

QUANTUM ANNEALING COMPLEXITY AND STOCHASTIC QUANTUM CONTROL

Aarón Villanueva



DONDERS
SERIES

**RADBOUD
UNIVERSITY
PRESS**

Quantum Annealing Complexity and Stochastic Quantum Control

Aarón Villanueva

The research described in this thesis was carried out at The Donders Institute, Faculty of Science, Radboud University, Nijmegen, The Netherlands, and at the Theoretical Division (T-4), Los Alamos National Laboratory, Los Alamos, New Mexico, USA, with financial support from the U.S. Department of Energy (DOE) through a quantum computing program sponsored by the Los Alamos National Laboratory (LANL) Information Science & Technology Institute.

Uziel Aarón Villanueva

Quantum Annealing Complexity and Stochastic Quantum Control

Radboud Dissertation Series

ISSN: 2950-2772 (Online); 2950-2780 (Print)

Published by RADBOUD UNIVERSITY PRESS

Postbus 9100, 6500 HA Nijmegen, The Netherlands

www.radbouduniversitypress.nl

Design: Aarón Villanueva

Cover image: Aylin Villanueva | aylin98zahira@gmail.com

Printing: ProefschriftMaken/Pumbo

ISBN: 9789465153179

DOI: 10.54195/9789465153179

Free download at: <https://doi.org/10.54195/9789465153179>

© 2026 Aarón Villanueva

**RADBOUD
UNIVERSITY
PRESS**

This is an Open Access book published under the terms of Creative Commons Attribution-Noncommercial-NoDerivatives International license (CC BY-NC-ND 4.0). This license allows reusers to copy and distribute the material in any medium or format in unadapted form only, for noncommercial purposes only, and only so long as attribution is given to the creator, see <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

Quantum Annealing Complexity and Stochastic Quantum Control

Proefschrift ter verkrijging van de graad van doctor
aan de Radboud Universiteit Nijmegen
op gezag van de rector magnificus prof. dr. J.M. Sanders,
volgens besluit van het college voor promoties
in het openbaar te verdedigen op

woensdag 18 november 2026
om 16.30 uur precies

door

Uziel Aarón Villanueva
geboren op 1 december 1993

te San Salvador de Jujuy, Argentinië

Promotor:

Prof. dr. Hilbert J. Kappen

Copromotor:

Dr. E.J. Dominguez Vázquez

Manuscriptcommissie:

Prof. dr. A.C.C. Coolen (voorzitter)

Prof. dr. V. Dunjko, Universiteit Leiden

Dr. R. Wiersema, Flatiron Institute, Verenigde Staten



Quantum Annealing Complexity and Stochastic Quantum Control

Dissertation to obtain the degree of doctor
from Radboud University Nijmegen
on the authority of the Rector Magnificus prof. dr. J.M. Sanders,
according to the decision of the Doctorate Board
to be defended in public on

Wednesday, November 18, 2026
at 4.30 pm

by

Uziel Aarón Villanueva
born on 1 December 1993

in San Salvador de Jujuy, Argentina

Supervisor:

Prof. dr. Hilbert J. Kappen

Co-supervisor:

Dr. E.J. Dominguez Vázquez

Manuscript Committee:

Prof. dr. A.C.C. Coolen (chair)

Prof. dr. V. Dunjko, Universiteit Leiden

Dr. R. Wiersema, Flatiron Institute, United States

The research described in this thesis was carried out at The Donders Institute, Faculty of Science, Radboud University, Nijmegen, The Netherlands, and at the Theoretical Division (T-4), Los Alamos National Laboratory, Los Alamos, New Mexico, USA, with financial support from the U.S. Department of Energy (DOE) through a quantum computing program sponsored by the Los Alamos National Laboratory (LANL) Information Science & Technology Institute.

Cover layout

ProefschriftOntwerp

Cover background design

Aylin Villanueva

contact: aylin98zahira@gmail.com

Thesis layout

Aarón Villanueva

Printed by

RUP / ProefschriftOntwerp

*Soñaba un océano y no el cielo,
el terror de un espacio incontable
y un tiempo de arena viva.
Despierto.
El rizoma deviene en verbo,
en la epopeya del héroe
y el júbilo sensible de la pregunta develada.*

Preface

This thesis is the culmination of the research conducted over the past four years during my PhD at Radboud University. The thesis is structured in a publication-based format: each chapter corresponds to a single paper and is designed to be self-contained.

I would like to thank my family, whose love and support form a long and uninterrupted thread of nurturing events in spacetime. Special thanks to my parents, for the freedom of thought and the sense of responsibility. I thank Bert for letting me pursue this adventure. For the support over these years, for the wisdom, and the critical eye. I thank my friends in The Netherlands for all the fun, the music and the experience. I thank Luis Pedro García-Pintos for the amazing collaboration and to Los Alamos National Laboratory for hosting me.

My endless gratitude to my beloved Elena, the person who wielded love and patience throughout both the highs and lows of this journey. Thank you for the trips, the Italian food and the temper. We composed together a beautiful song over these years!

Aarón
Nijmegen, The Netherlands
January 2026

Contents

Preface	7
1 Introduction	12
2 Why Adiabatic Quantum Annealing is unlikely to yield speed-up	16
2.1 Introduction	17
2.2 Adiabatic QA: preliminary definitions	20
2.3 An adiabatic algorithm for combinatorial optimization	21
2.3.1 The model	22
2.3.2 Fluctuations in the minimal gap location	24
2.3.3 Quadratic speedup for generic optimization problems	26
2.4 Comparison with the common mixing Hamiltonian	29
2.5 Conclusions and discussion	32
2.6 Acknowledgments	33
2.7 Appendices	34
2.7.1 Model analysis	34
2.7.2 Quadratic speed-up: proof of existence	43
2.7.3 Quadratic speed-up: proof of optimality	45
2.7.4 Random 3-SAT: density of states approximation	47
3 Stochastic optimal control of open quantum systems	52
3.1 Introduction	53
3.1.1 Contributions	55
3.2 The path integral control method	58
3.2.1 Parametrized controls and optimality	62

3.2.2	Adaptive importance sampling and open-loop control	64
3.3	Path integral control of quantum systems	65
3.3.1	Problem definition	65
3.3.2	Stochastic mapping	67
3.3.3	The Path integral Quantum Control algorithm for open systems	68
3.3.4	Annealed diffusion control for closed systems	69
3.3.5	A proposal for on-hardware PiQC	71
3.4	Numerical experiments	74
3.4.1	Control of a noisy qubit	74
3.4.2	Leveraging open-system solutions for unitary dynamics	81
3.4.3	GHZ states on open NMR systems	83
3.4.4	Application of annealed diffusion control	86
3.5	Discussions and Outlook	87
3.6	Acknowledgements	90
3.7	Data and code availability	90
3.8	Appendices	91
3.8.1	Unravelings of the controlled Lindblad equation	91
3.8.2	Path integral control theory	92
3.8.3	Open GRAPE	95
3.8.4	PiQC robustness example	96
3.8.5	Open NMR control theory	98
3.8.6	A 10 qubits example	103

4 Hamiltonian and Double-Bracket Flow Formulations of Quantum Measurements 106

4.1	Introduction	107
4.2	Dynamics of continuously monitored quantum systems	108
4.3	Hamiltonian formulation of quantum measurements	110
4.4	Connection with double-bracket gradient flows	114
4.4.1	Background	115
4.4.2	Results	116
4.5	Applications to feedback control	122

4.5.1	Feedback-induced DB flows	122
4.5.2	State-agnostic feedback for state preparation	124
4.5.3	Gradient flow interpretation of feedback dynamics	128
4.6	Discussion	131
4.7	Appendices	132
4.7.1	Monitored dynamics in the Itô picture	132
4.7.2	Monitored dynamics in the Stratonovich picture	134
4.7.3	Itô-Stratonovich conversion formula	136
4.7.4	A stochastic Hamiltonian that describes monitored dynamics	137
4.7.5	Double-bracket gradient flows	139
4.7.6	A simple proof of Theorem 6	143
4.7.7	Stability analysis of the measurement-based feedback dynamics	144
5	Conclusion	147
6	Nederlandse samenvatting	151
7	Research data management	153
7.1	Scope, ethics and governance	153
7.2	Ethics, privacy, and funding	154
7.3	Findability and accessibility	154
7.4	Interoperability and reusability	156
8	Curriculum Vitae	158
9	List of publications and collaborations	166
10	Donders Graduate School	167

Chapter 1

Introduction

The field of quantum computing has seen significant advances in recent years, with major breakthroughs in quantum simulation and quantum error correction [1–8]. Despite these milestones, the search for practically relevant applications of quantum computers remains a vital question [9, 10].

In this context, this work examines two fundamental areas that play a pivotal role in extending the practical relevance of quantum computing—combinatorial optimization and quantum optimal control. Combinatorial optimization (CO) seeks the minimum-cost solution within a discrete space of exponentially growing possibilities. By encoding these many possibilities into the space of accessible states of finite quantum systems, it is possible to design analog and digital quantum algorithms that solve combinatorial optimization (hopefully) faster than classical algorithms [11, 12]. We aim to characterize the runtime performance or speedup of specific quantum algorithms relative to classical methods for solving combinatorial optimization problems. On the other hand, quantum optimal control (QOC) concerns the development of theoretical and experimental techniques to optimally steer quantum systems to achieve a desired behavior [13–16]. In this regard, we aim to develop and analyze novel protocols for the optimal control of both closed and open quantum systems.

Concretely, for combinatorial optimization, in Chapter 2 we focus on a specific model for analog quantum computing within the adiabatic quantum

computation (AQC) framework [17–19]. AQC works by encoding the solution to the optimization problem into the ground state of a target Hamiltonian. In AQC, the quantum computer is represented by a time-dependent Hamiltonian that interpolates between an initial Hamiltonian, whose ground state is easy to prepare, and the target Hamiltonian encoding the solution to the optimization problem. The computation proceeds through the unitary time evolution of the quantum hardware, whose dynamics are governed by the Schrödinger equation. The adiabatic theorem ensures that, if the evolution is sufficiently slow and the system is initialized in the ground state of the initial Hamiltonian, the state remains arbitrarily close to the instantaneous ground state throughout the evolution. Specifically, we study the quantum speedup of an annealing algorithm to solve the 3-SAT problem whose initial Hamiltonian is agnostic of the structure of the problem instance. We establish rigorous lower bounds to the time complexity of the quantum algorithm. We establish a universal quadratic quantum advantage over classical, and prove existence of an annealing schedule that saturates this bound. Moreover, we argue that saturating this bound is hard in practice due to the exponential precision of the schedule construction.

For quantum control, in Chapter 3 we introduce a novel approach for the optimal control of quantum systems that combines path integral control with quantum trajectories. Path integral (PI) control is a powerful sampling-based trajectory optimization approach for stochastic optimal control [20–22]. It enables the design of sampling schemes where the optimal controls are computed as weighted averages over sample trajectories (path integrals), thus avoiding the need to solve the Hamilton-Jacobi-Bellman equation [23]. PI control finds extensive and successful use in multiple fields of industrial significance [24–33]. Ref. [34] is a recent review. Conversely, the quantum trajectories approach characterizes the dynamics of open quantum systems in terms of stochastic realizations of their time evolution [35–38]. We combine PI control with quantum trajectories to build a new framework for optimal quantum control of closed and open systems. The proposed framework computes optimal open-loop controls for a broad class of Lindblad dynamics using adaptive importance sampling. The mechanism consists of iteratively

updating our knowledge of the control fields by probing the control landscape with quantum trajectories and collecting relevant statistics dictated by the PI control theory. We further introduce an annealed variant of the algorithm for computing optimal controls in closed systems, in which the environmental couplings are gradually reduced during optimization to approximate closed-system dynamics. We name the resulting framework Path integral Quantum Control (PiQC). We validate the theory on realistic quantum systems such as nuclear magnetic resonance molecules.

Continuing our investigation of quantum control, in Chapter 4 we analyze a stochastic Hamiltonian formulation of continuous quantum measurements. As in our development of PiQC, we again employ quantum trajectories, now within the framework of monitored quantum systems [39, 40], which describe the evolution of open quantum systems undergoing continuous, randomly occurring projective von Neumann measurements. We can think of this process as performing short projective measurements on the composite state arising from the system–environment interaction. The system transitions from the pre-measurement state to a post-measurement state, while the environment state is projected onto an eigenstate of the measuring apparatus. In this way, we obtain information about the system without necessarily projecting it onto an eigenspace of the probe. Although, in general, the dynamics of open quantum systems are governed by equations that extend beyond purely Hamiltonian evolution, the equations for the monitored dynamics describe stochastic unitary processes when the state is initially pure. This enables a stochastic Hamiltonian description of the aforementioned dynamics. We provide explicit expressions for the Hamiltonian dynamics and establish a new connection with the theory of double-bracket flows [41, 42] a framework of growing relevance in quantum computing [43–45]. This connection allows us to reinterpret the Hamiltonian dynamics as a gradient flow on the unitary orbit of the system’s state. Finally, we show that the stochastic Hamiltonian formulation suggests novel pathways for feedback control. Concretely, we design feedback control protocols for state preparation and/or stabilization. We analyze the Lyapunov stability of the resulting dynamics and propose protocols that converge to the target state while remaining agnostic to the

system's current state, thereby addressing a common limitation in optimal control.

In Chapter 5 we conclude with an overall assessment of work and an outlook on future prospects, and in Chapter 9 we list all the publications, including collaborations, of the author.

Chapter 2

Why Adiabatic Quantum Annealing is unlikely to yield speed-up

Aarón Villanueva¹, Peyman Najafi¹, Hilbert J. Kappen¹

¹Faculty of Science, Radboud University, Heyendaalseweg 135, 6525 AJ Nijmegen, The Netherlands

Correspondence: aaronv@science.ru.nl

We study quantum annealing for combinatorial optimization with Hamiltonian $H = H_0 + zH_f$ where H_f is diagonal, $H_0 = -|\phi\rangle\langle\phi|$ is the equal superposition state projector and z the annealing parameter. We analytically compute the minimal spectral gap, which is $\mathcal{O}(1/\sqrt{N})$ with N the total number of states, and its location z_* . We show that quantum speed-up requires an annealing schedule which demands a precise knowledge of z_* , which can be computed only if the density of states of the optimization problem is known. However, in general the density of states is intractable to compute, making quadratic speed-up unfeasible for any practical combinatorial optimization problems. We conjecture that it is likely that this negative result also applies for any other instance independent transverse Hamiltonians such as $H_0 = -\sum_{i=1}^n \sigma_i^x$.

2.1 Introduction

Combinatorial optimization problems emerge in many real world scenarios and solving them efficiently is of great interest and importance in science and engineering [46]. Simulated Annealing (SA) is a well-known heuristic classical algorithm, which is general purpose and easy to implement [47]. The ergodicity of the Markov chain (thermal fluctuations) combined with a sufficiently slow annealing schedule, ensures that SA converges asymptotically to an optimal solution of the optimization problem. In practice, the total time and the efficiency of SA is determined by the mixing time of the Markov chain, which diverges for hard optimization problems [48].

The idea of quantum annealing (QA) [17, 49, 50] was introduced motivated by the conjecture that quantum fluctuations can be more efficient than classical thermal fluctuations to solve hard optimization problems [51, 52]. This would constitute in a quantum advantage over classical algorithms. Today, a well known example of quantum advantage is unstructured (Grover) search of one out of N items which can be done in $\mathcal{O}(\sqrt{N})$ time using a quantum computer [53]. It was shown in [54] that the same quadratic speed-up can be obtained using quantum annealing (QA).

NP-complete problems are defined by an universality property: the existence of an efficient algorithm to solve them would solve any other hard problem instantly. The question of whether QA can be used to solve efficiently NP-complete problems is a difficult topic that has remained open. An archetypal example is 3-satisfiability (3-SAT). Early applications of QA to 3-SAT indicated speed-up over classical algorithms [17, 18, 55–57]. However, such studies were restricted either using a small number of qubits or considering only a subset of typical instances while ignoring worst cases [58].

The question was also studied using Simulated QA, which uses quantum Monte Carlo. Initially, a better scaling of Simulated QA relative to SA was

reported in [59, 60] to find the lowest energy configuration for the 2d random Ising model and in [61] for random 3-SAT. However, it was reported later that this advantage was because of time discretization (Trotterization) artifacts, and in the continuous time limit there is no superiority for QA [62]. Around the same time, it was argued that QA is exponentially slow for solving NP-hard problems due to exponential closing of the spectral gap with increasing system size [63, 64].

One can build examples in which QA is exponentially faster than SA. For instance, the "Hamming weight with a spike" problem [19, 65], which has a high thin barrier that makes SA get stuck in a local minimum. It is designed to make the SA algorithm fail while QA can tunnel through the spike and is effective. Nevertheless, this problem is not NP-hard and can be efficiently solved by other classical algorithms [66, 67].

Further numerical studies show that adiabatic QA performs worse than SA on 3-SAT [68] and the antiferromagnet spin model on a 3-regular graph [69]. An experimental comparison of SA and QA on the D-Wave [70, 71] quantum annealing for random spin glass instances also reported no quantum speed-up [72].

Based on these negative results, efforts have been made to improve QA by instance specific initialization [73], the use of a more informed choice of the initial Hamiltonian [74], the use of an instance specific biasing σ^z field [75], other strategies which might allow an escape from bottlenecks or trap states [76], increase of the dimensionality of the Hamiltonian parametrization [77] or change the adiabatic path by adding a catalyst Hamiltonian which vanishes at the edges of the annealing [78]. Although these techniques exhibited cases of success in simple or ad-hoc problems, so far this has not yielded a demonstration of quantum speed for larger instances.

Due to the uncertainty surrounding adiabatic methods, in recent years there has been explored other strategies such as non-adiabatic or *diabatic* quantum computing [66, 79–82]. In the diabatic case, transitions between the ground and excited states are allowed in regions where the gap becomes exponentially small, therefore avoiding the time complexity carried by adiabatic evolution. An example is the glued-tree problem [83], for which there

exists a mixed adiabatic and diabatic evolution schedule that offers to date the only *provable* [72] quantum (exponential) speedup over classical algorithms [84]. It is worth noticing that a superpolynomial quantum speed-up was proven in [85] for stoquastic adiabatic computation in the oracle setting. More recently, following the same lines of [85], in [86] a sub-exponential quantum speed-up was proven in the stoquastic adiabatic framework for a problem in close relation with glued-trees.

While the removal of the adiabatic condition gives more freedom in building quantum algorithms and has proven advantage in some cases, the question whether diabatic methods can solve NP-hard problems in general is still unanswered.

Even in the adiabatic case, the question whether it can give quantum speedup for hard combinatorial optimization problems remains an open question to date [84]. Apart from some specific simplified cases, such as Grover search [54], there is no convincing demonstration of quantum advantage using adiabatic methods for generic hard problems.

It was shown for instance that QA can not give better than quadratic speed-up for a very general model Hamiltonian with linear schedule [74, 87]. Since this result is a lower bound on the time complexity, it leaves open the question of whether quadratic speed-up using QA for generic combinatorial optimization problems is possible at all. In this paper we explore this question and respond positively. We study a model Hamiltonian composed of a diagonal term plus the equal superposition state projector. We compute the minimal spectral gap and its location z_* in the annealing interval. We provide an analytical expression for the spectral gap, which is valid in a vicinity of the minimal value. We prove that the minimal gap is $\mathcal{O}\left(1/\sqrt{N}\right)^1$ which happens at $z_* = Z_1$ with Z_1 a function of the density of states. When the density of states is known, we can design a schedule that achieves $T = \mathcal{O}\left(\sqrt{N}\right)$, i.e. quadratic speedup over naive search for any combinatorial optimization problem. This generalizes earlier results in the case of Grover search [54]. Furthermore, we prove that this complexity is

¹We make the convention of omitting non-exponential corrections in complexity estimations, for instance $\mathcal{O}(\text{poly}(n))$ or $\mathcal{O}(\log n)$ factors.

optimal, generalizing a previous result that holds in the linear regime [74] to the case of arbitrary non-linear schedules.

However, the annealing schedule depends sensitively on the precise location of the gap $z_* = Z_1$. In general the density of states is instance dependent and is intractable to compute. Therefore Z_1 cannot be computed and we cannot build the optimized schedule. This issue was explored in [88] in the context of adiabatic unstructured search, leading to similar conclusions about the sensitivity of the annealing schedule. We argue that our result is likely to hold also for the more commonly used mixing Hamiltonian $H_0 = -\sum_{i=1}^n \sigma_i^x$ and we support this claim with numerical results (see Sec. 2.4).

The paper is organized as follows: In Section 2.2 we give a short introduction on adiabatic quantum annealing and define pertinent concepts that will be used later. In Section 2.3 we define the adiabatic model subject to study and describe its main properties including the spectral gap. In Section 2.3.3 we take advantage of the precise knowledge about the gap to prove the existence of a schedule that guarantees quadratic speedup over naive search for any problem instance. Furthermore, we prove that quadratic speedup is optimal, meaning there is no schedule that achieves better time complexity than $\mathcal{O}(\sqrt{N})$. In Section 2.4 we support our claims with numerical simulations using a transverse field Hamiltonian as the initial Hamiltonian. In Section 2.5 we summarize our results.

2.2 Adiabatic QA: preliminary definitions

In a standard adiabatic QA algorithm one specifies a time-dependent Hamiltonian $H(z)$ that evolves according to the Schrödinger equation. The parameter $z = z(\beta)$ is the *annealing schedule* and is a function of the dimensionless variable $\beta := t/T$, where t is the time and T the total evolution time. In the range $\beta \in [0, 1]$ the path $H(z)$ connects two Hamiltonians, $H_0 := H(z_i)$ and $H_f := H(z_f)$, where $z_i = z(0)$ and $z_f = z(1)$. The initial *mixing* Hamiltonian H_0 has a ground state that is easy to prepare. The final *target* Hamiltonian H_f encodes in its ground state the solution to the optimization problem. The

goal of QA is to transform the ground state of H_0 to a state which is close to the ground state of H_f at time $t = T$ by evolving the Schrödinger equation.

We say that the evolution is adiabatic if it is carried out slowly. The error induced by this adiabatic evolution is given by the adiabatic theorem (AT) [89, 90]. The story of the progress towards a mathematically clean AT is full of twists and turns. The reader can consult [84] for a review. In this work we use a rigorous version of the AT due to Jansen, Ruskai and Seiler [91] as it appears in [84]. The AT gives a bound on the adiabatic error in terms of the gap and derivatives of the Hamiltonian. It states that for achieving an error less than ε in the comparison between the evolved state and the desired final ground state, the total evolution time T must satisfy

$$T \geq \frac{C}{\varepsilon}, \quad C := 2 \max_{\beta} \frac{\|\dot{H}(\beta)\|}{g(\beta)^2} + \int_0^1 \left(\frac{\|\ddot{H}\|}{g^2} + \frac{7\|\dot{H}\|^2}{g^3} \right) d\beta, \quad (2.1)$$

where g denotes the spectral gap of H and where the dot represents derivative with respect to β . The value of C is a measure of the time complexity of the model, but it doesn't work as a general definition of the complexity/cost of the algorithm since it is not scale-invariant (one can make the complexity arbitrarily small by changing the time scale). A proper measure of a scale-invariant cost is $T \max_z \|H(z)\|$ [84], where the maximum is taken in the interval $[z_i, z_f]$. In cases where the maximum norm of the Hamiltonian is $\mathcal{O}(\text{poly}(n))$ with n the number of qubits, as it is for the present model, we can simply take $\text{cost} = T$ (we ignore non-exponential corrections in our complexity estimations). Hereby, we take T as a well-defined measure of the cost of the model.

2.3 An adiabatic algorithm for combinatorial optimization

2.3.1 The model

Consider an arbitrary optimization problem of the form

$$s^* = \operatorname{argmin}_s E(s)$$

with $s = (s_1, \dots, s_n)$ a vector of n binary variables and $E(s)$ is an integer in the range $0 \leq E(s) \leq m$ with $m = \operatorname{poly}(n)$. Define N_E the number of states with energy E as $N_E = \sum_s \delta_{E(s), E}$, subject to $\sum_{E=0}^m N_E = N$ and $N = 2^n$. An example is 3-satisfiability (3-SAT) with $m \propto n$ clauses where $E(s) = \sum_{a=1}^m e_a(s)$ and each $e_a(s) = 0, 1$ depends on 3 spins. In this work we assume that $N_0 > 0$, i.e. we have at least one satisfying assignment.

We encode the optimization problem in an adiabatic model defining a spin Hamiltonian ²

$$H = H_0 + zH_f \quad (2.2)$$

where $H_f = \sum_s E(s) |s\rangle \langle s|$ is diagonal in the σ^z -basis of n spins, and $H_0 = -|\phi\rangle \langle \phi|$ with $|\phi\rangle = \frac{1}{\sqrt{N}} \sum_s |s\rangle$ the equal superposition state. The parameter z is the annealing parameter that changes with time t varied from $z = 0$ to its final value z_f . For $z = 0$, H has ground state $|\phi\rangle$. For large z , the ground state of H encodes the minimal energy solution.

The spectral analysis of model (2.2) can be performed with some precision. For details consult Appendix 2.7.1. Here we outline the main points of the proof and refer the reader to Appendix 2.7.1 for details. The problem Hamiltonian $H(z)$ has a special form that allows to specify the first and second eigenvalues needed to compute the gap as the first two roots of a characteristic non-linear equation. Due to a permutation symmetry present in $H(z)$ we can prove that the gap is given by $g = \lambda_1 - \lambda_0$, where λ_0 and λ_1 are the first two roots of the characteristic equation and satisfy $\lambda_0 < 0 < \lambda_1$ when $z > 0$. When the gap becomes small, each eigenvalue becomes small and close to zero. We can then perform a Taylor approximation of the char-

² H is related to the usual form for quantum annealing $H = (1 - A)H_0 + AH_f$ with $0 \leq A \leq 1$ by defining $z = \frac{A}{1-A}$.

acteristic equation and compute λ_0 and λ_1 with high precision for large n . Then, given an instance with density of states $\{n_E\}$ where $n_E := \frac{N_E}{N}$ and, by assumption, $N_0 > 0$, we obtain the following approximate expression for the spectral gap:

$$g(z) = \frac{z}{Z_2} \sqrt{(z - Z_1)^2 + 4 \frac{N_0}{N} Z_2} \quad (2.3)$$

where $Z_p := \sum_{E=1}^m \frac{n_E}{E^p}$ is a partition sum encoding information about the problem instance.

From (2.1) we see that the time complexity T is roughly dictated by inverse powers of the gap and is dominated when the gap reaches its minimum value. Eq. (2.3) implies that the minimal gap occurs at $z_* := Z_1 + \mathcal{O}(n_0)$ and is approximately given by

$$g_* := \min_z g(z) = 2Z_1 \sqrt{\frac{N_0}{NZ_2}} \quad (2.4)$$

to leading relative order in $\frac{1}{N}$. See Appendix 2.7.1 for details. From the expression of the minimal gap g_* we see that to know its exact location we need detailed knowledge about the problem instance in the form of a partition sum Z_1 (see Section 2.3.2). This characterizes the complexity of model (2.2) in term of the spectral gap and its minimum value. Eq. (2.4) shows that the gap closes to a smallest value g_* that scales exponentially in n , at a point z_* which depends on detailed information (the density of states $\{n_E\}$) about the problem instance. Then, the position of g_* in the annealing interval fluctuates instance by instance. This poses a problem at the moment of designing the schedule for achieving speedup as we will see in the next section.

To complete the definition of the algorithm we need a stopping condition. In adiabatic QA this is simply the final evolution time T where $z = z_f$. The annealing point z_f can be specified by requiring that at the end of the evolution the probability of finding the ground state solution is high enough, which can be encoded through a free parameter in the model. We find that a sufficient stopping condition is $z_f := z_* + \Delta$ with $\Delta := \sqrt{\frac{4N_0Z_2}{N\delta}}$ with $\delta \ll 1$ and




Figure 2.1 **Left:** Gap g of (2.3) versus z for 20 random 3-SAT instances with density of states drawn from a Gaussian distribution $n_E \sim \mathcal{N}(p_E, \Sigma)$ (see Appendix 2.7.4) for $n = 100$ spins. The clause/spin ratio $\alpha := m/n = 4.2$ close to the SAT-UNSAT phase transition point [92]. **Right:** Zoom-in on the minimal gap region. The lower black dots with minimal gaps $\approx 10^{-15}$ correspond to instances with $N_0 = 1$ whereas the others with minimal gaps $\approx 10^{-9}$ correspond to $N_0 > 1$.

arbitrary constant independent of n (see Appendix 2.7.1). This guarantees a high probability of finding the ground state solution at the final time. It is sufficient to stop just outside the range of minimal gap to achieve a high probability ratio of finding the solution.

2.3.2 Fluctuations in the minimal gap location

Eq. (2.3) explicitly gives the minimal gap and its location in terms of Z_1 and Z_2 , which depend on the density of states N_E and thus on the problem instance. Note that the range of z where the gap is small ($\mathcal{O}(1/\sqrt{N})$) is very narrow and of width $\mathcal{O}(1/\sqrt{N})$. Instance by instance fluctuations in the density of states δn_E induces fluctuations of $\mathcal{O}(n^{-5/2})$ in the minimal gap location $z_* = Z_1$ (see Fig. 2.5 in Appendix 2.7.4), while the depth of the gap stays of $\mathcal{O}(1/\sqrt{N})$. For random instances n_E drawn from arbitrary distributions, the fluctuations in the minimal gap location z_* is thus very large compared to the width of the minimal gap. We illustrate this in Fig. 2.1, where we computed the gap $g(z)$ dependency with z for 20 random 3-SAT instances with $n = 100$ spins. The gap depends on the instance only through the density of states n_E . For random 3-SAT, the mean density of states $\langle n_E \rangle$ is Binomial-distributed (see Eq. (2.33)

in Appendix 2.7.4). In the large n limit, n_E becomes Gaussian-distributed with covariance matrix given by Eq. (2.34) in Appendix 2.7.4.

In order to obtain a quantum speedup, the annealing schedule (how z varies with β) should be optimized such that annealing is fast when g is large and (very) slow when g is small. Due to the strong instance dependence of $g(z)$, the schedule is also instance dependent. This is possible, but requires the exact numerical value of Z_1 , since that determines the location of the minimal gap. When Z_1 is known, or can be efficiently computed, one can construct an annealing schedule that computes the solution in $\mathcal{O}(\sqrt{N})$ and thus achieves quadratic speed-up relative to a naive exhaustive search. An example is Grover search where $E(s)$ has only values 0, 1. The optimal schedule has $\dot{z} = \mathcal{O}(1/\sqrt{N})$ in an exponentially small range of $\mathcal{O}(1/\sqrt{N})$ around $z = Z_1$ and large $\dot{z}$ elsewhere. For general optimization problems, however, Z_1 is intractable to compute [93] and an efficient algorithm cannot be designed. The problem of computing the partition sum or, equivalently, the density of states, is a counting problem that is at least as hard as the corresponding NP problem [94]. An example is 3-SAT, which is NP-complete. Computing the partition sum requires computing the density of states $\{n_E\}$. In particular, it requires computing $n_0 = N_0/N$. Computing the total number of solutions N_0 can be a harder task than computing the solutions themselves, since an efficient method to count solutions can be iteratively applied order n times to find a particular solution in $\text{poly}(n)$ runtime.

A possible way to avoid this negative conclusion, is to note that the cost function $0 \leq E(s) \leq m$ can always be transformed to a new cost function $0 \leq E'(s) \leq 1$ where $E'(s) = 1$ iff $E(s) > 0$. Then the density of state is known: $N_{E'=0} = N_0, N_{E'=1} = N - N_0$ in terms of the number of optimal solutions N_0 . In terms of E' , the problem is equivalent to Grover's unstructured search, for which we already know that quadratic speed-up is possible [54]. But formulating the optimization problem in this way, yields a Hamiltonian H_f that is no longer the sum of local terms, which may prevent an efficient circuit implementation³ Essentially, we have reduced the problem to Grover-like

³The Hamiltonian H_0 can be implemented on a quantum circuit by using an ancilla qubit with $\mathcal{O}(\text{poly}(n))$ controlled-NOT and single-qubit gates [95].

search at the expense of the locality of the target Hamiltonian. As mentioned earlier, the problem of computing N_0 can be a challenging task for classical algorithms. In the quantum circuit model, the original Grover search algorithm does not encounter the fine-tuning problem inherent to analog quantum computing. However, it requires a precise knowledge of N_0 to determine when to stop the algorithm due to the soufflé problem [96]. Gate-based algorithms exist to address this challenge and achieve quadratic speed-up for unknown N_0 [97, 98]. In the continuous case, we can circumvent the need to compute the location of the minimal gap by recasting the problem in the form of the original adiabatic Grover algorithm introduced in [54]. This algorithm features a fixed location of the minimal gap within the annealing interval and is independent of n . Nevertheless, we still rely on knowing N_0 to design the schedule that achieves quadratic speed-up. Even more, this algorithm already faces some precision requirements that can become computationally demanding with increasing number of qubits [84]. Setting aside precision issues, it was proven in [99] that the adiabatic Grover algorithm, together with a tailored schedule [54], can achieve quadratic speed-up for estimates N'_0 such that $N'_0 < N_0$, which translates into runtimes which are longer than necessary.

2.3.3 Quadratic speedup for generic optimization problems

From the AT (2.1) we see that the time complexity of QA depends on the schedule $z(\beta)$. An example is adiabatic Grover search where a naive schedule results in $T \sim N$ while an optimized version results in $T \sim \sqrt{N}$, the celebrated quadratic speedup [54]. Here we aim to generalize this result to arbitrary optimization problems. In this section we design a schedule that achieves quadratic speedup for the general case, from which Grover search is a particular case. Later we prove that this speedup is optimal in the sense that this schedule minimizes T .

For the construction of the schedule we use a local-optimization ansatz

that was first introduced in [54] (see also [100]) and further analyzed in [91]. Assume a schedule $z(\beta)$ that satisfies

$$\dot{z} = cg^2 \quad (2.5)$$

where $g = g(z)$ is the gap of model H in (2.2) and c is a normalization constant. Note that Eq. (2.5) forces the schedule to slow down where the gap becomes small, which obviously requires knowing the position of its minimum. Eq. (2.5) represents a boundary-value problem subject to conditions $z(0) = 0$ and $z(1) = z_f$. Hence, c is completely determined:

$$c = \int_0^{z_f} \frac{dz}{g(z)^2}. \quad (2.6)$$

We use the AT stated in Eq. (2.1) to compute the time complexity of model H . The complexity C in (2.1) can be bounded using the ansatz (2.5). It can be shown that [84]

$$C \leq c \left(2E_{\max} + 28 \int_0^{z_f} \frac{dz}{g(z)} \right) \quad (2.7)$$

with c given by (2.6) and $E_{\max} := \|H_f\| = \mathcal{O}(n)$. Note that unlike (2.1), equation (2.7) depends on the schedule only through the gap, meaning that we have gotten rid of schedule parametrization details and focus instead on the form of the gap along the adiabatic line $H(z)$. We use this bound and the analytical approximation for the gap (2.3) to show the following.

Lemma 1. *Given the model $H = H_0 + zH_f$, there exist a schedule $z(\beta)$, solution to the boundary-value problem (2.5), that achieves a time complexity*

$$T = \mathcal{O} \left(\sqrt{N/N_0} \right).$$

Proof. See Appendix 2.7.2. □

Lemma 1 states that for any optimization problem that can be encoded in H there exists a schedule achieving quadratic speedup. This was previously shown for Grover search [54], for which an analytical solution to (2.5) is

available. Here we generalize this to the case of arbitrary target Hamiltonians that are diagonal in the computational basis.

Now we can ask whether this is optimal, or whether there exists a schedule that yields better than quadratic speedup, by optimizing the schedule using a different strategy than (2.5). Previously, it was shown that quadratic speedup is optimal for the Hamiltonian $H = (1 - A)H_0 + AH_f$ using a schedule $A(\beta)$ that is linear in β [74]. We generalize this result and prove that quadratic speedup is optimal for arbitrary schedules.

Lemma 2. *Let T the total evolution time of the adiabatic model $H = H_0 + zH_f$ and P the projector onto the ground state subspace of H_f with degeneracy N_0 . Define $p := \langle \psi(T) | P | \psi(T) \rangle > 0$. Then*

$$\sqrt{\frac{N}{N_0}} \gamma \leq T \quad (2.8)$$

with $\gamma := \frac{p}{2} \frac{\left(1 - \sqrt{\frac{N_0}{Np}}\right)^2}{1 + \sqrt{p}}$ a constant of $\mathcal{O}(1)$.

Proof. See Appendix 2.7.3. □

Note that Lemma 2 states that the optimal time complexity scales as $T \sim \frac{1}{g_*}$, see (2.4), while a naive estimate based on the adiabatic theorem predicts a scaling $T \sim \frac{1}{g_*^2}$ [100]. Applied to our case, this would yield $T \sim N/N_0$, i.e. no speedup at all. Also, note that this lower bound on the annealing time applies regardless we have detailed knowledge of the problem instance or if the quantum annealing is performed non adiabatically.

The time complexity in (2.8) depends on the number of solutions N_0 . This number is constant and independent of n for Grover search [54] for which N_0 is fixed. In this case the complexity is $\mathcal{O}(\sqrt{N})$. The same complexity is obtained for the case of *unique satisfying assignments* in random 3-SAT [101] where one considers instances with only one solution. For random 3-SAT, the number of solutions N_0 is a random variable that has mean value $\langle N_0 \rangle = Np_0$, with $p_0 = \left(\frac{7}{8}\right)^{\alpha n}$ and $\alpha = m/n$ a constant (see Appendix 2.7.4). Then $\langle N_0 \rangle \sim N^\gamma$ with $\gamma = 1 - \alpha \left| \log \frac{7}{8} \right|$. Knowing that g_* goes like $\sqrt{N_0/N}$ we can estimate the mean value of the minimal gap as $\langle g_* \rangle \sim N^{-\frac{\alpha}{2k}}$ where

$k := 1/|\log \frac{7}{8}| \approx 5.19$. This value of k was the first reported as an upper bound on the phase transition point α_c for random 3-SAT [102, 103]. Today, this value can be accurately estimated and is known to be around $\alpha_c = 4.26$ [92, 104, 105]. For $\alpha < k$ the mean number of solutions $\langle N_0 \rangle$ grows exponentially fast with n , while for $\alpha > k$ it decreases with the same rate. Using the optimal schedule, the relation $T \sim 1/g_*$ implies $T \sim N^{\frac{\alpha}{2k}}$. Thus, the speedup for the average instance with $N_0 = \langle N_0 \rangle$ depends monotonically on α , reaching a $\sqrt{N}$ complexity when $\alpha = k$, just where $\langle N_0 \rangle = 1$, i.e. the case of unique satisfying assignments. Although it is interesting to note that the speedup can be better than $\sqrt{N}$ for $\alpha < k$, it still implies an exponentially large computation time. For smaller α there are also efficient classical algorithms to solve random 3-SAT instances [92].

2.4 Comparison with the common mixing Hamiltonian

One could argue that our negative result only holds for the rank one projector Hamiltonian H_0 that we used, and much more optimistic results could be obtained using the common transverse field Hamiltonian $H_0 = -\sum_{i=1}^n \sigma_i^x$. This is a valid concern. In this section we numerically investigate the spectral gaps up to $n = 20$ qubits by using the transverse field Hamiltonian as our mixing Hamiltonian and producing the diagonal Hamiltonian (H_f) energies with two different models. The first one is the 3-SAT model, and the second one belongs to a class of Ising spin glass models described by 3-spin interactions. Numerically we see that in the worst case, for 3-SAT instances the minimal gap approximately scales like $\frac{1}{\sqrt{N}}$, similar to the case with rank one projector Hamiltonian, however for the spin glass instances, the minimal gap vanishes faster than $\frac{1}{\sqrt{N}}$, indicating a longer adiabatic annealing time for the common transverse field Hamiltonian compared to the case in which the rank one projector Hamiltonian is the mixing Hamiltonian. Early studies on small instances indicated that the median spectral gap in this case scales as $1/\text{poly}(n)$, suggesting exponential speedup [17, 18, 55–57]. However, later it




Figure 2.2 Scaling of the g_* for $H = (1 - A)H_0 + AH_f$, with $H_0 = -\sum_{i=1}^n \sigma_i^x$. **Left** corresponds to 3-SAT instances with unique solutions ($N_0 = 1$). **Right**, to 3-spin model instances with unique ground state. For each n , we computed the minimal gap $\min_A g(A)$ for a batch of different instances. We show the median and minimum gap over instances, denoted by g_{mdn} and g_{min} respectively, and the dashed line corresponds to a $1/\text{poly}(n)$ fit for g_{mdn} . We generated 200 instances for $n \leq 13$ and 50 instances for $n > 13$. The figure indicates that up to 20 qubits, minimal gap of the 3-SAT instances in the worst case scales approximately as $1/\sqrt{N}$ and for 3-spin model instances it indicates a close to $1/N$ scaling as the number of qubits is increased.

was shown that for the worse case instances the gap closes exponentially with n and that these instances dominate for large n [58]. In agreement with these findings, we find that the majority of small instances have large gaps, but the worse case instances have exponentially small gap. Since the worse case instances dominate at large n , we expect that the spectral gap does not scale better than $1/\sqrt{N}$ for large instances. Therefore, in order to reach quadratic speedup with this Hamiltonian, one also needs an optimized annealing schedule that requires precise knowledge of the location of the minimal gap. There exists no known analytical expression for this location, but there is no reason to assume that it would be less intractable than in the case of Eq. (2.2). From this we conclude that adiabatic quantum annealing also cannot yield better than quadratic speed up with this mixing Hamiltonian.

Delving into details, we numerically studied the minimal gap of the Hamil-

tonian $H(A) = (1 - A)H_0 + AH_f$ where the schedule $A \in [0, 1]$. This setup is equivalent to model (2.2) as we stated in Sec. 2.3. We use as the mixing Hamiltonian the transverse field $H_0 = -\sum_{i=1}^n \sigma_i^x$. We consider target Hamiltonians H_f associated to two different problems. The first one corresponds to 3-SAT with $\alpha = m/n = 4.26$. The second, to an Ising spin glass Hamiltonian described by 3-spin interactions with couplings J uniformly sampled from $\{\pm 1\}$:

$$H_f = \sum_{i < j < k} J_{ijk} \sigma_i^z \sigma_j^z \sigma_k^z, \quad J_{ijk} = \pm 1. \quad (2.9)$$

We refer to this model shortly as the 3-spin model.

For 3-SAT we consider instances with unique solutions ($N_0 = 1$). For the 3-spin model we only consider instances with unique ground state. At each n and for each instance, we compute $g_* = \min_A g(A)$.

