.

In figure 6.10, we show the ratio between the average fitness of the chromosome population and the optimal fitness as a function of the generation of the evolutionary algorithm. We compare the results relative to three banded OCT matrices, i. e., M_0 (blue line), M_1 (red line), and M_{D-1} (green line), for a typical run of the algorithm. The chromosome sizes in these three cases are $L = 51, 151$ and 2601 , respectively. At the beginning of the procedure, the average fitness is small because the starting population is random. With form M_0 , the fitness value eventually saturates to a value that is one order of magnitude smaller than the optimal fitness $f_{\text{opt}} = 50$. This means that we cannot improve quantum annealing significantly by merely affecting the diagonal entries of the Hamiltonian. By contrast, both M_1 and M_{D-1} can produce close-to-optimal solutions, with fitness values very close to f_{opt} . Convergence with M_{D-1} is faster than with M_1 , as the larger number of free parameters allows for a more efficient sampling of the solution space.

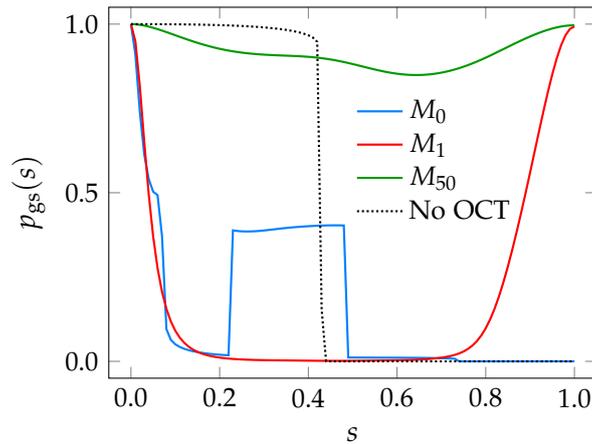


Figure 6.11 | Ground state probability as a function of time, using the OCT operators optimized via the genetic routine, following the same color scheme as of figure 6.10. We compare with the same quantity without STA (black dotted line).

In figure 6.11, we show the ground state occupation probability $p_{\text{gs}}(s)$ as a function of time, both with and without the OCT operator, following the same color scheme as in figure 6.10. In the absence of OCT, $p_{\text{gs}}(s)$ is large until it reaches the quantum critical point, at $s \approx 0.43$. After that, diabatic excitations deplete the ground state, which never regains population due to the large level spacing after the minimal gap. This is shown by using a black dotted line.

With optimal control theory, on the other hand, the fidelity becomes nonmonotonic. The exact counterdiabatic potential would produce $p_{\text{gs}}(s) = 1$ at all times. Our approach maximizes the ground state probability at $t = t_f$ but allows for diabatic transitions during the evolution. In the case of M_1 , this results in a completely depleted ground state in the middle part of the dynamics, and a sudden recovery around $t = t_f$. In the case of M_{D-1} , the ground state probability has richer features, showing oscillations that are absent in the other cases. Moreover, the ground state is never completely depleted and the probability remains large for the entire dynamics. By contrast, for matrix M_0 the system jumps multiple times between the ground state and the first excited state, and the final fidelity is small.

We could also have tested local OCT operators obtained by superposing 1- and 2-body operators. However, in order not to break the spin symmetry of the p -spin model, we would have to restrict to the following set of Hermitian operators, $\{S_x, S_y, S_z, S_x^2, S_y^2, S_z^2, S_x S_y + \text{h.c.}, S_x S_z + \text{h.c.}, S_y S_z + \text{h.c.}\}$, thus the chromosome size would be rather small ($L = 9$). For this reason, we have decided not to adopt the local strategy here.

The reason why the genetic routine works so well with the p -spin model is twofold. First of all, energy gaps at $t = t_f$ are large, therefore minimizing the mean energy of the final state always results in the system being in the objective ground state. Indeed, if we repeated the same optimization using the fidelity as the fitness function, the results would be quantitatively similar to those shown in this section. Secondly, the minimal gap of the p -spin system is found approximately in the

middle of the dynamics, where the coefficient $C(t)$ of the OCT operator has its maximum value. Thus, the effect of the optimal control theory operator is focused on the region of the dynamics where most of the diabatic transitions involving the ground state take place.

In these regards, hard test beds for our genetic approach are models having small spectral gaps at $t = t_f$ in the low-energy spectrum, and whose minimal gap is found around $t = t_f$, where the coefficient $C(t)$ of the OCT potential is small. In these cases, due to the large probability of getting trapped in excited states, the minimization of the average energy at $t = t_f$ will not necessarily imply a large fidelity and convergence could be slow. Further work in this direction is already in place and will be the subject of a forthcoming publication.

6.7 CONCLUSIONS

In this chapter, we reviewed the variational approach to CD driving and applied it to improve the annealing performance of the ferromagnetic p -spin model. We showed that the nested commutators ansatz by Claeys et al. [45] is able to enhance the success probability for systems of moderate sizes, however there are other successful ansätze that are more well-suited to the specific model at hand. In particular, we show that the cyclic ansatz leads to a size-independent high-fidelity quantum computation even for very short annealing times. What is more, we have also shown that our CA is general enough to be applied also to other variants of the p -spin model.

Finally, we proposed an alternative approach to STA based on a genetic algorithm, which is rather general but can be specialized if additional information about the system is known. Future work will explore other variants of this approach and dig more into the cases where the energy minimization alone is not enough to build a competitive optimal control operator.

CONCLUSIONS

In this work, we discussed quantum annealing and some of its variants in the presence of dissipation. Using numerical tools such as the adiabatic master equation and the Monte Carlo wave function approach, we numerically investigated the dissipative dynamics of the p -spin model, a simple fully-connected system that resembles the hardness of more complex optimization problems due to presence of a first-order quantum critical point in its phase diagram for $p > 2$.

We found numerous evidences that, within the Born-Markov approximation, the presence of a low-temperature environment can have a positive impact on the performance of quantum annealing for this system. Standard quantum annealing benefits from low-temperature dephasing by means of an observed decrease in the residual energy with respect to the true ferromagnetic ground state when the system-environment coupling strength is not extremely weak.

Something similar is observed when a pause is inserted mid-annealing around an optimal pausing point: the probability of ending up in the ferromagnetic ground state is increased thanks to incomplete thermal recombination that favors the relaxation of the p -spin system and brings it back to its ground state after the avoided crossing. This fidelity enhancement would not be possible without the presence of the environment.

Analogously, iterated reverse annealing, a more recent variant of QA, owes its ability of refining trial solutions to the presence of the environment: in a unitary setting, IRA would not be effective in improving the quality of the starting state, both in the diabatic and in the adiabatic regime. Conversely, we have shown that adiabatic reverse annealing, another proposed kind of reverse annealing, while very promising in a unitary setting, is not as effective in the presence of dissipation. In a closed system, there is a clear advantage in starting from a state that is already closer to the ferromagnetic solution compared to the uniform quantum superposition state in standard QA, as proven by the study of the time to solution. However, this advantage, originating from the avoidance of the first-order phase transition, is lost in the presence of dissipation.

All these results clearly show that in order to properly describe quantum annealing theoretically and numerically, it is paramount to take the environment into account. In some cases, the environment is even the main responsible for the

success or failure of a given annealing protocol and its effect on the annealing must be carefully addressed.

The results discussed in this thesis all concern the p -spin model. An immediate ramification of this work would be to validate these findings for other models that can be easily simulated numerically. Similarly, except for a brief parenthesis concerning short-iterative Lanczos, here we mostly focused on the adiabatic master equation approach to describe the dissipative dynamics of the p -spin model. Numerous other master equations could indeed be tested to go beyond the Born-Markov weak coupling approximation. Non-Markovian master equations could provide a more quantitative agreement with experimental data, an interesting point that we did not touch in this work. For example, it is well known that flux qubits are mostly affected by $1/f$ noise, which however cannot be incorporated easily in the Markovian adiabatic master equation approach.

Concerning shortcuts to adiabaticity, in this work we have focused on the counterdiabatic driving approach for closed quantum systems. As the environment plays such an important role in quantum annealing, a future line of research, already in place, will be devoted to develop a variational formulation to CD driving for Lindbladian dynamics that will complement the results discussed here for closed systems and will be the object of a future publication. In addition, another possible way of implementing a shortcut to adiabaticity is to act on the scheduling functions, $A(s)$ and $B(s)$. There are already proposals for variationally-optimized annealing schedules that are worth exploring as an alternative to, or in conjunction with, CD driving. In this regard, the genetic optimization method described in this thesis might be a good starting point as it allows one to optimize both the CD operator and the scheduling functions at the same time by redefining the chromosome appropriately. This will be the object of a future publication.

