Imperial College London

CONTROLLING OPEN QUANTUM SYSTEMS

Benjamin Dive

September 2017

Thesis submitted for the partial fulfilment of the degree of PhD

Centre for Doctoral Training in Controlled Quantum Dynamics
Department of Physics
Imperial College London

Supervised by Dr Daniel Burgarth & Dr Florian Mintert

DECLARATIONS

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

The research presented in this thesis was carried out over the last three years by me under the supervision of Daniel Burgarth and Florian Mintert. I hereby declare that the work is my own, except where otherwise acknowledged and suitably referenced. Some of the results contained have been published in the following papers:

- B. Dive, F. Mintert & D. Burgarth. "Quantum simulations of dissipative dynamics: Time dependence instead of size" *Phys. Rev. A* **92**, 032111 (2015).
- B. Dive, D. Burgarth & F. Mintert "Isotropy and control of dissipative quantum dynamics" *Phys. Rev. A* **94**, 012119 (2016).
- B. Dive, A. Pitchford, F. Mintert & D. Burgarth "*In situ* upgrade of quantum simulators to universal computers" *arXiv:1701.01723v2* (2017).
- B. Dive, A. Pitchford, F. Mintert & D. Burgarth "Learning tailored error correcting codes *in situ*" *in preparation* (2017).

ABSTRACT

All physical quantum systems are in contact with the external world which is inevitably a source of noise. It is therefore necessary to take into account this dissipation when designing controls that accomplish useful tasks in quantum information processing. This thesis is on that overlap of open quantum systems and control theory; it looks at what dynamics can happen with the use of external driving, and how this driving can be chosen to accomplish a desired goal.

The first result looks at how a given quantum dissipator can be manipulated, using coherent controls, into replicating the action of a different type of noise. This results in no-go theorems for how noise can be transformed based on its isotropy, with applications in simulating open systems. Another way of doing simulations is to use only unitary dynamics over the system and a finite dimensional ancilla, and it is proved that there always exists a dilation Hamiltonian that replicates the noisy dynamics continuously in time. This also highlights the fact that adding controls on noise can result in a different evolution than if the controls were done on the underlying system-environment level. A conjecture is introduced and studied which states that both approaches are equivalent in the special case of the controls commuting with the Lindbladian. A way around this difficulty is to use a quantum system to compute *in situ* which controls are best for achieving a desired task on itself. This problem is studied in the context of reaching entangling gates on a quantum simulator in order to upgrade it into a quantum computer. The experimental cost of doing so is found to be polynomial in the number of qubits in simulations. The same underlying principle is also used to find error correcting codes tailored to the dissipation in a system.

ACKNOWLEDGEMENTS

I would not have been able to complete this PhD and achieve what I did without the continued guidance, advice, and mentoring of my supervisors. Daniel Burgarth taught me everything I know about control theory and much beyond that; the opportunity to escape London and work in Aberystwyth several times a year is one I relished. Florian Mintert was in the unenviable position of being in the same floor as me and so being my first port of call when I was stuck on a problem, yet was always helpful. Working with both of them was delightful, and the occasional differences of opinion only encouraged more creative research and writing.

I am grateful to all my colleagues in the controlled quantum dynamics CDT and on the 12^{th} floor of EEE at Imperial, past and present, for the camaraderie and advice they provided. I am particularly thankful for the opportunity to get to know and learn from: Gregoire Brebner, Cristina Cirstoiu, Farhang Haddadfarshi, Zoe Holmes, Pavel Hrmo, David Jennings, Kamil Korzekwa, Marcel Langer, Matteo Lostaglio, Anthony Milne, David Newman, Thomas Nutz, Doug Plato, Terry Rudolph, John Selby, Jimmy Stammers, Jamie Tarlton, and Andrew Tranter. From Aberystwyth I am glad I have been able to work, and enjoy a beach barbecue, with Christian Arenz, Jukka Kiukas, and Alexander Pitchford. Special mention must go to Alex for our direct collaborations and for his invaluable work on developing the optimal control packages I have used repeatedly.

The opportunity I have had to engage with the research community at large has also given rise to many ideas and collaborations. I am grateful to Martin Fraas for fruitful discussions on the finer points of perturbation theory, to Rafal Demkowicz-Dobrzanski for proposing the Lindbladian and control commutativity conjecture, and to Paulo Facchi for his correspondence on the Friedrichs-Lee model. I am thankful to Stephen Glaser and David Leiner for their time in trying to develop an experimental realisation of our ideas in NMR, and to Syndey Schreppler for promising exchanges on a transmon implementation.

I extend my thanks to my parents for nurturing an intellectual curiosity in me and providing me with the means to do research. Being a graduate student was however not a constant source of fun; I am deeply appreciative of the kind patience of my supervisors and Richard Thompson, and I am grateful to my friends both inside and outside imperial who supported me through difficult times. I would not be where I am today without them.

CONTENTS

List Of Figures **8**

Glossary **9**

I Introduction **11**

I.A Overview 11

I.B Background on Open Quantum Systems 14

I.C Background on Quantum Control 28

II Isotropy and Control of Open Quantum Systems **37**

II.A Introduction 37

II.B Theorems 43

II.C Proofs 45

II.D Discussion 51

II.E Numerical Examples 54

II.F Relations to Thermodynamics 58

II.G Conclusion 59

III Minimal Dilations **61**

III.A Introduction 61

III.B Minimal Dilation Construction 64

III.C Proofs 69

III.D Examples 72

III.E Conclusion 75

IV Lindbladian and Control Commutativity Conjecture **77**

IV.A Introduction 77

IV.B Partial Results 81

IV.C Detailed Example 84

IV.D Conclusion 95

V *In Situ* Upgrade of Quantum Simulators to Computers 97

V.A Introduction 97

V.B The Scheme 101

V.C Local Gate Fidelity 103

V.D Numerical Investigation 108

V.E Conclusion 115

VI *In Situ* Error Correction 117

VI.A Introduction 117

VI.B The Scheme 121

VI.C Numerical Investigation 124

VI.D Conclusion 129

VII Conclusion 133

VII.A Summary 133

VII.B Outlook 134

Bibliography 137

LIST OF FIGURES

I.1 Action of a CPT map and its generator in state space 19

I.2 Illustration of dephasing and amplitude damping noise on the Bloch sphere 27

I.3 Steps in optimal control 35

II.1 Open quantum evolution as a flow in state space 38

II.2 Effective simulation of generators 54

II.3 Strength of reachability conditions on a generalised amplitude damping Lindbladian 55

II.4 Non-symmetric qutrit in a Λ configuration 56

II.5 Reachability of non-symmetric Λ qutrit systems 57

III.1 Method for constructing minimal dilations 64

III.2 Effect of introducing a finite cut off to the dilation of the phase-flip channel 73

III.3 Dilation for an amplitude damping channel with resonant driving 75

IV.1 Non-commutativity of tracing away the bath and adding controls 78

IV.2 Truncating Taylor expansion for calculating modified Friedrichs-Lee model . 89

IV.3 Commutativity of modified Friedrichs-Lee model with qubit rotation 90

IV.4 Violation of the LCC conjecture in modified Friedrichs-Lee model 91

V.1 Classical and quantum computer performing *in situ* control 98

V.2 Steps in finding an optimal set of controls *in situ* 102

V.3 Comparison of the gate fidelity with the local estimator of the fidelity . . . 107

V.4 Optimisation cost for different couplings and topologies 111

V.5 Scaling of optimisation cost in an Ising chain 112

V.6 Comparison of optimisation cost between neighbouring and distant qubits . 113

VI.1 Three-qubit Shor Code 119
VI.2 Encoding step for the five-qubit code 120
VI.3 Collective dephasing energy structure 121
VI.4 Steps for *in situ* error correction. 122
VI.5 Recovering the three-qubit Shor code 125
VI.6 Correcting unknown errors 126
VI.7 Recovering the five-qubit stabiliser code 127
VI.8 Obtaining a DFS from continuous collective dephasing noise 128
VI.9 Performing a Hadamard gate on a logical qubit 129

GLOSSARY

The following notation is used through out the thesis, and is summarised here for convenience.

Symbol	Name	Description
$\mathbb{H}$	Hilbert Space	
$\|\psi\rangle$	Pure State	$\|\psi\rangle \in \mathbb{H}$
d	Dimension	Dimension of the Hilbert space and states in it
ρ	Quantum State	$\rho = \sum p_i \|\psi_i\rangle \langle\psi_i\|$
$\mathcal{T}$	Time-ordering Operator	$\mathcal{T}[x(t_1)x(t_2)] = x(t_2)x(t_1) \quad t_1 \leq t_2$
$+h.c.$	Hermitian Conjugate	Shorthand to include the Hermitian conjugate
σ_i	Pauli Matrices	Canonical Pauli matrices, $\sigma_0 = \mathbb{I}$
$\sigma_+(\sigma_-)$	Raising (Lowering) Operator	$\sigma_+ \|0\rangle = \|1\rangle, \sigma_+ \|1\rangle = 0 \quad (\sigma_- \|0\rangle = 0, \sigma_- \|1\rangle = \|0\rangle)$
$\mathbb{I}$	Identity Operator	$\mathbb{I} \|v\rangle = \|v\rangle \quad \forall \ \|v\rangle$
$\mathcal{I}$	Identity Map	$\mathcal{I}(\rho) = \rho \quad \forall \ \rho$
$M(\cdot)$	Quantum Channel	$\rho(t) = M[\rho(0)]$
$\mathcal{U}(\cdot)$	Unitary Channel	$\mathcal{U}(\rho) = U\rho U^\dagger$
$G(\cdot)$	Generator	$\dot{\rho} = G(\rho)$
$\mathcal{H}(\cdot)$	Hamiltonian Generator	$\mathcal{H}(\rho) = -i[H, \rho]$
$L(\cdot)$	Lindbladian Generator	$L(\rho) = -i[H,\rho] + \sum_k \gamma_k \left[V_k \rho V_k^\dagger - \frac{1}{2}\{V_k^\dagger V_k, \rho\}\right]$
G_0, H_0, L_0	Drift	Uncontrollable part of the generator
K_k	Kraus Operator	$M(\rho) = \sum K_k \rho K_k^\dagger \qquad \sum K_k^\dagger K_k = \mathbb{I}$
$\|\Omega\rangle$	Bell Entangled State	$\|\Omega\rangle = \frac{1}{\sqrt{d}} \sum_{i=1}^{d} \|i\rangle \otimes \|i\rangle$
ρ_M	Choi State	$\rho_M = (M \otimes \mathcal{I}) \|\Omega\rangle \langle\Omega\|$
$\underline{S}$	Vectorised	Vectorised super-operator
$\tilde{S}$	Frame Change	Operator/super-operator in a different frame
$\overline{S}$	Unital	Unital block of super-operator
$\mathbb{U}(d)$ ($\mathbb{SU}(d)$)	(Special) Unitary Group	$UU^\dagger = \mathbb{I} \qquad \lvert\det U\rvert = 1 \quad (\det U = 1)$
$\mathbb{O}(d)$ ($\mathbb{SO}(d)$)	(Special) Orthonormal Group	$OO^T = \mathbb{I} \qquad \lvert\det O\rvert = 1 \quad (\det O = 1)$
$\mathfrak{u}(d)$ ($\mathfrak{su}(d)$)	(Special) Unitary Algebra	Lie algebra of $\mathbb{U}(d) \quad (\mathbb{SU}(d))$
$\mathfrak{o}(d)$ ($\mathfrak{so}(d)$)	(Special) Orthonormal Algebra	Lie algebra of $\mathbb{O}(d) \quad (\mathbb{SO}(d))$
$\boldsymbol{\lambda}(\cdot)$	Eigenvalues	Eigenvalues of an operator/super-operator
$\boldsymbol{\sigma}(\cdot)$	Singular values	Singular values of an operator/super-operator
$\prec$	Majorization	Partial order between sets (defined in Eq.(II.4))
$F(M,U)$	Gate Fidelity	$F(M,U) = \mathrm{Tr}[\rho_M \rho_U] = \langle\psi\|\rho_M\|\psi\rangle$
$F_{LE}(M,U)$	Local Fidelity Estimator	$F_{LE}(M,U) = 1 - \sum(1 - F(M_i, U_i))$

CHAPTER I

INTRODUCTION

I.A Overview

I.A.1 General Motivation

The laws of quantum mechanics are weird and wonderful. They predict things, that have been experimentally verified to astonishing precision, which cannot be explained by any physical model that obeys basic common sense principles [HBDRK15]. This is because such common sense about how the world ought to behave is based on experience of the macroscopic world of billiard balls, springs, and planets. Delving into the microscopic world of atoms, electrons, and photons requires a fundamentally new understanding of physics, and even logic.

Quantum mechanics has been around from the start of the previous century, but since Feynman's landmark paper [Fey82] there has been a drive to understand it from a computational viewpoint and to harness it to do useful things. This new field is loosely called 'Quantum Information' or 'Quantum Technology' depending on whether the focus is on a better understanding of quantum behaviour or on its application. It has had a remarkable growth since the turn of the millennium and is still accelerating; new journals in this area are appearing almost every year [EPJ14, IOP15, NPJ15, Ope17] and it has attracted large amounts of funding from both the public and private sector [EPS14, Eur16, MRNBD17, IBM17, Mic17].

The potential uses of quantum technology fall into three broad categories: cryptography, sensing, and computation. Of these cryptography is the most advanced as it is already commercially available [SGGRZ02] and has even been used in satellites [LYLSL17]. The basic idea is to generate a shared random key between two parties (which can then be used to perfectly encrypt information sent between them) in such a way that any potential eavesdropper leaves a statistical trail that gives them away. Quantum sensing involves the construction of exotic quantum states that are highly sensitive to some external parameter, such as gravity, allowing it to be measured far more accurately and efficiently than

otherwise possible. It is less developed but has already been used in gravitational wave detectors [AAAAA13] and is seen as an achievable short-term goal for more prosaic tasks such as detecting pipes beneath the ground [BBCFH16]. Lastly is quantum computation, the holy grail of the field. Algorithms have been developed for a quantum computer that can solve certain problems vastly faster than even the most powerful conventional computer, doing in a fraction of a second what would otherwise take billions of years [DE98]. A universal quantum computer that could run any such algorithm reliably is probably not going to be available in the immediate future, but more specialised ones that can simulate other quantum systems would also be of use, and already exist on a small scale [BR12, LBRdLM16, SSBCM16, OBKRM16, MMHHL16]. Even this more limited goal is of revolutionary significance as it would allow molecular or solid-state systems to be simulated, understood, and even designed to have specific properties.

Making machines capable of achieving these tasks is of course difficult; there are numerous theoretical and experimental hurdles to overcome. Many of these centre around two questions, "How do imperfections affect the system?" and "How can the system be manipulated?" The first question brings in the notion of open systems. These take into account the imperfections of the real world: noise arising from vibrations in the experimental set up, stray electric or magnetic fields, and general thermal effects from the rest of the universe interacting with the system at hand. Such effects can never be eliminated completely so it is important to understand their impact and how to mitigate them or even use them beneficially. The second question is the topic of control theory which addresses both controllability ("what control knobs does a machine need to accomplish a task?") and optimal control ("what is the best way of using these control knobs to do it?"). This thesis investigates the overlap of these two questions and asks what a quantum system can do in the presence of noise, and how it can be made to do it. Answering these questions not only paves the way to better quantum technology, but also sheds new light on natural quantum processes.

I.A.2 Context

As the research we did is on the overlap between open systems and control theory, it is worthwhile to first discuss what both terms mean in this context. Open quantum systems are those which are not isolated from the rest of the universe and so do not solely obey the Schrödinger equation. The system is governed by more than just a Hamiltonian and can undergo a broader range of evolutions which, crucially, can be used to take into account the myriad of ways that the environment affects it [BP02, RH11]. This normally acts to deteriorate and destroy the fragile quantum states that are useful for information processing and so is considered noise. The environment is everything that the system is in contact with but external to it. By its very nature it is something that cannot be controlled directly, if it could be then it would be considered part of the system itself rather than

an outside influence. For example, if an experimentalist could remove or set to a desired value all electromagnetic fields in the laboratory, then those fields would be taken as part of the system and not the environment. In this thesis, therefore, controlling open quantum systems means controlling systems in the presence of noise, rather than directly changing the noise in the system itself (as others have considered [SHMFKL12, BWSH16]).

There are nevertheless a lot of different ways of controlling a quantum system. A control is anything that can be tuned or applied to the system in order to modify its behaviour. The type of control most used in this thesis is Hamiltonian based: adding an extra (time-dependent) term to the Hamiltonian of the system in order to make it evolve in the desired way [D'A08]. This is usually done experimentally by creating an extra classical field, such as a laser pulse, that interacts with the system in a known and tunable way. The other type of control considered here is the use of ancillas; these are additional systems initialised in a certain state that interact with the system and are then discarded. For a set of allowed controls, whatever they may be, a natural aim is to find all possible tasks that can be accomplished: this is the topic of controllability or reachability [D'A08, Ell09]. As it is generally difficult to implement lots of different controls in an experiment, finding the minimal set which can do the required tasks is of practical use, as is showing that some systems are more flexible than others. The second aspect of control theory is optimal control, which finds a set of controls that reach a desired objective best in terms of fidelity or some other constraint. A common use for this is to find control pulses, such as the shape of a laser, used to drive a quantum system to do a particular logic gate [GBCKK15].

In this thesis, the control of open quantum systems is focussed on finite dimensional systems. That is, although the source of dissipation may be a bath in an infinite dimensional Hilbert space, the system being controlled itself resides in finite dimensions. This allows some of the results from the control of closed quantum systems to be applied, as well as the use of a matrix description of quantum mechanics. As most proposals for quantum technology make use of discrete degrees of freedom, this is also the most relevant for the field. The control tasks that are investigated are all on the level of operators (unitaries, quantum channels, or Lindbladians), rather than on creating a particular state. This is the most powerful for quantum information processing as it gives the ability to act on any input state in the desired way, allowing such operations to be chained together arbitrarily. It also reproduces the results of state control provided that the initial state is known.

The remainder of this chapter covers the general background necessary to follow the thesis. This includes how to describe open quantum systems mathematically, and some of their properties in §I.B; while major results and common methods in quantum control theory are introduced in §I.C. More specialised discussion of existing state of the art results are left to the beginning of the relevant chapters where our original research is presented. In §II the question of controllability in the presence of fixed noise is investigated and

results in no-go theorems where the anisotropy of the dissipative dynamics is a resource which can be used to reach different target operations. This can be used as a method to simulate other open systems, a problem which is also tackled in §III. In that case there is no intrinsic noise in the system, but the use of time-dependent Hamiltonians and ancillas allows open dynamics to be replicated continuously in time. This is also an example of how unitary controls can modify the noise felt by a system. A conjecture about when this does not happen, which also investigates some of the earlier assumptions made, is proposed in §IV where our investigation of it is discussed. The last two chapters take these, and previous, ideas about what can be done with quantum control and finds a powerful way of calculating what the required values of the optimal controls are. The idea is to let quantum machines learn how to control themselves; in §V this is used to find logic gates for a quantum simulator, and in §VI to implement tailor made error correcting codes.

I.B Background on Open Quantum Systems

Open quantum systems are, as mentioned above, those whose evolutions are not simply governed by unitary propagators and Hamiltonian generators due to their interaction with an external environment. This causes more general dynamics to take place than that described by the Schödinger equation, including loss of coherence and dissipation of information. Such lossy behaviour is termed noise and is usually an unwanted affect in experiments investigating quantum behaviour or in attempts to build quantum technologies. It is nevertheless present and hard to remove, making it important to take into account such effects explicitly when working with models of quantum systems. This allows protocols to be designed which either mitigate the effect of noise or use it as a resource, for example to cool the system or to simulate other open dynamics.

A review of the formalism used to describe open systems in terms of super-operators is given here in order to introduce to the reader the key physical and mathematical concepts that are used thereafter; this also serves to fix the notation that will be used. This introduction is based on a number of standard texts [BP02, BZ06, RH11]. It begins by discussing how to generalise the description of how closed quantum systems evolve in order to take into account external noise and loss of quantum information. An axiomatic definition of such quantum channels is given, which is then showed to be related to standard unitary evolution over a larger space. The question of whether such channels can be broken into small pieces such that a continuous evolution in time is well defined is addressed, which naturally gives rise to a definition of Markovianity for quantum systems. A derivation of how such dynamics can be calculated from an interaction between a system and its bath is then presented. Finally, some methods used to describe these linear operators as matrices, and their properties, are given. There exist other methods to describe open quantum systems that are not used in presenting our results later on, and so are not introduced,

including: quantum stochastic equations [HP84a, Gre01] and the SLH formalism [CKS17], and hierarchical equations of motion [XCLMY05, XXCLY07].

I.B.1 Generalising Unitary Evolution

In closed quantum systems, the state at a given time $|\psi(t)\rangle$ is related to the initial state $|\psi(0)\rangle$ according to $|\psi(t)\rangle = U|\psi(0)\rangle$, where U is a unitary propagator obeying $UU^\dagger = \mathbb{I}$. To generalise this to open systems, it is useful to first consider the same equation acting on a density operator ρ. These represent mixed states and are a convex mixture of pure states,

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i| \qquad \text{where} \quad p_i > 0, \quad \sum_i p_i = 1. \tag{I.1}$$

They can be interpreted as being a probabilistic mixture: the system has probability p_i of being in state $|\psi_i\rangle$, and the expectation value of some operator P over it is simply given by $\mathrm{Tr}[P\rho]$. From this definition, ρ has three important properties. It is Hermitian, positive semi-definite (all its eigenvalues are non-negative real numbers), and has trace one. Furthermore, all operators with these properties are valid states. Under this formalism, unitary evolution is expressed as

$$\rho(t) = U\rho(0)U^\dagger \equiv \mathcal{U}[\rho(0)], \tag{I.2}$$

where the last equality defines the super-operator $\mathcal{U}$, which is the Adjoint action of U on ρ. One of the most fundamental questions of open systems is to find what other propagators (also known as dynamical maps or quantum channels) are allowed that give physically sensible evolutions.

For a dynamical map M to be physical, it needs to map any quantum state to another quantum state. In addition to it being linear that requires three properties, similar to those given for quantum states,

$$M(\rho) = M(\rho)^\dagger \tag{I.3a}$$

$$\mathrm{Tr}[M(\rho)] = 1 \tag{I.3b}$$

$$M \otimes \mathcal{I}_d\,(\rho \otimes \mathbb{I}_d) \geq 0, \tag{I.3c}$$

for all ρ and where $\mathbb{I}$ is the identity matrix, $\mathcal{I}$ is the identity map, and in the last line an additional ancillary system of dimension d is added. The first condition simply states that M maps Hermitian operators to Hermitian operators, this is clearly necessary in order for $M(\rho)$ to be a valid state as defined above. Similarly for the second property of trace preservation.

The last condition is the most interesting one, in the case of $d = 1$ (such that the ancillary system is trivial) it defines a positive map: one that maps positive operators

to positive operators. The extension of it to an additional ancillary system of arbitrary dimension defines a completely positive map. The physical justification for it is that it is not enough to require that M maps an isolated physical state to another physical state; it must also be the case that applying the map to one system while doing nothing to another system also results in a physical state. It is a surprising result of quantum mechanics that not all positive maps are completely positive. Fortunately, it is not necessary to check over all possible multipartite input states and over all possible extensions of the dynamical map to determine complete positivity, it is sufficient to consider the single case of the ancilla having the same dimension as the system, and for the larger map to act trivially ($\mathcal{I}$) on it. Dynamical maps that satisfy all these properties are called completely positive and trace preserving (CPT).

A property that some, but far from all, quantum channels have is unitality. This is the preservation of the maximally mixed state ($\rho = \frac{1}{d}\mathbb{I}$): a quantum channel is unital if and only if

$$M(\mathbb{I}) = \mathbb{I}. \tag{I.4}$$

As the maximally mixed state represents a system with maximum uncertainty, and is invariant under unitary transformations, this is a property which comes up regularly in quantum information.

One way of expressing a quantum channel is in terms of its Kraus representation. This states that a dynamical map can always be written in the form

$$M(\rho) = \sum_{k=1}^{R} K_k \, \rho \, K_k^\dagger \qquad \text{where} \quad \sum_{k=1}^{R} K_k^\dagger \, K_k = \mathbb{I}, \tag{I.5}$$

where the K are the Kraus operators and R is the Kraus rank of the map and is upper bounded by d^2. This relation holds both ways; any set of Kraus operators for which $\sum K_k^\dagger K_k = \mathbb{I}$ gives rise a CPT map. From this it is easy to see that unitary evolution also falls under this framework, it consists of CPT maps with a Kraus rank of one. If the quantum channel is unital, then the Kraus operators also obey $\sum K_k K_k^\dagger = \mathbb{I}$.

An important relation between quantum channels and states is given by the Choi-Jamiołkowski isomorphism. This states that there is a one-to-one correspondence between the space of CPT maps and quantum states on the doubled space. The meaning of this is more readily apparent from the transformation that goes from a channel to an (unnormalised) state, which is

$$\rho_M = (M \otimes \mathcal{I}) \, |\Omega\rangle \langle\Omega| \qquad \text{where} \quad |\Omega\rangle = \sum_{i=1}^{d} |i\rangle \otimes |i\rangle \,, \tag{I.6}$$

where $\{|i\rangle\}$ is a basis set for the system and ancilla. The Choi state ρ_M of a CPT map M is therefore given by M acting on one half of a maximally entangled state. The isomorphism

can be interpreted as linking correlations in time (the dynamical map) to correlations in space (the state). One practical benefit of this isomorphism is that it provides a simple way to determine whether M is completely positive: M is completely positive if and only if ρ_M is positive semi-definite [BZ06].

I.B.2 Unitary Dilation

The previous section provided an axiomatic approach to determining which dynamical maps are physically sensible. Another approach is to follow the 'Church of the Larger Hilbert Space' as coined by John Smolin. This states that, ignoring measurements, the Schrödinger equation is all there is and that all evolutions are unitary. Things appear to be open because only a restricted subsystem is being observed, the underlying physics driving the system are unitary on a larger Hilbert space.

To formulate what this means for open systems it is necessary to first introduce the partial trace. This is defined over a bipartite system as the only linear operator for which

$$\mathrm{Tr}_B[X_A \otimes Y_B] = X \,\mathrm{Tr}[Y] \qquad \forall \quad X \in \mathbb{H}_A,\ Y \in \mathbb{H}_B, \tag{I.7}$$

where $\mathbb{H}_A$, $\mathbb{H}_B$ are the Hilbert spaces of the two parts of the system. That is, it 'traces out' one subsystem while preserving the other. Applied to a product state $\rho = \rho_A \otimes \rho_B$ it simply removes one of the subsystems while leaving the other unchanged. For more general states, such as entangled ones, it provides an average of the remaining state weighed by the one being traced out. Physically, it represents loss of information about the space being traced out and causes an entangled state to become a mixed one. It is a useful tool to focus the mathematical description of a quantum state to the part of the system that can be directly manipulated in the lab and remove the inaccessible environment. This also reduces the complexity of the description from one over potentially infinite degrees of freedom to the smaller system of interest.

With this, it is possible to formulate the open evolution of one system as a Stinespring dilation

$$M(\rho_S) = \mathrm{Tr}_A \left[U \rho_S \otimes \rho_A U^\dagger \right], \tag{I.8}$$

where ρ_S is the system of interest, ρ_A is an ancillary system and U is a joint unitary over both. This precisely formulates a quantum channel as unitary evolution on a large space followed by loss of information about part of the system. One restriction of this approach is that it requires the system and ancilla to start in a product state; this ensures that the resulting map is linear and well defined. There are attempts to go around these assumptions in [RRMAG10, Pec94], but this gives rise to new problems and remains controversial (as discussed in comments on the latter reference [Ali95, Pec95]).

A key result in the theory of open quantum systems is that the definitions of CPT

maps given in Eq.(I.3) and Eq.(I.8) are equivalent. Every CPT map that obeys the conditions given in Eq.(I.3) has a Stinespring dilation which allows it to be expressed as Eq.(I.8), where the dimension of the ancillary space is at least the Kraus rank of the channel. Moreover, the two descriptions are truly interchangeable in that all M which are of the form Eq.(I.8) are CPT maps. This is a reassuring result: it confirms that complete positivity and trace preservation are indeed sensible properties to ask for open quantum systems, and that everything can be considered as unitary on a large enough space.

I.B.3 Generators of CPT Maps

The previous two sections extended the concept of a propagator from the unitaries of closed quantum dynamics to the dynamical maps of open quantum systems. It is natural to do the same thing for another class of important operators - Hamiltonians. In closed systems these describe the dynamics of the system, that is, not the evolution from one point in time to another single point in time, but the whole path it takes in the Hilbert space. They are related to the evolution of unitaries according to ($\hbar$ is set to 1 throughout the thesis)

$$\frac{d}{dt}U(t) = -iHU(t) \qquad \text{or equivalently} \qquad U(t) = e^{-iHt}, \tag{I.9}$$

and give the evolution of states according to

$$\frac{d}{dt}\rho(t) = -i[H, \rho(t)] \equiv \mathcal{H}[\rho(t)], \tag{I.10}$$

where the last equality defines the super-operator $\mathcal{H}$, which is the adjoint action of H on ρ. In this case the unitaries are a function of time, so a Hamiltonian generates a one parameter group of unitary matrices (that is any set $\{U(t)\}$ which has the property $U(t_2)U(t_1) = U(t_1+t_2)$ for all $t \in \mathbb{R}$). Due to the Lie group property of unitaries (discussed in more detail in §I.C.1) the statement can also be reversed: any one parameter group of unitary matrices can be expressed as an exponential of a Hamiltonian. The situation is slightly more complex if the Hamiltonian is time-dependent

$$\frac{d}{dt}U(t) = -iH(t)U(t) \qquad \text{or equivalently} \qquad U(t) = \mathcal{T}\exp\left(-i\int_0^\tau d\tau H(\tau)\right), \tag{I.11}$$

where $\mathcal{T}$ is the time-ordering operator, but the principle result is the same: all unitaries are exponentials of Hamiltonians.

It is desirable to have an operator with similar properties in open systems, a generator which when exponentiated gives a CPT map, as shown in Fig.I.1. However this cannot always be done; the relation between propagator and generator is lost in the case of dynamical maps due to the requirement for complete positivity. In order to adapt the idea of generators to open systems it is necessary to see which quantum channels belong

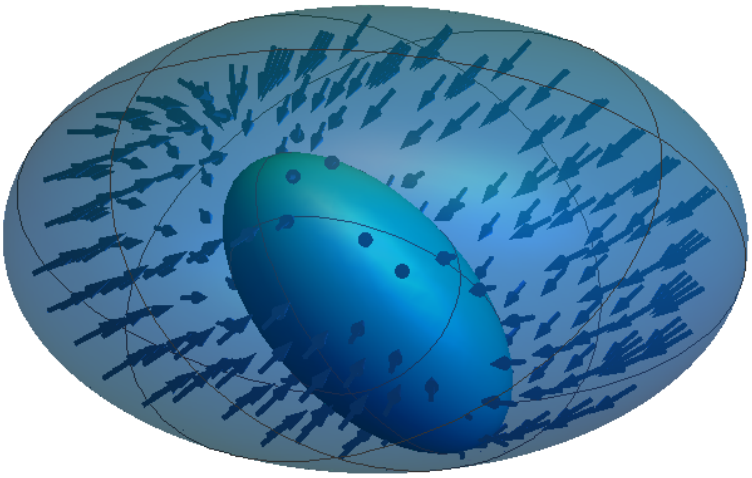


Figure I.1 | Illustration of the action of a CPT map and its generator on a high dimensional state space. The outer ellipsoid represents the state space at $t = 0$, while the inner ellipsoid is the same state space at a later time, after it has been acted on by a CPT map. The space has undergone a linear transformation and has shrunk. The generator for that quantum channel is represented by the arrows, it shows how the state space evolves at $t = 0$; integrating this flow leads to the recovery of the dynamical map.

to a one parameter semigroup

$$\{M_t : M_{t_2} M_{t_1} = M_{t_1+t_2} \ \ \forall\, t_1, t_2 \geq 0\}\,, \tag{I.12}$$

where $M_{t_2} M_{t_1}$ is a concatenation of maps. They will not form a group as non-unitary quantum channels cannot be undone by another quantum channels because loss of information is irreversible. This semigroup property is studied in [WC08] where different classes of CPT maps are introduced based on their divisibility. The maps that are most similar to unitaries are the infinitely divisible ones which have the property

$$M = M_n^n \quad \forall \ \ n \in \mathbb{N}, \tag{I.13}$$

stating that the map M is identical to a different map M_n (itself CPT) acting n times. It might be expected that such a property would imply that the map is of the form $M_t = e^{Lt}$. This is indeed the case. Furthermore the generator L has a specific structure,

$$L(\rho) = -i[H, \rho] + \sum_k \gamma_k \left[V_k \rho V_k^\dagger - \frac{1}{2}\{V_k^\dagger V_k, \rho\} \right], \tag{I.14}$$

where H is an ordinary Hamiltonian, $\gamma_k > 0$, the V_k are referred to as Lindblad or jump operators, and L itself is the Lindbladian [Kos72, GKS76, Lin76]. The form of the dynamical map as the exponential of this gives rise to the celebrated Gorini-Kossakowski-

Sudarshan-Lindblad equation on either the operator or state level

$$\frac{d}{dt}M_t = LM_t \qquad\qquad \frac{d}{dt}\rho(t) = L\left(\rho(t)\right). \tag{I.15}$$

Such evolution is called Markovian, although different definitions of the word exist based on different concepts [BLP09, RHP14, BD16]. For the purpose of this thesis, open systems described by Eqs.(I.14-I.15) are called time-independent Markovian.

A slight generalisation of the infinitely divisible channels are the infinitesimally divisible ones. These can also be expressed as an infinite concatenation,

$$M = \prod_{i=1}^{n} M_i \quad \forall \;\; n \in \mathbb{N}, \tag{I.16}$$

but the maps into which they are broken down are no longer required to be identical (but are still CPT). Such dynamical maps also belong to a semigroup and have a generator, but it is now time-dependent. In the same way as for time-dependent Hamiltonians, the evolution of the propagated can be described by

$$\frac{d}{dt}M_t = L_t M_t \qquad \text{or equivalently} \qquad M(t) = \mathcal{T}\exp\left(-i\int_0^t d\tau L_\tau\right). \tag{I.17}$$

The generator, L, has the same structure as in Eq.(I.14) except that the γ and V_k are allowed to be time-dependent. Unsurprisingly, such systems are referred to as time-dependent Markovian.

Not every quantum channel, however, has these properties. Even for qubits (quantum bits with $d = 2$), the majority of maps are not Markovian [WECC08]. The most extreme are the non-divisible CPT maps, they are those which cannot be expressed as the concatenation of two CPT maps

$$M \neq M_2\, M_1 \tag{I.18}$$

where neither M_1 or M_2 are purely unitary.

From the divisibility properties required of them, the Lindblad generators given here give rise to CPT maps between any two points in time, that is, the propagator that maps the system from time t_1 to t_2,

$$M_{t_2,t_1} = \mathcal{T}\exp\left(-i\int_{t_1}^{t_2} d\tau L_\tau\right), \tag{I.19}$$

is CPT for all $t_2 > t_1$. It is possible to impose a weaker requirement, that only the CPT map acting from $t = 0$ to a later time is CPT and that the dynamics are time-local

$$M_t = \mathcal{T}\exp\left(-i\int_0^t d\tau G_\tau\right), \tag{I.20}$$

where G is a more general generator not necessarily of Lindblad form. The physical reasoning behind considering such dynamics is that it allows memory effects, where the evolution of the system at t is influenced by its past evolution. Although a much broader ranged of CPT maps admit such a form (though still not all of them); adding a Hamiltonian $\mathcal{H}$ on the system on top of G in Eq.(I.20) can result in loss of complete positivity for certain evolution times. It is therefore generally not clear what the physical meaning of G is in this case and, especially in a control scenario where Hamiltonians are commonly added, need to be treated with care. Nevertheless, they are a useful way of describing a wide-class of non-Markovian systems.

Unitality as defined in Eq.(I.4) was a useful property for quantum channels as it categorised the maps that left the maximally mixed state invariant. It is similarly convenient to define it for generators. The desired behaviour is for unital generators to give rise to only unital CPT maps, which naturally leads to the definition

$$G(\mathbb{I}) = 0, \tag{I.21}$$

for unital generators, and the same for unital Lindbladians. Hamiltonian evolution is, of course, always unital.

I.B.4 Microscopic derivation of Lindbladians

The generators of Lindblad type, introduced in the previous section in Eq.(I.14), give rise to semigroups of CPT maps and therefore occupy a similar role in open quantum dynamics as Hamiltonians play for closed systems. The physical meaning of this generator, or how it arises from an interaction with another system in the same way as CPT maps arise from a unitary on a larger space, is not at all evident from the definition given previously. It is however possible, with certain assumptions, to start with a system interacting with a bath via a Hamiltonian and derive a Lindbladian on the system alone; one way of doing so is shown below following well known methods [BP02, RH11].

The starting point is to consider the system to be in contact with a bath. A bath in this context means an environment or an ancilla with a great number of degrees of freedom, such as a quantum field, with which the system is coupled. It is also typically in a thermal state. The system and the bath have the joint Hamiltonian

$$H = H_0 + H_I = H_S \otimes \mathbb{I} + \mathbb{I} \otimes H_B + H_I, \tag{I.22}$$

where the first two terms act solely on the system and bath respectively while the latter is an interaction between them. This interaction term is treated separately because the isolated dynamics of the system and bath are assumed to be known and solved. As in the case of Eq.(I.8), the initial state is assumed to be in a product state, $\rho(0) = \rho_S(0) \otimes \rho_B(0)$.

Without loss of generality, it is also possible to assume that $\mathrm{Tr}[H_I\rho(0)] = 0$ by shifting terms between the Hamiltonians.