In Fig. 2.2 we plot the median g_{mdn} and the minimum g_{min} of the minimal gap over instances as a function of n for each model. We see that in both cases and up to $n = 20$ qubits, g_{mdn} scales as $1/\text{poly}(n)$, indicating polynomial median complexity, whereas g_{min} decreases exponentially, indicating exponential complexity for the worst cases. In the case of 3-SAT the gap decreases like $1/\sqrt{N}$ for worst cases. This indicates that the complexity of adiabatic QA for 3-SAT with the transverse field H_0 is similar to the projector H_0 . This supports our argument that the exponential shrinkage of the gap for hard instances will not disappear by using a different instance-independent mixing Hamiltonian. However, with the transverse field H_0 , the prospects of constructing a schedule that achieves the quadratic speedup are dim, because we cannot compute z_* . In simulations we observe that $g(z)$ may even have multiple local minima as a function of z .

For the 3-spin model, our simulations show a scaling of the minimal gap close to $1/N$ which is worse than projector H_0 , where the gap does not vanish faster than $1/\sqrt{N}$. This results disagrees with previous studies [106], who reported exponentially small gaps of order $1/\sqrt{N}$ close to the phase transition point. In any case, the point is that the results with a transverse field H_0 are not better than with the projector H_0 . Therefore, it seems

implausible that the use of the transverse field H_0 rather than a projector H_0 will improve the worst case performance of an adiabatic QA algorithm.

2.5 Conclusions and discussion

In this paper we have addressed the question of whether quantum speed-up is possible with adiabatic quantum annealing. For this purpose, we have studied a class of Hamiltonians Eq. (2.2) which allows us to compute the gap $g(z)$ around its minimal value $g_* = g(z_*)$. The minimum location z_* is given in terms of a partition sum that is in general intractable to compute. Given z_* we can construct an annealing schedule for which the time complexity scales as $\mathcal{O}(\sqrt{N})$. This was first shown for adiabatic Grover search in [54], which is a particular instance of our model. Consequently our construction generalizes this result to arbitrary target Hamiltonians that are diagonal in the computational basis. Furthermore, we prove that for this Hamiltonian quadratic speed-up is optimal. This extends earlier results from [74] in which it was shown the same optimality for linear schedules, to the case of arbitrary non-linear schedules. Thus, adiabatic quantum annealing with Hamiltonian given by (2.2) gives at best quadratic speed-up, but only when z_* is tractable to compute, which is not the general case.

The fact that quantum annealing does not yield exponential speedup should come as no surprise. The same holds in fact for classical Simulated Annealing. For instance, [48] shows that in order to guarantee an optimal solution, the temperature should be decreased very slowly and the time of the algorithm scales exponential with n . This does not say anything about the potential practical value of Simulated Annealing. In fact, it is among the most powerful methods to find approximate solutions for combinatorial optimization problems. The same could be true for quantum annealing.

A further issue that complicates adiabatic quantum annealing is the fact that, even when z_* is known or can be computed, it needs to be specified to an exponential precision, which prevents practical implementation on a circuit [88]. The precision problem was also pointed out in [107] which

discusses a classical algorithm for unstructured search that realizes quadratic speedup at the expense of an increasing precision in the initial conditions with system size.

Recently, there has been a proposal of quantum adiabatic spectroscopy to empirically find the location z_* where the spectral gap is minimal [108]. It may be interesting to investigate whether this method can be used to estimate the partition functions that we used in this paper.

It is also worth mentioning that the quadratic speed-up obtained in this work corresponds to a model that does not encode the structure of the problem instance in the driver Hamiltonian H_0 . Driver Hamiltonians harnessing such structure might improve the speed-up over the more unstructured/unbiased $H_0 = -|\phi\rangle\langle\phi|$. How much structure one needs to encode in the model to gain considerable speed-ups is a question that remains open in the quantum computing community [109].

Although we have shown that the prospect of quantum speed-up using adiabatic QA is unlikely, this does not preclude the pursuit of non-adiabatic QA or hybrid approaches to achieve speedup. Analytical methods to explore the complexity of general non-adiabatic algorithms are not fully available and more research needs to be done to fill this breach.

In that line, it is an interesting question to consider what is the response of the model considered here if we use a hybrid protocol. Would it be possible to find a general time evolution that leads us towards finding the ground state solution with better scaling or at least states that represent quasi-solutions? We leave the answer to these questions for future work.

2.6 Acknowledgments

We thank the anonymous referees for their fruitful comments and suggestions on a previous draft of this work. PN acknowledge support from the 'Quantum Inspire – the Dutch Quantum Computer in the Cloud' project (NWA.1292.19.194) of the NWA research program 'Research on Routes by Consortia (ORC)', which is funded by the Netherlands Organization for Scien-

tific Research (NWO). All authors made equal contribution to this work.

2.7 Appendices

2.7.1 Model analysis

We complement the paper with analytical results related to model (2.2)

$$H = H_0 + zH_f$$

where z is the annealing parameter, H_f is diagonal in the computational basis and $H_0 = -|\phi\rangle\langle\phi|$ is the rank-one Hamiltonian with $|\phi\rangle$ the equal superposition state. We perform the exact computation of the spectral and gap properties of H .

Energy spectrum and eigenstates

Denote $\lambda_0 \leq \dots \leq \lambda_{N-1}$ the eigenvalues of H . Sort the $m' \leq m$ distinct energies for which $N_E > 0$ as $E_1 < E_2 < \dots < E_{m'}$ ⁴. Remember that N_E is the number of states with energy E and E_i can take values in the set $\{0, 1, \dots, m\}$.

Lemma 3. *H has a unique ground state eigenvalue $\lambda_0 < zE_1$. In addition, H has one non-degenerate eigenvalue in each interval $zE_i < \lambda < zE_{i+1}$, $i = 1, \dots, m'-1$ and $N_{E_i} - 1$ eigenvalues $\lambda = zE_i$, $i = 1, \dots, m'$.*

Proof. Given an arbitrary real vector $v \in \mathbb{R}^N$ and an invertible matrix $M_0 \in \mathbb{R}^{N \times N}$, the determinant of $M = M_0 + vv^\top$ satisfies $\det(M) =$

⁴Note that these energies are different from the eigenvalues of H .

$(1 + v^\top M_0^{-1} v) \det(M_0)$ [110]. Thus,

$$\det(\lambda I - H) = X(\lambda) \prod_{E=0}^m (\lambda - zE)^{N_E} \quad (2.10)$$

$$= p(\lambda) \prod_{i=1}^{m'} (\lambda - zE_i)^{N_{E_i}-1} \quad (2.11)$$

with

$$X(\lambda) = 1 + \frac{1}{N} \sum_{E=0}^m \frac{N_E}{\lambda - zE} = 1 + \frac{1}{N} \sum_{i=1}^{m'} \frac{N_{E_i}}{\lambda - zE_i} = \frac{p(\lambda)}{\prod_{i=1}^{m'} (\lambda - zE_i)} \quad (2.12)$$

$p(\lambda)$ is a polynomial of order m' and has thus m' solutions that are given by the zero crossings of $X(\lambda)$. $X(\lambda)$ is singular for $\lambda = zE_i$, $i = 1, \dots, m'$ and $\frac{dX}{d\lambda} < 0$ for all $\lambda \neq zE_i$. Thus $X(\lambda)$ has a unique zero crossing in the intervals $\lambda < zE_1$ and $zE_i < \lambda < zE_{i+1}$, $i = 1, \dots, m' - 1$. All these eigenvalues are non-degenerate. The remaining $N - m'$ eigenvalues are solutions of $\prod_{i=1}^{m'} (\lambda - zE_i)^{N_{E_i}-1} = 0$. For each E_i , there are $N_{E_i} - 1$ eigenvalues $\lambda = zE_i$. $\square$

We can analytically compute the eigenvectors of H for those eigenvalues that satisfy $X(\lambda) = 0$ (i.e. those that satisfy $zE_i < \lambda < zE_{i+1}$, not the eigenvalues $\lambda = zE_i$). In particular,

Lemma 4. *Let λ be the eigenvalues of H that satisfy $X(\lambda) = 0$. The corresponding eigenvectors $|v\rangle = \sum_s v(s) |s\rangle$ have components*

$$v(s) \propto \frac{1}{zE(s) - \lambda}, \quad (2.13)$$

where $E(s)$ defines the target Hamiltonian $H_f = \sum_s E(s) |s\rangle \langle s|$.

Proof. The eigen-equation is

$$(H_0 + zH_f) |v\rangle = \lambda |v\rangle, \quad (2.14)$$

which implies

$$|v\rangle = -(zH_f - \lambda)^{-1}H_0|v\rangle = (zH_f - \lambda)^{-1}\langle\phi|v\rangle|\phi\rangle. \quad (2.15)$$

Note that the inverse in (2.15) is well defined. Since $|\phi\rangle = \frac{1}{\sqrt{N}}\sum_s|s\rangle$ and H_f is diagonal in the computational basis,

$$|v\rangle \propto \sum_s (zE(s) - \lambda)^{-1} |s\rangle.$$

□

We wish to guarantee that the adiabatic algorithm finds a solution s with $E(s) = 0$ when such a solution exists. We therefore restrict ourselves to problem instances for which $N_0 > 0$.

By definition, the spectral gap is the energy difference between the first excited state and the ground state. In our model, this might pose an apparent dilemma. For it, let's take a look at the pattern

$$\lambda_0 < 0 < \lambda_1,$$

where λ_0, λ_1 now represent the first and second roots of $X(\lambda) = 0$. The value $\lambda = 0$ is an eigenvalue of H only when $N_0 > 1$. In that case it carries degeneracy $N_0 - 1$. When $N_0 = 1$, $\lambda = 0$ is not an eigenvalue and we can safely define the gap as $\lambda_1 - \lambda_0$. However, when $N_0 > 1$, it is raised the question of whether we should take the gap as $0 - \lambda_0$ or $\lambda_1 - \lambda_0$. It can be proven that the 0 eigenvalue corresponds to states that belong to the solution space. This implies that the true energy gap, the one that characterizes the separation between the solution space and first excited non-optimal states, is given by $\lambda_1 - \lambda_0$. Next, we give a proof of this making use of a permutation symmetry present in H .

Lemma 5. *The spectral gap of H is given by $g := \lambda_1 - \lambda_0$, where λ_0, λ_1 are the first and second roots of $X(\lambda) = 0$.*

Proof. The proof starts by noticing that the whole quantum dynamics can be

restricted to a space of dimension $m' < N$. To see this, define the normalized states

$$|a_i\rangle := \frac{1}{\sqrt{N_{E_i}}} \sum_{s|E(s)=E_i} |s\rangle, \quad i = 1, \dots, m'. \quad (2.16)$$

Denote the span of these states as $\mathcal{H}_a$. Consider an eigenvalue λ of H that is a root of $X(\lambda) = 0$ with eigenvector $|v\rangle$. For each energy E for which $N_E > 1$ define permutation operators R_E that interchange two arbitrary states s corresponding to the same $E = E(s)$ and leaves the rest unchanged. For energies E for which $N_E = 1$ define R_E as the identity operator. By construction, we have $[R_E, H] = 0$ for all E . Since λ is a non-degenerate eigenvalue, its eigenvector $|v\rangle$ is invariant under the action of R_E , i.e. $R_E |v\rangle = |v\rangle$ for all E . So we conclude that $|v\rangle$ should be a linear combination of $|a_i\rangle$ states, therefore $|v\rangle \in \mathcal{H}_a$. Since we have m' eigenvectors with eigenvalues λ that satisfy $X(\lambda) = 0$, these eigenvectors also span $\mathcal{H}_a$. Then, the Hamiltonian H can be written as

$$H = H_a + H_a^\perp,$$

where H_a is an operator that acts only on $\mathcal{H}_a$ and $H_a^\perp$ acts on the orthogonal complement space $\mathcal{H}_a^\perp$.

Finally, note that the initial state $|\phi\rangle$ can be written as

$$|\phi\rangle = \sum_i \sqrt{n_{E_i}} |a_i\rangle,$$

where $n_{E_i} := N_{E_i}/N$ for $i = 1, \dots, m'$, so it also belongs to $\mathcal{H}_a$. Since $[H_a, H_a^\perp] = 0$ the entire quantum evolution becomes restricted to $\mathcal{H}_a$. The initial state has no component in $H_a^\perp$ and the dynamics will not develop a component in $H_a^\perp$ at any later time. Thus, effectively, the Hamiltonian is given by

$$H = H_a = \sum_{i=0}^{m'-1} \lambda_i |v_i\rangle \langle v_i|.$$

From this we conclude that the gap is given by $\lambda_1 - \lambda_0$.

□

Gap

Consider a given instance with density of states $\{n_E\}$ with $n_E = \frac{N_E}{N}$ and $N_0 > 0$ and define

$$Z_p := \sum_{E=1}^m \frac{n_E}{E^p}. \quad (2.17)$$

We show the following result.

Lemma 6. *For a given instance, the lowest two eigenvalues of H , $\lambda_0 < 0 < \lambda_1$, are approximately given by*

$$\lambda'_{1,0} = z - Z_1 \pm \frac{\sqrt{(z - Z_1)^2 + 4 \frac{N_0}{N} Z_2}}{2Z_2}. \quad (2.18)$$

where prime denote approximate quantities. The approximation for λ_0 has relative error $|\delta\lambda_0/\lambda_0| := \left| \frac{\lambda_0 - \lambda'_0}{\lambda_0} \right| \leq \delta$, for $z \geq z_0 := Z_1 - \frac{Z_2^2}{Z_3} \delta$. The approximation for λ_1 has relative error $|\delta\lambda_1/\lambda_1| := \left| \frac{\lambda_1 - \lambda'_1}{\lambda_1} \right| \leq \delta$ for $z \leq z_1 := Z_1 + \frac{Z_2^2}{Z_3} \delta$. The spectral gap is approximately given by

$$g' = \lambda'_1 - \lambda'_0 = \frac{z}{Z_2} \sqrt{(z - Z_1)^2 + 4 \frac{N_0}{N} Z_2}. \quad (2.19)$$

The approximation for g has relative error $|\delta g/g| := \left| \frac{g - g'}{g} \right| \leq \delta$ for $z_0 \leq z \leq z_1$. The minimal spectral gap g_* occurs for $z = z_* := Z_1 + \mathcal{O}(n_0)$ where $n_0 = N_0/N$ and is given by

$$g_* = \min_z g(z) = 2Z_1 \sqrt{\frac{N_0}{N Z_2}} + \mathcal{O}(n_0^{3/2}) \quad (2.20)$$

.

Proof. Since $g = \lambda_1 - \lambda_0$, we must solve for the two lowest eigenvalues λ_0, λ_1 of the characteristic equation. From Lemma 10, $\lambda_0 < 0 < \lambda_1$. When the gap $g = \lambda_1 - \lambda_0$ is small, both λ_0, λ_1 are close to 0. For $z > 0$ and to leading order

in λ ,

$$X(\lambda) = 1 + \frac{N_0}{N\lambda} + \frac{1}{N} \sum_{E \geq 1} \frac{N_E}{\lambda - zE} = \frac{1}{Nz\lambda} \left(N_0z + b\lambda - a\lambda^2 - \frac{NZ_3}{z^2}\lambda^3 + \mathcal{O}(\lambda^4) \right). \quad (2.21)$$

with $b := N(z - Z_1)$ and $a := N\frac{Z_2}{z}$. The solution to the quadratic approximation to $X(\lambda) = 0$ is

$$\lambda'_{1,0} = \frac{b \pm \sqrt{D}}{2a}, \quad D = b^2 + 4N_0NZ_2. \quad (2.22)$$

We easily verify that this solution satisfies $\lambda_0 < 0 < \lambda_1$. The gap is $g' = \sqrt{D}/a$ which yields Eq. (2.19). It is minimized when $z = z_* = Z_1 + \mathcal{O}(n_0)$ which gives Eq. (2.20).

The error in the gap estimate is due to the errors in λ'_0 and λ'_1 . For $z_* - \Delta z \leq z \leq z_* + \Delta z$ both errors are small for some sufficiently small $\Delta z > 0$ and the gap is accurately estimated by Eq. (2.19). Denote $\delta\lambda_i$ the error in λ'_i for $i = 0, 1$. To estimate $\delta\lambda_i$ we solve the third order correction to the quadratic approximation to $X(\lambda) = 0$ as follows. Consider

$$X'(\lambda) - N\frac{Z_3}{z^2}\lambda^3 = 0$$

where $X'(\lambda) = -a\lambda^2 + b\lambda + c$ is the quadratic approximation in (2.21) and $\lambda = \lambda' + \delta\lambda$ with $X'(\lambda') = 0$. By solving to first order in $\delta\lambda$, we have

$$\delta\lambda = \frac{N\frac{Z_3}{z^2}}{-2a\lambda' - 3N\frac{Z_3}{z^2}\lambda'^2 + b} \lambda'^3$$

Using (2.22) and ignoring the $\mathcal{O}(\lambda'^2)$ term in the denominator we have

$$\delta\lambda = \frac{N\frac{Z_3}{z^2}}{-2a\lambda' + b} \lambda'^3 = \mp \frac{NZ_3}{z^2\sqrt{D}} \lambda'^3.$$

Therefore,

$$\left| \frac{\delta\lambda}{\lambda'} \right| = \frac{NZ_3}{z^2\sqrt{D}}\lambda'^2. \quad (2.23)$$

We wish to find a validity interval $z_0 \leq z \leq z_1$ in which the relative error remains bounded as

$$\left| \frac{\delta\lambda}{\lambda'} \right| \leq \delta, \quad (2.24)$$

for $\lambda = \lambda_0, \lambda_1$ and some $\delta \ll 1$. It can be easily shown that the relative error $|\delta\lambda/\lambda'|$ increases with z for λ'_1 and decreases for λ'_0 . Therefore the maximum relative error is achieved at the boundaries z_0 and z_1 . For λ_0 the maximum is achieved at $z = z_0$ and for λ_1 , at $z = z_1$. Then, it is enough to find a suitable $\varepsilon > 0$ such that $z_0 = Z_1 - \varepsilon$ and $z_1 = Z_1 + \varepsilon$ in order to satisfy condition (2.24). Combining (2.23) with condition (2.24) and assuming that $\varepsilon \gg n_0$ we obtain

$$\varepsilon \leq \frac{Z_2^2}{Z_3} \delta.$$

Therefore, by defining $\varepsilon := \frac{Z_2^2}{Z_3} \delta$ we ensure that (2.24) is satisfied in the interval $Z_1 - \varepsilon \leq z \leq Z_1 + \varepsilon$. Finally, note that the bound in Eq. (2.24) directly bounds the relative error in the gap $|\delta g/g'| \leq \delta^5$ where $\delta g = g - g'$.

It remains to determine the dependency of δ with respect to n . For consistency, δ has to be sufficiently small to satisfy $z_0 > 0$. This requires that $\frac{Z_1 Z_3}{Z_2^2} > \delta$, meaning that if the quotient $\frac{Z_1 Z_3}{Z_2^2}$ decreases with n , δ must decrease with n as well. Using the definition of Z_p and the fact that $1 \leq E \leq m = \mathcal{O}(n)$, it can be shown that $\delta = \mathcal{O}(1/n^2)$ is a sufficient condition to ensure $\frac{Z_1 Z_3}{Z_2^2} > \delta$. Nevertheless, we give a typicality argument to support the claim that $\delta \ll 1$ independent of n is enough to maintain consistency. Assume that Z_p can be approximated by $1/\langle E \rangle^p$ for large n . This is true, in particular, for 3-SAT, where $\langle E \rangle = m/8$ with $m = \alpha n$ and α a constant. Then, for large n , the quantity $\frac{Z_2^2}{Z_3 Z_1} \rightarrow 1$. Numerical results for 3-SAT show that the quotient $\frac{Z_1 Z_3}{Z_2^2}$

⁵We prove this by noticing that $\frac{\delta g}{g'} = \frac{\delta \lambda'_1 - \delta \lambda'_0}{g'}$. Then, $\left| \frac{\delta g}{g'} \right| \leq \delta \frac{|\lambda'_1| + |\lambda'_0|}{g'} \leq \delta$, where in the last step we used the formulas for λ'_0, λ'_1 and g' .

asymptotically approaches one from above for large n . Therefore, $\frac{Z_1 Z_3}{Z_2^2}$ is order 1 in the typical case, and taking δ sufficiently small and independent of n is enough to ensure the consistency of the approximation. $\square$

In Lemma 6 we specified the annealing interval $z \in [z_0, z_1]$ for which both λ_0 and λ_1 (and hence the gap) are accurately estimated. We compare the approximation Eq. (2.19) with the exact gap for a random 3-SAT instance with $n = 20$ spins in Fig. 2.3. It confirms that the errors in $\lambda_{0,1}$ and g are small in their range of validity.




Figure 2.3 **Left top:** Comparison between approximate (solid line) and true (dashed line) λ_0 (red) and λ_1 (blue) eigenvalues (see (2.18)) for one random 3-SAT instance with $n = 20$ and $\alpha = 4.2$. Dashed magenta line is detection bound $z_f = z_* + \sqrt{\frac{4N_0 Z_2}{N\delta}}$. **Left bottom:** Approximate vs true gap $g = \lambda_1 - \lambda_0$ (see (2.19)) in the range $z \in [0, 1]$ and minimum location (z_*, g_*) (black dot). **Top right:** Relative error in the eigenvalues $|\delta\lambda_0/\lambda_0|$ (red) and $|\delta\lambda_1/\lambda_1|$ (blue). **Bottom right:** Relative error in the gap $|\delta g/g|$. Vertical black dashed lines indicate the validity intervals $z \geq z_0$ for λ_0 and $z \leq z_1$ for λ_1 with $z_{1,0} = Z_1 \pm \delta \frac{Z_2^2}{Z_3}$ where the relative errors $|\delta\lambda_{0,1}/\lambda_{0,1}| \leq \delta = 0.01$, respectively.

The next lemma assesses the behavior of the gap in the remaining of the annealing interval and it shows that it is not exponential.

Lemma 7. *For $0 \leq z \leq z_0$ the spectral gap $g > \mathcal{O}(1/\text{poly}(n))$.*

Proof. Define $\lambda = z\lambda'$. Then the characteristic equation $X(\lambda) = 0$ becomes $F(\lambda') = \sum_{E=1}^m \frac{n_E}{\lambda' - E} = -z$. F is a decreasing function of λ' . Increasing z , decreases the intercept $F(\lambda') = -z$ and thus increases λ' . Thus all solutions λ of $X(\lambda) = 0$ are increasing with z . The gap is $g(z) = \lambda_1(z) - \lambda_0(z) > -\lambda_0(z)$. Since λ_0 is increasing function of z we have for $0 \leq z \leq z_0$ that $\lambda_0(z) < \lambda_0(z_0) = -\frac{z_0 Z_2}{Z_3} \Delta z$. Thus $g(z) > \frac{z_0 Z_2}{Z_3} \Delta z = \mathcal{O}(1/\text{poly}(n))$. $\square$

This completes the characterization of the gap behavior along the annealing interval until the location of the minimal gap. It shows the existence of a unique phase transition point at z_* where the minimal gap $g_* \sim \sqrt{N_0/N}$ whereas in the rest of the interval the gap is non-exponential in n .

When to stop the adiabatic evolution?

We compute the annealing point z_f for which it is sufficient to achieve a high probability ratio of finding a solution at the end of the evolution.

Lemma 8. *Define $z_f := z_* + \Delta$ with $\Delta := \sqrt{\frac{4N_0 Z_2}{N\delta}}$ and $\delta \ll 1$ a constant independent of n . The probability ratio in the ground state to detect the optimal solution $E = 0$*

$$\frac{p(E = 0)}{p(E > 0)} \geq \frac{4Z_2}{\delta}. \quad (2.25)$$

Proof. It is reasonable to assume that $z_f > z_*$, that is, stopping the evolution after the phase transition point z_* . Then, we can use Eq. (2.18), which is accurate up to $\mathcal{O}(N_0/N)$ corrections, to derive $|\lambda_0| \leq \frac{zN_0}{N\Delta}$ for $z \geq z_f$, where $\Delta = \sqrt{\frac{4N_0 Z_2}{N\delta}}$.

Using Lemma 4 the probability amplitudes are

$$p(E = 0) \propto \frac{N_0}{\lambda_0^2} \geq \frac{N^2 \Delta^2}{N_0 z^2} \quad (2.26)$$

$$p(E > 0) \propto \sum_{E \geq 1}^m \frac{N_E}{(zE - \lambda_0)^2} \leq \frac{N}{z^2} \quad (2.27)$$

$$\frac{p(E = 0)}{p(E > 0)} \geq \frac{N \Delta^2}{N_0} = \frac{4Z_2}{\delta}. \quad (2.28)$$

□

The final point z_f is indicated in Fig. 2.3 as the magenta dashed line. Note that for large n we have $z_f < z_1$ so that the annealing terminates within the range of z values where the approximation of the gap is accurate.

2.7.2 Quadratic speed-up: proof of existence

Proof. Here we provide a proof of Lemma 1. The integrals

$$\int_0^{z_f} \frac{dz}{g(z)^p}, \quad p = 1, 2$$

can be bounded as follows. From Lemma 6, in the interval $[z_0, z_1]$ the gap is accurately approximated by (2.3). Set $z_f = z_* + \Delta < z_1$ with $\Delta = \sqrt{\frac{d}{\delta}}$, $d := 4n_0 Z_2$ and $\delta \ll 1$ independent of n . Therefore in $[z_0, z_f]$ and up to terms of the same order as the integral,

$$\int_{z_0}^{z_f} \frac{dz}{g(z)^p} = \int_{z_0}^{z_f} \frac{Z_1^p}{z^p} \frac{dz}{\left(\sqrt{(z - Z_1)^2 + 4n_0 Z_2}\right)^p} \leq \frac{Z_1^p}{z_0^p} \int_{z_0}^{z_f} \frac{dz}{\left(\sqrt{(z - Z_1)^2 + 4n_0 Z_2}\right)^p}.$$

Hence, for $p = 1$ ⁶

$$\begin{aligned} \int_{z_0}^{z_f} \frac{dz}{\sqrt{(z - Z_1)^2 + 4n_0 Z_2}} &= \frac{1}{2} \ln \left(\frac{\sqrt{x^2 + d} + x}{\sqrt{x^2 + d} - x} \right) \Big|_{z_0 - Z_1}^{\Delta} \\ &= \frac{1}{2} \ln \left(\frac{\sqrt{1 + \delta} + 1}{\sqrt{1 + \delta} - 1} \right) - \frac{1}{2} \ln \left(\frac{\sqrt{(z_0 - Z_1)^2 + d} + z_0 - Z_1}{\sqrt{(z_0 - Z_1)^2 + d} - (z_0 - Z_1)} \right). \end{aligned}$$

Note that $|z_0 - Z_1| = \delta \frac{Z_2^2}{Z_3} \gg \sqrt{d} = 2\sqrt{Z_2} \sqrt{\frac{N_0}{N}}$ for large n when the number of optimal solutions N_0 is $\mathcal{O}(N^\gamma)$ with $\gamma < 1$. Therefore, to leading order in d

$$\ln \left(\frac{\sqrt{(z_0 - Z_1)^2 + d} + z_0 - Z_1}{\sqrt{(z_0 - Z_1)^2 + d} - (z_0 - Z_1)} \right) = -\ln \left(\frac{d}{4(z_0 - Z_1)^2} \right) = \mathcal{O}(n).$$

Thus

$$\int_{z_0}^{z_f} \frac{dz}{\sqrt{(z - Z_1)^2 + 4n_0 Z_2}} = \mathcal{O}(n).$$

For $p = 2$

$$\begin{aligned} \int_{z_0}^{z_f} \frac{dz}{(z - Z_1)^2 + 4n_0 Z_2} &= \frac{1}{\sqrt{d}} \tan^{-1} \left(\frac{x}{\sqrt{d}} \right) \Big|_{z_0 - Z_1}^{\Delta} \\ &= \frac{1}{\sqrt{d}} \left(\tan^{-1}(1/\sqrt{\delta}) - \tan^{-1} \left(\frac{z_0 - Z_1}{\sqrt{d}} \right) \right) \\ &\approx \frac{1}{\sqrt{d}} \pi = \mathcal{O}(\sqrt{N/N_0}). \end{aligned}$$

⁶For $p = 1$ and $p = 2$ we have the following primitives

$$\int \frac{dx}{\sqrt{x^2 + d}} = \frac{1}{2} \ln \frac{\sqrt{x^2 + d} + x}{\sqrt{x^2 + d} - x} + \text{ct.},$$

and

$$\int \frac{dx}{x^2 + d} = \frac{1}{\sqrt{d}} \tan^{-1} \left(\frac{x}{\sqrt{d}} \right) + \text{ct.},$$

respectively.

Putting all together in (2.7),

$$\begin{aligned} C &\leq \left(\int_0^{z_0} \frac{dz}{g(z)^2} + \int_{z_0}^{z_f} \frac{dz}{g(z)^2} \right) \left(2E_{max} + 28 \int_0^{z_0} \frac{dz}{g(z)} + 28 \int_{z_0}^{z_f} \frac{dz}{g(z)} \right) \\ &= \left(\int_0^{z_0} \frac{dz}{g(z)^2} + \mathcal{O}(\sqrt{N/N_0}) \right) \left(2m + 28 \int_0^{z_0} \frac{dz}{g(z)} + \mathcal{O}(n) \right). \end{aligned}$$

We know from Lemma 7 that $g > 1/\text{poly}(n)$ in the interval $[0, z_0]$. Since $z_0 \leq 1$, then

$$\int_0^{z_0} \frac{dz}{g(z)^p} < \mathcal{O}(\text{poly}(n)).$$

Thus, $C = \mathcal{O}(\sqrt{N/N_0})$ up to factors polynomial in n . Therefore the implicit schedule implemented by Eq. (2.5) achieves a time complexity

$$T = \mathcal{O}(\sqrt{N/N_0}).$$

□

2.7.3 Quadratic speed-up: proof of optimality

Proof. For the proof of Lemma 2 we take advantage of a permutation symmetry present in model (2.2). Define the set of solution states $S = \{s | E(s) = 0\}$. The Hilbert space $\mathcal{H}_S := \text{Span}\{|s\rangle | s \in S\}$ is the associated solution space. Define the permutation operator R which acts in the solution space $\mathcal{H}_S$ by interchanging two solution state vectors. Since $[R, H] = 0$, if we start the evolution with the state $|\phi\rangle$, which is invariant under the action of R , the symmetry is preserved in the solution space along the entire evolution. Therefore, $|\psi_t\rangle := U(t)|\phi\rangle$, with $U(t)$ the propagator for H , is invariant under permutation in $\mathcal{H}_S$. Since the entire dynamics of the problem is restricted to $\mathcal{H}_a$ (see proof of Lemma 5), we can write $|\psi_t\rangle$ as

$$|\psi_t\rangle = \sum_{i=0}^{m'-1} c_i(t) |a_i\rangle, \quad c(t) \in \mathbb{C}. \quad (2.29)$$

Define the projector $P = \sum_{s \in S} |s\rangle \langle s|$ onto the solution space $\mathcal{H}_S$ ⁷. We can relate $c_0(t)$ in (2.29) with the detection probability p at $t = T$ as

$$p = \langle \psi_T | P | \psi_T \rangle = |c_0(T)|^2 \rightarrow c_0(T) = \sqrt{p} e^{i\theta_T}$$

for some θ_T real.

Define the *unnormalized* state

$$|\omega_t\rangle := \frac{1}{\sqrt{p}} P |\psi_t\rangle = \frac{c_0(t)}{\sqrt{p}} |a_0\rangle .$$

Note that $|\omega_T\rangle$ is normalized. Define the error measure

$$E(t) := \|\omega_T - \omega_t\|^2 = \|e^{i\theta_T} |a_0\rangle - \omega_t\|^2 = 1 + \langle \omega_t | \omega_t \rangle - 2 \operatorname{Re} (e^{-i\theta_T} \langle a_0 | \omega_t \rangle) . \quad (2.30)$$

We differentiate the error (2.30) wrt t and obtain

$$\partial_t E = \frac{i}{p} \langle \psi_t | [H, P] | \psi_t \rangle - \frac{2}{\sqrt{p}} \operatorname{Im} (e^{-i\theta_T} \langle a_0 | PH | \psi_t \rangle) . \quad (2.31)$$

where we use that $|\psi_t\rangle$ satisfies the Schrödinger equation $i\partial_t |\psi_t\rangle = H |\psi_t\rangle$.

For $s \in S$,

$$H |s\rangle = \left(z \sum_{s'} E(s') |s'\rangle \langle s'| - |\phi\rangle \langle \phi| \right) |s\rangle = -\frac{1}{\sqrt{N}} |\phi\rangle .$$

The commutator can be computed from

$$HP = \sum_{s \in S} H |s\rangle \langle s| = -\sqrt{\frac{N_0}{N}} |\phi\rangle \langle a_0|$$

as

$$[H, P] = HP - (HP)^\dagger = \sqrt{\frac{N_0}{N}} (|a_0\rangle \langle \phi| - |\phi\rangle \langle a_0|) .$$

⁷Note that P is different from the projector onto the particular state $|a_0\rangle = \frac{1}{\sqrt{N_0}} \sum_{s \in S} |s\rangle$ (see Eq. (2.16)).

Inserting the commutator in (2.31) and taking the absolute value we get

$$|\partial_t E| \leq 2\sqrt{\frac{N_0}{N}} \left(\frac{1}{p} + \frac{1}{\sqrt{p}} \right)$$

where we used that $P|a_0\rangle = |a_0\rangle$. By integrating in time and using $|E(T) - E(0)| \leq \int_0^T |\partial_t E| dt$,

$$|E(T) - E(0)| \leq 2T\sqrt{\frac{N_0}{N}} \left(\frac{1}{p} + \frac{1}{\sqrt{p}} \right).$$

From definition (2.30) we have

$$E(T) = 0, \quad E(0) = \left| e^{i\theta_T} - \sqrt{\frac{N_0}{Np}} \right|^2. \quad (2.32)$$

Therefore

$$\left| e^{i\theta_T} - \sqrt{\frac{N_0}{Np}} \right|^2 \leq 2T\sqrt{\frac{N_0}{N}} \left(\frac{1}{p} + \frac{1}{\sqrt{p}} \right)$$

from which we obtain the final result

$$\sqrt{\frac{N}{N_0}} \frac{p}{2} \frac{\left(1 - \sqrt{\frac{N_0}{Np}}\right)^2}{1 + \sqrt{p}} \leq T.$$

□

2.7.4 Random 3-SAT: density of states approximation

For 3-SAT on n spins with m clauses define $E(s) := E_s = \sum_{a=1}^m e_a(s)$ with

$$e_a(s) = \prod_{i \in \mathcal{S}_a} \frac{1 + J_{ai}s_i}{2}$$

and $\mathcal{S}_a$ a subset of 3 spins. A specific instance of 3-SAT defines a density of states $N_E = \sum_{s=1}^N \delta_{E, E_s}$ with $N = 2^n$.

In general, N_E is intractable to compute. However, for random 3-SAT we can compute its expected value and the fluctuations, as follows. $J_{ai} = \pm 1$ are independent uniform random variables. Therefore, the $e_a(s) = 0, 1$ are independent binary variables with $p(e_a = 1) = \frac{1}{8}$ and the probability that $E(s) = E$ is binomial distributed

$$p_E = \binom{m}{E} p^E (1-p)^{m-E}, \quad p = \frac{1}{8}. \quad (2.33)$$

Therefore the expected density of states satisfies $\langle n_E \rangle = p_E$ where $n_E = \frac{N_E}{N}$.

To compute the variance in n_E due to instance by instance fluctuations, we need to compute the correlation matrix $\Sigma_{EE'} := \langle n_E n_{E'} \rangle - p_E p_{E'}$.

Lemma 9. *For random 3-SAT*

$$\langle n_E n_{E'} \rangle = \frac{p_E}{N} \left(\delta_{E,E'} + \sum_{d=1}^n \binom{n}{d} p(E'|E, d) \right) \quad (2.34)$$

with $p(E'|E, d)$ given by Eq. (2.38) with p_0, p_1 depending on d as given by Eqs. (2.36) and (2.37).

Proof. Since $N_E = \sum_s \delta_{E_s, E}$ we find

$$\begin{aligned} \langle N_E N_{E'} \rangle &= \sum_{s, s'} \langle \delta_{E_s, E} \delta_{E_{s'}, E'} \rangle = \delta_{E, E'} \sum_s \langle \delta_{E_s, E} \rangle + \sum_{s, s' \neq s} \langle \delta_{E_s, E} \delta_{E_{s'}, E'} \rangle \\ &= p_E N \delta_{E, E'} + \sum_{s, s' \neq s} \langle \delta_{E_s, E} \delta_{E_{s'}, E'} \rangle. \end{aligned} \quad (2.35)$$

In order to evaluate the second term, we need to compute the joint probability $p(E(s) = E, E(s') = E')$. For given $s \neq s'$,

$$e_a(s) e_a(s') = \prod_{i \in \mathcal{S}_a} \frac{1}{4} (1 + J_{ai} s_i) (1 + J_{ai} s'_i).$$

$e_a(s) e_a(s') = 1$ iff $s_i = s'_i$ for all $i \in \mathcal{S}_a$, ie. when $s_a = s'_a$, with s_a, s'_a the three

bits of state s, s' in clause a . Thus, by randomizing over the J_{ai} ,

$$p(e_a(s)e_a(s') = 1 | s_a = s'_a) = \text{Prob} \left(\prod_{i \in S_a} \frac{1}{2} (1 + J_{ai} s_i) = 1 \right) = \frac{1}{8}.$$

and zero for $s_a \neq s'_a$. Denote d the number of bits that s, s' differ. The probability that s, s' are identical on the three bits in S_a is

$$p_1 := p(s_a = s'_a | d) = \frac{\binom{n-d}{3}}{\binom{n}{3}} \quad d = 1, \dots, n. \quad (2.36)$$

Note that $p(s_a = s'_a | d) = 0$ when $d \geq n - 3$. Thus

$$\begin{aligned} p(e_a(s) = 1, e_a(s') = 1 | d) &= p(e_a(s) = e_a(s') = 1 | s_a = s'_a) p(s_a = s'_a | d) = \frac{1}{8} p_1 \\ p(e_a(s) = 0, e_a(s') = 1 | d) &= p(e_a(s') = 1 | d) - p(e_a(s) = 1, e_a(s') = 1 | d) = \frac{1}{8} (1 - p_1) \\ p(e_a(s') = 1 | e_a(s) = 1, d) &= \frac{p(e_a(s) = 1, e_a(s') = 1 | d)}{p(e_a(s) = 1)} = p_1 \\ p(e_a(s') = 1 | e_a(s) = 0, d) &= \frac{p(e_a(s) = 0, e_a(s') = 1 | d)}{p(e_a(s) = 0)} = \frac{1}{7} (1 - p_1) = p_0. \end{aligned} \quad (2.37)$$

We write $p(E(s) = E, E(s') = E' | d) = p(E)p(E' | E, d)$ with $p(E)$ given by Eq. (2.33). We condition the probability of the outcomes of $e_a(s')$ on the outcomes of $e_a(s)$. We sort the outcomes of $E_{a=1:m}(s)$ as

$$\begin{aligned} e_a(s) &= \underbrace{1, \dots, 1}_E, \underbrace{0, \dots, 0}_{m-E} \\ e_a(s') &= \underbrace{1, \dots, 1}_x, \underbrace{0, \dots, 0}_{E-x}, \underbrace{1, \dots, 1}_{E'-x}, \underbrace{0, \dots, 0}_{m-E-(E'-x)}. \end{aligned}$$

For $a = 1 : E$, $e_a(s') = 1$ x times and zero $E - x$ times. For $a = E(s) + 1 : m$, $e_a(s') = 1$ $E' - x$ times and zero $m - E - (E' - x)$ times. x can range as $0 \leq x \leq E$ and $0 \leq E' - x \leq m - E$. The latter expression is equivalent to $E + E' - m \leq x \leq E'$. Therefore, $\max(E + E' - m, 0) \leq x \leq \min(E, E')$. Define

the binomial distribution $B(k|n, p) = \binom{n}{k} p^k (1-p)^{n-k}$. Then

$$p(E'|E, d) = \sum_{x=\max(E+E'-m, 0)}^{\min(E, E')} B(x|E, p_1) B(E' - x|m - E, p_0). \quad (2.38)$$

Finally, define

$$\begin{aligned} \sum_{s, s' \neq s} \langle \delta_{E_s, E} \delta_{E_{s'}, E'} \rangle &= \sum_{s, s' \neq s} p(E_s = E) p(E_{s'} = E' | E_s = E) \\ &= N p(E) \sum_{d=1}^n \binom{n}{d} p(E'|E, d). \end{aligned}$$

Substitution in Eq. (2.35) gives the desired result. $\square$

We compare the theoretical estimate of $\Sigma_{EE'}$ with a numerical estimate in Fig. 2.4 for $n = 15$, showing excellent agreement.




Figure 2.4 Comparison of analytical estimate of $\Sigma_{EE'}^{(th)} = \langle n_E n_{E'} \rangle - p_E p_{E'}$ using Eqs. 2.33 and (2.34) with numerical estimates $\Sigma_{EE'}^{(num)}$ using 10000 random 3-SAT instances with $n = 15, \alpha = 4.2$. **Left:** Diagonal Σ_{EE} versus E . **Right:** $\Sigma_{EE'}$ versus E, E' . Maximal absolute error $\max_{E, E'} |\Sigma_{EE'}^{(th)} - \Sigma_{EE'}^{(num)}| = 2.9 \times 10^{-6}$.

With Lemma 9 we numerically compute the variance in Z_1 and Z_2

$$\text{Var}(Z_1) = \sum_{E, E'=1}^m \frac{\langle n_E n_{E'} \rangle}{E E'} - \langle Z_1 \rangle^2, \quad \text{Var}(Z_2) = \sum_{E, E'=1}^m \frac{\langle n_E n_{E'} \rangle}{E^2 E'^2} - \langle Z_2 \rangle^2.$$

We find numerically that $\sqrt{\text{Var}(Z_1)}$ scales as $\mathcal{O}(n^{-5/2})$ and $\sqrt{\text{Var}(Z_2)}$ scales

as $\mathcal{O}(n^{-7/2})$ (see Fig. 2.5).




Figure 2.5 Scaling of $\sqrt{\text{Var}(Z_1)}$ (blue) and $\sqrt{\text{Var}(Z_2)}$ (red) versus n for 3-SAT problem with $\alpha = 4.2$. Dotted lines are $n^{-5/2}$ (blue) and $n^{-7/2}$ (red).