Regarding this last point, genetic algorithms are the simplest representative of a broader family of optimization methods based on evolutionary strategies. As discussed in this work, they are global algorithms that look for the best solution to a problem in the whole parameter space. Like all global methods, genetic algorithms have a tendency of getting stuck in local minima of the cost function. This issue can partially be lifted by implementing local searches as intermediate steps during the genetic optimization, giving rise to the so-called memetic algorithms. We have already tested the efficiency of memetic algorithms for minor embedding and the preliminary results are promising so far. The related publication is already under review. In the future, we will test the feasibility of memetic algorithms for the optimization of counterdiabatic driving operators as well.

Finally, a technical note concerning the computational methods used for this thesis. Our codes for the numerical integration of the AME and MCWF make use of multi-threaded libraries and parallelized loops with the hybrid MPI/openMP paradigm. In addition, quantum operators are represented as sparse objects to save memory. However, in recent years there has been a strong interest of the scientific community in *general-purpose graphics processing units* (GPGPUs) for heavily parallelized tasks using GPUs. Graphics processing units allow for the usage of

highly-efficient routines to speed up matrix-vector and/or matrix-matrix operations, which are typically the bottleneck of numerical simulations of quantum dynamics. Therefore, in the near future we plan to port our programs to this architecture. Since the amount of VRAM per GPU is typically much smaller compared to the CPU RAM available on high-performance computing facilities, this porting will likely involve smart compression strategies for quantum states and operators such as tensor networks, MPS's and MPO's, or quantum machine learning-inspired representations such as restricted or deep Boltzmann machines. A code refactoring will be necessary to reduce the amount of CPU-GPU communication, but the analysis of this issue is beyond the scope of the present work.

MEAN FIELD ANALYSIS



In this appendix, we are going to discuss the mean field analysis of the ferromagnetic p -spin model. We will discuss both the standard transverse field quantum annealing and the adiabatic reverse annealing Hamiltonians.

A.1 QUANTUM ANNEALING

In order to evaluate the partition function \mathcal{Z} of the p -spin model, we can apply the Suzuki-Trotter formula to decouple the two noncommuting terms in the Hamiltonian (3.4), therefore mapping the quantum system to a classical system with an extra dimension, the imaginary time [154]. Equivalently, we can expand the operator function $g(m^z) = (m^z)^p$ around the mean value $m = \langle m^z \rangle$, keeping terms up to the first order in the difference $m^z - m$:

$$g(m^z) = g(m) + \left. \frac{dg(m^z)}{dm^z} \right|_{m^z=m} (m^z - m) + \mathcal{O}\left((m^z - m)^2\right). \quad (\text{A.1})$$

Replacing equation (A.1) into the Hamiltonian (3.4), we are left with N independent 2×2 problems that can be diagonalized exactly, thus allowing us to evaluate the partition function \mathcal{Z} and the free energy $f(\beta, m) = (1/\beta N) \lim_{N \rightarrow \infty} \log \mathcal{Z}(\beta, m)$. The final result is

$$f(\beta, m) = s(p-1)m^p - \frac{1}{\beta} \log \left[2 \cosh \left(\beta \sqrt{\Gamma^2(1-s)^2 + s^2 p^2 m^{2p-2}} \right) \right]. \quad (\text{A.2})$$

The stationary condition $\partial f(\beta, m) / \partial m = 0$ yields the self-consistent equation

$$m = \frac{spm^{p-1}}{\sqrt{\Gamma^2(1-s)^2 + s^2 p^2 m^{2p-2}}} \tanh \left(\beta \sqrt{\Gamma^2(1-s)^2 + s^2 p^2 m^{2p-2}} \right), \quad (\text{A.3})$$

which includes both the classical ($\Gamma \rightarrow 0$) and the fully quantum ($\beta \rightarrow \infty$) cases.

In figure A.1, we report the phase diagram [panel A.1(a)] and the mean magnetization m [panel A.1(b)] in the mean-field approximation for $p = 2$ and $\Gamma = 1$. In the phase diagram, plotted as a function of s and $T = 1/\beta$, the darker color represents the disordered phase, where $m = 0$, while the lighter color represents the ordered phase, with $m \neq 0$. The lighter is the color, the larger is m . The two phases

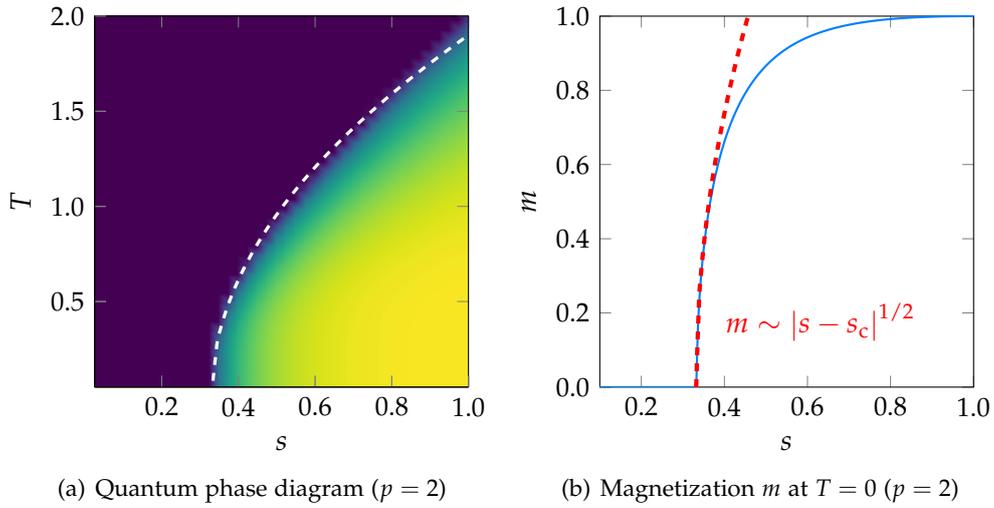


Figure A.1 | Phase diagram and magnetization in the mean-field approximation, for $p = 2$. In the left panel, the white dashed line is the critical line.

are separated by the critical line, indicated by the dashed line in figure A.1(a). The x axis is the zero temperature limit ($T = 0$), while the line $s = 1$ is the fully classical limit, where the transverse field is zero. The point $m = 0$ is always a solution of the self-consistent equation (A.3). In the paramagnetic phase, it is the global minimum of the free energy. When the system approaches the critical line, this minimum becomes broader and splits into two other minima, connected by the symmetry $m \rightarrow -m$. At the same time, $m = 0$ becomes a local maximum. Focusing on $m \geq 0$, the mean magnetization continuously grows from zero to a saturation value that tends to $m = 1$ when $T \rightarrow 0$. Therefore, the order parameter is continuous and the quantum phase transition is second-order. The noticeable smooth gradient from the dark, blue color to the yellow, light color in figure A.1(a) is the signature of the continuous second-order QPT. At $T = 0$, a series expansion of the order parameter around $s = s_c = 1/3$ (see later) reveals the usual mean-field behavior $m \sim |s - s_c|^{1/2}$. This is represented by the red dashed line in figure A.1(b).

In figure A.2(a), we show the phase diagram and mean magnetization in the case $p = 3$ and $\Gamma = 1$. For all $p > 3$, the system behaves similarly to the case $p = 3$, except for the fact that for even p the free energy is even as for $p = 2$, and there are two minima $m \rightarrow -m$. Therefore, we will only discuss the case $p = 3$. Moreover, in the following we will always avoid accidental degeneracies of the ground state by restricting our analysis to odd values of p , without any loss of generality. Also in this case, $m = 0$ is always a solution of the self-consistent equation (A.3). However, as opposed to the case $p = 2$, for $p = 3$ the paramagnetic solution is (meta)stable for every value of T and s . By approaching the critical line, indicated by the white dashed line in figure A.2(a), the free energy density starts to develop another minimum $m_1 \neq 0$, although this minimum is such that $f(\beta, m_1) > f(\beta, 0)$ until the critical line is eventually crossed and the inequality is

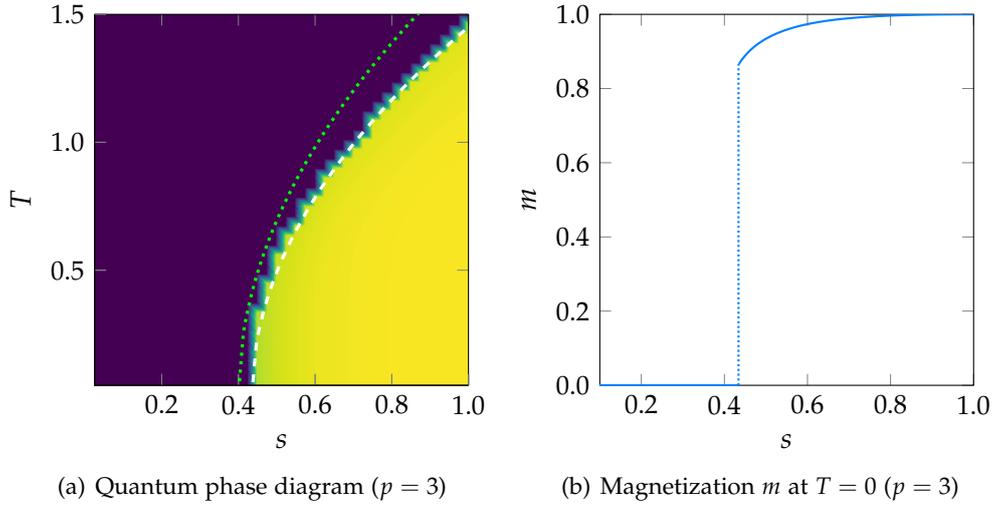


Figure A.2 | Phase diagram and magnetization in the mean-field approximation, for $p = 3$. In the left panel, the white dashed line is the critical line; the dotted line is the spinodal line, i. e., the limit of existence of a ferromagnetic metastable state in the paramagnetic phase.