The first step to calculate an equation of motion for $\rho_S(t)$ only is to move to the interaction picture, denoted by a tilde,

$$\tilde{H}_I(t) = e^{iH_0t}H_Ie^{-iH_0t}, \tag{I.23}$$

and $\tilde{\rho}(0) = \rho(0)$. The Schrödinger equation applied to density operators give rise to the Liouville-Von Neumann equation, Eq.(I.10); using this in the interaction picture and formally integrating it with respect to time gives

$$\tilde{\rho}(t) = \rho(0) - i\int_0^t [\tilde{H}_I(\tau), \tilde{\rho}(\tau)]d\tau. \tag{I.24}$$

This can then be substituted back into the Liouville-Von Neumann equation, essentially doing a series expansion up to the second term

$$\frac{d}{dt}\tilde{\rho}(t) = -i[\tilde{H}_I(t), \rho(0)] - \int_0^t [\tilde{H}_I(t), [\tilde{H}_I(\tau), \tilde{\rho}(\tau)]]d\tau \tag{I.25}$$

$$\frac{d}{dt}\tilde{\rho}_S(t) = -i\mathrm{Tr}_B\left[[\tilde{H}_I(t), \rho(0)]\right] - \int_0^t \mathrm{Tr}_B\left[[\tilde{H}_I(t), [\tilde{H}_I(\tau), \tilde{\rho}(\tau)]]\right]d\tau, \tag{I.26}$$

where the bath was traced out to get to the last line.

To proceed further a number of different assumptions have to be made which together form the Born-Markov approximation. First, the bath is assumed to be in a fixed state for during the entire evolution: $\rho(t) = \rho_S(t) \otimes \rho_B$. This cannot be exactly true if the two systems are truly interacting and ρ_S is becoming less pure, but it is nevertheless a good approximation in many cases. This is because the bath is considered to be large compared to the system, and so any back-action caused to it is small. Due to this, the system is memoryless as correlations built up between the system and the bath decay much faster that the dynamics of the system. This is similar to the discussion in the previous section where the evolution of the open system depended only on its current state and not on the past; this allows $\rho_S(\tau)$ inside the integral of Eq.(I.26) to be replaced by $\rho_S(t)$ and extends the limit of the integral to $+\infty$. Finally, by noting that $\mathrm{Tr}_B[[\tilde{H}_I(\tau), \rho(0)]] = 0$ due to the definition of H_I, the dynamics of $\rho_S(t)$ can be expressed as the Markovian master equation

$$\frac{d}{dt}\tilde{\rho}_S(t) = -\int_0^\infty \mathrm{Tr}_B\left[[\tilde{H}_I(t), [\tilde{H}_I(t-\tau), \tilde{\rho}_S(t) \otimes \rho_B]]\right]d\tau. \tag{I.27}$$

This is an equation of motion for $\rho_S(t)$ (albeit in the interaction picture) which only depends on its current state, the initial state of the bath, and the Hamiltonian coupling them together. This is not quite yet sufficient to show that this equation describes the generator of a CPT semigroup.

Some further manipulations, as well as a secular approximation, is required to obtain the Lindblad equation, which are sketched out here. The interaction Hamiltonian, H_I can always be decomposed as a sum of tensor products of Hermitian operators

$$H_I = \sum_\alpha A_\alpha \otimes B_\alpha. \tag{I.28}$$

This caries forward in the interaction picture

$$\tilde{H}_I(t) = \sum_\alpha e^{iH_S t} A_\alpha e^{-iH_S t} \otimes e^{iH_B t} B_\alpha e^{-iH_B t} = \sum_\alpha \tilde{A}_\alpha(t) \otimes \tilde{B}_\alpha(t). \tag{I.29}$$

It is advantageous to transform the system operators $A_\alpha(t)$ to be functions of frequency, ω, such that

$$[H_S, A_\alpha(\omega)] = -\omega A_\alpha(\omega) \qquad \text{and} \qquad [H_S, A_\alpha^\dagger(\omega)] = \omega A_\alpha^\dagger(\omega), \tag{I.30}$$

which allows the interaction Hamiltonian to be recast as

$$\tilde{H}_I(t) = \sum_{\alpha,\omega} \tilde{A}_\alpha(\omega) \otimes \tilde{B}_\alpha(t), \tag{I.31}$$

where

$$\langle B_\alpha(t) \rangle = \mathrm{Tr}[B_\alpha(t)\rho_B] = 0, \tag{I.32}$$

from the original constraint that $\mathrm{Tr}[H_I \rho(0)] = 0$. Substituting this into Eq.(I.27) gives, after some algebraic manipulation,

$$\frac{d}{dt}\tilde{\rho}_S(t) = -\sum_{\omega,\omega'}\sum_{\alpha,\beta} e^{-i(\omega'-\omega)t}\Gamma_{\alpha,\beta}(t,\omega)\left(A_\beta(\omega)\tilde{\rho}_S(t)A_\alpha^\dagger(\omega') - A_\alpha^\dagger(\omega')A_\beta(\omega)\tilde{\rho}_S(t)\right) + h.c., \tag{I.33}$$

where $h.c.$ is the Hermitian conjugate and the Γ are the one sided Fourier transform of the bath correlation functions

$$\Gamma_{\alpha,\beta}(t,\omega) = \int_0^\infty d\tau e^{i\omega\tau}\mathrm{Tr}[B_\alpha^\dagger(t)B_\beta(t-\tau)\rho_B]. \tag{I.34}$$

From the Born-Markov approximation, the state of the bath is fixed, therefore the correlation function are unchanged by a shift in time and so the previous expression simplifies to

$$\Gamma_{\alpha,\beta}(\omega) = \int_0^\infty d\tau e^{i\omega\tau}\mathrm{Tr}[B_\alpha^\dagger(\tau)B_\beta(0)\rho_B]. \tag{I.35}$$

The final approximation to make is to only consider the terms in Eq.(I.33) where $\omega = \omega'$. This is a rotating-wave (also called secular) approximation similar to Born-Markov made earlier: terms oscillating rapidly compared to the evolution of the system are averaged out

and ignored. Put together this yields

$$\frac{d}{dt}\tilde{\rho}_S(t) = -\sum_{\omega}\sum_{\alpha,\beta}\Gamma_{\alpha,\beta}(\omega)\left(A_\beta(\omega)\tilde{\rho}_S(t)A_\alpha^\dagger(\omega') - A_\alpha^\dagger(\omega')A_\beta(\omega)\tilde{\rho}_S(t)\right) + h.c. \tag{I.36}$$

The last step is to transform out of the interaction picture and split the Γ into Hermitian and anti-Hermitian parts

$$\Gamma_{\alpha,\beta}(\omega) = \frac{1}{2}\left(\gamma_{\alpha,\beta} + iS_{\alpha,\beta}(\omega)\right). \tag{I.37}$$

Finally, this gives

$$\begin{aligned}
&\frac{d}{dt}\rho_S(t) = -i[H_{LS} + H_S, \rho_S(t)] + D(\rho_S(t)] \qquad \text{(I.38)}\\
&\text{where} \quad H_{LS} = \sum_{\omega}\sum_{\alpha,\beta} S_{\alpha,\beta}(\omega)A_\alpha^\dagger(\omega)A_\beta(\omega)\\
&\qquad\qquad D(\rho_S) = \sum_{\omega}\sum_{\alpha,\beta}\gamma_{\alpha,\beta}(\omega)\left[A_\beta(\omega)\rho_S A_\alpha^\dagger(\omega) - \frac{1}{2}\{A_\alpha^\dagger(\omega A_\beta(\omega), \rho_S\}\right].
\end{aligned}$$

H_{LS} is the Lamb shift, which is the extra energy the system has due to its interaction with the environment. The dissipative part of the generator can be diagonalised and brought into the standard form of Eq.(I.14). It can be shown from properties of the Fourier transform and bath correlation functions that $\gamma_{\alpha,\beta}(\omega)$ is positive semi-definite, and therefore the decay rates are all positive. The result of Eq.(I.38) is therefore an evolution in time of the reduced system governed by a generator of Lindblad type, which was precisely the class of generators that gave rise to CPT semigroups.

This derivation shows that there is a clear link between the mathematical properties required by the generators of CPT semigroups, and the evolution of open systems interacting with a large bath that satisfies certain approximations. In the same way as Stinespring dilations provide a link between CPT maps and unitaries on a larger space, this is a reassuring result that shows mathematical and physical approaches converge. It also highlights the assumptions that go into such a model and the regimes where open systems are not Markovian. Indeed, it should be noted that the Born-Markov approximations as described above are not fully satisfactory. One problem is that they are an approximation and not an expansion, it is therefore not clear how higher order terms could be calculated in order to check that they are sufficiently small. This can be circumvented by more detailed calculations in some models [DL05]. A different problem is that the approximation made leads to exponential decay, which is approximately linear for short times. Yet any unitary dynamics can only lead to quadratic decay of population, as diagonal terms in the density matrix vary only in second order. This shows that the Born-Markov approximation result in some coarse-graining of the dynamics and should not be expected to resolve details happening on time scales shorter than the bath correlation time.

I.B.5 Vectorisation

The formalism used for finite dimensional closed quantum systems is usually linear algebra, specifically: states are column vectors and operators are matrices. This is a very convenient formulation both for algebraic manipulations and for numerical simulations. The way to do this for open quantum systems is to transform density matrices into vectors, and the super-operators that act on them (quantum channels and generators) into matrices. This can be done with a technique called vectorisation [BZ06], which has close links with tensor network representations [CD11, WBC15].

The starting point is to work in the generalised Bloch representation, where a quantum state ρ is represented as the real vector

$$|\rho\rangle = (x_1, x_2, x_3, ..., x_{d^2})^T \qquad \text{where} \quad x_i = \mathrm{Tr}[\sigma_i \rho], \tag{I.39}$$

are the expectation values over an orthonormal set $\{\sigma_i\}$ which forms a basis over all operators. The inverse map is simply given by

$$\rho = \sum_i x_i \sigma_i. \tag{I.40}$$

The same thing can be done for a super-operator S, where its vectorised form is denoted by $\underline{S}$, according to

$$\underline{S} = S_{ij} \left|i\right\rangle \left\langle j\right| \tag{I.41a}$$

$$S_{ij} = \mathrm{Tr}[\sigma_i^\dagger S(\sigma_j)] \tag{I.41b}$$

$$S(\rho) = \sum_{i,j} \sigma_i S_{ij} \rho_j \tag{I.41c}$$

$$|S(\rho)\rangle = \underline{S} \left|\rho\right\rangle . \tag{I.41d}$$

This is a useful representation as the vectorised form $\underline{S}$ has the same linear algebra properties as the parent super-operator S. The spectral properties (eigenvalues, singular values, trace and determinant) of S and $\underline{S}$ are the same, the dual of a super-operator $S^\dagger$, defined according to $\mathrm{Tr}[\mu S(\rho)] = \mathrm{Tr}[S^\dagger(\mu)\rho]$, has as its matrix representation the Hermitian conjugate of the matrix representation of the original super-operator, such that $\underline{(S^\dagger)} = (\underline{S})^\dagger$, and the concatenation of two super-operators is given by the matrix product of their vectorised form.

One simple choice of basis for the $\{\sigma_i\}$ is row-vectorisation. All the elements of σ_i are 0 except for a single one, which has value 1; the basis is enumerated by the location of that element, going along each row and then downwards. Using such a basis, a general

qubit state is

$$\rho = \begin{pmatrix} \rho_{00} & \rho_{01} \\ \rho_{10} & \rho_{11} \end{pmatrix} \quad \longrightarrow \quad |\rho\rangle = (\rho_{00}, \rho_{01}, \rho_{10}, \rho_{11})^T, \tag{I.42}$$

such that the rows of the matrix are simply listed one after the other. A nice property of this basis, known as Roth's Lemma, is

$$ABC \longrightarrow A \otimes C^T |B\rangle\,, \tag{I.43}$$

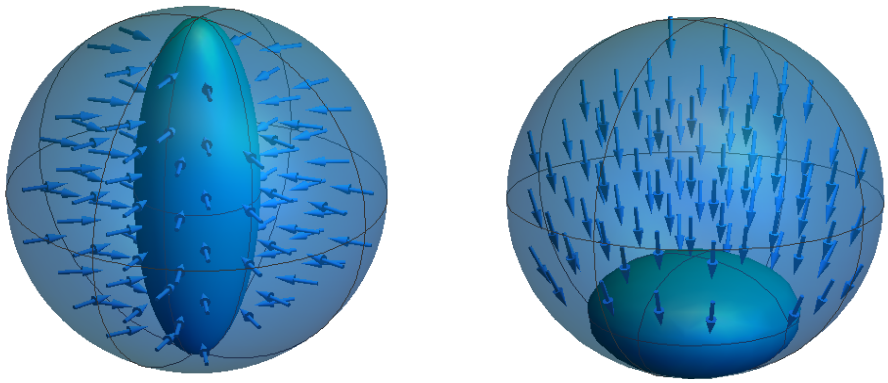
which is particularly powerful to vectorise entire equations. The Choi-Jamiołkowski isomorphism of Eq.(I.6) is also very straightforward in this representation and can be done by 'reshuffling' elements, $\rho_M = M^R$, where reshuffling is defined as

$$S^R_{m\mu,n\nu} = S_{mn,\mu\nu}, \tag{I.44}$$

where the first two indices label the row element over the bipartite system, and the last two the column element. This is eminently more clear with an example for a qubit channel and a two-qubit Choi state:

$$\rho_M = \underline{M}^R = \begin{pmatrix} M_{00} & M_{01} & M_{10} & M_{11} \\ M_{02} & M_{03} & M_{12} & M_{13} \\ M_{20} & M_{21} & M_{30} & M_{31} \\ M_{22} & M_{23} & M_{32} & M_{33} \end{pmatrix}. \tag{I.45}$$

Another particularly useful choice of for the $\{\sigma_i\}$ are traceless Hermitian matrices for $i = 2, ..., d^2$, and $\sigma_1 = \frac{1}{\sqrt{d}}\mathbb{I}$. The Pauli matrices and the identity are the natural choice for such a basis, which highlights that vectorisation is just a Bloch representation for higher dimensions. A consequence of this is that Hermitian operators and super-operators are real when vectorised this way. As they are particularly important, the form that the super-operators of closed dynamics take in this representation are described here, which touches on elements of Lie groups discussed below in §I.C.1. Unitary propagators, $\mathcal{U}(\cdot) = U \cdot U^\dagger$, become matrices in the defining representation of the rotation group $\mathbb{I}_1 \oplus \mathbb{SO}(d^2 - 1)$. Hamiltonian generators, $\mathcal{H}(\cdot) = -i[H, \cdot]$, are in the corresponding Lie algebra, $0_1 \oplus \mathfrak{so}(d^2 - 1)$, which consist of real antisymmetric (and therefore traceless) matrices. In the case of $d = 2$, unitary propagators form all such matrices (and like-wise for Hamiltonians), but in higher dimension they only form a subgroup. This means that in $d = 2$ only, the set of every unitary on the system corresponds to all possible rotations of the state vector $|\rho\rangle$. In higher dimensions however, there exist rotations of this vector which cannot be induced by any Hamiltonian controls on the system; this is analoguous to saying that not every vector in the generalised Bloch space is a valid quantum state.



(a) Dephasing **(b)** Amplitude Damping

Figure I.2 | Illustration of the Bloch sphere under two common types of qubit noise. The first is a unital dephasing channel, which preserves mixtures of the $|0\rangle$ and $|1\rangle$ states but destroys superpositions. The second is an amplitude damping channel where all states decay to the ground state $|0\rangle$.

I.B.6 Common types of qubit noise

There are two different types of common qubit noise that are typical of dissipative processes: dephasing and amplitude damping. These, or variants of these, will be used throughout this thesis as examples and it is useful to have a visualisation of what they do. A convenient way of picturing their effect is to plot how they transform the Bloch sphere, as is done in Fig.I.2.

Possibly the simplest type of qubit noise is dephasing, where the probability of being in two states (usually picked as $|0\rangle$ and $|1\rangle$) is conserved, but phase information between them is lost. This type of noise is also sometimes called pure decoherence as it does not degrade classical information, as such it is unital, and in many cases is the dominant type of noise. One physical way it can occur is when the energy splitting in a qubit fluctuates in some unknown way, and it corresponds to the T_2 decoherence time commonly quoted in NMR. The Lindbladian for it is

$$L(\rho) = -\gamma[\sigma_z, [\sigma_z, \rho]], \tag{I.46}$$

and a state $|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$ evolving under it will tend to

$$\rho = |\alpha|^2\,|0\rangle\langle 0| + |\beta|^2\,|1\rangle\langle 1|\,, \tag{I.47}$$

with the off-diagonal terms decaying exponentially, as is sketched in Fig.I.2a. In a microscopic derivation of the type done in §I.B.4, such noise typically arises from the interaction between the system and the low frequency part of the bath.

Another common type of noise is amplitude damping, where all points in the Bloch sphere decay to the ground state. Unlike dephasing, such noise has a single fixed point and

destroys classical, as well as quantum, information. It is the typical example of non-unital qubit noise and in many ways is as different to dephasing as possible; phase information is lost by the system at the minimum rate, constrained by the decay to the ground state and complete positivity. In can be interpreted as a qubit cooling down to absolute zero, or as an atom undergoing spontaneous emission. The Lindbladian for it is

$$L(\rho) = -\gamma\left(\{\sigma_+\sigma_-, \rho\} - 2\,\sigma_-\rho\sigma_-\right), \tag{I.48}$$

and a state evolves, as shown in Fig.I.2b, according to

$$\begin{pmatrix} |\alpha|^2 & \alpha\beta^* \\ \alpha^*\beta & |\beta|^2 \end{pmatrix} \longrightarrow \begin{pmatrix} |\alpha|^2 + |\beta|^2(1-e^{-2\gamma t}) & \alpha\beta^* e^{-\gamma t} \\ \alpha^*\beta e^{-\gamma t} & |\beta|^2 e^{-2\gamma t} \end{pmatrix}. \tag{I.49}$$

While the T_2 noise described above comes from the low frequency part of the bath, energy conservation constraints show that the T_1 type of noise described by amplitude damping arises from the modes of the bath which are close to resonance with the qubit internal energy splitting.

I.C Background on Quantum Control

Control theory, in the context of quantum mechanics, is about how the ability to vary certain parameters in a system allows different operations to be reached. There are two main types of controls considered throughout this thesis: Hamiltonians that can be varied in time, and the ability to add and discard ancillary systems. The latter is completely binary and is used exclusively at the start or end of the dynamics; it is therefore considered separately from the ability to vary Hamiltonians during the evolution of the system.

In this formulation, a control system consists of a Hamiltonian

$$H(t) = H_0 + \sum_k f_k(t) H_k, \tag{I.50}$$

where H_0 is the drift Hamiltonian, the H_k are the control Hamiltonians, and the $f_k(t)$ are called the control parameters or pulse interchangeably. Control theory investigates how the choice of the controls affect the unitary that can be reached

$$U(t) = \mathcal{T} \exp\left(-i \int_0^t H(\tau) d\tau\right). \tag{I.51}$$

In the common case of piecewise-constant controls, the Hamiltonian itself is also piecewise-constant and so the above expression simplifies to

$$U(t) = \prod_{n=0}^{N} \exp\left[-iH\left(\frac{nt}{N}\right)\frac{t}{N}\right], \tag{I.52}$$

where N is picked such that the Hamiltonian is constant over the interval T/N and the n are sorted with the lowest value to the right in the product.

Control theory can also be used in the context of state control, where the objective is to map a state ρ_1 to ρ_2. The work in this thesis focusses entirely on operator control so this branch of control theory is not presented here, the methods used in any case are very similar and the result is that operator controllability implies state controllability (but not the reverse). It also possible to do feedback control, where the system is continuously measured as it is evolves and the controls are adjusted based on the results [DJ99, DHJMT00, WD05]. This is not a method considered in this thesis, although the *in situ* method proposed in §V and §VI does involve measuring the system after it has finished evolving.

The principle results of controllability of closed quantum systems, preceded by a brief introduction of the key requisite notions of Lie groups and algebras, are presented below following the lines of [D'A08, Ell09, Nak03] and the author's MRes thesis on the same project [Div14]. A short discussion of why this approach does not work in the case of either open or infinite dimensional systems is also included. This is followed by a way of finding what values the control parameters need to take to achieve a given goal, the topic of optimal control, focussing on a particular algorithm GRAPE [KRKSHG05] implemented in Python on QuTiP [JNN13, JNPGG16] that we used repeatedly in our research.

I.C.1 Some properties of Lie groups and algebras

The dynamics of closed quantum systems are, as has already been touched on in §I.B, highly dependent on Lie properties linking Hermitian operators (the Hamiltonians) to unitary operators (the propagators). The fundamental representation of the unitary group $\mathbb{U}(d)$ consists of d dimensional complex matrices U, with the property that $UU^\dagger = \mathbb{I}$ and $|\det U| = 1$. Hermitian matrices, H, are trivially those for which $H = H^\dagger$. An important concept for control theory is that $\mathbb{U}$ is a Lie group, rather than just a group, and therefore the corresponding Hamiltonians form a Lie algebra.

A Lie group is a group where every element is also a point on a smooth manifold, such that the group operation and the inverse are smooth maps on the manifold. The matrix groups, such as $\mathbb{U}(d)$, have this property. To define a lie algebra it is necessary to consider a curve $X(t)$ on a Lie group G such that $X(t) \in G \; \forall \; t \in \mathbb{R}$ and $X(0) = \mathbb{I}$. The direction explored by the curve at $t = 0$ is defined by

$$x = \frac{d}{dt} X(t)|_{t=0}. \tag{I.53}$$

Taking all possible curves $X(t)$ on G leads to many different x which together form the Lie algebra $\mathfrak{g}$, also known as the tangent space of G. This Lie algebra is a vector space

over $\mathbb{R}$ which has an additional map

$$[\cdot,\cdot] : \mathfrak{g} \times \mathfrak{g} \to \mathfrak{g}, \tag{I.54}$$

called the Lie bracket. This obeys a series of properties

$$[\alpha x + y, z] = \alpha[x, z] + [y, z] \tag{I.55a}$$

$$[x, y] = -[y, x] \tag{I.55b}$$

$$[x, [y, z]] + [y, [z, x]] + [z, [x, y]] = 0, \tag{I.55c}$$

which are all satisfied by the usual matrix commutator $[A, B] = AB - BA$. It is worth pointing out that the dimension of G as group and as a manifold is equal to that of $\mathfrak{g}$.

In the same way that Eq.(I.53) provides a map from a curve on the Lie group to an element of the associated algebra, the reverse can be done by defining the exponential map as

$$\exp(xt) = X(t), \tag{I.56}$$

where $X(t)$ is a one parameter subgroup of G. The usual matrix exponential satisfies this property. Just as all curves on G give all elements of $\mathfrak{g}$, the Lie group can be generated by the Lie algebra according to

$$G = e^{\mathfrak{g}} = \{e^{x_1} e^{x_2} ... e^{x_m} \mid x_1, x_2, ..., x_m \in \mathfrak{g}\}. \tag{I.57}$$

A visualisation of this is given by a heuristic geometrical picture. Each element of the Lie group G corresponds to a point on a smooth differentiable manifold. The Lie algebra $\mathfrak{g}$ corresponds to the set of tangent vectors at the identity element of G: the different directions that can be explored from the identity. As such, an element of $\mathfrak{g}$ leads to a vector flow in G; by starting at the identity and moving along it a one parameter subgroup of G is generated. The commutator $[x_1, x_2]$ is the vector displacement of moving an infinitesimal amount along x_1 followed by x_2, compared to moving along x_2 followed by x_1; this itself is a direction in G and therefore an element of $\mathfrak{g}$.

The link between this and quantum mechanics comes by identifying G with $\mathbb{U}(d)$, the trajectory $X(t)$ with the propagator $U(t)$, and elements of $\mathfrak{g}$ with Hamiltonians $-iH$. The presence of the $-i$ is one that depends on convention, mathematicians often prefer to work directly with elements of the Lie algebra and so define Hamiltonians to be anti-Hermitian; this will not be done here in order to maintain the usual physical significance of Hamiltonians as energy operators. The result is that, to apply Lie theory to quantum mechanics, $-i$ needs to be inserted in front of x in Eq.(I.53) and Eq.(I.56). Indeed, doing this shows how similar those equations are to the operator version of the Schrödinger equation.

Adopting this convention, the Lie algebra associated with the unitary group $\mathbb{U}(d)$ is $\mathfrak{u}(d)$ which is composed of all d dimensional Hermitian matrices. As global phases are of no physical importance in quantum mechanics, it is often useful to work with slightly different groups. These are the special unitaries $\mathbb{SU}(d)$ which have the stronger requirement that $\det U = 1$. The associated algebra is $\mathfrak{su}(d)$ and is formed of all traceless Hermitian matrices. The equivalent groups also exist in the case of real matrices (as encountered in §I.B.5): the special orthogonal group $\mathbb{SO}(d)$ consists of real rotation matrices O such that $OO^T = \mathbb{I}$ and $\det O = 1$, and its algebra $\mathfrak{so}(d)$ made up of antisymmetric real matrices.

I.C.2 Reachability of Closed and Finite Quantum Systems

The definitions above, and especially Eq.(I.57) state that, if all Hamiltonians can be implemented with controls, then every propagator can be reached. This is an enormously strong requirement that is very rarely realised in any physical situation. In practice, a much smaller set of control Hamiltonians can be used to reach the space of all unitaries. A physically motivated illustration of what this set is, although not a formal proof, is outlined below.

It is instructive to begin by formally writing down the reachable set

$$\mathcal{R} = \left\{ \mathcal{T} e^{-i\int H(t)dt} |\ H(t) = H_0 + \sum_k f_k(t) H_k \ \ \forall\, f_k(t) \text{ and } t \geq 0 \right\}, \tag{I.58}$$

which is all the unitaries at any (positive) time that can can be generated by the Hamiltonian. A trivial observation is that the exponential of any Hamiltonian with constant controls $f_k(t) = f_k$ is in the reachable set. By increasing the evolution time, positive multiples of the Hamiltonian can be created. While the controls may allow for the $f_k(t)$ to be negative, it is not immediately obvious that the sign of H_0 can be reversed. This can in fact also be done by using the quantum recurrence theorem [BL57] which states that, because $\mathbb{U}(d)$ is a compact group,

$$\forall H, \epsilon \ \ \exists\, t_r > T \ : \quad ||e^{-iHt_r} - \mathbb{I}|| < \epsilon, \tag{I.59}$$

where $||.||$ is any matrix norm. This states that there always exists a time (above any minimum time) such that the system has evolved arbitrarily close to the identity. This permits the system to effectively evolve backwards in time by evolving for long enough forward in time

$$e^{-iH(-t)} \simeq e^{-iH(t_r - t)} \qquad \text{where} \quad t_r > t, \tag{I.60}$$

up to some small error ϵ. Thus, if e^{-iH} is in the reachable set, then so is e^{-iHt} for all t.

Another observation can be made by invoking the Trotter formula (a lemma of the

Baker-Campbell-Hausdorff formula)

$$e^{-i(A+B)t} = \lim_{n\to\infty} \left(e^{\frac{-it}{n}A} e^{\frac{-it}{n}B}\right)^n, \tag{I.61}$$

stating that the sum of two Hamiltonians can be exponentiated by alternating between the exponential of the two individual Hamiltonians. Combining this with the previous result, the exponential of any Hamiltonian which is in the linear space spanned by $H(t)$ is in the closure of the reachable set, where the closure is used to include elements that are ϵ close to being reachable. This space can be further expanded by using another lemma of the BCH formula, the commutator relation

$$e^{-[A,B]t^2} = \lim_{n\to\infty} \left(e^{\frac{-it}{n}A} e^{\frac{-it}{n}B} e^{\frac{it}{n}A} e^{\frac{it}{n}B}\right)^{n^2}, \tag{I.62}$$

which states that the exponential of the commutator of A and B can be reached by alternating between exponentiating those two directions forward and backwards in time. As it has already been shown that this can be done, the reachable set includes the exponential of the commutator of $H(t)$.

By combining the ability to scale Hamiltonians, take linear combinations, and commutators iteratively the whole Lie algebra

$$\mathcal{L} = \left\langle H(f_k) = H_0 + \sum_k f_k H_k \right\rangle_{\text{Lie}}, \tag{I.63}$$

(rather than just the vector space) spanned by the control Hamiltonians is in the closure of the reachable set. Furthermore, this is exactly the reachable set [D'A08, Ell09], such that

$$\overline{\mathcal{R}} = e^{\mathcal{L}}, \tag{I.64}$$

is the set of all unitary matrices that can be reached arbitrarily well by the control system of Eq.(I.50). In order to determine whether $\mathbb{U}(d)$ is reachable it is therefore sufficient to calculate the dimension of the Lie algebra $\mathcal{L}$; if this is d^2 or d^2-1 (the dimension of $\mathbb{U}(d)$ and $\mathbb{SU}(d)$ respectively) then the system is fully controllable. In the case that the dimension is lower only a subgroup of $\mathbb{U}(d)$ can be reached, as these are limited in number, the dimension of the Lie algebra is a good indicator of what the reachable set is. This dimension is easy to compute as it is composed of nothing but linear combinations and commutators.

The Lie algebra criterion is not, however, a practical constructive method for how the $f_k(t)$ need to be chosen in order to reach a desired unitary. The use of the quantum recurrence theorem implies waiting a significant amount of time to reach negative times, and both Eq.(I.61) and Eq.(I.62) require an unbounded number of repetitions. This would therefore require huge evolution time before getting close to a desired unitary and for the

controls to be very high frequency, neither of which are desirable. Different methods, typically numerical, have to be used to find the required pulse. Before giving an introduction to such a method, a brief explanation of why this criteria fails for open or infinite systems is presented.

The approach taken here fails for open systems at the first hurdle. The space of CPT maps, even when limited to just the Markovian ones which admit a generator, form a semigroup but not a group. The physical reason for this is that CPT maps represent a loss of coherence, there is information flowing out of the system. Any concatenation of CPT maps cannot only lead to information being recovered. Without this group structure, the whole construction presented here falls apart. It can be partially recovered by defining Lie semigroups and related objects, this gives rise to a limited set of results for open systems [DHKSH09, ODSH12]; these are discussed in the context of a different method to determine controllability introduced in §II.

One way around such difficulties would be to treat the open system as a larger unitary one. The environments which are of physical interest are almost always infinitely dimensional, thus the problem of open control can be mapped to infinite dimensional control, although the question of which dilation should be used is subtle. The key step in the construction of the reachable set presented above that fails for such systems is in evolving for negative time using the recurrence theorem of Eq.(I.59). In mathematical terms, this is because unitary groups in infinite dimensions are not compact. Geometrically this can be seen as the Lie group in finite dimensions being curved (like the Bloch sphere in $d = 2$) such that evolution in any direction eventually leads back to the identity, while in infinite dimension the evolution goes on forever without curving back. A simple physical explanation of this is to consider a particle in one dimension evolving according to the momentum operator, $U(t) = e^{-ipt}$, there is no value of $t > 0$ for which this approaches the identity. Nevertheless, in the case of quadratic Hamiltonians, the rank criteria described here can be used with only small modifications [GSKB12].

I.C.3 Methods in Optimal Control

As discussed above, the Lie algebraic condition is an efficient way of determining which unitaries can be reached by a quantum control system. It is, however, not constructive and provides little insight as to what $f_k(t)$ in Eq.(I.50) needs to be in order to reach a target unitary U_T. This problem is the topic of optimal control, where iterative numerical methods are used to find a control pulse that achieves the desired goal, and the description here is based around well established methods [MSGdFG11, GBCKK15]. As this approach does not use the analytic properties of Lie groups, it can also be used for open or infinite dimensional systems, and with minimal modifications also applies to state control (where many of the algorithms were actually first developed).

The steps for finding the controls for a unitary are outlined in Fig.I.3. The starting point is an initial set of values for control parameters, including the total evolution time and value of $f_k(t)$. In order to have a finite number of control parameters, those are usually taken as piecewise-constant [KRKSHG05], or are expressed in a Fourier basis with a finite number of components [DCM11, CCM11, MTWA15] . If there is no a priori knowledge as to what the control pulse ought to be, the values of the parameters are chosen randomly. The evolution of the system with these parameters is then calculated in order to obtain $U_{f_k(t)}$, the unitary that is generated by the controls. This requires obtaining an accurate model for the system (the values of H_0 and H_k, as well as the generator of the noise if the system is open) and then solving the time-dependent Schrödinger equation numerically for a model of the system. If the system is fully controllable, such a simulation will be exponentially difficult as it scales with the dimension of the entire Hilbert space. In small systems, this can however be done without undue difficulty. The algorithm then needs a cost functional

$$F[f_k(t), U_T], \tag{I.65}$$

which determines how far the current controls are from realising the desired operation. The simplest choice for this is the gate fidelity

$$F_U = \mathrm{Tr}[U_{f_k(t)} U_T^\dagger] \qquad \text{or} \quad F_{SU} = |\mathrm{Tr}[U_{f_k(t)} U_T^\dagger]|, \tag{I.66}$$

depending on whether the global phase of the evolution matters or not respectively; this common choice is why the cost functional is often called the fidelity. Additional terms to the cost functional can be added in order to encourage the pulse to have certain properties, for example adding a cost proportional to the L_1 norm results in sparser control pulses or a maximum value can be imposed; this is used to include physical constraints on the desired pulses. If the value of $F[f_k(t), U_T]$ is above a certain threshold, the target fidelity, then the scheme terminates and the values of the control parameters which achieved the desired operation are recorded. If this is not the case, the value of the control pulses (either all of them simultaneously or a subset at a time) are updated, and the process loops around. If a random initial pulse was taken, the whole process would generally be repeated with different initial pulses, as the final fidelity reached (and the number of iterations required to find it) depends strongly on the initial values.

How to update the control pulse is a completely classical problem and countless techniques that have been developed in computer science and statistics can be used here. One class of methods are called simplex, such as Nelder-Mead [NM65], which starts by evaluating the functional at n points in the parameter space that form a simplex and picks the next point to evaluate as a linear combination of the previous ones, weighed by their fidelity. The protocol then repeats itself, taking the best n points in each iteration when guessing the next one. A different approach is to use a genetic algorithm, where different values of the control parameters are randomly mutated and the most successful are

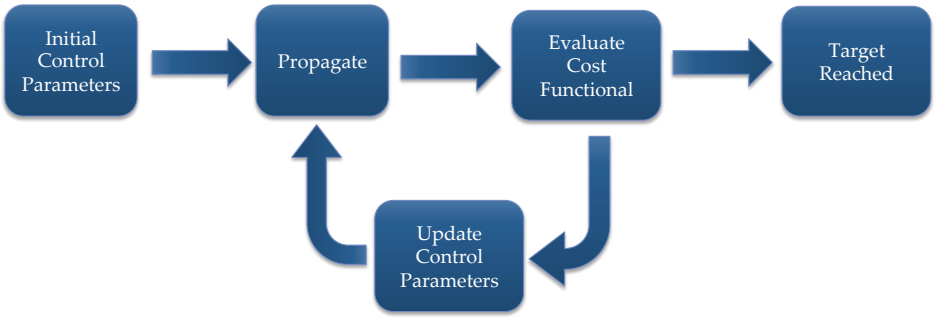


Figure I.3 | Outline of the process used in optimal control. An initial choice for the control parameters, $f_k(t)$, are propagated in time to find the resultant unitary, $U_{f_k(t)}$. The distance between this and the target U_T are evaluated using a cost functional F. An optimiser uses this distance, and previously calculated distances, to update the control pulse. The whole processes is repeated until the distance to the target is below a threshold, at which point the algorithm outputs the final control parameters.

'bred' together [Gol89]. Machine learning methods can also be used which use artificial neural networks and reinforcement learning to predict what the best value of the control parameters are [PWZVS16]. Possibly the most popular methods are gradient based, such as the well known Newton-Raphson. Rather than using the fidelity at different points in parameter space to decide how to modify the controls, these use the fidelity at a single point plus the derivative of the fidelity

$$\frac{\partial}{\partial f_k(t)} F[f_k(t), U_T], \tag{I.67}$$

to determine how to modify the pulse. Although gradient based methods generally converge faster than others, calculating the gradients can be computationally expensive, and they can get stuck in local extrema of the cost functional. The presence of such extreme are hotly debated and the subject of much research [HDR09, PT11, RRW16], although they are generally believed to not be a problem in finite dimensional closed systems when the standard gate fidelities of Eq.(I.66) are used.

The method used repeatedly in this thesis is GRadient Ascent Pulse Engineering (GRAPE) [KRKSHG05, dFSGK11], which was originally designed for state transfer protocols. This uses piecewise-constant controls and a gradient based method to update the entire control pulse simultaneously. As well as the first derivative of the fidelity, it also uses the second derivatives in order to update the pulse; this guarantees convergence in a single step for a quadratic fidelity landscape and, even when this is not the case, greatly improves the convergence rate. Calculating so many derivatives would however be very costly, so the algorithm estimates the second derivatives based on the history of the first derivatives. These are themselves calculated exactly (rather than by finite difference methods) efficiently using [Fun04]. The particular implementation of GRAPE used by us is in Quantum Toolbox in Python (QuTiP) [JNN13, JNPGG16], a library for python

which allows easy and fast quantum computation. The control packages for QuTiP were written by Alexander Pitchford at the same time as this PhD was started, and we have contributed to developing the code as our needs evolved.

CHAPTER II

ISOTROPY AND CONTROL OF OPEN QUANTUM SYSTEMS

II.A Introduction

This chapter is a direct, if only partial, answer to the question of controllability of open quantum systems; the main results stem by considering how controls can average out the decay of the state space induced by the noise. This provides necessary but not sufficient conditions for what CPT maps (or generators) an open quantum system with a fixed dissipative generator and Hamiltonian controls can reach. The work presented here follows closely a paper authored by us during the PhD [DBM16].

II.A.1 Overview

The ability to control open quantum systems is of importance both for the insight it provides to fundamental science and for its potential applications in technology, as was covered in §I. Before attempting to find a specific control sequence that makes a system do a desired task, it is useful to know if this task can be achieved to begin with. The question of controllability is often easier to determine than finding the controls, as was discussed in §I.C.2 for the case of closed quantum systems, therefore it is natural to answer this question first for open systems.