Chapter 3

Stochastic optimal control of open quantum systems

Aarón Villanueva¹ and Hilbert Kappen¹

¹Faculty of Science, Radboud University, Heyendaalseweg 135, 6525 AJ Nijmegen, The Netherlands

Correspondence: aaronv@science.ru.nl

We address the generic problem of optimal quantum state preparation for open quantum systems. It is well known that these systems can be simulated by quantum trajectories described by a stochastic Schrödinger equation. In this context, the state preparation becomes a stochastic optimal control problem. The latter requires the solution of a Bellman equation that is, in general, challenging to solve. A notable exception are the so-called path integral control problems, for which one can estimate the optimal control solution by direct sampling of the cost objective. In this work, we formulate a class of quantum state preparation problems that are amenable to path integral control techniques, and propose a new quantum control algorithm, Path integral Quantum Control (PiQC). PiQC is a gradient-free method that employs adaptive importance sampling in the optimization step, a technique where the controls are iteratively improved based on global averages over quantum trajectories. Moreover, we demonstrate that an annealing variant of PiQC, where the environmental couplings are progressively modified across optimization, finds high accuracy solutions for unitary control problems. We further discuss potential hardware implementations of PiQC. We illustrate the effectiveness of our approach through examples of open-loop control for single- and multi-qubit systems.

Citation: Villanueva, A., & Kappen, H. (2026). Path integral control of open quantum systems. *Physical Review A*, 114(1), 012426.

DOI: <https://doi.org/10.1103/kwxb-3jy1>

3.1 Introduction

The effective and efficient manipulation of quantum systems to accomplish specific behaviors underpins a broad spectrum of technological and scientific applications, including quantum error correction [111], quantum simulation [112], quantum metrology [113], quantum sensing [114], state preparation [115], quantum state transfer [116], and quantum gate synthesis [117]. Recent advancements in experimental capabilities for measuring and manipulating quantum systems at the atomic, molecular, and optical levels [118–122] have further fueled the development of efficient and scalable quantum control techniques [16, 123].

A number of quantum optimal control techniques have been developed for open-loop control problems, i.e. problems in which the controls are time-dependent deterministic functions that are independent of the state [15,

124]. In the unitary regime, diverse approaches have been developed for open-loop quantum optimal control [125], including the gradient-based GRAPE algorithm [126] and its variants [127], search-based methods such as CRAB [128], variational approaches like Krotov optimization [129–132], and techniques derived from the Pontryagin Maximum Principle (PMP) [133]. These methods have been successfully applied to a broad range of quantum control tasks; for a review see [16, 123].

At the same time, the problem of controlling open quantum systems received substantial attention in recent years [16, 134], although with less development than in the closed systems case. Approaches relying on dynamical simulation of master equations, such as the Lindblad equation [135–137], become significantly more computationally intensive with increasing system size, as they necessitate the evolution of density matrices rather than state vectors. In this context, the theory of stochastic *quantum trajectories* [35–38] constitutes a useful framework for simulating open system dynamics in presence or absence of measurements, and has found numerous applications [138–144]. Quantum trajectories can be thought as noisy paths in the Hilbert space of the system, and serve to model the system interactions with the environment. In particular, in quantum control, quantum trajectories have been employed for approximating gradients in local optimization, benefiting from the quadratic scaling advantage of propagating wave functions of size $N = 2^n$, over density matrices of size N^2 [145].

A class of quantum trajectories are continuous diffusions described by stochastic Schrödinger equations (SSEs) [37, 40, 146]. Within the language of SSEs, the problem of optimal control of open quantum systems naturally leads to the definition of a stochastic optimal control (SOC) problem [147–150]. Approaching this class of problems generally requires solving a Bellman equation or applying a stochastic variant of the PMP via backward stochastic differential equations [151, 152], both of which may pose significant computational challenges. Notable exceptions constitute the so-called linear-quadratic-Gaussian (LQG) problems, where the dynamics is linear, the cost is quadratic in the control, and the noise is additive Gaussian. While SSEs with coherent controls are not in LQG form due to the bi-linearity introduced

by terms such as $u\psi$ with u the control and ψ the quantum state, certain cases admit LQG formulations in the Heisenberg picture [150], enabling applications in feedback control [119, 122, 153, 154]. In general, however, the control of open quantum systems governed by SSEs is a nonlinear control problem, for which no efficient solution method has been widely available.

In this line, the path integral (PI) control method [22, 155, 156] was proposed to efficiently solve a large class of non-linear SOC problems, in which the cost objective is quadratic in the control and allows for nonlinearities in the state; this includes LQG problems as a special case. One key advantage of the PI formulation is that an optimal control solution can be written in closed-form as a path integral over dynamical trajectories in the configuration space, circumventing the need for solving the Bellman or PMP equations. This further enables leveraging well known statistical estimation methods to approximate the optimal solution. An important example is adaptive importance sampling (IS) [22], which iteratively improves the estimation of the solution based on self-generated trajectories that are computed using the control estimations of previous steps, constituting in method that is highly parallelizable, and with small computational overhead. PI control has demonstrated remarkable success in robotics, enabling solutions to high-dimensional, nonlinear problems with real-time constraints, where traditional methods often fail [157]. See [158] for a recent review on PI control and related advances.

3.1.1 Contributions

In this work we apply the path integral control theory [22, 155, 156] to address the problem of optimal open-loop control of open quantum systems governed by a class of Lindblad equations. We formulate the deterministic quantum control problem as a stochastic optimal control problem, which can be addressed using path integral control techniques. This connection allows us to integrate the path integral control method in the existing ecosystem of quantum control algorithms, as a new approach to approximate optimal

controls for noisy quantum systems. We provide expressions to approximate control solutions as weighted averages over quantum trajectories. We use adaptive importance sampling [22] to approximate the optimal controls efficiently. In adaptive importance sampling, the estimation of the controls at each optimization step is iteratively improved by computing global statistics extracted from trajectories generated using the controls computed in previous steps. See Fig. 3.1 for an illustration of the workflow of the algorithm.




Figure 3.1 Schematic representation of the workflow of our control algorithm in the case of state preparation of a single qubit. Consider an initial state ψ_0 and a target state ϕ . At the start of the optimization, generate a batch of quantum trajectories (blue solid lines) controlled by an initial control guess $u^{(1)}$. Through adaptive importance sampling, all trajectories contribute collectively to update $u^{(1)}$ to a new control $u^{(2)} = u^{(1)} + \Delta u^{(1)}$, in such a way that the trajectories of the next step will be a bit closer to the target. In the next step, new trajectories are sampled using $u^{(2)}$, and the information of the samples is used to compute the next control $u^{(3)}$. This cycle is repeated until the target is reached. Additionally, in the annealing variant of the algorithm (see Sec. 3.3.4) the noise in the batch of trajectories is simultaneously reduced with the optimization steps, which further reduces the variance between trajectories and approaches the average dynamics to that one of a close system.

We provide a framework to compute controls that can be written as linear combinations of basis functions, which includes pulse-based controls. Pulse-based control may offer several potential advantages over circuit-based

control. For instance, on four transmon qubits, a reduction in state preparation time of roughly three orders of magnitude compared to gate based strategies was demonstrated [159]. More recently, authors [160] have recently applied pulse-based optimal control for ground state preparation, demonstrating that pulse-based control required less time than an equivalent digital circuit. These approaches are limited to closed dynamics and open loop control. PiQC provides a realistic and computationally tractable approach to extend previous work in the direction of open quantum systems and feedback control.

This work is significant for several reasons. We apply the path integral control technique, as formulated in [22, 155, 156], to the control of a particular class of open quantum systems for the first time. We refer to this method as PiQC. While PiQC can be applied in principle to both open-loop and feedback control, in this work we restrict to the open loop setting. PiQC demonstrates competitive, and in some contexts, superior performance and robustness in numerical experiments compared to traditional gradient-based methods such as GRAPE.

Our method differs fundamentally from previous methods that use unravelings to compute optimal controls for open quantum systems. In general, those methods solve a deterministic control problem with the dynamics given by the Lindblad equation. Unravelings are then used to approximate the evolution of the density matrix and the gradients of the cost objective. Here, instead, we formulate and solve a stochastic optimal control problem, and avoid the need for gradient computations by using an adaptive importance sampling scheme in the optimization.

In addition, we demonstrate that integrating PiQC with annealing in the environmental coupling yields a highly effective approach for obtaining high-quality solutions to the unitary control problem in the absence of dissipation. Furthermore, we outline how PiQC could, in principle, be implemented on quantum hardware, offering a continuous-time analog to variational quantum algorithms designed for digital circuits.

The rest of the paper is organized as follows. Section 3.2 presents the main results from path integral control theory, which is meant to be self-

contained but expect the reader to be acquainted with basics on the theory of stochastic processes and Itô calculus; auxiliary proofs are provided in the Supplementary Material. Section 3.3 combines quantum trajectories with path integral control to formulate a control algorithm for Lindblad equations that belong to a certain class. Section 3.3.3 summarizes our control algorithm for open quantum systems, and Section 3.3.4 presents an annealing variant of the algorithm tailored for controlling closed quantum systems. Section 3.3.5 discusses the possibility of a hardware implementation of the algorithm. Section 3.4 contains proof-of-principle numerical examples of state preparation for single and multi-qubit open systems. We benchmark PiQC with OpenGRAPE –a variant of GRAPE algorithm for open systems– for the single qubit scenario and find that PiQC outperforms the latter. For the multi-qubit case we choose Nuclear Magnetic Resonance systems as testbed for state preparation of the open system control problem (Section 3.4.3) and the annealing variant for the closed system control problem (Section 3.4.4). Section 4.6 discusses the benefits and limitations of the algorithm, and prospects of future work.

3.2 The path integral control method

We review the path integral control theory [155] in the most general form that enables optimal feedback control solutions. This Section is meant to be self-contained, but the reader may consult [156] for further details.

The path integral control framework is naturally expressed in the language of stochastic calculus using the Itô convention [161]. Consider a finite time interval $[t_0, T]$, where t_0 is the initial time and T the final time. Let W_t denote an n_c -dimensional real-valued stochastic process with vanishing mean and variance proportional to the elapsed time $t - t_0$. Such processes are referred to as Wiener processes, or equivalently Brownian diffusions [161]. Intuitively, one may regard a Wiener process as a temporal sequence of independent Gaussian increments dW_t , each of which with zero mean and variance proportional to dt . In differential form, the Wiener increments satisfy $\langle dW_t \rangle = 0$

and $\langle dW_t^a dW_t^b \rangle = D_{ab}dt$, where D is a time-independent real, symmetric, positive semidefinite matrix, and the expectation $\langle \cdot \rangle$ is taken over all sample paths of W_t on $[t_0, T]$. For compactness, we often write $\langle dW_t dW_t^T \rangle = Ddt$, and refer to this as the Itô structure of the process. When needed, we will write $\langle \cdot \rangle_{X_t=x}$ to denote expectation over trajectories conditioned on $X_t = x$ at time t .

Consider a dynamical system of the form

$$dX_t = f(t, X_t)dt + g(t, X_t)(u(t, X_t)dt + dW_t) \quad (3.1)$$

for $t_0 \leq t \leq T$ with X_t an M -dimensional real stochastic variable, with initial condition $X_{t_0} = x_0$. Here W_t is a n_c -dimensional real Brownian motion with $\langle dW_t^a \rangle = 0$ and $\langle dW_t^a dW_t^b \rangle = D_{ab}dt$ with D a $n_c \times n_c$ positive symmetric matrix. Define the spacetime domain $\Omega = [t_0, T] \times \mathbb{R}^M$. Define the functions $f : \Omega \rightarrow \mathbb{R}^M$, $g : \Omega \rightarrow \mathbb{R}^{M \times n_c}$ and the feedback control $u : \Omega \rightarrow \mathbb{R}^{n_c}$ with the sufficient regularity such that a solution to (3.1) exists. The reader may consult [162] for further details on these conditions. Denote by $\mathcal{U}$ the set of admissible feedback controls u . Define the cost ¹

$$S^u(t) = \Phi(X_T) + \int_t^T \left[V(s, X_s)ds + \frac{1}{2}u(s, X_s)^T R u(s, X_s)ds + u(s, X_s)^T R dW_s \right] \quad (3.2)$$

with X_t the stochastic solution of (3.19) at time t . The matrix R is positive $n_c \times n_c$ symmetric such that $R = \lambda D^{-1}$ with $\lambda > 0$ a constant. The function V is a state-dependent path cost, and Φ denotes the end cost.

The stochastic optimal control problem consists of finding the optimal feedback control $u^*(t, x) \in \mathcal{U}$ that minimizes the expected cost

$$C[u] = \langle S^u(t_0) \rangle \quad (3.3)$$

¹Note that the expectation of the last term in (3.2) vanishes, since it represents a Martingale [161]. However, we keep it in the definition for the sake of completeness, as this term is needed when deriving the optimal solution of the path integral theory [22, 156].

with initial condition $X_{t_0} = x_0$ at time t_0 , in symbols

$$u^* = \underset{u \in \mathcal{U}}{\operatorname{argmin}} C[u] \quad (3.4)$$

We say that the stochastic differential equation (3.1) is in PI form or PI-compatible, and together with (3.3) and (3.4) they define the PI control problem. The generic approach to solve this would be to derive a partial differential equation for the so called cost-to-go, known as Bellman equation, and derive the optimal control from the solution (see Lemma 10 in the Supplementary Material). However, for the class of PI-compatible control problems defined by Eqs. (3.1), (3.2) and (3.4), one bypasses this requirement by establishing the following result.

Theorem 1. *The optimal expected cost $C^* = C[u^*]$ and optimal control u^* for the control problem (3.4) with dynamics (3.1) and cost (3.3) is given by*

$$C^* = -\lambda \log \left\langle e^{-S^u(t_0)/\lambda} \right\rangle \quad (3.5)$$

$$u^*(t, x) = u(t, x) + \lim_{dt \rightarrow 0} \frac{\langle \omega_t^u \Delta W_t \rangle_{X_t=x}}{dt} \quad \forall (t, x) \in \Omega \quad (3.6)$$

$\omega_t^u := \frac{e^{-S^u(t)/\lambda}}{\langle e^{-S^u(t)/\lambda} \rangle_{X_t=x}}$, $u(t, x) \in \mathcal{U}_t$ is an arbitrary control sampler, and $\Delta W_t = W_{t+\Delta t} - W_t$.

Proof. The proof is given in [156] and reproduced in the Supplementary Material. □

Theorem 1 gives an explicit solution of the optimal expected cost C^* and the optimal control for any $(t, x) \in \Omega$ in terms of path integrals, i.e. as weighted expectations over u -controlled trajectories traced in the spacetime Ω . This implies that the optimal control solution u^* can be approximated using known sampling methods. In this regard, we make an important remark.

Remark 1. The control function u in (3.5) and (3.6) is referred to as the sampling

control or, in short, the sampler. Note that the l.h.s. of (3.5) and (3.6) are *independent* of u : the same result holds for any sampler u , in particular, for $u(x, t) = 0$. However, naive sampling strategies can yield estimates with large statistical variance, for which the accuracy of one-step sampling estimates of u^* heavily relies on a relatively good choice of the sampler u . •

In practice, we must approximate the expectations in (3.5) and (3.6) using a time discretization and a finite number of samples. Denote by $S_i^u(t)$ the cost of the i -th sample trajectory $X_{t:T}$ generated by a particular noise realization $W_{t:T}$. Denote with N_{traj} the number of noise realizations/trajectories. Each trajectory is weighted by $\omega_t^i = \frac{e^{-S_i^u(t)/\lambda}}{\frac{1}{N_{\text{traj}}} \sum_{j=1}^{N_{\text{traj}}} e^{-S_j^u(t)/\lambda}}$. When $S_i^u(t)$ has large variance, the batch of samples is dominated by one sample: the sample with lowest $S_i^u(t)$, in particular for low λ (i.e. low noise). The quality of the sampling can be quantified by the effective sample size [156]

$$ESS := N_{\text{eff}}/N_{\text{traj}}, \quad N_{\text{eff}} := \left(\sum_{i=1}^{N_{\text{traj}}} (\omega_t^i)^2 \right)^{-1} \quad (3.7)$$

where $1/N_{\text{traj}} \leq ESS \leq 1$ and $ESS = 1$ is maximal when $N_{\text{eff}} = N_{\text{traj}}$ with all samples having equal weight $\omega_t^i = 1/N_{\text{traj}}$. In [156], it was shown that sampling controls with lower control cost (3.2) are better samplers in the sense of smaller statistical variance and larger ESS . It turns out that the optimal sampler u , in the sense of minimal variance, is given by the optimal control solution u^* , for which the variance of $S^u(t)$ is zero, and $\omega_t^i = 1/N_{\text{traj}}$, $ESS = 1$. We make an important remark.

Remark 2. Any sampler u induces an unbiased estimator for C^* and u^* , but with different quality in terms of statistical variance. When the sampler $u = u^*$, the variance of $S^u(t)$ becomes identically zero, which further implies that $\langle S^{u^*}(t) \rangle = S^{u^*}(t)$ and shows the self-consistency of (3.6), since $\langle \omega_t^u \Delta W_t \rangle = \omega_t^u \langle \Delta W_t \rangle = 0$.

•

The remark above justifies the use of adaptive importance sampling for control estimation, as we discuss in Sec. 3.2.2.

Finally, the quality of the control solution is in principle given by the total expected cost $\langle S^u(t_0) \rangle$. Because of the monotonic relation between control cost and effective sample size [156], one can also measure the quality of the control solution in terms of ESS instead of $\langle S^u(t_0) \rangle$, which turns out to be a more sensitive measure of optimality. See Sec. 3.4.1.

3.2.1 Parametrized controls and optimality

The practicality of Equation (3.6) is limited to the computational complexity of estimating the optimal control $u^*(t, x)$ for each spacetime point $(t, x) \in \Omega$. For high-dimensional problems, this task can turn challenging due to the curse of dimensionality. Because of this, we must resort on a different strategy to approximate u^* . One strategy is to assume that the control can be described by parametrized models in terms of known functions. Here, we consider linear parametrizations in terms of basis functions. Define a control model $\hat{u}$ as a linear parametrization of the form

$$\hat{u}_a(t, x) = \sum_{k \in I} \hat{A}_{ak} h_k(t, x) \quad a = 1, \dots, n_c \quad (3.8)$$

where $\hat{A}_{ak}$ are time-independent constants and $h_k : \Omega \rightarrow \mathbb{R}$ are fixed basis functions. The symbol I denotes the index set for labeling basis functions.

For the moment, assume that the optimal control u^* and the sampler control u in (3.6) belong to the Span of $\{h_k\}_{k \in I}$, i.e. u^* and u can be written in the form (3.8) as

$$u_a^*(t, x) = \sum_{k \in I} A_{ak}^* h_k(t, x), \quad u_a(t, x) = \sum_{k \in I} A_{ak} h_k(t, x) \quad a = 1, \dots, n_c \quad (3.9)$$

for some $A_{ak}^*, A_{ak} \in \mathbb{R}$. Under this assumption, one can derive, using Theorem 2 from the Supplementary Material with $f(t, x) = h_{k'}(t, x)$, the following

matrix equation

$$\sum_k (A_{ak}^* - A_{ak}) B_{kk'} = C_{ak'} \quad (3.10)$$

with

$$\begin{aligned} B_{kk'} &= \left\langle e^{-S^u(t_0)/\lambda} \int_{t_0}^T h_k(t, X_t) h_{k'}(t, X_t) dt \right\rangle \\ C_{ak} &= \left\langle e^{-S^u(t_0)/\lambda} \int_{t_0}^T h_k(t, X_t) dW_t^a \right\rangle \end{aligned} \quad (3.11)$$

Solving (3.10) for $\hat{A}$ in terms of the matrices A, B and C gives

$$A^* = A + CB^{-1} \quad (3.12)$$

where we assumed B to be invertible.

We remark that (3.12) can be derived as a particular case of Corollary 4 in [156]. Equation (3.12) shows that when u^* is of the parametrized form (3.9), u^* can be approximated in terms of statistics C, B , which in turn can be computed once for a single batch of N_{traj} noise realizations $W_{t_0:T}$. This is in stark contrast to the case of direct estimation of u^* , where we must solve for $u^*(t, x)$ for each $(t, x) \in \Omega$, which implies a samplings complexity $\sim N_{\text{traj}}L$, with L the total number of spacetime points in the discretization of Ω .

What if u^* does not belong to the Span of $\{h_k\}_{k \in I}$? In this case, one can prove that the solution given by (3.12) is optimal in a precise sense. A complete explanation of this case would take us farther than the scope of this work, therefore we will limit ourselves in briefly articulating the main ideas, while referring the reader to [22] for further details. It turns out that the optimal control problem stated in (3.4) can be formulated as a Kullback-Leibler divergence minimization problem between the distribution p^* over state trajectories $X_{t_0:T}$ controlled by the optimal feedback control u^* , and the distribution p^u over trajectories controlled by u [22]. In symbols, we want to minimize $\text{KL}(p^*|p^u)$.

The problem of minimizing $\text{KL}(p^*|p^u)$ is convex in p^u , and the optimal

solution is reached when $u = u^*$. This corresponds to the unconstrained case in which the search space for u is the whole $\mathcal{U}$. If the search space for u is constrained, such as in the linear parametrization given by (3.8), then the solution (3.12) is optimal not in the $\text{KL}(p^*|p^u)$ divergence sense, but in the flipped version $\text{KL}(p^u|p^*)$. Obviously, minimizing the KL is not the same problem as minimizing its flipped version, but the two solutions approach each other when u^* is sufficiently well approximated by a parametrized solution of the form (3.8), because the difference $\text{KL}(p^*|p^u) - \text{KL}(p^u|p^*) = \mathcal{O}(\delta p^3)$ with $\delta p = p^u - p^*$. This is easy to see by Taylor-expanding the difference around p^* . In analogy with classical mechanics, in this formulation the problem of finding the optimal control to the SOC problem can be seen as variational problem where the KL divergence plays the role of an action.

The approach of minimizing the flipped version of the KL divergence is generally known in the literature as the Cross-Entropy (CE) method [163], and it is commonly used as a proxy for approximating the solution to the original problem. Along the text, we will adopt this point of view and frame the optimality of the solution given by (3.12) within the context of minimizing the KL divergence $\text{KL}(p^u|p^*)$.

3.2.2 Adaptive importance sampling and open-loop control

Using (3.12) we can build an optimization scheme in which the control amplitudes A are iteratively improved starting from an arbitrary guess. One can iteratively apply (3.12) using the A that is computed at iteration p as the sampler for iteration $p + 1$. This is called adaptive importance sampling and the update step in the optimization takes the form

$$A^{(p+1)} = A^{(p)} + C^{(p)}(B^{(p)})^{-1} \quad (3.13)$$

where $A^{(p)}, C^{(p)}, B^{(p)}$ are the statistics computed with $u^{(p)}$ specified by $A^{(p)}$ in the previous step. We can initialize the algorithm with e.g. $A^{(0)} = 0$, or

with smarter initializations².

In this work, we are interested in open-loop control parametrizations, i.e. parametrizations of u that are independent of the state x . In concrete, motivated by quantum control applications, we are interested in pulse-based control, i.e. in controls that are piece-wise constant in time. Consider a partition of the time interval $[t_0, T]$ in K sub-intervals $I_k = [\tau_{k-1}, \tau_k]$ such that $t_0 = \tau_0, \tau_1, \dots, \tau_K = T$. A piece-wise constant open-loop parametrization can be defined by choosing the the basis functions h_k to adopt the form of indicator functions, $h_k(t) = \delta_{t \in I_k}$, where $\delta_{t \in I_k} = 1$ if $t \in I_k$, and zero otherwise. The adaptive importance sampling equation (3.13) takes de form

$$A_{ak}^{(p+1)} = A_{ak}^{(p)} + \left\langle \omega_p \frac{\Delta W_k^a}{\Delta \tau_k} \right\rangle_p \quad k = 1, \dots, K \quad (3.14)$$

where

$$\omega_p := \frac{e^{-S^u(t_0)/\lambda}}{\left\langle e^{-S^u(t_0)/\lambda} \right\rangle_p} \quad (3.15)$$

and $\Delta W_k^a := \int_{I_k} dW_t^a$, $\Delta \tau_k := \tau_k - \tau_{k-1}$. The average $\langle \cdot \rangle_p$ is computed w.r.t. to trajectories generated by $u^{(p)}$ with initial condition x_0 at time t_0 .

3.3 Path integral control of quantum systems

3.3.1 Problem definition

We now turn to quantum control and consider the general state preparation problem. From now on, quantum states will be denoted by Greek letters such as ψ and ϕ , with their inner products written explicitly in the standard linear-algebraic form, e.g. $\psi^\dagger \phi$.

We want to prepare a quantum state from a given initial state in total time T , taking into account the interactions between system and environment. For

²For instance, with the solution of the corresponding deterministic control solution with $D = 0$.

this, we assume the system dynamics follows the Lindblad equation [164]

$$\dot{\rho} = -i[H, \rho] + \mathcal{D}[D', \mathcal{C}'](\rho) \quad \mathcal{D}[D', \mathcal{C}'](\rho) := D'_{ab} \left(C'_a \rho C'^{\dagger}_b - \frac{1}{2} \{C'^{\dagger}_b C'_a, \rho\} \right) \quad (3.16)$$

where, from now on, we adopt the convention of repeated indices for summation. The Hamiltonian H takes the form $H = H_0 + u_a H_a$ with $H_a (a = 1, \dots, n_c)$ Hermitian operators (or control Hamiltonians), and $u_a (a = 1, \dots, n_c)$ real-valued control fields. The noise matrix D is positive symmetric, and the dissipators $\mathcal{C}' = \{C'_a\}_{a=1}^{n_c}$ satisfy a property to be determined.

Given the total time interval $[0, T]$, consider the cost objective

$$C_{\text{det}}[u] = -\frac{Q}{2} \mathcal{F}(\rho_T) + \frac{1}{2} \int_0^T u(t)^T R u(t) dt \quad (3.17)$$

where ρ_T is the density operator at final time T , $Q > 0$ and R is a real symmetric positive $n_c \times n_c$ matrix. The first term corresponds to the fidelity $\mathcal{F}(\rho_T) = \text{Tr}(\rho_T \phi \phi^\dagger)$ of the final state ρ_T with respect to a target (pure) state ϕ . The second term is an energy constraint, and it penalizes the total magnitude of the control fields. The task is to find the optimal open-loop control $u^*(t) \in \mathcal{U}_o$, given a suitable open-loop control space $\mathcal{U}_o$, that minimizes the objective function (3.17), i.e.

$$u^* = \underset{u \in \mathcal{U}_o}{\text{argmin}} C_{\text{det}}[u] \quad (3.18)$$

The Lindblad equation (3.16) and the cost (3.17) define a deterministic control problem. It can be approached using, for example, the PMP method [133]. Here, instead, we propose to use path integral control, and for this we need to establish a suitable mapping between the deterministic setting given by (3.17) and (3.18), and a stochastic control problem.

3.3.2 Stochastic mapping

Consider the controlled stochastic Schrödinger equation

$$d\psi_t = -iH_0\psi_t dt - \frac{1}{2}D_{ab}H_aH_b\psi_t dt - iH_a\psi_t(u_a(t)dt + dW_t^a) \quad (3.19)$$

with Itô structure $\langle dW_t dW_t^T \rangle = Ddt$, open-loop control $u \in \mathcal{U}_o$, and Hermitian operators H_a as in (3.16). Consider the expected cost

$$C[u] = \left\langle -\frac{Q}{2}\mathcal{F}(\psi_T) + \frac{1}{2} \int_0^T u(t)^T R u(t) dt + \int_0^T u(t)^T R dW_t \right\rangle \quad (3.20)$$

where ψ_t is the quantum state satisfying the SSE (3.19), $\mathcal{F}(\psi_t) := \text{Tr}(\psi_t \psi_t^\dagger \phi \phi^\dagger)$, and $\langle \cdot \rangle$ is the expectation over all quantum trajectories generated by (3.19) with initial state ψ_0 at time $t = 0$. Define the SOC problem

$$u^* = \underset{u \in \mathcal{U}_o}{\text{argmin}} C[u] \quad (3.21)$$

Finally, define the density operator $\rho(t) := \langle \psi_t \psi_t^\dagger \rangle$. Then, the SOC problem given by (3.21) and the deterministic problem given by (3.18) are equivalent in the sense that $C_{\text{det}}[u] = C[u]$ for all $u \in \mathcal{U}_o$, provided that the dissipators $\mathcal{C}'$ are related to the operators $C_a := -iH_a$ through the following transformation

$$C_b = C'_a G_{ab} \quad D' = G D G^\dagger \quad (3.22)$$

with G an invertible matrix $G \in \mathbb{C}^{n_c \times n_c}$ such that D' is real symmetric.

The proof follows directly by recognizing that (3.19) is an unraveling of the Lindblad equation with dissipation operator $\mathcal{D}[D, \mathcal{C}]$ —this connection is well-established in the literature [40] and we include a derivation in Appendix 3.8.1 for completeness—combined with the fact that the dissipation operator remains unchanged under transformations of the form (3.22), i.e. $\mathcal{D}[D, \mathcal{C}] = \mathcal{D}[D', \mathcal{C}']$.

Note that the equivalence between the SOC formulation and the deterministic problem only holds in the case where u is an open-loop control.

When u is a feedback control, the mapping does not hold because control and state become correlated.

The stochastic mapping enables the exploitation of established methodologies for to address the latter. Specifically, we can implement path integral control techniques by recognizing that (3.19) is PI-compatible, i.e. it is written in the form of (3.1) ³. In summary, we can solve the deterministic problem (3.18) using PI control provided that the condition $R = \lambda D^{-1}$ with $\lambda > 0$ is satisfied and the dissipators $\mathcal{C}'$ are linearly related to anti-Hermitian operators through transformation (3.22). Dissipation processes obeying this property arise naturally in practical contexts, see Section 4.6. While the anti-Hermitian condition defines the range of applicability of the control framework developed here, in Section 4.6 we discuss possible pathways to circumvent such a limitation.

3.3.3 The Path integral Quantum Control algorithm for open systems

To finalize the definition of the diffusion algorithm, we must specify a suitable open-loop search space $\mathcal{U}_o$. As for quantum control applications, we are interested in pulse-based control. In consequence, we choose $\mathcal{U}_o$ to be the space of piece-wise constant functions as defined in Sec. 3.2.2. We refer to this algorithm as Path integral Quantum Diffusion Control and summarize it as follows. Set an initial sampler u .

1. Sample N_{traj} noise trajectories $W_{0:T}$
2. Generate N_{traj} state trajectories $\psi_{0:T}$ using (3.19) and record the final states ψ_T

³To see this, note that the dynamics in (3.19) can be written in the form

$$d\psi_t = f(\psi_t)dt + g_a(\psi_t)(u_a(t)dt + dW_t^a)$$

with f and g_a complex valued vector functions. Split the equation in real and imaginary components and define $X_t = (\psi_t^R, \psi_t^I)^T$, $F(\psi_t) = (f(\psi_t)^R, f(\psi_t)^I)^T$, $G_a(\psi_t) = (g_a(\psi_t)^R, g_a(\psi_t)^I)^T$. Then, equation

$$dX_t = Fdt + G_a(u_a(t)dt + dW_t^a)$$

is PI-compatible.

3. Compute N_{traj} weights ω using (3.15)

4. Update u using (3.14)

5. Repeat

In practice, due to the statistical nature of the algorithm, the controls can present high variance. We can add a smoothing procedure in Step 4 in order to reduce such variance and improve convergence. This is done as follows. In Step 4, instead of estimating the next control $u^{(p+1)}$ using $u^{(p)}$, use a recent average of controls $u^{(i)}$ computed over a window of past importance sampling steps. Operationally, this means replacing A_{ak}^p on the right hand side of (3.14) (in both terms) by the average $\frac{1}{w} \sum_{p'=p-w+1}^p A_{ak}^{p'}$, where w is a user-defined window size. This smoothing ansatz can improve the quality of solutions at the cost of slowing down the convergence.

3.3.4 Annealed diffusion control for closed systems

In this Section we propose a variation of the diffusion control algorithm of Sec. 3.3 for approximating optimal controls for closed systems.

Consider the deterministic control problem with dynamics given by the Liouville equation resulting from setting the dissipation term in (3.16) to zero. This can be achieved e.g. by setting the noise matrix D to zero. More explicitly, consider the closed dynamics

$$\dot{\rho} = -i[H_0 + u_a H_a, \rho] \quad (3.23)$$

together with the cost objective (3.17). If the initial state $\rho_0 = \psi_0 \psi_0^\dagger$, (4.5) describes a purely unitary evolution.

We can approach the problem of minimizing the cost objective above given dynamics (4.5) using the PiQC algorithm by modifying the noise matrix D across optimization steps, such that $D \rightarrow 0$ at the end of the optimization. Intuitively, this means approaching the closed dynamics by progressively switching off the interactions with the environment.

More concretely, consider a series of noise matrices

$$D_1, D_2, \dots, D_{n_{\text{IS}}}$$

with n_{IS} the number of importance sampling steps. Recall that, for the PI control method to be applicable, the relation $R = \lambda D^{-1}$ should be satisfied. Then, $D_{n_{\text{IS}}}$ cannot vanish. In consequence, we will choose $D_{n_{\text{IS}}}$ such that $\|D_{n_{\text{IS}}}\| = \varepsilon$ with ε sufficiently small but not zero. The set $\{D_p\}$ will start from relatively high noise strength values at $p = 1$ and will gradually reduce such that $\|D_{n_{\text{IS}}}\| = \varepsilon$. We denote such a set $\{D_p\}$ as a control annealing schedule or simply, annealing schedule.

The reason for using the word *annealing* stems from the fact that the progression from high-noise values at the beginning of the optimization towards low-noise values at the end of it resembles the process of *cooling down* in classical annealing. At $p = 1$, we start with very noisy quantum trajectories in the Hilbert space induced by high values of the noise strength D_1 . As the optimization progresses, these trajectories become less noisy, decreasing their relative variance, due to a weaker interaction with the environment induced by a smaller noise strength D_p . At the end of the optimization at $p = n_{\text{IS}}$, the quantum trajectories tend to coalesce into a single trajectory dictated by the unitary dynamics (4.5). In this sense, we *cool down* the stochastic trajectories to finally merge into a single unitary evolution. We remark that the analogy with classical annealing is not related to the temperature of the device, nor to any type of physical temperature. A picture of temperature arises only from the use of synthetic noise induced by the schedule D_p , which controls the amplitude of the noise manifested in the quantum trajectories.

If we choose a good enough schedule $\{D_p\}$, the algorithm will anneal towards a solution $u^{(n_{\text{IS}})}$ that is a good approximation to the true solution of the unitary problem given by (4.5) and (3.17). In this variation, PiQC can be interpreted as a quantum control annealer. We illustrate this process in Fig. 3.2.

In Section 3.4.3 we demonstrate the utility of this approach to compute optimal controls for a multi-qubit spin system.




Figure 3.2 Schematic representation of the annealing diffusion control algorithm process. In early stages of the optimization, the amplitudes of the noise matrix D_p are relatively high and the trajectories (solid blue lines) present high variance w.r.t. to the solution corresponding to $D = 0$ (green solid line). By the end of the optimization, the noise amplitudes are very small, and the trajectories tend to align with the trajectory corresponding to $D = 0$ due to a reduction of variance, while reaching the target ϕ .

3.3.5 A proposal for on-hardware PiQC

The PiQC algorithm includes a subroutine for simulating quantum trajectories with computational complexity that scales exponentially on classical hardware. This problem is common to all existing continuous time optimal control methods. However, in the case of PiQC one can accelerate the computation by running the quantum state evolutions on real quantum hardware. It can be seen as a hybrid quantum-classical optimization protocol [165]. The crucial observation is that the stochastic dynamics dictated by (3.19) corresponds to a random unitary evolution [166]. The latter can be defined by a time-dependent Hamiltonian

$$H(t) = H_0 + \sum_a Y_a(t) H_a \quad (3.24)$$

driven by a random control signal of the form $Y_a(t) = u_a(t) + \dot{W}_t^a$, where $u_a(t)$ is the deterministic part and $\dot{W}_t^a$ the stochastic part. Here the time derivative of the Wiener process is formal, and should be interpreted in the discrete sense, by choosing at time discretization sufficiently small. Importantly, it is assumed that the quantum dynamics is noiseless, and $\dot{W}_t^a$ is provided by the experimenter. This requirement is needed because the control update (3.14)

requires knowledge of the noise trajectory $W_{0:T}$. In this way, we replace classical simulations of the open dynamics (3.19) with evolutions/simulations on the quantum hardware. In Fig. 3.3 we illustrate the on-hardware PiQC workflow for a general end cost

$$\Phi(\psi_T) = \psi_T^\dagger H_{\text{target}} \psi_T \quad (3.25)$$

with H_{target} an arbitrary Hermitian matrix representing an observable. A key




Figure 3.3 At each IS step p , simulate in the quantum computer (QC) N_{traj} unitary trajectories with input control signal $Y^{(p,i)}(t) = u^{(p)}(t) + \dot{W}_t^{(p,i)}$ per trajectory i , where $u^{(p)}$ is the deterministic control at step p , and $\dot{W}_t^{(p,i)}$ is a self-generated noise at step p and trajectory i . At time T perform a quantum measurement, producing N_{traj} measurement outcomes $\Phi^{(p,i)} = \psi_T^{(p,i)\dagger} H_{\text{target}} \psi_T^{(p,i)}$ for $i = 1, \dots, N_{\text{traj}}$. Input the current control, the noise trajectories and the measurements into the IS formula (3.14) and update the control to $u_a^{(p+1)}$. Repeat.

step of the algorithm involves accurate estimation of the end cost through iterative state preparations and measurements. The tomographic complexity of this task follows a naive scaling $\propto \frac{\alpha}{\varepsilon_q^2 \varepsilon_c^2}$ [165]. Here, $\alpha = \mathcal{O}(\text{poly}(n))$ with n the system's size, and ε_q and ε_c denote the precision requirements for quantum and classical expectation value estimations, respectively. This naive scaling could be substantially improved using better estimation techniques and considering the parallelization strategies discussed in Sec. 4.6. For example, in the case of ground state preparation, H_{target} is the Hamiltonian whose ground state we want to prepare. In this case, energy measurements can be

more efficiently estimated than full state tomography using sophisticated methods [167].

As defined here, the hardware implementation of PiQC can be realized either, at the circuit level, or at the pulse level. At the circuit level one optimizes parametrized quantum circuits that simulate the quantum dynamics. The control parameters become gate parameters within the quantum circuit. Variational quantum algorithms (VQAs) operate at this level. The pulse level operates below the circuit level, and involves direct optimization of the Hamiltonian dynamics through successive applications of electromagnetic pulses, analogous to the NMR control case described in Sec. 3.4.3. Here, the control parameters are the actual physical fields driving the quantum system. The pulse level is the natural layer of abstraction for the PiQC algorithm. Methods for relating these two levels have been extensively studied [168], establishing the circuit level as a well-defined abstraction layer above the physical pulse level [165].

Simulation-based PiQC, as considered in this work, assumes precise knowledge of the Hamiltonian describing the system of interest. By comparison, on-hardware PiQC only requires that the control Hamiltonian can be written in the form (3.24), without the need to specify the detailed structure of H_0 or the operators H_a . What ultimately matters are the measurement outcomes at the end of the evolution generated by H . In this sense, when implemented directly on hardware, the algorithm becomes partially agnostic about the precise form of the system model. In the Supplementary Material we provide an experiment for the single qubit that tests PiQC robustness against imperfect knowledge of the system and when the system suffers from uncontrolled drifts during the optimization. Related model-free approaches have been investigated e.g. using Reinforcement Learning to compute optimal control policies for quantum systems [169] although we remark that a lack of a model often results in poor performance [170].

The PiQC framework assumes Gaussian white noise with a certain strength. It also assumes that the corresponding Lindblad operator is known. In practice, these assumptions may be violated and, even if they are approximately correct, they must be tested and calibrated on the particular

hardware. This is left for future work. Instead, the on-hardware PiQC proposal assumes that the device itself is noiseless, which is a common assumption in the field of quantum control. We add noise artificially for two reasons: 1) to make the problem fit in the path integral framework, and 2) because, in our experiments, adding noise improves the overall optimization process, and yields better results. It is clear that ignoring the device's noise leads to an overestimation of the quality of the computed control solution.

There are other problematic aspects related to scalability in the field of quantum control, such as barren plateaus [171], algorithmic slowdown [172] and hardware constraints [173] that we do not explore here. We ignore how these potential issues would affect the applicability of the on-hardware proposal presented here for large systems, and leave for future work the exploration of this question.

3.4 Numerical experiments

In this section we apply the PiQC algorithm to several examples to showcase its effectiveness. In what follows, we assume open-loop control parametrizations in the form of piece-wise constant control pulses over K time intervals. We parametrize the controls according to (3.8), and write

$$u_a(t) = \sum_{k=1}^K u_{ak} \delta_{t \in I_k} \quad a = 1, \dots, n_c \quad (3.26)$$

with $I_k = [\tau_{k-1}, \tau_k]$, u_{ak} the k -th pulse corresponding to control u_a , and n_c the total number of controls. To simulate the stochastic dynamics (3.19), we use the Euler-Maruyama integration scheme. We made the code implementation supporting this Section publicly available at [174].

3.4.1 Control of a noisy qubit

Consider a single-qubit system evolving according to the Lindblad equation (3.16) with $H = u_x H_x + u_y H_y$, where $H_x = \sigma_x$ and $H_y = \sigma_y$. The dissipation

part is given by the two non-Hermitian operators $C'_1 = \sigma^+$ and $C'_2 = \sigma^-$, where $\sigma^\pm = \frac{1}{2}(\sigma_x \pm i\sigma_y)$. This system is commonly used for modeling the emission and absorption of light quanta in a two-level system coupled to an electromagnetic field, e.g. a cavity resonator [38]. Assume a diagonal noise matrix $D'_{ab} = D\delta_{ab}$ with D real positive. The Lindblad equation is

$$\dot{\rho} = -i[H, \rho] + D(\sigma^+ \rho \sigma^- + \sigma^- \rho \sigma^+ - \rho) \quad (3.27)$$

To propose an unraveling suitable for PI formulation, we transform the dissipators to a pair of anti-Hermitian operators using (3.22) with $C_1 = -iH_x$, $C_2 = -iH_y$,

$$G = \begin{pmatrix} -i & -1 \\ -i & 1 \end{pmatrix} \quad (3.28)$$

The unraveling (3.19) becomes

$$d\psi_t = -D\psi_t dt - i \sum_{a=x,y} \sigma_a \psi_t (u_a(t) dt + dW_t^a) \quad (3.29)$$

As stated in Sec. 3.3.2, unraveling (3.29) is norm-preserving, which implies that all the generated stochastic trajectories will lie on the Bloch sphere.

We use as cost Eq. (3.20), subject to $RD = \lambda$. We consider the state preparation problem from an initial state $\psi_0 = |X\rangle := \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ to a target state $\phi = |Y\rangle := \frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle)$. Although the transformation $X \rightarrow Y$ can be realized by a simple σ_z rotation, this task is non-trivial, since σ_z is not one of the control primitives in H . Instead, σ_x, σ_y must coordinate so as to realize the desired rotation. As a result, the optimal trajectories do not lie on the equator.