suddenly reversed. This metastable minimum shows up when (s, T) crosses the spinodal line of the p -spin model [the green dotted line in figure A.2(a)], i. e., a line in the phase diagram that represents the limit of existence of the ferromagnetic phase as a metastable state. After the critical line, the metastable ferromagnetic phase becomes stable and the mean magnetization jumps from $m = 0$ to $m = m_1$. In figure A.2(a), this corresponds to the fact that the transition from the dark to the light color is very sharp, as opposed to the case $p = 2$. In figure A.2(b), the transition is indicated by the vertical dotted line at $s_c = 0.435$, i. e., the value of the critical time at $T = 0$. Thus, the order parameter is discontinuous and the phase transition is first-order.

The ground state energy is the limit $\beta \rightarrow \infty$ of the free energy density of equation (A.2). Using equation (A.3), we can write it as

$$e_{\text{gs}}(s) = \inf_m \left[-sm^p - \Gamma(1-s)\sqrt{1-m^2} \right], \quad (\text{A.4})$$

where the inf is used to choose the correct minimum according to the phase. The paramagnetic energy is $e_{\text{pm}}(s) = -\Gamma(1-s)$. For $p = 2$ and $p > 2$ (past the spinodal line), the ferromagnetic energy is $e_{\text{fm}}(s) = -sm_1^p - \Gamma(1-s)\sqrt{1-m_1^2}$, where $m_1 \neq 0$ is the magnetization in the ordered phase at time s . The critical magnetization m_c is found by noting that, at the quantum critical point, $e_{\text{fm}}(s_c) = e_{\text{pm}}(s_c)$. Simple but tedious math finally yields the critical magnetization and critical time s_c in the zero-temperature limit:

$$m_c = \sqrt{\frac{p(p-2)}{(p-1)^2}}, \quad s_c = \frac{\Gamma}{\Gamma + \frac{p}{p-1} \left[\frac{p(p-2)}{(p-1)^2} \right]^{(p-2)/2}}, \quad e_c = -\Gamma(1-s_c), \quad (\text{A.5})$$

which, in the large- p limit, yield

$$m_c \sim 1 - \frac{1}{2p^2} + \mathcal{O}(p^{-3}), \quad s_c \sim \frac{\Gamma}{1+\Gamma} \left[1 + \frac{1}{2(1+\Gamma)p} \right] + \mathcal{O}(p^{-2}). \quad (\text{A.6})$$

Note that equations (A.5) are valid (as a limit) also for $p = 2$, where

$$m_c(p = 2) = 0, \quad s_c(p = 2) = \frac{\Gamma}{2+\Gamma}. \quad (\text{A.7})$$

A.2 ADIABATIC REVERSE ANNEALING

The mean-field free energy in ARA can be derived following the same procedure as for standard quantum annealing. We can consider s and λ in equation (3.5) as fixed parameters and either perform a Suzuki-Trotter decomposition in the exponent of the partition function \mathcal{Z} or use the same series expansion of equation (A.1). The final result for the free energy density in this static approximation is

$$f(\beta, m) = s(p-1)m^{p-1} - \frac{1}{\beta} \left\langle \log \left[2 \cosh \left(\beta \sqrt{\Gamma^2(1-s)^2\lambda^2 + [spm^{p-1} + (1-s)(1-\lambda)\epsilon_i]^2} \right) \right] \right\rangle_i, \quad (\text{A.8})$$

where the angular brackets imply the average over all sites. Here, we focus on the low temperature limit, where equation (A.8) simplifies to

$$f(\beta \rightarrow \infty, m) = s(p-1)m^{p-1} - c \sqrt{[spm^{p-1} + (1-s)(1-\lambda)]^2 + \Gamma^2(1-s)^2\lambda^2} - (1-c) \sqrt{[spm^{p-1} - (1-s)(1-\lambda)]^2 + \Gamma^2(1-s)^2\lambda^2}. \quad (\text{A.9})$$

The self-consistent equation for the average magnetization reads

$$m = c \frac{spm^{p-1} + (1-s)(1-\lambda)}{\sqrt{[spm^{p-1} + (1-s)(1-\lambda)]^2 + \Gamma^2(1-s)^2\lambda^2}} + (1-c) \frac{spm^{p-1} - (1-s)(1-\lambda)}{\sqrt{[spm^{p-1} - (1-s)(1-\lambda)]^2 + \Gamma^2(1-s)^2\lambda^2}}. \quad (\text{A.10})$$

We recover equations (A.2) and (A.3) for $\lambda = 1$. Instead, for $\lambda = 0$, i.e., in the absence of the transverse field term, the model undergoes a first-order QPT where the magnetization jumps from a value close to $m = 2c - 1$ to a value close to $m = 1$ at a critical time $s = s_c$ given by

$$s_c(\lambda = 0) = \frac{2(1-c)}{1 - (2c-1)^p + 2(1-c)}. \quad (\text{A.11})$$

In figure A.3, we show the mean-field phase diagram of Hamiltonian (3.5) at $T = 0$ as a function of s and λ . We fix $p = 3$ and discuss the two cases

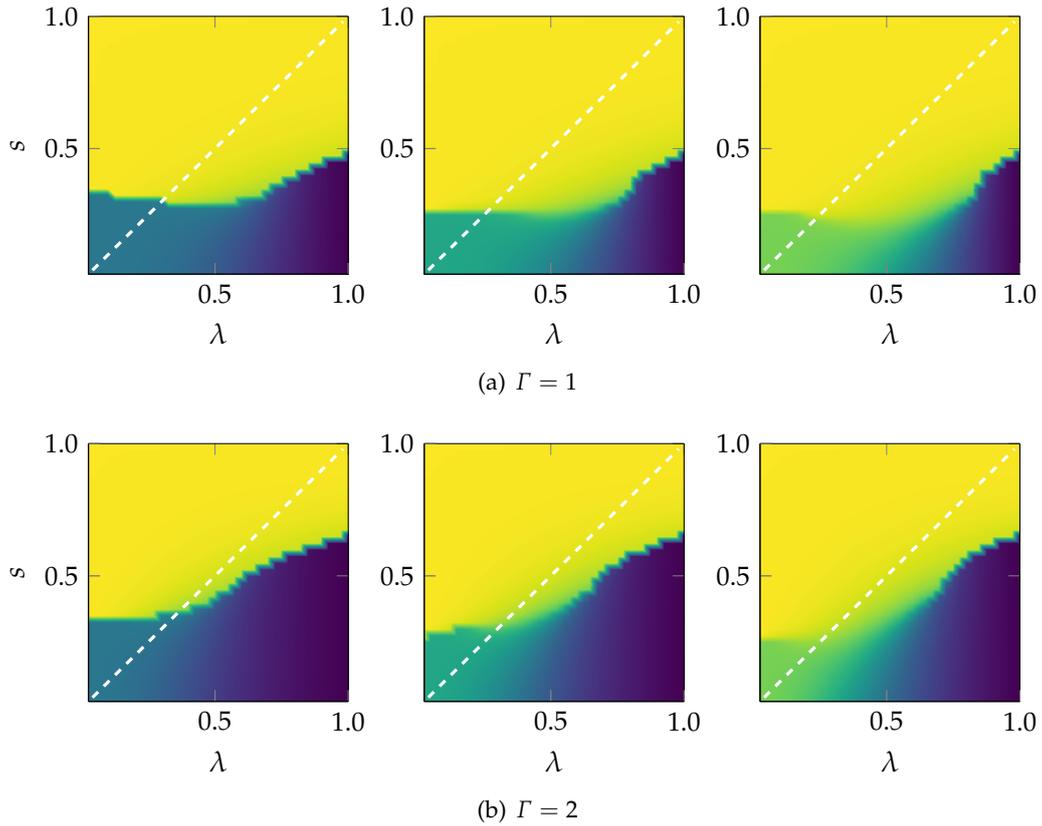


Figure A.3 | Phase diagram of ARA for $p = 3$. The top row corresponds to $\Gamma = 1$ and the bottom row corresponds to $\Gamma = 2$. The left panel is $c = 0.7$, the mid panel is $c = 0.8$, and the right panel is $c = 0.9$. (Top row) In the case $c = 0.9$, the line $\lambda(s) = s$ (the white dashed line) does not cross the line of the first-order quantum phase transition. For $c = 0.7$, no functions $\lambda(s)$ allow to avoid the first-order QPT. For $c = 0.8$, the linear function $\lambda(s) = s$ crosses the first-order critical line. However, as $c > c_c = 0.74$, it is still possible to find other paths in the parameter space that do not cross the critical line. (Bottom row) The annealing path $\lambda(s) = s$ does not cross the first-order QPT for $c = 0.8$ and $c = 0.9$, but for $c = 0.7$ the critical line is continuous and the first-order quantum phase transition cannot be avoided.