This chapter covers our work on investigating this controllability question for open quantum systems described by a time-local Master equations. Some bounds are constructed which determine whether a given quantum channel can be reached by the system by using any time-dependent Hamiltonian, while keeping the noise fixed. These are a series of conditions which exclude channels from being reachable by any unitary controls. These primarily arise by defining decay rates for the generator of the dynamics of the open system, and then showing that controlling the system can only make these rates more isotropic. This forms a series of constraints on the shape and non-unitality of al-

lowed evolutions, as well as an expression for the time required to reach a given target that are easy to compute. Numerical examples of the usefulness of these criteria are given, and a brief exploration of some similarities the results have with thermodynamics.

In order to get an intuitive understanding for the principle idea behind this chapter, it is useful to consider dissipative quantum processes as a flow in state space, as illustrated in Fig.II.1. The noise can act in a variety of ways on the state space, including rotating and shrinking the space (corresponding to decay) in a potentially anisotropic way. Hamiltonian controls allow additional rotations to be imposed on the system such that different parts of the state space feel different contraction rates at different times. This results in the ability to mix the decay rates together and leads to the overall evolution obeying some averaged rates. Under such controls the final state space, which represents the total evolution, is more isotropic than in the absence of controls.

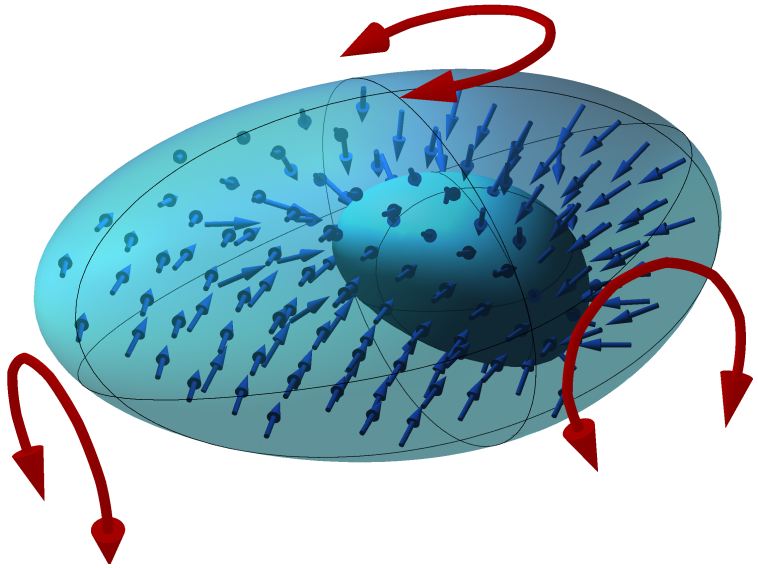


Figure II.1 | Illustration of the anisotropic flow of states in an open system. The principle ideas of this chapter are exemplified by a cross section of a high dimensional state space at two different times, with the small arrows indicating the direction of flow induced by a Lindbladian, as in Fig.I.1. The red arrows correspond to the available controls. The different parts of the space are contracting at different rates; by rotating the system in time with Hamiltonian controls some of these decay rates can be averaged together.

Our main results relate directly to this, and state that a quantum operation cannot be reached if it has a more ordered structure than the noise acting on the system. This chapter begins with introducing the specific problem being addressed in §II.A.2, followed by a brief comparison of previous methods in §II.A.3 and an overview of the relevant mathematics in §II.A.4. A definition of the decay rates of an open system §II.B is presented along with a series of necessary but not sufficient criteria for reachability showing their use. This is followed by a proof of these conditions in §II.C and a detailed discussion of their meaning and utility in §II.D. The strength of these criteria is tested numerically in §II.E

for common examples of noise showing that, for qubit and qutrit systems at least, the necessary conditions are strong and tight enough to provide a major restriction on what evolutions are possible in realistic situations. Both the language and the mathematics used to describe these relations are reminiscent of thermodynamics, a link which is explored in §II.F. Finally §II.G concludes with a summary of the results, their implications, and interesting open questions.

II.A.2 Problem Definition

Before the intuition described in Fig.II.1 can be used to form specific control theorems, and extended to more general settings, it is necessary to state exactly the problem at hand. The systems considered are finite dimensional with Hamiltonian controls interacting with an environment, under the requirement that the reduced system obeys a time-local master equation

$$\frac{d}{dt}\rho = G_t(\rho) = G_t^0(\rho) - i[H_t, \rho] \tag{II.1}$$

$$\rho_T = M_T(\rho_0) \equiv \mathcal{T} e^{\int_0^T dt G_t(\cdot)} \rho_0, \tag{II.2}$$

where the terms are as defined previously in §I.B with G_t being the generator for the system, decomposed into its drift G_t^0 and a controllable Hamiltonian H_t, and M_T is the resultant dynamical map after an evolution time T. If G_t^0 is restricted to be a Lindblad operator, denoted by L_t, then the allowed solutions are Markovian, completely positive trace preserving maps, as discussed in §I.B.3. Although the Lindbladian case is the most commonly used and the one with the clearest physical interpretation, the key results of this chapter do not rely on the specific form of the Lindblad operator and hold for a more general generator which gives rise to non-Markovian dynamics. In such cases there can be substantial additional restrictions beyond the ones presented here, since the dynamics induced by control Hamiltonian and non-Markovian drift does not necessarily induce completely positive dynamics, even if the uncontrolled evolution is completely positive [Lin76]. The validity of the conditions presented in this chapter are, however, not impaired by the additional intricacies of non-Markovian dynamics provided that the dynamics are still describable by a time-local linear generator. The effect on the control criteria of possible limitations on H_t due to experimental constraints are not considered; any reduction on the allowed controls can only reduce the set of reachable operations but, as the results here only exclude maps, the validity of the conditions detailed here is not affected.

An implicit assumption in Eq.(II.1) is that the control Hamiltonian does not affect the dissipative component of the generator. While there are cases where this approximation holds exactly [DJR14] (a discussion of our research into another case where this may hold is in §IV), this is not always the case, and it is possible for controls to modify

the decoherence induced by the environment [AHFB15], a point touched on in §III. In such circumstances a variety of different methods have been developed [FB07, RSSHW09, CBK10, CBK12, AGB14], but these typically require detailed knowledge about the environment or additional assumptions about finite dimensionality or bounded interactions. They may also call for experimentally difficult regimes such as strong or rapidly oscillating control fields. For these reasons, it is highly desirable to explore what can be achieved with controls if the dissipative component of the dynamics is not modified, which is the regime studied here.

Previous work on the controllability of open systems typically use the same definition of the problem and make similar assumptions, as discussed below.

II.A.3 Prior work

While writing up this research for publication we came across work by Yuan [Yua11] which explored similar ideas but with a focus on the special case of time-independent, Markovian, and unital quantum systems. In this restricted case, the results presented in this chapter overlap directly with that paper. The substantial difference lies in more general and higher dimensional systems; Yuan said little about these other than showing that a direct extension does not hold. This chapter shows that they can be generalised, but at the cost of going from a complete characterisation of the reachable set to necessary but not sufficient conditions for being in it. In addition, conditions that take into consideration (and are a constraint on) the non-unitality of a system are introduced; an aspect Yuan did not investigate.

A completely different approach on approximating the reachable set of Markovian systems was carried out in [DHKSH09, ODSH12] using Lie algebraic ideas, as introduced in §I.C.1. This is an attempt to directly generalise the controllability conditions of unitary quantum systems, as discussed in §I.C.2, to Markovian systems. Similar assumptions about the underlying system as ours were made, with the additional constraint that the Lindbladian must be unital and time-independent. The characterisation of the reachable set produced by that method has the benefit of being able to limit the available controls from all Hamiltonians to a reduced set.

The approach consists of finding the Lie wedge, which is the set of all possible directions (in the manifold of the dynamics of the system) that can be explored with controls. These are loosely given by all commutators of the control Hamiltonians added to the Lindbladian, in direct analogy with the controllability results of closed quantum systems. For an element $\mathfrak{w}$ in the wedge all maps $e^{\mathfrak{w}t}$ can be reached; in closed systems this is also the set of all reachable maps. However this is not the case for open systems. The underlying mathematical reason is that for any two Hermitian operators H_1, H_2 it is always true that $\exp(-iH_1)\exp(-iH_2) = \exp(-iH_3)$ where H_3 is also Hermitian. This is not

the case for two Lindblad operators, $\exp(L_1)\exp(L_2) \neq \exp(L_3)$ with the requirement that L_3 is also a Lindbladian, unless L_1, L_2 satisfy additional algebraic constraints linked to Baker-Campbell-Hausdorff multiplication. The physical meaning of this is that there exist directions that cannot be explored by the control system at $t = 0$, but can be at later times. This means that (unless the Lie wedge is global), the reachable set is not given by $e^{\mathfrak{w}t}$, but by the semigroup closure of the product of such terms.

The end result of this approach is a characterisation of the reachable set of Markovian systems that cannot be computed exactly, except in special cases such as for qubits. Instead an inner approximation can be calculated using an iterative strategy. Such a method, as the name implies, will only identify a subset of the quantum maps that can be reached. It therefore provides a sufficient but not necessary condition on a target map being reachable by the control system. This is the opposite to the work outlined in this chapter which, by providing necessary but not sufficient criteria, is an outer approximation to the reachable set. The two methods are thus complementary as they tackle the problem from opposite sides.

A different use of the method of [DHKSH09, ODSH12] is to characterise how large the space that can be explored by a Markovian system is by dimension counting. Different Lie algebras can be calculated, and the size of the different algebras contains some information about the range of possible dynamics. The first two measures are the size of the Lie algebra of the controls, and the size of the Lie algebra of the controls and the Hamiltonian drift. This has the same interpretation as in closed systems and characterise the size of the space that the system could explore in the absence of any dissipation. The last measure quantifies accessibility and is the Lie algebra of the controls and the total drift (Hamiltonian plus Lindbladian) of the system. The closer these dimensions are to the size of the full dynamical system, given by $d^4 - d^2$, the broader the range of dynamics possible. However it is not possible to use these measures to assess whether a given map is reachable or not.

II.A.4 Mathematical Background

In order to make the picture introduced in Fig.II.1 rigorous, and to obtain the desired results, it is instructive to highlight some of the key properties of the super-operators of Eq.(II.1). This is best done by representing them as matrices using the vectorisation method discussed in §I.B.5. An orthonormal Hermitian basis set is used here where the basis matrices are traceless (such as the Pauli matrices), except for the first which is proportional to $\mathbb{I}$. For a trace preserving dynamical map the vectorised expression for M

using such a basis is given by

$$\underline{M} = \begin{pmatrix} 1 & 0 & 0 & ... \\ v_1 & \overline{M}_{11} & \overline{M}_{12} & ... \\ v_2 & \overline{M}_{21} & \overline{M}_{22} & ... \\ ... & ... & ... & ... \end{pmatrix}, \tag{II.3}$$

where the top row is fixed, all the elements are real if M is Hermiticity preserving (which it is for the vast majority of physically sensible cases), and the bar refers to the reduced matrix. This form has the advantage of explicitly separating the unital and non-unital part of the dynamics. Unitality is an important property in this context as Hamiltonian controls are necessarily unital, but the noise may not be. The left hand column consisting of the elements v_i fully describes the non-unital part of the map and quantifies how much the maximally mixed state is translated by the dynamical map. For unital operations, such as unitary evolution, these vanish and the dynamical map reduces to $\underline{M} = \underline{1}_1 \oplus \overline{M}$. The reduced matrix $\overline{M}$ thus describes solely the unital part of the evolution. These are partially decoupled from the non-unital dynamics in the sense that $\overline{(\underline{M}_B \underline{M}_A)} = \overline{M}_B \overline{M}_A$, meaning that the total unital dynamics of a concatenation is given by the concatenation of the unital part of the individual super-operators. This can easily be seen by noting that the concatenation of super-operators is given by the product of their matrix representations.

As unitality is a key property, it is also useful to have a measure of how unital or non-unital a dynamical map is. A convenient one is $\mathrm{Tr}\left[M(\frac{1}{d}\mathbb{I})^2\right]$, the purity of the state obtained by applying the map to the maximally mixed state. This value is maximised at one if the maximally mixed state is mapped to a pure state and is minimised to $\frac{1}{d}$ if the map is unital. In a similar fashion the non-unitality can also be quantified by higher moments [SGB14], $\mathrm{Tr}\left[M(\frac{1}{d}\mathbb{I})^n\right]$, and the first d moments are linearly independent. These should all be seen as describing roughly the same physical quantity: the ability of the map to purify states.

The matrix form of a trace preserving generator G is identical to Eq.(II.3), except that the entire top row vanishes. Its unital and non-unital parts can be separated in a similar way and the same results on concatenation holds. Because of this, $\overline{(e^{\underline{G}})} = e^{(\overline{G})}$ and, hence, if M is the dynamical map generated by G_t then $\overline{M}$ is the one given by $\overline{G}_t$. This means that contained inside every non-unital problem is a unital one with a dimension one smaller, and any solution to the control problem of generating M must also solve $\overline{M}$. It is possible for this reduced problem to be non-Markovian even if the original one is Markovian, but this does not affect the validity of the approach.

The intuitive picture described in the introduction relies on a notion of averaging a set (the decay rates) to obtain another. A natural way to describe this process is the majorization relation [Bha97] which tests if one real set is more uniformly distributed than an other. A set $\boldsymbol{a}$ of real numbers is majorized by another such set $\boldsymbol{b}$, written as

$\boldsymbol{a} \prec \boldsymbol{b}$, if and only if

$$a_1^{\downarrow} \leq b_1^{\downarrow} \tag{II.4a}$$
$$a_1^{\downarrow} + a_2^{\downarrow} \leq b_1^{\downarrow} + b_2^{\downarrow} \tag{II.4b}$$
$$\cdots$$
$$\sum a_i^{\downarrow} = \sum b_i^{\downarrow}, \tag{II.4c}$$

where $\downarrow$ signifies that the elements of the set are sorted in decreasing order. Another way of stating this is that $\boldsymbol{a}$ is majorized by $\boldsymbol{b}$ if and only if an ordering of $\boldsymbol{a}$ can be obtained by a convex sum of different orderings of $\boldsymbol{b}$. It is this property which makes it suitable to describe an averaging procedure. Another useful property is that it is conserved under scaling: $\boldsymbol{a} \prec \boldsymbol{b}$ also implies $x\,\boldsymbol{a} \prec x\,\boldsymbol{b}$ for all real (including negative) x. A point to note is that, unlike standard inequalities on the reals, majorization provides only a partial order.

II.B Theorems

The formalism developed above allows the intuition developed in Fig.II.1 to be refined and given a specific mathematical meaning. The picture was of a Lindbladian (or a more general drift) acting on the state space such that it flowed from one shape to another. By coherently controlling the system, the space can be rotated so that some of the decay rates are averaged together. For a Markovian two-level system where the Lindbladian has no Hamiltonian component, these decay rates are the eigenvalues of the Lindbladian, which are always non-positive. In higher dimensions however these eigenvalues can be complex, which gives rise to two problems: they may not faithfully quantify the contraction of the space, and it is not clear what averaging them would signify. The situation is exacerbated in both the non-unital and non-Markovian case where there are even fewer constraints on the spectrum of the generator. This points to the need for a different way of quantifying the decay rates of an open system than naively taking the spectrum of the drift.

In §II.C it is shown that the correct decay rates to consider are the eigenvalues of the sum of the drift and its dual, $\underline{G} + \underline{G}^{\dagger}$, corresponding to the Hermitian part of the drift. Their physical relevance is supported by two important properties. Firstly, as the components of $\underline{G}$ that generate rotations are anti-Hermitian (whether they correspond to a Hamiltonian degree of freedom or not), these rates capture only the decay or growth component of the flow. Secondly they are real, so averaging them corresponds to a more uniform flow in a way which can be naturally defined using majorization. These suggest that the eigenvalues of the Hermitian part of the drift capture some of the key aspects of the controllability of the system, and their prominence in the criteria detailed below show that this is indeed the case. In a similar fashion the anisotropy of the dynamical map

is described by its singular values (as they are rotationally invariant and non-negative), which loosely correspond to the characteristic lengths of the final state space.

This results in precise mathematical formulations of conditions which must all be satisfied for a dynamical map M to be reachable by a system with drift G_t^0. Using the notation defined previously these are: an expression for the state space volume reached at a given time (independent of the controls)

$$\det(M) = e^{\int_0^T \mathrm{Tr}[G_t^0]dt}, \tag{II.5}$$

which extends prior work on the determinant of quantum channels [WC08]; a constraint on the anisotropy of the dynamical map

$$\log[\boldsymbol{\sigma}(M)] \prec \int_0^T \boldsymbol{\lambda}\left(\frac{G_t^0 + G_t^{0\dagger}}{2}\right) dt, \tag{II.6}$$

where $\boldsymbol{\lambda}$ and $\boldsymbol{\sigma}$ refer to the set of eigenvalues and singular values of the super-operator respectively, and the log acts element-wise on the set; a unital version of this condition (for trace preserving drifts)

$$\log[\boldsymbol{\sigma}(\overline{M})] \prec \int_0^T \boldsymbol{\lambda}\left(\frac{\overline{G}_t^0 + \overline{G}_t^{0\dagger}}{2}\right) dt, \tag{II.7}$$

which was shown in [Yua11] for a time-independent Lindbladian; and bounds on the maximal non-unitality that can be reached

$$\mathrm{Tr}[M(\tfrac{1}{d}\mathbb{I})^n] \leq \sup_{\rho}\left\{\mathrm{Tr}[\rho^n] \mid \exists\, t : \mathrm{Tr}[\rho^{n-1}G_t^0(\rho)] = 0\right\}, \tag{II.8}$$

which gives rise to independent conditions for $n = 2, ..., d$. In order to highlight the systematic similarities of the conditions, and stress that it is the Hermitian part of the drift that matters (which in the Bloch representation removes any Hamiltonian or other rotational component), G_t^0 can be replaced by $\frac{1}{2}(G_t^0 + G_t^{0\dagger})$ in Eq.(II.5) and Eq.(II.8) without changing the results.

Instead of seeing if the system can reach a target map, the question can be modified to ask if a time-independent generator G^0 and controls can approximately simulate another generator G' arbitrarily well. In this case the constraint is

$$\boldsymbol{\lambda}\left(G' + G'^{\dagger}\right) \prec \boldsymbol{\lambda}\left(G^0 + G^{0\dagger}\right), \tag{II.9}$$

whose unital version

$$\boldsymbol{\lambda}\left(\overline{G}' + \overline{G}'^{\dagger}\right) \prec \boldsymbol{\lambda}\left(\overline{G}^0 + \overline{G}^{0\dagger}\right), \tag{II.10}$$

also holds provided both generators are trace preserving.

A proof of these conditions is presented here, followed by a detailed discussion.

II.C Proofs

II.C.1 Evolution time

The starting point for the derivation of Eq.(II.5) is to note that the formal solution for M in Eq.(II.1) is given in terms of time-ordered matrix exponential, which can be expressed according to the Magnus expansion [BCOR09]

$$\underline{M} = \mathcal{T} e^{\int_0^T \underline{G}(t)dt} \tag{II.11}$$

$$= e^{\int_0^T \underline{G}(t_1)dt_1 + \frac{1}{2}\int_0^t dt_1 \int_0^{t_1} dt_2 [\underline{G}(t_1),\underline{G}(t_2)] + \ldots}, \tag{II.12}$$

where all higher order terms in the series consist of nested commutators. As the determinant of a matrix exponential is the exponential of the trace this can be rewritten as

$$\det(\underline{M}) = e^{\mathrm{Tr}\left[\int_0^T \underline{G}(t_1)dt_1 + \frac{1}{2}\int_0^t dt_1 \int_0^{t_1} dt_2 [\underline{G}(t_1),\underline{G}(t_2)] + \ldots\right]} \tag{II.13}$$

$$= e^{\mathrm{Tr}\left[\int_0^T \underline{G}(t)dt\right]}, \tag{II.14}$$

where the tracelessness of the commutators has been used. As discussed previously, the control Hamiltonians appear in the equation of motion Eq.(II.1) as commutators, therefore their Bloch representation are also traceless giving $\mathrm{Tr}[\underline{G}(t)] = \mathrm{Tr}[\underline{G}^0(t)]$. As the trace and determinant of M and G are identical to those of $\underline{M}$ and $\underline{G}$, this gives the desired expression

$$\det(M) = e^{\int_0^T \mathrm{Tr}[G_t^0]dt}. \tag{II.15}$$

If the Magnus expansion does not converge (which may happen if $\int_0^T ||\underline{G}(t)||_2 dt > \pi$ [BCOR09]), then the proof can be extended by splitting the propagator into sufficiently many terms

$$\underline{M} = \mathcal{T} e^{\int_{t_n}^T \underline{G}(t)dt} \ldots \mathcal{T} e^{\int_{t_1}^{t_2} \underline{G}(t)dt} \mathcal{T} e^{\int_0^{t_1} \underline{G}(t)dt}, \tag{II.16}$$

such that the Magnus expansion converges for each term. Applying the same steps as before to each term and using the fact that the determinant of a product is the product of the determinants results in

$$\det(M) = \det(\mathcal{T} e^{\int_{t_n}^T \underline{G}(t)dt}) \ldots \det(\mathcal{T} e^{\int_0^{t_1} \underline{G}(t)}) \tag{II.17}$$

$$= e^{\int_{t_n}^T \mathrm{Tr}[G_t^0]dt} \ldots e^{\int_{t_1}^{t_2} \mathrm{Tr}[G_t^0]dt} e^{\int_0^{t_1} \mathrm{Tr}[G_t^0]dt} \tag{II.18}$$

$$= e^{\left\{\int_{t_n}^T \mathrm{Tr}[G_t^0]dt + \ldots + \int_{t_1}^{t_2} \mathrm{Tr}[G_t^0] + \int_{t_0}^{t_1} \mathrm{Tr}[G_t^0]dt\right\}} \tag{II.19}$$

$$= e^{\int_0^T \mathrm{Tr}[G_t^0]dt}, \tag{II.20}$$

as before.

II.C.2 Anisotropy of the dynamical map

The proofs of Eqs.(II.6-II.7) arise from two observations. Firstly, the evolution can always be decomposed into infinitesimal time steps in a Trotter-like way, alternating between coherent and incoherent evolution. Secondly, the controls only affect the coherent steps which are all rotation matrices and so do not modify the singular values of the incoherent time steps, because singular values of a matrix depend only on the product of that matrix with its Hermitian adjoint. To prove Eq.(II.6) the time-ordered exponential is expanded in terms of short time steps

$$\underline{M} = \mathcal{T} e^{\int_0^T \underline{G}(t)dt}, \tag{II.21}$$

$$= \lim_{\delta t \to 0} \left(e^{\underline{G}^0(T)\delta t} e^{\underline{H}(T)\delta t} ... e^{\underline{G}^0(0)\delta t} e^{\underline{H}(0)\delta t} \right). \tag{II.22}$$

The next step is to consider the singular values of both sides of the equation, denoted by the operator $\boldsymbol{\sigma}$. Specifically, the expression [Bha97]

$$\log \boldsymbol{\sigma}(AB) \prec \log \boldsymbol{\sigma}(A) + \log \boldsymbol{\sigma}(B), \tag{II.23}$$

for the majorization relation between the singular values of matrices and their products. Here the log is understood as acting on each element in the set, and the sum on the right hand side acts on the elements of the sets ordered by magnitude. Generalising this to the case of multiple sums and applying it to Eq.(II.22) gives

$$\log \boldsymbol{\sigma}(\underline{M}) \prec \lim_{\delta t \to 0} \left\{ \log \boldsymbol{\sigma} \left(e^{\underline{G}^0(T)\delta t} e^{\underline{H}(T)\delta t} \right) + ... + \log \boldsymbol{\sigma} \left(e^{\underline{G}^0(0)\delta t} e^{\underline{H}(0)\delta t} \right) \right\}. \tag{II.24}$$

As mentioned above, the coherent steps corresponds to rotation matrices and therefore do not affect the singular values. This allows the expression for the singular values to be simplified to

$$\log \boldsymbol{\sigma}(\underline{M}) \prec \lim_{\delta t \to 0} \left\{ \log[\boldsymbol{\sigma}(e^{\underline{G}^0(T)\delta t})] + ... + \log[\boldsymbol{\sigma}(e^{\underline{G}^0(0)\delta t})] \right\}. \tag{II.25}$$

Recalling that singular values are obtained by $\boldsymbol{\sigma}(A) = \boldsymbol{\lambda}\left(\sqrt{AA^\dagger}\right)$, where $\boldsymbol{\lambda}$ signifies the eigenvalues, each term in the previous equation can be expressed for small δt as

$$\log \boldsymbol{\sigma}(e^{\underline{G}^0(t)\delta t}) = \log \left[\boldsymbol{\lambda} \left(e^{\underline{G}^0(t)\delta t} e^{\underline{G}^{0\dagger}(t)\delta t} \right)^{\frac{1}{2}} \right] \tag{II.26}$$

$$\approx \log \left[\boldsymbol{\lambda} \left(e^{(\underline{G}^0(t) + \underline{G}^{0\dagger}(t))\delta t + [\underline{G}^0(t), \underline{G}^{0\dagger}(t)]\delta t^2} \right)^{\frac{1}{2}} \right] \tag{II.27}$$

$$= \tfrac{1}{2} \boldsymbol{\lambda} \left((\underline{G}^0(t) + \underline{G}^{0\dagger}(t))\delta t + [\underline{G}^0(t), \underline{G}^{0\dagger}(t)]\delta t^2 \right), \tag{II.28}$$

where higher order terms can be calculated using the Baker-Campbell-Hausdorff formula. The first term is of order δt and, as the number of terms in Eq.(II.25) is $\frac{T}{\delta t}$, it contributes to the integral in the limit $\delta t \to 0$ while all the higher order terms vanish. This gives as the final expression

$$\log\left[\boldsymbol{\sigma}(\underline{M})\right] \prec \lim_{\delta t\to 0}\left\{\boldsymbol{\lambda}\left(\frac{\underline{G}^0(T)+\underline{G}^{0\dagger}(T)}{2}\right)\delta t + ... + \boldsymbol{\lambda}\left(\frac{\underline{G}^0(0)+\underline{G}^{0\dagger}(0)}{2}\right)\delta t + \frac{T}{\delta t}O(\delta t^2)\right\} \tag{II.29}$$

$$\prec \int_0^T \boldsymbol{\lambda}\left(\frac{\underline{G}^0(t)+\underline{G}^{0\dagger}(t)}{2}\right)dt, \tag{II.30}$$

which is independent of the representation used and so holds for the super-operators themselves. The proof for condition Eq.(II.7), the unital version of this for trace preserving generators, follows immediately from the fact that $\underline{M} = \mathcal{T}\exp\{\int_0^T \underline{G}(t)dt\}$ implies $\overline{\underline{M}} = \mathcal{T}\exp\{\int_0^T \overline{\underline{G}}(t)dt\}$, as was discussed in §II.A.2. A different proof of this latter unital result was shown in [Yua11] for time-independent Lindbladians only, and relied on similar mathematical ideas.

II.C.3 Maximal non-unitality

The proof for the non-unitality bounds, Eq.(II.8), begins with the formal expression for the non-unitality and ends with an easily evaluable bound for it. To do this the additional assumption that G_t is continuous is required. The maximal non-unitality of an open system is quantified by

$$\mathrm{Tr}[M(\tfrac{1}{d}\mathbb{I})^n] \leq \sup_{t,H_\tau} \mathrm{Tr}[\rho^n(t,H_\tau)] \tag{II.31}$$
$$\text{where}\quad \rho(t,H_\tau) = \mathcal{T}e^{\int_0^t (G_\tau^0+H_\tau)(\cdot)d\tau}\tfrac{1}{d}\mathbb{I},$$

and the supremum is over all possible evolution times and all possible controls. It is sufficient to consider only H_τ which are defined for $\tau \in [0,\infty)$ as Hamiltonians defined only up to finite τ can be extended without affecting the dynamics up to τ. Eq.(II.31) appears as difficult to calculate as solving the control problem, and is therefore of limited use. However an upper bound for it that can be calculated more readily is derived below. This requires noting a property that the supremum must satisfy as a function of t for any H_τ, and then reformulating the constraint that ρ has been evolved from the maximally mixed state into one which is easier to work with.

For a given continuous H_τ finding the supremum of Eq.(II.31) reduces to finding the supremum of a scalar function which is bounded between $\frac{1}{d}$ and 1 and differentiable everywhere. There are several cases in which this could happen. Firstly, the supremum being reached at $t = 0$ can be immediately excluded as the function is a minimum at that point.

Secondly, the supremum being reached for some finite time leads to $\frac{d}{dt}\mathrm{Tr}[\rho^n(t)] = 0$ at that point in time. Lastly, if the supremum is not reached for finite t then either: it is reached in the limit $t \to \infty$ and so the derivative also goes to 0 in this limit (due to the function being bounded from above), or the limit is undefined because the function does not converge. In the latter case there are many local maxima which form a series, the supremum of which gives the supremum of the original function. As the gradient of each these maxima is 0, the largest value that can be reached by the function also occurs when the gradient vanishes. Hence, a necessary condition for the supremum of Eq.(II.31) is

$$\tfrac{d}{dt}\mathrm{Tr}[\rho^n] = 0, \tag{II.32}$$

which can be transformed using the equation of motion of the system and elementary properties of the trace to

$$\mathrm{Tr}[\rho^{n-1}\, G_t^0(\rho)] - i\mathrm{Tr}\left[\rho^{n-1}\,[H_t, \rho]\right] = 0 \tag{II.33}$$

$$\mathrm{Tr}[\rho^{n-1}\, G_t^0(\rho)] - i\mathrm{Tr}\left[\,[H_t, \rho^n]\,\right] = 0 \tag{II.34}$$

$$\mathrm{Tr}[\rho^{n-1}\, G_t^0(\rho)] = 0. \tag{II.35}$$

Since the gradient as calculated above depends solely on G_t^0 and not on the controls, the condition on G_t being continuous can be relaxed to G_t^0 being continuous. Instead of calculating the supremum over all controls, the supremum can be calculated over all states instead. As this is necessarily a larger set, it provides a valid upper bound to the largest possible non-unitality of a drift G_t^0 and is given by

$$\mathrm{Tr}[M(\tfrac{1}{d}\mathbb{I})^n] \leq \sup_{\rho}\left\{\ \mathrm{Tr}[\rho^n]\ \mid\ \exists\, t : \mathrm{Tr}[\rho^{n-1}G_t^0(\rho)] = 0\right\}, \tag{II.36}$$

which, as desired, does not require any propagators to be calculated. This provides up to d constraints (including the trivial case for $n = 1$), as higher moments of ρ are not independent.

II.C.4 Generator anisotropy

The criterion for generator anisotropy (applicable only to time-independent generators) is similar to the one for the anisotropy of the dynamical map, with the important difference that the target is a flow in state space that is to be achieved continuously in time, rather than a snapshot of the evolution at a single instance, as shown schematically in Fig.II.2. The derivation for this condition begins with Eq.(II.6), where the target M has been replaced by $e^{G'T}$, and limited to the system drift being time-independent, such that the

starting point is

$$\log\left[\boldsymbol{\sigma}(e^{G'T})\right] \prec \int_0^T \boldsymbol{\lambda}\left(\frac{G^0+G^{0\dagger}}{2}\right) dt \;\; \forall\, T. \tag{II.37}$$

If this condition is satisfied for infinitesimal δt then, by concatenation, it holds for all time T and G^0 can effectively simulate G'. The term 'effectively' is used to emphasise that although the generator G' cannot be reached exactly, it is possible to follow a trajectory in state space which is arbitrarily close to the one generated by it. Under these conditions Eq.(II.37) simplifies to

$$\log \boldsymbol{\sigma}(e^{G'\delta t}) \prec \tfrac{1}{2}\boldsymbol{\lambda}\left(G^0+G^{0\dagger}\right)\delta t, \tag{II.38}$$

by substituting $T \to \delta t$ in Eq.(II.37). Calculating the singular values on the left hand side of that expression in term of the eigenvalues leads to

$$\log \boldsymbol{\lambda}(e^{G'\delta t}e^{G'^{\dagger}\delta t})^{\frac{1}{2}} \prec \tfrac{1}{2}\boldsymbol{\lambda}\left(G^0+G^{0\dagger}\right)\delta t \tag{II.39}$$

$$\log \boldsymbol{\lambda}\left(e^{(G'+G'^{\dagger})\frac{\delta t}{2}+O(\delta t^2)}\right) \prec \tag{II.40}$$

$$\boldsymbol{\lambda}\left(\frac{G'+G'^{\dagger}}{2}\delta t+O(\delta t^2)\right) \prec \tag{II.41}$$

$$\implies \boldsymbol{\lambda}\left(G'+G'^{\dagger}\right) \prec \boldsymbol{\lambda}\left(G^0+G^{0\dagger}\right), \tag{II.42}$$

where only the leading order in δt contributes in the last line. The unital version of this relation also holds provided both generators are trace preserving for the reasons discussed in §II.A.2. As might be expected from the previous results, this condition implies that $\mathrm{Tr}[G'] = \mathrm{Tr}[G^0]$ is also required.

There are cases where the time it takes to simulate the dynamics is not of concern, which corresponds to the traces of G and G' not being equal. In such cases, the condition above can be relaxed to

$$\tfrac{1}{2\mathrm{Tr}[G']}\boldsymbol{\lambda}\left(G'+G'^{\dagger}\right) \prec \tfrac{1}{2\mathrm{Tr}[G^0]}\boldsymbol{\lambda}\left(G^0+G^{0\dagger}\right), \tag{II.43}$$

by a rescaling of time in Eq.(II.37). This comes by replacing T with $\frac{\mathrm{Tr}[G]}{\mathrm{Tr}[G']}T$ on the left hand side that that equation. This is a relaxation of Eq.(II.9), as it holds even if $\mathrm{Tr}[G'] \neq \mathrm{Tr}[G^0]$.

II.C.5 Unital qubit Lindbladians

For the case of qubits undergoing unital Lindbladian dynamics, Eq.(II.9) can be simplified to

$$\boldsymbol{\lambda}\left(L'\right) \prec \boldsymbol{\lambda}\left(L^0\right), \tag{II.44}$$

and is also sufficient, as was previously shown in [Yua11]. We derived this independently and using slightly different methods, and include our proof here for completeness. The proof that it is a necessary condition follows directly from Eq.(II.9) as the dissipative part of the Bloch representation of any unital qubit Lindbladian is symmetric [Len87] and drift Hamiltonians are not of interest (they can be cancelled out by controls). To prove its sufficiency, an explicit way to reach $\overline{L}'$ using a drift $\overline{L}_0$ and unrestricted Hamiltonian controls is provided. For simplicity, the time scale of the target dynamics is picked such that $\mathrm{Tr}[\overline{L}'] = \mathrm{Tr}[\overline{L}_0] = 1$. By using the eigenvalue decomposition, the target can be expressed as

$$\underline{\overline{M}} = e^{\underline{\overline{L}}'t} = UDU^\dagger = U \begin{pmatrix} e^{-\nu_1 t} & 0 & 0 \\ 0 & e^{-\nu_2 t} & 0 \\ 0 & 0 & e^{-\nu_3 t} \end{pmatrix} U^\dagger, \tag{II.45}$$

where the $-\nu_i$ are the eigenvalues of $\underline{\overline{L}}'$, and U is an element of $\mathbb{O}(3)$ (the definition of this group, and others, is given in §I.C.1). Furthermore, as $\underline{M}$ has a positive determinant and the diagonal block is positive, U can be chosen to have determinant $+1$, thereby restricting it to $\mathbb{SO}(3)$. In a similar way, the free evolution of the system for time t can be expressed as

$$e^{\underline{\overline{L}}_0 t} = WF(t)W^\dagger = W \begin{pmatrix} e^{-\mu_1 t} & 0 & 0 \\ 0 & e^{-\mu_2 t} & 0 \\ 0 & 0 & e^{-\mu_3 t} \end{pmatrix} W^\dagger, \tag{II.46}$$

where the $-\mu_i$ are the eigenvalues of $\underline{\overline{L}}_0$. As $e^{\underline{\overline{L}}_0 t}$ has a positive determinant, W can also be chosen to be in $\mathbb{SO}(3)$. Controls on the system allow the implementation of any $R = e^{\underline{\overline{H}}}$ which, as noted previously, corresponds to any matrix in $\mathbb{SO}(3)$. The control scheme to reach the target map corresponds to alternating free evolution and instantaneous controls as

$$UDU^\dagger = R_1 WF(t_1)W^\dagger R_2 .... WF(t_n)W^\dagger R_{n+1}. \tag{II.47}$$

Due to the group structure it is possible to choose $R_1 = UW^\dagger$, $R_{n+1} = WU^\dagger$ and relabel $W^\dagger R_j W = R'_j$. Next the R' are picked to be permutation matrices (which all lie in $\mathbb{SO}(3)$) such that Eq.(II.47) consists solely of diagonal matrices where every term is an exponential. This lets the previous matrix equation be expressed in the simple form

$$\begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} = \begin{pmatrix} \mu_1 \\ \mu_2 \\ \mu_3 \end{pmatrix} t_1 + \begin{pmatrix} \mu_1 \\ \mu_3 \\ \mu_2 \end{pmatrix} t_2 + ... + \begin{pmatrix} \mu_3 \\ \mu_2 \\ \mu_1 \end{pmatrix} t_6, \tag{II.48}$$

where $\sum t_i = 1$ due to the scale of L_0 that was chosen.This control scheme thus allows any $\underline{\overline{M}} = e^{\underline{\overline{L}}'t}$ to be reached provided the eigenvalues of $\underline{\overline{L}}'$ are a convex combination of those of $\underline{\overline{L}}^0$; which is equivalent to saying that they are majorized by them [Bha97]. This

means that $\boldsymbol{\lambda}(L') \prec \boldsymbol{\lambda}\left(L^0\right)$ is a sufficient, as well as necessary, condition for reachability in unital qubit systems with unconstrained Hamiltonian control.