In Fig. 3.4 (A) we show the adaptive importance sampling. We plot the average fidelity over trajectories F_{avg} , the worse case fidelity over trajectories F_{min} and the average control cost C in (3.20) versus IS steps. In addition, we plot the effective sample size ESS in (3.7), which is a sensitive measure of the quality of the optimal control solution. The ESS indicates how close the

control solution is to the optimal (feedback) control, for which $ESS = 1$. We observe that while the fidelity and control cost converge fast to constant values, the ESS still increases indicating that the quality of the control solution is still improving (shaded region). We observe that the control solution becomes smoother in these later IS iterations. The asymptotic average fidelity for $K = 128$ is $F^* = 0.9759 \pm 0.0006$. Since the optimal control solution is a compromise between the final fidelity and the fluence $U := \int_0^T u^T(t)u(t)dt$, higher fidelity solutions can be obtained by lowering R in Eq. (3.20). Indeed, by reducing $R = 1$ to $R = 0.1$ the average asymptotic fidelity increases to $F^* = 0.9963 \pm 0.0002$. Fig. 3.4 (B) shows one of the two optimal control solutions $u_{x,y}(t)$ after convergence of the IS training. The other solution is obtained by taking $u_{x,y}(t) \rightarrow -u_{x,y}(t)$. In Fig. 3.4 (C) we show how the quality of the optimal control solution in terms of asymptotic ESS depends on the number of pulses K . The ESS increases monotonically until reaching an asymptote at around $ESS \sim 0.21$, showing the sub-optimality of the open-loop control compared to the optimal feedback solution (for which $ESS = 1$). Fig. 3.4 (D) shows quantum trajectories on the Bloch sphere under optimal control.

Benchmark with Open GRAPE.— As benchmark, we compare the performances of PiQC and (first order) Open GRAPE algorithm [135], which computes coherent controls for the Lindblad equation. See the Supplementary Material for details on the algorithm and implementation. We solve the state preparation problem $X \rightarrow Y$ for the single qubit using $K = 128$ pulses. The control problem is defined by the Lindblad equation (3.27) and the cost function (3.17).

The convergence of GRAPE-type algorithms depend on the specific problem and initial guesses, therefore it is common practice to run the algorithm with multiple random control pulses and select only those that are successful [175]. We executed the PiQC and Open GRAPE algorithms 505 times, utilizing random control pulse initializations sampled from a normal distribution (centered at 1, 0, and -1, with standard deviation 2), and from a uniform distribution (integers between -10 and 10, with a scale factor of 1 and 0.5). Regarding the stopping condition for Open GRAPE, the optimization stops




Figure 3.4 Control of a noise qubit from $X \rightarrow Y$. **(A)** Average fidelity F_{avg} , minimal fidelity F_{min} , effective sample size ESS and cost C in 3.20 versus importance sampling iterations p . The converged average fidelity is $F^* = 0.9759 \pm 0.0006$. **(B)** Optimal control solution $u_{x,y}(t)$ after convergence of the algorithm. The optimal solution is two-fold degenerate. The two solutions are related by a global sign $u_{x,y} \rightarrow -u_{x,y}$. **(C)** Dependence of the quality of the optimal control, measured by the ESS , on the number of pulses K . **(D)** Optimally controlled trajectories on the Bloch sphere. Parameters: final time $T = 1$, noise coupling $D = 0.005$, control weight $R = 1$, fidelity weight $Q = 10$, number of pulses $K = 128$, number of trajectories $N_{traj} = 400$ per IS step, maximum number of IS steps $n_{IS} = 1000$ and time discretization $dt = T/N_T$ with $N_T = 128$. IS smoothing window $w = 40$ (see Section 3.2.2).

when the cost cannot be further optimized by decreasing the learning rate; see App. 3.8.3 for details. On the other hand, PiQC converges within 800 IS steps utilizing a single control initialization, i.e. $u^{(0)} = 0$.

In Fig. 3.5 we show the distribution of control costs (along the fidelities and fluences) for Open GRAPE (main panels) and PiQC (insets). We mark the average cost, fidelity and fluence of PiQC solutions by a yellow dot in each graph. The PiQC solutions have lower cost than the Open GRAPE solutions. The width of the dashed vertical lines locating each mean is much larger than the standard deviation of each PiQC distribution, showing that the variance of PiQC solutions is very narrow compared to that of Open GRAPE ones. Open GRAPE converges to many different solutions with different costs




Figure 3.5 Histograms of optimal control costs obtained with Open GRAPE (main panel) and PiQC (inset) on 505 runs. For Open GRAPE, the average cost, fidelity and fluence (and standard deviations) are $\langle C \rangle = -4 \pm 1$, $\langle F \rangle = 0.8 \pm 0.2$ and $\langle U \rangle = 3 \pm 2$, respectively. The minimum cost achieved is $C_{min} = -4.07$. For PiQC, the corresponding average values (yellow dot) are $\langle C \rangle = -4.171 \pm 0.003$, $\langle F \rangle = 0.9759 \pm 0.0006$ and $\langle U \rangle = 1.4170 \pm 0.0009$. In the main panels, the dashed vertical lines at the PiQC average are wider than the actual standard deviations of the PiQC distributions. This illustrates the large difference in variance between the two methods. All PiQC runs converge to solutions with the same cost, fidelity, fluence and effective sample size up to small statistical fluctuations due to the stochastic nature of PiQC algorithm. Parameter setting: $D = 0.005$, $T = 1$, $R = 1$, $Q = 10$, $K = 128$, $w = 20$.

depending on the initial condition. Instead, PiQC converges consistently to one of two solutions with same cost, fidelity, fluence and effective sample size, which are related by a global sign (Fig. 3.4 (B)). The difference between individual solutions has RMS error less than $\mathcal{O}(10^{-4})$. By direct inspection we can also verify that these profiles correspond to two different underlying smooth solutions. We conclude that the different PiQC solutions, rather than representing local minima, correspond to the same underlying solution up to small statistical fluctuations due to the stochastic nature of the algorithm.

We further compare these two algorithms' solutions by performing a t-SNE embedding [176] over Open GRAPE solutions together with PiQC solutions. In Fig. 3.6 (A) we plot some of the most regular Open GRAPE control solutions. In Fig. 3.6 (B) we show the t-SNE embeddings for $[u_x, u_y]$, considered as a vector of length $2K$. In Fig. 3.6 (C) we plot the average profiles of the Open GRAPE solutions of cluster 1 (red) and cluster 2 (blue) together with the average PiQC profiles. Each dot represent a different Open GRAPE solution and we represent the (average) PiQC solutions with yellow dots. We further perform a K-means decomposition to cluster the different solutions into 3 clusters (colored with red, blue and black). The different Open GRAPE solu-

tions tend to cluster around each ground truth PiQC solution. The qualitative resemblance between the averages and the PiQC solution indicates that the Open GRAPE local minima tend to cluster around the ground truth solutions represented by PiQC solutions. The cluster colored with black consists of local minima control solution of Open GRAPE with high control cost C . We excluded the black cluster from the analysis because they correspond to high cost solutions.




Figure 3.6 (A) Examples of Open GRAPE solutions with diverse levels of regularity for different seeds. (B) 2-D t-SNE dimension reduction representation followed by a K-means decomposition of Open GRAPE solutions (red, blue and black dots) and the two average PiQC solutions (yellow dots). The cluster of black dots correspond to local minima of relatively high cost C . (C) Average Open GRAPE solutions for cluster 1 (red) and cluster 2 (blue). In the bottom panels, the PiQC solutions corresponding to each cluster. (D) Robustness $\Delta C = C(u + \xi) - C(u)$ versus σ with u the high-fidelity high-ESS PiQC solution (red) and the high-fidelity low-ESS Open Grape solution (black).

During the IS iterations of the PiQC algorithm, the optimization keeps improving in terms of ESS, even when maximum fidelities are reached; see shaded area in Fig. 3.4 (A). Instead, Open GRAPE solutions, while having similar fidelities as the PiQC solution, have low $ESS = \mathcal{O}(10^{-3})$. We expect that the high fidelity, high-ESS PiQC solutions are more robust to random perturbations than the high fidelity Open GRAPE solutions. We test this by

perturbing the optimal control solution $u(t) \rightarrow u(t) + \xi(t)$ with $\xi(t)$ with Gaussian perturbations with mean zero and standard deviation σ and compute the maximal cost difference $\Delta C = C(u + \xi) - C(u)$ for 20 different noise realizations. We compare the high-fidelity Open GRAPE solutions, defined by having a cost $C < -3.8$ with the high fidelity high-ESS PiQC solution in Fig. 3.6 (D) for different σ . For completeness, we also include in the experiment the high-fidelity but low-ESS PiQC solutions corresponding to the shaded area of Fig. 3.4 (A). It can be observed that the PiQC solution (in red) is more robust to random perturbations than the Open GRAPE solutions (in black). Therefore, we conclude that the ESS is a useful additional criterion to assess the quality in terms of the robustness of the (open loop) control solution against perturbations in the control fields. See also the Supplementary Material for an experiment on robustness that suggests that PiQC is more robust under imperfect knowledge of the model Hamiltonian H_0 .

Due to discretization artifacts and a control model consisting of a finite number of pulses, the control problem is generally non-convex, and most gradient-based algorithms suffer from local minima. The yield of GRAPE-type algorithms, i.e. the probability of finding quality solutions, depends on the implementation and is usually modest compared to the number of trials [175]. For the single qubit system in this section and the NMR experiments in Sec. 3.4.3, the yield of PiQC can be regarded as approximately one, as we consistently reached quality solutions using only a single control seed. One might note that the optimization process of PiQC requires sampling N_{traj} trajectory evolutions at each optimization step to update the control, contrary to Open GRAPE that requires only two evolutions per step, forward and backward, for updating the controls. However, in PiQC the trajectory evolutions are embarrassingly parallelizable, as they are computed using a linear SSE that can be easily vectorized, considerably reducing the computation overhead. The high parallelizability of PiQC speeds up the optimization and enhances the unique feature that separates the path integral control method from standard approaches, which is the adaptive importance sampling scheme. At each optimization step, an ensemble of trajectories probes the objective landscape and collect the necessary information to update the

controls for the next step. This collective search driven by noise proves to be helpful in avoiding local minima, as we have seen in this Section for the single qubit.

3.4.2 Leveraging open-system solutions for unitary dynamics

In general, it is more difficult to compute the optimal control solution at low noise than at high noise. The reason is that $\lambda \rightarrow 0$ in the low noise case so that the effective sample size decreases. Here we explore the possibility to use the high noise solution as a proxy for the deterministic control solution. We explore this question for the noisy qubit. We compute control solutions for the open case for different noise values D , and compare the infidelities $1 - F$ scored in the open and closed cases for each control solution. We consider the $X \rightarrow Y$ state preparation problem. In the open system case, we compute for each D a control u that achieves maximal average fidelity F_{open} . We use this control to steer the unitary dynamics corresponding to $D = 0$ and computing the final fidelity F_{closed} . Then, we compare the infidelity $1 - F_{open}$ with the infidelity $1 - F_{closed}$ achieved in the unitary case. See Fig. 3.7 (A).




Figure 3.7 (A) Single qubit, $X \rightarrow Y$. In the left y-axis, infidelities $1 - F_{open}$ and $1 - F_{closed}$ vs D ; in the right y-axis, normalized fluence U/U_{max} , vs D , with $U_{max} = 1.86$. Parameters: $T = 1$, $K = 50$, $R = 0.5$, $N_T = 100$, $N_{traj} = 2000$ and Q iterates over 5, 50 and 100, and stopping when the fidelity threshold condition $1 - F_{avg} < 0.02$ is met. (B) $X \rightarrow Haar$. Distribution of D_{max} for 100 random Haar target states. Average $\langle D_{max} \rangle = 1 \pm 0.3$.

Surprisingly, the control that is estimated for non-zero D works well for the unitary case, reaching high fidelities F_{closed} for a wide range of noise values, even when F_{open} deteriorates fast as D increases. Moreover, the infidelity $1 - F_{closed}$ achieves its minimum at a higher value of D rather than at the start of the interval. This is not surprising, as the sampling becomes harder and needs more trajectories at close-to-zero values of D to maintain estimation efficiency. This high-noise effect is closely related to the annealing PiQC technique that we explore in Sec. 3.4.3.

Both infidelities $1 - F_{open}$ and $1 - F_{closed}$, show a sharp transition in the fidelity at a certain noise value between $0.8 - 0.9$. This point also corresponds with a fast decrease in the fluence. The explanation of this qualitative change in the control is that for relatively low values of D , there is a non-trivial solution that minimizes the objective due to weak coupling with the environment. When D increases beyond certain point, the noise becomes dominant and dissipation overwhelms coherent control, leading to marginal fidelity gains compared to energy costs. In that regime the open loop controller can no longer yield a high fidelity solution. As a result, the optimal solution is to minimize U instead and the final quantum state becomes maximally mixed. We remark that this phase transition is particular to the $X \rightarrow Y$ problem, and is not necessarily present for other control paths $A \rightarrow B$, for arbitrary states A and B .

We demonstrate the robustness of this effect by repeating the previous experiment where we choose target states from a random Haar measure, with fixed initial state $\psi_0 = |0\rangle$ ($X \rightarrow \text{Haar}$). For each problem, we record the maximum noise value D_{max} for which the fidelity F_{closed} is above 0.98. That is, for each instance we compute the maximum D for which the stochastic control solution gives a good fidelity in the corresponding noiseless $D = 0$ case. In Fig. 3.7 (B), we show the distribution of D_{max} for 100 random Haar states. We observe that D_{max} concentrates around 1, suggesting that the steering of the unitary dynamics is robust up to relatively high values of D . These results indicate that PiQC can be used as a noise-assisted control protocol to find or improve the quality of solutions for coherent control of unitary dynamics. This idea is closely related to the annealed diffusion control

protocol introduced in Section 3.3.4, for which we show a case example later, see Sec. 3.4.4.

3.4.3 GHZ states on open NMR systems

In this section, we test the capability of our algorithm in handling systems with multiple qubits. We demonstrate that our PiQC algorithm can successfully prepare states for realistic systems such as NMR molecules. NMR quantum computing has gained popularity in the past years as a testbed technology for experimental and theoretical research due to the long coherence times attained compared to the dynamical scales of the system [177]. NMR systems equipped with transverse field controls have been demonstrated to be fully controllable for state preparation [178], which initially positioned them as promising candidates for universal quantum simulators [179]. We select NMR systems due to their simplicity, which enables a straightforward setup for the control problem, and allows us to compare with existing published results. We remark that same approach could also be applied to the coherent control of superconducting qubits [180], ion traps [181, 182] or neutral atoms [183] systems. NMR systems are defined by a total Hamiltonian $H = H_0 + H_c$, where H_0 is the drift and H_c is the control part. The drift and control Hamiltonians are further defined as

$$H_0 = H_Z + H_I \quad H_c = \sum_{i=1}^n u_{ix} \sigma_i^x + u_{iy} \sigma_i^y \quad (3.30)$$

where $H_Z = \sum_{i=1}^n \pi D_i \sigma_i^z$ and $H_I = \sum_{i<j}^n \frac{\pi}{2} J_{ij} \sigma_i^z \sigma_j^z$. The constants D_i denote self-energies, commonly called chemical shifts, and J_{ij} are the coupling strengths between the different spins in the system. Pulse-based control is a common approach to control NMR systems. It consists of the sequential application of external radio-frequency pulses across the x - y plane. These pulses are represented by time-dependent fields u_{ia} acting on each qubit i along x and y directions.

Due to the typical large difference in magnitude between H_Z and H_I , it is

common practice to work in a rotating frame where the dynamics induced by H_Z is absent. Operationally, the rotating frame has the effect of setting $H_Z = 0$ in the drift Hamiltonian H_0 and inducing a rotation of the controls. The optimized pulses in the rotating frame can be easily transformed back to the laboratory frame. We write the details about the rotation frame approach and the proof of control equivalence in the Supplementary Material.

We write the Lindblad equation for the NMR system as

$$\dot{\rho} = -i[H, \rho] + D \sum_{i=1}^n (\sigma_i^- \rho \sigma_i^+ + \sigma_i^+ \rho \sigma_i^- - \rho) \quad (3.31)$$

where $H = H_I + H_c$ and dissipation $\sigma_i^\pm$ of equal strength $D > 0$ acting on individual spins. The control cost is defined as

$$C = \left\langle -\frac{Q}{2} \mathcal{F}(\psi_T) + \frac{R}{2} \int_0^T \left(\sum_{i=1}^n \sum_{a=x,y} u_{ia}^2(t) \right) dt \right\rangle \quad (3.32)$$

Here $R > 0$ and satisfies $R = \lambda D^{-1}$ with $\lambda > 0$. As in the case of one qubit, we use the transformation (3.28) to convert the non-Hermitian operators $C'_{ia} = \sigma_i^+$, σ_i^- in Eq. (3.31) into anti-Hermitian operators

$$C_{ib} = C'_{ia} A_{ab} \quad \text{for all } i \quad (3.33)$$

This leads to the following expression for the transformed operators

$$C_{ib} = -i\sigma_i^b \quad \text{for all } i \text{ and } b = x, y. \quad (3.34)$$

The corresponding unraveling is

$$d\psi = -iH_I\psi dt - i \left(\sum_{i=1}^n \sum_{a=x,y} (u_{ia}dt + dW_{ia})\sigma_i^a \right) \psi - nD\psi dt \quad (3.35)$$

with $dW_{ia}dW_{jb} = \delta_{ij}\delta_{ab}Ddt$ for $i, j = 1, \dots, n$ and $a, b = x, y$. Equations (3.32) and (3.35) define a path integral control problem. The initial state is set to $\psi_0 = |0\rangle^n = Z^n$. We choose the target state ϕ to be a Green-

berger-Horne-Zeilinger (GHZ) state defined by $|GHZ_n\rangle := \frac{1}{\sqrt{2}}(|0\rangle^n + |1\rangle^n)$, which represents a maximally entangled state in the global entanglement (or Meyer-Wallach) measure for multipartite systems [184]. In the rotating frame, the target state takes the form $\phi = \frac{1}{\sqrt{2}}(e^{i2\pi\omega T}|0\rangle^n + e^{-i2\pi\omega T}|1\rangle^n)$ with $\omega = \frac{1}{2}\sum_j D_j$ (see Supplementary Material).

We consider the 4-qubit NMR example of [185]. We add in the Supplementary Material a toy example with $n = 10$ qubits to showcase the applicability of PiQC to handle larger systems. We take the shifts D_i and couplings J_{ij} from Table I therein. In [185] authors use an iterative version of GRAPE, called i-GRAPE, to optimize the pulses. GRAPE-like methods need, in general, a high number of pulses to validate the approximation of the optimization step and reach high fidelities. As a result, those solutions tend to be very irregular. In addition, the variation is high between solutions for different runs, and the reported solutions for GRAPE and i-GRAPE are very different between each other. PiQC, on the other hand, is not constrained by the number of pulses, which can be chosen arbitrarily. Whether the solutions achieve high fidelities or not depends on other factors, such as system controllability or cost parameters, such as R and Q . This difference arises from the distinct approach in dealing with time correlations in PiQC compared to gradient-based methods; we return to this point in Sec. 4.6.

We couple the 4-qubit NMR molecule to an environment. The dynamics corresponds to the Lindblad equation (3.31). Following [185], we choose the final time $T = 8.8$ and set the noise coupling $D = 1 \times 10^{-3}/T_r$, with $T_r = 0.61$ the minimum relaxation time of the system. We compute the infidelity $1 - F_{avg}$ reached after convergence for different values of K . This is shown in red in Fig. 3.8 (A). In the unitary case, authors in [185] report solutions with infidelity 2×10^{-3} and $K = 1760$ GRAPE pulses. In the noisy case (non-zero D), we reach infidelities less than 1×10^{-2} with just $K = 5$ pulses. The same solutions achieve an infidelity less than 8×10^{-4} for the unitary case ($D = 0$). This showcases the advantage of PiQC in requiring less pulses to achieve high fidelity solutions. In Fig. 3.8 (B-C) we show the solution profiles corresponding to $K = 5$ and $K = 500$ pulses, respectively. In the $K = 500$ case, the controls are iteratively smoothed during the importance




Figure 3.8 4-NMR control. **(A)** Infidelity $1 - F$ vs K . Curve open corresponds to the open system case, with fixed noise coupling D . Curves anneal1, anneal2 and anneal3 correspond are the infidelities using annealing control. The rest of parameters are: $N_T = 500$, $n_{IS} = 1000$. We show the control solutions for the open system case for $K = 5$ pulses **(B)** and $K = 500$ pulses **(C)**. In panel **(D)** we show an example of annealing training for $K = 500$ pulses and logarithmic end cost. We plot the average and the worst case infidelities as function of the IS step, and in panel **(E)** we plot the corresponding final solutions. These solutions achieve an infidelity of 6×10^{-11} . The annealing training is performed using an averaging window $w = 10$ in the first 200 IS steps, then setting $w = 1$ for the rest of the training.

sampling using cubic splines. For large K , this enforces smooth solutions without sacrificing accuracy, which can be preferred in analog settings such as in neutral-atom computers or quantum annealers.

3.4.4 Application of annealed diffusion control

In this Section we compare PiQC and GRAPE solutions in the unitary regime. For this, we use the annealing variant of PiQC presented in Section 3.3.4 to compute control solutions for unitary dynamics. We apply this procedure to the 4-NMR case. We parametrize the noise value D in (3.35) in the importance sampling step p , interpolating between two noise values: a high initial noise D_0 , and a much lower noise $D_f \ll 1$. The form of the annealing schedule $D(p)$ is an ansatz that we chose empirically. We find that good schedules allow the algorithm to sample the control landscape more effectively in early stages

of the optimization. We choose a piece-wise constant schedule composed of D_i ($i = 1, \dots, N_{steps}$) values distributed equidistantly in the IS interval. In the i -th IS interval, the noise value $D_i := D_0 \left(\frac{D_f}{D_0} \right)^{(i-1)/(N_{steps}-1)}$. We set $N_{steps} = 20$, $D_0 = 3 \times 10^{-3}$ and $D_f = 1 \times 10^{-12}$; see dashed line in Fig. 3.8 (D).

We compute infidelities $1 - F$ vs K , for different values of Q . In Fig. 3.8 (A) we show the infidelity curves for two values of Q , with fixed $R = 1$. By increasing Q , we expect to obtain lower infidelities, and this is shown in panel (A), where curve `annea11` corresponds to $Q = 400$, and `annea12` to $Q = 4000$. We readily observe that PiQC solutions for these two cases produce an improvement in infidelity of 2 and 4 orders of magnitude, respectively, with respect to GRAPE [185]. We can push further this advantage by changing the end cost in (3.32) to a log-infidelity defined by $\Phi(\psi_T) = \frac{Q}{2} \log(1 - \mathcal{F}(\psi_T))$. The log-infidelity is more sensitive to fine tuning when the fidelity is close to one, forcing the algorithm to keep improving the controls even when high fidelities are reached. Applying the log-infidelity cost ($K = 500$ and $Q = 400$) we observe an improvement of 7 orders of magnitude with respect to GRAPE. This is shown by the infidelity curve `annea13` in panel Fig. 3.8 (A). For this last case, we show in panel (D) the infidelity training curves (solid red and black) and the corresponding smooth solutions in panel (E). As for the open case, the smoothing was performed during training, using cubic splines. These results demonstrate the effectiveness of the annealing variant of PiQC to find high quality solutions for closed systems, outperforming by several orders of magnitude those found by GRAPE for the same setting.

3.5 Discussions and Outlook

In this paper we have introduced Path integral Quantum Diffusion Control (PiQC), a novel method for computing optimal control solutions for open and closed quantum systems. The approach uses a stochastic optimal control formulation based on unravelings of the open quantum dynamics. Unlike traditional approaches that rely on solving the Hamilton-Jacobi-Bellman equation or adjoint methods, PiQC reformulates coherent quantum control

of open systems as a stochastic optimal control problem with solutions expressed directly as path integrals. This enables efficient approximation of optimal controls by sampling trajectories in Hilbert space, bypassing otherwise prohibitive analytical and/or computational calculations. Through adaptive importance sampling, PiQC transforms an analytically intractable problem into a highly parallelizable, gradient-free optimization scheme.

The main feature of PiQC is that it uses an optimization approach that is very different from standard gradient-based optimization approaches. First of all, PiQC uses adaptive importance sampling, which can be regarded as a *gradient-free* method, since the control updates are computed directly by sampling the cost objective, with no need of gradient computations. Secondly, for any control problem, the primary computational challenge lies in quantifying the influence of early control actions on the system's state and cost at later times. For instance, GRAPE and PMP use some form of backward evolution of the co-state to estimate this dependence. PiQC, instead, uses an entirely different mechanism, by generating forward-sampled stochastic-particle trajectories that are weighted with their exponentiated control cost. These *global statistics* are then used to update the parameters in each time bin independently. This is a key feature of the path integral control approach, and our experiments show its effectiveness in computing optimal controls. Another feature of PiQC is that it is a *state-centered* approach: the quantum trajectories are obtained by propagating state vectors of size N , opposed to density-based approaches that require propagating density matrices of size N^2 ; in this way, the PiQC algorithm inherently comprises a quadratic advantage in algorithmic runtime compared to alternative approaches such as those based on the PMP or adjoint method [137]. Our method is *embarrassingly parallel*: the importance sampling updates Eqs. (3.11) and (3.13) can be efficiently parallelized by noticing that the gradient consists of a sum over the data and that the computation of this sum can be distributed over m machines so that the updated parameters $\theta_{p+1} = \theta_p + \frac{1}{m} \sum_{i=1}^m d\theta_p^{(i)}$ with $d\theta_p^{(i)}$ the gradient update computed on machine i . No large volumes of simulated raw data need to be transferred between machines, only the parameter updates. See Fig. 3.9.




Figure 3.9 During each iteration, samples are generated using control functions $u(\theta^{(p)})$. These samples are used to compute new parameters $\theta^{(p+1)}$. These two steps can be distributed on many machines with minimal communication overhead, since no large volumes of simulated raw data need to be transferred between machines, only the parameter updates.

We remark that for linear control parametrizations (an example is the pulse-based control scheme treated in this work) our method provides closed-form expressions for control updates computed as stochastic averages over quantum trajectories. For more general non-linear parametrizations –e.g., when the control model $u(t; \theta)$ is represented by a neural network– the Path Integral Cross-Entropy (PICE) method can be employed [22].

The main limitation of our algorithm rests in the anti-Hermitian constraint that requires the dissipators C_a to live in a space spanned by anti-Hermitian operators (see Sec. 3.3.2). This is a direct consequence of using a linear dynamics for the unraveling (see Eqs. (3.19) and (3.36)). However, there are many real world scenarios where this constraint is satisfied. For example, PiQC can be applied for controlling dissipation in atom-light interactions [38], relaxation and detuning in modern quantum computing platforms based on, e.g., superconducting qubits, ion-traps or neutral atoms [180–182]. For instance, see [186] for different examples of noise that affect neutral atom platforms which satisfy the anti-Hermitian constraint. The same framework can also be applied for optimal error mitigation of coherent noise [187–189] in quantum computing platforms based on superconducting qubits, trapped-ion or neutral atoms, operating at finite temperature [180]. It can also be used for error mitigation of Pauli errors, such as dephasing and depolarizing noises

and, therefore, for the design of more resilient compilation protocols for quantum algorithms [188]. The anti-Hermitian constraint could be potentially avoided by resorting to more advanced techniques such as in [190], which enables the exploration of more general forms of unravelings that do not necessarily follow the PI form. We leave these questions for future work.

We further introduce an annealing variant of PiQC that leverages synthetic noise to compute high-quality solutions for closed quantum systems. Our numerical validation on NMR state preparation demonstrates that noise-assisted control can deliver superior solutions compared to established methods such as GRAPE, with improvements spanning several orders of magnitude in fidelity. We remark that in the annealing variant of PiQC, the anti-Hermitian constraint becomes irrelevant, as the environment coupling turns into an ansatz for computing optimal controls in absence of dissipation. In this sense, the annealing variant of PiQC can be viewed as a noise-assisted optimal quantum control algorithm.

PiQC naturally accommodates feedback control solutions, but this requires full observability of the quantum state. Extending PiQC to feedback control requires incorporating measurements in the optimization process, which requires continuous measurement theory [191]. This extension is currently under development and will be presented in future work.

3.6 Acknowledgements

We thank Christiane Koch, Luis Pedro Garcia-Pintos, Antonio Sannia, Martin Larocca, Peter Komar for helpful discussions and advice, and to the anonymous reviewers for fruitful comments on an earlier version of this manuscript.

3.7 Data and code availability

The data supporting the findings of this study can be fully reproduced by following the methodology described in the article. In addition, the code

necessary for replicating the PiQC results is available in a GitHub public repository [174].

3.8 Appendices

3.8.1 Unravelings of the controlled Lindblad equation

In the theory of stochastic quantum unravelings [149, 192–197], the controlled Lindblad equation (3.16) can be derived as the average dynamics over all particular time realizations of a stochastic quantum state ψ_t that follows the SSE

$$d\psi_t = -iH_0\psi_t dt - \frac{1}{2}D_{ab}C_b^\dagger C_a\psi_t dt + C_a\psi_t(u_a(t)dt + dW_t^a) \quad (3.36)$$

with Itô structure $\langle dW_t dW_t^T \rangle = Ddt$, with D a positive symmetric matrix, and u an open-loop control, and $C_a = -H_a$. The proof is straightforward using Itô calculus, and the reader may consult [37, 150, 166, 191] for derivations corresponding to diagonal D in the absence of control. We give a proof for general non-diagonal D' and $u \neq 0$.

Define the stochastic density operator $P_t := \psi_t\psi_t^\dagger$. Compute the dynamics of P_t as

$$\begin{aligned} dP_t &= d\psi_t\psi_t^\dagger + \psi_t d\psi_t^\dagger + d\psi_t d\psi_t^\dagger \\ &= -i[H_0, P_t]dt + \mathcal{D}[D, \mathcal{C}](P_t)dt + [C_a P_t(u_a(t)dt + dW_t^a) + \text{h.c.}] \end{aligned} \quad (3.37)$$

On the other hand, it is easy to show that the change in the norm $\|\psi_t\|^2 = \psi_t^\dagger\psi_t$ is given by

$$d\|\psi_t\|^2 = \psi_t^\dagger(C_a + C_a^\dagger)\psi_t(u_a(t)dt + dW_t^a)$$

The average dynamics should correspond to a completely-positive trace-preserving (CPT) map, a general physical requirement for quantum chan-

nels [40]. In the stochastic setup this means that the norm of the state ψ_t must preserve on average, i.e. $\langle d\|\psi_t\|^2 \rangle = 0$. This applies for all u iff C_a are anti-Hermitian, i.e. $C_a = -iH_a$ for some Hermitian H_a . Note that, under this condition, the norm of ψ_t conserves deterministically, and all the quantum trajectories remain pure. After inserting this into the average dynamics (3.37), taking the expectation over all noise realization, and setting $\rho(t) = \langle P_t \rangle$, we obtain the lindblad equation

$$\dot{\rho} = -i[H_0 + u_a H_a, \rho] + \mathcal{D}[D, \mathcal{C}](\rho)$$

Do to invariance of the dissipator $\mathcal{D}[D, \mathcal{C}] = \mathcal{D}[D', \mathcal{C}']$ under transformation (3.22), we conclude that (3.36) with $C_a = -iH_a$ unravels the class of Lindblad equations with D' and $\mathcal{C}'$ generated by transformation (3.22).

3.8.2 Path integral control theory

In this section we describe some basic results regarding PI control theory. For further details the reader may consult the original papers [20, 198].

Define the optimal cost-to-go as the minimal expected cost starting in any x at time t with $t_0 \leq t \leq T$

$$J(t, x) = \min_{u \in \mathcal{U}_t} \langle S^u(t) \rangle_{X_t=x} \quad u^*(t, x) = \operatorname{argmin}_{u \in \mathcal{U}_t} \langle S^u(t) \rangle_{X_t=x} \quad (3.38)$$

where $\mathcal{U}_t$ is the restriction of the space of valid feedback controls $\mathcal{U}$ to the time domain $[t, T]$.

Define

$$\psi(t, x) = e^{-J(t, x)/\lambda} \quad \phi(t) = e^{(S^u(t) - S^u(t_0))/\lambda} \quad (3.39)$$

where $\lambda I = RD$. Furthermore we define the stochastic processes $\psi(t) = \psi(t, X_t)$, $u(t) = u(t, X_t)$ and $u^*(t) = u^*(t, X_t)$

Lemma 10. For all t in $t_0 \leq t \leq T$

$$e^{-S^u(t)/\lambda} - \psi(t) = \frac{1}{\lambda\phi(t)} \int_t^T \phi(s)\psi(s)(u^*(s) - u(s))^T R dW_s \quad (3.40)$$

Proof. The HJB equation for the control problem Eqs. (3.1) and (3.2) in the paper is

$$-\partial_t J = \min_u \left(V + \frac{1}{2} u^T R u + (f + gu)^T \partial_x J + \frac{1}{2} \text{Tr}(g D g^T \partial_{xx} J) \right) \quad (3.41)$$

with the boundary condition $J(T, x) = \Phi(x)$. We can solve for u , which gives $u^* = -R^{-1} g^T \partial_x J$ and

$$-\partial_t J = V - \frac{1}{2} \partial_x J^T g R^{-1} g^T \partial_x J + f^T \partial_x J + \frac{1}{2} \text{Tr}(g D g^T \partial_{xx} J) \quad (3.42)$$

We assume that the matrices R, D are related such that $R = \lambda D^{-1}$ with $\lambda > 0$. In terms of ψ , Eq. (3.42) becomes linear

$$\partial_t \psi + f^T \partial_x \psi + \frac{1}{2} \text{Tr}(g D g^T \partial_{xx} \psi) = \frac{V}{\lambda} \psi \quad (3.43)$$

Alternatively, we consider the stochastic process $\psi(t) = \psi(t, X_t)$. Using Itô calculus we obtain

$$d\psi = \left(\partial_t \psi + \partial_x \psi^T (f + gu) + \frac{1}{2} \text{Tr}(g D g^T \partial_{xx} \psi) \right) dt + \partial_x \psi^T g dW \quad (3.44)$$

$$= \frac{V}{\lambda} \psi + \partial_x \psi^T g (u dt + dW) \quad (3.45)$$

where we used Eq. (3.43), with boundary condition $\psi(T, x) = e^{-\Phi(x)/\lambda}$. Using the definition of $\phi(t)$ we obtain

$$d\phi = -\frac{1}{\lambda} (V dt + u^T R dW) \quad (3.46)$$

with initial condition $\phi(t_0) = 1$. Using the Itô product rule we obtain

$$d(\phi\psi) = d\phi\psi + \phi d\psi + d\phi d\psi = -\frac{1}{\lambda}\psi\phi u^T RdW + \phi\partial_x\psi^T g dW \quad (3.47)$$

$$= \frac{1}{\lambda}\phi\psi(u^* - u)^T RdW \quad (3.48)$$

where in the last step we used $u^* = \lambda R^{-1}g^T\partial_x\psi/\psi$. Integrating from t to T we obtain

$$\phi(T)\psi(T) - \phi(t)\psi(t) = \frac{1}{\lambda} \int_t^T \phi(s)\psi(s)(u^*(s) - u(s))^T RdW_s \quad (3.49)$$

We use $\phi(T) = \phi(t)e^{(S^u(T) - S^u(t))/\lambda}$ and divide by $\phi(t)$ to obtain Eq. (3.40). $\square$

Corollary 1. *We take the expectation of Eq. (3.40) with respect to $dW_{t:T}$ and condition on events up to time t , commonly called a filtration $\mathcal{F}_t$. We obtain*

$$\psi(t) = \left\langle e^{-S^u(t)/\lambda} \right\rangle_{\mathcal{F}_t} \quad (3.50)$$

The following path integral control theorem is useful to estimate these parameters for all times from one set of samples $X_{t_0:T}$.

Theorem 2. *Let $f : \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}$ and define the process $f_t = f(t, X_t)$. Then*

$$\left\langle e^{-S^u(t_0)/\lambda} \int_{t_0}^T (u^*(s) - u(s))f(s)ds \right\rangle = \left\langle e^{-S^u(t_0)/\lambda} \int_{t_0}^T f(s)dW_s \right\rangle \quad (3.51)$$

where the expectation is with respect to the stochastic process (3.1) in the paper.

Proof. Consider Lemma lemma 10 for $t = t_0$. Multiply both sides by $\int_{t_0}^T f(s)dW_s$ and take the expectation value

$$\begin{aligned} \left\langle e^{-S^u(t_0)} \int_{t_0}^T f(s)dW_s \right\rangle &= \\ \left\langle \int_{t_0}^T f(s)dW_s \int_{t_0}^T \phi(s')\psi(s')(u^*(s') - u(s'))^T \frac{R}{\lambda} dW_{s'} \right\rangle &\quad (3.52) \end{aligned}$$

Using $R = \lambda D^{-1}$ and Itô isometry we have

$$\left\langle e^{-S^u(t_0)} \int_{t_0}^T f(s) dW_s \right\rangle = \int_{t_0}^T \langle f(s) \phi(s) \psi(s) (u^*(s) - u(s)) \rangle ds \quad (3.53)$$

Using the definition of $\psi(s)$ and $\phi(s)$ we write

$$\begin{aligned} \left\langle e^{-S^u(t_0)} \int_{t_0}^T f(s) dW_s \right\rangle = \\ \int_{t_0}^T \left\langle f(s) e^{(S^u(s) - S^u(t_0))/\lambda} \left\langle e^{-S^u(t_0)} \right\rangle_{\mathcal{F}_s} (u^*(s) - u(s)) \right\rangle ds \end{aligned} \quad (3.54)$$

Finally, by applying the law of total expectation we obtain (3.51). □

Corollary 2. *By setting $f(t) = \delta_{t \in I}$ with $I = [t_0, t_0 + dt]$, diving by dt and taking the limit $dt \rightarrow 0$ we establish the proof of (3.5) in the paper.*

3.8.3 Open GRAPE

Here we summarize the details of our implementation of Open GRAPE for the benchmarking part of Section 3.4.1 in the paper. For further details concerning the definition of Open GRAPE algorithm we refer the reader to [135]. Given the cost objective (3.17) and the control model Eq. (3.8) in the paper, with R a positive scalar, the first order Open GRAPE update rule (ignoring terms of $\mathcal{O}(dt^2)$ and higher) becomes

$$A_{ak} \rightarrow A'_{ak} = A_{ak} - \varepsilon \frac{\partial C}{\partial A_{ak}}$$

with

$$\frac{\partial C}{\partial A_{ak}} = \frac{iQ}{2} \text{Tr}(\lambda_k [H_a, \rho_k]) dt + R A_{ak} dt$$

where ρ_k is the forward-propagated state from ρ_0 to time t_k and λ_k is the backward-propagated state from ρ_{target} to time t_k , ε is the learning rate, and

dt is the time step. We implement the algorithm in Python language. For the dynamics simulation of the backward and forward propagations we use the open-source Python package `dynamics` [199].

The learning rate ε usually needs to be carefully chosen in order to avoid overshooting. In the case of GRAPE for closed systems, work [200] reported an adaptive learning rate at each j step of ε/j^α with $\alpha = 1/2$ was enough to avoid overshooting of local minima/saddle points. We found this approach not suitable for our benchmarking purposes as it required an excessive fine tuning of α for different types of seeds. We choose instead the following protocol: every time the cost $C(A') \geq C(A)$, where A the current solution and A' the new solution, the learning rate is reduced by a factor 10 and $C(A')$ is recomputed until either $C(A') < C(A)$ or ε is reduced a maximum number of times that we set to 10. This stopping condition allows us to capture local minima and detect when the first order approximation of the gradient update starts to fail. At the start of each run, we sample random pulse seeds u from Gaussian distributions (centered at zero, 1, and -1, with standard deviation 2) and uniform distributions (integers between -10 and 10, with a scale factor of 1 and 0.5).

3.8.4 PiQC robustness example

We show an experiment for the single qubit that tests PiQC robustness and algorithmic resilience under imperfect knowledge of the system model. We consider the single qubit system of Section 3.4.1 in the presence of dissipation, with Hamiltonian $H = H_0 + H_c(t)$, where the *true* drift Hamiltonian model is given by $H_0 = 0$ and $H_c(t)$ is the control Hamiltonian with transverse controls in x and y . We test the algorithm response under random perturbations of the drift Hamiltonian which encode our ignorance about the exact form of H_0 . For this, we consider random perturbations in time $\delta H_0(t)$ such that $H = \delta H_0(t) + H_c(t)$ such that the components of $\delta H_0(t)$ are sampled from a Gaussian distribution with a given standard deviation σ , and then taking the Hermitian part of the resulting matrix. In consequence, the

quantum trajectories will depend on the individual realizations of the Wiener increments dW_t as well as $H_0(t)$. The control problem consists of the preparation of the state $|Y\rangle$ starting from $|X\rangle$ with dissipation strength $D = 0.005$. After convergence, we use the control solutions in the true dynamics with $H_0 = 0$ and compute the average fidelity $F(\sigma)$ across trajectories. We repeat this multiple times to extract some statistics and for a number of σ values. In parallel, we repeat the same exercise with Open GRAPE. We introduce in the dynamics a time varying drift $\delta H_0(t)$ drawn from a Gaussian distribution with standard deviation σ as before, and use the control solutions for the ground truth dynamics with $H_0 = 0$. The results are shown in Fig. 3.10 and demonstrate PiQC's enhanced robustness under Hamiltonian uncertainty, showing mean infidelity values one order of magnitude lower than Open GRAPE across the tested perturbation amplitudes.




Figure 3.10 Infidelity $1 - F$ vs perturbation amplitude σ . The parameters for PiQC are: $N_{\text{traj}} = 400$, $n_{\text{IS}} = 500$, $R = 1$, $Q = 100$, $T = 1$, $K = 50$ and 100 time steps. The parameters for Open GRAPE are the same as PiQC to match the control problem definition, i.e. R, Q, T, K . For each σ we solve the control problem for each algorithm 10 times and compute the mean and variance of the infidelities. In the case of Open GRAPE the solutions are chosen based on best performance criteria: for each σ we run the algorithm for 10 different control seeds and choose the one with highest fidelity. The control seeds are normally distributed with standard deviation 5. The learning rate adapts according to Sec. 3.8.3 with adaptive factor 0.5. The results demonstrate that PiQC is more robust than Open GRAPE, achieving one order of magnitude lower mean infidelity.

3.8.5 Open NMR control theory

Rotating frame transformation

Define the unitary transformation

$$U = \exp(iH_Z t) \quad H_Z = \pi \sum_j D_j \sigma_j^z \quad (3.55)$$

It represents a transformation from the lab frame to a multiple rotating frame. Given the state ρ in the lab frame, define the transformed state to the rotating frame by $\rho' = U\rho U^\dagger$. Assume ρ follows the generic Lindblad equation

$$\dot{\rho} = -i[H_Z + H_I + H_c, \rho] + \mathcal{D}[D, \mathcal{C}]\rho \quad (3.56)$$

where, for simplicity, D is a real scalar, H_I and H_c are defined as in (3.30), and the $\mathcal{C} = \{C_a\}_{a=1}^{n_c}$ are arbitrary dissipators.