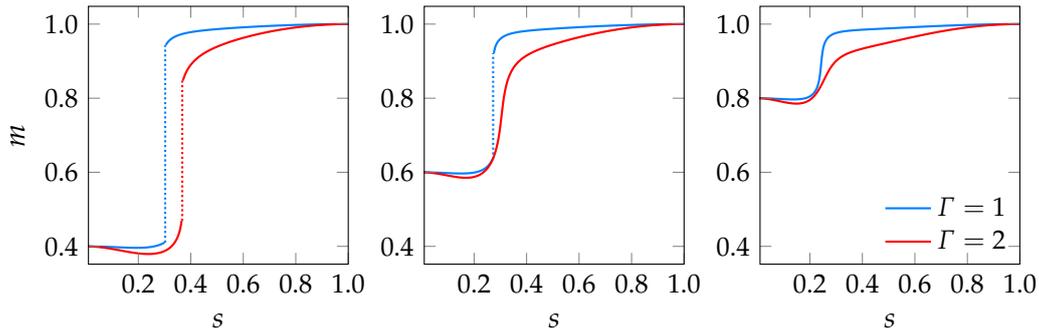


Figure A.4 | Magnetization in ARA for $p = 3$. The annealing path is $\lambda(s) = s$. The left panel is $c = 0.7$, the mid panel is $c = 0.8$, and the right panel is $c = 0.9$. In the case $c = 0.9$, the magnetization is continuous for both $\Gamma = 1$ and $\Gamma = 2$. For $c = 0.8$, the annealing path does not cross the quantum critical line for $\Gamma = 2$ and the magnetization is continuous. For $c = 0.7$, the quantum phase transition is always first-order.

$\Gamma = 1$ (top row) and $\Gamma = 2$ (bottom row). The left panel is for $c = 0.7$, the mid panel is for $c = 0.8$, and the right panel is for $c = 0.9$. The color represents the mean magnetization corresponding to the minimum of the free energy density of equation (A.9), obtained by numerically solving the self-consistent equation (A.10).

Consider the case $\Gamma = 1$. For $c = 0.7$ (left panel), the magnetization jumps from zero to a finite value when (s, λ) crosses a critical line. The critical line is continuous, meaning that there are no paths $(s, \lambda(s))$ in the parameter space that allow to avoid the first-order QPT. Visually, this is represented by the sharp transition from the lighter to the darker color. By contrast, when c is above a critical threshold c_c , the critical line breaks around the center of the phase diagram. Visually, this corresponds to a smooth color gradient from the darker to the lighter color. Therefore, it is possible to identify some paths $(s, \lambda(s))$ that do not cross the quantum critical point. Hence, the first-order quantum phase transition does not occur and is turned into a second-order QPT. For $p = 3$ and $\Gamma = 1$, the critical value is $c_c = 0.74$. For $c = 0.8$ (mid panel), we indeed see that the transition from the starting phase to the final ferromagnetic phase is very sharp for λ close to zero and to one, but becomes softer around $\lambda = 0.5$, which indicates that the first-order transition can be circumvented. In the following, we will consider as a path in the parameter space the curve $\lambda(s) = s$. This curve crosses the critical line for $c = 0.7$ and also for $c = 0.8$, as seen from the left and middle panels of figure A.3. By contrast, for $c = 0.9$ the line $\lambda(s) = s$ does not cross the critical line and the average magnetization grows continuously from $m = 2c - 1$ to $m = 1$.

For $\Gamma = 2$, the phase diagram is similar to that for $\Gamma = 1$, except for the fact that for $c = 0.8$ the annealing path $\lambda(s) = s$ does not cross the first-order quantum critical line. Therefore, the magnetization is continuous during the annealing. These features are seen more clearly in figure A.4, where we plot the solution of equation (A.10) as a function of s , for $\lambda(s) = s$. We observe a finite jump of the magnetization for $c = 0.7$ and $c = 0.8$ ($\Gamma = 1$); by contrast, for $c = 0.9$ and $c = 0.8$ ($\Gamma = 2$) the magnetization is continuous.

AME AND MCWF CPU COST

In this appendix, we discuss a comparison between the computational costs of the adiabatic master equation and of Monte Carlo wave function. In the AME approach, the most expensive operation is the evaluation of the Lindbladian given a density matrix $\rho(t)$. This is done $s_1(D)$ times per each simulation. In MCWF, on the other hand, the most computationally expensive operation is the matrix-vector product $H_{\text{eff}}(t) |\psi(t)\rangle$ needed for the time evolution operator, which is evaluated $s_2(D)$ times per trajectory. If we denote by $N_T(D)$ the number of trajectories needed to achieve a certain standard error in MCWF, we then have that the serial CPU times needed to simulate the dissipative dynamics of a quantum system using the two approaches are

$$T_{\text{AME}} = k_1 s_1(D) D^\beta, \quad (\text{B.1})$$

$$T_{\text{MCWF}} = k_2 N_T(D) s_2(D) D^\alpha, \quad (\text{B.2})$$

with $\beta \approx \alpha + 1$ as building the Lindbladian involves matrix-matrix operations rather than matrix-vector operations as in MCWF. By contrast, the number of operations s_1 and s_2 grow approximately in the same way with D .

Due to the fact that $\lambda_{\mathcal{O}}$ decreases with D , we expect that a threshold dimension D^* exists after which we can obtain the desired accuracy with just one trajectory, i. e., $N_T(D \geq D^*) = 1$. In particular, $D^* = (\Lambda_{\mathcal{O}}/\sigma_f^2)^{1/x}$. Putting all together, we find that

$$\tau = \frac{T_{\text{AME}}}{T_{\text{MCWF}}} \begin{cases} = \frac{k_1}{k_2} D & \text{for } D \geq D^*, \\ \in \left[\frac{k_1}{k_2'} D, \frac{k_1}{k_2} D^2 \right] & \text{for } D < D^*, \end{cases} \quad (\text{B.3})$$

where $k_2' = k_2(\Lambda_{\mathcal{O}}/\sigma_f^2)$. This number is typically large, hence there is an advantage in using MCWF over AME ($\tau > 1$) only when D is sufficiently large or when trajectories are computed in parallel over C cores, in which case we must replace $k_2' \rightarrow k_2'/C$ in equation (B.3) and redefine the critical dimension as $D^* = [\Lambda_{\mathcal{O}}/(C\sigma_f^2)]^{1/x}$ [33].

As an example, in figure B.1 we show the scaling of τ as a function of N in a typical case, a transverse field Ising model. Here, we have fixed $t_f = 1$ in appropriate units and $dt = 5 \times 10^{-3}$. Each qubit is coupled via σ_z to an Ohmic bath. The number of trajectories, simulated serially, is $N_T = 100$. The shaded

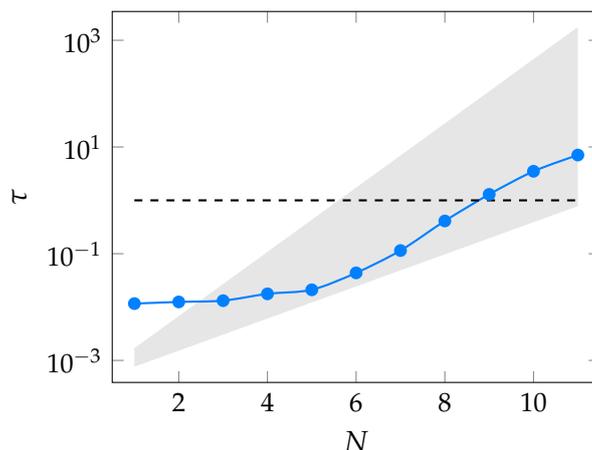


Figure B.1 | Scaling of τ as a function of N . The shaded area represents the interval of equation (B.3) when $D < D^*$, while the dashed line indicates $\tau = 1$.

area represents the interval of equation (B.3) when $D < D^*$, while the dashed line indicates $\tau = 1$. For small systems, the AME is faster than MCWF, but when the system dimension grows the Monte Carlo wave function algorithm becomes more efficient. These simulations have been performed using the two solvers `mcsolve` and `mesolve` of the QuTiP toolbox in Python [138, 139]. To obtain the results of this thesis, we have used custom-made Fortran 90 codes both for AME and for MCWF that have the advantage of being faster and more memory-efficient than the Python implementation. These codes have been thoroughly tested and compared with reliable sources such as QuTiP itself to ensure the validity of our results.