II.D Discussion

II.D.1 Evolution time

The first criterion, Eq.(II.5), is an equality which appeared in [WC08] for the Markovian and time-independent case. Extending it to the time-dependent case gives it an important use in control theory: it states that a target map may only be reached by a dynamical system for the times which satisfy Eq.(II.5). This can be pictured by noting that the modulus of the determinant of M is the volume occupied by its image, and that the trace of G_t is the rate at which this state space is growing (this will be non-positive, unless the system is non-Markovian, leading to a contraction of the space). The interpretation of this result is thus that the total rate at which volume is lost in state space is independent of the Hamiltonian controls. More insight can be gained by noting that the trace of the drift is always positive in the physically sensible case of the generator being Hermiticity preserving. This means that the evolution of such a system can only reach maps with positive determinant. As this is not the case for every completely positive trace preserving map, this condition immediately rules out large sections of the space as unreachable for a broad class of dynamical systems [WC08, DHKSH09].

Further restricting the drift to be a Lindbladian (with a non-vanishing dissipative part) at all times confines the trace of the generator to be negative, signifying that a target map can only be reached at a single instant in time (if at all), and that this time can be easily calculated as it is independent of the controls. If the drift is non-Markovian (where the interplay between memory effects and controls has received much recent attention [MP14, RKK15]) the trace of the drift can be positive for certain times, leading to revivals in the determinant. Indeed, this has already been suggested as an indicator of non-Markovianity [LPP13]. One consequence of this particular feature is the possibility of there being several solutions to Eq.(II.5) for a given target map. In both cases, the precise information about the required evolution time given by this condition is in stark contrast to the case of closed systems, where in general very little is known about the time required to reach a target without explicitly solving for the evolution of the controlled system.

II.D.2 Anisotropy of the dynamical map

The conditions of Eqs.(II.6-II.7) are a refinement of the intuition that controlling the system allows the decay rates of the drift to be averaged together. This is most easily seen

by considering the unital case. The left hand side of the relation contains the singular values of the dynamical map which, in the same way as the determinant is the volume in state space, are the characteristic lengths of the final state space. Thus, while Eq.(II.5) determines the volume reached, Eq.(II.7) provides a constraint on the anisotropy of the dynamical maps that can be reached.

The non-unital majorization relation has broadly the same interpretation, although the overall shift caused by the non-unitality manifests itself in the decay rates and singular values in a complex way. Indeed, one of the eigenvalues of $G_t^0 + G_t^{0\dagger}$ will typically be positive in the non-unital case, even if the drift is Markovian. While the idea of a positive decay rate in a Markovian system may appear counterintuitive, it only signifies that some states become purer under such a Lindbladian. This is most easily seen by considering the Bloch sphere: negative eigenvalues correspond to states moving towards the maximally mixed state, but if the system is decaying to the state $|0\rangle$ then there is also a dynamic evolution away from the centre towards a pure state on the surface of the sphere. Thus, despite their similar form, Eqs.(II.6-II.7) give very different results and there are many dynamical maps that satisfy one but not the other for a given drift (as is shown in Fig.II.3). It is also worth noting that Eq.(II.5) is recovered, up to a modulus, by the last term in the majorization relations.

II.D.3 Maximal non-unitality

This last condition on dynamical maps, Eq.(II.8), is conceptually very different from the others. Rather than restricting the shape of the dynamical map, it provides a series of constraints on how much the maximally mixed state can be displaced, corresponding to where the centre of the image of the dynamical map lies in state space. Although the right hand side of Eq.(II.8) is independent of controls, the maximisation over all states (and over all t if the drift is time-dependent) makes this criterion somewhat harder to evaluate in higher dimensions. It is worth noting that the constraint is less strict than $G_t^0(\rho) = 0$, which means that the non-unitality is not bounded by the fixed points of the drift. The interpretation of this criterion is therefore that it is possible to increase the ability of noise to purify states by using controls, but only up to the limits given.

A physical example of this is a three-level system in a Λ configuration (such as in Fig.II.4), with the excited state decaying into the two low level states. In the absence of controls, the system has some non-unitality as the maximally mixed state over the three levels will decay to a mixed state over only two levels. With the use of controls, however, the population can be coherently transferred back from one of the two ground states to the excited state where it will once again decay. Doing this many times results in the total population being transferred to the other ground state and the total action of the dynamics is to map everything to a single pure state. Thus, this dynamical map induced

by a specific set of controls has maximal non-unitality. This is the principle behind optical pumping and shows that non-unitality can be increased with controls. If the two lower levels had some decay between them, however, this scheme may not work perfectly and Eq.(II.8) provides bounds for how well it can be done.

II.D.4 Generator anisotropy

Instead of investigating if the system can reach a target map, the issue of whether it can be made to approximate a different drift continuously in time is examined by Eq.(II.9). To do this, the restricted case of time-independent G_t^0 is taken and the question is if it can give rise to evolutions arbitrarily close to $M = e^{G't}$ for all t. If it can, then G^0 is defined as being able to effectively simulate G' as it can replicate the same dynamics arbitrarily well in a time-continuous fashion. Necessary but not sufficient criteria to do this are given by Eqs.(II.9-II.10). These are stricter than the anisotropy conditions on dynamical maps; it imposes not only a target map but the whole trajectory in time to it. That this can be done at all is at first hand surprising, as the only generators which can be reached exactly from a given drift are given precisely by the drift plus all possible controls. However, by quickly rotating the system it is possible to get arbitrarily close to the required trajectory by winding tightly around the desired path without ever moving exactly along it as illustrated in Fig.II.2.

In the case of the generator G^0 being a unital qubit Lindbladian, Eq.(II.9) can be simplified further to $\boldsymbol{\lambda}(L') \prec \boldsymbol{\lambda}\left(L^0\right)$ and is sufficient [Yua11]. Furthermore the eigenvalues of such unital qubit Lindbladians, $\boldsymbol{\lambda}(L) = -(\lambda_1, \lambda_2, \lambda_3)$, are constrained by complete positivity to obey [GKS76]

$$\lambda_i + \lambda_j \geq \lambda_k \qquad \text{for } i,j,k \in \{1,2,3\} \text{ all unequal,} \tag{II.49}$$

and normalising them adds the additional constraint that $\lambda_1 + \lambda_2 + \lambda_3 = 1$. Combining these two expressions together allows the spectrum of any qubit unital Lindbladian to be written as

$$\begin{aligned} \boldsymbol{\lambda}(L) &= -(\lambda_1, \lambda_2, 1-\lambda_1-\lambda_2) \\ \text{where} \quad \lambda_1 \leq \tfrac{1}{2} \quad \lambda_2 \leq \tfrac{1}{2} \quad \lambda_1 + \lambda_2 &\leq \tfrac{1}{2}. \end{aligned} \tag{II.50}$$

Any spectrum can therefore be written as convex combinations of $\boldsymbol{\lambda}(L) = -(\frac{1}{2}, \frac{1}{2}, 0)$, which is equivalent to saying that it majorizes the spectrum of all other such Lindbladians. Hence, one with such a spectrum, such as the dephasing Lindbladian introduced in §I.B.6, is universal and it can simulate all other unital qubit Lindbladians. Conversely, the completely depolarising channel [NC00] with eigenvalues $\boldsymbol{\lambda}(L) = -(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is majorized by all other Lindbladians and is therefore at the bottom of the hierarchy defined by ma-

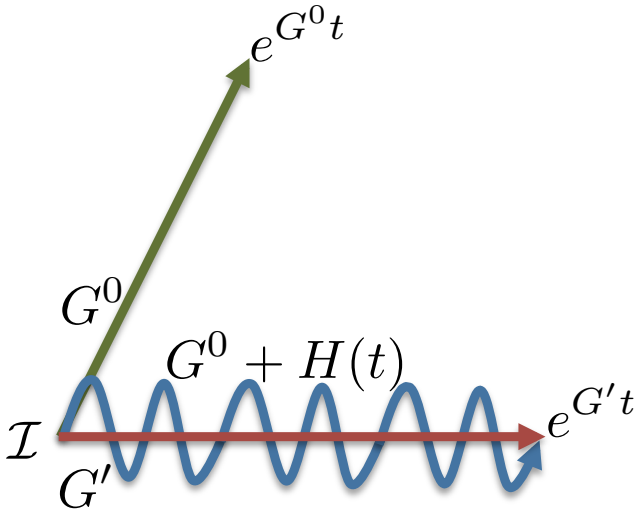


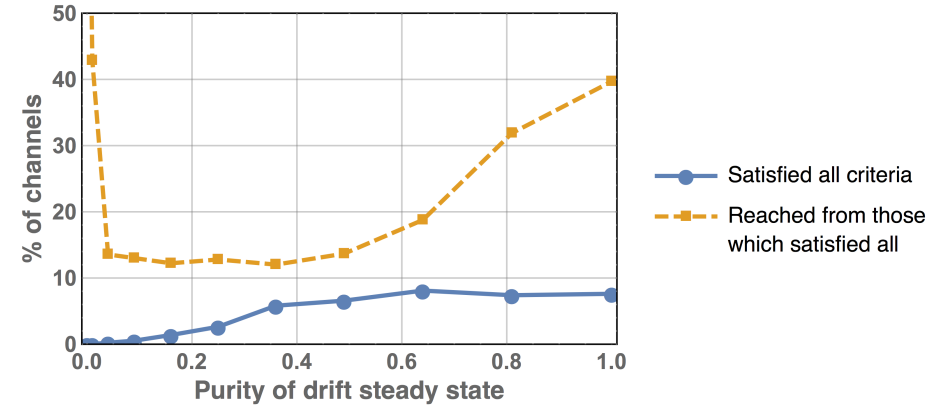
Figure II.2 | Sketch of the action of generators in the space of CPT maps showing the notion of effective simulation. The arrows represent evolution in time under different generators. They all start at the identity map, $\mathcal{I}$, at $t = 0$. The green arrow is the evolution of the system in the absence of control; at time t the resulting dynamical map is $M = e^{G^0 t}$. Similarly, the red arrow shows the desired evolution under a different generator G'. The aim of simulating a generator is to steer the system not just towards $M = e^{G' t}$ for a fixed t, but to follow the trajectory of the red line at each point in time. This cannot be done exactly unless G^0 and G' differ only by a Hamiltonian component. However, it is more generally possible to effectively simulate G' by $G^0 + H(t)$ in a way that the system evolves in a trajectory which is arbitrarily close to the desired one (such that it crosses it every δt), as shown by the blue arrow.

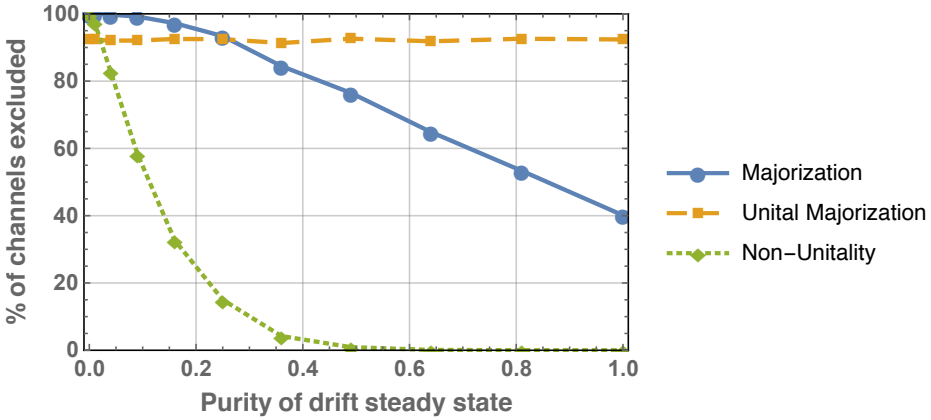
jorization and can be effectively simulated by any other unital qubit Lindbladian.

II.E Numerical Examples

The conditions Eqs.(II.6-II.8) are necessary but not sufficient, so the question of how tight they are is important. Our numerical simulations quantify this with two different measures: firstly by how often the criteria forbid a target from being reached, and secondly by how often a target can be reached when it is not excluded. The issue with doing this is that, for the very reason that the criteria derived in this paper are useful, it is computationally very difficult to test if a dynamical map can be reached with a given drift and controls. The closest to a definite method is solve the optimal control with numerical techniques. The size of the simulation itself scales with the fourth power of the dimension of the system, and the cost of optimising the control pulses scales far worse [MSGdFG11]. Nevertheless, results for a class of common non-unital qubit and qutrit Lindbladians are presented here that were found using the optimal control packages in QuTiP introduced in §I.C.3.

(a) Fraction of channels passing all conditions that could be reached





(b) Fraction of channels that satisfy each condition

Figure II.3 | Plot of the strength of Eqs.(II.6-II.8) for generalised amplitude damping Lindbladians. The x-axis is the purity of the drift steady state, a function of the bath temperature. The top graph shows what fraction of the randomly generated channels satisfied all of Eqs.(II.6-II.8) and, out of those, how many could be reached to within a distinguishability (given by the diamond norm) of at least 0.1% by optimal control. The bottom graph shows what fraction of the same channels are ruled out by each of the conditions individually.

For the qubit system the Lindbladian investigated is generalised amplitude damping [NC00], a ubiquitous type of noise, which is a slight modification of amplitude damping as introduced in §I.B.6. Rather than decaying to the ground state, the system relaxes to a mixture of the ground and excited state, corresponding to a qubit which can exchange an excitation with a bath at finite (instead of zero) temperature. This can be thought of as a spin which has a finite rate for transitioning from the excited to the ground state and from the ground to the excited state, where the ratio between the two is a function of temperature. The temperature determines the non-unitality of the noise: at zero temperature the steady state is the pure ground state, while at infinite temperature it is the maximally mixed state. As the non-unitality of the Lindbladian is an important aspect of its controllability, the purity of the steady state is used as a parametrisation of tempera-

ture. Our results show the ratio of randomly generated time-dependent Markovian maps [WECC08] which could be reached numerically, and whether they satisfied the criteria Eqs.(II.6-II.8), for different values of the non-unitality of the drift. The evolution time used when looking for solutions was determined by Eq.(II.5).

The results for such a Lindbladian at different temperatures are shown in Fig.II.3. Taken together, the criteria state that over 90% of the space is unreachable at each temperature considered, showing that the conditions are useful as they cut out the large majority of dynamical maps as impossible to achieve. The insufficiency of the criteria manifests itself in that - at some temperatures - only 10% of those not ruled out can be reached. This number, however, approaches 100% in the unital case, which is expected as it has been proved that the majorization condition is sufficient in unital qubit systems. That this figure rises again for highly non-unital, low temperature baths shows that the criteria are increasingly useful in this limit too. It is also interesting to note that the relative importance of the different conditions varies with temperature: when the noise has a pure fixed point the unitality criterion provides no information and the unital majorization criterion is the most restrictive, while their importance is reversed when the fixed point is maximally mixed.

The second example is another common type of noise, sketched in Fig.II.4, is a qutrit in a non-symmetric Λ configuration where the top level decays to the two lower levels according to the Lindbladian

$$L(\rho) = \sum_{i=1,2} \gamma_i \left(L_i \rho L_i^\dagger - \tfrac{1}{2}\{L_i^\dagger L_i \,,\, \rho\} \right), \tag{II.51}$$

where $L_i = |i\rangle\langle 3|$ [SS04]. The focus is on how the controls and the skew, γ_1/γ_2, can influence the asymmetry of the final evolution. To do this a drift is picked with a fixed skew, and how close the system could get to maps generated by a similar drift but with a different skew is investigated. Fig.II.5 shows that those which were as or more symmetric (a skew

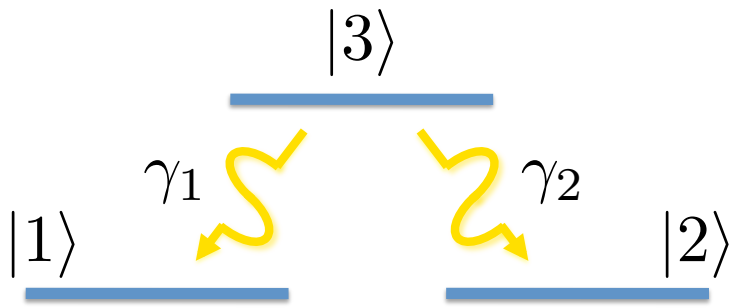


Figure II.4 | Energy levels of a qutrit in a Λ configuration with asymmetric decay rates. The Lindbladian causes the excited level to decay to the two ground state at different rates; the skew, γ_1/γ_2, characterises the asymmetry in the system.

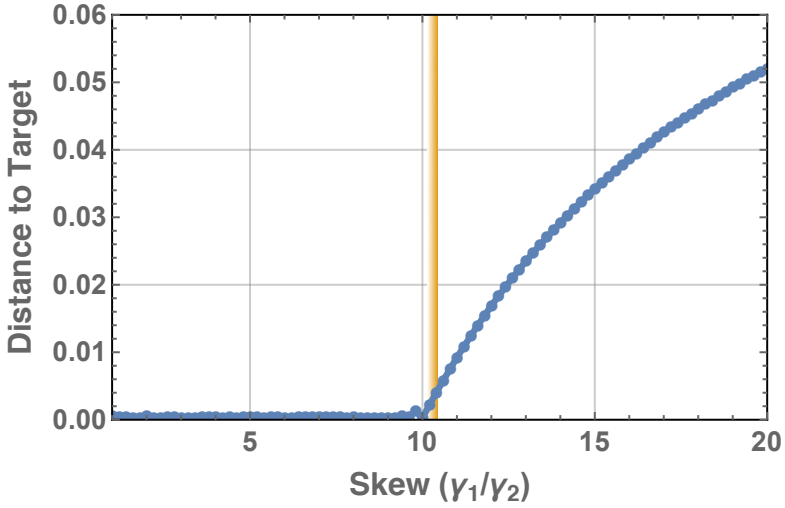


Figure II.5 | Plot of the reachability of CPT maps generated by non-symmetric Λ qutrit Lindbladians. The system drift, defined in Eq.(II.51), has a skew of 10 and the target maps were generated by the same Lindbladian but with skews between 1 and 20. Plotted is the minimal distance (given by the diamond norm) to targets with different skews that could be reached by optimal control. The thick vertical line at 10.3 is the non-reachable boundary given by the majorization criteria - everything to the right of it is excluded - in excellent agreement with the numerical results. The slight bump just below 10 is due to the difficulty of numerically finding solutions which require rapidly oscillating control Hamiltonians.

closer to 1) as the drift could be reached with a very high fidelity, and increasingly poorly those which were less symmetric. This is in excellent agreement with the majorization criteria as plotted. The tightness of the necessary conditions in this scenario show how useful they are in cases where there is a clear measure of non-uniformity, demonstrating that controls can be used to go from a highly ordered evolution to a less ordered one, but not the other way around.

This result may at first hand appear to contradict the conclusion arrived at in §II.D when this example was also discussed on maximal non-unitality in the context of optical pumping. There the conclusion was that a qutrit in a Λ configuration could have a pure fixed point regardless of the ratio of the decay rates; while here the results are that skew cannot be increased. The resolution of this problem is that although the fixed point of the dynamics can be chosen independently of the skew, this only determines the evolution at $t \to \infty$, at all other times the state space occupies a finite volume and the shape of this volume is what is constrained by the skew.

In general it is expected that the tightness of the criteria would decrease as the system size increases. The main argument for this is that the number of rotations in the state space that do not correspond to Hamiltonian degrees of freedom grows with dimension. That the conditions were tighter in the qutrit than the qubit example, however, suggests that there are specific cases where they remain an excellent approximation to the

allowed operations. Even when less tight, the potential usefulness of the conditions may be greater for larger systems as it is substantially harder to simulate these and therefore to learn about their controllability via other means.

II.F Relations to Thermodynamics

II.F.1 Super-operator lift of the Second Law

The focus of this chapter has been the majorization conditions Eqs.(II.6-II.7) and Eqs.(II.9-II.10) which provide limits on the operations that can be reached on an open quantum system. It is interesting to note that majorization also plays a key role on the related question of state, rather than operator, controllability. A central result there is that a state σ can be reached from a state ρ by a unital completely positive trace preserving map if and only if $\boldsymbol{\lambda}(\sigma) \prec \boldsymbol{\lambda}(\rho)$ [BZ06]. This is a stricter form of the Second Law of Thermodynamics as the majorization relation $\boldsymbol{\lambda}(\sigma) \prec \boldsymbol{\lambda}(\rho)$ imposes the constraint $S(\sigma) \geq S(\rho)$ on the von-Neumman entropy [NC00] of the two states, but the converse is not always true. This means that, under unital evolution, the eigenvalues of a state cannot become more ordered and therefore the entropy cannot decrease.

The natural extension of this formulation of the Second Law to the results on generator controllability, Eqs.(II.9-II.10), is that it lifts it from applying to states to super-operators. This is a restatement of what this chapter details: that generators of dissipative dynamics cannot become more ordered by the presence of coherent controls. These relations imply that the process is irreversible, controls can be used to make an existing generator of noise G^0 arbitrarily close to a different generator G', but the reverse cannot be done even approximately. This is a surprising result as the controls themselves are coherent and therefore fully reversible. The rise of irreversibility from purely reversible pieces is a long standing puzzle of quantum mechanics and a key aspect of the Second Law, the criteria developed here shows that it applies to super-operators as well as states.

Irreversibility in classical thermodynamics comes from averaging or coarse-graining the state as it evolves; if there was perfect knowledge of the initial and final micro-state then entropy would play no role. The same thing happens in this context. On the super-operator lift of the Second Law everything is similarly reversible when only exact simulation of a generator by another is considered. The irreversibility arrises when effective simulation is taken into account, which can be interpreted as having imperfect knowledge of the dynamics as the evolution under the controls and G^0 only match the desired generator G' at intervals δt apart (as shown in Fig.II.2). Thus, effective simulation and its irreversibility can be thought of as exact simulation on top of coarse-graining, the latter being the source of the irreversibility.

II.F.2 Entropy generation

A more explicit link between the present results and thermodynamics can be found by considering the rate of change of the entropy of a system as it undergoes unital evolution. An expression for this in terms of the spectral properties of a channel is given in [Rag02] which, modified using Eq.(II.6), gives for Markovian channels:

$$\frac{d}{dt}S[\rho(t)] \geq \frac{\lambda_1}{2}||\rho(t) - \tfrac{\mathbb{I}}{d}||_2^2, \tag{II.52}$$

where $\rho(t)$ is a state evolving under a unital Lindbladian (possibly under the presence of controls), $||\cdot||_2$ is the L_2 norm, and λ_1 is the smallest (in magnitude) eigenvalue of the unital part of the Hermitian component of the Lindbladian. While the left hand side is any channel generated by a Lindbladian and control, the lower bound is the smallest decay rate of the Lindbladian and independent of the controls. This shows that the minimal rate of entropy production cannot be lowered. The physical picture is that rotating a system as it decays cannot increase how well the best protected states are shielded from the production of entropy.

II.G Conclusion

The principle idea behind this chapter is that the decay rates of the generator of an open system, the eigenvalues of the sum of the generator and its dual, provides limitations as to the operations the system can achieve with coherent controls, resulting in the conditions of Eqs.(II.5-II.10). Our results show that the decay rates can be made more isotropic with the use of coherent controls, corresponding to the rates being averaged out by rotations. The total rate of decay however cannot be changed, and it is not possible to create a more ordered structure or to increase the non-unitality beyond a given limit. These hold for a range of open quantum systems, going beyond Markovian ones, within some assumptions which are discussed earlier. This work is therefore a direct, if incomplete, answer to one of the main questions of this thesis: the controllability of open quantum systems.

On their own our results are a useful tool in the quest for designing quantum systems to achieve desired non-unitary tasks, as they rule out some dynamics as impossible without the high cost of simulation and optimisation. Two examples of the use of the criteria were investigated numerically highlighting that, although they are necessary but not sufficient, they still give a practical approximation to the allowed operations. Due to the partial order induced by majorization, it shows that some types of noise are 'more useful' for simulations than others as they can be used to replicate all the same evolutions, plus others.

There are numerous possible directions that this work could be developed further

that we did not pursue. One would be to apply it to different types of systems, such as harmonic oscillators. A restriction to purely Gaussian states and operations is likely to make such a task trackable. Another avenue would be to consider the system as being multipartite (for example, two qubits) where the only controls available are local, and the subsystems are coupled by either sharing the environment or having an (uncontrollable) coupling Hamiltonian. The reachable set of such a system would contain information about whether it is possible to protect some information in the system by shuffling the action of the noise, such as by using decoherence free subspaces. Our initial attempts at doing this met with the difficulty of how restricting the set of allowed operations affects the allowed 'mixing' of the different decay rates, but this remains an enticing avenue of research. Finally, further investigations into the link with thermodynamics may also provide fresh insights into microscopic heat engines.

CHAPTER III

MINIMAL DILATIONS

III.A Introduction

This chapter covers our work on how closed systems in finite dimensions can be used to replicate the dynamics of an open system in a time-continuous way. It can be used as a method for simulating Lindbladians or other types of noise using ancillas and time-dependent Hamiltonians, rather than with another open system as was considered in the previous chapter. By showing that such Hamiltonians exist and what their properties are, these results can also be seen as a Stinespring theorem (introduced in §I.B.2) on the level of generators rather than propagators. This work was started during the MRes that preceded the PhD, so there is some overlap with our work in [Div14]; the completed project was also published independently in [DMB15], and the presentation of the results here is drawn from that paper.

III.A.1 Overview

The central idea of this chapter is to find, for a time-dependent quantum channel which describes the evolution of an open system, a finite closed quantum system with a well behaved Hamiltonian such that the evolution matches at all times. The key result is that it is always possible to construct the dilation Hamiltonian to be continuous in time and bounded while matching the reduced dynamics arbitrarily well. This extends previous results that two channels which are close by have unitary dilations which are also close by [KSW08], to show when smoothly varying channels result in smoothly varying dilations.

Although the methods presented here are not limited to Markovian channels, particular attention is paid to these as they constitute some of the most common types of noise encountered in quantum information. In the case of a non-trivial Markovian channel, our results show that the dilation must diverge at $t = 0$, although this can be avoided while causing only arbitrary small errors. More generally, this work can be understood as raising dilations from the level of propagators to the level of generators; mapping Lindbladians

into Hamiltonians.

Having such a dilation allows the dynamics of a given open system to be simulated continuously in time by acting on a finite system with a physically sensible Hamiltonian. The motivation for this is to perform quantum simulations, a key aim of quantum technologies [CZ12, GAN14], for open systems. As every physical system invariably interacts with its environment, any accurate simulation must take into account dissipative effects. The question of how this can be done efficiently is one which has received significant research interest in recent years [KBAL05, ZŠB05, KBA05, KBGKE11, RFZB12, SSP14]. Being able to do the simulation continuously in time provides additional benefits compared to a snapshot simulation, particularly when the time at which the system will be measured is not known beforehand. It is also applicable in cases where the system is monitored continuously via weak measurements [Jac14]. This allows information about the behaviour of the system over an interval of time to be recovered; a situation where prior approaches which evolve the system to a fixed point in time would fail.

As is discussed below in §III.A.2, previous work on dilating Lindbladians typically results in the ancillary space being infinite dimensional, and the coupling Hamiltonian being time-independent. The work here requires only a finite dimensional ancilla, but has a time-dependent Hamiltonian. By shifting the complexity from infinite space to time-dependence, it provides a model on which open systems can be studied easily which is used to investigate how the system-environment interaction is affected by adding controls to the system.

The chapter begins with §III.A.2 where prior methods for generating dilations of open dynamics are reviewed. After this our new method for constructing dilations that are both in finite dimensions and continuous is presented §III.B, including both the core method and extensions for more complex systems. Properties of the dilation Hamiltonian, such as its continuity and boundedness, are also discussed. §III.C contains details of specific aspects of the proofs. Examples of this procedure for the spin-boson model and amplitude damping are given in §III.D. The conclusion in §III.E discusses of the use of such dilations for quantum simulations, as well as possible links between them and infinite dimensional dilations.

III.A.2 Prior Methods

A dilation is a closed quantum system that replicates the dynamics of a noisy one. This is always possible to do because, up to the approximations made in deriving their dynamics (discussed in §I.B.4), open systems arise by the environment being traced out and deliberately removed from the quantum system being considered; the universe as a whole is thought to be unitary. One dilation is therefore the original environment of an open system with its original interaction, although it is far from only the one. There are broadly

speaking two different classes of dilations, ones which replicate the evolution of the system at a fixed time t, and ones which do so from all times. The former is any $\{U, \rho_B\}$ that solves

$$M(\rho) = \mathrm{Tr}_B[U\rho \otimes \rho_B U^\dagger], \tag{III.1}$$

for a target map M and for all ρ. As long as M is completely positive and trace preserving, Stinespring's theorem guarantees that this can be done, as was discussed in §I.B.2. This chapter focusses on the more involved case where M becomes time-dependent, and therefore so does U; the task then becomes to find the Hamiltonian that generates such a unitary. Before introducing a method which gives such a dilation with a finite dimensional bath and a time-dependent Hamiltonian, existing ways of creating a dilation with an infinite dimensional bath and a time-independent Hamiltonian are discussed.

The traditional way of approaching this problem is to treat the dynamics of the system not in terms of quantum channels, but as a quantum stochastic differential equation [HP84a, Par92]. The dilation arises by considering the stochastic action of the noise as an interaction with a sample space which determines the underlying probability distribution. The solution to the equation governing the system and bath ensemble obeys some group theoretic properties (specifically, it has the cocycle property) and so can be associated to a unitary group, which can in turned be mapped to a time-independent Hermitian operator on the larger space [Fri85, Maa85]. Furthermore, it can be shown that the generator of this group is self-adjoint and therefore a valid Hamiltonian [Gre01]. The end result of this is that a time-independent Markovian evolution can be dilated to unitary evolution, with a time-independent Hamiltonian on the system and an infinite dimensional bath. The bath, however, does not have any direct physical interpretation and the Hamiltonian is not physically well behaved.

A different way of thinking about a dilation that matches the dynamics at all times, provided once again that the reduced dynamics are time-independent, is to use collision models [ZŠB05, KBA05]. The starting point is a dilation that matches the evolution at a fixed time only, as in Eq.(III.1). If this time δt is picked to be short, the unitary U naturally gives rise to the Hamiltonian $H = \frac{i}{\delta t}(U - \mathbb{I})$. To simulate the evolution for large times, the system interacts with a sequence of ancillas one at a time, each in the state ρ_B and for time δt. In the limit that $\delta t \to 0$, this replicates the reduced dynamics exactly. While each ancilla can itself be finite dimensional, as there must be a continuum of them, the total ancilla space also becomes infinite. This is thus a different method to construct dilations that requires access, and control, to an infinite bath, or for a single qubit to be reused many times [SSP14].

III.B Minimal Dilation Construction

A sketch of our process to calculate the dilation Hamiltonian is shown in Fig.III.1. The individual steps are individually already well known and used in many different problems. The novelty is in putting them together, and in tracking the behaviour of all the elements as a function of time in order to derive properties of the dilation Hamiltonian.

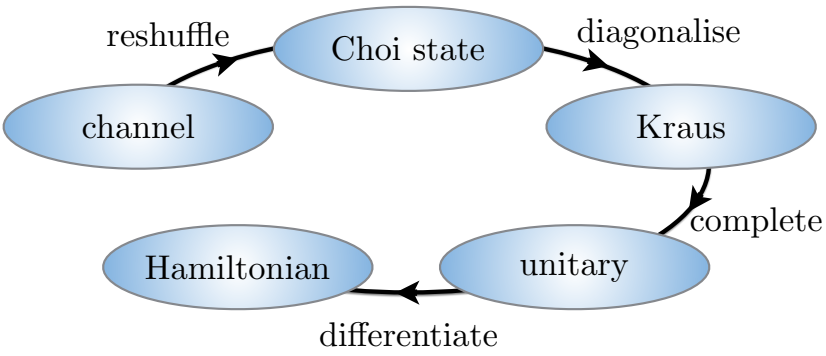


Figure III.1 | Schematic of how to construct a dilation for a family of quantum channels. All the steps preserve the analyticity in time apart from separating the diagonalised Choi state into Kraus operators which introduces a square root. It is differentiating this square root that introduces the possibility of discontinuities and divergences.

The starting point is a family of quantum channels, M_t, which is analytic in time. As has been discussed in §I.B.3, for Markovian channels, these are generated by a Lindbladian according to $M_t = \mathcal{T} \exp\left(\int_0^t d\tau L_\tau\right)$. The method derived here however works for any analytic family of quantum channels, generated by a time-local master equation or not. The first step is to construct the Choi state as given by the Choi-Jamiołkowski isomorphism (as discussed in §I.B.2), which is explicitly written as $\rho_M(t) = (M_t \otimes \mathcal{I}) |\Omega\rangle \langle\Omega|$ where $|\Omega\rangle$ is an (unnormalised) maximally entangled state between the system's Hilbert space and its duplicate. This is equivalent to reshuffling the elements of the channel represented as a matrix as following the method introduced in §I.B.5. This is just a linear transformation, therefore $\rho_M(t)$ is also analytic in t.

As $\rho_M(t)$ is Hermitian and positive, it can be decomposed into its eigenvalues and eigenvectors, $\lambda_k(t)$ and $|v_k(t)\rangle$ respectively, where the index k runs from 1 to the Kraus rank of the channel, R, which is upper bounded by the dimension of the system squared d^2. It is known via perturbation theory that, as $\rho_M(t)$ is Hermitian and analytic, its eigenvalues and vectors are also analytic for real t [Kat80]. This allows the channel to be expressed by its Kraus representation, $M_t(\rho) = \sum K_k(t)\, \rho\, K_k^\dagger(t)$, where the Kraus operators are

$$K_k(t) = \sum_{ij} \sqrt{\lambda_k(t)}\, \langle i \otimes \Omega | v_k(t) \otimes j \rangle\, |i\rangle \langle j| \,, \tag{III.2}$$

for any choice of basis $\{|i\rangle\}$, $\{|j\rangle\}$ over the initial Hilbert space. All the terms preserve analyticity, except for the square root which is not well behaved at $\lambda_k(t) = 0$. This results in the Kraus operators being continuous everywhere, and smooth everywhere apart possibly from individual points.

The final step in finding a unitary dilation is solving

$$\sum_{k=1}^{R} K_k(t)\,\rho\,K_k^\dagger(t) = \mathrm{Tr}_B[U(t)\,\rho_A \otimes \omega_B\,U^\dagger(t)], \tag{III.3}$$

where the system space has subscript A and the ancilla B. There is freedom in picking the initial ancilla state and so it can be chosen to be $\omega = |0\rangle\langle 0|$, which gives as a solution

$$\langle k_B|\,U(t)\,|0_B\rangle = K_k(t) \;\; \forall\, k. \tag{III.4}$$

This provides a dilation where the dimension of the ancilla is R, which is the smallest sized ancilla which may be required. The dilation unitary has dimensions $Rd \times Rd$, and Eq.(III.4) constrains d of its columns. These can be thought of as forming d orthonormal vectors in an Rd dimensional space. As the desired dilation is unitary, there is the additional constraint that $U(t)U^\dagger(t) = \mathbb{I}$, which is equivalent to requiring that the remaining $d(R-1)$ columns complete the orthonormal space. Our proof, given below in §III.C.1 using Gram-Schmidt and perturbation theory, shows that this can always be done in a smooth way whenever the $K_k(t)$ are smooth. In practice, it is often possible to fill in the extra columns by making use of symmetries in the Kraus operators without resorting to numerical techniques. The desired Hamiltonian satisfies the Schrödinger equation, and is therefore given by

$$H(t) = i\,\dot{U}(t)U^\dagger(t). \tag{III.5}$$

The Hamiltonian involves the derivative of U, which is continuous but in general not analytic in t, thus there is the possibility of it being discontinuous or divergent due to the behaviour of the derivative of $\sqrt{\lambda(t)}$. It is therefore important to analyse potential singularities of

$$\frac{d}{dt}\sqrt{\lambda(t)} = \frac{\dot{\lambda}(t)}{2\sqrt{\lambda(t)}}. \tag{III.6}$$

There are two potential problems. Firstly, it can be singular when $\lambda_k(t \neq 0) = 0$. At such a point, $\dot{\lambda}_k = 0$ as the eigenvalues of the Choi state are non-negative. It follows from the analyticity of $\rho_M(t)$ that such singularities lead to discontinuities in $\dot{K}_k(t)$, but it remains bounded. The second type of singularity occurs when $\lambda_k(t = 0) = 0$ and $\dot{\lambda}_k(t = 0) \neq 0$, which can lead to $\dot{K}_k(t)$ being unbounded. These properties are inherited by the Hamiltonian, bar some accidental cancellation. The first case, where the Hamiltonian has a step change at some later time, corresponds to the Kraus rank of the channel

decreasing at a single point in time. We are not aware of any interesting physical type of noise which leads to such behaviour.

More interesting is the second case. In §III.C.2 below is a proof that the divergence of the dilation Hamiltonian at $t = 0$, is avoided if and only if the dissipative part of $\dot{M}(t = 0)$ vanishes. This gives the immediate corollary that all non-trivial time-independent Markovian channels lead to a divergent dilation. Such channels necessarily have that the survival probability of certain states decays linearly at short times. However, it is known from the quantum Zeno effect that the survival probability of a state in any unitary system with a bounded Hamiltonian must decay quadratically for short times [Pas14]. This linear decay might be a signature of the unbounded system-environment interaction, but often it is merely a consequence of approximations, like infinitely fast relaxation within the environment, which were made in the derivation of the original Master equation (as was done in §I.B.4). In both cases, the divergence of the dilation is inherited from this and is an indicator of time-independent Markovianity.