Lemma 11. *The transformed state ρ' follows the Lindblad equation*

$$\dot{\rho}' = -i[H_I + H'_c, \rho'] + \mathcal{D}[D, \mathcal{C}']\rho' \quad (3.57)$$

where

$$H'_c = \sum_i u'_{ix} \sigma_i^x + u'_{iy} \sigma_i^y \quad \mathcal{C}'_a = U C_a U^\dagger \quad (3.58)$$

and the controls in the rotating frame, $u'_i = (u'_{ix}, u'_{iy})^T$, are related to those in the lab frame, $u_i = (u_{ix}, u_{iy})^T$, by a real rotation

$$u'_i = S_i(t) u_i \quad S_i(t) = \begin{pmatrix} \cos(2\pi D_i t) & \sin(2\pi D_i t) \\ -\sin(2\pi D_i t) & \cos(2\pi D_i t) \end{pmatrix} \quad (3.59)$$

Proof. By differentiating ρ' w.r.t. time we obtain

$$\begin{aligned}\dot{\rho}' &= \dot{U}\rho U^\dagger + U\dot{\rho}U^\dagger + U\rho\dot{U}^\dagger \\ &= -i[H_I + H'_c, \rho'] + \mathcal{D}[D, \mathcal{C}']\rho'\end{aligned}$$

where $H'_c = UH_cU^\dagger$ and $\mathcal{C}'_a = UC_aU^\dagger$.

Next, compute

$$H'_c = UH_cU^\dagger = \sum_i u_{ix}(U\sigma_i^xU^\dagger + u_{iy}U\sigma_i^yU^\dagger)$$

For each spin i we obtain

$$\begin{aligned}U\sigma_i^xU^\dagger &= e^{i\pi D_it\sigma_i^z}\sigma_i^xe^{-i\pi D_it\sigma_i^z} \\ &= \cos(2\pi D_it)\sigma_i^x - \sin(2\pi D_it)\sigma_i^y\end{aligned}$$

and

$$\begin{aligned}U\sigma_i^yU^\dagger &= e^{i\pi D_it\sigma_i^z}\sigma_i^ye^{-i\pi D_it\sigma_i^z} \\ &= \cos(2\pi D_it)\sigma_i^y + \sin(2\pi D_it)\sigma_i^x\end{aligned}$$

Then, we can write H'_c as

$$H'_c = \sum_i u'_{ix}\sigma_i^x + u'_{iy}\sigma_i^y$$

where

$$\begin{aligned}u'_{ix} &= \cos(2\pi D_it)u_{ix} + \sin(2\pi D_it)u_{iy} \\ u'_{iy} &= -\sin(2\pi D_it)u_{ix} + \cos(2\pi D_it)u_{iy}\end{aligned}\tag{3.60}$$

By defining

$$S_i(t) = \begin{pmatrix} \cos(2\pi D_i t) & \sin(2\pi D_i t) \\ -\sin(2\pi D_i t) & \cos(2\pi D_i t) \end{pmatrix}, \quad (3.61)$$

$u'_i = (u'_{ix}, u'_{iy})^T$ and $u_i = (u_{ix}, u_{iy})^T$, we rewrite (3.60) as $u'_i = S_i(t)u_i$ and the Lemma is proved. $\square$

Assume the NMR system is in interaction with a EM field bath, which we encode in the dissipators $C_a = \sigma_i^+, \sigma_i^- (i = 1, \dots, n)$.

Lemma 12. *The dissipator in (3.30) is invariant under the action of U , i.e. $\mathcal{D}[D, C'] = \mathcal{D}[D, C]$.*

Proof. Apply the unitary transformation U to the Lindblad operators.

$$\begin{aligned} U\sigma_i^+U^\dagger &= e^{i\pi D_i t \sigma_i^z} \sigma_i^+ e^{-i\pi D_i t \sigma_i^z} \\ &= e^{i2\pi D_i t} \sigma_i^+ \end{aligned}$$

where we used that $[\sigma_i^z, \sigma_i^+] = 2\sigma_i^+$. Analogously for σ_i^- , we have $U\sigma_i^-U^\dagger = e^{-i2\pi D_i t} \sigma_i^-$. Therefore, the Lindblad operators C'_a in the rotating frame are the same as in the lab frame up to phase factors. As the dissipator is invariant under phase transformations, the Lemma is thus proven. $\square$

Control problem equivalence

The main purpose of the unitary transformation U is to factor out potentially fast frequencies D from the definition of the control problem that are not necessary for computing the optimal control and may even hinder the efficient search for a solution. By defining and solving the control problem in the rotating frame, we can uniquely map the optimal solution back to the lab frame through the action of the unitary transformation in control space. In this Section we describe a general invariance present in the control problem under certain family of transformations $\mathcal{T}$.

Assume the Lindblad equation for an open NMR system

$$\dot{\rho} = -i[H, \rho] + \mathcal{D}[D, \mathcal{C}]\rho \quad (3.62)$$

with $H = H_Z + H_I + H_c$ and dissipator

$$\mathcal{D}[D, \mathcal{C}]\rho = \sum_j D_{jab} \left(C_{ja}\rho C_{jb}^\dagger - \frac{1}{2}\{C_{jb}^\dagger C_{ja}, \rho\} \right).$$

where index j labels qubit j , and the dissipation operators C_{ja} act on qubit j . The noise matrix D is in block diagonal form $D := \bigoplus_j D_j$ with $(D_j)_{ab} = D_{jab}$.

Define the cost

$$C[x] = -\frac{Q}{2}\text{Tr}(\rho_T \rho_{tar}) + \frac{1}{2} \sum_j \int_0^T u_j^T R_j u_j dt \quad (3.63)$$

where the tuple $x := (\rho_T, \rho_{tar}, u, R)$ with $R := \bigoplus_j R_j$ block diagonal.

Consider U as in (3.55), and define the transformation $\mathcal{T} = (U, S)$ to the rotating frame by it's action on pair (ρ, u) , with $u = u_i (i = 1, \dots, n)$, as $\mathcal{T}(\rho, u) = (U\rho U^\dagger, Su)$, and it's action on pair (D, R) as $\mathcal{T}(D, R) = (S^T D S, S^T R S)$.

Under $\mathcal{T}$, the dissipators C_{ja} transform as

$$UC_{ja}U^\dagger = e^{i\alpha_j} \sum_b S_{jab} C_{jb} \quad (3.64)$$

where $S_j (j = 1, \dots, n)$ is an orthogonal matrix with $S_j S_j^T = 1$, α_j arbitrary real numbers, and define $S = \bigoplus_j S_j$.

Lemma 13. *The dissipator $\mathcal{D}[D, \mathcal{C}]$ is invariant under $\mathcal{T}$, i.e. $\mathcal{D}[D', \mathcal{C}'] = \mathcal{D}[D, \mathcal{C}]$.*

Proof. The results follows by applying $\mathcal{T}D = SDS^T$ and definition (3.64) to the dissipators C_{ja} . $\square$

Theorem 3. *The cost $C[x]$ in (3.63) is invariant under $\mathcal{T}$ and the optimal control solutions u'^* and u^* in the rotating and lab frames, respectively, are related by $u'^* = Su^*$.*


Proof. Write the action of $\mathcal{T}$ on C as

$$\begin{aligned}
(\mathcal{T}C)[x] &= C[\mathcal{T}x] \\
&= -\frac{Q}{2}\text{Tr}(\rho'_T \rho'_{tar}) + \frac{1}{2} \sum_j \int_0^T u'_j{}^T R'_j u'_j dt \\
&= -\frac{Q}{2}\text{Tr}(U \rho_T U^\dagger U \rho_{tar} U^\dagger) + \frac{1}{2} \sum_j \int_0^T (Su_j)^T S R_j S^T (Su_j) dt \\
&= C[x]
\end{aligned}$$

where in the last line we used the unitarity/orthogonality of U and S_i , respectively. Thus, $C[x] = C[x']$ and the invariance of C follows.

Next, fix ρ_{tar} and R , and consider the optimal solution in the lab frame $u^* = \text{argmin}_u C[u]$, which satisfies $C[u^*] \leq C[u] \forall u$. By invariance of C under $\mathcal{T}$, and given that $\mathcal{T}$ is invertible, the transformation Su^* satisfies $C[Su^*] \leq C[u'] \forall u'$ where $u' = Su$. Thus, the optimal solution in the rotating frame u'^* corresponds to a rotation of the optimal solution in the lab frame, $u'^* = Su^*$. $\square$

Provided that the dissipators C' in the rotating frame can be transformed to anti-Hermitian operators, we can solve the control problem in the rotating frame using the PiQC algorithm. By virtue of Theorem 3, The optimal control u'^* in the rotating frame can be mapped back to the lab frame. Define $\mathcal{L}^{rot}$, $\mathcal{L}^{lab}$ as the Lindbladian operators in the rotating and lab frames, respectively. The following diagram illustrates the end-to-end workflow of the transformation.



Example 1. A particular example of $\mathcal{T}$ is the case in which the dissipators

$C_a = \sigma_i^+, \sigma_i^-$. Then, S corresponds to the rotation defined in (3.59) and we can use the transformation A defined in (3.28) to transform to anti-hermitian operators and simulate the problem using PiQC. •

Example 2. When the target state in the lab frame is a GHZ state, i.e. $\phi = |GHZ\rangle$, the corresponding state in the rotating frame becomes $\phi' = U_T |GHZ\rangle = \frac{1}{\sqrt{2}}(e^{i2\pi\omega T} |0\rangle^n + e^{-i2\pi\omega T} |1\rangle^n)$ with $\omega = \frac{1}{2} \sum_j D_j$. In particular, note that $\phi' = \phi$ whenever $\omega T \in \mathbb{Z}$. •

3.8.6 A 10 qubits example

Here we provide a complimentary example applying PiQC to a 10-qubit toy model example to test the scalability of our algorithm.

Consider a 1-D spin chain of n qubits with single and pair-wise controls. Since this is a toy model designed to test the scalability of our algorithm for a larger number of qubits, we do not assert its feasibility in a lab setting, but rather emphasize the simplicity of defining such control problems within our framework. For simplicity, we set $H_0 = 0$. The control Hamiltonian is

$$H_c = \sum_{i=1}^n \sum_{a=x,y} u_{ia} \sigma_i^a + \sum_{i=1}^{n-1} \sum_{a,b=x,y} u_{iab} \sigma_i^a \sigma_{i+1}^b \quad (3.65)$$

where the controls u_{ia} act on spin i , and controls u_{iab} on the pair $i, i + 1$. Note, that H_c is an extension of typical NMR control Hamiltonians [177] with the addition of pair-wise controls.

The open dynamics is described by the Lindblad equation

$$\begin{aligned} \dot{\rho} = & -i[H, \rho] + D_1 \sum_{i=1}^n (\sigma_i^- \rho \sigma_i^+ + \sigma_i^+ \rho \sigma_i^- - \rho) \\ & + D_2 \sum_{i=1}^{n-1} \sum_{v,w=\pm} (\sigma_i^v \sigma_{i+1}^w \rho \sigma_i^{-v} \sigma_{i+1}^{-w} - \rho) \end{aligned} \quad (3.66)$$

with dissipation $\sigma^\pm$ of equal strength D_1 acting on individual spins and with strength D_2 on neighboring pairs of spins.

The control cost is defined as

$$C = \left\langle -\frac{Q}{2}\mathcal{F}(\psi_T) + \frac{1}{2} \int_0^T \left(R_1 \sum_{i=1}^n \sum_{a=x,y} u_{ia}^2 + R_2 \sum_{i=1}^{n-1} \sum_{a,b=x,y} u_{iab}^2 \right) dt \right\rangle \quad (3.67)$$

Analogous to the previous section, we perform the transformation (3.28) to map non-Hermitian operators into anti-Hermitian. We transform the operators $C'_{ia} = \sigma_i^a, a = \pm$ and $C'_{iab} = \sigma_i^a \sigma_{i+1}^b, a, b = \pm$ in Eq. (3.66) into anti-Hermitian operators

$$C_{ib} = C'_{ia} A_{ab} \quad C_{icd} = C'_{iab} A_{abcd} \quad \text{for all } i \quad (3.68)$$

with $A_{abcd} := iA_{ac}A_{bd}$. This leads to the following expression for the transformed operators

$$C_{ib} = -i\sigma_i^b \quad C_{icd} = -i\sigma_i^c \sigma_{i+1}^d \quad \text{for all } i \text{ and } b, c, d = x, y. \quad (3.69)$$

The corresponding unraveling is

$$\begin{aligned} d\psi = & -i \left(\sum_{i=1}^n \sum_{a=x,y} (u_{ia} dt + d\tilde{W}_{ia}) \sigma_i^a + \sum_{i=1}^{n-1} \sum_{a,b=x,y} (u_{iab} dt + dW_{iab}) \sigma_i^a \sigma_{i+1}^b \right) \psi dt \\ & - (nD_1 + 2(n-1)D_2) \psi dt \end{aligned} \quad (3.70)$$

with $dW_{ia}dW_{jb} = \delta_{ij}\delta_{ab}D_1dt$ for $i, j = 1, \dots, n$ and $a, b = x, y$, and $dW_{iab}dW_{jcd} = \delta_{ij}\delta_{ac}\delta_{bd}D_2dt$ for $i, j = 1, \dots, n-1$ and $a, b, c, d = x, y$. Equations (3.67) and (3.70) define a path integral control problem when

$$\lambda = D_1 R_1 = D_2 R_2 \quad (3.71)$$

The initial state is $\psi_0 = |0\rangle^n$. We choose the target state ϕ to be a Greenberger–Horne–Zeilinger (GHZ) state, which is defined as

$|GHZ_n\rangle := \frac{1}{\sqrt{2}}(|0\rangle^n + |1\rangle^n)$ which represents a maximally entangled state in the global entanglement (or Meyer-Wallach) measure for multipartite systems [184]. We set $R_1 = R_2 = R/n_c$ with $n_c = 2n + 4(n - 1)$ the total number of controls. This normalization defines a fluence cost per unit control. The PI condition (3.71) then implies that $D_2 = D_1 \equiv D$. We take as free parameters R and D_1 . We run the algorithm for $n = 10$ qubits. The results are presented in Fig. 3.11, highlighting the potential of PiQC for controlling large quantum systems.




Figure 3.11 Control of a $n = 10$ qubit 1-D spin chain. The IS training (panel A) has two parts (marked by the vertical dashed line), where in the first part we set the IS window $w = 1$ (no smoothing), and in the second part we change it to $w = 10$. This is to speedup convergence in the first part, and later to improve the statistics. The average fidelity reached $F_{avg} = 0.982 \pm 0.001$. Runtime ~ 6.8 hr on a desktop computer. In panel (B) we illustrate the control solutions for qubits 1, 5, 9 and 10. The x/y components of each control are marked by blue and red colors, respectively. The parameters are $K = 64$, $T = 4$, $R=0.1$, $D = 0.001$, $Q=100$, $n_{IS}=4000$, and $N_{traj}=100$.

It is worth noting that a direct implementation of this problem using Open GRAPE would necessitate the computation and storage of matrix exponentials of dimensions $N^2 \times N^2$ at each step K , which could become prohibitive in terms of memory requirements. This challenge underscores the need for an efficient implementation of the algorithm (see, for instance, [135, 145]) and/or the use of specialized software. In contrast, the PiQC algorithm can run efficiently on a standard desktop computer for this number of qubits.

Chapter 4

Hamiltonian and Double-Bracket Flow Formulations of Quantum Measurements

Aarón Villanueva^{1,2} and Luis Pedro García-Pintos²

¹Faculty of Science, Radboud University, Heyendaalseweg 135, 6525 AJ Nijmegen, The Netherlands

²Theoretical Division (T-4), Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA

Correspondence: aaronv@science.ru.nl

We introduce a framework that unifies quantum measurement dynamics, Hamiltonian dynamics, and double-bracket gradient flows. We do so by providing explicit expressions for stochastic Hamiltonians that produce state dynamics identical to those that happen during continuous quantum measurements. When such dynamical processes are integrated over sufficiently long time intervals, they yield the same results and statistics as during wavefunction collapse. That is, wavefunction collapse can be interpreted as coarse-grained (stochastic) Hamiltonian dynamics. Alternatively, wavefunction collapse can be interpreted as double-bracket gradient flows determined by derivatives of (stochastic) potentials defined in terms of observables with direct physical interpretations. The gradient flows minimize the variance of the monitored observable. Our derivations hold for general monitoring described by non-Hermitian jump processes. We show that such reinterpretations of measurement dynamics facilitate the design of feedback processes. In particular, we introduce feedback processes that yield deterministic double-bracket flow equations that prepare ground states of a target Hamiltonian, and feedback processes for state preparation. We conclude by re-interpreting feedback processes as gradient flows with tilted fixed points.

Citation: Aarón Villanueva and Luis Pedro García-Pintos 2026 Rep. Prog. Phys. 89 067602

DOI: 10.1088/1361-6633/ae6fe9

4.1 Introduction

Traditional interpretations of quantum theory divide dynamics into two distinct classes: Hamiltonian processes and measurement processes [95, 201]. In the former, the state ψ_t of isolated quantum systems evolves unitarily, governed by Schrödinger's equation $d\psi_t/dt = -iH\psi_t$. In such a regime, the system's Hamiltonian H completely determines its dynamics.

Hamiltonian dynamics are in sharp contrast to what occurs during quantum measurements. In the latter, the collapse postulate says that, upon measuring an observable $A = \sum_j a_j \Pi_j$ with eigenvalues a_j and eigenprojectors Π_j , the outcome a_j is obtained with probabilities $\psi_t^\dagger \Pi_j \psi_t$. The collapse postulate also states that, immediately after observing such an outcome, the system's state is $\Pi_j \psi_t / \sqrt{\psi_t^\dagger \Pi_j \psi_t}$ [95, 201]. Open system dynamics due to interactions with an environment, or stochastic master equations aris-

ing from conditioning on measurement outcomes, can be derived from a combination of the two aforementioned regimes.

In this letter, we show that despite their distinct physical origins, measurement processes can be mathematically cast as Hamiltonian processes, which, in turn, can be interpreted as gradient flows. In Section 4.3, we show that arbitrary pure-state dynamics can be expressed in terms of single- and double-bracket Hamiltonian processes. We derive explicit expressions for stochastic Hamiltonians that enforce dynamics equivalent to measurement processes. Such Hamiltonians, which depend on the measured observable, the random measurement outcomes, and the system's state, reveal that measurement dynamics involve a double-bracket form contribution interpretable as gradient flows, as shown in Section 4.4. In Section 4.5, we leverage such reformulation of measurement dynamics to design novel Hamiltonian-based techniques for quantum feedback control that modify the minimum of the gradient flow.

4.2 Dynamics of continuously monitored quantum systems

Continuously monitored quantum systems [40] that are subject to a single noise channel can be described by the following stochastic master equation (SME) in Itô form:

$$d\rho_t = \mathcal{L}(\rho_t)dt + \{C - \langle A \rangle_t, \rho_t\}_c dW_t, \quad (4.1)$$

where ρ_t is the density operator. The right-hand side involves the Lindblad super-operator $\mathcal{L}(\rho_t) := -i[H_0, \rho_t] + \mathcal{D}[C](\rho_t)$, where H_0 is the system's Hamiltonian. The dissipation part is $\mathcal{D}[C](\rho_t) := \gamma (C\rho_t C^\dagger - \frac{1}{2} \{C^\dagger C, \rho_t\})$ for an arbitrary jump operator C with damping rate $\gamma > 0$, which we assume to be time-independent. In the diffusion term, we define $\{X, Y\}_c := XY + Y^\dagger X^\dagger$ for arbitrary operators X, Y . The quantity $\langle A \rangle_t = \text{Tr}(A\rho_t)$ is the quantum expectation value at time t of the Hermitian part $A := (C + C^\dagger)/2$. The quantity

dW_t is a stochastic Wiener increment satisfying $dW_t^2 = \gamma dt$.

The state $\bar{\rho}(t) := E(\rho_t)$ obtained by averaging over all noise realizations [40] satisfies the Lindblad equation [202, 203]:

$$\frac{d\bar{\rho}}{dt} = \mathcal{L}(\bar{\rho}). \quad (4.2)$$

We assume from now on that ρ_t describes a pure quantum state, that is, $\rho_t = \psi_t \psi_t^\dagger$ where ψ_t is a normalized state vector. For pure states, one can split C into Hermitian and anti-Hermitian parts and re-write the diffusion term in (4.1) as $\{C - m_t, \rho_t\}_c = [[A, \rho_t] + B, \rho_t]$, where $B := (C - C^\dagger)/2$. By defining the state-dependent Hamiltonian $H_t := i([A, \rho_t] + B)$, Eq. (4.1) can be equivalently written as

$$d\rho_t = \mathcal{L}(\rho_t)dt - i[H_t, \rho_t]dW_t. \quad (4.3)$$

The form of (4.3) is somewhat unconventional, as the noise term appears as a Liouville contribution with Hamiltonian H_t . This reformulation is a direct consequence of the fact that the state remains pure for all times. See App. 4.7.1 for a derivation.

The continuous measurement process associated to Eqs. (4.1) and (4.3) is characterized by the output

$$dy_t = 2\gamma \langle A \rangle_t dt + dW_t, \quad (4.4)$$

where y_t is the accumulated measurement outcome [40]. Since dW_t is of zero mean, upon averaging measurement instances over noise processes on a fixed state ρ_t , the noise-averaged measurement output uncovers the observable's expected value, $\overline{y_t} \propto \langle A \rangle_t$.

Consider the pure-measurement case where $B = 0$ and $H_0 = 0$. The monitoring process is then characterized by the Hermitian observable A and the rate γ . We refer to this as the pure-measurement regime as, in such a case, subspaces corresponding to the eigen-projectors Π_j of A are fixed points of the dynamics. Moreover, the system's state reaches such eigenspaces with probabilities given by the Born rule. The rate at which the full measurement

process happens is determined by the rate γ which is, in turn, dictated by interaction with the measurement apparatus. See, e.g., Refs. [204–206] for experimental demonstrations of continuous measurement processes as described above. Next, we relate the stochastic measurement dynamics described above to purely Hamiltonian processes.

4.3 Hamiltonian formulation of quantum measurements

Note that Eq. (4.3) satisfies $\text{Tr}(d\rho_t) = d(\|\psi_t\|^2) = 0$, that is, the norm of the state is conserved throughout the entire evolution. As a consequence, the system's state evolves unitarily, although stochastically, across the Hilbert space. This suggests that a Hamiltonian interpretation of the monitored dynamics (4.3) should be possible.

Hamiltonian formulations that reproduce the paths traced by a set of connected pure states were studied in the deterministic case in the context of the quantum brachistochrone problem [207, 208] and counterdiabatic driving [209]. In the stochastic case, a Hamiltonian formulation of the monitored dynamics is not feasible within the Itô picture, as non-Hermitian corrections to the dynamics prevent it from being expressed as a quantum Liouville equation

$$d\rho_t = -i[dH_t, \rho_t] \quad (4.5)$$

for some stochastic Hermitian generator process dH_t .

However, the situation changes in the Stratonovich picture, where the rules of stochastic calculus take the same form as those of standard calculus. Although computations are generally less straightforward than in the Itô picture, the Stratonovich picture offers descriptions of stochastic processes that align more naturally with physical intuition [210, 211]. In particular, the Stratonovich formulation of (4.3) allows us to define a stochastic Hamiltonian governing the monitored dynamics. A recent precedent is Ref. [212] where,

for the Hermitian case $C = C^\dagger$, the authors derive a Hamiltonian that retraces the monitored dynamics and use it to define feedback protocols to steer the quantum arrow of time. Our result generalizes that of [212] to the case of arbitrary jump operators. Moreover, we recast the dynamics in terms of single- and double-bracket commutators, offering a new connection between open quantum dynamics and double-bracket flows [213]. In what follows, we focus on processes whose dynamics can be written in Stratonovich form.

We illustrate a straightforward connection between double brackets and quantum dynamics by noting that any stochastic dynamics that preserves a state's purity satisfies

$$d\rho_t = [[d\rho_t, \rho_t], \rho_t]. \quad (4.6)$$

This follows directly by expanding the commutators and using $\rho_t d\rho_t \rho_t = 0$, which itself is obtained by differentiating both sides of $\rho_t^2 = \rho_t$ [214].

It is straightforward to identify relation (4.6) as a quantum Liouville equation of the form (4.5), with Hamiltonian generator given by $dH_t := i[d\rho_t, \rho_t]$. Note that the same reasoning holds true in the deterministic case. Moreover, if ρ_t is smooth, we have $\dot{\rho}_t = -i[H_t, \rho_t]$ with Hamiltonian $H_t = i[\dot{\rho}_t, \rho_t]$.

Equations (4.5)–(4.6) give a recipe for building the Hamiltonian generator of a predefined quantum trajectory $\rho_{0:T}$, where T is the final time. This procedure is the Liouville-space analog, for stochastic processes, of constructing Schrödinger equations that reproduce trajectories in Hilbert space [207, 208]. Conversely, a predefined Hamiltonian generator dH_t yields the corresponding pure-state dynamics; for instance, setting $dH_t = H_0 dt$ gives the well-known quantum Liouville equation $\dot{\rho}_t = -i[H_0, \rho_t]$.

Note that a generator process like the last example, $dH_t = H_0 dt$, reduces the dynamics to a single-bracket commutator form. An immediate question is whether higher-order bracket contributions can be constructed that still preserve the purity of the state. Under the constraint of purity preservation $\rho_t^2 = \rho_t$, it follows that single- and double-bracket terms constitute the most general possible contributions. To see this, consider the following

decomposition

$$d\rho = \mathcal{C}_\rho(d\Lambda_1) + \mathcal{C}_\rho^2(d\Lambda_2) + \mathcal{C}_\rho^3(d\Lambda_3) + \dots \quad (4.7)$$

where, for the moment, we omit the time subscript. In the equation above, we denote the super-operator $\mathcal{C}_Q(\bullet) := [\bullet, Q]$ for an arbitrary operator Q . The quantities $d\Lambda_i$ ($i = 1, \dots$) are arbitrary (not necessarily Hermitian) generator processes. When ρ is pure, it is easy to show that the super-operator $\mathcal{C}_\rho$ satisfies the periodic property $\mathcal{C}_\rho^{2k+1} = \mathcal{C}_\rho$ ($k = 0, 1, \dots$). This further implies that $\mathcal{C}_\rho^{2k} = \mathcal{C}_\rho^2$ ($k = 0, 1, \dots$). Then, these two relations lead to the equivalent expansion

$$d\rho = \mathcal{C}_\rho(d\Lambda_1 + d\Lambda_3 + \dots) + \mathcal{C}_\rho^2(d\Lambda_2 + d\Lambda_4 + \dots). \quad (4.8)$$

We conclude that the most general multi-bracket expansion for pure dynamics is equivalent to a single- and double-bracket decomposition of the form

$$d\rho_t = -i[dH_t^{\text{SB}}, \rho_t] + [[dH_t^{\text{DB}}, \rho_t], \rho_t], \quad (4.9)$$

where dH_t^{SB} and dH_t^{DB} are some Hermitian generator processes.

Next, we derive explicit expressions for the generators dH_t^{SB} and dH_t^{DB} , when the dynamics describes monitored open quantum systems governed by (4.3). These expressions are in terms of the Lindblad jump operator C and the measurement process dy_t . We establish the following result.

Theorem 4. *The monitored dynamics (4.3) in Stratonovich form is given by the stochastic quantum Liouville equation*

$$d\rho_t = -i[dH_t^{\text{SB}}, \rho_t] + [[dH_t^{\text{DB}}, \rho_t], \rho_t], \quad (4.10)$$

where the single-bracket and double-bracket Hamiltonian generator processes

are, respectively,

$$dH_t^{SB} := H_0 dt - i \frac{\gamma}{2} \{A, B\} dt + i B dy_t \quad (4.11)$$

and

$$dH_t^{DB} := -\gamma A^2 dt - \frac{\gamma}{2} [A, B] dt + A dy_t \quad (4.12)$$

Equation (4.10) should be interpreted according to the midpoint evaluation rule [211, 215] characteristic of the Stratonovich picture. We leave the proof of Theorem 4 and its extension to multiple noise channels to the Appendices 4.7.1-4.7.4.

The Hamiltonian process $H_t^{SB} = \int_0^t dH_\tau^{SB}$ can be interpreted as a modification of the system Hamiltonian H_0 by the noise back-action due to the system-environment interaction conditioned on the measurement process. Conversely, the Hamiltonian process $H_t^{DB} = \int_0^t dH_\tau^{DB}$ encodes the measurement back-action resulting from the measurement process. Note that both processes H_t^{SB} and H_t^{DB} depend on the state ρ_t only via the measurement process dy_t .

We note that, in the Hermitian case $C = C^\dagger$, Eq. (4.10) reduces to the monitored dynamics derived in Ref. [212]. See App. 4.7.2 for the explicit correspondence.

It is important to highlight the physical significance of Eq. (4.10). We successfully expressed what essentially represents an open quantum dynamics as a stochastic Hamiltonian evolution of the form of (4.5) with Hamiltonian generator $dH_t = dH_t^{SB} + i[dH_t^{DB}, \rho_t]$. This formulation is valid for *any Lindblad operator* C .

In general, quantum Liouville equations are written in Lax form [216], which are isospectral flows preserving the eigenvalues of the density operator. On the other hand, it is known that Itô diffusions of the form (4.3) tend to asymptotically purify initially mixed states [40]. This implies that the spectrum of mixed density operators is not preserved during the evolution, and a Hamiltonian formulation of the dynamics is no longer applicable.




Figure 4.1 Illustration of the Hamiltonian evolution given by (4.10). The evolution takes place on the unitary orbit $\mathcal{O}$ of the density operator. In the tangent space $\mathcal{T}_\rho \mathcal{O}$ at point ρ , the increment $d\rho$ is composed of two (not necessarily orthogonal) contributions, $d\rho^{\text{SB}} = -i[dH_t^{\text{SB}}, \rho_t]$ and $d\rho^{\text{DB}} = [[dH_t^{\text{DB}}, \rho_t], \rho_t]$, corresponding to the single-bracket and double-bracket Hamiltonians, respectively.

We remark that, while the monitored state ψ_t traces a unitary path across the Hilbert space, it does not mean that the corresponding stochastic propagator [see (4.63) in the Appendices] is a valid unitary operator for arbitrary states, since its definition depends explicitly on the monitored state.

Equation (4.10) comes equipped with a direct geometrical interpretation. The evolution takes place in the so-called unitary orbit of the state ρ , defined by $\mathcal{O}_\rho = \{U\rho U^\dagger | U \text{ unitary}\}$. At any given time, the increment $d\rho$ takes place in the tangent space of the orbit at point ρ , $\mathcal{T}_\rho \mathcal{O}_\rho$ [217–219], and is composed of single- and double-bracket contributions, $d\rho^{\text{SB}}$ and $d\rho^{\text{DB}}$, respectively, each of which is a commutator with the state. Thus, the Hamiltonian processes are the generators of dynamical flows in the orbit of the state. We illustrate this geometrical picture in Fig. 4.1. We return to this picture in the next section in the context of gradient flows.

4.4 Connection with double-bracket gradient flows

4.4.1 Background

Gradient flows represent a powerful approach to constrained optimization problems [42, 220]. Generically defined as evolutions on a constrained manifold, gradient flows follow paths of steepest descent (or ascent) of a given quality function defined on that manifold. In quantum applications, gradient flows play a crucial role in enabling advanced methods for quantum control [43, 221] and quantum machine learning [222, 223].

On the other hand, double-bracket (DB) flows have, in recent years, gained increasing attention within the quantum computing community for the design of efficient quantum algorithms [43–45, 224]. DB flows are commonly introduced as equations of the form $dX/dt = [[N, X(t)], X(t)]$ where N is a diagonal matrix with unique entries. These systems exhibit the remarkable property that the unique stable fixed point $X_\infty = \lim_{t \rightarrow \infty} X(t)$ of the dynamics is also a diagonal matrix with elements following the same sorting as those of N . In other words, if $\text{diag}(N)$ is sorted in increasing order, $\text{diag}(X_\infty)$ is in increasing order as well [213].

In geometrical terms, DB flow equations can be formulated as gradient flows on the adjoint orbit of a Lie group [43, 216, 225]. This connection was recently exploited in Refs. [45, 226] to frame imaginary-time evolutions [227] as DB flows for ground state preparation, enabling coherent implementations on quantum hardware. In this context, DB flows provide the crucial theoretical bridge between the fundamentally non-unitary operation defined by the imaginary-time evolution and a circuit implementation in terms of unitary operations.

More broadly, it is noteworthy that connections between DB gradient flows and quantum-mechanical treatments arise in the literature exclusively via imaginary-time evolution, rather than real-time quantum dynamics. Nonetheless, in Theorem 4 of the previous section, we make a precise connection between DB flows and the latter. Thus, it is not surprising that this connection can be taken a step further by framing the quantum Liouville equation (4.10) in the language of gradient flows. For this, we write each DB quantity in (4.10) as a Riemannian gradient on the unitary orbit of ρ_t .

4.4.2 Results

The Hilbert-Schmidt (HS) metric on the unitary group $U(N)$ is given by the inner product $\langle X, Y \rangle_{\text{HS}} := \text{Tr}(X^\dagger Y)$, where X, Y belong to the Lie algebra of $U(N)$. The HS metric restricted to the unitary orbit [43] defines a Riemannian metric $\langle \bullet, \bullet \rangle_{\mathcal{O}}$ such that $\langle X, Y \rangle_{\text{HS}} = \langle X', Y' \rangle_{\mathcal{O}}$, where $X' = [X, \rho]$ and $Y' = [Y, \rho]$ are tangent vectors on the state's orbit.

To motivate the subsequent discussions, consider Eq. (4.6) in the deterministic case, that is, $\dot{\rho}_t = [[\dot{\rho}_t, \rho_t], \rho_t]$. By defining the scalar potential $F_t(\sigma) = \text{Tr}(\dot{\rho}_t \sigma)$ with σ in the state's orbit, one can prove that the dynamics can be written as the gradient flow $\dot{\rho}_t = \nabla F_t(\rho_t)$, where $\nabla F_t(\rho_t)$ is the Riemannian gradient induced by the HS metric on the state's orbit.

A more general result turns out to be true for stochastic processes.

Proposition 1. *A DB equation of the form $d\rho_t = [[dO_t, \rho_t], \rho_t]$, with dO_t some Hermitian generator process, admits the gradient flow interpretation $d\rho_t = \nabla F_t(\rho_t)dt$, with potential given by $F_t(\sigma) = \text{Tr}(\dot{O}_t \sigma)$ and $\dot{O}_t = dO_t/dt$.*

This general result is, to the best of our knowledge, new, and is a direct consequence of the double-bracket form of the dynamics. We refer the reader to App. 4.7.5 for the full derivation.

Note that the deterministic case $\dot{\rho}_t = [[\dot{\rho}_t, \rho_t], \rho_t]$ is a particular instance of Prop. 1 in the absence of noise and when $O_t \equiv \rho_t$.

Proposition 1 is to be interpreted in the sense of stochastic differential equations. A general Hermitian generator process dO_t is of the form $dO_t = M_t dt + \Sigma_t dW_t$, for some matrix-valued processes, M_t and Σ_t . As a consequence, $F_t(\sigma)dt = \text{Tr}(M_t \sigma)dt + \text{Tr}(\Sigma_t \sigma) dW_t$, and we write the dynamics as $d\rho_t = \nabla F_t^{(1)}(\rho_t)dt + \nabla F_t^{(2)}(\rho_t)dW_t$ with $F_t^{(1)}(\sigma) = \text{Tr}(M_t \sigma)$ and $F_t^{(2)}(\sigma) = \text{Tr}(\Sigma_t \sigma)$.

A direct application of Prop. 1 to the Hamiltonian dynamics (4.10) leads to the following interpretation of measurement processes as gradient flows.

Theorem 5. *The monitored dynamics (4.10) can be written as a gradient flow*

on the unitary orbit as

$$d\rho_t = \nabla F_t(\rho_t)dt, \quad (4.13)$$

where the potential F_t consists of the sum of two parts,

$$F_t^{SB}(\sigma) = \text{Tr} \left(-i[\dot{H}_t^{SB}, \rho_t]\sigma \right) \quad (4.14)$$

and

$$F_t^{DB}(\sigma) = \text{Tr} \left(\dot{H}_t^{DB}\sigma \right). \quad (4.15)$$

$\dot{H}_t^{SB} \equiv dH_t^{SB}/dt$ and $\dot{H}_t^{DB} \equiv dH_t^{DB}/dt$ are the instantaneous single- and double-bracket Hamiltonians, respectively, and σ belongs to the state's orbit.

Consult App. 4.7.5 for the proof.

In Fig. 4.2, we illustrate the flow formulation given by Theorem 5. To facilitate an intuitive picture of the underlying geometry, we can think of the state's orbit as an embedded submanifold of the unitary group, with the latter being represented by a sphere. The state's orbit is represented by a circumference, whose tangent vectors are orthogonal to those vectors X in the Lie algebra that commute with the state, i.e., the generators of the stabilizer group [218]. These vectors represent directions in the unitary group that do not generate any dynamics, that is, unitary operations such that $U\rho_t U^\dagger = \rho_t$.

To gain a deeper physical understanding of Eq. (4.13), let's write the flow equation in the equivalent form

$$d\rho_t = \nabla F_t^{SB}(\rho_t)dt - \gamma \nabla V_t(\rho_t)dt + \nabla K(\rho_t)dW_t, \quad (4.16)$$

where we define the potentials $V_t(\sigma) := \text{Tr}(\Delta A_t^2 \sigma) + \frac{1}{2} \text{Tr}([A, B]\sigma)$, with $\Delta A_t := A - \langle A \rangle_t$, and $K(\sigma) := \text{Tr}(A\sigma)$. The equivalence between both flows is due to the gauge freedom present in the commutator $\mathcal{C}_\rho(dO_t + d\alpha_t) = \mathcal{C}_\rho(dO_t)$ for an arbitrary scalar process α_t .

The potentials V_t and K are quantum expectations of various observables




Figure 4.2 Schematic representation of the gradient flow formulation in Theorem 5. We illustrate the embedding of the state's orbit $\mathcal{O}_\rho$ (circumference), in the unitary group $U(N)$ (sphere). The tangent vectors to the orbit correspond to vectors in the Lie algebra that are orthogonal to the stabilizer algebra of ρ , which is generated by vectors X that commute with ρ . In particular, the gradient flow generates displacements $\delta\rho_t = \nabla F_t \delta t$ tangent to the orbit.

defined in terms of A and B . We write V_t with a time subscript to emphasize that its definition depends on ρ_t , with domain in the unitary orbit of ρ_t . The value $V_t(\rho_t) = \text{Tr}(\Delta A_t^2 \rho_t) + \frac{1}{2} \text{Tr}([A, B] \rho_t)$ has two contributions with direct physical interpretation. The first term corresponds to the variance of the observable A . The second term can be interpreted as the lower bound of the Robertson inequality [228] between A and B , that is, $-\langle \Delta A_t^2 \rangle \langle \Delta B_t^2 \rangle \geq |\langle [A, B] \rangle|^2 / 4$, where $\Delta B_t := B - \langle B \rangle_t$ (the minus sign is due to the anti-hermiticity of B).

In the diffusion term, the potential $K(\rho_t)$ represents the quantum expectation of the observable A . At an arbitrary point ρ , the gradients of these functions are vectors tangent to the state's orbit. Because (4.16) is stochastic, one can interpret the gradient $\nabla V_t(\rho_t)$ as inducing non-zero drifts in the orbit in directions where the variation in the quantum expectations are the steepest. Conversely, the gradient corresponding to the diffusion term does not have a preferred direction due to the undefined sign of the Wiener increment dW_t .

Let us next focus on the pure-measurement case where $B = 0$, and assume, for simplicity, that $H_0 = 0$ as well. The single-bracket term in (4.16)

vanishes, and the dynamics, expressed in terms of DBs, are described by

$$d\rho_t = -\gamma[[\Delta A_t^2, \rho_t], \rho_t]dt + [[A, \rho_t], \rho_t]dW_t. \quad (4.17)$$

The stationary points of (4.17) coincide with the eigenstates of A , as they should for pure measurement processes. One can show by computing the evolution equation of the overlap between the state and any of the eigenspaces of A that these subspaces are attractors of the dynamics, and therefore stable. Moreover, when the state is close to an eigenspace of A , the variance in the overlap tends to vanish, making it harder for the system to escape from the vicinity of the attractor subspace.

Alternatively, these observations can be drawn from the gradient interpretation of (4.17). The drift term, which is equivalent to $-\gamma\nabla V_t(\rho_t)dt$, steers the state in directions that minimize the variance $V_t(\rho_t)$ associated to the operator A . This can be seen more clearly by computing the Stratonovich differential of $V(t) := \text{Tr}(\Delta A_t^2 \rho_t)$, which gives

$$dV(t) = -\gamma\|\nabla V_t(\rho_t)\|_{\mathcal{O}}^2 dt + \langle \nabla V_t(\rho_t), \nabla K(\rho_t) \rangle_{\mathcal{O}} dW_t, \quad (4.18)$$

where we used the compatibility condition of the metric (see App. section 4.7.5 for definitions) and the pure-measurement dynamics $d\rho_t = -\gamma\nabla V_t(\rho_t)dt + \nabla K(\rho_t)dW_t$. See App. 4.7.5 for the derivation. Equation (4.18) says that the variance of the measurement operator tends to decrease along the solution trajectory with rate proportional to the squared norm of its gradient on the orbit. Since the minima of the potential V_t correspond to the eigenspaces of A , the long-time dynamics will localize the state in one of these eigenspaces, and the measurement apparatus will retrieve the corresponding eigenvalue almost surely. See Fig. 4.3 for a schematic illustration of this idea. One must note the resemblances with Lyapunov stability [211], where the potential V_t can be related to the Lyapunov function of the dynamics.

We note that the monitored dynamics tends to minimize the variance of the measurement operator is known from previous work; see, for instance,




Figure 4.3 Illustration of the pure-measurement case $C = A$ with $H_0 = 0$. The increment $d\rho$ is composed only of the double-bracket term $d\rho^{\text{DB}}$, with drift dictated by the Riemannian gradient of the the variance $V_t(\rho_t) = \text{Tr}(\Delta A_t^2 \rho_t)$. The variance vanishes when ρ_t corresponds to an eigenstate $|a\rangle$ of A .

Refs. [229–232]. The novelty of Eq. (4.17), or its gradient flow equivalent, is that it gives a streamlined interpretation of the monitored dynamics as a continuous collapse process, where the asymptotic fate of the state is a stationary density operator with the correct Born probabilities [39, 233, 234].

When the system Hamiltonian H_0 is non-zero, the asymptotic collapse analysis remains valid under the assumption of quantum non-demolition measurements, i.e. when H_0 commutes with the measurement operator A [39, 235]. When H_0 and A do not commute, localization is still possible in non-trivial decoherence-free subspaces [236].