In particular, the MCWF code is built using the Intel compiler and MKL libraries for multi-threaded linear algebra operations, in conjunction with MPI for the parallel evaluation of quantum trajectories. All quantum operators are represented as sparse matrices using the sparse BLAS utilities built in the MKL library. Numerical diagonalization is performed using one of three different subroutines, i. e., exact diagonalization using LAPACK's `zhpevx` if the system size is sufficiently small, ARPACK's `znaupd` and `zneupd` for Arnoldi diagonalization with reverse communication to set the starting condition, and a more efficient custom Lanczos routine with full Gram-Schmidt reorthogonalization for the ODE solver for the waiting-time distribution when the system size is large. In order to propagate the unnormalized state $|\tilde{\psi}\rangle$ via H_{eff} , we use the following approximation,

$$|\tilde{\psi}(t + dt)\rangle \approx e^{-iH_{\text{eff}}(t+dt/2)dt} |\tilde{\psi}(t)\rangle, \quad (\text{B.4})$$

which is more accurate and less prone to convergence issues compared to a first order approximation. The downside is that at each step we have to diagonalize the Hamiltonian $H_0(t + dt/2)$, with a cost $\mathcal{O}(dD^2)$ if d is the number of Lanczos iterations^a. Even for large system sizes ($D \sim 10^6$), $d = 20$ to 100 is sufficient to achieve convergence if the time step dt is small enough.

^aEquivalently, we could evaluate the matrix exponential, but the computational cost is similar.

BIBLIOGRAPHY

- [1] G. Acampora, V. Cataudella, P. R. Hegde, P. Lucignano, **G. Passarelli**, and A. Vitiello. “An evolutionary strategy for finding effective quantum 2-body Hamiltonians of p -body interacting systems”. In: *Quantum Machine Intelligence* 1.3 (2019), pp. 113–122.
- [2] L. M. Cangemi, **G. Passarelli**, V. Cataudella, P. Lucignano, and G. De Filippis. “Beyond the Born-Markov approximation: Dissipative dynamics of a single qubit”. In: *Phys. Rev. B* 98 (18 2018), p. 184306.
- [3] L. M. Cangemi, **G. Passarelli**, V. Cataudella, P. Lucignano, and G. De Filippis. “A Numerically Exact Method for Dissipative Dynamics of Qubits”. In: *Proceedings* 12.1 (2019).
- [4] **G. Passarelli**, G. De Filippis, V. Cataudella, and P. Lucignano. “Dissipative environment may improve the quantum annealing performances of the ferromagnetic p -spin model”. In: *Phys. Rev. A* 97 (2 2018), p. 022319.
- [5] **G. Passarelli**, G. De Filippis, V. Cataudella, and P. Lucignano. “May a Dissipative Environment Be Beneficial for Quantum Annealing?” In: *Proceedings* 12.1 (2019).
- [6] **G. Passarelli**, V. Cataudella, and P. Lucignano. “Improving quantum annealing of the ferromagnetic p -spin model through pausing”. In: *Phys. Rev. B* 100 (2 2019), p. 024302.
- [7] **G. Passarelli**, K.-W. Yip, D. A. Lidar, H. Nishimori, and P. Lucignano. “Reverse quantum annealing of the p -spin model with relaxation”. In: *Phys. Rev. A* 101 (2 2020), p. 022331.
- [8] **G. Passarelli**, K.-W. Yip, D. A. Lidar, H. Nishimori, and P. Lucignano. *Adiabatic reverse annealing of the p -spin model with dissipation*. In preparation.
- [9] **G. Passarelli**, V. Cataudella, R. Fazio, and P. Lucignano. “Counterdiabatic driving in the quantum annealing of the p -spin model: A variational approach”. In: *Phys. Rev. Research* 2 (1 2020), p. 013283.
- [10] T. Kadowaki and H. Nishimori. “Quantum annealing in the transverse Ising model”. In: *Phys. Rev. E* 58 (5 1998), pp. 5355–5363.

- [11] G. E. Santoro, R. Martoňák, E. Tosatti, and R. Car. “Theory of Quantum Annealing of an Ising Spin Glass”. In: *Science* 295.5564 (2002), pp. 2427–2430.
- [12] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. “Optimization by simulated annealing”. In: *Science* 220.4598 (1983), pp. 671–680.
- [13] T. Kato. “On the Adiabatic Theorem of Quantum Mechanics”. In: *Journal of the Physical Society of Japan* 5.6 (1950), pp. 435–439.
- [14] H. P. Breuer and F. Petruccione. *The Theory of Open Quantum Systems*. OUP Oxford, 2007.
- [15] *D-Wave System Documentation. Technical Description of the D-Wave QPU*. https://docs.dwavesys.com/docs/latest/doc_qpu.html. Accessed: 2020-07-23.
- [16] R. Harris et al. “Quantum annealing with manufactured spins”. In: *Nature* 473 (7346 2011), pp. 194–198.
- [17] R. Harris et al. “Compound Josephson-junction coupler for flux qubits with minimal crosstalk”. In: *Phys. Rev. B* 80 (5 2009), p. 052506.
- [18] Yuki Bando et al. “Probing the universality of topological defect formation in a quantum annealer: Kibble-Zurek mechanism and beyond”. In: *Phys. Rev. Research* 2 (3 2020), p. 033369.
- [19] Luca Arceci, Simone Barbarino, Davide Rossini, and Giuseppe E. Santoro. “Optimal working point in dissipative quantum annealing”. In: *Phys. Rev. B* 98 (6 2018), p. 064307.
- [20] Anirban Dutta, Armin Rahmani, and Adolfo del Campo. “Anti-Kibble-Zurek Behavior in Crossing the Quantum Critical Point of a Thermally Isolated System Driven by a Noisy Control Field”. In: *Phys. Rev. Lett.* 117 (8 2016), p. 080402.
- [21] Phillip Weinberg et al. “Scaling and Diabatic Effects in Quantum Annealing with a D-Wave Device”. In: *Phys. Rev. Lett.* 124 (9 2020), p. 090502.
- [22] Vicky Choi. “Minor-embedding in adiabatic quantum computation: I. The parameter setting problem”. In: *Quantum Information Processing* 7.5 (2008), pp. 193–209.
- [23] Vicky Choi. “Minor-embedding in adiabatic quantum computation: II. Minor-universal graph design”. In: *Quantum Information Processing* 10.3 (2011), pp. 343–353.
- [24] Kotaro Tanahashi, Shinichi Takayanagi, Tomomitsu Motohashi, and Shu Tanaka. “Application of Ising Machines and a Software Development for Ising Machines”. In: *Journal of the Physical Society of Japan* 88.6 (2019), p. 061010.
- [25] J. Cai, W. G. Macready, and A. Roy. “A practical heuristic for finding graph minors”. In: *arXiv e-prints*, arXiv:1406.2741 (2014), arXiv:1406.2741. arXiv: 1406.2741 [quant-ph].