Both the discontinuities and the divergence are, when they happen, benign. The divergence at $t = 0$ has only a finite impact on the dynamics as the dilation can always be picked such that $U \to \mathbb{I}$ as $t \to 0$. This implies that the evolution can be arbitrarily well approximated by replacing the exact Hamiltonian with a bounded one. This can be seen from the relation between the error in the unitary (which is the error in the evolution) and that of the Hamiltonian [NDGD06]:

$$||U(t) - U_E(t)|| \leq \int_0^t ||H(t') - H_E(t')||dt', \tag{III.7}$$

where $||\cdot||$ is the operator norm and H_E and U_E are the exact Hamiltonian and unitary respectively. The right hand side is bounded by $\int_0^t ||H_E(t')||dt'$ for any reasonable choice of dilation Hamiltonian that approximates it. From its construction as the derivative of $||U_E(t)||$ this term is finite, and thus the error in the unitary is always finite and can be decreased arbitrarily by having the applied Hamiltonian differ from the exact Hamiltonian for a sufficiently short time. Our calculation of this error for a specific dilation (plotted in Fig.III.2), highlights that the approximation error can be made small easily. In the case of isolated discontinuities, the dilation Hamiltonian itself can be arbitrary well approximated by a continuous one, which leads to the evolution being arbitrarily well approximated. This gives the principle analytic result of this chapter:

> *For an analytic family of quantum channels acting on d dimensional states there always exists a continuous and bounded Hamiltonian acting on at most d^3 dimensional states such that the dynamics on the reduced system are arbitrarily well matched at all times.*

The method given above works well, but it requires the diagonalisation of the Choi state and, in the case that the starting point is a master equation, the calculation of

the channel. In practice, one or both of these may be very difficult to do analytically. Indeed, systems where these are hard to do are the ones where it is the most interesting to find a dilation for and be able to simulate on a quantum computer, as they are precisely those which are difficult to simulate classically. Therefore, three alternate methods were also developed by us to construct dilations for complicated systems which rely on having found a dilation for a simpler system. The methods are: changing frames, separating into commuting parts, and perturbation theory.

III.B.1 Changing Frames

Changing frames (such as moving to the interaction picture) is a common tool in solving quantum problems. This method is applied here to the finding of dilations. The starting point is the equation of motion of the open system which is described by a master equation:

$$\frac{d\rho}{dt} = G_t(\rho). \tag{III.8}$$

This can be represented in a different frame by the transformation $\tilde{\rho}(t) = U_0^\dagger(t)\rho(t)U_0(t)$ where $U_0(t) = e^{-iH_0 t}$. The equation of motion in this frame is given by:

$$\begin{aligned} \frac{d\tilde{\rho}}{dt} &= \tilde{G}_t(\tilde{\rho}) \\ \text{where:}\quad \tilde{G}_t(\cdot) &= U_0^\dagger G_t(U_0 \cdot U_0^\dagger)U_0 + i[H_0, \cdot]. \end{aligned} \tag{III.9}$$

The dilation of $\tilde{M}_t$ has unitary $\tilde{U}(t)$ and Hamiltonian $\tilde{H}(t)$, and can be related to the dilation of M_t. The channel in the two different frames are related by

$$M_t(\cdot) = U_0(t)\tilde{M}_t(\cdot)U_0^\dagger(t), \tag{III.10}$$

as $\tilde{\rho}(0) = \rho(0)$. This implies that

$$M_t(\rho) = U_0 \mathrm{Tr}_B \left[\tilde{U}\ \rho \otimes \omega\ \tilde{U}^\dagger\right] U_0^\dagger \tag{III.11}$$

$$\equiv \left[U\, \rho \otimes \omega\, U^\dagger\right], \tag{III.12}$$

which shows that the dilation in the original frame is given by $U = (U_0 \otimes \mathbb{I})\tilde{U}$. Differentiating this to find the Hamiltonian results in

$$H = H_0 \otimes \mathbb{I} + (U_0 \otimes \mathbb{I})\, \tilde{H}\, (U_0^\dagger \otimes \mathbb{I}). \tag{III.13}$$

This is a very simple relation between a dilation Hamiltonian in one frame, and a dilation in a different frame. Indeed, the relation is the same as when changing frames in normal unitary dynamics with the small additional step of extending $U_0 \to U_0 \otimes \mathbb{I}$ in order to make the dimensions match up correctly.

III.B.2 Separation into commuting parts

Another useful method for simplifying quantum problems is to split them into commuting parts. Our results show how a dilations can be calculated by separating the channel into commuting parts, at the cost of increasing the ancilla space. This requires starting with two different channels which commute at all times such that

$$M_t^{(12)}(\cdot) = M_t^{(1)}\left(M_t^{(2)}(\cdot)\right). \tag{III.14}$$

Performing the dilations one after the other, and on different ancilla spaces, leads to

$$M_t^{(12)}(\rho) = M^{(1)}\left(\mathrm{Tr}_B\left[U_{AB}^{(2)}\,\rho_A\otimes\omega_B^{(2)}\,U_{AB}^{(2)\dagger}\right]\right) \tag{III.15}$$

$$= \mathrm{Tr}_{BC}\left[U_{AC}^{(1)}U_{AB}^{(2)}\,\rho_A\otimes\omega_B^{(2)}\otimes\omega_C^{(1)}\,U_{AB}^{(2)\dagger}U_{AC}^{(1)\dagger}\right], \tag{III.16}$$

where the simplified notation $U_{AC}\equiv U_{AC}\otimes\mathbb{I}_B$ is used for this and other operators. This gives as the Hamiltonian

$$H_{ABC}^{(12)} = H_{AC}^{(1)} + U_{AC}^{(1)}H_{AB}^{(2)}U_{AC}^{(1)\dagger}. \tag{III.17}$$

Due to the commutativity, the order in Eq.(III.14) can be inversed, giving

$$H_{ABC}^{(21)} = H_{AB}^{(2)} + U_{AB}^{(2)}H_{AC}^{(1)}U_{AB}^{(2)\dagger}, \tag{III.18}$$

which is, in general, a different Hamiltonian but leads to the same dynamics on the reduced system.

III.B.3 Perturbative Expansion

Thirdly, perturbation theory can be applied to calculating dilations in cases where the exact channel cannot be calculated. In this case the Lindbladian for the system is $L_t^{(0)} + \delta L_t^{(1)}$ with $\delta \ll 1$, a dilation for $L_t^{(0)}$ is already known, and the task is to find a new dilation which gives the correct dynamics to first order in δ. The quantum channel for such a Lindbladian is

$$M_t = M_t^{(0)} + \delta\int_0^t M_{(t,\tau)}^{(0)}L_t^{(1)}M_{(\tau,0)}^{(0)}d\tau + O(\delta^2) \tag{III.19}$$

$$\approx M_t^{(0)} + \delta M_t^{(1)}, \tag{III.20}$$

where $M_{(t_2,t_1)}^{(0)} = \mathcal{T}e^{\int_{t_1}^{t_2}L_\tau^{(0)}d\tau}$. Constructing the Choi state from the channel is a linear transformation which can be done separately for $M_t^{(0)}$ and $M_t^{(1)}$. The eigenvalues and vectors of $\Lambda_t^{(0)} + \delta\Lambda_t^{(1)}$ can be found to first order in δ using standard methods from perturbation theory. This is much easier to do than diagonalising the Choi state exactly,

although some of the advantage is lost if the $\Lambda^{(0)}$ has a high degree of degeneracy which is broken. From this, the Kraus operators are found by expanding Eq.(III.2) to first order in δ. It is worth noting that if $K_k^{(0)} = 0$ it is sufficient to find K_k to $O(\sqrt{\delta})$, as the equation of motion is quadratic in the Kraus operators and only dynamics up to first order in δ are important. In the case that the Kraus rank of the channel is unaffected by the perturbation, so that the above condition does not hold, $K_k = K_k^{(0)} + \delta K_k^{(1)} + O(\delta^2)$. In that case the correction to the dilation unitary satisfies

$$\langle k_B | U^{(1)}(t) | 0_B \rangle = K_k^{(1)}(t) \ \ \forall \, k \tag{III.21}$$

$$U^{(1)} U^{(0)\dagger} + U^{(0)} U^{(1)\dagger} = 0. \tag{III.22}$$

The last expression is equivalent to requiring the unitarity condition to hold to first order in δ. This reduces the problem of finding the dilation unitary to a system of linear equations. The Hamiltonian is then given by

$$H^{(1)} = \left(i \frac{dU^{(1)}}{dt} - H^{(0)} U^{(1)} \right) U^{(0)\dagger}, \tag{III.23}$$

such that $H = H^{(0)} + \delta H^{(1)} + O(\delta^2)$.

In the case that the Kraus rank of the channel does change, similar expressions have been found by us, although care must be taken to ensure that both terms $O(\sqrt{\delta})$ and $O(\delta)$ are properly accounted for. This perturbative method can be extended to take into account second, or higher, order effects.

III.C Proofs

In the general method for constructing a minimal dilation given above, two statements were made without proof in order to maintain the flow of the description. These were that Kraus operators that are smooth in time always lead to a unitary which is smooth in time, and a condition for the dilation Hamiltonian to diverge. Our proofs of both statements are given here.

III.C.1 Completing the Unitary

In order to find the dilation unitary from the Kraus operators, it is necessary to solve Eq.(III.4):

$$\langle k_B | U(t) | 0_B \rangle = K_k(t) \ \ \forall \, k$$
$$U(t) U^\dagger(t) = \mathbb{I}.$$

It is always possible to find a solution for $U(t)$ which is smooth whenever the $K_k(t)$ are smooth, by following the method given here.

The first condition constrains d columns of the unitary as being orthonormal vectors and the second condition requires picking the remaining $d(R-1)$ columns such that they form a complete orthonormal set. At $t = 0$ Gram-Schmidt is used from an arbitrary basis such that the none of the $d(R-1)$ columns are 0. In order to ensure at later times that these extra columns vary continuously in time the vectors at t are used as an initial basis for Gram-Schmidt at $t + \delta t$. The columns of $U(t)$ can be expressed as the Rd vectors $|v_n(t)\rangle$. At time $t + \delta t$ the first d of these vectors are transformed according to the change of the Kraus operators. They are still orthonormal to each other, but no longer orthogonal to the other $R(d-1)$ vectors. Using Gram-Schmidt, the $d+1$ vector is updated by first calculating:

$$|v'_{d+1}(t+\delta t)\rangle = |v_{d+1}(t)\rangle - \sum_{n=1}^{n=d} \langle v_{d+1}(t)|v_n(t+\delta t)\rangle \ |v_n(t+\delta t)\rangle\,, \tag{III.24}$$

For $n \leq d$ the change can be expressed as

$$|v_n(t+\delta t)\rangle = |v_n(t)\rangle + \delta t\,|\Delta_n(t)\rangle + O(\delta t^2), \tag{III.25}$$

whenever the $|v_n(t)\rangle$ are a smooth function of t, that is, whenever the $K_k(t)$ are smooth. This enables Eq.(III.24) to be rewritten as

$$|v'_{d+1}(t+\delta t)\rangle = |v_{d+1}(t)\rangle - \sum_{n=1}^{n=d} \delta t\,\langle v_{d+1}(t)|\Delta_n(t)\rangle \ |v_n(t+\delta t)\rangle + O(\delta t^2), \tag{III.26}$$

due to the orthogonality at t. This explicitly shows that $|v'_{d+1}(t)\rangle$ is a smooth function of t. From this the normalised vector is found to be

$$\begin{aligned} |v_{d+1}(t+\delta t)\rangle = &|v_{d+1}(t)\rangle \\ &+ \delta t \sum_{n=1}^{n=d} \mathrm{Re}\left(\langle v_{d+1}(t)|\Delta_n(t)\rangle\,\langle v_{d+1}(t)|v_n(t+\delta t)\rangle\right)|v_{d+1}(t)\rangle \\ &- \delta t \sum_{n=1}^{n=d} \langle v_{d+1}(t)|\Delta_n(t)\rangle \ |v_n(t+\delta t)\rangle + O(\delta t^2), \end{aligned} \tag{III.27}$$

which is also a smooth function of t whenever the $K_k(t)$ are smooth. By induction, that this is true for all $|v(t)\rangle$ and all t. Hence it is always possible to construct a unitary which is smooth whenever the Kraus operators are smooth.

III.C.2 Divergence at $t = 0$

We prove that the dilation of $M(t)$ diverges at $t = 0$ if and only if $\dot{M}(0)$ is not a Hamiltonian operator

$$\dot{M}(0) \neq -i[H, \cdot]. \tag{III.28}$$

The first case to consider is where the dilation does not diverge, which implies

$$M_t(\rho) = \mathrm{Tr}_B[U(t)\rho \otimes \omega U^\dagger(t)] \tag{III.29}$$

$$\dot{M}_t(\rho) = \mathrm{Tr}_B\left\{-i[H(t), U(t)\rho \otimes \omega U^\dagger(t)]\right\} \tag{III.30}$$

$$\dot{M}_0(\rho) = -i\mathrm{Tr}_B\left\{[H(0), \rho \otimes \omega]\right\}, \tag{III.31}$$

as $M_0(\rho) = \rho$ means that there exists a dilation such that $U(0) = \mathbb{I}$. It is important to note that the last line would not hold if $H(0)$ was unbounded. $H(0)$ can be expanded in terms of separable Hermitian operators as

$$H(0) = \sum_k A_k \otimes B_k. \tag{III.32}$$

This enables the partial trace to be calculated according to

$$\dot{M}_0(\rho) = -i\mathrm{Tr}_B\left[\sum_k (A_k\rho \otimes B_k\omega - \rho A_k \otimes \omega B_k)\right] \tag{III.33}$$

$$= -i\sum_k [\lambda_k A_k, \rho] \tag{III.34}$$

$$= -i[H', \rho], \tag{III.35}$$

where $\lambda_k = \mathrm{Tr}(\omega B_k)$. These are necessarily real as B_k is Hermitian and ω is a state, from which it follows that H' must also be Hermitian. This proves that if the dilation is finite at $t = 0$, then $\dot{M}_0(\rho)$ must be purely Hamiltonian.

The proof of the converse statement, that if $\dot{M}(0)$ only contains Hamiltonian parts the dilation will not diverge, begins with the analysis of the eigenvalues of the Choi state that was done in Eq.(III.6). Specifically, that at least one of the eigenvalues must obey $\lambda(0) = 0$ but $\dot{\lambda}(0) \neq 0$. The Choi state is pure if and only if the quantum channel is unitary. If this is the case it is clear that the eigenvalues do not change, so the Hamiltonian part of the channel can be safely discarded. Assuming that $\dot{M}_0 = 0$, the quantum channel for short times is

$$M_t \approx \mathbb{I} + X\frac{t^2}{2} + O(t^3), \tag{III.36}$$

where X is an unknown super-operator. As previously stated, reshuffling the elements of M_t, written as a matrix, gives the Choi state $\rho_M(t)$. This means that the elements of $\rho_M(t)$ have, potentially, terms of every order in t *apart* from first order. As $\rho_M(t)$ is a

state, though not quite normalised, its eigenvalues are all non-negative and sum to d. It is also a pure state at $t = 0$ and, as the eigenvalues are analytic in t, they can be expressed as:

$$\lambda_0 = d + \alpha_0 t + \beta_0 t^2 + O(t^3) \tag{III.37}$$

$$\lambda_i = 0 + \alpha_i t + \beta_i t^2 + O(t^3) \quad i = 1, ..., d^2 - 1. \tag{III.38}$$

The eigenvalues are also obtained from the characteristic equation for $\rho_M(t)$, implying

$$\lambda^{d^2} + A\lambda^{d^2-1} + B\lambda^{d^2-2} + ... = 0, \tag{III.39}$$

where A, B, C,... are products and sums of the elements of $\rho_M(t)$. This means that none of these coefficients can have terms which are linear in t. The basic properties of the roots of polynomial equations give

$$B = \sum_{\mu \neq \nu} \lambda_\mu \lambda_\nu \quad \mu, \nu = 0, ..., d^2 - 1 \tag{III.40}$$

$$= d \sum_i \alpha_i t + O(t^2). \tag{III.41}$$

As already noted, B cannot have a term proportional to t. This implies that $\sum_i \alpha_i = 0$ and, as each are non-negative, that $\alpha_i = 0$. The sum of the eigenvalues hence also means that $\alpha_0 = 0$. Therefore, all the eigenvalues obey $\dot{\lambda}(t = 0) = 0$ and, as explained in §III.6, this is a sufficient condition for the dilation to not diverge.

III.D Examples

The methods of the previous sections are examined by looking at specific examples of quantum channel acting on qubits which represent some of the canonical decoherence models in quantum information. As shown below, in many cases the resulting Hamiltonians are sufficiently simple to be directly experimentally realisable.

III.D.1 Spin-Boson Model with Dilation Hamiltonian Cut Off

In the spin-boson model a single spin interacts with a bath of bosons via a constant Hamiltonian, such that the reduced dynamics are given by $\dot{\rho} = -\gamma(t)[\sigma_z, [\sigma_z, \rho]]$, where the decay rate is a function of time [DZWKH08, PS11]. This is the dephasing Lindbladian discussed in §I.B.6, except that the decay rate γ is time-dependent. The exact time-dependence of the decay rate depends on the spectral density and initial state of the bosonic bath, as well as the type of frequency cut off (here is a sharp cut off ω_c is used). In the case of an Ohmic spectral density with a bath temperature much higher than the cut

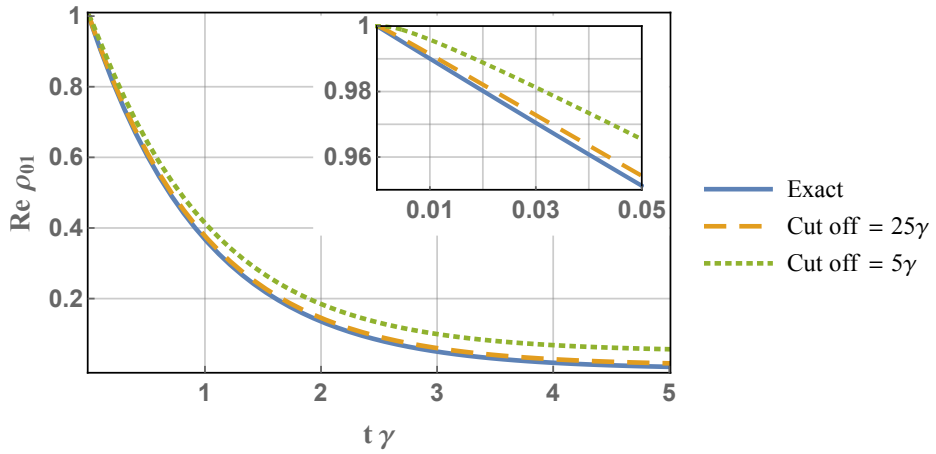


Figure III.2 | Plot of the decay of the $|+\rangle\langle+|$ state when subjected to a dilation of the dephasing channel. The case with exact divergent dilation Hamiltonian is shown, as well as when this Hamiltonian is capped to a maximum magnitude for short times. The inset shows the behaviour for short times.

off frequency, and for evolution times longer than $1/\omega_c$, the system becomes Markovian.

Performing the steps outlined in §III.B, one possible dilation is

$$H_{s.b.}(t) = \frac{\gamma(t)}{2\sqrt{e^{2\int_0^t \gamma(t')dt'} - 1}}\sigma_z \otimes \sigma_y, \tag{III.42}$$

where σ_z acts on the system and σ_y acts on the ancilla (which consists of a single qubit). For a typical bath spectral density $\gamma(0) = 0$ and, for certain values of t, it becomes negative. In such a case the channel is non-Markovian and the dilation is bounded and continuous for all times. In the case of constant γ, however, this is the Markovian dephasing Lindbladian whose dilation diverges at $t = 0$ as expected. If this Hamiltonian is approximated by a constant one by replacing the scalar pre-factor by C for short times (that is, when $\frac{\gamma}{2\sqrt{e^{2\gamma t}-1}} > C$) the error between the target unitary and the unitary reached is upper bounded by $\frac{\gamma}{8C} + O(\frac{\gamma^2}{C^2})$. Thus, with a sufficiently large C the error can be made arbitrarily small. Another way of looking at the errors is to see how the dynamics of a state depends on C, as is plotted in Fig.III.2. This shows that the main effect of introducing a cut off is to make the behaviour quadratic, rather than linear, at short time, and that even a modest value for the cut off is enough to reduce errors to insignificance.

III.D.2 Amplitude Damping with Driving

Another common type of noise is amplitude damping, as described in §I.B.6. In addition to decaying to the $|0\rangle$ state, the qubit also has an internal energy structure such that the equation of motion is

$$\dot{\rho} = -\gamma(\{\sigma_+\sigma_-, \rho\} - 2\sigma_-\rho\sigma_+) - i\tfrac{1}{2}\omega_0[\sigma_z, \rho], \tag{III.43}$$

where $\sigma_\pm$ is the raising/lowering operator. A minimal dilation for this system is

$$H_{a.d.}(t) = \frac{i\gamma}{\sqrt{e^{2t\gamma}-1}}(\sigma_- \otimes \sigma_+ - \sigma_+ \otimes \sigma_-) + \tfrac{1}{2}\omega_0\,\sigma_z \otimes \mathbb{I}, \tag{III.44}$$

where once again the ancilla consists of a single qubit. It is interesting to note that, although this has a very different physical origin to the dephasing channel, the time dependency is almost identical. The comparison highlights some of the most common features of dilations of simple systems. The Hamiltonian is strongest at $t = 0$ (possibly even diverging, as discussed previously), and the terms corresponding to decay fall to 0 for large t; which is to be expected for the system to settle to its steady state.

As a more involved example, an additional constant driving term,$-i\Omega[\sigma_x,\,\rho]$, can be added on top of the amplitude damping noise. In this case the coherent and incoherent part of the evolution no longer commute, which makes a direct calculation of the exact dynamics difficult. Nevertheless, by taking the limit where the driving strength Ω is much smaller than the decay rate γ, perturbative methods can be used to find the dilation

$$H(t) = H_{a.d.}(t) + \Omega\frac{2}{1+e^{\gamma t}}\sigma_x \otimes \mathbb{I} + \Omega\frac{\sqrt{e^{2\gamma t}-1}}{\left(e^{\gamma t}+1\right)^2}\sigma_z \otimes \sigma_x, \tag{III.45}$$

where $\omega_0 = 0$ for simplicity. This dilation has two potentially unexpected features. Firstly, the driving term has gone from being constant to decaying in time. For large t it does this at same rate as $H_{a.d.}$, which, as already stated, is necessary for the system to settle to a single fixed point. Secondly, there is the appearances of a third term, which is back-action caused by dilating the control and it has a more complex structure in time, although it too decays at the same rate for large t. This term is caused by $H_{a.d.}(t)$ building up correlations between the system and the ancilla. The emergence of complicated time structure induced by a simple control field is closely linked to the fact that Master equations are changed in a non-trivial way by the addition of a Hamiltonian acting on the subsystem [DJR14]; the question of when such effects occur is the foundation of §IV.

The significantly more complex case involving the driving is time-dependent, $-i\Omega\cos(\omega t)[\sigma_x, \rho]$, can also be dilated in a perturbative method. By making the rotating-wave approximation the dilation is, in the resonant case:

$$H(t) = iH_0(t)\,\sigma_- \otimes \sigma_+ + \frac{\omega_0}{4}\sigma_z \otimes \mathbb{I} + \Omega\,f(t)\,\sigma_- \otimes \mathbb{I} + \Omega\,g(t)\,\sigma_z \otimes \sigma_x + h.c, \tag{III.46}$$

$$\text{where} \quad H_0(t) = \frac{e^{-i\omega_0 t}\gamma}{\sqrt{e^{2t\gamma}-1}}, \quad f(t) = \frac{e^{-i\omega_0 t}}{(1+e^{\gamma t})}, \quad \text{and} \quad g(t) = \frac{\sqrt{e^{2\gamma t}-1}}{\left(4\left(e^{\gamma t}+1\right)^2\right)},$$

are plotted in Fig.III.3. The increase in complexity of the Hamiltonian is directly related to the time-dependence of the original equation of motion, but the dilation could still be constructed which shows that this approach is applicable to a wide range of problems.

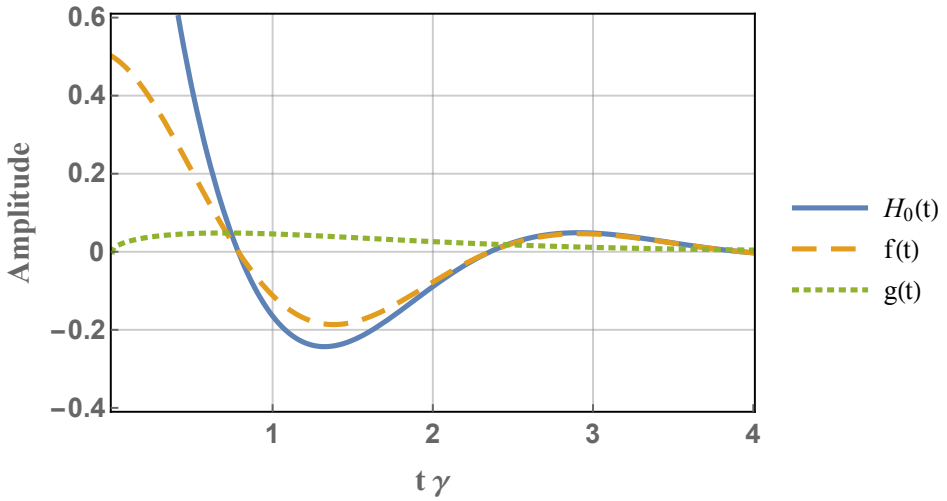


Figure III.3 | Plot of the dilation for a qubit subjected to an amplitude damping Lindbladian and a resonant sinusoidal driving field. The real part of the different components of the Hamiltonian detailed in Eq.(III.46) are shown here with $\omega_0/\gamma = 2$. Only one of the terms diverges at $t = 0$, while all the terms decay to 0 quickly at large times.

III.E Conclusion

A core method for constructing the dilation of a family of CPT maps to a time-dependent Hamiltonian acting on the system and a finite dimensional ancilla was presented. This was demonstrated on two typical examples of qubit noise, dephasing and amplitude damping, and gave simple Hamiltonians on two qubits in both the Markovian and non-Markovian case. The case of extra rotations on the system was also investigated perturbatively. These all showed that the dilation Hamiltonian was well behaved everywhere but at $t = 0$, a consequence of needing to circumvent the quantum Zeno effect. Truncating the divergent Hamiltonian to a fixed maximum value however had negligible impact on the total dynamics.

By rescaling time in a nonlinear way, some of these channels can even be dilated to constant Hamiltonians. In general, this is true whenever the dilation Hamiltonian is of the form $H(t) = h(t)X$ such that it has a single time-dependent factor. This allows a constant Hamiltonian $H = h_0 X$ to be applied for time τ to simulate the real dynamics evolving for a time t where $\tau = \frac{1}{h_0}\int_0^t h(t')dt'$. As $h(t)$ is continuous and bounded, this is always well defined. In cases where there are several different time dependencies, this method can be used to remove one of them. This is particularly useful in eliminating divergences, which would otherwise be problematic to implement experimentally. It also has the advantage that, in many cases, the evolution for an infinite amount of time t can be simulated with a finite τ.

As well as being useful in the single qubit case, these results can also be directly used in the case of a system of many qubits all subjected to independent channels, leading to a

super-polynomial speed up from what could be achieved classically. Prominent questions which could be investigated include the decay of multipartite entanglement of an initially highly entangled state [LM13], or the performance of one way computation [RB01] with a cluster state affected by local noise. For such systems of N qubits the dilation can be calculated once classically and the dynamics of the entire system then simulated on a quantum computer using a maximum of 2N ancilla qubits and Hamiltonians that affect, at most, three qubits. Simulating this classically would require applying the quantum channel up to 2^N times for an initial state which is highly entangled or, equivalently, solving the dynamics of the complete channel which would have 2^N Kraus operators.

This method also provides the ability to do something which a normal Stinespring dilation cannot do at all. By simulating the dynamics continuously in time, the evolution of the state follows the 'true' path that is being simulated in the Hilbert space, rather than just reaching the required goal for a single instance in time. This means that the full information about the behaviour of the system over an extended interval of time is accessible, allowing simulations where the time at which measurement takes place is not known a priori. In our scheme, such a scenario creates no additional difficulty, as the system follows the correct continuous dynamics. In a standard Stinespring dilation approach, however, this is either completely impossible or introduces substantial errors. These can be quantified by considering the snapshot Stinespring dilation as being a constant Hamiltonian (the logarithm of the unitary dilation) applied for different durations, leading to an error which grows in time.

In our approach, the complexity of open system dynamics is condensed in the time-dependent system-environment interaction, allowing a simulation to be implemented using state of the art methods; whereas in natural systems, and in previous approaches to simulating the dynamics of open systems, the complexity resides in the dynamics of the infinitely large environment [Gre01, ZŠB05, KBA05]. These two perspectives can be understood as the two ends of the spectrum of quantum simulations. Since any time-dependence can be understood as originating from the dynamics of an additional system (fundamentally all interactions are time-independent), the approached described in this chapter suggests very clearly how to access the entire spectrum: expand the ancilla system while gradually reducing the time-dependence of the interaction. A general procedure for doing this is the clock-construction [MSK15]. Such a continuous variation will give valuable insight, for example in the controllability of open quantum dynamics, as seen in Eqs.(III.45-III.46). There the back-action of an external control caused by the environment interaction makes itself transparent, whereas such effects are extremely hard to unravel in a model based on an infinitely large environment, as can be seen in §IV. The suggested transition would allow the study of this back-action in its entire range of manifestations, opening up a completely new angle on the investigation of control on open quantum systems.

CHAPTER IV

LINDBLADIAN AND CONTROL COMMUTATIVITY CONJECTURE

IV.A Introduction

This idea for this chapter came up during the Quantum Information Processing Conference (QIPC Leeds, 2015) where we were discussing with Rafal Demkowicz-Dobrzanski how, in control theory, Hamiltonians are often naively added on top of a Lindbladian without considering how it might modify the dissipative dynamics, as was done in §II. He suggested that it ought to be valid to take this approach if the control Hamiltonian and the Lindbladian commute, which we call the Lindbladian and Control Commutativity (LCC) conjecture. This chapter details our efforts in proving whether this is true or not.

IV.A.1 Motivation

In §II, no-go theorems for which CPT maps could be reached by an open system using coherent controls were derived, under the assumption that the addition of controls did not modify the dissipative part of the dynamics. However, in §III, it was clear that adding controls on top of the noise leads to a different dilation. This is a particular example of something which is well known (if commonly ignored): adding an extra Hamiltonian to the Schrödinger equation describing the unitary evolution of the system and the bath, and then reducing this to a Lindblad type equation of motion is not equivalent to reducing the Schrödinger equation and then adding the same Hamiltonian to the system [DJR14]. The two different ways of getting to the reduced dynamics in the presence of controls are shown diagrammatically in Fig.IV.1.

Naively adding the control Hamiltonian on top of the Lindbladian (the clockwise path in Fig.IV.1) is often the method taken, as in §II, although the opposite order would be more rigorous. The principle reason for doing this is that it is simpler both mathematically and physically. Tracing out the bath in order to obtain a Lindblad type master equation,

as was done in §I.B.4, is hard and may not be tractable if complex dynamics happen to the system [DS78, DJR14] because this may cause Markovianity to be lost leading to the reduced system not obeying a simple Master equation. Having to redo this calculation for every potential control of interest would rarely be feasible; the very reason for deriving a master equation is to be able to abstract away the bath as much as possible. Furthermore, the noise operator in the absence of control is often known (approximately) experimentally by observing how the system decays. It is harder, and not always possible, to calculate from this the exact interaction Hamiltonian between the system and the bath [Cyw14, ZÁK16]. For these reasons, it is desirable to know when the two paths in Fig.IV.1 are (approximately) equivalent. It is already known that in the singular coupling limit this is a valid approximation [DJR14], but in many cases it is not.

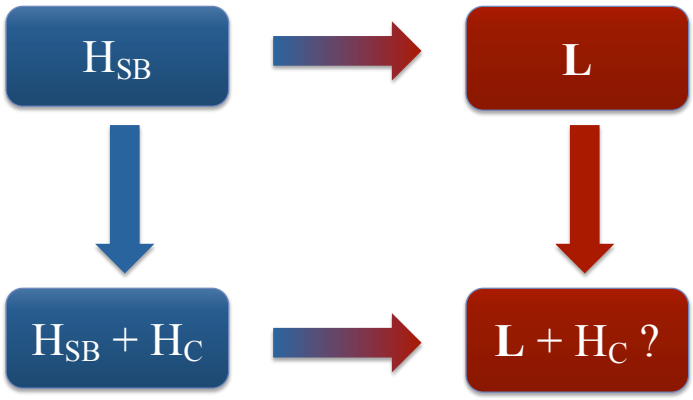


Figure IV.1 | Flowchart of the two different ways of adding controls to a Lindblad master equation. The starting point is always the exact Hamiltonian describing the system and bath. In the anticlockwise path the control Hamiltonian is added to it, and the bath is then traced out to give a reduced equation of motion on the system that incorporates the controls. It is however often easier and desirable to take the clockwise the path, where the bath is first traced out, giving a fixed Lindbladian for the system, and the control Hamiltonian then added. In general these two approaches give different answers: the order of the operations do not commute.

An interesting case when the two paths of Fig.IV.1 are not equivalent, and even maximally different, is in the Shallow-Pocket model [AHFB15, ABFH17]. In this case the use of dynamical decoupling type of controls [VKL99, LB13] in the clockwise path leads to periodic rotations on top of decoherence, while in the anti-clockwise path the controls undo the effect of the noise. The reason for this can be seen by solving the model. It consists of a qubit interacting with a bath made up of a single free particle in 1-D via the Hamiltonian $H_{\text{s-p}} = \frac{1}{2} g \sigma_z \otimes x$. If the free particle is originally in the state $|\psi\rangle$ where

$$\langle x|\psi\rangle = \sqrt{\frac{\gamma}{\pi}} \frac{1}{x + i\gamma}, \tag{IV.1}$$

the reduced dynamics on the qubit are

$$L(\rho) = -\frac{\gamma}{4} g[\sigma_z, [\sigma_z, \rho]], \tag{IV.2}$$

which corresponds exactly to the dephasing Lindbladian of §I.B.6 where the coherence of the qubit decays exponentially in time. The controls used are an instantaneous σ_x rotation on the qubit, $\exp(-i\frac{\pi}{2}\sigma_x) = \sigma_x$, applied at fixed intervals of τ. If such controls are applied in the reduced picture the final state is

$$\rho(2\tau) = \sigma_x \left(e^{L\tau}\left[\sigma_x e^{L\tau}(\rho(0))\sigma_x\right]\right)\sigma_x \tag{IV.3}$$

$$= e^{L\tau}\left[e^{L\tau}(\rho(0))\right], \tag{IV.4}$$

because

$$e^{L\tau}\left(\sigma_x \cdot \sigma_x\right) = \sigma_x e^{L\tau}\left(\cdot\right)\sigma_x, \tag{IV.5}$$

therefore the evolution of the qubit after time 2τ is the same as not applying any controls. Doing this in the microscopic picture, however, completely removes the decoherence as the same control sequence acts as [AHFB15]

$$\left(\sigma_x \otimes \mathbb{I}_B\right) e^{-iH_{\text{s-p}}\tau} \left(\sigma_x \otimes \mathbb{I}_B\right) e^{-iH_{\text{s-p}}\tau} = e^{iH_{\text{s-p}}\tau}\, e^{-iH_{\text{s-p}}\tau} \tag{IV.6}$$

$$= \mathbb{I}. \tag{IV.7}$$

This shows how drastic the difference between the two different ways of adding controls can be.

This sort of calculation, however, relies on knowing the precise interaction between the system and the bath. The aim of this chapter is to see what can be said when this is not known, but extra constraints on the allowed controls are imposed. The Lindbladian and Control Commutativity (LCC) conjecture is that if the control Hamiltonian commutes with the Lindbladian on the system at all times, then it is exactly valid to add the control on top of the noise no matter what the microscopic interaction between the system and the bath is. This conjecture was suggeted by Rafal Demkowicz-Dobrzanski, who has used it in some of his work [SKHDD16]. The intuition for why this should be true is because if the controls and noise commutes then it is possible to consider the controls acting instantaneously after the evolution under the Lindbladian. In this case, the noise is 'unaware' that there will be controls and so should remain unchanged. A rigorous proof of this however has alluded us and we have been unable to prove or disprove this conjecture either way. The heart of the problem is that, although the controls commute with the Lindbladian, they may not commute with the microscopic system-bath interaction, so it is not necessarily valid to consider the two acting separately.

The exact definition of the LCC conjecture is given below. This is followed by §IV.B

where several of our attempts to prove it result in the partial results given. In §IV.C a particularly interesting model is investigated which is numerically shown to violate the conjecture for some choice of parameters, yet is analytically proved to obey it in the Markovian limit. Finally, this chapter is concluded in §IV.D with a summary of our results, along with their implications and limitations.

IV.A.2 Problem Definition

The conjecture involves a finite dimensional system ρ_S whose evolution in the absence of controls is effectively described by a Lindbladian L. This arises from a time-independent interaction with a bath

$$\begin{aligned}\rho_S(t) &= e^{Lt}(\rho_S) && \text{(IV.8)}\\ &= \mathrm{Tr}_B\left[e^{-iH_{SB}t}(\rho_S\otimes\rho_B)e^{iH_{SB}t}\right] && \text{(IV.9)}\\ &= \mathrm{Tr}_B\left[e^{\mathcal{H}_{SB}t}(\rho_S\otimes\rho_B)\right], && \text{(IV.10)}\end{aligned}$$

where H_{SB} is the complete system-bath Hamiltonian, $\mathcal{H}_{SB}$ is its adjoint action ($\mathcal{H}(\cdot) = -i[H,\,\cdot\,]$), and ρ_B is the initial state of the bath. Although L is known, H_{SB} and ρ_B are not, and there are generally a great many different dilations that give the same reduced dynamics on ρ_S. The conjecture states that given an additional Hamiltonian H_C on the system alone, such that $[\mathcal{H}_C, L] = 0$, then, for all $\{H_{SB},\,\rho_B\}$ that satisfy Eq.(IV.10):

$$e^{(L+\mathcal{H}_C)t}(\rho_S) = \mathrm{Tr}_B\left[e^{(\mathcal{H}_{SB}+\mathcal{H}_C\otimes\mathbb{I})t}(\rho_S\otimes\rho_B)\right], \qquad \text{(IV.11)}$$

for all t and all initial ρ_S. In the case that the controls are time-dependent, the equivalent statement is

$$\mathcal{T}e^{\int_0^t L+\mathcal{H}_C(\tau)d\tau}(\rho_S) = \mathrm{Tr}_B\left[\mathcal{T}e^{\int_0^t \mathcal{H}_{SB}+\mathcal{H}_C(\tau)\otimes\mathbb{I}d\tau}(\rho_S\otimes\rho_B)\right]. \qquad \text{(IV.12)}$$

For convenience, the shorthand $\mathcal{H}_C\otimes\mathbb{I} = \mathcal{H}'_C$ is used below.