It is straightforward to generalize Eq. (4.17) to multiple measurement operators A_i as a linear combination of DB flows contributions, one per noise channel. In this case, the different flows compete to reduce the variance of each measurement, with localization being possible in the common eigenspaces associated with all measurement operators.

In the most general case where $H_0 \neq 0$, $A \neq 0$, and $B \neq 0$, Eq. (4.10) involves various competing terms that prevent collapse unless special conditions are met. For instance, when A and B commute. Assuming the existence of non-trivial common eigenspaces between H_0 and A , one can show that these subspaces are attractors of the dynamics by studying the behavior of

the overlap of the system's state with any of the subspaces. In particular, if the state enters a common eigenspace corresponding to an eigenvalue a , the dynamics becomes the quantum Liouville equation (4.5) with Hamiltonian $dH_t := iB_a(-\gamma a dt + dy_t)$, where $B_a = \Pi_a B \Pi_a$ is the orthogonal projection of B onto the eigenspace of A with eigen-projector Π_a . Thus, the effect of the anti-Hermitian operator B is to unitarily scramble the state within the selected eigenspace. Setting the hitting time in the selected eigenspace to $t = 0$, with corresponding state ρ_0 , we can integrate the dynamics and write the solution as $\rho_t = U(t)\rho_0 U(t)^\dagger$, with the propagator $U(t) = \exp(B_a(-\gamma at + y_t))$ and the accumulated measurement outcome $y_t = \int_0^t dy_s$.

For general non-commuting A, B , the variance of A is lower bounded by the Robertson inequality $-\langle \Delta A_t^2 \rangle \langle \Delta B_t^2 \rangle \geq |\langle [A, B] \rangle|^2/4$, which might prevent further localization. Nevertheless, this lower bound is exactly one of the quantities that the dynamics tends to minimize, together with the variance of the observable A . As a consequence, a synergy emerges between the two DB terms such that the variance in A gets minimized as much as the minimization of the Robertson bound allows. We already described a related phenomenon in the previous paragraph assuming the stronger condition of commuting $[A, B] = 0$. In that case, the dynamics is free to localize the state in any of the common eigenspaces of H_0 and A . In the non-commuting case, the dynamics will still promote (although incomplete) localization by simultaneously minimizing the Robertson bound and the variance in A . This inevitably resembles the case of coherent states for continuous measurement operators. For the harmonic oscillator with jump operator $C = X + iP$ where X, P are the position and momentum operators, respectively, the monitored dynamics asymptotically converges to coherent states [39]. In this case, the Robertson inequality with $A = X$ and $B = iP$ becomes increasingly saturated, and the state becomes increasingly coherent, simultaneously reducing the uncertainty in A and B .

4.5 Applications to feedback control

4.5.1 Feedback-induced DB flows

We present a notable case where the stochastic dynamics (4.10) can be turned into a deterministic DB flow equation after incorporating feedback into the dynamics.

Consider the pure-measurement case where $B = 0$. Assume that the measurement apparatus couples to the system's energy operator, i.e., $A = H_0$. The measurement process is characterized by a measurement record $dy_t = 2\gamma \langle H_0 \rangle_t dt + dW_t$. After each measurement at time t , we inject into Eq. (4.10) the coherent feedback Hamiltonian $dH_{fb} = i[H_0, \rho_t]u_t dt$, where u_t is a real number representing the control signal. We design the feedback such that it precisely counteracts the measurement's influence on the system dynamics, achieved by setting $u_t dt = -dy_t$. Under this feedback scheme, the system evolves according to the equation

$$\frac{d\rho_t}{dt} = -i[H_0, \rho_t] - \gamma [[H_0^2, \rho_t], \rho_t] . \quad (4.19)$$

The effect of counteracting the measurement process is to suppress the stochasticity in the dynamics, which subsequently leads to the deterministic equation (4.19). This scheme can be seen as a noise-canceling protocol. In related work, Ref. [237] uses coherent feedback to cancel the background noise due to the environment, and Ref. [212] uses coherent feedback to stabilize the state of a system, thereby withdrawing energy from the monitoring process. We use instead a feedback control to suppress the measurement back-action, which leads to a deterministic feedback-induced DB dynamics that, as we see below, can be used for ground state preparation.

Equation (4.19) comprises a deterministic isospectral evolution for which we can immediately apply known results from double-bracket flow theory. Without loss of generality, consider the spectrum of H_0 to be $0 \leq E_0 < E_1 \leq \dots \leq E_N$ with unique ground-state energy E_0 . We establish the following result.

Theorem 6. *The unique stable fixed point of the feedback dynamics (4.19) is given by the ground state of H_0 .*

One can prove this in a number of ways. One is transforming to the interaction picture, and invoking the results by Brockett [213] and Helmke [42] on asymptotic stability of DB flows. Another way is analyzing the behavior of populations for the different energy levels and show that the ground state is the only stable asymptotic state; we show this in App. 4.7.6.

An important feature of double-bracket flow evolutions is the exponential convergence to the stable solution [224]. In particular, the feedback protocol described above converges exponentially fast to the ground state of the system [42] provided there is an initial non-zero overlap with the ground state [44, 45].

We illustrate the exponential convergence of the feedback-induced DB dynamics (4.19) for the 1-D antiferromagnetic Heisenberg model with Hamiltonian $H_0 = \sum_{i=1}^{n-1} (\sigma_i^x \sigma_{i+1}^x + \sigma_i^y \sigma_{i+1}^y + \sigma_i^z \sigma_{i+1}^z)$ where $\sigma_i^x, \sigma_i^y, \sigma_i^z$ are Pauli strings. This model exhibits a unique ground state for even numbers of qubits [238, 239]. As initial condition we use (instead of a random state) a physically motivated [45] tensor product state given by $|\psi_0\rangle := 2^{-n/4} (|10\rangle - |01\rangle)^{\otimes n/2}$. We evolve the system using the dynamics (4.19) over a certain time interval for different qubit numbers ($n = 4, 6, 8, 10$), and monitor the error in computing the ground state energy E_0 . We show the results in Fig. 4.4. We observe that, for the same damping rate value γ , the error decreases exponentially fast with time across different qubit numbers, reaching in all cases machine precision.

The feedback protocol described above leads to a concrete realization of dynamical ground-state preparation by harnessing properties of DB flows. However, the fact that preparing the feedback Hamiltonian dH_{fb} requires precise knowledge of the state can limit the applicability of the protocol, as its implementation would require real-time parallel state estimation. This is feasible for small dimensional systems, and has been the subject of active research [240–245].

One possible alternative approach is to entirely forgo state estimation,




Figure 4.4 Energy errors for the 1-D Heisenberg model. We monitor the error between the energy expectation $\text{Tr}(H_0 \rho_t)$ and the true ground state energy E_0 obtained from exact diagonalization. We do this for a total time $T = 2$ and 300 time steps. The damping rate is $\gamma = 4 \times 10^{-1}$. To ensure a positive spectrum, we shift the system Hamiltonian by adding a constant offset of 18 for all qubit numbers. The y -axis is in logarithmic scale, showcasing the exponential convergence of the DB dynamics with errors reaching machine precision.

and consider feedback Hamiltonians that, apart from knowledge of the measurement outcomes, remain completely agnostic of the state.

4.5.2 State-agnostic feedback for state preparation

Let's suppose we aim to prepare/stabilize an arbitrary pure state $Q = |q\rangle\langle q|$. This time, we assume we only have access to the measurement outcomes dy_t , i.e., we don't know the system's state. We assume a setup with ideal measurement efficiency and no delay in the feedback loop. In what follows, we present a general blueprint for measurement-based feedback control that is agnostic of the state and produces dynamics that have as fixed point the state Q we want to prepare or stabilize.

Consider a system evolving under the monitored dynamics (4.10). Immediately after each measurement, we add to the dynamics two feedback Hamiltonians that depend, apart from system specifications, only on the current measurement outcome and the target state. The first feedback is a counter-driving Hamiltonian process $dH_t^{\text{fb}1} = -dH_t^{\text{SB}}$ whose effect is to

exactly cancel the action of the single-bracket Hamiltonian on the dynamics. The second feedback is $dH_t^{\text{fb}2} = -i[dH_t^{\text{DB}}, Q]$, i.e., the negative of the DB Hamiltonian process with the state ρ being replaced by Q . Given these definitions, the effective dynamics is

$$d\rho_t = -\gamma [[A^2, \rho_t - Q], \rho_t] dt - \frac{\gamma}{2} [[[A, B], \rho_t - Q], \rho_t] dt + [[A, \rho_t - Q], \rho_t] dy_t. \quad (4.20)$$

The generalization of (4.20) to multiple noise channels is straightforward.

Q is a fixed point of the dynamics, meaning that whenever the state reaches the target, the dynamics immediately comes to a halt. A natural question concerns the stability of Q as a fixed-point solution. The answer is problem-dependent, and we illustrate this below with examples.

Equation (4.20) can be interpreted in two different yet complementary ways. In the first view, we are given an uncontrollable environment that can be probed to gather information about the system. A stability analysis determines which regions of the Hilbert space are suitable for state preparation and stabilization. In the second view, we are given a target state Q , and the task is to design suitable system-environment interactions such that the target state is a fixed point of the dynamics and can be certified as stable. This perspective is closely related to dissipative reservoir engineering (also called bath or environment engineering), where one tailors the system-environment coupling so that the dissipative dynamics autonomously drives the system toward and stabilizes the desired target state [246–251]. We illustrate the interplay of these two views with an example.

Assume a single-qubit system undergoing amplitude damping at a rate $\gamma_- > 0$ and described by the jump operator $C = \sigma^- = (\sigma_x - i\sigma_y)/2$, where σ_x, σ_y are Pauli operators. The measuring apparatus is coupled to the Hermitian part of C with measurement record given by $dy_t^- = \gamma_- \langle \sigma_x \rangle_t dt + dW_t^-$ and $(dW_t^-)^2 = \gamma_- dt$. Assume we want to prepare a state $|q\rangle$ defined by Bloch angles (θ, ϕ) . Ref. [252] analyzed a similar example with an alternative control scheme. Consider the overlap process $dQ_{\parallel}(t) = \text{Tr}(Q d\rho_t)$ that quantifies the support of the state along the target. A stability analysis of the Lyapunov ex-

ponent of the overlap process shows that the region of stability corresponds to states satisfying $\cos(\theta) < 0$ (see App. 4.7.7 for details). This means that the feedback protocol can prepare and stabilize arbitrary states in the lower hemisphere of the Bloch sphere.

That the upper-half of the sphere is unstable is a natural consequence of the energy loss induced by the amplitude damping. To counteract this loss, we introduce a second feedback mechanism. Assume that the system is coupled to an energy reservoir. We model this by augmenting the previous setting with a jump operator $\sigma_+ = (\sigma_x + i\sigma_y)/2$, with damping rate γ_+ . The measurement record is given by $dy_t^+ = \gamma_+ \langle \sigma_x \rangle_t dt + dW_t^+$ with $(dW_t^+)^2 = \gamma_+ dt$. Note that, although both jump operators lead to the same observable, we have two independent measurement streams representing different sources. The feedback-induced dynamics for this configuration is

$$d\rho_t = \frac{\Delta\gamma}{4} [[\sigma_z, \rho_t - Q], \rho_t] dt + \frac{1}{2} [[\sigma_x, \rho_t - Q], \rho_t] dy_t, \quad (4.21)$$

where $\Delta\gamma := \gamma_+ - \gamma_-$ and $dy_t = dy_t^+ + dy_t^-$. By analyzing the stability of the fixed-point solution Q , we find that the stable region satisfies

$$\Delta\gamma \cos(\theta) + 2(\gamma_+ + \gamma_-) \sin^2(\theta) \cos^2(\phi) > 0. \quad (4.22)$$

For target states with angle $\phi = \pi/2$, stabilization in the upper hemisphere ($\theta < \pi/2$) requires $\Delta\gamma > 0$; that is, energy injection must exceed dissipation due to damping. For angles $\phi \neq \pi/2$, the condition $\Delta\gamma > 0$ is merely sufficient to achieve stability in the upper hemisphere. The equatorial line $\theta = \pi/2$ also represents fixed points of the dynamics. There, stability is achieved trivially for $\phi \neq \pi/2$. This implies, as shown by simulations, that states close to (and including) the eigenstates of σ_x are increasingly more stable than those near the eigenstates of σ_y . Conversely, a sufficient condition for stabilization in the lower hemisphere is $\Delta\gamma < 0$; that is, energy renewal must not exceed the losses due to dissipation.

The stabilization timescale τ_s under small perturbations around the target scales like $\tau_s \sim |\Delta\gamma \cos(\theta) + 2(\gamma_+ + \gamma_-) \sin^2(\theta) \cos^2(\phi)|^{-1}$. Thus, stabilization




Figure 4.5 Measurement-based feedback state preparation and stabilization for the single-qubit. We display (top panel) the time evolution of the Bloch coordinates $m_i(t) = \text{Tr}(\sigma_i \rho_t)$ ($i = x, y, z$) of 50 trajectories along their ensemble averages $\bar{m}_x, \bar{m}_y, \bar{m}_z$. The parameters are $\gamma_- = 1$ and $\gamma_+ = 4$, with final time $T = 8$, for a total of 800 time steps. All trajectories stabilize relatively fast around the target, with a final average infidelity (bottom panel) $1 - \overline{\text{Tr}(Q \rho_t)} \sim 10^{-6}$, and median $\sim 10^{-10}$.

is generally more effective for larger damping gaps $\Delta\gamma$. This means that, to accelerate convergence in the upper hemisphere, one must increase the pumping rate γ_+ . For the lower hemisphere, one should set $\gamma_+ = 0$, thereby decoupling the energy reservoir from the system. This is illustrated in Fig. 4.5. We prepare a state in the upper hemisphere with Bloch angles $(\theta, \phi) = (\pi/4, \pi/4)$, for an initial condition equal to $|1\rangle$. By considering a non-negligible gap $\Delta\gamma$, we achieve a relatively fast convergence and stabilization around the target of an ensemble of quantum trajectories.

4.5.3 Gradient flow interpretation of feedback dynamics

Next, we interpret general state-agnostic feedback protocols within the context of gradient flows. A general feedback of this kind can be implemented by adding a Hamiltonian generator dH_t^{fb} to the dynamics (4.10), so that the latter is transformed into

$$d\rho_t = -i\mathcal{C}_\rho(dH_t^{\text{SB}}) + \mathcal{C}_\rho^2(dH_t^{\text{DB}}) - i\mathcal{C}_\rho(dH_t^{\text{fb}}). \quad (4.23)$$

Using the periodic property $\mathcal{C}_\rho^3 = \mathcal{C}_\rho$, we can write the feedback equation in DB form

$$d\rho_t = \mathcal{C}_\rho^2 \left(-i\mathcal{C}_\rho(dH_t^{\text{SB}} + dH_t^{\text{fb}}) + dH_t^{\text{DB}} \right). \quad (4.24)$$

As a direct application of Prop. 1 to Eq. (4.24) gives the following.

Corollary 3. *The feedback dynamics (4.23) is equivalent to a gradient flow $d\rho_t = \nabla F_t(\rho_t)dt$ with potential*

$$F_t(\sigma) = \text{Tr} \left(\left(-i\mathcal{C}_\rho(\dot{H}_t^{\text{SB}} + \dot{H}_t^{\text{fb}}) + \dot{H}_t^{\text{DB}} \right) \sigma \right). \quad (4.25)$$

Enforcing fixed-point dynamics would demand the design of a feedback process such that $\nabla F_t(\rho_t = \text{fixed point}) = 0$. This is indeed the case for the state-agnostic protocol treated in the previous section, where the target is an arbitrary state $Q = |q\rangle\langle q|$. In that case, the feedback process is given by $dH_t^{\text{fb}} = -dH_t^{\text{SB}} - i\mathcal{C}_Q(dH_t^{\text{DB}})$, where $\mathcal{C}_Q(\bullet) = [\bullet, Q]$, and we write the dynamics as

$$d\rho_t = \mathcal{C}_\rho \mathcal{C}_{\rho-Q}(dH_t^{\text{DB}}). \quad (4.26)$$

When $\rho = Q$, $d\rho_t \equiv 0$, as expected. The corresponding potential is $F_t(\sigma) = \text{Tr} \left(\mathcal{C}_\rho \mathcal{C}_{\rho-Q}(\dot{H}_t^{\text{DB}}) \sigma \right)$, with flow equation

$$d\rho_t = -\gamma \nabla F^{(1)}(\rho_t)dt + \nabla F^{(2)}(\rho_t)dy_t, \quad (4.27)$$

where

$$F_t^{(1)}(\sigma) := \text{Tr} \left(\mathcal{C}_\rho \mathcal{C}_{\rho-Q} (A^2 + [A, B]/2) \sigma \right) \quad (4.28)$$

and

$$F_t^{(2)}(\sigma) := \text{Tr} \left(\mathcal{C}_\rho \mathcal{C}_{\rho-Q} (A) \sigma \right) . \quad (4.29)$$

In the previous section, we showed that a Lyapunov stability analysis determines whether the solution $\rho = Q$ is a stable fixed point and provides sufficient conditions under which this is achieved. With Corollary 3, we can re-interpret the stability of Q as that of a stationary solution of a certain loss landscape. For this, let's compute the change in the potential increment

$$F_t(\sigma)dt = \text{Tr} \left(\mathcal{C}_\rho \mathcal{C}_{\rho-Q} (dH_t^{\text{DB}}) \sigma \right) \quad (4.30)$$

under a small dynamical update $\rho_t \rightarrow \rho_t + d\rho_t$ in the direction of the flow. The result is

$$(F_t dt)(\rho_t + d\rho_t) = \gamma \|\nabla F_t^{(2)}(\rho_t)\|_{\mathcal{O}}^2 dt + \mathcal{O}(dt^2) . \quad (4.31)$$

This shows that the potential increment $F_t dt$ is non-decreasing when the state is updated according to the flow. Note that, in particular, for $\rho = Q$, the gradient vanishes $\nabla F_t^{(2)}(Q) = 0$, which implies that the potential has reached a stationary solution. Whether the solution is stable or not depends on a second order analysis of $F_t(\sigma)dt$ around $\sigma = Q$. An approach is the following.

Note that, in general, the increment (4.30) quantifies the change in the potential when we move along the orbit of ρ_t . Since the increment is non-negative, if Q is a stable fixed-point, we must have $F_t(\rho_t + \tau \delta \rho)dt = F_t(Q)dt > 0$ for any non-zero difference $Q - \rho_t = \tau \delta \rho$, with $\tau > 0$ sufficiently small and $\delta \rho$ being any fixed traceless operator such that ρ_t remains in the orbit of Q .

Because

$$F_t(Q)dt = \tau \underbrace{\text{Tr} \left(Q C_Q C_{\delta\rho} (dH_t^{\text{DB}}) \right)}_{=0} + \tau^2 \text{Tr} \left(Q C_{\delta\rho}^2 (dH_t^{\text{DB}}) \right), \quad (4.32)$$

the condition becomes $\text{Tr} \left(Q C_{\delta\rho}^2 (dH_t^{\text{DB}}) \right) > 0$. We can satisfy this in expectation if

$$\text{Tr} \left(Q C_{\delta\rho}^2 \{ A^2 + [A, B]/2 - 2\text{Tr}(QA) A \} \right) < 0, \quad (4.33)$$

where the average $E(\cdot)$ is over all noise realizations. After replacing dH_t^{DB} by its definition, the condition becomes

$$\text{Tr} \left(Q C_{\delta\rho}^2 (A^2 + [A, B]/2) \right) < 0 \quad (4.34)$$

for all traceless $\delta\rho$ such that ρ is in the orbit of Q .

Since the diffusion potential $F_t^{(2)}(\sigma)$ in (4.29) tends to vanish as the state approaches the target Q , the variance in the trajectories becomes increasingly suppressed, increasing the chances that individual trajectories are trapped in the surroundings of Q . Condition (4.34) is equivalent to the stability condition that one derives in the context of Lyapunov stability (see App. 4.7.7).

The addition of the state-agnostic feedback with target Q introduces an additional fixed-point solution to the landscape of the potential, but at the same time it modifies it. For instance, when $B = 0$ and in the absence of feedback ($Q \equiv 0$), the stable fixed points of the dynamics correspond to the eigenspaces of A . If we switch on the feedback ($Q \neq 0$), we introduce a new fixed point equal to Q , but we also remove the eigenspaces of A from the set of fixed points unless $[Q, A] = 0$. This indicates that the effect of the feedback process in the potential landscape is to tilt the fixed points in favor of the target Q .

4.6 Discussion

There tends to be a tension between the unitary evolution of isolated quantum systems and measurement processes. The former describes smooth, differentiable dynamics, with a Hamiltonian well-determined by fields and interactions. In contrast, quantum measurements are typically presented as discontinuous jumps that break from Hamiltonian dynamics. Our results provide an alternative to the latter picture, one in which measurements are also described by stochastic Hamiltonians that, in turn, lead to double-bracket gradient flows. In this way, we unify double-bracket gradient flows, measurement processes, and Hamiltonians.

One could thus conceive of a picture where the collapse postulate is replaced by a ‘measurement Hamiltonian postulate’, in which measurements are described by H_{SB} and H_{DB} in Eqs. (4.11) and (4.12), with a stochastic measurement output as in Eq. (4.4). Alternatively, one could assume a gradient flow with the stochastic potentials defined in Eqs. (4.14) and (4.15).

Single- and double-bracket gradient flows describe the ensuing dynamics. In the pure-measurement case ($B = 0$ and $\|H_0\| \ll \{\|H^{SB}\|, \|H^{DB}\|\}$), the coarse-grained dynamics and statistics of measurement outcomes match the Born rule. Naturally, such approaches would still entail an uncomfortable pitfall shared with collapse-based interpretations: nothing in the theory says when a measurement occurs [253] (i.e., when dynamics is dictated by $H^{SB} + H^{DB}$ rather than H_0). In this respect, interpretations that leverage environmental decoherence to explain measurement occurrences benefit from criteria for defining when a measurement occurs [254]. However, interpretations based on environmental decoherence often suffer from the “*and/or problem*”, as no definite outcome is selected [253, 254]. In contrast, the stochastic measurement Hamiltonians introduced in this paper produce a single (random) outcome, the minimum of a gradient flow.

We have leveraged the generalized framework to describe measurement processes that inform novel quantum feedback protocols. In particular, we provided a feedback mechanism that results in deterministic dynamics for a double-bracket flow form. In such a process, the system ends up in the

ground state of a Hamiltonian. We also showed a state-agnostic feedback process that can prepare arbitrary states of a qubit. More generally, the framework allows interpreting feedback processes for state preparation as gradient flows with a feedback-modified stable fixed point.

Acknowledgments

A.V. acknowledges support from the U.S. Department of Energy (DOE) through a quantum computing program sponsored by the Los Alamos National Laboratory (LANL) Information Science & Technology Institute. This material is based upon work supported by the U.S. DOE, Office of Science, National Quantum Information Science Research Centers, Quantum Science Center, and the Beyond Moore's Law project of the Advanced Simulation and Computing Program at LANL.

4.7 Appendices

4.7.1 Monitored dynamics in the Itô picture

We provide a self-contained derivation of the SME (4.3) when the density operator describes pure states. The non-linear stochastic Schrödinger equation (SSE) for a single noise channel is given by

$$d\psi_t = -iH_e(t)\psi_t dt + (C - m_t)\psi_t dW_t, \quad (4.35)$$

where dW_t is such that $dW_t^2 = \gamma dt$ with $\gamma > 0$. We define the parameter $m_t := \psi_t^\dagger C^H \psi_t$, with $C^H = (C + C^\dagger)/2$ the Hermitian part of C , and the time-dependent operator

$$H_e(t) := H_0(t) - \frac{i}{2}\gamma(C^\dagger C - 2m_t C + m_t^2). \quad (4.36)$$

Note that the operator H_e is not Hermitian. Equation (4.35) satisfies $d\|\psi_t\|^2 = 0$, i.e., the norm of ψ_t is conserved deterministically. Equation (4.35) represents the dynamics of the state ψ_t conditioned on the measurement process y_t [40], which is characterized by the measurement output

$$dy_t = 2\gamma m_t dt + dW_t. \quad (4.37)$$

We compute the SME for the density operator $\rho_t = \psi_t \psi_t^\dagger$ by applying the Itô rule [215] $d\rho_t = d(\psi_t \psi_t^\dagger) = d\psi_t \psi_t^\dagger + \psi_t d\psi_t^\dagger + d\psi_t d\psi_t^\dagger$. We obtain

$$\begin{aligned} d\rho_t &= -iH_e(t)\rho_t dt + (C - m_t)\rho_t dW_t + \text{h.c.} + \gamma(C - m_t)\rho_t(C^\dagger - m_t)dt \\ &= -i(H_e(t)\rho_t - \rho_t H_e^\dagger(t))dt + ((C - m_t)\rho_t + \rho_t(C^\dagger - m_t))dW_t \\ &\quad + \gamma(C - m_t)\rho_t(C^\dagger - m_t)dt, \end{aligned} \quad (4.38)$$

where we used that $\rho_t^\dagger = \rho_t$. Consider the conjugate anticommutator between operators X, Y given by $\{X, Y\}_c := XY + Y^\dagger X^\dagger$ (which is just twice the Hermitian part of XY). Using this notation, we write the resulting dynamics as

$$d\rho_t = \mathcal{L}(\rho_t)dt + \{C - m_t, \rho_t\}_c dW_t \quad (4.39)$$

where the Lindbladian $\mathcal{L}(\rho_t) = -i[H_0, \rho_t] + \mathcal{D}[C](\rho_t)$, with dissipation part $\mathcal{D}[C](\rho_t) = \gamma(C\rho_t C^\dagger - \frac{1}{2}\{C^\dagger C, \rho_t\})$.

We can recast the diffusion term in (4.39) as a commutator by splitting the jump operator C in Hermitian and anti-Hermitian parts. Therefore,

$$\{C - m_t, \rho_t\}_c = [[C^H, \rho_t] + C^{AH}, \rho_t], \quad (4.40)$$

where $C^{AH} = (C - C^\dagger)/2$. This is easily shown by noting that $\rho_t^2 = \rho_t$ and $\rho_t C^H \rho_t = m_t \rho_t$. Using the previous equivalence we can write (4.39) as

$$d\rho_t = \mathcal{L}(\rho_t)dt - i[H_t, \rho_t]dW_t \quad (4.41)$$

where the second term is interpreted as a stochastic coherent action with

Hamiltonian H_t defined by $H_t := i([C^H, \rho_t] + C^{AH})$.

The generalization of (4.41) to n_c multiple noise channels is straightforward. Consider n_c Wiener processes $W_a(t)$ ($a = 1, \dots, n_c$) such that $dW_a(t)dW_b(t) = \delta_{ab}\gamma_a dt$ with $\gamma_a > 0$, and define n_c arbitrary jump operators C_a . The SSE is

$$d\psi_t = -iH_e(t)\psi_t dt + \sum_a (C_a - m_a(t))\psi_t dW_a(t). \quad (4.42)$$

with $H_e(t) = H_0(t) - \frac{i}{2} \sum_a \gamma_a (C_a^\dagger C_a - 2m_a(t)C_a + m_a^2(t))$. Then, the unraveling for the state ρ_t is given by

$$d\rho_t = \mathcal{L}(\rho_t)dt + \sum_a \{C_a - m_a(t), \rho_t\}_c dW_a(t) \quad (4.43)$$

where $\mathcal{L}(\rho_t) = -i[H_0, \rho_t] + \sum_a \mathcal{D}_a(\rho_t)$ with $\mathcal{D}_a(\rho_t) = \gamma_a (C_a \rho_t C_a^\dagger - \frac{1}{2} \{C_a^\dagger C_a, \rho_t\})$, and $m_a(t) = \psi_t^\dagger C_a^H \psi_t$.

4.7.2 Monitored dynamics in the Stratonovich picture

To write the SME (4.39) in the Stratonovich picture, we use the Itô-Stratonovich transformation from App. 4.7.3.

First, let's define $\sigma_t = \sigma(\rho_t) = \{C - m_t, \rho_t\}_c$, which depends on ψ_t only through ρ_t since $m_t \equiv \text{Tr}(\rho_t C^H)$.

Use the Itô-Stratonovich transformation (4.54) to write

$$\sigma_t^{ab} dW_t = \sigma_t^{ab} \circ dW_t - \frac{1}{2} \sigma_t^{cd} \partial_{cd} \sigma_t^{ab} \gamma dt \quad (4.44)$$

where the derivatives ∂_{cd} are w.r.t. the coordinates $(\rho_t)_{cd}$, and $\circ$ denotes the mid-point evaluation rule defining the Stratonovich picture. Repeated indices denote summation.

Note that the contraction $\sigma^{cd} \partial_{cd} \sigma_t$ can be written as the application of the linear operator $\partial_\rho \sigma_t$ on σ_t , where the partial derivative is taken component-

wise. Then,

$$\begin{aligned}
(\partial_P \sigma_t)(\sigma_t) &= (C - m_t)\sigma_t + \sigma_t(C^\dagger - m_t) - 2\text{Tr}(C^H \sigma_t) \rho_t \\
&= C\sigma_t + \sigma_t C^\dagger - 2m_t\sigma_t - 2\text{Tr}(C^H \sigma_t) \rho_t \\
&= \{C^2, \rho_t\}_c - 4m_t\{C, \rho_t\}_c + 2C\rho_t C^\dagger + 8m_t^2\rho_t - 2\langle\{C^H, C\}_c\rangle \rho_t,
\end{aligned} \tag{4.45}$$

where $\langle \bullet \rangle$ denotes the quantum expectation $\text{Tr}(\bullet \rho_t)$.

Inserting (4.45) into transformation (4.44) we obtain

$$\begin{aligned}
\sigma(\rho_t)dW_t &= \sigma(\rho_t) \circ dW_t - \frac{\gamma}{2}\{C^2, \rho_t\}_c dt + 2\gamma m_t\{C, \rho_t\}_c dt - \gamma C\rho_t C^\dagger dt \\
&\quad - 4\gamma m_t^2\rho_t dt + \gamma \langle\{C^H, C\}_c\rangle \rho_t dt
\end{aligned} \tag{4.46}$$

Using that $\sigma(\rho_t) = \{C, \rho_t\}_c - 2m_t\rho_t$ we can write

$$\begin{aligned}
\sigma(\rho_t)dW_t &= \sigma(\rho_t) \circ (2\gamma m_t dt + dW_t) - \frac{\gamma}{2}\{C^2, \rho_t\}_c dt - \gamma C\rho_t C^\dagger dt \\
&\quad + \gamma \langle\{C^H, C\}_c\rangle \rho_t dt.
\end{aligned} \tag{4.47}$$

Finally, by replacing (4.47) in the SME (4.39) given by

$$d\rho_t = \mathcal{L}(\rho_t)dt + \sigma(\rho_t)dW_t, \tag{4.48}$$

we obtain the dynamics in Stratonovich form

$$\begin{aligned}
d\rho_t &= -i[H_0, \rho_t]dt - \frac{\gamma}{2}\left(\{C^\dagger C, \rho_t\} + \{C^2, \rho_t\}_c - 2\langle\{C^H, C\}_c\rangle \rho_t\right) dt \\
&\quad + \sigma(\rho_t) \circ (2\gamma m_t dt + dW_t).
\end{aligned} \tag{4.49}$$

Using the definition of $\sigma(\rho_t)$ and $dy_t = 2\gamma m_t dt + dW_t$, we can write the SME conditioned on the measurement process,

$$\begin{aligned}
d\rho_t &= -i[H_0, \rho_t]dt - \frac{\gamma}{2}\left(\{C^\dagger C, \rho_t\} + \{C^2, \rho_t\}_c - 2\langle\{C^H, C\}_c\rangle \rho_t\right) dt \\
&\quad + \{C - m_t, \rho_t\}_c \circ dy_t.
\end{aligned} \tag{4.50}$$

For Hermitian $C = C^\dagger$, the equation above reduces to

$$d\rho_t = -i[H_0, \rho_t]dt - \gamma\{C^2, \rho_t\} + 2\gamma\langle C^2 \rangle \rho_t dt + \{C - m_t, \rho_t\} \circ dy_t \quad (4.51)$$

which matches the monitored dynamics derived in [212]. To see this, define $\gamma = \frac{1}{4\tau}$, with τ as in [212], and $dW_t = \sqrt{\gamma}dB_t$ with $dB_t^2 = dt$. The measurement process in [212] is defined as $r_t dt = m_t dt + \sqrt{\tau} dB_t$ with $m_t = \langle C \rangle$, while in our formalism it is defined as $dy_t = 2\gamma m_t dt + dW_t$. These definitions are equivalent by setting $dy_t = \frac{r_t}{2\tau} dt$.

4.7.3 Itô-Stratonovich conversion formula

Consider the following complex-valued SDE in Stratonovich form

$$dY_a(t) = f(t, X_t) \circ dW_a(t) \quad (4.52)$$

where $dW_a(t)$, ($a = 1, \dots, n_c$) are real Wiener increments satisfying $dW_a(t)dW_b(t) = \Gamma_{ab}dt$ with $\Gamma > 0$ a real symmetric matrix. The process $X_t \in \mathbb{C}^N$ satisfies $dX_t = \mu(t, X_t)dt + \sigma(t, X_t)dW_t$ with $\mu : [0, T] \times \mathbb{C}^N \rightarrow \mathbb{C}^N$ and $\sigma : [0, T] \times \mathbb{C}^N \rightarrow \mathbb{C}^{N \times n_c}$, and dW_t is the vector of Wiener increments. We assume σ to be a smooth function so it allows a Taylor expansion. The symbol $\circ$ in (4.52) means mid-point evaluation of $f(t, X_t)$, and defines the Stratonovich convention for stochastic integrals [215].

We can derive the Itô form of (4.52) by applying the definition of mid-point evaluation followed by a stochastic Taylor expansion of f [211]:

$$\begin{aligned} dY_a(t) &= f(t, X_t) \circ dW_a(t) \\ &= f(t + dt/2, X_t + dX_t/2) dW_a(t) \\ &= \left(f(t, X_t) + \frac{1}{2} \partial_t f(t, X_t) dt + \frac{1}{2} \partial_b f(t, X_t) dX_t^b + \dots \right) dW_a(t) \\ &= f(t, X_t) dW_a + \frac{1}{2} \partial_b f(t, X_t) dX_t^b dW_a \\ &= f(t, X_t) dW_a(t) + \frac{1}{2} \partial_b f(t, X_t) \sigma^{bc}(t, X_t) dW_c(t) dW_a(t) \end{aligned} \quad (4.53)$$

Using $dW_c(t)dW_a(t) = \Gamma_{ca}dt$ we obtain the desired Itô-Stratonovich transformation in vector-matrix form

$$f(t, X_t) \circ dW(t) = f(t, X_t)dW(t) + \frac{1}{2} \nabla f(t, X_t) \sigma_t(t, X_t) \Gamma dt \quad (4.54)$$

4.7.4 A stochastic Hamiltonian that describes monitored dynamics

Since we work in the Stratonovich picture, we omit from now on the $\circ$ notation.

Note that the last term in (4.50) can be written as a commutator with the state ρ_t using the identity $\{C - m_t, \rho_t\}_c = [[C^H, \rho_t] + C^{AH}, \rho_t]$ [see discussion around (4.40)]. Therefore, we write

$$\{C - m_t, \rho_t\}_c dy_t = -i[H_M(t), \rho_t]dt \quad (4.55)$$

where we define the measurement Hamiltonian $H_M(t)$ by

$$H_M(t) = i \left([C^H, \rho_t] + C^{AH} \right) \frac{dy_t}{dt} \quad (4.56)$$

The term proportional to γ in (4.50) can be written in terms of double-brackets as follows. Note that, for an arbitrary operator Ω , the following identity holds:

$$[[\Omega, \rho_t], \rho_t] = \{\Omega, \rho_t\} - 2 \langle \Omega \rangle \rho_t. \quad (4.57)$$

In particular, for $\Omega = \{C^H, C\}_c$ we have

$$\begin{aligned} [[\{C^H, C\}_c, \rho_t], \rho_t] &= \frac{1}{2} (\{C^2, \rho_t\}_c + \{\rho_t, C^2\}_c) + \{C^\dagger C, \rho_t\} - 2 \langle \{C^H, C\}_c \rangle \rho_t \\ &= \frac{1}{2} (\{\rho_t, C^2\}_c - \{C^2, \rho_t\}_c) + \{C^2, \rho_t\}_c + \{C^\dagger C, \rho_t\} \\ &\quad - 2 \langle \{C^H, C\}_c \rangle \rho_t. \end{aligned} \quad (4.58)$$

Rearranging,

$$\begin{aligned} & [[\{C^H, C\}_c, \rho_t], \rho_t] - \frac{1}{2} (\{\rho_t, C^2\}_c - \{C^2, \rho_t\}_c) = \\ & \{C^2, \rho_t\}_c + \{C^\dagger C, \rho_t\} - 2 \langle \{C^H, C\}_c \rangle \rho_t. \end{aligned} \quad (4.59)$$

Writing $\{\rho_t, C^2\}_c - \{C^2, \rho_t\}_c = [\rho_t, C^2] + [(C^\dagger)^2, \rho_t]$ we obtain

$$\begin{aligned} & [[\{C^H, C\}_c, \rho_t] + \frac{1}{2} (C^2 - (C^\dagger)^2), \rho_t] = \\ & \{C^2, \rho_t\}_c + \{C^\dagger C, \rho_t\} - 2 \langle \{C^H, C\}_c \rangle \rho_t \end{aligned} \quad (4.60)$$

Note that the r.h.s. of the equation above is identical to the term proportional to the damping rate γ in (4.50). Define the dissipation Hamiltonian

$$H_D(t) = -i \frac{\gamma}{2} \left([\{C^H, C\}_c, \rho_t] + \frac{1}{2} (C^2 - (C^\dagger)^2) \right) \quad (4.61)$$

Collecting the Hamiltonians H_D (4.61) and H_M (4.56), we can write the monitored dynamics (4.50) as a stochastic Liouville equation

$$d\rho_t = -i[H_0 + H_D(t) + H_M(t), \rho_t]dt. \quad (4.62)$$

Because ρ_t is pure, the equation above corresponds to a stochastic unitary evolution with propagator in the form of a Dyson-Stratonovich time-ordered exponential

$$U = \mathcal{T} \exp \left\{ -i \int_0^T (H_0 + H_D(t) + H_M(t)) dt \right\}. \quad (4.63)$$

Note that there are no Itô corrections in the expression above as we are working in the Stratonovich picture, where the common rules of calculus apply likewise.

Let's write (4.62) explicitly in terms of the Hermitian and anti-Hermitian parts of C . Expanding (4.62) using the Hamiltonian definitions above, we

obtain

$$d\rho_t = -i[H_0, \rho_t]dt - \frac{\gamma}{2} \left[[\{C^H, C\}_c, \rho_t], \rho_t \right] dt - \frac{\gamma}{4} [C^2 - (C^\dagger)^2, \rho_t] dt \\ + [[C^H, \rho_t], \rho_t]dy_t + [C^{AH}, \rho_t]dy_t \quad (4.64)$$

Denote $A := C^H$ and $B := C^{AH}$. Using the identities $\{C^H, C\}_c = 2A^2 + [A, B]$ and $C^2 - (C^\dagger)^2 = 2\{A, B\}$, the dynamics in terms of A and B is given by

$$d\rho_t = -i[H_0, \rho_t]dt - \gamma[A^2, \rho_t]dt + [[A, \rho_t], \rho_t]dy_t \\ - \frac{\gamma}{2} [[[A, B], \rho_t], \rho_t]dt - \frac{\gamma}{2} [\{A, B\}, \rho_t]dt + [B, \rho_t]dy_t \quad (4.65)$$

Grouping in single- and double-bracket contributions we prove Theorem 4 in the main text. The generalization to multiple noise channels is straightforward following the definitions of App. 4.7.1.

4.7.5 Double-bracket gradient flows

Mathematical background

It is well known that double-bracket flows can be interpreted as Riemannian gradient flows in the unitary orbit [41–43]. In order to prove Theorem 5 in the main text, we provide here a simplified derivation based on [43], invoking minimal concepts from differential geometry [218] and Lie group theory [217, 219].

The Lie group $\mathfrak{u}(N)$ of the unitary group $U(N)$ is given by all anti-Hermitian matrices $X^\dagger = -X$. Define a Riemannian metric on the unitary group $U(N)$ given by the Hilbert-Schmidt (HS) inner product $\langle X, Y \rangle_{\text{HS}} = \text{Tr}(X^\dagger Y)$, where $X, Y \in \mathfrak{u}(N)$.

Given a state ρ , define the action map $\phi_\rho : U(N) \rightarrow \mathcal{O}_\rho$ such that $\phi_\rho(U) = U\rho U^\dagger$. Consider the tangent space to the unitary orbit composed of all commutations of the form $[X, \rho]$ with X in the Lie algebra $\mathfrak{u}(N)$ [217, 219]. The action map ϕ_ρ , via its differential map $D\phi_\rho$, induces a (normal) metric




Figure 4.6 The derivative of the action map ϕ_ρ maps tangent vectors on the unitary group—in this case, at the identity—to tangent vectors on the orbit of ρ . This induces a HS metric on the state’s orbit such that $\langle X, Y \rangle_{\text{HS}} = \langle [X, \rho], [Y, \rho] \rangle_{\mathcal{O}}$.

$\langle \bullet, \bullet \rangle_{\mathcal{O}}$ in the unitary orbit [218], such that for two arbitrary tangent vectors $[X, \sigma]$ and $[Y, \sigma]$ at a point σ in the orbit, we have $\langle [X, \sigma], [Y, \sigma] \rangle_{\mathcal{O}} = \langle X, Y \rangle_{\text{HS}}$. Strictly speaking, X, Y are orthogonal to the null directions of σ , that is, to the vectors V in the algebra that do not generate changes in the state, i.e., $[V, \sigma] = 0$. This restriction allows identifying one-to-one tangent vectors X on the unitary group with tangent vectors $[X, \sigma]$ on the state’s orbit. See Fig. 4.6 for a schematic illustration of the various concepts introduced so far.