- [26] Bernard Derrida. "Random-energy model: An exactly solvable model of disordered systems". In: *Phys. Rev. B* 24 (5 1981), pp. 2613–2626.
- [27] D.J. Gross and M. Mezard. "The simplest spin glass". In: *Nuclear Physics B* 240.4 (1984), pp. 431–452.
- [28] V. Bapst and G. Semerjian. "On quantum mean-field models and their quantum annealing". In: *Journal of Statistical Mechanics: Theory and Experiment* 2012.06 (2012), Po6007.
- [29] Jérémie Roland and Nicolas J. Cerf. "Quantum search by local adiabatic evolution". In: *Phys. Rev. A* 65 (4 2002), p. 042308.
- [30] Vittorio Gorini, Andrzej Kossakowski, and E. C. G. Sudarshan. "Completely positive dynamical semigroups of N-level systems". In: *Journal of Mathematical Physics* 17.5 (1976), pp. 821–825.
- [31] G. Lindblad. "On the generators of quantum dynamical semigroups". In: *Commun. Math. Phys.* 48.2 (1976), pp. 119–130.
- [32] T. Albash, S. Boixo, D. A. Lidar, and P. Zanardi. "Quantum adiabatic Markovian master equations". In: *New Journal of Physics* 14.12 (2012), p. 123016.
- [33] Ka Wa Yip, Tameem Albash, and Daniel A. Lidar. "Quantum trajectories for time-dependent adiabatic master equations". In: *Phys. Rev. A* 97 (2 2018), p. 022116.
- [34] Matteo M. Wauters, Rosario Fazio, Hidetoshi Nishimori, and Giuseppe E. Santoro. "Direct comparison of quantum and simulated annealing on a fully connected Ising ferromagnet". In: *Phys. Rev. A* 96 (2 2017), p. 022326.
- [35] Jeffrey Marshall, Davide Venturelli, Itay Hen, and Eleanor G. Rieffel. "Power of Pausing: Advancing Understanding of Thermalization in Experimental Quantum Annealers". In: *Phys. Rev. Applied* 11 (4 2019), p. 044083.
- [36] Huo Chen and Daniel A. Lidar. "Why and When Pausing is Beneficial in Quantum Annealing". In: *Phys. Rev. Applied* 14 (1 2020), p. 014100.
- [37] Alejandro Perdomo-Ortiz, Salvador E. Venegas-Andraca, and Alán Aspuru-Guzik. "A study of heuristic guesses for adiabatic quantum computation". In: *Quantum Information Processing* 10.1 (2011), pp. 33–52.
- [38] Yu Yamashiro, Masaki Ohkuwa, Hidetoshi Nishimori, and Daniel A. Lidar. "Dynamics of reverse annealing for the fully connected p -spin model". In: *Phys. Rev. A* 100 (5 2019), p. 052321.
- [39] Masaki Ohkuwa, Hidetoshi Nishimori, and Daniel A. Lidar. "Reverse annealing for the fully connected p -spin model". In: *Phys. Rev. A* 98 (2 2018), p. 022314.
- [40] Mustafa Demirplak and Stuart A. Rice. "Adiabatic Population Transfer with Control Fields". In: *The Journal of Physical Chemistry A* 107.46 (2003), pp. 9937–9945.

- [41] M. V. Berry. “Transitionless quantum driving”. In: *Journal of Physics A: Mathematical and Theoretical* 42.36 (2009), p. 365303.
- [42] Adolfo del Campo, Marek M. Rams, and Wojciech H. Zurek. “Assisted Finite-Rate Adiabatic Passage Across a Quantum Critical Point: Exact Solution for the Quantum Ising Model”. In: *Phys. Rev. Lett.* 109 (11 2012), p. 115703.
- [43] Kazutaka Takahashi. “Transitionless quantum driving for spin systems”. In: *Phys. Rev. E* 87 (6 2013), p. 062117.
- [44] Dries Sels and Anatoli Polkovnikov. “Minimizing irreversible losses in quantum systems by local counteradiabatic driving”. In: *Proceedings of the National Academy of Sciences* 114.20 (2017), E3909–E3916.
- [45] P. W. Claeys, M. Pandey, D. Sels, and A. Polkovnikov. “Floquet-Engineering Counteradiabatic Protocols in Quantum Many-Body Systems”. In: *Phys. Rev. Lett.* 123 (9 2019), p. 090602.
- [46] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser. “Quantum Computation by Adiabatic Evolution”. In: *arXiv e-prints*, quant-ph/0001106 (2000), quant-ph/0001106. arXiv: [quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106) [quant-ph].
- [47] Satoshi Morita and Hidetoshi Nishimori. “Mathematical foundation of quantum annealing”. In: *Journal of Mathematical Physics* 49.12 (2008), p. 125210.
- [48] Tameem Albash and Daniel A. Lidar. “Adiabatic quantum computation”. In: *Rev. Mod. Phys.* 90 (1 2018), p. 015002.
- [49] M. Born and V. Fock. “Beweis des Adiabatenatzes”. In: *Zeitschrift für Physik* 51.3 (1928), p. 165.
- [50] M. H. S. Amin. “Consistency of the Adiabatic Theorem”. In: *Phys. Rev. Lett.* 102 (22 2009), p. 220401.
- [51] M. S. Sarandy and D. A. Lidar. “Adiabatic Quantum Computation in Open Systems”. In: *Phys. Rev. Lett.* 95 (25 2005), p. 250503.
- [52] M. S. Sarandy and D. A. Lidar. “Adiabatic approximation in open quantum systems”. In: *Phys. Rev. A* 71 (1 2005), p. 012331.
- [53] Florian Neukart, Gabriele Compostella, Christian Seidel, David von Dollen, Sheir Yarkoni, and Bob Parney. “Traffic Flow Optimization Using a Quantum Annealer”. In: *Frontiers in ICT* 4 (2017), p. 29.
- [54] Michał Borowski et al. “New Hybrid Quantum Annealing Algorithms for Solving Vehicle Routing Problem”. In: *Computational Science – ICCS 2020*. Ed. by Valeria V. Krzhizhanovskaya et al. Cham: Springer International Publishing, 2020, pp. 546–561.
- [55] Yongcheng Ding, Xi Chen, Lucas Lamata, Enrique Solano, and Mikel Sanz. “Implementation of a Hybrid Classical-Quantum Annealing Algorithm for Logistic Network Design”. In: *SN Computer Science* 2.2 (2021), p. 68.
- [56] T. Stollenwerk et al. “Quantum Annealing Applied to De-Conflicting Optimal Trajectories for Air Traffic Management”. In: *IEEE Transactions on Intelligent Transportation Systems* 21.1 (2020), pp. 285–297.

- [57] A. Perdomo-Ortiz, N. Dickson, M. Drew-Brook, A. Aspuru-Guzik, and G. Rose. “Finding low-energy conformations of lattice protein models by quantum annealing”. In: *Scientific Reports* 2.571 (2012).
- [58] Richard Y. Li, Rosa Di Felice, Remo Rohs, and Daniel A. Lidar. “Quantum annealing versus classical machine learning applied to a simplified computational biology problem”. In: *npj Quantum Information* 4.1 (2018), p. 14.
- [59] Ryan Babbush, Peter J. Love, and Alán Aspuru-Guzik. “Adiabatic Quantum Simulation of Quantum Chemistry”. In: *Scientific Reports* 4.1 (2014), p. 6603.
- [60] L. K. Grover. “A Fast Quantum Mechanical Algorithm for Database Search”. In: *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing*. STOC '96. ACM, 1996, pp. 212–219.
- [61] B. Hayes. “The Easiest Hard Problem”. In: *American Scientist* 90 (2 2002), p. 113.
- [62] Tobias Graß, David Raventós, Bruno Juliá-Díaz, Christian Gogolin, and Maciej Lewenstein. “Quantum annealing for the number-partitioning problem using a tunable spin glass of ions”. In: *Nature Communications* 7.1 (2016), p. 11524.
- [63] Sergey Knysh. “Zero-temperature quantum annealing bottlenecks in the spin-glass phase”. In: *Nature Communications* 7.1 (2016), p. 12370.
- [64] F. Barahona. “On the Computational Complexity of Ising Spin Glass Models”. In: *Journal of Physics A: Mathematical and General* 15.10 (1982), p. 3241.
- [65] A. Lucas. “Ising formulations of many NP problems”. In: *Front. Phys.* 2 (2014), p. 5.
- [66] J. D. Biamonte. “Non Perturbative k -body to two-body Commuting Convergence Hamiltonians and embedding problem instances into Ising spins”. In: *Phys. Rev. A* 77 (5 2008), p. 052331.
- [67] S. Sachdev. *Quantum Phase Transitions*. 2nd ed. Cambridge University Press, 2011.
- [68] L. Landau. “Zur theorie der energieübertragung”. In: *Physikalische Zeitschrift der Sowjetunion* 2 (1932), pp. 46–51.
- [69] C. Zener. “Non-Adiabatic Crossing of Energy Levels”. In: *Proc. R. Soc. London, Ser. A* 137.833 (1932), pp. 696–702.
- [70] Gernot Schaller, Sarah Mostame, and Ralf Schützhold. “General error estimate for adiabatic quantum computing”. In: *Phys. Rev. A* 73 (6 2006), p. 062307.
- [71] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martinez-Garaot, and J. G. Muga. “Shortcuts to adiabaticity: Concepts, methods, and applications”. In: *Rev. Mod. Phys.* 91 (4 2019), p. 045001.