There are two observations worth making at this point. Firstly, the conjecture is true if $[H'_C, H_{SB}] = 0$. This makes intuitive physical sense, if the controls commute with the microscopic interaction between the system and the bath, then the two ignore each other. Therefore any attempts to disprove the conjecture, or find a counter example, must rely on the control not commuting with the system-bath Hamiltonian even though it commutes with the Lindbladian. However, the result can hold even in such a case; a toy model is detailed in §IV.C where the controls do not commute with the full Hamiltonian, but the LCC conjecture holds.

Secondly, for L to be a true Lindbladian it is necessary for H_{SB} to be unbounded, an observation that was discussed in §III.C.2. For this reason, the work in this chapter

involves unbounded operators in infinite dimensional Hilbert spaces so care must be taken to not mistakenly adopt results that only hold in finite dimensions [Sch12].

One of the reasons why this conjecture is hard to prove is that it must hold for any $\{ H_{SB}, \rho_B \}$ that give rise to a chosen L. It is not sufficient to find such a Hamiltonian and bath state that give the correct reduced dynamics and test whether the conjecture holds as there exist infinitely many different dilations that give rise to the same Lindbladian (see §III.B for the many degrees of freedom that exist even when picking a finite dimensional dilation). There exist classes of dilations, based on stochastic Schrödinger equations [HP84b, Par92, Gre01], such that all commuting controls can simply be added on top of the Lindbladian, but this does not prove that all dilations for that L follow this property.

IV.B Partial Results

We attempted a direct proof of the conjecture using the intuition described above; that the commutativity of noise and controls allows them to be moved in time arbitrarily. The failure of the proof lies in two technical points that are detailed below. Nevertheless, important insights into the LCC conjecture can be gleamed from the partial results given here, including a standalone result on the partial commutativity of H_{SB} and H'_C. Only the case of time-independent controls is presented in order to keep the notation simple, no significant differences arise in the time-dependent case.

The system as defined in Eq.(IV.10) is finite dimensional, therefore it follows directly from the condition $[\mathcal{H}_C, L] = 0$ that $[e^{\mathcal{H}_C t}, e^{Lt}] = 0$. This means that the unitary generated by the controls can be distributed before or after the dissipative channel as desired

$$e^{(L+\mathcal{H}_C)t}(\rho_S) = e^{(1-\lambda)\mathcal{H}_C t} e^{Lt} e^{\lambda \mathcal{H}_C t}(\rho_S) \quad \forall \ \lambda. \tag{IV.13}$$

Replacing the Lindbladian by the system-bath interaction and shuffling operators leads to

$$e^{(L+\mathcal{H}_C)t}(\rho_S) = e^{(1-\lambda)\mathcal{H}_C t} \mathrm{Tr}_B \left[e^{\mathcal{H}_{SB} t} \left(e^{\lambda \mathcal{H}_C t}(\rho_S) \otimes \rho_B \right) \right] \tag{IV.14}$$

$$= \mathrm{Tr}_B \left[e^{(1-\lambda)\mathcal{H}'_C t} e^{\mathcal{H}_{SB} t} e^{\lambda \mathcal{H}'_C t}(\rho_S \otimes \rho_B) \right]. \tag{IV.15}$$

From this relation, two separate statements about the dynamics of the system and bath can be deduced which are relevant to the conjecture.

IV.B.1 Commutativity relation inside the partial trace

An interesting result about the commutativity of H'_C and H_{SB} can be found from the expression above. First, as the lefthand side of Eq.(IV.15) does not depend on λ, the

righthand side at $\lambda = 0$ and $\lambda = 1$ can be equated with each other

$$\mathrm{Tr}_B\left[e^{\mathcal{H}'_C t}e^{\mathcal{H}_{SB}t}(\rho_S \otimes \rho_B)\right] = \mathrm{Tr}_B\left[e^{\mathcal{H}_{SB}t}e^{\mathcal{H}'_C t}(\rho_S \otimes \rho_B)\right]. \tag{IV.16}$$

Differentiating this with respect to time twice leads to

$$\mathrm{Tr}_B\left[e^{\mathcal{H}'_C t}(\mathcal{H}'_C + \mathcal{H}_{SB})e^{\mathcal{H}_{SB}t}(\rho_S \otimes \rho_B)\right] = \mathrm{Tr}_B\left[e^{\mathcal{H}_{SB}t}(\mathcal{H}_{SB} + \mathcal{H}'_C)e^{\mathcal{H}'_C t}(\rho_S \otimes \rho_B)\right] \tag{IV.17}$$

$$\mathrm{Tr}_B\left[e^{\mathcal{H}'_C t}(\mathcal{H}'^2_C + 2\mathcal{H}'_C\mathcal{H}_{SB} + \mathcal{H}^2_{SB})e^{\mathcal{H}_{SB}t}(\rho_S \otimes \rho_B)\right] = \mathrm{Tr}_B\left[e^{\mathcal{H}_{SB}t}(\mathcal{H}^2_{SB} + 2\mathcal{H}_{SB}\mathcal{H}'_C + \mathcal{H}'^2_C)e^{\mathcal{H}'_C t}(\rho_S \otimes \rho_B)\right], \tag{IV.18}$$

with an important caveat. The time derivative is only well defined if $e^{\mathcal{H}_{SB}t}\rho_S \otimes \rho_B$ is in the domain of $\mathcal{H}_{SB}$. But, for it to give rise to a Lindbladian, the initial state must not be in its domain, making this is major obstacle. It may be possible to surmount it by assuming that physical systems are not exactly Lindbladian (due to the Born-Markov or coarse-graining approximations discussed in §I.B.4 typically made). Under this framework $||H_{SB}\rho_S \otimes \rho_B||$ would be large but finite, leading to a small change in the evolution of the state. Keeping track of this with a maximum amplitude cut off, in a similar way as was done in §III, may lead to a circumvention of the domain problem here. This possibility is not investigated in depth as the direction of the main proof sketched out in §IV.B.2 suffers a more fundamental problem.

Ignoring this problem here, Eq.(IV.18) at $t = 0$ can be simplified to

$$\mathrm{Tr}_B\left[\mathcal{H}'_C\mathcal{H}_{SB}(\rho_S \otimes \rho_B)\right] = \mathrm{Tr}_B\left[\mathcal{H}_{SB}\mathcal{H}'_C(\rho_S \otimes \rho_B)\right] \tag{IV.19}$$

$$\mathrm{Tr}_B\left[\left[\mathcal{H}'_C, \mathcal{H}_{SB}\right](\rho_S \otimes \rho_B)\right] = 0. \tag{IV.20}$$

This states that the system-bath Hamiltonian commutes with the control Hamiltonian within the partial trace when acting on an appropriate state (that is, any system state and the initial bath state of the dilation). This is a much weaker condition than the two Hamiltonians commuting, but perhaps not a surprising one. The restriction on the bath state is due to it being as important a part of the dilation of L as H_{SB} is, while the partial trace arises because only the evolution of ρ_S is relevant to the conjecture. As such, Eq.(IV.20) can be seen as a lift of the condition $[\mathcal{H}_C, L] = 0$ (which acts on the system alone), to one that acts on the system-bath jointly. That this is a weaker condition leaves the possibility that the LCC conjecture is not true because, as discussed previously, if $[\mathcal{H}'_C, \mathcal{H}_{SB}] = 0$ a proof would be immediate.

IV.B.2 Generator of joint Unitary

A more wholistic result about how the unitary describing the system-bath propagation evolves in time can be obtained by returning to Eq.(IV.15). The left hand side is independent of λ, therefore differentiating with respect to it gives

$$0 = \mathrm{Tr}_B \left[-\mathcal{H}'_C t e^{(1-\lambda)\mathcal{H}'_C t} e^{\mathcal{H}_{SB} t} e^{\lambda \mathcal{H}'_C t} + e^{(1-\lambda)\mathcal{H}'_C t} e^{\mathcal{H}_{SB} t} \mathcal{H}'_C t e^{\lambda \mathcal{H}'_C t} (\rho_S \otimes \rho_B) \right]. \quad \text{(IV.21)}$$

Setting $\lambda = 1$ and rearranging yields

$$\mathrm{Tr}_B \left[\mathcal{H}'_C e^{\mathcal{H}_{SB} t} e^{\mathcal{H}'_C t} (\rho_S \otimes \rho_B) \right] = \mathrm{Tr}_B \left[e^{\mathcal{H}_{SB} t} \mathcal{H}'_C e^{\mathcal{H}'_C t} (\rho_S \otimes \rho_B) \right]. \quad \text{(IV.22)}$$

This expression shows that H'_C commutes with $e^{H_{SB} t}$ 'inside the partial trace' and when acting on the initial state (the term $e^{\mathcal{H}'_C t}$ can be absorbed into ρ_S without loss of generality). The next step is to once again return to Eq.(IV.15) but, rather than differentiating with respect to λ, setting $\lambda = 1$ and taking the time derivative as before. This gives

$$\begin{aligned}
\frac{d}{dt} e^{(L+\mathcal{H}_C)t} (\rho_S) &= \mathrm{Tr}_B \left[\tfrac{d}{dt} e^{\mathcal{H}_{SB} t} e^{\mathcal{H}'_C t} (\rho_S \otimes \rho_B) \right] && \text{(IV.23)} \\
&= \mathrm{Tr}_B \left[\mathcal{H}_{SB} e^{\mathcal{H}_{SB} t} e^{\mathcal{H}'_C t} + e^{\mathcal{H}_{SB} t} \mathcal{H}'_C e^{\mathcal{H}'_C t} (\rho_S \otimes \rho_B) \right] && \text{(IV.24)} \\
&= \mathrm{Tr}_B \left[\left(\mathcal{H}_{SB} + \mathcal{H}'_C \right) e^{\mathcal{H}_{SB} t} e^{\mathcal{H}'_C t} (\rho_S \otimes \rho_B) \right], && \text{(IV.25)}
\end{aligned}$$

where the result from Eq.(IV.22) was used in the last line. Equating the right hand side of the top and bottom equations above gives

$$\mathrm{Tr}_B \left[\tfrac{d}{dt} e^{\mathcal{H}_{SB} t} e^{\mathcal{H}'_C t} (\rho_S \otimes \rho_B) \right] = \mathrm{Tr}_B \left[\left(\mathcal{H}_{SB} + \mathcal{H}'_C \right) e^{\mathcal{H}_{SB} t} e^{\mathcal{H}'_C t} (\rho_S \otimes \rho_B) \right], \quad \text{(IV.26)}$$

where the same caveat as before about taking the time derivative applies to Eq.(IV.24).

Putting aside this difficulty, Eq.(IV.26) is a first order linear differential equation describing the dynamics of ρ_S. Were the partial trace over the bath not present and for it to hold for all $\rho_S \otimes \rho_B$ then it could be recast as

$$\frac{d}{dt} \mathcal{U}(t) = \left(\mathcal{H}_{SB} + \mathcal{H}'_C \right) \mathcal{U}(t). \quad \text{(IV.27)}$$

From Stone's theorem [Sch12] the solution would be (even for unbounded Hamiltonians)

$$\mathcal{U}(t) \equiv e^{\mathcal{H}_{SB} t} e^{\mathcal{H}'_C t} = e^{\left(\mathcal{H}_{SB} + \mathcal{H}'_C \right) t}, \quad \text{(IV.28)}$$

which would prove the LCC conjecture. However, the partial trace and fixed bath state changes the nature of the problem. Specifically, the partial trace is a non-invertible function. Therefore while it is true that $e^{\mathcal{H}_{SB} t} e^{\mathcal{H}'_C t} = e^{\left(\mathcal{H}_{SB} + \mathcal{H}'_C \right) t}$ is a valid substitution in Eq.(IV.26), this does not imply that it is a valid substitution in Eq.(IV.15) which defines

the evolution of ρ_S. In other words, our result for the generator of the joint unitary is consistent with the conjecture being true but does not prove it.

Applying Stone's theorem inside the partial trace was a problem that we could not overcome, and prevented us from proving the LCC conjecture. However it is an interesting end point in its own right. It states that, when the controls and Lindbladian commute, the unitary operator that governs the evolution of the system and bath obeys almost the same differential equation as if the underlying interaction and control Hamiltonian commuted, except that it does so 'inside the partial trace' and for a fixed initial bath state.

The same result can be found in the case that the controls are time-dependent using the same method as above; due to the significant similarities and lack of new concepts the reader is spared the repetition.

IV.C Detailed Example

As well as approaching the problem with general analytical techniques in the way done above, we also studied specific instances of the LCC conjecture in the hope of finding a counter example. This is a surprisingly difficult task. Our first attempt to do this was with the Shallow-Pocket model described in IV.A.1. This, however, is not of much interest because the only control that commutes with the Lindbladian is σ_z and this also commutes with the system-bath interaction, meaning that it necessarily obeys the conjecture. This is because the H_{SB} term in the Shallow-Pocket model is of the form $S \otimes B$ which leads to dephasing of the qubit around S [BLPV15]. To commute with the Lindbladian the controls therefore need to commute with S, which therefore implies that they also commute with H_{SB}.

A model with a more complicated structure is therefore required, for a qubit the natural candidate for this is a dilation of the amplitude damping Lindbladian discussed in §I.B.6. The Friedrichs-Lee model ([Fac16, BF17] and references in latter) does exactly this. This model is described and solved below, followed by some of our numerical calculations which imply that the LCC conjecture is broken in a specific non-Markovian coupling. Finally, our proof that the LCC conjecture is obeyed at all times in the Markovian limit of the model is presented.

IV.C.1 Friedrichs-Lee Model

The Friedrichs-Lee model consists of a one-particle bath that obeys precise exponential decay. When coupled with a qubit appropriately, this results in exactly an amplitude damping channel without any approximations being made in the derivation. This makes it a very useful tool in which to study the LCC conjecture. The description given here is heavily drawn from private communications with Paolo Facchi [Fac16].

The bath Hilbert space is $\mathbb{H}_B = \mathbb{C} \oplus L^2(\mathbb{R})$. This can be thought of as a Fermionic field (or a Bosonic field confined to a single excitation) in 1-D, and therefore not quite as unphysical as it looks. Vectors in this space are $|\psi\rangle = x\,|a\rangle \oplus \int \xi(\omega)\,|\omega\rangle\,d\omega$ which can be written as the column vector

$$|\psi\rangle = \begin{pmatrix} x \\ \xi(\omega) \end{pmatrix}. \tag{IV.29}$$

The free Hamiltonian of this model is

$$H_0 = \omega_a\,|a\rangle\langle a| + \int_{\mathbb{R}} d\omega\;\omega\,|\omega\rangle\langle\omega|\,, \tag{IV.30}$$

which can equivalently be expressed as the matrix

$$H_0 = \begin{pmatrix} \omega_a & 0 \\ 0 & \omega \end{pmatrix}. \tag{IV.31}$$

As it stands, this system does not do anything interesting. To describe the decay of an unstable vacuum into the one-particle sector, an interaction term of the form $V = \int_{\mathbb{R}} d\omega\, g(\omega)\,(|a\rangle\langle\omega| + |\omega\rangle\langle a|)$ needs to be added. This gives as the total Hamiltonian

$$H = \begin{pmatrix} \omega_a & \langle g| \\ g(\omega) & \omega \end{pmatrix}, \tag{IV.32}$$

where $|g\rangle = \int d\omega\, g(\omega)\,|\omega\rangle$. This operator is indeed a valid Hamiltonian [BF17], a non-trivial observation for unbounded operators.

This model can be solved exactly for the initial condition $|\psi\rangle = |a\rangle$. The Schrödinger equation reads

$$i\dot{x} = \omega_a x(t) + \int_{\mathbb{R}} d\omega\, g(\omega)\xi(\omega,t) \tag{IV.33}$$

$$i\dot{\xi}(\omega,t) = \omega\xi(\omega,t) + g(\omega)x(t). \tag{IV.34}$$

The former can changed to an integral equation with the substitution $x(t) = e^{-i\omega_a t}y(t)$

$$ie^{-i\omega_a t}\dot{y}(t) = \int_{\mathbb{R}} d\omega\, g(\omega)\xi(\omega,t) \tag{IV.35}$$

$$\frac{dy(t)}{dt} = -ie^{i\omega_a t}\int_{\mathbb{R}} d\omega\, g(\omega)\xi(\omega,t) \tag{IV.36}$$

$$y(t) - y(0) = -i\int_0^t ds\, e^{i\omega_a s}\int_{\mathbb{R}} d\omega\, g(\omega)\xi(\omega,s) \tag{IV.37}$$

$$x(t) = e^{-i\omega_a t} - i\int_0^t ds\, e^{-i\omega_a(t-s)}\int_{\mathbb{R}} d\omega\, g(\omega)\xi(\omega,s), \tag{IV.38}$$

using the initial condition $y(0) = 1$. The same steps follow analogously for ξ, with the

exception that the initial condition is $\xi(\omega, 0) = 0$, and gives

$$\xi(\omega, t) = -i \int_0^t e^{-i\omega(t-s)} g(\omega) x(s) ds. \tag{IV.39}$$

These equations are still coupled. To solve them it is useful to substitute one into the other which, keeping track of dummy variables, gives

$$x(t) = e^{-i\omega_a t} - \int_0^t ds e^{-i\omega_a(t-s)} \int_{\mathbb{R}} d\omega\, g(\omega) \int_0^s d\tau\, e^{-i\omega(s-\tau)} g(\omega) x(\tau) \tag{IV.40}$$

$$= e^{-i\omega_a t} - e^{-i\omega_a t} \int_0^t ds \int_{\mathbb{R}} d\omega\, e^{-i(\omega-\omega_a)s} \int_0^s d\tau\, e^{i\omega\tau} x(\tau) g^2(\omega). \tag{IV.41}$$

Using Fubini's theorem [DiB02] to switch the order of integration, the integral in s can be done first as x is now a function of τ. Care must be taken however, as s appears both in the exponential and as the limit of the last integral. Treating them as two different functions of s and integrating by parts carefully gives

$$\begin{aligned} x(t) =& e^{-i\omega_a t} - e^{-i\omega_a t} \int_{\mathbb{R}} d\omega \Bigg\{ \\ & \left[\frac{i}{\omega - \omega_a} e^{-i(\omega-\omega_a)s} \int_0^s d\tau\, e^{i\omega\tau} x(\tau) \right]_{s=0}^{s=t} - \int_0^t ds \frac{i}{\omega - \omega_a} e^{-i(\omega-\omega_a)s} e^{i\omega s} x(s) \Bigg\} g^2(\omega) \end{aligned} \tag{IV.42}$$

$$x(t) = e^{-i\omega_a t} - e^{-i\omega_a t} \int_{\mathbb{R}} d\omega \frac{i}{\omega - \omega_a} \int_0^t d\tau \left[e^{-i(\omega-\omega_a)t} e^{i\omega\tau} - e^{i\omega_a\tau} \right] x(\tau) g^2(\omega), \tag{IV.43}$$

where in getting to the last line the dummy variable in the last integral was changed from s to τ and like terms collected. From here the equation for x can be expressed in terms of a kernel

$$x(t) = e^{-i\omega_a t} - i \int_0^t e^{-i\omega_a(t-\tau)} \int_{\mathbb{R}} \frac{e^{-i(\omega-\omega_a)(t-\tau)} - 1}{\omega - \omega_a} g^2(\omega) d\omega\, x(\tau) d\tau \tag{IV.44}$$

$$x(t) = e^{-i\omega_a t} - \int_0^t e^{-i\omega_a(t-\tau)} G(t-\tau) x(\tau) d\tau, \tag{IV.45}$$

where the kernel is

$$G(t-\tau) = i \int_{\mathbb{R}} \frac{e^{-i(\omega-\omega_a)(t-\tau)} - 1}{\omega - \omega_a} g^2(\omega) d\omega \tag{IV.46}$$

$$= (t-\tau) \int_{\mathbb{R}} e^{-i(\frac{\omega-\omega_a}{2})(t-\tau)} \mathrm{sinc}\left(\frac{\omega - \omega_a}{2} (t-\tau) \right) g^2(\omega) d\omega. \tag{IV.47}$$

Although seemingly more complicated than Eq.(IV.33), this is a decoupled equation for $x(t)$ that does not depend on ξ at all. In particular the time-dependency of $G(t)$ determines the behaviour of $x(t)$. In order to get exponential decay from $t = 0$ rather than an initial period of quadratic (Zeno) decay, it is necessary to have a constant kernel [Fac16], which in this case is achieved by having a constant coupling function $g^2(\omega) = \frac{\gamma}{2\pi}$. Substituting

this expression into Eq.(IV.47) with the identity $\int e^{ix}\mathrm{sinc}(x) = \frac{\pi}{2}$ reduces the kernel to

$$G(t-\tau) = \frac{\gamma}{2}\,\mathrm{sign}(t-\tau). \tag{IV.48}$$

Plugging this result back into Eq.(IV.45) allows the evolution in time of x to be simply solved with an exponential ansatz

$$x(t) = e^{-i\omega_a t} - \mathrm{sign}(t)\,\frac{\gamma}{2}\int_0^t e^{-i\omega_a(t-\tau)}x(\tau)d\tau \tag{IV.49}$$

$$x(t) = \exp\Big(-\frac{\gamma}{2}t - i\omega_a t\Big), \tag{IV.50}$$

(dropping the constraint on the sign of t as only positive times are of interest) which is the desired exponential decay. Thus, the Friedrichs-Lee model leads to the vacuum state of the field decaying exponentially onto a one-particle state.

In order to map this exact exponential behaviour onto a qubit it is sufficient to expand the Hilbert space and associate the vacuum field state to the upper qubit state, and the one-particle field state to the lower level. That is, the Hilbert space becomes $\mathbb{H} = \mathbb{C}^2 \otimes \mathbb{H}_B$ with the mapping $|a\rangle \to |1\rangle \otimes |a\rangle$ and $|\omega\rangle \to |0\rangle \otimes |\omega\rangle$. A general state can thus be written as

$$|\psi\rangle = x_1\,|1\rangle\,|a\rangle + \int d\omega\,\xi_1(\omega)\,|1\rangle\,|\omega\rangle + x_0\,|0\rangle\,|a\rangle + \int d\omega\,\xi_0(\omega)\,|0\rangle\,|\omega\rangle\,, \tag{IV.51}$$

and the Hamiltonian on the joint system is

$$H_{SB} = \begin{pmatrix} \omega_a & 0 & 0 & \langle g| \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ g(\omega) & 0 & 0 & \omega \end{pmatrix}, \tag{IV.52}$$

where $|g\rangle = \int d\omega\, g(\omega)\,|\omega\rangle$ as before.

Due to the block structure of the Hamiltonian, the solution derived previously maps directly to it. The initial states of interests are ones where the qubit is in an arbitrary superposition and the bath is, as before, in the vacuum state

$$|\psi\rangle = \begin{pmatrix} \alpha \\ 0 \\ \beta \\ 0 \end{pmatrix}, \tag{IV.53}$$

which evolves to

$$|\psi(t)\rangle = \begin{pmatrix} \alpha \exp\left(-\frac{\gamma}{2}t - i\omega_a t\right) \\ 0 \\ \beta \\ \xi_0(\omega, t) \end{pmatrix}, \tag{IV.54}$$

where the last term has not been calculated as it is not of interest. To check that the dynamics of the reduced state are indeed those of an amplitude damping channel, the partial trace needs to be calculated. This can be done as per normal, noting that all the bath modes need to be integrated over. This gives as the reduced state of the qubit

$$\rho(t) = \begin{pmatrix} |\alpha|^2 e^{-\gamma t} & \beta\alpha^* e^{-\left(\frac{\gamma}{2}t + i\omega_a t\right)} \\ \beta^*\alpha e^{-\left(\frac{\gamma}{2}t - i\omega_a t\right)} & |\beta|^2 + |\alpha|^2(1 - e^{-\gamma t}) \end{pmatrix}, \tag{IV.55}$$

using the normalisation condition to eliminate ξ_0. This is precisely the action of an amplitude damping channel on a qubit (as described in §I.B.6), therefore the Friedrichs-Lee model is doing exactly what is desired. It is important to stress that no approximations, Born-Markov or otherwise, were made in this derivation. While the physicality of a flat coupling constant may be questionable, the lack of approximations makes this a very useful model.

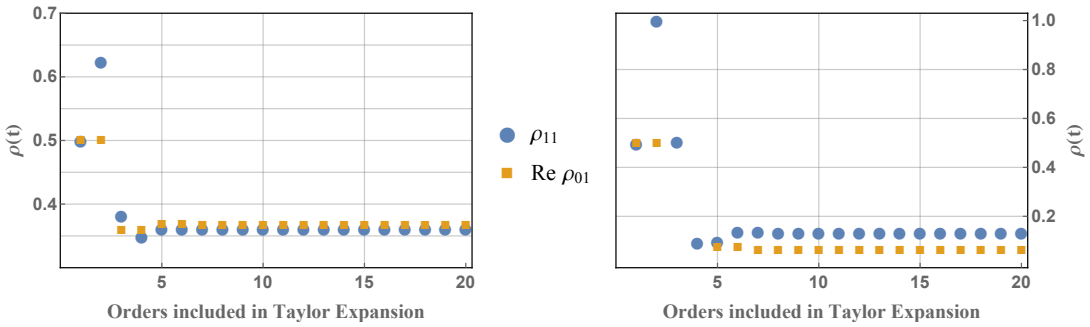
IV.C.2 Numerical Violations for non-flat coupling

The Friedrichs-Lee model gives rise to an amplitude damping channel, as shown above, which can be described by the Lindbladian

$$L(\rho) = -\gamma(\{\sigma_+\sigma_-, \rho\} - 2\sigma_-\rho\sigma_+) - i\frac{\omega_a}{2}[\sigma_z, \rho], \tag{IV.56}$$

as introduced in §I.B.6. This Lindbladian describes a qubit decaying to the state $|0\rangle\langle 0|$ while having an internal energy splitting of ω_a. It is easily shown that this Lindbladian commutes with $H_C = \Omega\,\sigma_z$. However, $H'_C = H_C \otimes \mathbb{I}_B$ does not commute with the Friedrichs-Lee Hamiltonian of Eq.(IV.52). That makes this model a particularly interesting one to test the LCC conjecture.

Specifically the question answered numerically is if Eq.(IV.11) holds with L defined as in Eq.(IV.56), H_{SB} in Eq.(IV.52) with the initial bath state $|\psi_B\rangle = |a\rangle$, and $H_C = \Omega\,\sigma_z$. This model involves an unbounded Hamiltonian which couples the qubit to all frequencies equally, therefore a numerical simulation of it is very difficult and unstable. To circumvent this, the coupling function $g(\omega)$ is replaced from a constant to a Gaussian, $g(\omega) = e^{-(\frac{\omega}{\omega_a})^2}$. The Hamiltonian is still unbounded due to the one-particle energy ω but, as discussed below, is well enough behaved to be approximately simulated numerically. This also enables the integral over ω present in the Hamiltonian, and used when taking the partial trace, from being over the interval $[+R, -R]$ rather than over all frequencies. Due to the



(a) Free Evolution **(b)** Evolution with Rotation

Figure IV.2 | Plot showing the approximate evolution obtained by Taylor expanding the unitary operator. The $|1\rangle\langle 1|$ population and the real part of the coherence is plotted for the state evolved up to $t = \frac{1}{2}\omega_a$ using Eq.(IV.57) with the Taylor expansion truncated to different orders. The series converges rapidly both for the free evolution (left) and with an additional σ_z rotation (right). Based on this, the rest of the calculations were done with a Taylor expansion cut off at the 10^{th} order.

rapidly decaying coupling for large ω, $R = 100\,\omega_a$ is more than large enough to get an excellent approximation for the purpose of this work.

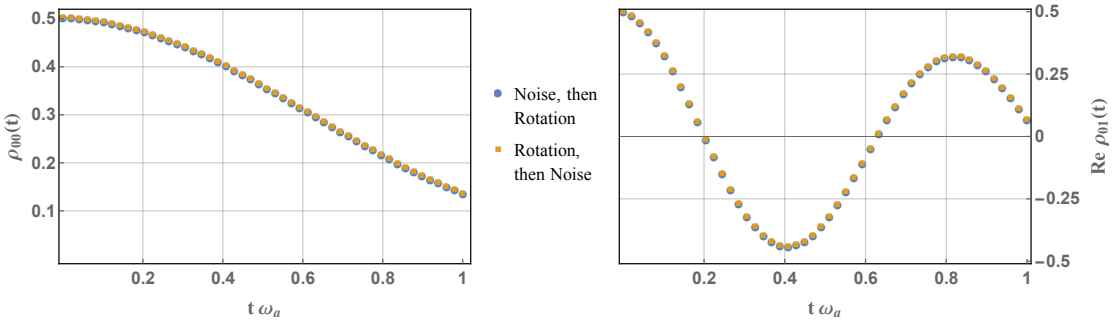
The state at some time t is found by expanding the unitary according to

$$|\psi(t)\rangle = U(t)\,|\psi(0)\rangle = \sum_k (-it)^k H_{SB}^k\,|\psi(0)\rangle\,. \tag{IV.57}$$

The state $H_{SB}^k\,|\psi(0)\rangle$ is calculated here by acting on $|\psi(0)\rangle$ with H_{SB} k times, rather than calculating H_{SB}^k and applying it to $|\psi(0)\rangle$ as this simplified the numerics. This could be done because the initial state, $|\psi(0)\rangle = (\alpha, 0, \beta, 0)^T$, has a finite norm when acted on repeatedly by H_{SB}; this is not the case if a flat coupling constant is used. For specificity, in our numerical simulations the qubit is initialised in the $|+\rangle$ state corresponding to $\alpha = \beta = \frac{1}{\sqrt{2}}$, and the control strength is $\Omega = \pi\,\omega_a$.

The first question to address is how many orders are needed in the Taylor expansion in Eq.(IV.57) for the solution to converge. This is explored in Fig.IV.2a where 10 is seen as amply sufficient. A second problem is that, as the coupling has been changed, the model is no longer exactly Friedrichs-Lee and so does not give rise to exactly amplitude damping. Furthermore, as was shown in §IV.C.1, getting an closed expression for the evolution of the qubit in time for a non-flat coupling $g(\omega)$ is a formidable task even without the presence of controls, and one that we did not attempt. For this modified model to be a test of the LCC conjecture, however, it is necessary for the reduced dynamics to commute with the controls. In the absence of a well defined generator for the reduced system, this is best described on the level of propagators

$$e^{-iH_C t}\mathrm{Tr}_B\left[U(t)\,|\psi(0)\rangle\langle\psi(0)|\,U^\dagger(t)\right]e^{iH_C t} = \mathrm{Tr}_B\left[U(t)e^{-iH_C' t}\,|\psi(0)\rangle\langle\psi(0)|\,e^{iH_C' t}U^\dagger(t)\right], \tag{IV.58}$$



(a) Population Evolution **(b)** Coherence Evolution

Figure IV.3 | Plot of a qubit evolving showing the commutativity of the modified Friedrichs-Lee model with a σ_z rotation. The equality in Eq.(IV.58) is tested numerically for the Friedrichs-Lee model with Gaussian coupling. Each plot shows how one element of the density matrix of the qubit state evolves in time, in one case by acting first with H'_C and then with the joint system-bath interaction, and in the other case by commuting the order of the operations. The points overlap so completely that they are distinguishable, from which it can be inferred that the noisy operation and σ_z rotation on the qubit approximately commute for this model.

where

$$U(t) = \sum_{k=0}^{10} (-itH_{SB})^k. \tag{IV.59}$$

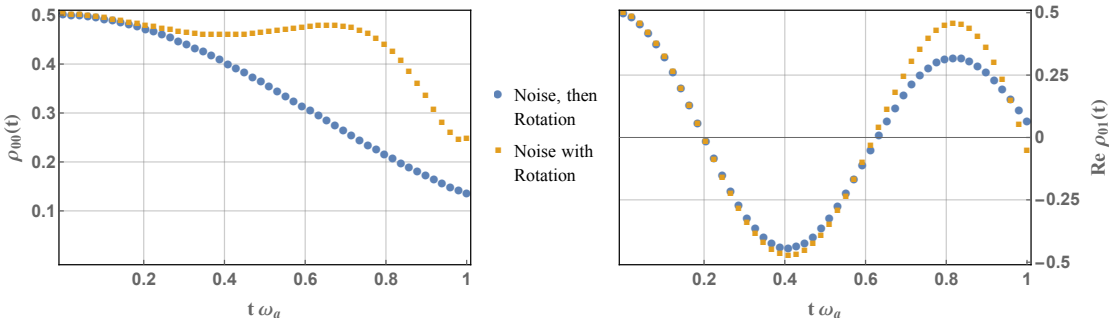
We tested this numerically and the agreement is excellent, as shown in Fig.IV.3. In addition, the plot shows that the modified Friedrichs-Lee model gives qualitatively a similar evolution as a qubit undergoing amplitude damping: the population of the excited state decays without sign of recurrence (for the evolution time considered), and the decoherence also decays while oscillating at a frequency of $\sim \omega_a + 2\Omega$. The main difference that can be observed is in the behaviour at short t: in the modified model the decay is quadratic at short times while true amplitude damping would be exponential. This is to be expected as H_{SB} has been modified precisely such that $||H_{SB}\,|\psi(0)\rangle\,||$ is bounded, which forces the Zeno limit to hold.

The evidence thus strongly suggests that this is a model which is numerically solvable on which to test the LCC conjecture. If the conjecture is true, the equality of Eq.(IV.11) holds which, in this context, can be expressed as

$$e^{-iH_C t}\mathrm{Tr}_B\left[U(t)\,|\psi(0)\rangle\,\langle\psi(0)|\,U^\dagger(t)\right]e^{iH_C t} = \mathrm{Tr}_B\left[U_C(t)\,|\psi(0)\rangle\,\langle\psi(0)|\,U^\dagger(t)\right], \tag{IV.60}$$

where $U(t)$ is as defined in Eq.(IV.59) and $U_C(t)$ is the equivalent expression when the controls are applied at the same time as system-bath interaction,

$$U_C(t) = \sum_{k=0}^{10} \left(-it(H_{SB} + H'_C)\right)^k. \tag{IV.61}$$



(a) Population Evolution **(b)** Coherence Evolution

Figure IV.4 | Plot of the evolution of the qubit with controls added simultaneously with the Gaussian coupled Friedrichs-Lee model, compared to the controls acting afterward. This shows a violation of the LCC conjecture. Each side of Eq.(IV.60) is shown in each plot, with the former showing the population of the excited state and the latter the real part of the coherence. The population shows drastically different behaviour as it is held close to constant for a large amount of time when the controls are applied at the same time as the interaction with the bath. The coherence has a very similar behaviour as before, the only change being less decay in the oscillations.

We also tested this relation numerically and a major departure from the LCC conjecture is seen in Fig.IV.4. The large deviation between the two data sets implies a rejection of the hypothesis that Eq.(IV.60) is true for this specific model. This is especially surprising as H_C is as simple as it could be: a constant rotation. In order to see if this was due to a numerical error, we checked whether enough orders were taken in the Taylor expansion of $U_C(t)$. Fig.IV.2b shows that 10, the same number as was used in the absence of controls, is amply sufficient. The results were also stable under an increase of R, the range of integration of ω.

This implies that, in this specific model, the LCC conjecture is violated. Before refuting the conjecture completely, however, it is important to note that it was defined in terms of a Lindbladian on a reduced system, but the Friedrichs-Lee model with Gaussian coupling does not give rise to a Lindbladian even though the evolution of the reduced system is well defined and CPT.

IV.C.3 Flat coupling obeys the LCC Conjecture

The numerical result shown above shows that, when the coupling between the qubit and the bath is Gaussian, the LCC conjecture is not obeyed when the control is a σ_z rotation of constant strength in time. That is, even though the total control commutes with the evolution of the reduced system, having the controls present at the same time as the system-bath interaction leads to very different dynamics on the reduced system than having them act one after the other. As unexpected as this observation might be, a more surprising result is that this cannot hold in the case of a flat coupling. This can be seen

by returning to the exact dynamics of the Friedrichs-Lee model solved in §IV.C.1.