Consider an arbitrary Hermitian operator O . Define the scalar function $F(\sigma) = \text{Tr}(O\sigma)$ where σ is an arbitrary state in the unitary orbit of ρ . As a function of σ , $F(\sigma)$ admits a differential $DF|_\sigma$ acting in the tangent space of the orbit of ρ , such that $DF|_\sigma([X, \sigma]) \in \mathbb{R}$ for arbitrary anti-Hermitian X in the Lie algebra. The induced metric allows us to define a unique representative vector of the differential DF in the tangent space of the orbit. It defines a Riemannian gradient induced by the HS metric, and we denote it by $\nabla F(\sigma)$. To derive the form of the $\nabla F(\sigma)$ we invoke the compatibility condition just described. Since $\nabla F(\sigma)$ is a tangent vector in the unitary orbit, it has a form $[Z, \sigma]$ for some Z in the Lie algebra. Then, for an arbitrary X in the Lie algebra

we have

$$\begin{aligned}
\langle \nabla F(\sigma), [X, \sigma] \rangle_{\mathcal{O}} &\equiv DF|_{\sigma}([X, \sigma]) \\
\langle Z, X \rangle_{\text{HS}} &= \text{Tr}(O[X, \sigma]) \\
-\text{Tr}(ZX) &= \text{Tr}([\sigma, O]X)
\end{aligned} \tag{4.66}$$

where in the second line we used the compatibility condition of the induced metric and that $DF|_{\sigma}([X, \sigma]) = \frac{d}{d\tau}(F(\sigma(\tau)))|_{\tau=0} = \text{Tr}(O[X, \sigma])$, where $\sigma(\tau) := e^{\tau X} \sigma e^{-\tau X}$ is the integral curve corresponding to the tangent vector $[X, \sigma]$ [217–219]. In the third line we used the definition of the HS product and the trace property. Because the last line holds for arbitrary X , we must have $Z = [O, \sigma]$. Thus, we arrive at the final form of the gradient

$$\nabla F(\sigma) = [Z, \sigma] = [[O, \sigma], \sigma]. \tag{4.67}$$

Now that we established the main connection between double-brackets and gradient flows, we apply it to prove the results of Section 4.4.

Gradient flow formulations

The relation between double-brackets and gradient flows derived in the previous section implies the following. Suppose a general double-bracket dynamics of the form

$$d\rho_t = \mathcal{C}_{\rho}^2(dO_t) = [[dO_t, \rho_t], \rho_t], \tag{4.68}$$

where $\mathcal{C}_{\rho}$ is the commutator with the state, and dO_t is some Hermitian generator process. We write its gradient flow formulation as

$$d\rho_t = \nabla F_t(\rho_t) dt, \tag{4.69}$$

where $F_t(\sigma) := \text{Tr}(\dot{O}_t \sigma)$ with $\dot{O}_t := dO_t/dt$, and σ in the state's orbit.

In the main text, we stated that pure states satisfy

$$d\rho_t = [[d\rho_t, \rho_t], \rho_t] . \quad (4.70)$$

Setting $dO_t \equiv d\rho_t$ and using the gradient flow formulation given by (4.68)–(4.69), it is immediate to write the relation above as a gradient flow

$$\dot{\rho}_t = \nabla F_t(\rho_t) \quad (4.71)$$

where the potential $F_t(\sigma) = \text{Tr}(\dot{\rho}_t \sigma)$.

Along the same lines, we prove Theorem 5 in the main text as follows. The full dynamics (4.10) can be written as

$$\begin{aligned} d\rho_t &= \mathcal{C}_\rho(-idH_t^{\text{SB}}) + \mathcal{C}_\rho^2(dH_t^{\text{DB}}) \\ &= \mathcal{C}_\rho^2\left(\mathcal{C}_\rho(-idH_t^{\text{SB}}) + dH_t^{\text{DB}}\right) \\ &= \nabla F_t(\rho_t)dt . \end{aligned} \quad (4.72)$$

In the second line we used the periodic property $\mathcal{C}_\rho^3 = \mathcal{C}_\rho$. The third line is obtained by setting $dO_t \equiv \mathcal{C}_\rho(-idH_t^{\text{SB}}) + dH_t^{\text{DB}}$ and using the gradient flow formalism derived before. By splitting F_t into single- and double-bracket terms we arrive at the desired result.

The evolution equation of the variance $V(t)$

We derive Eq. (4.18). Consider $V(t) = \text{Tr}(\Delta A_t^2 \rho_t)$ as defined in the text and compute

$$dV(t) = \text{Tr}(d\Delta A_t^2 \rho_t + \Delta A_t^2 d\rho_t) = \text{Tr}(\Delta A_t^2 d\rho_t) , \quad (4.73)$$

where we used that $\text{Tr}(d\Delta A_t^2 \rho_t) = 0$. The increment $d\rho_t$ is a tangent vector on the state's orbit. Then, we can equivalently write the r.h.s. of (4.73) as the

differential map of $V_t(\sigma) = \text{Tr}(\Delta A_t^2 \sigma)$ applied to $d\rho_t$, that is,

$$\begin{aligned} dV(t) &= DV_t(d\rho_t) = \langle \nabla V_t(\rho_t), d\rho_t \rangle_{\mathcal{O}} \\ &= -\gamma \|\nabla V_t(\rho_t)\|_{\mathcal{O}}^2 dt + \langle \nabla V_t(\rho_t), \nabla K_t(\rho_t) \rangle_{\mathcal{O}} dW_t, \end{aligned} \quad (4.74)$$

where in the last two steps we used the compatibility condition of the metric and the dynamics $d\rho_t = -\gamma \nabla V_t(\rho_t) dt + \nabla K(\rho_t) dW_t$.

4.7.6 A simple proof of Theorem 6

Let $H_0 := \sum_j E_j \Pi_j$ with Π_j the j -th eigen-projector of H_0 . From Eq. (4.19),

$$\frac{d}{dt} \langle \Pi_j \rangle = -2\gamma (\text{Tr}(H_0^2 \rho_t \Pi_j) - \text{Tr}(H_0^2 \rho_t \Pi_j \rho_t)) = -2\gamma \langle \Pi_j \rangle (E_j^2 - \langle H_0^2 \rangle). \quad (4.75)$$

This equation implies that the populations of all levels with energy $E^2 \leq \langle H_0^2 \rangle$ increase in time (which, in turn, leads to a decrease in $\langle H_0^2 \rangle$). Moreover, the rate at which the population of the minimum energy eigenstate increases is the highest. Thus, the lowest energy eigenstate is the only stable fixed point of the dynamics.

Note that if $\langle \Pi_0 \rangle \equiv 0$ we cannot guarantee the exponential convergence to the ground state of H_0 . As a consequence, a sufficient condition for exponential convergence is that the initial state has non-zero overlap with the ground state. This is in agreement with literature [44].

4.7.7 Stability analysis of the measurement-based feedback dynamics

Consider the feedback-induced dynamics (4.20), which we reproduce here as

$$d\rho_t = -\gamma [[A^2, \rho_t - Q], \rho_t] dt - \frac{\gamma}{2} [[[A, B], \rho_t - Q], \rho_t] dt + [[A, \rho_t - Q], \rho_t] dy_t. \quad (4.76)$$

Given the target $Q = |q\rangle\langle q|$, we want to analyze the stability of the fixed-point solution Q . That is, whether small perturbations around the solution tend to converge to the solution or escape it's surroundings. For this, we study the behavior of the overlap $Q_{\parallel}(t) := \text{Tr}(Q\rho_t)$ by computing $dQ_{\parallel}(t) = \text{Tr}(Qd\rho_t)$. Note that we can write

$$dQ_{\parallel}(t) = \sum_Z dQ_{\parallel}^Z(t), \quad (4.77)$$

where

$$dQ_{\parallel}^Z(t) = \text{Tr}([dZ_t, \rho_t - Q], \rho_t)Q = \text{Tr}([dZ_t, \rho_t][\rho_t, Q]) - \text{Tr}([dZ_t, Q][\rho_t, Q]), \quad (4.78)$$

for Hermitian $dZ_t \in \{-\gamma A^2 dt, -\frac{\gamma}{2}[A, B]dt, A dy_t\}$.

For clarity of exposition we adopt the Dirac bracket notation. The first trace above is $\text{Tr}([dZ_t, \rho_t][\rho_t, Q]) = \langle \{dZ_t, Q\} \rangle_t - 2 \langle dZ_t \rangle_t Q_{\parallel}$, where $\langle \cdot \rangle_t$ is the quantum expectation w.r.t. $|\psi_t\rangle$. The second trace is $\text{Tr}([dZ_t, Q][\rho_t, Q]) = 2 \langle dZ_t \rangle_q Q_{\parallel} - \langle \{dZ_t, Q\} \rangle_t$, where $\langle \cdot \rangle_q$ is the quantum expectation w.r.t. $|q\rangle$. Thus,

$$dQ_{\parallel}^Z(t) = 2 \langle \{dZ_t, Q\} \rangle_t - 2(\langle dZ_t \rangle_t + \langle dZ_t \rangle_q)Q_{\parallel}. \quad (4.79)$$

We can always parametrize $|\psi_t\rangle = \sqrt{Q_{\parallel}(t)}|q\rangle + \sqrt{Q_{\perp}(t)}|q_{\perp}(t)\rangle$, where $Q_{\perp} := 1 - Q_{\parallel}$ and $|q_{\perp}\rangle$ is a unit vector resulting from projecting $|\psi_t\rangle$ onto the orthogonal space of $|q\rangle$. When the projection is null, we simply set $|q_{\perp}\rangle = 0$.

We use this decomposition to write (4.79) as

$$dQ_{\parallel}^Z(t) = 2Q_{\parallel}Q_{\perp} \left(\langle dZ_t \rangle_q - \langle dZ_t \rangle_{q_{\perp}} + 2\sqrt{\frac{Q_{\perp}}{Q_{\parallel}}} \operatorname{Re} \langle q | dZ_t | q_{\perp} \rangle \right) \quad (4.80)$$

From this we see, as expected, that $Q_{\parallel} = 1$ is a fixed point of the dynamics, but also $Q_{\parallel} = 0$, meaning that when the system enters the orthogonal subspace of $|q\rangle$, it remains there for the rest of the evolution.

To analyze the stability of $Q_{\parallel} = 1$, consider the system state is subject to a small perturbation around the solution $|q\rangle$. Thus, we can write the overlap $Q_{\parallel} = 1 - \varepsilon$ and $Q_{\perp} = \varepsilon$ for some $\varepsilon > 0$ small. Then, we write the linearization in ε of (4.77) as

$$\begin{aligned} d\varepsilon = & 2\gamma\varepsilon \left\{ \sum_Z \left(\langle Z \rangle_q - \langle Z \rangle_{q_{\perp}} \right) - 2 \left(\langle A \rangle_q - \langle A \rangle_{q_{\perp}} \right) \langle A \rangle_q \right\} dt \\ & - 2\varepsilon \left(\langle A \rangle_q - \langle A \rangle_{q_{\perp}} \right) dW_t \end{aligned} \quad (4.81)$$

where $Z \in \{A^2, \frac{1}{2}[A, B]\}$.

Note that $d\varepsilon^2 \propto \varepsilon^2 dt$, that is, the variance in the overlap tends to vanish close to the boundary $Q_{\parallel} = 1$. This means it becomes increasingly harder for the system to leave the surroundings of $|q\rangle$.

Recall we are dealing with Stratonovich equations. For this solution to be (stochastically asymptotically) stable [211] the Lyapunov exponent associated to the linear equation (4.81) must be negative. This implies that

$$\sum_{Z \in \{A^2, \frac{1}{2}[A, B]\}} \left(\langle Z \rangle_q - \langle Z \rangle_{q_{\perp}} \right) - 2 \left(\langle A \rangle_q - \langle A \rangle_{q_{\perp}} \right) \langle A \rangle_q < 0. \quad (4.82)$$

Given a target $|q\rangle$, we can engineer suitable A, B such that the inequality above is met. Conversely, if A, B are given, we can study the stability regions of the feedback dynamics (4.76) by solving the inequality (4.82).

In the case of multiple n_c noise channels the condition above generalizes

to

$$\sum_{a=1}^{n_c} \gamma_a \left(\sum_{Z_a} \left(\langle Z_a \rangle_q - \langle Z_a \rangle_{q\perp} \right) - 2 \left(\langle A_a \rangle_q - \langle A_a \rangle_{q\perp} \right) \langle A_a \rangle_q \right) < 0, \quad (4.83)$$

where $Z_a \in \{A_a^2, \frac{1}{2}[A_a, B_a]\}$ for the a -th noise channel with jump operator C_a .

Let's apply the previous analysis to the single-qubit example of the main text. The jump operators are σ_+ , σ_- with damping rates γ_+ , γ_- , respectively. Decompose the jump operators in Hermitian and anti-Hermitian parts as $A_+ = A_- = \frac{\sigma_x}{2}$ and $B_{\pm} = \pm \frac{i\sigma_y}{2}$. The inequality (4.83) implies

$$-\frac{\Delta\gamma}{4} (\langle \sigma_z \rangle_q - \langle \sigma_z \rangle_{q\perp}) - (\gamma_+ + \gamma_-) (\langle \sigma_x \rangle_q - \langle \sigma_x \rangle_{q\perp}) \langle \sigma_x \rangle_q < 0 \quad (4.84)$$

where $\Delta\gamma = \gamma_+ - \gamma_-$. Writing the target state in polar form as $|q\rangle = \cos(\frac{\theta}{2})|0\rangle + e^{i\phi} \sin(\frac{\theta}{2})|1\rangle$, and the (unique up to a global phase) orthogonal state as $|q\perp\rangle = \sin(\frac{\theta}{2})|0\rangle - e^{i\phi} \cos(\frac{\theta}{2})|1\rangle$ and computing the expectations in (4.84), we arrive to the desired condition stated in the main text,

$$\Delta\gamma \cos(\theta) + 2(\gamma_+ + \gamma_-) \sin^2(\theta) \cos^2(\phi) > 0. \quad (4.85)$$

Conclusion

In Chapter 2, we demonstrated that quantum annealing algorithms employing instance-agnostic driver Hamiltonians can yield, at most, a quadratic speedup in solving general combinatorial optimization problems. This observation is consistent with the established understanding in quantum computing that achieving super-polynomial quantum speedups typically necessitates exploiting structural properties specific to the problem instance [109]. Moreover, we show that building the schedule that achieves quadratic speedup requires precise knowledge of the density of states of the target Hamiltonian, and this scales exponentially with the number of qubits. Moreover, resolving the location of the minimum gap requires exponential precision as the number of qubits increase. These no-go results are rigorously proven for a drive Hamiltonian corresponding to the equal superposition state projector. Furthermore, we provide compelling numerical evidence that the same holds for transverse Hamiltonians. To date, no practical realizations of quadratic speedups exist in the literature of quantum annealing, and the prospect for it seems elusive due to the aforementioned complexity caveats. Although there is evidence of algorithmic advantage against a state-of-the-art classical heuristic [255], no conclusive experimental demonstration of quadratic speedup is available to date.

Structured driver Hamiltonians in quantum annealing have been shown to enhance performance metrics such as minimum gap scaling, success

probabilities, and effective time-to-solution on hard optimization problems and structured tasks [256–258]. These results support the view that moving beyond instance-agnostic transverse-field drivers and exploiting non-stoquasticity, non-locality, or problem structure is a promising route to achieving speedups that surpass the Grover-like quadratic regime in relevant problem families [256–258].

In Chapter 3, we introduced Path Integral Quantum Control (PiQC), a novel framework for optimal open-loop control of quantum systems whose dynamics are described by the Lindblad master equation. By combining path integral control techniques with stochastic quantum trajectories, PiQC yields an optimization scheme for computing open-loop controls that minimize a prescribed cost functional and can be applied in a highly generic manner, for instance to quantum state preparation, as in the present text, or ground-state energy estimation [259].

Not all Lindblad evolutions are amenable to simulation within the PiQC framework, since the corresponding Lindblad operators must lie in a subspace spanned by anti-Hermitian operators. Nevertheless, this requirement does not constitute a severe limitation in practice, as we show that a broad class of practically relevant control tasks can be formulated and solved within PiQC. Furthermore, the annealing variant of PiQC determines optimal open-loop controls for closed quantum systems, where the environmental couplings serve as an auxiliary mechanism that facilitates the optimization procedure. In this case, the anti-Hermitian constraint does not represent a problem as the environment is progressively switched off across the optimization.

PiQC provides a foundation for developing more general quantum control theories that integrate path integral control with quantum trajectories. In particular, extending this framework to incorporate quantum measurements offers an immediate and promising research direction, while its application to the optimization of parametrized quantum circuits—recently explored in a publication co-authored by the author of this thesis [259]—illustrates its potential in quantum algorithm design. Other promising avenues for

extending PiQC include considering noise models that go beyond white noise, such as Lévy processes [260] or colored noise [166, 186]. Other possibilities are exploring PI control beyond affine noise [190].

In Chapter 4, we derive a Hamiltonian formulation for the dynamics of continuously monitored open quantum systems. Under continuous measurement by the environment, the system's state evolves stochastically along trajectories in Hilbert space that connect pure states, driven by the randomizing effect of quantum jumps and diffusive measurement records. We obtain an equation in the Stratonovich picture that governs the conditioned (post-measurement) state and demonstrate that this evolution admits an equivalent stochastic Hamiltonian description.

The Hamiltonian formulation admits an interpretation in terms of double-bracket flows. Before this work, double-bracket flows had appeared in the quantum-information literature solely through quantum algorithmic applications to diagonalization and ground-state preparation [44, 45]. In this work, we establish for the first time an explicit connection between double-bracket flows and the dynamics of monitored quantum systems, which in turn motivates novel coherent feedback-control schemes for state preparation. In general, optimal feedback protocols rely on knowledge of the system state to determine the appropriate control action, a requirement that is often impractical because it entails running a parallel state-estimation procedure. Our feedback blueprint circumvents this limitation by dispensing with explicit state estimation and instead designing the dynamics so that the desired target state is realized as a fixed point. The resulting trade-off is that, for a given environment, one must identify regions of Hilbert space containing fixed points that are stable, in the sense of being robust against perturbations. This consideration is particularly relevant in experimental settings: if the target state is unstable under the current configuration but there is flexibility in the measurement scheme, the environment can be engineered to convert an unstable fixed point into a stable one, as exemplified by the single-qubit case discussed in this work.

The availability of double-bracket flow formulations of quantum dynamics

opens a promising research direction with applications that extend beyond quantum control. In particular, such flows can be leveraged to design error-correction and error-mitigation strategies for analog quantum computers, a line of work that is currently under development by the author of this thesis.

Chapter 6

Nederlandse samenvatting

Dit proefschrift behandelt twee kernproblemen binnen de kwantuminformaticawetenschap: de complexiteit van quantum annealing voor combinatorische optimalisatie en de controle van open en gemonitorde kwantumsystemen. Het combineert ideeën uit quantum computing, stochastische processen en regeltheorie om zowel de beperkingen als het potentieel van kwantumalgoritmen in realistische settings te onderzoeken.

Ten eerste analyseert het proefschrift adiabatische quantum annealing voor problemen zoals 3-SAT. Het laat zien dat een kwadratische quantumversnelling in principe mogelijk is, maar alleen wanneer het annealingschema wordt afgestemd met behulp van gedetailleerde spectrale informatie die voor moeilijke instanties in het algemeen niet tractabel te berekenen is, waardoor praktische versnelling voor deze klasse algoritmen onwaarschijnlijk is.

Ten tweede introduceert het proefschrift Path-integral Quantum Control (PIQC), een nieuw raamwerk voor de controle van open kwantumsystemen met behulp van stochastische quantumtrajecten en adaptieve importance sampling. Deze methode is van nature paralleliseerbaar, toepasbaar op een brede klasse van Lindblad-dynamica, en presteert competitief in numerieke experimenten, waaronder ruisachtige NMR-omgevingen. Daarnaast

ontwikkelt het proefschrift een op annealing gebaseerde variant van PiQC voor de controle van gesloten systemen, die controleoplossingen van hoge kwaliteit oplevert en in de gerapporteerde voorbeelden beter presteert dan benchmarks op basis van GRAPE.

Ten slotte ontwikkelt het proefschrift een Hamiltoniaanse formulering van continu gemonitorde kwantumsystemen, legt het een verband tussen gemonitorde dynamica en double-bracket-gradientstromen, en leidt het feedbackcontroleprotocollen af voor toestandspreparatie en stabilisatie.

In het geheel levert het proefschrift enerzijds een rigoureuze complexiteitsanalyse die de beperkingen van adiabatische quantum annealing verduidelijkt, en anderzijds nieuwe stochastische en geometrische methoden voor de controle van open kwantumsystemen. Samen versterken deze resultaten de theoretische grondslagen van kwantumoptimalisatie, kwantumcontrole en door metingen aangedreven dynamica.

Research data management

7.1 Scope, ethics and governance

This thesis is based on theoretical and computational research in quantum information, quantum annealing, and quantum control. The research output consists primarily of mathematical derivations, numerical simulations, analysis scripts, and research software, rather than human-participant or clinical data. Data and software management were carried out in line with the FAIR principles, with the aim of ensuring that the results are findable, accessible, interoperable, and reusable.

No human participants, personal data, or animal data were involved in the research presented in this thesis. Consequently, no ethical approval, informed consent procedure, pseudonymisation key, or data-access restriction on privacy grounds was required.

Research was conducted at the Donders Institute, Faculty of Science, Radboud University, and at Los Alamos National Laboratory. Financial support is acknowledged in the thesis front matter, including support through a U.S. Department of Energy quantum-computing program sponsored by the LANL Information Science & Technology Institute. Research-data management followed the FAIR principles and the applicable requirements of

Radboud University and the Donders Institute. Since the work contains no personal or confidential participant data, the research materials described below can be shared openly, subject only to third-party software licences and any institutional or publisher restrictions applicable to a specific software release.

7.2 Ethics, privacy, and funding

No human-participant, patient, personal, animal, or otherwise sensitive data were collected, generated, or processed in the research reported in this thesis. Consequently, GDPR-related privacy risks, informed-consent requirements, and pseudonymisation procedures are not applicable.

The research was carried out at the Donders Institute, Radboud University, and at Los Alamos National Laboratory. Financial support is acknowledged in the thesis and in the relevant publications.

7.3 Findability and accessibility

1. Chapter 2: Why Adiabatic Quantum Annealing is unlikely to yield speed-up

This chapter is based on the published article:

A. Villanueva, P. Najafi, and H. J. Kappen, "Why adiabatic quantum annealing is unlikely to yield speed-up," *Journal of Physics A: Mathematical and Theoretical* **56**, 465304 (2023).

The article is available through the publisher at <https://doi.org/10.1088/1751-SB8121/ad0439>. A preprint version is available at <https://arxiv.org/abs/2212.13649>.

Chapter 2 is primarily theoretical. Its main research outputs are the analytical derivations, proofs, definitions of the annealing model, and the numerical illustrations reported in the published article and its

appendices. The article specifies the Hamiltonians, the construction of the 3-SAT and 3-spin models, numerical procedures, parameter choices, and the resulting spectral-gap calculations.

No experimental, observational, personal, or sensitive data were acquired. No separate public code repository or standalone computational-data collection was created for this work. The published article and its openly available preprint constitute the available public research record for this chapter.

2. **Chapter 3: Stochastic optimal control of open quantum systems**

This chapter is based on the article:

A. Villanueva and H. J. Kappen, "Path integral control of open quantum systems," *Physical Review A* **114**, 012426 (2026).

A preprint version is available at <https://arxiv.org/abs/2410.18635>.

The code supporting this chapter is publicly available in the PiQC repository:

<https://github.com/aaronuv/piqc>

The repository contains MATLAB implementations of the Path integral Quantum Control algorithm used in the work. It provides the computational implementation underlying the numerical examples of noisy-qubit control, NMR control, annealed diffusion control, robustness calculations, and multi-qubit examples.

No experimental, observational, personal, or sensitive data were acquired. The computational outputs consist of generated quantum trajectories, optimisation results, control pulses, convergence statistics, and data used to create the figures. These outputs can be regenerated from the publicly available PiQC code using the model definitions, parameters, and numerical procedures reported in the article and thesis. Large intermediate trajectory ensembles are transient simulation outputs and are not maintained as a separate archival dataset.

3. **Chapter 4: Hamiltonian and Double-Bracket Flow Formulations of Quantum Measurements**

This chapter is based on the published article:

A. Villanueva and L. P. García-Pintos, "Hamiltonian and double-bracket flow formulations of quantum measurements," *Reports on Progress in Physics* (2026).

The published article is available through the publisher at <https://doi.org/10.1088/1361-SB6633/ae6fe9>. A preprint version is available at <https://arxiv.org/abs/2512.15412>.

This chapter is primarily theoretical. Its principal outputs are analytical derivations of the stochastic Hamiltonian formulation of continuous quantum measurements, the double-bracket flow formulation, theoretical results on feedback control, and the stability analysis. Any numerical illustrations and their defining model parameters are described in the article and thesis.

No experimental, observational, personal, or sensitive data were acquired. No separate public code repository or standalone computational-data collection is associated with this chapter. The published article and openly available preprint constitute the public research record for the theoretical results and the numerical illustrations.

7.4 Interoperability and reusability

The publications and preprints are available in PDF format and contain the mathematical definitions, analytical derivations, model parameters, numerical methods, and references needed to interpret and independently implement the reported work.

For Chapter 3, the PiQC source code is provided as MATLAB source files in a public GitHub repository. The repository and associated article document

the algorithmic implementation, model configurations, and numerical examples. Users of the code should cite the associated publication and respect the licence specified in the GitHub repository.

For Chapters 2 and 4, reproducibility relies on the complete mathematical specification of the models, derivations, numerical procedures, and reported results in the peer-reviewed articles and thesis. No third-party personal or restricted data are involved. Any third-party software or external resources used to reproduce calculations remain subject to their own licence terms.

Chapter 8

Curriculum Vitae

Aarón Villanueva

Email: aaronuzv@gmail.com, aaron.villanueva@donders.ru.nl

Website: aaronuv.github.io/aaronvillanueva.github.io

GitHub: github.com/aaronuv

Theoretical physics researcher with expertise in quantum computing, quantum control and optimization, and scientific computing. Develops quantum algorithms and optimal control methods for open quantum systems, with a focus on path integral, feedback control, and geometric optimization approaches. Brings strong scientific software experience in Python, Julia, and MATLAB, including open-source development for computational research. Proven ability to work across academic and national-laboratory environments and translate advanced theory into publishable, high-impact results.

Education

PhD in Physics

Radboud University, Nijmegen, The Netherlands

2022–2026

Master Physics Degree (Cum Laude)

National University of Córdoba, Córdoba, Argentina

2014–2019

Professional Experience

Graduate Researcher

Los Alamos National Laboratory (LANL), Los Alamos, New Mexico, USA

06/2026–11/2026

- Design of quantum error correction protocols using geometric optimization.
- Machine learning for quantum control applications.
- Quantum foundations.

Doctoral Candidate, Physics

Radboud University, Nijmegen, The Netherlands

2022–2026

- Developed quantum control methods for open quantum systems, including stochastic optimal control and path-integral quantum control frameworks.
- Worked on Lindblad dynamics, quantum trajectories, and continuous-measurement formalisms using both Itô and Stratonovich approaches.
- Investigated feedback-driven quantum state-preparation schemes and connections between quantum measurement dynamics and double-bracket gradient flows.
- Built numerical tools in Python and MATLAB for simulation, optimization, and control of open quantum systems.

- Worked on methods relevant to near-term quantum hardware, including decoherence-aware and measurement-aware state preparation.
- Performed quantum complexity analysis of quantum annealing algorithms.

Quantum Computing Summer Intern

Los Alamos National Laboratory (LANL), Los Alamos, New Mexico, USA
Summer 2025

- Selected for a LANL and U.S. Department of Energy-supported summer fellowship with a stated acceptance rate below 2%.
- Developed quantum feedback control methods for continuous-quantum-measurement dynamics and geometric connections with double-bracket gradient flows. Publication: arXiv:2512.15412 (2025).
- Developed software in Julia and MATLAB for calibrating D-Wave quantum annealers.
- Presented project results in the QCSS LANL 2025 seminar.

Doctoral Researcher

National University of Córdoba / CONICET, Córdoba, Argentina
2020–2021

- Conducted research in computational physics and reduced-order modeling.
- Contributed to surrogate-modeling and machine-learning-oriented scientific software, including the Arby package.

Visiting Research Student

Louisiana State University, Baton Rouge, Louisiana, USA
Jan 2020–Feb 2020

- Visiting student invited by Dr. Jorge Pullin for research collaboration.
- Worked on a Review about reduced-order modelings methods for gravitational waves, published in Living Reviews in Relativity.

Computer Skills

- **Languages and frameworks:** Python, Julia, MATLAB, PyTorch, JAX, Git, CI/CD, unit testing, profiling, parallel computing.

Code Projects

PiQC — Path Integral Control for Open and Closed Quantum Systems

GitHub: github.com/aaronuv/piqc

- Implementations of path-integral quantum control methods in MATLAB and Python.
- Quantum simulation and Monte Carlo optimization of open and closed quantum systems for control, including state preparation, gate synthesis, and noise mitigation.

Arby — Reduced Order Modeling Tools in Python

GitHub: github.com/aaronuv/arby

- Scientific software for reduced-order modeling and surrogate construction in Python.
- Production-level research code with software engineering best practices, including version control, testing, and CI/CD.

Selected Publications

- **Aarón Villanueva** & Luis Pedro García-Pintos, “Hamiltonian and double-bracket flow formulations of quantum measurements,” *Reports on Progress in Physics* **89**, 067602 (2026).
- Najafi, P., **Villanueva, A.**, & Kappen, H., “Path Integral Quantum Control for Quantum Chemistry Applications,” arXiv:2509.24104 (2025).

- **Villanueva, A.**, & Kappen, H., "Path integral control of open quantum systems," arXiv:2410.18635 (2024).
- **Villanueva, A.**, Najafi, P., & Kappen, H. J., "Why adiabatic quantum annealing is unlikely to yield speed-up," *Journal of Physics A: Mathematical and Theoretical* **56**(46), 465304 (2023).

Other Publications

- Barsotti, D., Cerino, F., Tiglio, M., & **Villanueva, A.**, "Gravitational wave surrogates through automated machine learning," *Classical and Quantum Gravity* **39**(8), 085011 (2022).
- Tiglio, M., **Villanueva, A.**, "On ab initio-based, free and closed-form expressions for gravitational waves," *Scientific Reports* **11**, 5832 (2021). doi:10.1038/s41598-021-85102-y
- Manuel Tiglio and **Aarón Villanueva**, "On the stability and accuracy of the empirical interpolation method and gravitational wave surrogates," *Classical and Quantum Gravity* **38**, 135011 (2021).
- Manuel Tiglio, **Aarón Villanueva**, "Reduced order and surrogate models for gravitational waves," *Living Reviews in Relativity* **25**, 2 (2022)
- **Aarón Villanueva** et al., "Arby — Fast data-driven surrogates," arXiv:2108.01305 [stat.CO] (2021).

Awards

- **LANL Summer Fellowship**, 2025. Supported by Los Alamos National Laboratory and the U.S. Department of Energy.
- **Junior Scientist Scholarship**, AQC 2023.
- **CONICET Doctoral Scholarship**, 2020–2021.
- **EVC-CIN Scholarship** for academic merit, 2019.

Teaching

- Teaching Assistant, Advanced Machine Learning, Radboud University, 2022–2024.
- Teaching Assistant, CDS Machine Learning, Radboud University, 2022–2024.
- Teaching Assistant, Mathematical Methods of Physics I, FaMAF, 2019.
- Teaching Assistant, Electromagnetism I, FaMAF, 2019.

Languages

- Spanish (native), English (C1 advanced), Italian (A2).

Academic Courses

Advanced undergraduate courses

- Mathematical Methods of Physics I, II
- Electromagnetism I, II
- Quantum Mechanics I, II
- Thermodynamics and Statistical Mechanics I, II

Graduate level courses

- Molecular Electronics — Prof. Horacio Pastawski — Score: 10/10
- Quantum Field Theory I — Prof. Carlos N. Kozameh — Score: 10/10
- Reduced Order Modeling — Prof. Manuel Tiglio — Score: 10/10
- Neural Networks — Prof. Francisco Tamarit — Score: 10/10

- Software Design for Scientific Computing — Prof. Juan B. Cabral — Score: 10/10
- Partial Differential Equations: Analytic and Numerical Methods — Prof. Oscar Reula — Approved
- General Relativity I — Prof. Omar Ortiz — Approved
- CDS Machine Learning — Prof. Hilbert Kappen — Approved
- Advanced Machine Learning — Prof. Hilbert Kappen — Approved

Extracurricular Courses and Certifications

- **English course** (3 months, 2018) — Facultad de Lenguas, UNC — Medium Level Certification.
- **DataRobot Essentials** (4 sessions, 2020) — DataRobot University — Essentials Certification: verify.skilljar.com/c/pway97gbocro

Conferences, Schools and Workshops

- **Quantum Computing Summer Fellowship** (Internship, 2025) — Los Alamos National Laboratory, Los Alamos, New Mexico, USA.
- **EQAI 2024** (Summer School, 2024) — Lignano, Italy.
- **AQC 2023** (Congress, 2023) — Center for Quantum Information and Control, The University of New Mexico, USA.
- **ELLIS unConference 2023** (Congress, 2023) — HEC, Paris.
- **QTML 2023** (Congress, 2023) — CERN, Switzerland.
- **Duke Machine Learning Winter School: Computer Vision** (Virtual school, Jan 2022).

- **The Einstein Toolkit Workshop** (Virtual Workshop, 2020) — CCT-LSU, Baton Rouge, USA.
- **GRAV 19** (Congress, 2019) — GRG-FaMAF, Córdoba, Argentina.
- **Expresiones cerradas ab Initio para ondas gravitacionales** (Seminar, presenter, 2019) — GRG-FaMAF, Córdoba, Argentina.
- **RAFA 2018** (Congress, 2018) — AFA, Buenos Aires, Argentina.
- **The Sound of Space-time: The Dawn of Gravitational Wave Science** (School, 2018) — ICTP-SAIFR, São Paulo, Brazil.

Seminars and Outreach

- **QCSS LANL 2025** (Seminar, July 2025) — Project presentation on calibrating the D-Wave machine.
- **QTML 2023** (Poster, Nov 2023) — CERN, Switzerland. Subject: quantum annealing and quantum complexity.
- **AQC 2023** (Poster, June 2023) — Poster presentation on quantum annealing and quantum complexity.
- **INQA Seminar** (Seminar, 6 June 2023) — Online seminar for the INQA network. Title: *Grover-like speedup for combinatorial optimization and the fine tuning problem in quantum annealing.*
- **Pint Of Science** (Talk, 23 May 2023) — Gave a talk on quantum computing. Title: *Finding needles in a haystack with quantum computers.*
- **DCN Seminar** (Seminar, 2023) — Weekly seminar at DCN, The Donders Institute. Title: *Finding needles in a haystack: how quantum computing can help in solving hard problems.*

Chapter 9

List of publications and collaborations

1. **A. Villanueva** and L. P. García-Pintos, "Hamiltonian and double-bracket flow formulations of quantum measurements," *Rep. Prog. Phys.* 89 067602 (2026).
2. P. Najafi, **A. Villanueva**, and H. Kappen, "Path integral quantum control for quantum chemistry applications," arXiv:2509.24104 (2025).
3. **A. Villanueva** and H. Kappen, "Path integral control of open quantum systems," arXiv:2410.18635 (2024).
4. **A. Villanueva**, P. Najafi, and H. J. Kappen, "Why adiabatic quantum annealing is unlikely to yield speed-up," *Journal of Physics A: Mathematical and Theoretical* **56**(46), 465304 (2023).

Chapter 10

Donders Graduate School

For a successful research Institute, it is vital to train the next generation of scientists. To achieve this goal, the Donders Institute for Brain, Cognition and Behaviour established the Donders Graduate School in 2009. The mission of the Donders Graduate School is to guide our graduates to become skilled academics who are equipped for a wide range of professions. To achieve this, we do our utmost to ensure that our PhD candidates receive support and supervision of the highest quality.

Since 2009, the Donders Graduate School has grown into a vibrant community of highly talented national and international PhD candidates, with over 500 PhD candidates enrolled. Their backgrounds cover a wide range of disciplines, from physics to psychology, medicine to psycholinguistics, and biology to artificial intelligence. Similarly, their interdisciplinary research covers genetic, molecular, and cellular processes at one end and computational, system-level neuroscience with cognitive and behavioural analysis at the other end. We ask all PhD candidates within the Donders Graduate School to publish their PhD thesis in the Donders Thesis Series. This series currently includes over 750 PhD theses from our PhD graduates and thereby provides a comprehensive overview of the diverse types of research performed at the Donders Institute. A complete overview of the Donders Thesis Series can be

found on our website: <https://www.ru.nl/donders/donders-series>

The Donders Graduate School tracks the careers of our PhD graduates carefully. In general, the PhD graduates end up at high-quality positions in different sectors, for a complete overview see <https://www.ru.nl/donders/destination-our-former-phd>. A large proportion of our PhD alumni continue in academia (>50%). Most of them first work as a postdoc before growing into more senior research positions. They work at top institutes worldwide, such as University of Oxford, University of Cambridge, Stanford University, Princeton University, UCL London, MPI Leipzig, Karolinska Institute, UC Berkeley, EPFL Lausanne, and many others. In addition, a large group of PhD graduates continue in clinical positions, sometimes combining it with academic research. Clinical positions can be divided into medical doctors, for instance, in genetics, geriatrics, psychiatry, or neurology, and in psychologists, for instance as healthcare psychologist, clinical neuropsychologist, or clinical psychologist. Furthermore, there are PhD graduates who continue to work as researchers outside academia, for instance at non-profit or government organizations, or in pharmaceutical companies. There are also PhD graduates who work in education, such as teachers in high school, or as lecturers in higher education. Others continue in a wide range of positions, such as policy advisors, project managers, consultants, data scientists, web- or software developers, business owners, regulatory affairs specialists, engineers, managers, or IT architects. As such, the career paths of Donders PhD graduates span a broad range of sectors and professions, but the common factor is that they almost all have become successful professionals. For more information on the Donders Graduate School, as well as past and upcoming defences please visit: <http://www.ru.nl/donders/graduate-school/phd/>

Bibliography

- [1] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter, et al., “Logical quantum processor based on reconfigurable atom arrays”, *Nature* **626**, 58–64 (2024).
- [2] V. V. Sivak, A. Eickbusch, B. Royer, S. Singh, I. Tsioutsios, S. Ganjam, A. Miano, B. Brock, A. Ding, L. Frunzio, et al., “Real-time quantum error correction beyond break-even”, *Nature* **616**, 50–55 (2023).
- [3] R. Acharya, I. Aleiner, R. Allen, T. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, J. Atalaya, R. Babbush, et al., “Suppressing quantum errors by scaling a surface code logical qubit”, *Nature* **614**, 676–681 (2023).
- [4] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, “High-threshold and low-overhead fault-tolerant quantum memory”, *Nature* **627**, 778–782 (2024).
- [5] H. J. Manetsch, G. Nomura, E. Bataille, X. Lv, K. H. Leung, and M. Endres, “A tweezer array with 6,100 highly coherent atomic qubits”, *Nature* **647**, 60–67 (2025).
- [6] King, Andrew D. et al., “Beyond-classical computation in quantum simulation”, *Science* **388**, 199–204 (2025).
- [7] Google Quantum AI and Collaborators, “Quantum error correction below the surface code threshold”, *Nature* **638**, 920–926 (2025).

- [8] Google Quantum AI and Collaborators, "Observation of constructive interference at the edge of quantum ergodicity", *Nature* **646**, 825–830 (2025).
- [9] J. Eisert and J. Preskill, "Mind the gaps: the fraught road to quantum advantage", 10.48550/arXiv.2510.19928 (2025).
- [10] R. Babbush, R. King, S. Boixo, W. Huggins, T. Khattar, G. H. Low, J. R. McClean, T. O'Brien, and N. C. Rubin, "The grand challenge of quantum applications", 10.48550/arXiv.2511.09124 (2025).
- [11] F. Gemeinhardt, F. C. Garmendía, et al., "Quantum combinatorial optimization in the nisq era: a systematic mapping study", *ACM Computing Surveys* **56**, 70:1–70:36 (2023).
- [12] F. Chicano, G. Luque, Z. A. Dahi, and R. Gil-Merino, "Combinatorial optimization with quantum computers", *Engineering Optimization*, 10.1080/0305215X.2024.2435538 (2025).
- [13] M. R. James, "Optimal quantum control theory", *Annual Review of Control, Robotics, and Autonomous Systems* **4**, 343–367 (2021).
- [14] C. Brif, R. Chakrabarti, and H. Rabitz, "Control of quantum phenomena: past, present, and future", *New Journal of Physics* **12**, 075008 (2010).
- [15] D. D'Alessandro, *Introduction to quantum control and dynamics*, 2nd ed. (Chapman and Hall/CRC, 2021).
- [16] C. P. Koch, U. Boscain, T. Calarco, G. Dirr, S. Filipp, S. J. Glaser, R. Kosloff, S. Montangero, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm, "Quantum optimal control in quantum technologies. Strategic report on current status, visions and goals for research in Europe", *EPJ Quantum Technology* **9**, 19 (2022).
- [17] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, "Quantum computation by adiabatic evolution", 10.48550/arXiv.quant-SBph/0001106 (2000).
- [18] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, "A quantum adiabatic evolution algorithm applied to random instances of an np-complete problem", *Science* **292**, 472–475 (2001).

- [19] E. Farhi, J. Goldstone, and S. Gutmann, "Quantum adiabatic evolution algorithms versus simulated annealing", 10.48550/arXiv.quant-SBph/0201031 (2002).
- [20] H. J. Kappen, "Path integrals and symmetry breaking for optimal control theory", Journal of Statistical Mechanics: Theory and Experiment **2005**, P11011 (2005).
- [21] H. J. Kappen, "Linear theory for control of nonlinear stochastic systems", Physical Review Letters **95**, 1–4 (2005).
- [22] H. J. Kappen and H. C. Ruiz, "Adaptive importance sampling for control and inference", Journal of Statistical Physics **164**, 109–129 (2016).
- [23] W. H. Fleming and H. M. Soner, *Controlled markov processes and viscosity solutions*, 2nd ed., Vol. 25, Applications of Mathematics (Springer, New York, 2006).
- [24] M. Raisi, A. Noohian, and S. Fallah, "A fault-tolerant and robust controller using model predictive path integral control for free-flying space robots", Frontiers in Robotics and AI **9**, 10.3389/frobt.2022.1027918 (2022).
- [25] L. Hou, H. Wang, H. Zou, and Y. Zhou, "Robotic manipulation planning for automatic peeling of glass substrate based on online learning model predictive path integral", Sensors **22**, 10.3390/s22031292 (2022).
- [26] K. Yamamoto, R. Ariizumi, T. Hayakawa, and F. Matsuno, "Path integral policy improvement with population adaptation", IEEE Transactions on Cybernetics **52**, 312–322 (2022).
- [27] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Aggressive driving with model predictive path integral control", in 2016 IEEE International Conference on Robotics and Automation (ICRA) (2016), pp. 1433–1440.