- [72] D. A. Battaglia, G. E. Santoro, and E. Tosatti. "Optimization by quantum annealing: Lessons from hard satisfiability problems". In: *Phys. Rev. E* 71 (6 2005), p. 066707.
- [73] Jiří Matoušek and Robin Thomas. "On the complexity of finding iso- and other morphisms for partial k -trees". In: *Discrete Mathematics* 108.1 (1992), pp. 343–364.
- [74] Yuki Susa, Yu Yamashiro, Masayuki Yamamoto, Itay Hen, Daniel A. Lidar, and Hidetoshi Nishimori. "Quantum annealing of the p -spin model under inhomogeneous transverse field driving". In: *Phys. Rev. A* 98 (4 2018), p. 042326.
- [75] Yuki Susa, Yu Yamashiro, Masayuki Yamamoto, and Hidetoshi Nishimori. "Exponential Speedup of Quantum Annealing by Inhomogeneous Driving of the Transverse Field". In: *Journal of the Physical Society of Japan* 87.2 (2018), p. 023002.
- [76] Zoe Gonzalez Izquierdo, Shon Grabbe, Stuart Hadfield, Jeffrey Marshall, Zhihui Wang, and Eleanor Rieffel. "Ferromagnetically Shifting the Power of Pausing". In: *Phys. Rev. Applied* 15 (4 2021), p. 044013.
- [77] Luc Fraiture. "A History of the Description of the Three-Dimensional Finite Rotation". In: *The Journal of the Astronautical Sciences* 57.1 (2009), pp. 207–232.
- [78] Andrew D. King et al. "Observation of topological phenomena in a programmable lattice of 1,800 qubits". In: *Nature* 560.7719 (2018), pp. 456–460.
- [79] Erica Grant, Travis S. Humble, and Benjamin Stump. "Benchmarking Quantum Annealing Controls with Portfolio Optimization". In: *Phys. Rev. Applied* 15 (1 2021), p. 014012.
- [80] V. Bapst, L. Foini, F. Krzakala, G. Semerjian, and F. Zamponi. "The Quantum Adiabatic Algorithm applied to Random Optimization Problems: the Quantum Spin Glass Perspective". In: *Phys. Rep.* 523.3 (2013), pp. 127–205.
- [81] Y. Seki and H. Nishimori. "Quantum annealing with antiferromagnetic fluctuations". In: *Phys. Rev. E* 85 (5 2012), p. 051112.
- [82] Shunji Matsuura, Hidetoshi Nishimori, Walter Vinci, Tameem Albash, and Daniel A. Lidar. "Quantum-annealing correction at finite temperature: Ferromagnetic p -spin models". In: *Phys. Rev. A* 95 (2 2017), p. 022308.
- [83] Shunji Matsuura, Hidetoshi Nishimori, Walter Vinci, and Daniel A. Lidar. "Nested quantum annealing correction at finite temperature: p -spin models". In: *Phys. Rev. A* 99 (6 2019), p. 062307.
- [84] Nathan Shammah, Shahnawaz Ahmed, Neill Lambert, Simone De Liberato, and Franco Nori. "Open quantum systems with local and collective incoherent processes: Efficient numerical simulations using permutational invariance". In: *Phys. Rev. A* 98 (6 2018), p. 063815.

- [85] H. J. Lipkin, N. Meshkov, and A. J. Glick. “Validity of many-body approximation methods for a solvable model: (I). Exact solutions and perturbation theory”. In: *Nuclear Physics* 62.2 (1965), pp. 188–198.
- [86] N. Meshkov, A. J. Glick, and H. J. Lipkin. “Validity of many-body approximation methods for a solvable model: (II). Linearization procedures”. In: *Nuclear Physics* 62.2 (1965), pp. 199–210.
- [87] A. J. Glick, H. J. Lipkin, and N. Meshkov. “Validity of many-body approximation methods for a solvable model: (III). Diagram summations”. In: *Nuclear Physics* 62.2 (1965), pp. 211–224.
- [88] S. Dusuel and J. Vidal. “Continuous unitary transformations and finite-size scaling exponents in the Lipkin-Meshkov-Glick model”. In: *Phys. Rev. B* 71 (22 2005), p. 224420.
- [89] S. Dusuel and J. Vidal. “Finite-Size Scaling Exponents of the LMG Model”. In: *Phys. Rev. Lett.* 93 (23 2004), p. 237204.
- [90] T. Caneva, R. Fazio, and G. E. Santoro. “Adiabatic quantum dynamics of the Lipkin-Meshkov-Glick model”. In: *Phys. Rev. B* 78 (10 2008), p. 104426.
- [91] C. R. Laumann, R. Moessner, A. Scardicchio, and S. L. Sondhi. “Quantum Adiabatic Algorithm and Scaling of Gaps at First-Order Quantum Phase Transitions”. In: *Phys. Rev. Lett.* 109 (3 2012), p. 030502.
- [92] Junichi Tsuda, Yuuki Yamanaka, and Hidetoshi Nishimori. “Energy Gap at First-Order Quantum Phase Transitions: An Anomalous Case”. In: *Journal of the Physical Society of Japan* 82.11 (2013), p. 114004.
- [93] T. Jörg, F. Krzakala, J. Kurchan, A. C. Maggs, and J. Pujos. “Energy gaps in quantum first-order mean-field-like transitions: The problems that quantum annealing cannot solve”. In: *EPL (Europhysics Letters)* 89.4 (2010), p. 40004.
- [94] B. Seoane and H. Nishimori. “Many-body transverse interactions in the quantum annealing of the p -spin ferromagnet”. In: *Jour. Phys. A* 45.43 (2012).
- [95] Elizabeth Crosson, Tameem Albash, Itay Hen, and A. P. Young. “De-Signing Hamiltonians for Quantum Adiabatic Optimization”. In: *Quantum* 4 (2020), p. 334.
- [96] N. Chancellor. “Modernizing quantum annealing using local searches”. In: *New Journal of Physics* 19.2 (2017), p. 023024.
- [97] Wolfgang Lechner, Philipp Hauke, and Peter Zoller. “A quantum annealing architecture with all-to-all connectivity from local interactions”. In: *Science advances* 1.9 (2015), e1500838–e1500838.
- [98] Andrea Rocchetto, Simon C Benjamin, and Ying Li. “Stabilizers as a design tool for new forms of the Lechner-Hauke-Zoller annealer”. In: *Science advances* 2.10 (2016), e1601246–e1601246.
- [99] Simon E. Nigg, Niels Lörch, and Rakesh P. Tiwari. “Robust quantum optimizer with full connectivity”. In: *Science advances* 3.4 (2017), e1602273–e1602273.

- [100] A. W. Glaetzle, R. M. W. van Bijnen, P. Zoller, and W. Lechner. “A coherent quantum annealer with Rydberg atoms”. In: *Nature communications* 8 (2017), pp. 15813–15813.
- [101] Martin Leib, Peter Zoller, and Wolfgang Lechner. “A transmon quantum annealer: decomposing many-body Ising constraints into pair interactions”. In: *Quantum Science and Technology* 1.1 (2016), p. 015008.
- [102] David E. Goldberg and John H. Holland. “Genetic Algorithms and Machine Learning”. In: *Machine Learning* 3.2 (1988), pp. 95–99.
- [103] Yorick Hardy and W-H. Steeb. “Genetic algorithms and optimization problems in quantum computing”. In: *International Journal of Modern Physics C* 21.11 (2010), pp. 1359–1375.
- [104] Félix-Antoine Fortin, François-Michel De Rainville, Marc-André Gardner, Marc Parizeau, and Christian Gagné. “DEAP: Evolutionary Algorithms Made Easy”. In: *Journal of Machine Learning Research* 13.70 (2012), pp. 2171–2175.
- [105] Xin Yao. “An empirical study of genetic operators in genetic algorithms”. In: *Microprocessing and Microprogramming* 38.1 (1993). Proceedings Euromicro 93 Open System Design: Hardware, Software and Applications, pp. 707–714.
- [106] F. Herrera, M. Lozano, and A. M. Sánchez. “A taxonomy for the crossover operator for real-coded genetic algorithms: An experimental study”. In: *International Journal of Intelligent Systems* 18.3 (2003), pp. 309–338.
- [107] Tameem Albash, Victor Martin-Mayor, and Itay Hen. “Temperature Scaling Law for Quantum Annealing Optimizers”. In: *Phys. Rev. Lett.* 119 (11 2017), p. 110502.
- [108] Tadashi Kadowaki. “Dynamics of open quantum systems by interpolation of von Neumann and classical master equations, and its application to quantum annealing”. In: *Phys. Rev. A* 97 (2 2018), p. 022312.
- [109] Walter Vinci and Daniel A. Lidar. “Scalable effective-temperature reduction for quantum annealers via nested quantum annealing correction”. In: *Phys. Rev. A* 97 (2 2018), p. 022308.
- [110] T. Albash and D. A. Lidar. “Decoherence in adiabatic quantum computation”. In: *Phys. Rev. A* 91 (6 2015), p. 062320.
- [111] Howard Carmichael. *An open systems approach to quantum optics*. Lecture notes in physics monographs. Berlin: Springer, 1993.
- [112] R. Dum, P. Zoller, and H. Ritsch. “Monte Carlo simulation of the atomic master equation for spontaneous emission”. In: *Phys. Rev. A* 45 (7 1992), pp. 4879–4887.
- [113] Klaus Mølmer, Yvan Castin, and Jean Dalibard. “Monte Carlo wave-function method in quantum optics”. In: *J. Opt. Soc. Am. B* 10.3 (1993), pp. 524–538.