To recap, the Hamiltonian was Eq.(IV.52)

$$H_{SB} = \begin{pmatrix} \omega_a & 0 & 0 & \langle g| \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ g(\omega) & 0 & 0 & \omega \end{pmatrix}.$$

The initial qubit-bath state, $|\psi(0)\rangle = (\alpha, 0, \beta, 0)^T$, evolved under this Hamiltonian gives the reduced qubit state at a later time Eq.(IV.55)

$$\begin{aligned} \rho(t) &= \mathrm{Tr}_B \left[e^{\mathcal{H}_{SB} t} \rho_S \otimes \rho_B \right] \\ &= \begin{pmatrix} |\alpha|^2 e^{-\gamma t} & \beta\alpha^* e^{-\left(\frac{\gamma}{2}t + i\omega_a t\right)} \\ \beta^*\alpha e^{-\left(\frac{\gamma}{2}t - i\omega_a t\right)} & |\beta|^2 + |\alpha|^2(1 - e^{-\gamma t}) \end{pmatrix}, \end{aligned}$$

The control being considered is $H_C = \Omega\,\sigma_z$ which, acting on the combined system and ignoring global phases, is

$$H_C' = \begin{pmatrix} 2\Omega & 0 & 0 & 0 \\ 0 & 2\Omega & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{IV.62}$$

The second entry of $|\psi\rangle$ is initially 0 and never couples to any other state under H_{SB} or H_C' and can therefore be ignored. As such, the joint Hamiltonian can be treated as

$$H_{SB} + H_C' = \begin{pmatrix} \omega_a + 2\Omega & 0 & 0 & \langle g| \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ g(\omega) & 0 & 0 & \omega \end{pmatrix}. \tag{IV.63}$$

This is exactly the same evolution as in the absence of controls, except that $\omega_a \to \omega_a + 2\Omega$. As such, the evolution of the reduced state under this Hamiltonian is simply

$$\rho_C(t) = \mathrm{Tr}_B \left[e^{(\mathcal{H}_{SB} + \mathcal{H}_C')t} \rho_S \otimes \rho_B \right] \tag{IV.64}$$

$$= \begin{pmatrix} |\alpha|^2 e^{-\gamma t} & \beta\alpha^* e^{-\left(\frac{\gamma}{2}t + i(\omega_a + 2\Omega)t\right)} \\ \beta^*\alpha e^{-\left(\frac{\gamma}{2}t - i(\omega_a + 2\Omega)t\right)} & |\beta|^2 + |\alpha|^2(1 - e^{-\gamma t}) \end{pmatrix}. \tag{IV.65}$$

It is easy to see that this can be rewritten as

$$\rho_C(t) = \begin{pmatrix} e^{-i\Omega t} & 0 \\ 0 & e^{i\Omega t} \end{pmatrix} \rho(t) \begin{pmatrix} e^{i\Omega t} & 0 \\ 0 & e^{-i\Omega t} \end{pmatrix} \tag{IV.66}$$

$$= e^{-iH_C t}\rho(t)e^{iH_C t} \tag{IV.67}$$

$$= e^{\mathcal{H}_C t}e^{Lt}\rho(0), \tag{IV.68}$$

which states that acting with the controls after the joint system-bath evolution gives the same final state as applying both together. Using $[L, \mathcal{H}_C] = 0$ implies that the defining relation of the LCC conjecture, Eq.(IV.11)

$$e^{(L+\mathcal{H}_C)t}(\rho_S) = \mathrm{Tr}_B\left[e^{(\mathcal{H}_{SB}+\mathcal{H}_C\otimes\mathbb{I})t}(\rho_S \otimes \rho_B)\right],$$

is obeyed here. This is in complete contrast to the numerical results above.

This is also true for more general time-dependent controls $H_C = \Omega(t)\,\sigma_z$ provided that the time-dependency is integrable. The method used is similar as in §IV.C.1, the principle difference is that $g(\omega)$ is taken as constant from the start which removes the need to calculate the general kernel. This greatly simplifies the calculation in this more difficult case of a time-dependent Hamiltonian.

The differential equations that govern the evolution of $|\psi(t)\rangle = (x_1, \xi_1, x_0, \xi_0)^T$, similarly to §IV.C.1, are

$$i\dot{x_1}(t) = \omega_a(t)x_1(t) + \int_{\mathbb{R}} d\omega\, g\xi_0(\omega, t) \qquad x_1(0) = \alpha \tag{IV.69}$$

$$i\dot{\xi_1}(\omega, t) = 2\Omega(t)\xi_1(\omega, t) \qquad \xi_1(\omega, 0) = 0 \tag{IV.70}$$

$$i\dot{x_0}(t) = 0 \qquad x_0(0) = \beta \tag{IV.71}$$

$$i\dot{\xi_0}(\omega, t) = \omega\xi_0(\omega, t) + gx_1(t) \qquad \xi_0(\omega, 0) = 0, \tag{IV.72}$$

where the shorthand $\omega_a(t) = \omega_a + 2\Omega(t)$ has been used, and only the case $g(\omega) = \sqrt{\frac{\gamma}{2\pi}} = g$ is being considered. As before, two of these equations are solved trivially

$$\xi_1(\omega, t) = 0 \tag{IV.73}$$

$$x_0(t) = \beta. \tag{IV.74}$$

The others can be turned into integral equations via a similar substitution

$$x_1(t) = e^{-i\int_0^t \omega_a(\tau)d\tau}y(t) \equiv e^{-i\theta(t,0)}y(t) \equiv e^{-i\theta(t)}y(t), \tag{IV.75}$$

as was used previously. This gives

$$ie^{-i\theta(t)}\dot{y}(t) = \int_{\mathbb{R}} d\omega\, g\xi_0(\omega, t) \tag{IV.76}$$

$$y(t) - y(0) = -ie^{-i\theta(t)} \int_{\mathbb{R}} d\omega\, g\xi_0(\omega, s)ds \tag{IV.77}$$

$$x_1(t) = \alpha e^{-i\theta(t)} - i\int_0^t ds\, e^{-i\theta(t,s)} \int_{\mathbb{R}} d\omega\, g\xi_0(\omega, s). \tag{IV.78}$$

The same steps follow analogously for ξ_0, with a different initial condition, and result in

$$\xi_0(\omega, t) = -i\int_0^t ds\, e^{-i\omega(t-s)} g x_1(s). \tag{IV.79}$$

Solving for $x_1(t)$ requires the substitution of Eq.(IV.79) into Eq.(IV.78) to obtain

$$x_1(t) = \alpha e^{-i\theta(t)} - \frac{\gamma}{2\pi}\int_0^t ds e^{-i\theta(t,s)} \int_{\mathbb{R}} d\omega \int_0^s d\tau\, e^{-i\omega(s-\tau)} x_1(\tau). \tag{IV.80}$$

Switching the order of integration, in an analogous way to Eq.(IV.42), gives

$$x_1(t) = \alpha e^{-i\theta(t)} - \frac{\gamma}{2\pi}\int_0^t ds e^{-i\theta(t,s)} \int_0^s d\tau\, x_1(\tau) \int_{\mathbb{R}} d\omega\, e^{-i\omega(s-\tau)} \tag{IV.81}$$

$$= \alpha e^{-i\theta(t)} - \gamma\int_0^t ds e^{-i\theta(t,s)} \int_0^s d\tau\, x_1(\tau)\delta(s-\tau) \tag{IV.82}$$

$$= \alpha e^{-i\theta(t)} - \tfrac{1}{2}\gamma\int_0^t ds e^{-i\theta(t,s)} x_1(s), \tag{IV.83}$$

where the factor of $\frac{1}{2}$ comes from the Dirac delta function being on the edge of the integral. While a delta function on the edge of integration is usually undefined, taking it as $\frac{1}{2}$ here is reasonable as it recovers the same result as was found in the case without controls in Eq.(IV.49). The last equation is solved by

$$x_1(t) = \alpha e^{-\frac{1}{2}\gamma t} e^{-i\theta(t)}, \tag{IV.84}$$

as can be verified by direct substitution. The reduced state can be found as before

$$\rho_C(t) = \begin{pmatrix} |\alpha|^2 e^{-\gamma t} & \alpha\beta^* e^{-\frac{1}{2}\gamma t - i\left[\omega_a t + 2\int_0^t \Omega(s)ds\right]} \\ \alpha^*\beta e^{-\frac{1}{2}\gamma t + i\left[\omega_a t + 2\int_0^t \Omega(s)ds\right]} & |\beta|^2 + |\alpha|^2(1 - e^{-\gamma t}) \end{pmatrix}. \tag{IV.85}$$

This is exactly the same expression as Eq.(IV.65), except that $\Omega t \to \int \Omega dt$. As such, this state can also be expressed as the evolution without controls followed by the controls on their own. This is a direct verification that, for a flat coupling $g(\omega)$ Eq.(IV.11)

$$\mathcal{T} e^{\int_0^t L + \mathcal{H}_C(\tau)d\tau}(\rho_S) = \mathrm{Tr}_B\left[\mathcal{T} e^{\int_0^t \mathcal{H}_{SB} + \mathcal{H}_C(\tau)\otimes\mathbb{I}d\tau}(\rho_S \otimes \rho_B)\right],$$

holds and so the LCC conjecture is obeyed by the Friedrichs-Lee model.

IV.D Conclusion

The picture on the LCC conjecture is a confused and mixed one. We were not able to prove that commutativity of a Lindbladian and a control Hamiltonian is a sufficient condition for the addition of the control Hamiltonian on the system-bath level to leave the Lindbladian unchanged. In other words, we could not determine whether, in general, the two paths of Fig.IV.1 are equivalent even when $[L, \mathcal{H}_C] = 0$.

A number of related results, however, were shown. The commutativity of the Lindbladian and the control does not imply the commutativity of the control with the system-bath Hamiltonian in general, but they do commute 'under the partial trace' Eq.(IV.20). Furthermore, the unitary describing the evolution of the joint system with the presence of controls obeys almost the same differential equation as if the Hamiltonians could be added naively, except that it does so, once again, 'under the partial trace' Eq.(IV.26).

A specific model, Friedrichs-Lee, for the system-bath interaction was also considered and gave almost paradoxical results. When the coupling between the system and the bath is constant for all bath modes, the evolution of the system alone is exactly an amplitude damping Lindbladian. This Lindbladian commutes with σ_z rotations and, although such controls do not commute with the joint system-bath Hamiltonian, the LCC conjecture is obeyed. On the other hand, we found a numerical example where this is explicitly *not* the case when the coupling is a Gaussian rather than flat. One way of understanding these seemingly contradictory result is to return to the point made in §I.B.6 that noise which leads to decay in the excited population arises from the coupling strength between the system and the bath at the resonant frequency of the system. In the example here the controls act to change that resonant frequency, which causes a large difference in $g(\omega_a)$ for the Gaussian coupling case (the strength of the control is comparable to the width of the Gaussian), but has no impact when the coupling is flat.

The conclusion of this is that, if the LCC conjecture is indeed true, then it can only hold for Markovian systems. Any proof of it must therefore make use of this property. The directions we have taken towards an analytic proof so far are unlikely to work well in this direction as the time derivative of $\exp(-iH_{SB}t)$ at $t = 0$ was used, yet this is not a well defined quantity for Markovian systems due to the unboundedness of operators. We therefore expect that, as interesting as the current results presented here are, a radically different approach will be needed to definitely settle the LCC conjecture.

CHAPTER V

IN SITU UPGRADE OF QUANTUM SIMULATORS TO COMPUTERS

V.A Introduction

Finding control pulses with optimal control invariably requires making a numerical model of the system to be controlled, and then solving this model. Neither of these steps are practical for quantum systems large enough to do something that cannot be done classically. This chapter presents an alternative way of finding control pulses that implement gates by using the quantum system to be controlled to simulate itself. These results are currently under consideration by a journal; a draft of that paper can be found at [DPMB17].

V.A.1 Overview

If quantum control is to be useful in the quest to build a quantum computer then it needs to go beyond determining whether something can be done, as in the previous three chapters, and provide a method for how to implement desired tasks. The problem at hand is optimal control: finding specific control pulses that make a quantum system do requisite operations. The conventional way of doing this is to create an accurate model of the quantum system to optimise, and to simulate this model on a classical computer. This simulation will take an initial set of control parameters and optimise it via a classical algorithm, in the way discussed in §I.C.3.

The problem with this approach is that both making an accurate model of a quantum system and simulating that model classically are very difficult to do. Unless the system can be broken down into a collection of non-interaction parts, the experimental and computational costs rise exponentially in the number of qubits. This makes such an optimal control approach unfeasible for a many-body system. This chapter presents a method we have devised to get around these difficulties. The key point is that it does not require a model of the quantum system and for the action of potential controls to be simulated

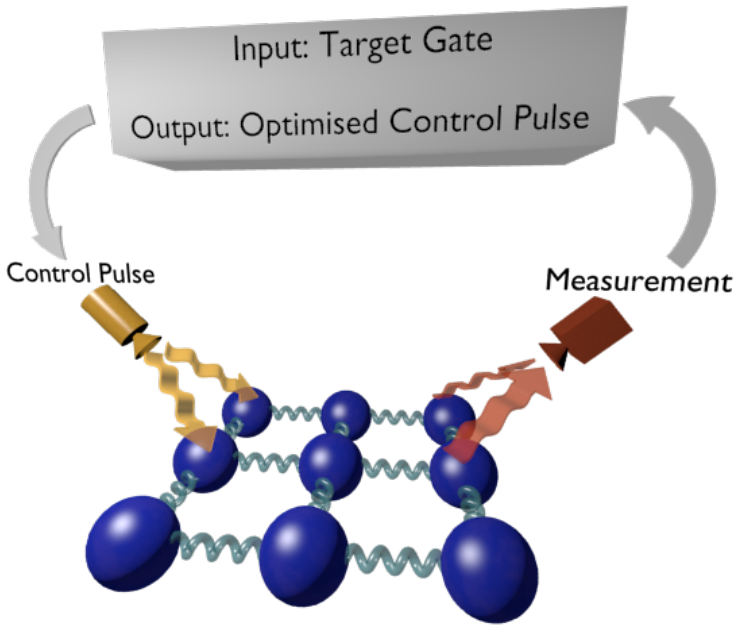


Figure V.1 | Cartoon of *in situ* control. A classical computer finds a control pulse which enables a quantum simulator to perform logic gates. It does this in an iterative process by applying a control pulse to the simulator and then improving it based on the result of measurements.

on that model, instead those controls are directly implemented on the system itself, as illustrated in Fig.V.1. The difficulty this raises is how to determine whether such a pulse drove the system in the desired way or not without making an exponential number of measurements. A way of doing this, along with investigations into the other experimental costs of such an *in situ* optimisation, are covered in this chapter.

The specific problem considered here is how to turn a many-body system into a universal quantum computer of the standard gate circuit type [NC00]. This requires the ability to perform a universal set of gates. Finding the control sequences that realise such gates is what *in situ* optimisation will be used for. In order for this to be possible, however, there are some requirements that the system must fulfil. It must be complex enough, in terms of its internal Hamiltonian and the control Hamiltonians that can be applied to it, for the system to be fully controllable in the sense of §I.C.2. There must also be enough control already known about the system to be able to perform some type of tomography such that information about how close an applied pulse has driven the system to the desired gate can be extracted. For these reasons, the scheme detailed below considers specifically quantum simulators and how they can be upgraded to quantum computers. For numerical simulations of up to nice qubits with Ising interactions, the *in situ* approach scales polynomially in the number of qubits in the simulator for different topologies, suggesting that this is a practical way of upgrading quantum simulators to computers.

Some background on quantum simulators is provided below in §V.A.2, and on prior

work on optimising controls using a physical quantum system in §V.A.3. The details of the proposed scheme are given in §V.B, the local fidelity that underpins the scalability is defined in §V.C, and is followed by the results of numerical simulations investigating other aspects of the scalability in §V.D. Finally the implication of the results and potential ways forward are discussed in §V.E

V.A.2 Quantum Simulators

Quantum simulators, machines that can replicate the dynamics of quantum systems, are seen as a stepping stone to universal quantum computers and are being built as useful machines in their own right. Recent and ongoing work with engineering large quantum systems is leading to such machines being realised at an ever larger size. They are able to model physical phenomena, allowing questions about the underlying science to be answered [CZ12, JCJ14, GAN14]. These machines are a register of quantum particles each storing some quantum information, typically a single bit as a qubit. The presence of interactions between these leads to dynamics that, by varying control parameters in the Hamiltonian or by applying external Hamiltonians, can replicate the quantum behaviour of systems of interest. This, however, is less general than a quantum computer which is able is to perform a universal set of logic gates on the qubits and some form of error correction [DiV00].

It is important to note that by simulator here an analog simulator is meant rather than a digital one, the latter being a full computer being used to run a program that simulates a system. The lesser requirements of the former compared to a computer means that it is highly likely that they will demonstrate quantum supremacy and will lead to useful products first [CZ12]. Nevertheless, they have a much more limited applicability and are of less general use than a quantum computer. The ability to turn them into such, therefore, would be of huge value as it would help bridge the divide between current or near future technology and the holy grail of a universal computer. The state-of-art of simulators is rapidly evolving and spans the full spectra of different technologies being researched as platforms for quantum technology; some of the simulators that follow the criteria the *in situ* scheme require are listed below.

Trapped ions is a well established avenues for quantum simulators [JVW09, LHNMG11, BR12], where the qubit is typically two hyperfine-levels of the ground state of an ion. Cold atoms [BDN12, LBRdLM16] has similar internal workings; the qubit there is two energy levels of the atom, either both close to the ground state or one in a highly excited Rydberg state. A much older technology is NMR [PS10, CRJP13, SSBCM16], where the spin state of the nucleus of atoms is used as qubits. This can be done either with molecules in a liquid or by defects in crystals, such as NV centres in diamond. A more recent but rapidly progressing platform is superconducting circuits [HTK12, OBKRM16] which are favoured by the tech giants. A specific type of simulator are Ising machines

[LHNMG11, LBRdLM16, MMHHL16], which only simulate different variants of the Ising model. Instead of finding the evolution in time of systems, these tend to focus on calculating the ground state of the system, a problem closely related to many classical optimisation tasks.

V.A.3 Prior Work

The core principle of this chapter is to use the system at hand, such as one of the quantum simulators discussed above, to implement and improve control pulses until they perform the desired gate. Such an adaptive approach to finding controls is naturally used in laboratory work; the author's personal experience of this is in working in an optics laboratory and adjusting the mirrors one-by-one to maximise the reading on a dial. The scheme proposed here goes beyond optimising in that sense and instead tries to learn how to achieve a gate without a good initial guess and with a large number of control parameters. Prior work in using physical quantum systems as a tool for optimal control has focussed on a range of different tasks, which are reviewed here.

The advantage of using *in situ* optimisation as a way of avoiding the need to have a very accurate model of the system parameters was recognised in [EW14] where the Ad-HOC (Adaptation by Hybrid Optical Control) scheme is introduced. The method there is to use a standard gradient based method to find an approximate control sequence and then to fine-tune this experimentally. They report significant gains in gate fidelity, but do not address the scalability problem as classical simulations are used as a first step. Simultaneously, a different group [KBCCC14] used a similar method named ORBIT (Optimised Randomised Benchmarking for Immediate Tuneup) to reduce control errors, and were also able to use it to correct for errors stemming from imperfections in the control pulses themselves. In that case randomised benchmarking was used to estimate the fidelity faster, which limited them to Clifford gates. The same group later extended their results to correct for parameter drift [KBFMJ16]. The relation between the final fidelity reached and the number of measurements needed is investigated in [FM15] using ACRONYM (Adaptive Control via Randomised Optimisation Nearly Yielding Maximisation), where stochastic gradient descent is used to greatly improve the scaling.

The same basic idea of testing controls by applying them to the system to be controlled can also be used for state preparation. This has been suggested in quantum chemistry [JR92, RdVRMK16] where it is called 'Learning Control' to distinguish it from open-loop and feedback control. It was investigated in [LLLKP17] for the creation of a highly entangled quantum state, concurrently to us developing it for gate synthesis. This concept can also be thought of as using a quantum system as an oracle for the classical simulator, as was done in [LYPS17], also to create an entangled quantum state. A different use for this type of oracle is to offer considerable speed up for more general optimisation problems [RSPL16].

The task of designing control pulses for a system which has not been completely characterised can also be viewed as a learning problem in a computer science sense. Machine learning as a classical field has had huge success in recent years due to the development of deep neural networks [LBH15], and is also being applied to quantum problems [PWZVS16, BWPRW17]. Such approaches have been used in a control context in two different ways: firstly by using classical networks to address quantum control problems, and secondly by making fully quantum networks. In both cases, techniques from machine learning are used to learn how to reach target operations more efficiently through a series of measurements and optimisation. These methods are used in [WDKGK17] to perform small scale quantum communication tasks, which also includes a proposal to scale this up using quantum neurones. Our *in situ* approach is focussed on implementing gates (and error correcting codes in §VI) without the machinery of neural networks, but share many of the steps in measurement and gradient based optimisation, and has been demonstrated numerically on a larger scale.

That there has been a considerable amount of work in the last three years in this area suggests that it is increasingly seen as a powerful tool in developing good control pulses. A key difference between this prior work and what we present is that the aim here is to discover a control sequence that implements a gate on a large quantum system. The focus is strongly on all aspects of the scaling relation with the number of qubits. Equally importantly, the results are used for a realistic near future quantum simulator and show that it is a viable proposal for turning it into a computer.

V.B The Scheme

The model considered here is a quantum simulator consisting of n qubits with some interactions between them such that they form a fully connected graph. Furthermore it is required that the timescale associated with this interaction is much shorter than the decoherence time in order for significant entanglement to be built up. In addition to this the ability to perform the following operations on each qubit individually is needed: preparation in a complete basis set of states, fast rotations by applying strong Hamiltonians, and measurement in a complete basis set. These requirements are significant, but much easier than demanding direct control over two-qubit operations, and can be achieved (or is close to being achieved) for the platforms described in §V.A.2.

Those requirements are more than enough for it to be, in principle, possible to do an arbitrary gate on a quantum simulator [Llo95, DNBT02, BBBG09], and to do process tomography [PCZ97]. While other systems satisfy these conditions and the approach detailed here would work, the focus is on this model for clarity. While it is easy to know that a given quantum simulator can do a universal set of gates (by using Lie algebraic methods as described in §I.C.2), finding the time-dependency for the control Hamiltonians

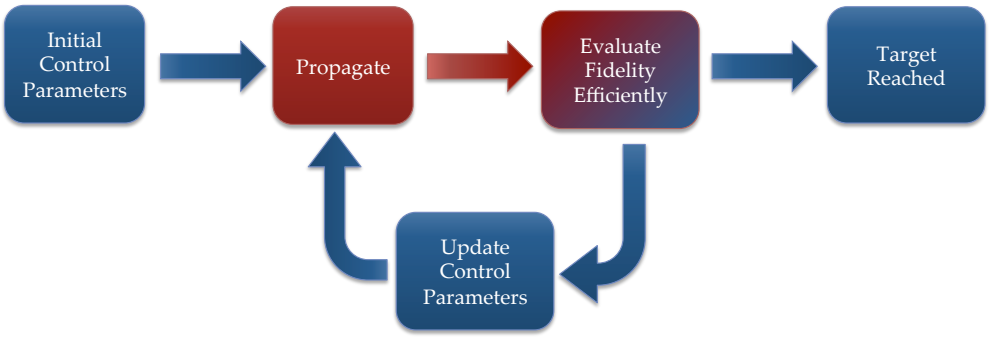


Figure V.2 | Outline of the process used in optimal control *in situ*, these are identical to Fig.I.3, except that the steps in red are done on the quantum simulator itself. The propagation is done completely *in situ*, while the fidelity estimation requires some classical processing too.

that give the desired gate is substantially harder.

The steps for finding such controls for a gate in a purely classical scheme were described in §I.C.3; the key differences between that approach and an *in situ* one are outlined in Fig.V.2. The starting point in both cases is an initial set of controls that parametrise the strength of the control Hamiltonians over the gate duration, in our examples these are generated randomly. The evolution of the system with these parameters is then calculated. On a classical computer this requires solving the time-dependent Schrödinger equation numerically for a model of the system, while in this scheme this is simply implementing the controls on the simulator. Evaluating the gate fidelity in the classical case is straightforward but, when done in situ, requires some form of tomography to measure it. Our derivation that a tight bound for this gate fidelity can be measured efficiently is given in §V.C. If this fidelity is above a threshold, the process terminates successfully, otherwise the control parameters are updated based on the results of the latest and previous runs, and the process repeats. There are many ways to update the controls classically that can also be used in this method. In the examples described in §V.D a steepest ascent method was used which requires the gradient of the fidelity in the control parameters to also be measured; this can also be done efficiently. The difficulty with doing these steps classically is in computing what unitary is produced by a given choice of control parameters; this requires both an accurate model of a high dimensional system and an exponentially large classical computer to solve it. Neither of those things can be done for a quantum system large enough to be an interesting quantum computer capable of doing things that said classical computer cannot.

These twin difficulties are eliminated by using the quantum simulator to compute the effects of the control pulse on itself. This works because the simulator with a trial set of controls is, evidently, guaranteed to be an accurate model of itself with those controls. The propagation step is therefore done *in situ*, but the method by which the control parameters are updated remains purely classical. This is because the information extracted from the

quantum simulator (the gate fidelity) and the parametrisation of the control pulses are purely classical. An upshot of this is that the myriad of different methods to do numerical optimisation that have already been developed and work for quantum systems, discussed in §I.C.3, can be used in this protocol directly. In order to use the quantum simulator as a universal computer, this optimisation procedure needs to be repeated for a universal set of gates.

As single qubit operations are assumed, the gates that controls are needed for are entangling ones, canonically the controlled-not (C-NOT) gate; these are vastly harder to perform using conventional methods and typically have much lower fidelities. The minimum number of these gates required, such that every quantum circuit can be implemented, is $n-1$. In practice it is more efficient, and produces shorter circuits, to find the $\frac{1}{2}n(n-1)$ C-NOT gates that act between every pair of qubits. Using the protocol detailed here, this is readily achieved even for pairs of qubits that are not directly interacting, as will be shown by numerical simulations in §V.D.

V.C Local Gate Fidelity

The measure used to gauge how close the system evolution is to the desired unitary is typically the gate fidelity generalised to CPT maps [GLN05]. This is a function between the dynamical map M, which describes the evolution of the system under a set of controls (including potential dissipation acting on the system), and the target unitary U. There are several subtly different definitions, the one taken here is

$$\begin{aligned} F(M,U) &= \langle\psi|\,\rho_M\,|\psi\rangle \\ \text{where} \qquad |\psi\rangle &= U\otimes\mathbb{I}\,|\Omega\rangle \\ \rho_M &= (M\otimes\mathcal{I})\,|\Omega\rangle\langle\Omega|\,, \end{aligned} \tag{V.1}$$

and $|\Omega\rangle = \sum_k \frac{1}{\sqrt{d}}|kk\rangle$ is a (normalised) maximally entangled state between the original Hilbert space and a copy of it. That is, the gate fidelity of two quantum maps is equal to the state fidelity of the Choi states (given by the Choi-Jamiołkowski isomorphism discussed in §I.B.2) of the two maps. This fidelity has several advantages: it is one of the most commonly used and so enables easy comparison of results, it is easy to compute numerically (compared to, for example, the diamond norm [AKN98]), and in the case that both maps are unitary, $M(\cdot) = V(\cdot)V^\dagger$, it simplifies to the standard form of the gate fidelity, $F(V,U) = |\frac{1}{d}\mathrm{Tr}[V^\dagger U]|^2$, mentioned in §I.C.3.

When the propagation step of Fig.V.2 is done classically, the whole unitary describing the evolution of the system is calculated as an exponentially large matrix from which the fidelity is computed. In the *in situ* scheme this is no longer the case, the only thing available is a quantum state after it has been evolved by the quantum simulator under the

control parameters. The standard method of extracting the gate fidelity is to perform a variant of process tomography known as certification. This requires preparing the system in a specific state, evolving it, and then performing a set of measurements. The number of different preparation-measurement combinations required scales exponentially [DLCP11]: $N_{\text{meas}} = O(d^2) = O(2^{2n})$. Such a scaling would immediately make the *in situ* scheme unscalable as it requires exponentially large experimental resources.

This problem can be circumvented, however, as the case of interest is not the distance to an arbitrary unitary but to a quantum gate. The notion of a gate implies that it acts non-trivially on a small (2, in the case of a C-NOT gate) number of qubits, and that this number does not scale with the overall system size, such that the unitary can be described as

$$U = \bigotimes U_i = \text{C-NOT}_1 \otimes \mathbb{I}_2 \otimes \mathbb{I}_3 \otimes ... \tag{V.2}$$

or similar. In such cases we have an exponentially faster estimate on the fidelity, the *local estimator* F_{LE}, which bounds the true fidelity according to

$$F(M,U) \geq F_{LE}(M,U) \equiv 1 - \sum_i (1 - F(M_i, U_i)), \tag{V.3}$$

where $M_i(\rho_i) = \text{Tr}_{\bar{i}}[M(\rho_i \otimes \frac{1}{d_{\bar{i}}}\mathbb{I}_{\bar{i}})]$ ($\bar{i}$ signifies all subsystems apart from i) is the reduced dynamical map acting on subsystem i where the other subsystems have been initialised in the maximally mixed state. This expression holds in all cases where U can be decomposed as in Eq.(V.2), and makes no assumptions about M beyond it being completely positive and trace preserving.

The proof of this bound uses methods from [CPFSG10]. The first step is to construct the projectors $h_i = \mathbb{I}_i - |\psi_i\rangle\langle\psi_i|$ for each U_i. These projectors have a very simple spectrum with a single 0 eigenvalue with corresponding eigenket $|\psi_i\rangle$, and a degenerate orthogonal space with eigenvalue 1. These projectors are summed together to from a Hamiltonian $H = \sum_i h_i \otimes \mathbb{I}_{\bar{i}}$ such that each h_i acts on its own part of the Hilbert space and is identity on the rest. This has a single $E_0 = 0$ eigenvalue with eigenstate $|E_0\rangle = |\psi\rangle$, while all its other eigenvalues are positive integers.

The expectation of this Hamiltonian with an arbitrary state is found by expanding it in its eigenbasis $\{E_k, |E_k\rangle\}$, giving

$$\text{Tr}[H\rho] = \sum_{k\geq 0} E_k \langle E_k|\rho|E_k\rangle \tag{V.4}$$

$$\geq \sum_{k>1} \langle E_k|\rho|E_k\rangle\,, \tag{V.5}$$

as $E_0 = 0$ and all the other energies are one or greater. The identity

$$\mathrm{Tr}[\rho] = \sum_{k\geq 0} \langle E_k|\rho|E_k\rangle = 1, \tag{V.6}$$

allows this to be expressed as

$$\mathrm{Tr}[H\rho] \geq 1 - \langle E_0|\rho|E_0\rangle \tag{V.7}$$

$$\geq 1 - \langle\psi|\rho|\psi\rangle\,. \tag{V.8}$$

If ρ is taken to be the Choi state of M, the definition of the gate fidelity gives

$$F(M,U) \geq 1 - \mathrm{Tr}[H\rho]. \tag{V.9}$$

The expectation value of the Hamiltonian with the Choi state of M can also be evaluated directly from its definition, which gives

$$\mathrm{Tr}[H\rho] = \sum_i \mathrm{Tr}[(\mathbb{I}_i - |\psi_i\rangle\langle\psi_i|) \otimes \mathbb{I}_{\bar{i}}\ \rho] \tag{V.10}$$

$$= \sum_i \mathrm{Tr}_i[(\mathbb{I}_i - |\psi_i\rangle\langle\psi_i|)\rho_i] \tag{V.11}$$

$$= \sum_i (1 - \langle\psi_i|\rho_i|\psi_i\rangle), \tag{V.12}$$

where $\rho_i = \mathrm{Tr}_{\bar{i}}[\rho]$ and ρ is the Choi state corresponding to M. The form of $\langle\psi_i|\rho_i|\psi_i\rangle$ suggests that it is also a gate fidelity. To write it as such, it is necessary to identify the quantum maps whose Choi states are $|\psi_i\rangle$ and ρ_i. From its construction, $|\psi_i\rangle$ is the Choi state corresponding to U_i. The dynamical map associated to ρ_i is

$$M_i(\cdot) \equiv \mathrm{Tr}_{\bar{i}}\left[M(\,\cdot \otimes \tfrac{1}{d_{\bar{i}}}\mathbb{I}_{\bar{i}})\right]. \tag{V.13}$$

This can be found by directly calculating the Choi state of M_i and showing it is ρ_i. To do so, $|\Omega\rangle$ needs to be expanded as

$$|\Omega\rangle = \sum_{p_i, p_{\bar{i}}} |p_i\rangle\,|p_{\bar{i}}\rangle\,|p_i\rangle\,\langle p_{\bar{i}}|\,, \tag{V.14}$$

where the ordering of the four terms is: the subsystem i, the rest of the Hilbert space of the system, the copy of the Hilbert space of the sub-system i, and the copy of the rest of

the Hilbert space. From this it follows from basic linear algebra that

$$\begin{aligned}
\rho_i &= \mathrm{Tr}_{\bar{i}}[\rho] && \text{(V.15)}\\
&= \mathrm{Tr}_{\bar{i}}\left[(M \otimes \mathcal{I}) |\Omega\rangle \langle\Omega|\right] && \text{(V.16)}\\
&= \frac{1}{d_i^2 d_{\bar{i}}^2} \sum_{p_i, p_{\bar{i}}, q_i, q_{\bar{i}}} \mathrm{Tr}_{\bar{i}}\left[M\left(|p_i\rangle \langle q_i| \otimes |p_{\bar{i}}\rangle \langle q_{\bar{i}}|\right) \otimes |p_i\rangle \langle q_i| \otimes |p_{\bar{i}}\rangle \langle q_{\bar{i}}|\right] && \text{(V.17)}\\
&= \frac{1}{d_i^2 d_{\bar{i}}^2} \sum_{p_i, q_i} \mathrm{Tr}_{\bar{i}}\left[M(|p_i\rangle \langle q_i| \otimes \mathbb{I}_{\bar{i}})\right] \otimes |p_i\rangle \langle q_i| \, d_{\bar{i}} && \text{(V.18)}\\
&= \frac{1}{d_i^2} \sum_{p_i, q_i} \mathrm{Tr}_{\bar{i}}\left[M(|p_i\rangle \langle q_i| \otimes \tfrac{1}{d_{\bar{i}}}\mathbb{I}_{\bar{i}})\right] \otimes |p_i\rangle \langle q_i| && \text{(V.19)}\\
&= \frac{1}{d_i^2} \sum_{p_i, q_i} (M_i \otimes \mathcal{I}) |p_i\rangle |p_i\rangle \langle q_i| \langle q_i| \,, && \text{(V.20)}
\end{aligned}$$

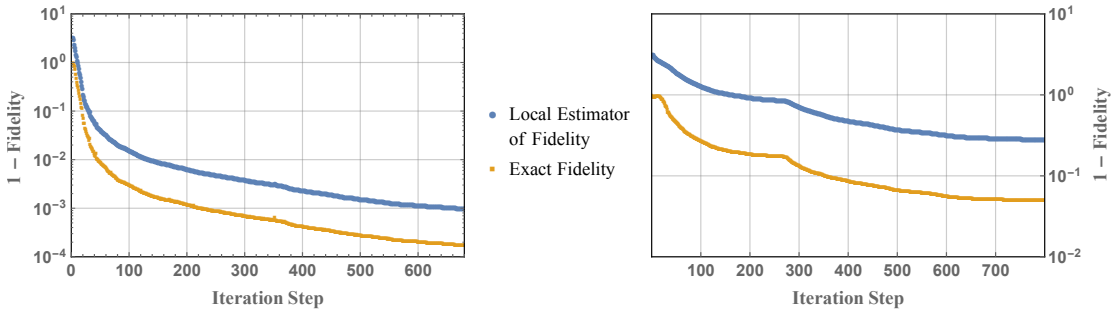
proving that ρ_i is indeed the Choi state of the map M_i as defined in Eq.(V.13). Putting this together with the expressions for $\mathrm{Tr}[\rho H]$ from Eq.(V.9) and Eq.(V.12) gives the desired expression for the bound the local estimator to the gate fidelity

$$F(M, U) \geq 1 - \sum_i (1 - F(M_i, U_i)),$$

where the lower bound is called the local estimator to the fidelity, F_{LE}, as previously stated in Eq.(V.3).

The name, local estimator, highlights the reason why this bound is useful. It is a function of operators, M_i and U_i, that are local to a single subsystem which is only one or two qubits large. It is therefore possible to construct a bound to the true gate fidelity using only local information, even though this information is only exponentially small compared to the size of the total Hilbert space. This means that certification needs to be performed on the two qubit subspace where the C-NOT gate is being performed and on each one qubit subspace where the identity is acting, but the correlations being built up between the subsystems are not needed to lower bound the fidelity.

An important question is whether this bound on the fidelity is a good one. Firstly it is a lower bound, which is useful as that means it guarantees that the gate is being done at least as well as the fidelity states, this is important to certify how well the quantum computer is performing. Another strong indication is its tightness: as $F(M, U) \to 1$, so does F_{LE}. As the fidelity is also upper bounded by 1, this means that F_{LE} constrains the true fidelity strongly. It is however possible that the local fidelity, when using it as the cost functional to optimise over, behaves erratically such that it hinders the optimisation algorithm. This is tested numerically in a few examples, shown in Fig.V.3, and is not found to be the case. Maximising the local fidelity also maximises the true gate fidelity almost completely monotonically while keeping the logarithmic distance constant. This is the desired behaviour to make F_{LE} a useful metric for optimal control.



(a) Ising Chain **(b)** Heisenberg Chain

Figure V.3 | Plot comparing the gate fidelity with the local estimator of the fidelity during an optimisation run. The gate fidelity and its local estimator, Eq.(V.3), are shown as a function of iteration step for one complete run of the *in situ* optimisation scheme. The system is a five qubit chain (Ising on the left, Heisenberg on the right) where the target is a C-NOT gate on the first two qubits and identity on the others. The algorithm minimised the infidelity of the local estimator. The exact infidelity is plotted at each step for comparison.

It remains to specify an explicit method for measuring F_{LE}. One way is to calculate each of the $F(M_i, U_i)$ sequentially by reconstructing the M_i sequentially. This can be done by initialising the state of one subsystem, i, in a specified pure state (for example, one of the three Pauli states if i is a qubit) and all the others in a maximally mixed state. After letting the system evolve by the dynamical map M, the state of i is measured in one basis of a complete set, and the others are ignored. This is repeated until this subsystem has been initialised and measured in each basis of a complete set; the number of combinations needed is $(d_i^2 - 1)^2$. This reconstructs M_i, from which $F(M_i, U_i)$ can be easily calculated. Doing this for every subsystem in turn leads to the overall number of measurements needed, N_{meas}, scaling as $O(\sum d_i^4) = O(n)$. The pre-factor for this scaling is the number of repetitions done for each preparation-measurement pair, which affects the accuracy of the estimated $F(M_i, U_i)$. How this affects the accuracy of F_{LE}, and its impact on optimisation, is investigated numerically in §V.D. This sequential approach is already an exponential improvement on previous methods, but the linear dependency can be removed.

Improving this scaling relies on two points: the method above is inefficient as most subsystems are not measured each time the system is evolved, and creating a maximally mixed state is equivalent to averaging over a sample of states drawn randomly from a complete basis set. Put together, this means that each of the M_i can be reconstructed simultaneously. This requires initialising the state of each of the subsystems in a random state drawn from a complete basis set, letting the system evolve, and then measuring each subsystem locally. The only thing recorded is how the measurement of the subsystems depended on their individual initial state; the potential correlations can be ignored which, after averaging, is equivalent to having initialised the other subsystems in a completely

mixed state. As before, this is repeated until each subsystems has been initialised and measured in a complete basis set. The scaling of this simultaneous approach thus depends only on the dimension of the largest subspace, giving

$$N_{\text{meas}} = O(\max d_i^4) = O(1). \tag{V.21}$$

This gives the very strong, and surprising, result that the local estimator to the fidelity, which bounds the gate fidelity, can be measured with a number of experiments that does not scale with the number of qubits. This result, however, does not take into account how the accuracy of this measurement of F_{LE} scales with the number of qubits, a question which is addressed below.