- [28] M. S. Gandhi, B. Vlahov, J. Gibson, G. Williams, and E. A. Theodorou, "Robust model predictive path integral control: analysis and performance guarantees", *IEEE Robotics and Automation Letters* **6**, 1423–1430 (2021).
- [29] I. S. Mohamed, K. Yin, and L. Liu, "Autonomous navigation of agvs in unknown cluttered environments: log-mppi control strategy", *IEEE Robotics and Automation Letters* **7**, 10240–10247 (2022).
- [30] B. van den Broek, W. Wiegierinck, and B. Kappen, "Graphical model inference in optimal control of stochastic multi-agent systems", *Journal of Artificial Intelligence Research* **32**, 95–122 (2008).
- [31] V. Gómez, S. Thijssen, A. Symington, S. Hailes, and H. J. Kappen, "Real-time stochastic optimal control for multi-agent quadrotor systems", in *Proceedings of the international conference on automated planning and scheduling*, Vol. 26 (2016), pp. 468–476.
- [32] P. Varnai and D. V. Dimarogonas, "Multi-agent stochastic control using path integral policy improvement", in *2022 american control conference (acc)* (2022), pp. 3406–3411.
- [33] N. Wan, A. Gahlawat, N. Hovakimyan, E. A. Theodorou, and P. G. Voulgaris, "Cooperative path integral control for stochastic multi-agent systems", in *2021 american control conference (acc)* (2021), pp. 1262–1267.
- [34] M. Kazim, J. Hong, M.-G. Kim, and K.-K. K. Kim, "Recent advances in path integral control for trajectory optimization: an overview in theoretical and algorithmic perspectives", *Annual Reviews in Control* **57**, 100931 (2024).
- [35] E. B. Davies, "Quantum stochastic processes", *Communications in Mathematical Physics* **15**, 277–304 (1969).
- [36] H. Carmichael, *An open systems approach to quantum optics: lectures presented at the université libre de bruxelles october 28 to november 4, 1991*, Vol. 18, *Lecture Notes in Physics Monographs* (Springer, 1993).
- [37] I. Percival, *Quantum state diffusion* (Cambridge University Press, 1998).

- [38] A. J. Daley, "Quantum trajectories and open many-body quantum systems", *Advances in Physics* **63**, 77–149 (2014).
- [39] K. Jacobs and D. A. Steck, "A straightforward introduction to continuous quantum measurement", *Contemp. Phys.* **47**, 279–303 (2006).
- [40] A. Barchielli and M. Gregoratti, *Quantum Trajectories and Measurements in Continuous Time: The Diffusive Case*, Vol. 782, *Lecture Notes in Physics* (Springer Berlin Heidelberg, Berlin, Heidelberg, 2009).
- [41] R. W. Brockett, "Differential geometry and the design of gradient algorithms. differential geometry", *Partial Differential Equations on Manifolds*, 69–92 (1993).
- [42] U. Helmke and J. B. Moore, *Optimization and dynamical systems* (Springer, 1994).
- [43] T. Schulte-Herbrüggen, S. J. Glaser, G. Dirr, and U. Helmke, "Gradient flows for optimization in quantum information and quantum dynamics: foundations and applications", *Rev. Math. Phys.* **22**, 597–667 (2010).
- [44] M. Gluza, "Double-bracket quantum algorithms for diagonalization", *Quantum* **8**, 1316 (2024).
- [45] M. Gluza, J. Son, B. H. Tiang, R. Zander, R. Seidel, Y. Suzuki, Z. Holmes, and N. H. Y. Ng, "Double-bracket quantum algorithms for quantum imaginary-time evolution", *Physical Review Letters* **136**, 020601 (2026).
- [46] M. R. Garey and D. S. Johnson, *Computers and intractability: a guide to the theory of NP-completeness* (W. H. Freeman, San Francisco, 1979).
- [47] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing", *Science* **220**, 671–680 (1983).
- [48] B. Hajek, "Cooling schedules for optimal annealing", *Mathematics of Operations Research* **13**, 311–329 (1988).
- [49] J. Brooke, D. Bitko, T. F. Rosenbaum, and G. Aeppli, "Quantum annealing of a disordered magnet", *Science* **284**, 779–781 (1999).

- [50] T. Kadowaki and H. Nishimori, "Quantum annealing in the transverse ising model", *Phys. Rev. E* **58**, 5355–5363 (1998).
- [51] B. Apolloni, C. Carvalho, and D. De Falco, "Quantum stochastic optimization", *Stochastic Processes and their Applications* **33**, 233–244 (1989).
- [52] B. Apolloni, N. Cesa-Bianchi, and D. De Falco, "A numerical implementation of "quantum annealing"", in *Stochastic processes, physics and geometry: proceedings of the ascona-locarno conference* (World Scientific, 1990), pp. 97–111.
- [53] L. K. Grover, "Quantum mechanics helps in searching for a needle in a haystack", *Physical review letters* **79**, 325 (1997).
- [54] J. Roland and N. J. Cerf, "Quantum search by local adiabatic evolution", *Physical Review A* **65**, 042308 (2002).
- [55] A. P. Young, S. Knysh, and V. N. Smelyanskiy, "Size dependence of the minimum excitation gap in the quantum adiabatic algorithm", *Physical Review Letters* **101**, 170503 (2008).
- [56] T. Hogg, "Adiabatic quantum computing for random satisfiability problems", *Physical Review A* **67**, 022314 (2003).
- [57] R. Schützhold and G. Schaller, "Adiabatic quantum algorithms as quantum phase transitions: first versus second order", *Physical Review A* **74**, 060304 (2006).
- [58] A. P. Young, S. Knysh, and V. N. Smelyanskiy, "First-order phase transition in the quantum adiabatic algorithm", *Physical Review Letters* **104**, 020502 (2010).
- [59] R. Martoňák, G. E. Santoro, and E. Tosatti, "Quantum annealing by the path-integral monte carlo method: the two-dimensional random ising model", *Physical Review B* **66**, 094203 (2002).
- [60] G. E. Santoro, R. Martoňák, E. Tosatti, and R. Car, "Theory of quantum annealing of an ising spin glass", *Science* **295**, 2427–2430 (2002).

- [61] D. A. Battaglia, G. E. Santoro, and E. Tosatti, "Optimization by quantum annealing: lessons from hard satisfiability problems", *Physical Review E* **71**, 066707 (2005).
- [62] B. Heim, T. F. Rønnow, S. V. Isakov, and M. Troyer, "Quantum versus classical annealing of ising spin glasses", *Science* **348**, 215–217 (2015).
- [63] T. Jörg, F. Krzakala, J. Kurchan, and A. C. Maggs, "Simple glass models and their quantum annealing", *Physical Review Letters* **101**, 147204 (2008).
- [64] B. Altshuler, H. Krovi, and J. Roland, "Adiabatic quantum optimization fails for random instances of np-complete problems", 10 . 48550 / *arXiv*.0908.2782 (2009).
- [65] L. Kong and E. Crosson, "The performance of the quantum adiabatic algorithm on spike hamiltonians", *International Journal of Quantum Information* **15**, 1750011 (2017).
- [66] S. Muthukrishnan, T. Albash, and D. A. Lidar, "Tunneling and speedup in quantum optimization for permutation-symmetric problems", *Physical Review X* **6**, 031010 (2016).
- [67] E. Crosson and A. W. Harrow, "Simulated quantum annealing can be exponentially faster than classical simulated annealing", in 2016 IEEE 57th annual symposium on foundations of computer science (focs) (2016), pp. 714–723.
- [68] T. Neuhaus, M. Peschina, K. Michielsen, and H. De Raedt, "Classical and quantum annealing in the median of three satisfiability", in *International conference on quantum information* (Optica Publishing Group, 2011), QMI23.
- [69] C.-W. Liu, A. Polkovnikov, and A. W. Sandvik, "Quantum versus classical annealing: insights from scaling theory and results for spin glasses on 3-regular graphs", *Physical Review Letters* **114**, 147203 (2015).

- [70] Z. Bian, F. Chudak, W. Macready, A. Roy, R. Sebastiani, and S. Varotti, "Solving sat (and maxsat) with a quantum annealer: foundations, encodings, and preliminary results", *Information and Computation* **275**, 104609 (2020).
- [71] M. Kowalsky, T. Albash, I. Hen, and D. A. Lidar, "3-regular three-xorsat planted solutions benchmark of classical and quantum heuristic optimizers", *Quantum Science and Technology* **7**, 025008 (2022).
- [72] T. F. Rønnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer, "Defining and detecting quantum speedup", *Science* **345**, 420–424 (2014).
- [73] A. Perdomo-Ortiz, S. E. Venegas-Andraca, and A. Aspuru-Guzik, "A study of heuristic guesses for adiabatic quantum computation", *Quantum Information Processing* **10**, 33–52 (2011).
- [74] E. Farhi, J. Goldstone, S. Gutmann, and D. Nagaj, "How to make the quantum adiabatic algorithm fail", *International Journal of Quantum Information* **6**, 503–516 (2008).
- [75] T. Graß, "Quantum annealing with longitudinal bias fields", *Physical Review Letters* **123**, 120501 (2019).
- [76] M. H. S. Amin, "Effect of local minima on adiabatic quantum optimization", *Physical Review Letters* **100**, 130503 (2008).
- [77] A. T. Rezakhani, W.-J. Kuo, A. Hamma, D. A. Lidar, and P. Zanardi, "Quantum adiabatic brachistochrone", *Phys. Rev. Lett.* **103**, 080502 (2009).
- [78] E. Farhi, J. Goldstone, and S. Gutmann, "Quantum adiabatic evolution algorithms with different paths", 10.48550/arXiv.quant-SBph/0208135 (2002).
- [79] B. Yan and N. A. Sinitsyn, "Analytical solution for nonadiabatic quantum annealing to arbitrary ising spin hamiltonian", *Nature Communications* **13**, 2212 (2022).

- [80] E. Crosson, E. Farhi, C. Y.-Y. Lin, H.-H. Lin, and P. Shor, “Different strategies for optimization using the quantum adiabatic algorithm”, 10.48550/arXiv.1401.7320 (2014).
- [81] E. J. Crosson and D. A. Lidar, “Prospects for quantum enhancement with diabatic quantum annealing”, *Nature Reviews Physics* **3**, 466–489 (2021).
- [82] A. Shi, H. Guan, and W. Zhang, “Efficient diabatic quantum algorithm in number factorization”, *Physics Letters A* **384**, 126745 (2020).
- [83] A. M. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D. A. Spielman, “Exponential algorithmic speedup by a quantum walk”, in *Proceedings of the thirty-fifth annual acm symposium on theory of computing* (2003), pp. 59–68.
- [84] T. Albash and D. A. Lidar, “Adiabatic quantum computation”, *Reviews of Modern Physics* **90**, 015002 (2018).
- [85] M. B. Hastings, “The power of adiabatic quantum computation with no sign problem”, *Quantum* **5**, 597 (2021).
- [86] A. Gilyén, M. B. Hastings, and U. Vazirani, “(sub)exponential advantage of adiabatic quantum computation with no sign problem”, in *Proceedings of the 53rd annual acm sigact symposium on theory of computing* (2021), pp. 1357–1369.
- [87] M. Žnidarič and M. Horvat, “Exponential complexity of an adiabatic algorithm for an NP-complete problem”, *Physical Review A* **73**, 022329 (2006).
- [88] M. Sluiskii, T. Albash, L. Barash, and I. Hen, “Analog nature of quantum adiabatic unstructured search”, *New Journal of Physics* **21**, 113025 (2019).
- [89] M. Born and V. Fock, “Beweis des Adiabatenatzes”, *Zeitschrift für Physik* **51**, 165–180 (1928).
- [90] T. Kato, “On the adiabatic theorem of quantum mechanics”, *Journal of the Physical Society of Japan* **5**, 435–439 (1950).

- [91] S. Jansen, M.-B. Ruskai, and R. Seiler, "Bounds for the adiabatic approximation with applications to quantum computation", *Journal of Mathematical Physics* **48**, 102111 (2007).
- [92] M. Mézard, G. Parisi, and R. Zecchina, "Analytic and algorithmic solution of random satisfiability problems", *Science* **297**, 812–815 (2002).
- [93] D. J. C. MacKay, *Information theory, inference, and learning algorithms* (Copyright Cambridge University Press, 2003).
- [94] S. Arora and B. Barak, *Computational complexity: a modern approach* (Cambridge University Press, 2009).
- [95] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge university press, 2010).
- [96] G. Brassard, "Searching a quantum phone book", *Science* **275**, 627–628 (1997).
- [97] M. Boyer, G. Brassard, P. Høyer, and A. Tapp, "Tight bounds on quantum searching", *Fortschritte der Physik* **46**, 493–505 (1998).
- [98] T. J. Yoder, G. H. Low, and I. L. Chuang, "Fixed-point quantum search with an optimal number of queries", *Physical Review Letters* **113**, 1–5 (2014).
- [99] A. M. Dalzell, T. J. Yoder, and I. L. Chuang, "Fixed-point adiabatic quantum search", *Physical Review A* **95**, 1–16 (2017).
- [100] W. van Dam, M. Mosca, and U. V. Vazirani, "How powerful is adiabatic quantum computation?", in *Proceedings of the 42nd annual symposium on foundations of computer science* (2001), pp. 279–287.
- [101] I. Hen and A. P. Young, "Exponential complexity of the quantum adiabatic algorithm for certain satisfiability problems", *Phys. Rev. E* **84**, 061152 (2011).
- [102] J. Franco and M. Paull, "Probabilistic analysis of the davis putnam procedure for solving the satisfiability problem", *Discrete Applied Mathematics* **5**, 77–87 (1983).

- [103] V. Chvátal and E. Szemerédi, “Many hard examples for resolution”, J. ACM **35**, 759–768 (1988).
- [104] A. Braunstein, M. Mézard, and R. Zecchina, “Survey propagation: an algorithm for satisfiability”, Random Structures & Algorithms **27**, 201–226 (2005).
- [105] S. Mertens, M. Mézard, and R. Zecchina, “Threshold values of random K-sat from the cavity method”, Random Structures & Algorithms **28**, 340–373 (2006).
- [106] 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”, Physics Reports **523**, 127–205 (2013).
- [107] I. Hen, “How quantum is the speedup in adiabatic unstructured search?”, Quantum Information Processing **18**, 183 (2019).
- [108] B. F. Schiffer, J. Tura, and J. I. Cirac, “Adiabatic spectroscopy and a variational quantum adiabatic algorithm”, PRX Quantum **3**, 020347 (2022).
- [109] S. Aaronson, “How much structure is needed for huge quantum speedups?”, 10.48550/arXiv.2209.06930 (2022).
- [110] J. Ding and A. Zhou, “Eigenvalues of rank-one updated matrices with some applications”, Applied Mathematics Letters **20**, 1223–1226 (2007).
- [111] D. A. Lidar and T. A. Brun, *Quantum error correction* (Cambridge University Press, 2013).
- [112] I. M. Georgescu, S. Ashhab, and F. Nori, “Quantum simulation”, Reviews of Modern Physics **86**, 153–185 (2014).
- [113] V. Giovannetti, S. Lloyd, and L. Maccone, “Advances in quantum metrology”, Nature Photonics **5**, 222–229 (2011).
- [114] C. L. Degen, F. Reinhard, and P. Cappellaro, “Quantum sensing”, Reviews of Modern Physics **89**, 035002 (2017).

- [115] J. M. Arrazola, V. Bergholm, K. Brádler, T. R. Bromley, M. J. Collins, I. Dhand, A. Fumagalli, T. Gerrits, A. Goussev, L. G. Helt, et al., “Quantum circuits with many photons on a programmable nanophotonic chip”, *Nature* **591**, 54–60 (2021).
- [116] S. L. N. Hermans, M. Pompili, H. K. C. Beukers, S. Baier, J. Borregaard, and R. Hanson, “Qubit teleportation between non-neighbouring nodes in a quantum network”, *Nature* **605**, 663–668 (2022).
- [117] P. Baßler, M. Zipper, C. Cedzich, M. Heinrich, P. H. Huber, M. Johanning, and M. Kliesch, “Synthesis of and compilation with time-optimal multi-qubit gates”, *Quantum* **7**, 984 (2023).
- [118] P. Bushev, D. Rotter, A. Wilson, F. Dubin, C. Becher, J. Eschner, R. Blatt, V. Steixner, P. Rabl, and P. Zoller, “Feedback Cooling of a Single Trapped Ion”, *Physical Review Letters* **96**, 043003 (2006).
- [119] M. Rossi, D. Mason, J. Chen, Y. Tsaturyan, and A. Schliesser, “Measurement-based quantum control of mechanical motion”, *Nature* **563**, 53–58 (2018).
- [120] C. S. Adams, J. D. Pritchard, and J. P. Shaffer, “Rydberg atom quantum technologies”, *Journal of Physics B: Atomic, Molecular and Optical Physics* **53**, 012002 (2019).
- [121] L. Henriët, L. Beguin, A. Signoles, T. Lahaye, A. Browaeys, G.-O. Raymond, and C. Jurczak, “Quantum computing with neutral atoms”, *Quantum* **4**, 327 (2020).
- [122] L. Magrini, P. Rosenzweig, C. Bach, A. Deutschmann-Olek, S. G. Hofer, S. Hong, N. Kiesel, A. Kugi, and M. Aspelmeyer, “Real-time optimal quantum control of mechanical motion at room temperature”, *Nature* **595**, 373–377 (2021).
- [123] S. J. Glaser, U. Boscain, T. Calarco, C. P. Koch, W. Köckenberger, R. Kosloff, I. Kuprov, B. Luy, S. Schirmer, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm, “Training Schrödinger’s cat: quantum optimal control: Strategic report on current status, visions and goals for research in Europe”, *The European Physical Journal D* **69**, 279 (2015).

- [124] S. Cong, *Control of quantum systems: theory and methods* (John Wiley & Sons, 2014).
- [125] T. S. Mahesh, P. Batra, and M. H. Ram, "Quantum optimal control: practical aspects and diverse methods", *Journal of the Indian Institute of Science* **103**, 591–607 (2023).
- [126] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, "Optimal control of coupled spin dynamics: design of nmr pulse sequences by gradient ascent algorithms", *Journal of Magnetic Resonance* **172**, 296–305 (2005).
- [127] Y. Chen, Y. Hao, Z. Wu, B.-Y. Wang, R. Liu, Y. Hou, J. Cui, M.-H. Yung, and X. Peng, "Accelerating quantum optimal control through iterative gradient-ascent pulse engineering", *Physical Review A* **108**, 052603 (2023).
- [128] T. Caneva, T. Calarco, and S. Montangero, "Chopped random-basis quantum optimization", *Physical Review A* **84**, 022326 (2011).
- [129] V. F. Krotov, "Global methods in optimal control theory", in *Advances in nonlinear dynamics and control: a report from russia* (Springer, 1996), pp. 74–121.
- [130] D. J. Tannor, V. Kazakov, V. Orlov, J. Broeckhove, and L. Lathouwers, "Time-dependent quantum molecular dynamics", *NATO ASI Series, Series B: Physics* **299**, 347–360 (1992).
- [131] D. M. Reich, M. Ndong, and C. P. Koch, "Monotonically convergent optimization in quantum control using krotov's method", *The Journal of Chemical Physics* **136**, 10.1063/1.3691827 (2012).
- [132] M. Abdelhafez, D. I. Schuster, and J. Koch, "Gradient-based optimal control of open quantum systems using quantum trajectories and automatic differentiation", *Physical Review A* **99**, 052327 (2019).
- [133] U. Boscain, M. Sigalotti, and D. Sugny, "Introduction to the pontryagin maximum principle for quantum optimal control", *PRX Quantum* **2**, 030203 (2021).

- [134] C. P. Koch, “Controlling open quantum systems: tools, achievements, and limitations”, *Journal of Physics: Condensed Matter* **28**, 213001 (2016).
- [135] S. Boutin, C. K. Andersen, J. Venkatraman, A. J. Ferris, and A. Blais, “Resonator reset in circuit qed by optimal control for large open quantum systems”, *Physical Review A* **96**, 042315 (2017).
- [136] M. H. Goerz and K. Jacobs, “Efficient optimization of state preparation in quantum networks using quantum trajectories”, *Quantum Science and Technology* **3**, 045005 (2018).
- [137] R. Gautier, É. Genois, and A. Blais, “Optimal control in large open quantum systems: the case of transmon readout and reset”, *Physical Review Letters* **134**, 070802 (2025).
- [138] K. W. Murch, S. J. Weber, C. Macklin, and I. Siddiqi, “Observing single quantum trajectories of a superconducting quantum bit”, *Nature* **502**, 211–214 (2013).
- [139] P. Campagne-Ibarcq, P. Six, L. Bretheau, A. Sarlette, M. Mirrahimi, P. Rouchon, and B. Huard, “Observing quantum state diffusion by heterodyne detection of fluorescence”, *Physical Review X* **6**, 011002 (2016).
- [140] Q. Ficheux, S. Jezouin, Z. Leghtas, and B. Huard, “Dynamics of a qubit while simultaneously monitoring its relaxation and dephasing”, *Nature Communications* **9**, 1926 (2018).
- [141] Z. K. Mineev, S. O. Mundhada, S. Shankar, P. Reinhold, R. Gutiérrez-Jáuregui, R. J. Schoelkopf, M. Mirrahimi, H. J. Carmichael, and M. H. Devoret, “To catch and reverse a quantum jump mid-flight”, *Nature* **570**, 200–204 (2019).
- [142] W. Wieczorek, S. G. Hofer, J. Hoelscher-Obermaier, R. Riedinger, K. Hammerer, and M. Aspelmeyer, “Optimal state estimation for cavity optomechanical systems”, *Physical Review Letters* **114**, 223601 (2015).

- [143] M. Rossi, D. Mason, J. Chen, and A. Schliesser, "Observing and verifying the quantum trajectory of a mechanical resonator", *Physical Review Letters* **123**, 163601 (2019).
- [144] R. A. Thomas, M. Parniak, C. Østfeldt, C. B. Møller, C. Bærentsen, Y. Tsaturyan, A. Schliesser, J. Appel, E. Zeuthen, and E. S. Polzik, "Entanglement between distant macroscopic mechanical and spin systems", *Nature Physics* **17**, 228–233 (2021).
- [145] M. R. Abdelhafez, "Quantum optimal control using automatic differentiation", PhD thesis (The University of Chicago, 2019).
- [146] N. Gisin and I. C. Percival, "The quantum-state diffusion model applied to open systems", *Journal of Physics A: Mathematical and General* **25**, 5677 (1992).
- [147] H. M. Wiseman and G. J. Milburn, "Quantum theory of optical feedback via homodyne detection", *Physical Review Letters* **70**, 548 (1993).
- [148] V. P. Belavkin, "Theory of the control of observable quantum systems", *Avtomatika i Telemekhanika* **44**, 178–188 (1983).
- [149] V. P. Belavkin, O. Hirota, and R. L. Hudson, *Quantum communications and measurement* (Springer Science & Business Media, 2013).
- [150] H. M. Wiseman and G. J. Milburn, *Quantum measurement and control* (Cambridge University Press, 2009).
- [151] H. Pham, *Continuous-time stochastic control and optimization with financial applications*, Vol. 61 (Springer Science & Business Media, 2009).
- [152] J. Yong and X. Y. Zhou, *Stochastic controls: hamiltonian systems and HJB equations* (Springer, 1999).
- [153] A. C. Doherty and K. Jacobs, "Feedback control of quantum systems using continuous state estimation", *Physical Review A* **60**, 2700–2711 (1999).
- [154] A. C. Doherty, S. Habib, K. Jacobs, H. Mabuchi, and S. M. Tan, "Quantum feedback control and classical control theory", *Physical Review A* **62**, 012105 (2000).

- [155] H. J. Kappen, "Linear theory for control of non-linear stochastic systems", *Physical Review Letters* **95**, 200201 (2005).
- [156] S. Thijssen and H. J. Kappen, "Path integral control and state-dependent feedback", *Physical Review E* **91**, 032104 (2015).
- [157] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Aggressive driving with model predictive path integral control", in 2016 IEEE International Conference on Robotics and Automation (ICRA) (2016), pp. 1433–1440.
- [158] M. Kazim, J. Hong, M.-G. Kim, and K.-K. K. Kim, "Recent advances in path integral control for trajectory optimization: an overview in theoretical and algorithmic perspectives", *Annual Reviews in Control* **57**, 100931 (2024).
- [159] O. R. Meitei, B. T. Gard, G. S. Barron, D. P. Pappas, S. E. Economou, E. Barnes, and N. J. Mayhall, "Gate-free state preparation for fast variational quantum eigensolver simulations", *npj Quantum Information* **7**, 155 (2021).
- [160] R. De Keijzer, O. Tse, and S. Kokkelmans, "Pulse based variational quantum optimal control for hybrid quantum computing", *Quantum* **7**, 908 (2023).
- [161] B. Øksendal, *Stochastic differential equations: an introduction with applications* (Springer Science & Business Media, 2013).
- [162] A. Bensoussan, *Estimation and control of dynamical systems*, Vol. 48 (Springer, 2018).
- [163] P.-T. De Boer, D. P. Kroese, S. Mannor, and R. Y. Rubinstein, "A tutorial on the cross-entropy method", *Annals of Operations Research* **134**, 19–67 (2005).
- [164] H.-P. Breuer and F. Petruccione, *The theory of open quantum systems* (Oxford University Press, 2002).

- [165] A. B. Magann, C. Arenz, M. D. Grace, T.-S. Ho, R. L. Kosut, J. R. McClean, H. A. Rabitz, and M. Sarovar, "From pulses to circuits and back again: a quantum optimal control perspective on variational quantum algorithms", *PRX Quantum* **2**, 010101 (2021).
- [166] I. Semina, V. Semin, F. Petruccione, and A. Barchielli, "Stochastic schrödinger equations for markovian and non-markovian cases", *Open Systems & Information Dynamics* **21**, 1440008 (2014).
- [167] A. Gresch and M. Kliesch, "Guaranteed efficient energy estimation of quantum many-body hamiltonians using shadowgrouping", *Nature Communications* **16**, 689 (2025).
- [168] A. Choquette, A. Di Paolo, P. K. Barkoutsos, D. Sénéchal, I. Tavernelli, and A. Blais, "Quantum-optimal-control-inspired ansatz for variational quantum algorithms", *Physical Review Research* **3**, 023092 (2021).
- [169] V. V. Sivak, A. Eickbusch, H. Liu, B. Royer, I. Tsioutsios, and M. H. Devoret, "Model-free quantum control with reinforcement learning", *Physical Review X* **12**, 011059 (2022).
- [170] R. Porotti, V. Peano, and F. Marquardt, "Gradient-ascent pulse engineering with feedback", *PRX Quantum* **4**, 030305 (2023).
- [171] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, "Barren plateaus in quantum neural network training landscapes", *Nature Communications* **9**, 4812 (2018).
- [172] K. W. Moore, R. Chakrabarti, G. Riviello, and H. Rabitz, "Search complexity and resource scaling for the quantum optimal control of unitary transformations", *Physical Review A* **83**, 012326 (2011).
- [173] J. Preskill, "Quantum computing in the nisc era and beyond", *Quantum* **2**, 79 (2018).
- [174] Aarón Villanueva, *Quantum diffusion control (qdc) - matlab implementations*, <https://github.com/aaronuv/qdc>, 2025.
- [175] Z.-J. Chen, H. Huang, L. Sun, Q.-X. Jie, J. Zhou, Z. Hua, Y. Xu, W. Wang, G.-C. Guo, C.-L. Zou, et al., "Robust and optimal control of open quantum systems", *Science Advances* **11**, eadr0875 (2025).

- [176] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne.", *Journal of Machine Learning Research* **9**, 2579–2605 (2008).
- [177] L. M. K. Vandersypen and I. L. Chuang, "Nmr techniques for quantum control and computation", *Reviews of Modern Physics* **76**, 1037–1069 (2004).
- [178] S. G. Schirmer, H. Fu, and A. I. Solomon, "Complete controllability of quantum systems", *Physical Review A* **63**, 063410 (2001).
- [179] J. A. Jones, R. H. Hansen, and M. Mosca, "Quantum logic gates and nuclear magnetic resonance pulse sequences", *Journal of Magnetic Resonance* **135**, 353–360 (1998).
- [180] S. Khatri, K. Sharma, and M. M. Wilde, "Information-theoretic aspects of the generalized amplitude-damping channel", *Physical Review A* **102**, 012401 (2020).
- [181] V. So, M. Duraisamy Suganthi, A. Menon, M. Zhu, R. Zhuravel, H. Pu, P. G. Wolynes, J. N. Onuchic, and G. Pagano, "Trapped-ion quantum simulation of electron transfer models with tunable dissipation", *Science Advances* **10**, eads8011 (2024).
- [182] V. C. Olaya-Agudelo, B. Stewart, C. H. Valahu, R. J. MacDonell, M. J. Millican, V. G. Matsos, F. Scuccimarra, T. R. Tan, and I. Kassal, "Simulating open-system molecular dynamics on analog quantum computers", *Physical Review Research* **7**, 023215 (2025).
- [183] J. Allen, M. Otten, S. K. Gray, and B. K. Clark, "Simulating large one-dimensional neutral-atom quantum systems", *Physical Review A* **112**, 022609 (2025).
- [184] D. A. Meyer and N. R. Wallach, "Global entanglement in multiparticle systems", *Journal of Mathematical Physics* **43**, 4273–4278 (2002).
- [185] Y. Chen, Y. Hao, Z. Wu, B.-Y. Wang, R. Liu, Y. Hou, J. Cui, M.-H. Yung, and X. Peng, "Accelerating quantum optimal control through iterative gradient-ascent pulse engineering", *Physical Review A* **108**, 052603 (2023).

- [186] R. J. de Keijzer, L. Visser, O. Tse, and S. Kokkelmans, “Qubit fidelity distribution under stochastic schrödinger equations driven by classical noise”, *Physical Review Research* **7**, 023063 (2025).
- [187] J. Berberich, D. Fink, and C. Holm, “Robustness of quantum algorithms against coherent control errors”, *Physical Review A* **109**, 012417 (2024).
- [188] L. P. García-Pintos, T. O’Leary, T. Biswas, J. Bringewatt, L. Cincio, L. T. Brady, and Y.-K. Liu, “Resilience–runtime tradeoff relations for quantum algorithms”, *Reports on Progress in Physics* **88**, 037601 (2025).
- [189] B. Shao, X. Yang, R. Liu, Y. Zhai, D. Lu, T. Xin, and J. Li, “Multiple classical noise mitigation by multiobjective robust quantum optimal control”, *Physical Review Applied* **21**, 034042 (2024).
- [190] S. Satoh, H. J. Kappen, and M. Saeki, “An iterative method for nonlinear stochastic optimal control based on path integrals”, *IEEE Transactions on Automatic Control* **62**, 262–276 (2016).
- [191] A. Barchielli and M. Gregoratti, *Quantum trajectories and measurements in continuous time: the diffusive case*, Vol. 782 (Springer, 2009).
- [192] H. M. Wiseman and L. Diósi, “Complete parameterization, and invariance, of diffusive quantum trajectories for Markovian open systems”, *Chemical Physics* **268**, 91–104 (2001).
- [193] R. L. Hudson and K. R. Parthasarathy, “Quantum Itô’s formula and stochastic evolutions”, *Communications in Mathematical Physics* **93**, 301–323 (1984).
- [194] K. R. Parthasarathy, *An introduction to quantum stochastic calculus*, Vol. 85 (Birkhäuser, 2012).
- [195] A. Barchielli and V. P. Belavkin, “Measurements continuous in time and a posteriori states in quantum mechanics”, *Journal of Physics A: Mathematical and General* **24**, 1495 (1991).
- [196] L. Diósi and W. T. Strunz, “The non-markovian stochastic schrödinger equation for open systems”, *Physics Letters A* **235**, 569–573 (1997).

- [197] J. Gambetta and H. M. Wiseman, "Non-markovian stochastic schrödinger equations: generalization to real-valued noise using quantum-measurement theory", *Physical Review A* **66**, 012108 (2002).
- [198] S. Thijssen and H. J. Kappen, "Path integral control and state-dependent feedback", *Physical Review E* **91**, 032104 (2015).
- [199] P. Guilmin, R. Gautier, A. Bocquet, and É. Genois, "Dynamics: an open-source python library for gpu-accelerated and differentiable simulation of quantum systems", <https://github.com/dynamics/dynamics>, arXiv:2405.06830, 2024.
- [200] M. Bukov, A. G. R. Day, D. Sels, P. Weinberg, A. Polkovnikov, and P. Mehta, "Reinforcement learning in different phases of quantum control", *Physical Review X* **8**, 031086 (2018).
- [201] D. J. Griffiths and D. F. Schroeter, *Introduction to quantum mechanics* (Cambridge university press, 2018).
- [202] V. Gorini, A. Kossakowski, and E. C. G. Sudarshan, "Completely positive dynamical semigroups of n-level systems", *Commun. Math. Phys.* **43**, 199–212 (1975).
- [203] G. Lindblad, "On the generators of quantum dynamical semigroups", *Commun. Math. Phys.* **48**, 119–130 (1976).
- [204] K. W. Murch, K. L. Moore, S. Gupta, and D. M. Stamper-Kurn, "Observation of quantum-measurement backaction with an ultracold atomic gas", *Nat. Phys.* **4**, 561–564 (2008).
- [205] K. W. Murch, S. Weber, C. Macklin, and I. Siddiqi, "Observing single quantum trajectories of a superconducting quantum bit", *Nature* **502**, 211–214 (2013).
- [206] P. Campagne-Ibarcq, P. Six, L. Bretheau, A. Sarlette, M. Mirrahimi, P. Rouchon, and B. Huard, "Observing quantum state diffusion by heterodyne detection of fluorescence", *Phys. Rev. X* **6**, 011002 (2016).
- [207] A. Carlini, A. Hosoya, T. Koike, and Y. Okudaira, "Time-optimal quantum evolution", *Phys. Rev. Lett.* **96**, 060503 (2006).

- [208] L. Hu and A. N. Jordan, “Quantum state driving along arbitrary trajectories”, *Phys. Rev. Res.* **5**, 033045 (2023).
- [209] S. Alipour, A. Chenu, A. T. Rezakhani, and A. del Campo, “Shortcuts to adiabaticity in driven open quantum systems: balanced gain and loss and non-markovian evolution”, *Quantum* **4**, 336 (2020).
- [210] K. Jacobs, *Stochastic processes for physicists: understanding noisy systems* (Cambridge University Press, 2010).
- [211] P. E. Kloeden, E. Platen, and H. Schurz, *Numerical solution of sde through computer experiments* (Springer Science & Business Media, 2012).
- [212] L. P. García-Pintos, Y.-K. Liu, and A. V. Gorshkov, “Reshaping the quantum arrow of time”, *Physical Review X* **16**, 011028 (2026).
- [213] R. W. Brockett, “Dynamical systems that sort lists, diagonalize matrices, and solve linear programming problems”, *Linear Algebra Appl.* **146**, 79–91 (1991).
- [214] S. Teufel, *Adiabatic Perturbation Theory in Quantum Dynamics*, Vol. 1821, *Lecture Notes in Mathematics* (Springer, Berlin, Heidelberg, 2003).
- [215] B. Øksendal, “Stochastic differential equations”, in *Stochastic differential equations: an introduction with applications* (Springer, 2003), pp. 38–50.
- [216] A. M. Bloch, R. W. Brockett, and T. S. Ratiu, “Completely integrable gradient flows”, *Commun. Math. Phys.* **147**, 57–74 (1992).
- [217] A. W. Knap, *Lie groups beyond an introduction*, *Progress in Mathematics* 140 (Birkhäuser, Boston, 1996).
- [218] M. Nakahara, *Geometry, Topology and Physics, Second Edition* (Taylor & Francis, June 2003).
- [219] W. Rossmann, *Lie Groups: An Introduction Through Linear Groups* (Oxford University Press, 2006).
- [220] A. Bloch, *Hamiltonian and Gradient Flows, Algorithms and Control* (American Mathematical Soc., 1994).

- [221] G. Dirr and U. Helmke, "Lie theory for quantum control", *GAMM-Mitteilungen* **31**, 59–93 (2008).
- [222] R. Wiersema and N. Killoran, "Optimizing quantum circuits with Riemannian gradient flow", *Phys. Rev. A* **107**, 062421 (2023).
- [223] R. Wiersema, D. Lewis, D. Wierichs, J. Carrasquilla, and N. Killoran, "Here comes the SU(N): multivariate quantum gates and gradients", *Quantum* **8**, 1275 (2024).
- [224] L. Xiaoyue, M. Robbiati, A. Pasquale, E. Pedicillo, A. Wright, S. Carrazza, and M. Gluza, "Strategies for optimizing double-bracket quantum algorithms", 10.48550/arXiv.2408.07431 (2024).
- [225] T.-Y. Tam, "Gradient flows and double bracket equations", *Differential Geometry and its Applications* **20**, 209–224 (2004).
- [226] N. A. McMahon, M. Pervez, and C. Arenz, "Equating quantum imaginary time evolution, Riemannian gradient flows, and stochastic implementations", 10.48550/arXiv.2504.06123 (2025).
- [227] A. Goldberg and J. L. Schwartz, "Integration of the schrödinger equation in imaginary time", *J. Comput. Phys.* **1**, 433–447 (1967).
- [228] H. P. Robertson, "The Uncertainty Principle", *Physical Review* **34**, 163–164 (1929).
- [229] S. L. Adler, D. C. Brody, T. A. Brun, and L. P. Hughston, "Martingale models for quantum state reduction", *J. Phys. A: Math. Gen.* **34**, 8795–8820 (2001).
- [230] D. C. Brody and L. P. Hughston, "Efficient simulation of quantum state reduction", *J. Math. Phys.* **43**, 5254–5261 (2002).
- [231] D. C. Brody and L. P. Hughston, "Quantum noise and stochastic reduction", *J. Phys. A: Math. Gen.* **39**, 833–876 (2006).
- [232] S. L. Adler and A. Bassi, "Collapse models with non-white noises", *J. Phys. A: Math. Theor.* **40**, 15083–15098 (2007).
- [233] V. P. Belavkin, "A new wave equation for a continuous nondemolition measurement", *Phys. Lett. A* **140**, 355–358 (1989).

- [234] A. Barchielli and V. P. Belavkin, "Measurements continuous in time and a posteriori states in quantum mechanics", *J. Phys. A: Math. Gen.* **24**, Preprint 1990, 1495–1514 (1991).
- [235] M. Bauer, T. Benoist, and D. Bernard, "Repeated quantum non-demolition measurements: convergence and continuous-time limit", *Annales Henri Poincaré* **14**, 639–679 (2013).
- [236] F. Schmolke, *Asymptotic Fate of Continuously Monitored Quantum Systems*, June 2025.
- [237] T. Karmakar, P. Lewalle, Y. Zhang, and K. B. Whaley, *Noise-Canceling Quantum Feedback: non-Hermitian Dynamics with Applications to State Preparation and Magic State Distillation*, July 2025.
- [238] E. H. Lieb and D. C. Mattis, "Ordering energy levels of interacting spin systems", *Journal of Mathematical Physics* **3**, 749–751 (1962).
- [239] A. Auerbach, *Interacting electrons and quantum magnetism* (Springer Science & Business Media, 2012).
- [240] A. C. Doherty and K. Jacobs, "Feedback-control of quantum systems using continuous state-estimation", *Phys. Rev. A* **60**, 2700–2711 (1999).
- [241] R. Ruskov and A. N. Korotkov, "Quantum feedback control of a solid-state qubit", *Phys. Rev. B* **66**, 041401 (2002).
- [242] J. Combes, H. M. Wiseman, and K. Jacobs, "Rapid measurement of quantum systems using feedback control", *Phys. Rev. Lett.* **100**, 160503 (2008).
- [243] S. Shankar, M. Hatridge, Z. Leghtas, K. Sliwa, A. Narla, U. Vool, S. M. Girvin, L. Frunzio, M. Mirrahimi, and M. H. Devoret, "Autonomously stabilized entanglement between two superconducting quantum bits", *Nature* **504**, 419–422 (2013).
- [244] S. Greenfield, L. Martin, F. Motzoi, K. B. Whaley, J. Dressel, and E. M. Levenson-Falk, "Stabilizing two-qubit entanglement with dynamically decoupled active feedback", *Phys. Rev. Appl.* **21**, 024022 (2024).

- [245] C. Sayrin, I. Dotsenko, X. Zhou, B. Peaudecerf, T. Rybarczyk, S. Gleyzes, P. Rouchon, M. Mirrahimi, H. Amini, M. Brune, et al., “Real-time quantum feedback prepares and stabilizes photon number states”, *Nature* **477**, 73–77 (2011).
- [246] F. Ticozzi and L. Viola, “Analysis and synthesis of attractive quantum markovian dynamics”, *Automatica* **45**, 2002–2009 (2009).
- [247] Y. Pan, V. Ugrinovskii, and M. R. James, “Ground-state stabilization of quantum finite-level systems by dissipation”, *Automatica* **65**, 147–159 (2016).
- [248] Z. Li, T. Roy, Y. Lu, E. Kapit, and D. I. Schuster, “Autonomous stabilization with programmable stabilized state”, *Nat. Commun.* **15**, 6978 (2024).
- [249] J. Guo, O. Hart, C.-F. Chen, A. J. Friedman, and A. Lucas, “Designing open quantum systems with known steady states: Davies generators and beyond”, *Quantum* **9**, 1612 (2025).
- [250] K. Aydo ğan, A. W. Schlimgen, and K. Head-Marsden, “Stabilizing steady-state properties of open quantum systems with parameter engineering”, *Phys. Rev. Res.* **7**, 023057 (2025).
- [251] K. R. Swain, R. K. Malla, and A. del Campo, “Noise-assisted feedback control of open quantum systems for ground state properties”, 10.48550/arXiv.2510.18984 (2025).
- [252] J. Wang and H. M. Wiseman, “Feedback-stabilization of an arbitrary pure state of a two-level atom”, *Phys. Rev. A* **64**, 063810 (2001).
- [253] J. Bell, “Against ‘measurement’”, *Physics World* **3**, 33 (1990).
- [254] M. Schlosshauer, “Decoherence, the measurement problem, and interpretations of quantum mechanics”, *Rev. Mod. Phys.* **76**, 1267–1305 (2005).
- [255] H. Munoz-Bauza and D. A. Lidar, “Scaling advantage in approximate optimization with quantum annealing”, *Physical Review Letters* **134**, 160601 (2025).

- [256] H. Schlömer, P. Bleicker, M. Dagrada, A. Klümper, J. Gemmer, and G. Morigi, “Quantum annealing with chaotic driver hamiltonians”, *Annals of Physics* **476**, 170042 (2025).
- [257] R. Ghosh, L. A. Nutricati, N. Feinstein, P. A. Warburton, and S. Bose, “Exponential speed-up of quantum annealing via n-local catalysts”, 10.48550/arXiv.2409.13029 (2024).
- [258] H. Leipold, F. M. Spedalieri, and N. Chancellor, “Quantum annealing with special drivers for circuit fault diagnosis”, *Scientific Reports* **12**, 10325 (2022).
- [259] P. Najafi, A. Villanueva, and H. Kappen, “Path integral quantum control for quantum chemistry applications”, 10.48550/arXiv.2509.24104 (2025).
- [260] Y. Okumura, K. Kashima, and Y. Ohta, “Iterative path integral approach to nonlinear stochastic optimal control under compound poisson noise”, *Asian Journal of Control* **19**, 781–792 (2017).