- [114] Mathias Michel, Ortwin Hess, Hannu Wichterich, and Jochen Gemmer. “Transport in open spin chains: A Monte Carlo wave-function approach”. In: *Phys. Rev. B* 77 (10 2008), p. 104303.
- [115] Jyrki Piilo, Sabrina Maniscalco, Kari Härkönen, and Kalle-Antti Suominen. “Non-Markovian Quantum Jumps”. In: *Phys. Rev. Lett.* 100 (18 2008), p. 180402.
- [116] A. J. Leggett, S. Chakravarty, A. T. Dorsey, Matthew P. A. Fisher, Anupam Garg, and W. Zwerger. “Dynamics of the dissipative two-state system”. In: *Rev. Mod. Phys.* 59 (1 1987), pp. 1–85.
- [117] Steven R. White and Adrian E. Feiguin. “Real-Time Evolution Using the Density Matrix Renormalization Group”. In: *Phys. Rev. Lett.* 93 (7 2004), p. 076401.
- [118] U. Schollwöck. “The density-matrix renormalization group”. In: *Rev. Mod. Phys.* 77 (1 2005), pp. 259–315.
- [119] Nancy Makri. “Numerical path integral techniques for long time dynamics of quantum dissipative systems”. In: *Journal of Mathematical Physics* 36.5 (1995), pp. 2430–2457.
- [120] Lothar Mühlbacher and Eran Rabani. “Real-Time Path Integral Approach to Nonequilibrium Many-Body Quantum Systems”. In: *Phys. Rev. Lett.* 100 (17 2008), p. 176403.
- [121] Ting Yu. “Non-Markovian quantum trajectories versus master equations: Finite-temperature heat bath”. In: *Phys. Rev. A* 69 (6 2004), p. 062107.
- [122] Heinz-Peter Breuer, Bernd Kappler, and Francesco Petruccione. “Stochastic wave-function method for non-Markovian quantum master equations”. In: *Phys. Rev. A* 59 (2 1999), pp. 1633–1643.
- [123] Bassano Vacchini and Heinz-Peter Breuer. “Exact master equations for the non-Markovian decay of a qubit”. In: *Phys. Rev. A* 81 (4 2010), p. 042103.
- [124] Rui-Xue Xu and YiJing Yan. “Dynamics of quantum dissipation systems interacting with bosonic canonical bath: Hierarchical equations of motion approach”. In: *Phys. Rev. E* 75 (3 2007), p. 031107.
- [125] Kiyoto Nakamura and Yoshitaka Tanimura. “Hierarchical Schrödinger equations of motion for open quantum dynamics”. In: *Phys. Rev. A* 98 (1 2018), p. 012109.
- [126] Inés de Vega, Ulrich Schollwöck, and F. Alexander Wolf. “How to discretize a quantum bath for real-time evolution”. In: *Phys. Rev. B* 92 (15 2015), p. 155126.
- [127] Ralf Bulla, Ning-Hua Tong, and Matthias Vojta. “Numerical Renormalization Group for Bosonic Systems and Application to the Sub-Ohmic Spin-Boson Model”. In: *Phys. Rev. Lett.* 91 (17 2003), p. 170601.

- [128] Matthias Vojta, Ning-Hua Tong, and Ralf Bulla. “Quantum Phase Transitions in the Sub-Ohmic Spin-Boson Model: Failure of the Quantum-Classical Mapping”. In: *Phys. Rev. Lett.* 94 (7 2005), p. 070604.
- [129] Yu-Yu Zhang, Qing-Hu Chen, and Ke-Lin Wang. “Quantum phase transition in the sub-Ohmic spin-boson model: An extended coherent-state approach”. In: *Phys. Rev. B* 81 (12 2010), p. 121105.
- [130] G. D. Mahan. *Many-particle physics*. 3rd ed. Springer, 2000.
- [131] A. O. Caldeira and A. J. Leggett. “Quantum tunnelling in a dissipative system”. In: *Annals of Physics* 149.2 (1983), p. 374.
- [132] Tameem Albash and Jeffrey Marshall. “Comparing Relaxation Mechanisms in Quantum and Classical Transverse-Field Annealing”. In: *Phys. Rev. Applied* 15 (1 2021), p. 014029.
- [133] S. Jansen, M. B. Ruskai, and R. Seiler. “Bounds for the adiabatic approximation with applications to quantum computation”. In: *Journal of Mathematical Physics* 48.10 (2007), p. 102111.
- [134] Lorenzo Campos Venuti, Tameem Albash, Daniel A. Lidar, and Paolo Zanardi. “Adiabaticity in open quantum systems”. In: *Phys. Rev. A* 93 (3 2016), p. 032118.
- [135] Lorenzo Campos Venuti and Daniel A. Lidar. “Error reduction in quantum annealing using boundary cancellation: Only the end matters”. In: *Phys. Rev. A* 98 (2 2018), p. 022315.
- [136] *Andree Richmond - Clay Artist*. <https://andreerichmond.com>. Accessed: 2020-07-29.
- [137] Erik Torrontegui et al. “Chapter 2 - Shortcuts to Adiabaticity”. In: *Advances in Atomic, Molecular, and Optical Physics*. Ed. by Ennio Arimondo, Paul R. Berman, and Chun C. Lin. Vol. 62. Advances In Atomic, Molecular, and Optical Physics. Academic Press, 2013, pp. 117–169.
- [138] J. R. Johansson, P. D. Nation, and F. Nori. “QuTiP: An open-source Python framework for the dynamics of open quantum systems”. In: *Computer Physics Communications* 183.8 (2012), pp. 1760–1772.
- [139] J. R. Johansson, P. D. Nation, and F. Nori. “QuTiP 2: A Python framework for the dynamics of open quantum systems”. In: *Computer Physics Communications* 184.4 (2013), pp. 1234–1240.
- [140] Michael H. Goerz et al. “Krotov: A Python implementation of Krotov’s method for quantum optimal control”. In: *SciPost Phys.* 7 (6 2019), p. 80.
- [141] Shai Machnes, Elie Assémat, David Tannor, and Frank K. Wilhelm. “Tunable, Flexible, and Efficient Optimization of Control Pulses for Practical Qubits”. In: *Physical Review Letters* 120.15 (2018).
- [142] J. J. W. H. Sørensen, M. O. Aramburu, T. Heinzl, and J. F. Sherson. “Quantum optimal control in a chopped basis: Applications in control of Bose-Einstein condensates”. In: *Physical Review A* 98.2 (2018).

- [143] Mustafa Demirplak and Stuart A. Rice. “Assisted Adiabatic Passage Revisited”. In: *The Journal of Physical Chemistry B* 109.14 (2005). PMID: 16851769, pp. 6838–6844.
- [144] Mustafa Demirplak and Stuart A. Rice. “On the consistency, extremal, and global properties of counterdiabatic fields”. In: *The Journal of Chemical Physics* 129.15 (2008), p. 154111.
- [145] Adolfo del Campo. “Shortcuts to Adiabaticity by Counterdiabatic Driving”. In: *Phys. Rev. Lett.* 111 (10 2013), p. 100502.
- [146] Hamed Saberi, Tomá š Opatrný, Klaus Mølmer, and Adolfo del Campo. “Adiabatic tracking of quantum many-body dynamics”. In: *Phys. Rev. A* 90 (6 2014), p. 060301.
- [147] A. del Campo and K. Sengupta. “Controlling quantum critical dynamics of isolated systems”. In: *The European Physical Journal Special Topics* 224.1 (2015), pp. 189–203.
- [148] Kohji Nishimura and Kazutaka Takahashi. “Counterdiabatic Hamiltonians for multistate Landau-Zener problem”. In: *SciPost Phys.* 5 (3 2018), p. 29.
- [149] Takuya Hatomura. “Shortcuts to Adiabaticity in the Infinite-Range Ising Model by Mean-Field Counter-Diabatic Driving”. In: *Journal of the Physical Society of Japan* 86.9 (2017), p. 094002.
- [150] Kazutaka Takahashi. “Shortcuts to adiabaticity for quantum annealing”. In: *Phys. Rev. A* 95 (1 2017), p. 012309.
- [151] T. Opatrný and K. Mølmer. “Partial suppression of nonadiabatic transitions”. In: *New Journal of Physics* 16.1 (2014), p. 015025.
- [152] Andreas Hartmann and Wolfgang Lechner. “Rapid counter-diabatic sweeps in lattice gauge adiabatic quantum computing”. In: *New Journal of Physics* 21.4 (2019), p. 043025.
- [153] Takuya Hatomura and Kazutaka Takahashi. “Controlling and exploring quantum systems by algebraic expression of adiabatic gauge potential”. In: *Phys. Rev. A* 103 (1 2021), p. 012220.
- [154] M. Suzuki. “Relationship between d -Dimensional Quantal Spin Systems and $(d + 1)$ -Dimensional Ising Systems: Equivalence, Critical Exponents and Systematic Approximants of the Partition Function and Spin Correlations”. In: *Progress of Theoretical Physics* 56.5 (1976), p. 1454.