V.D Numerical Investigation

The local fidelity detailed in the previous section shows that it is possible to estimate the fidelity of a quantum gate efficiently as the size of the system increases, removing one direct barrier from the scalability of the *in situ* optimisation scheme. There are, however, other factors that determine the time the protocol takes which need to be taken into account in order to assess its scalability. This requires an expression for the total time required to construct a control sequence for a gate in terms of the number of qubits in the system. As this is an optimisation problem that would be done 'numerically' on a hybrid classical-quantum computer, analytic expressions could not be obtained. In order to investigate this, numerical simulations of the protocol have to be done on a purely classical computer; due to the exponential difficulty of simulating quantum mechanics, we could not go beyond nine qubits.

The average time needed to find a control sequence for a gate can be expressed as:

$$T_{\text{total}} = T_{\text{run}} N_{\text{runs}} / p_{\text{succ}}, \tag{V.22}$$

where T_{run} is the time it takes to do one run of a control sequence on the quantum simulator, N_{runs} is the number of sequences that are run on the simulator until the protocol halts, and p_{succ} is the probability that the protocol halts with a control pulse that reaches the desired fidelity. T_{run} can be decomposed as $T_{\text{run}} = T_{\text{init}} + T_{\text{gate}} + T_{\text{meas}}$ which is the time to initialise the system, evolve the system under the interaction and control Hamiltonians, and then measure it respectively. T_{init} and T_{meas} are determined by the type of system being used; they are taken as fixed and independent of the number of qubits. The gate time, on the other hand, is a free parameter that must be decided before starting the *in situ* optimisation.

The total number of runs can be similarly decomposed as

$$N_{\text{runs}} = N_{\text{meas}}\, N_{\text{prec}}\, N_{\text{fids}}\, N_{\text{upds}}. \tag{V.23}$$

N_{meas} is the number of times the experiment with the same control pulse, but with different input states and measurement basis, must be repeated in order to measure the gate fidelity once. As the previous section showed, this is $O(1)$ for the local estimator to the fidelity, which is the cost function used here. N_{prec} is the number of times the fidelity must be measured to acquire sufficient statistics such that the fidelity is known to the desired precision. N_{fids} is the number of different fidelities that need to be measured for the optimisation algorithm to update the control sequence. It is one for gradient-free methods, while for steepest ascent methods it is one plus the number of gradients (when they are measured by finite difference). N_{upds} is the number of the times the control sequences must be updated, corresponding to the number of times the scheme goes around the loop of Fig.V.2.

The scaling relation of the terms in Eq.(V.22) depends on the underlying classical algorithm. The algorithm used is a steepest ascent gradient method similar to GRAPE [KRKSHG05] as described in §I.C.3. In this approach, each of the independent Hamiltonians that can be controlled are taken as piecewise-constant with N_{ts} time slots of equal widths that span the full gate time T_{gate}. The precision to which the fidelities need to be measured experimentally also has to be specified. This could be done by either fixing N_{prec} itself, or by repeatedly measuring the fidelity until the error of the mean is below a specified value. The latter approach is approximated numerically by rounding each local fidelity measurement to some numerical accuracy A_{num} and calculating the equivalent N_{prec}. The *in situ* scheme therefore requires T_{gate}, N_{ts} and A_{num} to be chosen beforehand, as well as a target fidelity F_{targ}, and to know the number of different control Hamiltonians which, multiplied by N_{ts}, gives the total number of controls N_{ctrl}. In order to simulate this completely numerically, the exact control Hamiltonians and the constant interaction Hamiltonians for the system need to be specified. This is not the case were this scheme done *in situ* experimentally.

As a gradient optimisation method is used, it is necessary to recover the gradient of the fidelity with respect to the control parameters. In GRAPE there is one gradient per control parameter and they can be calculated numerically to machine precision with ease. However, in the same way as the unitary describing the system evolution is not available when optimising *in situ*, neither are the operators describing the gradients of the evolution. They must therefore be measured experimentally, the most direct way of doing this is by finite difference. This involves evolving the system with a modified set of controls: they are all the same except in the one control parameter α where the amplitude of the Hamiltonian is increased by ϵ. The fidelity gradient is then approximately given by

$$\frac{\partial F}{\partial H_\alpha} \approx \frac{F(H_C + H_\alpha \epsilon) - F(H_C)}{\epsilon}, \tag{V.24}$$

where H_C is the control Hamiltonian being used. This method works provided that ϵ is small compared to the curvature of the fidelity and that the gradient is large compared to the error in the fidelity. It requires $N_{\text{fids}} = 1 + N_{\text{ctrl}}$. This step was not carried out explicitly for the simulations done (except for the three qubit NMR case discussed separately) due to the large computational cost that would be required. There are other methods of approximating the gradients, based on expressing it as an integral over different evolutions and sampling from that integral, that may be more costly but more robust to noise and be independent of ϵ.

The numerical simulations of the *in situ* optimisation were completed using the quantum optimal control modules in QuTiP introduced in §I.C.3. These provide methods for optimising a control pulse to some fidelity measure. The *local Choi fidelity* measure customisation, and a method for automating locating the p_{succ} threshold, were developed for this study by us and Alexander Pitchford respectively. A high-performance computing cluster was necessary for completing sufficient repetitions of the optimisation simulations of the larger systems in a reasonable time (the nine qubit optimisations each required around four days of Intel Xeon CPU E5-2670 0 2.90GHz core processing time and were repeated hundreds of times). This is because the processing time required to optimise a pulse scales exponentially with system size due to the need to exponentiate the Hamiltonians in order to compute propagators. This difficulty precisely highlights the need to optimise pulses *in situ* for quantum systems of the size that would perform a useful quantum computation.

We found numerically that, for a range of examples, there exist values of T_{gate} and N_{ts} such that the *in situ* scheme converges. Fig.V.4 shows typical values of the most important parameters for a variety of topologies and interaction types. Our results consistently show that Ising systems (with a $\sigma_z \otimes \sigma_z$ interaction) are easier to find controls for than Heisenberg systems (where the interaction is $\sigma_x \otimes \sigma_x + \sigma_y \otimes \sigma_y + \sigma_z \otimes \sigma_z$). In particular, they suggest that in Heisenberg systems a GRAPE based algorithm may require a T_{gate} that scales exponentially with the number of qubits in order for the optimisation to succeed. This discrepancy also exists in purely classical optimisation techniques. This suggests that Heisenberg systems are intrinsically harder to solve with optimal control methods than Ising ones, and that this does not depend on whether an *in situ* or classical approach is used. This is backed up in Fig.V.3 where the local estimator to the fidelity is compared with the true fidelity for a five qubit chain as a function of the iteration step as it is being optimised. The local estimator tracks the true fidelity steadily whether the underlying system is Heisenberg or Ising, the notable difference between the two plots is the plateau in the Heisenberg case. This shows that optimising this system is significantly harder and appears for both the local estimator and the true fidelity. An exponential scaling in

Topology	Coupling	T_{gate}	N_{ts}	N_{upds}
chain	Ising	π	12	60
star	Ising	π	12	214
fully connected	Ising	12π	160	295
chain	Heisenberg	16π	160	585
star	Heisenberg	12π	160	1043
fully connected	Heisenberg	12π	160	881

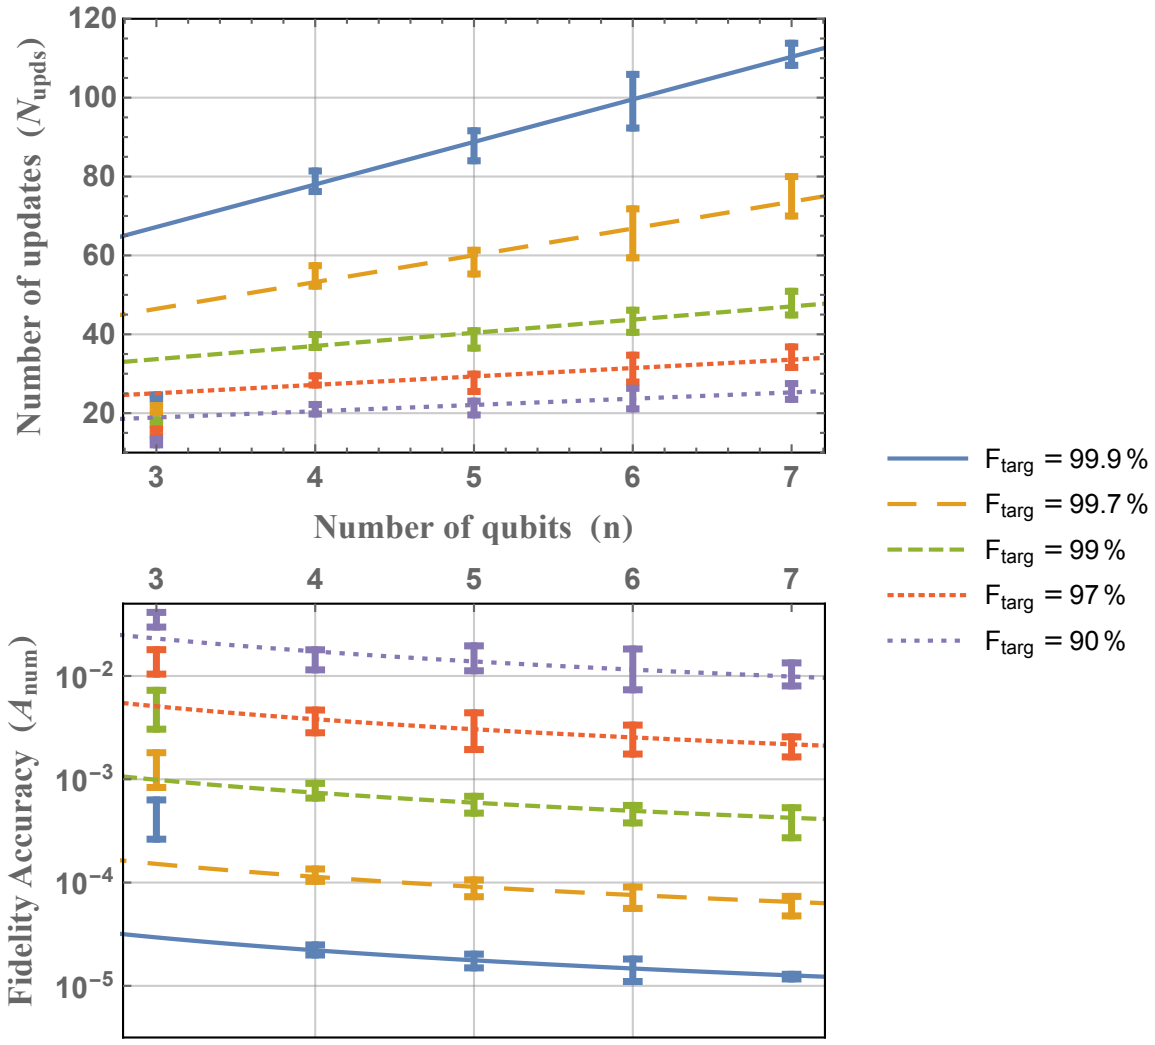
Figure V.4 | Summary of the cost of performing the *in situ* optimisation scheme for a range of different five qubit systems. The differences between the systems are their topology (a linear chain with nearest neighbour interactions, a star where all interact with a central qubit only, or fully connected where the interaction strengths are also randomised) and the interaction type. In each case the controls are σ_x and σ_y on each individual qubits such that $N_{\text{ctrl}} = 10$, the target operation is a C-NOT gate on two qubits and identity on the rest, $F_{\text{targ}} = 0.999$, and $p_{\text{succ}} > 0.98$. These simulations where done with full numerical precision. For five qubits, all six systems find the desired entangling gate, and do so at reasonable experimental cost. This indicates that the approach works for a range of possible quantum simulators.

the required gate time would make such an approach infeasible for a quantum computer. However, it can be shown [DNBT02] that it is possible to do fast entangling gates on such systems, so it is only the optimisation algorithm that struggles to find the solution. It may be the case that using a different algorithm inside the *in situ* protocol, such as parametrising the control Hamiltonians as a Fourier series rather than as piecewise-constant, would find control pulses for shorter gate times.

For the case of an Ising chain the number of qubits is also varied in order to hypothesise a scaling for T_{total} and the results are consistent with a polynomial scaling. In these simulations, whose results are shown in Fig.V.5 and Fig.V.6, the number of qubits goes from three up to seven or nine in some cases. The interaction is always $\sigma_z \otimes \sigma_z$ between nearest neighbours, with σ_x and σ_y controls at each qubit, gate time $T_{\text{gate}} = 4\pi$, and $N_{\text{ts}} = 48$ time slots. The target operation is a C-NOT between two qubits and identity on the rest and the figure of merit is F_{LE}.

In one of these parameter settings the impact of measurement noise is investigated by introducing a finite value of A_{num}, which characterises the sensitivity of the system to measurement noise, in order to obtain an expression for the total scaling. The expected scaling is $O(\epsilon/n)$ as, in order to reach a gate infidelity of ϵ, the fidelity ought to require a measurement accuracy $O(\epsilon)$; as this is calculated from the sum of the fidelities of the subsystem, we conjectured that they each need to be measured to an accuracy $O(\epsilon/n)$. The date shown in Fig.V.5b lie very close to a c/n curve, providing strong support for this argument. However, the constant c does not appear to have quite a linear relationship with ϵ; we did not investigate this further as it does not affect scalability. This relationship between A_{num} and n implies that $N_{\text{prec}} = O(n^2)$, due to the central limit theorem that states the number of repetitions required scales quadratically with the inverse desired

(a) N_{upds} scaling



(b) A_{num} scaling

Figure V.5 | Plot showing the number of updates and the fidelity accuracy required in a numerical simulation to find controls for a C-NOT gate. The system is an Ising chain and a steepest ascent *in situ* algorithm is used for optimisation. The target in each case is a C-NOT gate between two qubits in the middle of the chain, separated by one other qubit. The data at $n = 3$ has been omitted from all fits as it is qualitatively different from the rest.
Figure a) shows the how N_{upds} scales with the number of qubits for different target gate fidelities (error bars are twice the standard error). For this plot, the accuracy to which the fidelity is measured, A_{num}, is picked to give a $p_{\text{succ}} = 50\%$ success rate. There is a strong linear relation in the number of iterations required as a function of the number of qubits giving $N_{\text{upds}} = O(n)$.
Figure b) shows how the accuracy to which the local fidelity needs to be measured, A_{num}, scales with the number of qubits for different gate fidelities (error bars are five times the standard error). The points lie very close to a fit of $1/n$ for each target gate fidelity, as would be expected.

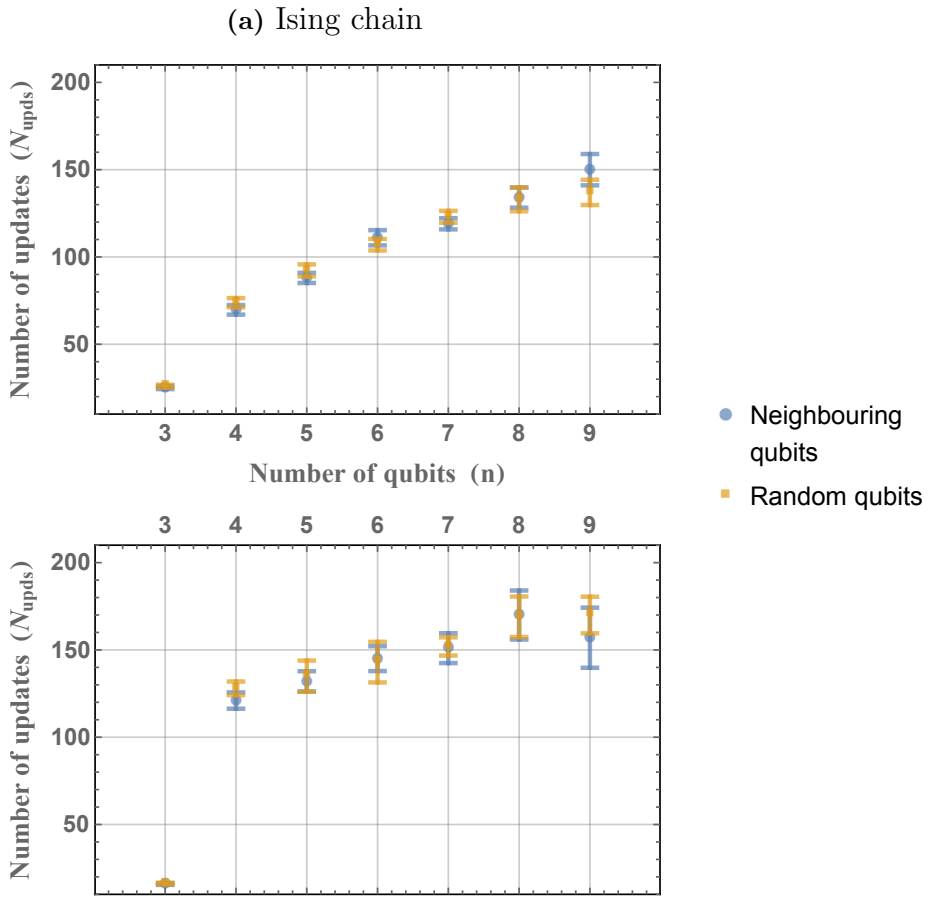


Figure V.6 | Plot of the number of control pulse updates needed to find a C-NOT gate in an Ising chain and ring. The number of updates required for the optimisation protocol to succeed is plotted for a chain (top) and a ring (bottom) with a gate fidelity of $F_{LE} = 0.999$ and using full numerical accuracy. Each data point represents repeated optimisations, 100 for $n < 8$, 96 for $n = 8$, and 30 for $n = 9$. The number of successful optimisations is $p_{\text{succ}} > 90\%$ in all cases. The error bars are twice the standard error.

accuracy: $N_{\text{prec}} \propto A_{\text{num}}^{-2} = O(n^2)$. The result for the number of times the control pulse needs to be updated is in Fig.V.5a and gives $N_{\text{upds}} = O(n)$. Putting this together with the previous results that $N_{\text{fids}} = O(n)$ for finite-difference gradient methods (as here the number of time slots is constant and the number of control Hamiltonian is linear in the number of qubits) and $N_{\text{meas}} = O(1)$ for the local estimator fidelity gives $N_{\text{runs}} = O(n^4)$. As this is done with a constant gate time and with a constant success probability, this implies that the time required to find a control sequence that implements a C-NOT gate on an Ising chain using a steepest-gradient *in situ* scheme scales as $T_{\text{total}} = O(n^4)$.

The other system investigated in depth is similar but compares the effect of the C-NOT being between two adjacent or two random qubits, and whether the difference between a chain and a ring is significant. Fig.V.6 shows the required N_{upds} for up to nine qubits for these systems. In both cases the scaling is sub linear, even disregarding the three qubit

data. As there is no obvious model to fit to these points apparent to us, no best-fit is shown. The required iterations is slightly higher for the ring than the chain, but appears to have a smaller gradient in n. For both graphs the results for nearest neighbour and random qubits are statistically indistinguishable. This highlights that the optimisation scheme operates equally well in both cases and works more efficiently than a naive dynamically decoupling protocol [DNBT02]. Specifically, it shows that the scalability of the scheme does not rely on boundary effects or on qubits being adjacent to each other. While a quantum computer could be built using only nearest neighbour gates, being able to entangle two arbitrary qubits in the time of a single gate drastically reduces the potential run time of algorithms and with it the negative effects of decoherence.

V.D.1 NMR Experiment

During the course of this project we were in contact with Steffen Glaser and David Leiner about an experimental realisation of these ideas in a liquid NMR system. The system was diethyl fluoromalonate-2-^{13}C . This is a long molecule which acts as a heterogeneous three qubit ring with effective Ising interactions with coupling strengths on the order of 100Hz and relaxation times of several seconds [LG16]. As NMR is done with large ensembles of (identical) molecules in a pseudo-pure state, experiments are naturally done with a great many repeats simultaneously. This means that the N_{prec} parameter can be taken to be large as that the finite-measurement noise is small. However, due to being in only a pseudo-pure state means that the signal-noise ratio limits the accuracy of the measurements. In practice, this limited the accuracy of the fidelity by $\sim \pm 20\%$. The task therefore was to find a parameter regime where the scheme would converge to a pulse (of maximum $\sim 80\%$ fidelity) under this much measurement noise with in a reasonable experimental time T_{total} as given in Eq.(V.22). Specifically, the proposal needed to complete within a week as that was the longest the necessary equipment could be used continuously.

The best compromise led to a required number of runs to complete the optimisation scheme of $N_{\text{runs}} \approx 10^4$ with a success probability as low as $p_{\text{succ}} \approx 10\%$. This would have been workable as the gate time is few miliseconds. This was not limiting factor, however, as the measurement took a few seconds and the relaxation time required to reinitialise the system was up to ten seconds. Due to these large time scales, the proposed experiment was not feasible. Surprisingly, it was the non-scaling part of Eq.(V.22) that proved the most detrimental. Nevertheless this showed that even when considering noisy systems and taking into account the effects of finite gradients, *in situ* optimisation can still converge to a working pulse.

V.E Conclusion

The polynomial scaling found in Fig.V.5 and Fig.V.6 suggests that the protocol can indeed be efficient. Some of the components that make up this scaling come from numerical data, so several fits are possible. However the points lie so close to a linear fit in Fig.V.5 that a different model, such as an exponential one, would diverge only slowly. Fig.V.6 suggests than corrections to it are actually more likely to make it sub-linear than more costly. For this reason it is reasonable to extrapolate from the data to moderately larger systems of a few tens of qubits. Systems of such a size are interesting as they correspond to the state of the art that can be realised experimentally for a range of systems mentioned in §V.A.2. Using the *in situ* scheme in such a situation would likely find control sequences for entangling gates that are not currently known, and where purely classical numerical optimisation schemes would fail due to the enormous computational requirements. Furthermore, testing these predictions in such experiments would extend these results to numbers of qubits that are completely unattainable for a purely classical computer to model. Doing such experiments would thus test this protocol closer to the scale needed for a useful universal quantum computer.

One potential difficulty in optimal control is the existence of traps: local maxima of the fidelity that optimisation algorithms converge to which are not the global maxima. The question of whether traps exist in unitary control using the standard gate fidelity has been well studied [HDR09, PT11, RRW16], and the conclusion is that generic quantum control landscapes are almost always trap free. This may also apply to the local estimator of the fidelity; traps were not a problem for the numerical simulations we performed and the behaviour of the local fidelity in Fig.V.3 suggests that the local fidelity does not add any extra traps.

The numerical results presented here used GRAPE, which decomposed the control pulses into piece-wise constant functions. A potentially more powerful approach, and one which is harder to do classically but may be easier to implement physically, would be to decompose them by frequency [BM13], such as in CRAB [DCM11, CCM11] and GOAT [MTWA15]. Such algorithms are slow to run classically due to the difficulty of exponentiating the time-dependent Hamiltonian, a step which is bypassed in the *in situ* scheme. They typically require fewer parameters to describe a successful control pulse and thus may prove faster than GRAPE when done experimentally. A different variation would be to change the type of classical optimiser used from a gradient based algorithm to a geometric or genetic one, or even to use machine learning algorithms to extract information out of the system, as discussed in §I.C.3 and §V.A.3.

A natural extension of this work is to apply it to another important aspect of quantum computation: error correction. Our work on this is presented in the following chapter.

CHAPTER VI

IN SITU ERROR CORRECTION

VI.A Introduction

This chapter extends the results of §V to finding error correcting codes. As errors in any quantum computation are never completely avoidable, being able to mitigate their effects without needing to first characterise them accurately is an important step in scaling up *in situ* optimisation. Our current results on this are still incomplete and we are investigating them further, but the current results are very promising. They show our *in situ* approach is capable of a range of interesting tasks in correcting errors in many-body quantum systems.

VI.A.1 Overview

The previous chapter provided a way to learn a control sequence that performs entangling gates on a quantum many-body system such as a simulator. Being able to implement a universal set of gates is of primary importance in using such a system as a quantum computer, but there is another crucial process required: error correction. In an analog simulator, like the ones in §V.A.2, errors happen slowly and often a small logical error corresponds to a small error in the result of the simulation because of locality constraints [GAN14]. This is not the case for a universal quantum computer, not only will distant qubits be entangled through gates but the computation may take arbitrary long. For this reason, it is accepted that some form of error correction is required for a large scale device [NC00, LB13].

This chapter presents work which shows that *in situ* optimisation can be used to discover a variety of error correcting codes for an ensemble of qubits. The underlying idea is a generalisation of the previous chapter. There, optimisation was used to find a control pulse that can perform a quantum gate on a system consisting of many qubits. Here, the task is to use optimisation in order to find controls that can protect a subset of the initial qubits (the carriers) from noise. The target operation has therefore moved from a unitary over all qubits (such as C-NOT$\otimes\mathbb{I}\otimes$...) to the identity over some of the

qubits only, with the others (the ancillas) being discarded. As the control pulse is found *in situ* it is guaranteed to give the same fidelity found during the optimisation as when applied during a 'real' algorithm as the two situations are indistinguishable. Because of this, coherent errors due to imperfect control pulses can be minimised by increasing the target fidelity used during optimisation. This leaves incoherent errors that arise from dissipative processes with an environment as the dominant ones to eliminate [LB13].

Our numerical simulations of this scheme were able to find control Hamiltonians that implemented well known procedures for preventing and correcting errors including the three-qubit Shor code, the five-qubit stabiliser code, and a two-qubit decoherence free subspace. On the latter they were also able to demonstrate that the protocol can be extended to find gates acting on the logical qubit which is spread out over both physical qubits. Furthermore, the protocol is able to go beyond recovering existing codes and, by changing the noise included in the simulation, can discover new codes that are tailored to the errors in that system.

These are exciting results that address two important difficulties in error correction. Firstly, it eliminates the need to measure experimentally, or make assumptions, about the possible errors that can occur in the system. Determining the possible errors is hugely costly and scales exponentially with the system size [MRL08]. For this reason, and to obtain analytic results, many protocols assume that noise will induce uncorrelated single qubit errors [LB13]. Using *in situ* optimisation naturally targets the real errors that are present in the system, whether they are correlated or not, without needing those errors to be identified beforehand. Secondly, the scheme here provides the control Hamiltonians that perform the operations involved in error correction. For the numerous platforms for quantum technologies discussed in §V.A.2 this allows the error correction to be implemented directly. This is an advantage compared to an analytic approach which gives a large unitary matrix that must be decomposed into a sequence of gates, which may then be implemented imperfectly.

Simple examples of error correcting schemes and decoherence free subspaces are explained below in §VI.A.2. The details of how *in situ* optimisation can be used to achieve such tasks is then given in §VI.B. This is followed by §VI.C detailing the results of numerical simulations of this approach. Finally, §VI.D discusses the implication of the results and potential future work, both theoretical and on an experimental implementation.

VI.A.2 Background

The discovery that quantum error correction is possible [Sho95] was the start of the quantum information revolution. It showed that a large scale quantum computer was, at least in theory, capable of mitigating the unavoidable effects of noise. Since then the field has developed hugely and it is now known that once errors are below a certain threshold,

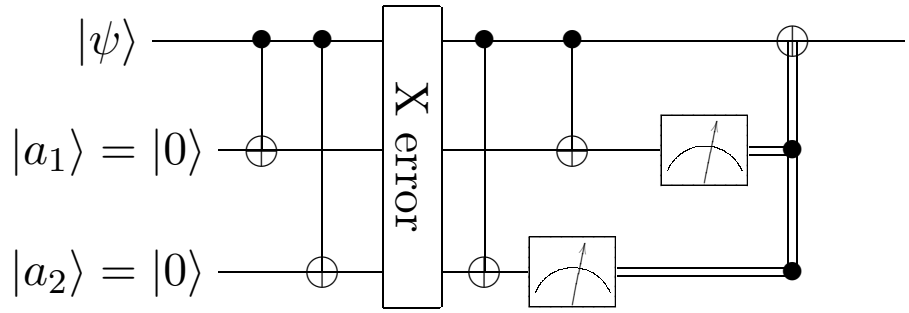


Figure VI.1 | Circuit diagram representation of the three-qubit Shor code. The first two gates encode the logical qubit across all three physical qubits. This is followed by an error consisting of a X rotation on any single qubit. The logical qubit is then moved back to the carrier. Finally, the ancillas are measured which determines which qubit (if any) was flipped, this is then used to correct the error.

it is possible to correct them faster than they occur, even if gates and other operations performed on the qubits are imperfect [Got07]. Such fault tolerant quantum computing is beyond what we have demonstrated numerically with the *in situ* optimisation scheme so far. The focus here is on performing simple error correction codes with a small number of qubits rather than on asymptotic limits. Before introducing our results, a few simple and well known error correcting protocols are reviewed [NC00, DMN13, LB13], as they are closely related to the approach we took.

The first quantum error correction was inspired by its classical analog where the state of one logical bit, 0 or 1, is copied to several bits. Any error on a single bit can be detected by observing the state of all the copies and taking the majority vote. This cannot be directly transferred to the quantum scenario due to the no-cloning theorem and collapse caused by measurement. Nevertheless a simplified three-qubit variant of Shor's original idea [Sho95], shown as a quantum gate circuit in Fig.VI.1, encodes one logical qubit across three physical qubits such that it is resistant to any single qubit flip (σ_x rotation by π).

The starting point is the carrier qubit in the logical state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ and two ancilla qubits in the state $|0\rangle$. These are entangled together via two C-NOT gates to create the state $\alpha|000\rangle + \beta|111\rangle$. The logical qubit, the state $|\psi\rangle$, is at this point spread out over all three physical qubits. There is then the possibility of having a single X error on any qubit, which flips $|0\rangle \leftrightarrow |1\rangle$. The gates and Z measurements on the ancillas after this perform a bell measurement on the system. This measures the parity of pairs of qubits without, crucially, measuring the logical state itself. This allows the presence of errors to be detected, which can then be corrected by an X gate on the first qubit classically controlled by the outcome of the measurements. It is also possible to remove the measurements and replace them with a fully quantum controlled-controlled-not gate instead, which pushes the error from the carrier qubit (if there was one) to an ancilla. Either way, the ancilla qubits must be discarded after error correction. This is the type

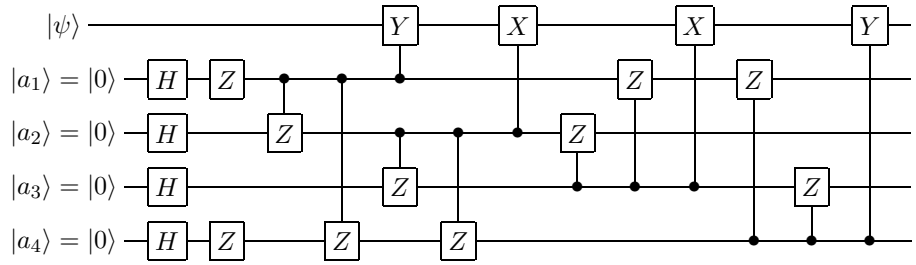


Figure VI.2 | Circuit diagram for the encoding step of the five-qubit code [Got97]. The substantial complexity increase compared to Fig.VI.1 is indicative of the need to prepare a special highly entangled state.

of protocol, unitary evolution - error - unitary evolution - ancilla discarded, that can be recovered by *in situ* optimisation.

Such a scheme, while being a nice theoretical tool, can only correct X errors. It can be generalised to the nine qubit Shor code, which protects against any single qubit X or Z (and therefore also Y) errors. A more economical form of this is the seven qubit Stean code [Ste96] and, the smallest possible code that corrects any single qubit error, the five-qubit stabiliser code [LMPZ96]. The latter is relevant here as it the largest space that we could simulate numerically in the *in situ* scheme, and is capable of correcting both bit and phase errors. It is normally described in terms of stabiliser states, but its fundamental workings are the same as for the Shor code. The encoding step is substantially more complicated than for the three-qubit code and is shown in Fig.VI.2 [Got97]; the same circuit run backwards decodes the logical qubit. Errors are then corrected by measuring stabilisers and performing the requisite single qubit correction. As with the Shor code, this can also be done with purely unitary operations.

Both of these codes worked against single errors that affect only one qubit. It is desirable however to think of different types of errors that acts on multiple qubits in a correlated way. The typical example of this is collective dephasing [Lid12]. This is an extension of the dephasing Lindbladian of §I.B.6 such that both qubits suffer correlated decay. It can be understood as two qubits gaining an identical but unknown energy difference, as sketched in Fig.VI.3. Such errors are common place in a large number of physical systems where the internal energy structure of the qubits fluctuates with an external field which is constant across several qubits. The information, $|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$, stored in a single physical qubit will gain an unknown phase difference between the two basis states. Averaging over this unknown phase leads to decoherence and the loss of phase information between α and β; a decoherence free subspace (DFS) is a way to mitigate this by encoding the logical qubit in a subspace of a larger system where the unknown phase is the same.

The two-qubit case works by mapping the state of the logical qubit from the carrier

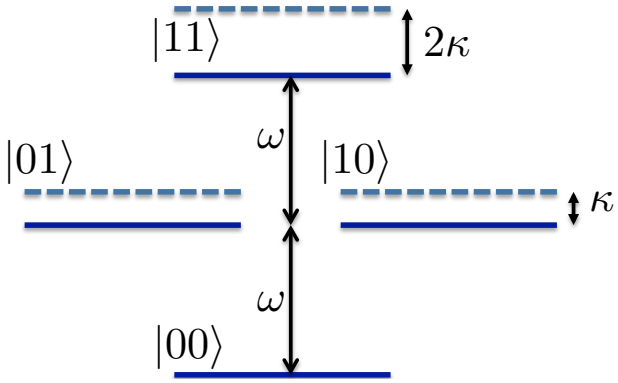


Figure VI.3 | Energy levels involved in two-qubit collective dephasing. Each qubit has an unknown energy term κ that, when averaged over, causes loss of phase information. As the states in the $\{|01\rangle, |10\rangle\}$ submanifold have the same energy dependency in κ, it can be used as a DFS to protect a logical qubit from such noise.

to an entangled state of the carrier and ancilla

$$|\psi\rangle|0\rangle \xrightarrow{U_{\text{enc}}} 0|00\rangle + \alpha|01\rangle + \beta|10\rangle + 0|11\rangle. \tag{VI.1}$$

The evolution after time t in a frame rotating with the known internal energy ω gives

$$0|00\rangle + e^{-i\kappa t}(\alpha|01\rangle + \beta|10\rangle) + e^{-2i\kappa t}0|11\rangle, \tag{VI.2}$$

where κ is the fluctuating internal energy. While the three terms have different (and unknown) phases that ought to be averaged over, only the middle term is non-zero. The unknown phase thus becomes a global phase that can be ignored, such that averaging over it has no impact. The quantum information can be mapped back onto the original qubit

$$\alpha|01\rangle + \beta|10\rangle \xrightarrow{U_{\text{dec}}} |\psi\rangle|0\rangle, \tag{VI.3}$$

such that this scheme also falls into the same unitary evolution - error - unitary evolution framework as the previous examples.

The rest of this chapter shows that these three methods can be learnt by *in situ* optimisation as well as variants which specifically target whatever errors are physically present.

VI.B The Scheme

In order to recover and extend the schemes discussed above to the *in situ* scenario, the model considered is a quantum many-body system subjected to some unknown noise. This system is made up of qubits, divided into carriers and ancillas. For simplicity, and to keep

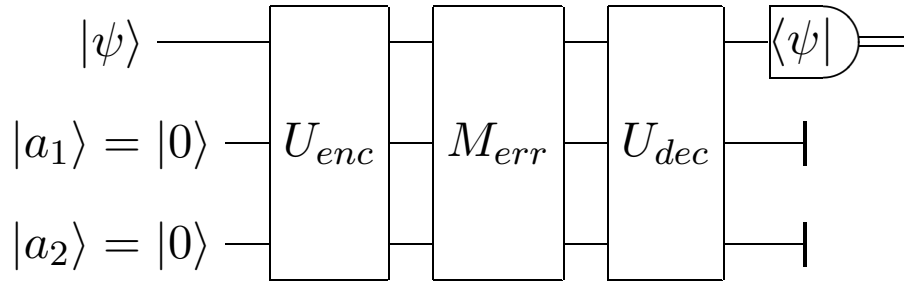


Figure VI.4 | Outline of the steps involved in *in situ* error correction. The carrier qubit in an unknown state $|\psi\rangle$ is to be preserved through an, also unknown, lossy quantum channel M_{err}. This is done with the help of ancilla qubits, initialised in a known pure state, that evolve with the information carrier through two unitaries, an encoding one before the errors occur and a decoding one afterwards. The optimisation task is to find which control Hamiltonians give the unitaries which best put the carrier qubit back into the state $|\psi\rangle$. The final states of the ancillas are not important and are discarded without measurement; the carrier qubit is measured while the scheme is being optimised but not when used to correct errors in a real application.

numerical simulations small, only systems with a single carrier qubit and up to four ancilla qubits are considered. The task is to preserve the state of the carrier from the effects of dissipation; this is done by varying control Hamiltonians in time. Such controls may act only locally on individual qubits while an interaction Hamiltonian is constant, or all parts of the Hamiltonian can be controllable. Either way, the system must be complex enough to be fully controllable (as discussed in §I.C.2). In addition, the ability to prepare the ancilla qubits in a pure state is required, as well as the means to prepare and measure the carrier qubit in states from a complete basis set. These are all very similar requirements to those put forward in the previous chapter in §V.B.

The basic workings of the scheme are shown diagrammatically in Fig.VI.4. Initially, the system has some quantum information stored, a logical qubit encoded in the carrier, while the ancillas are initialised to a known state. It then must be evolved for some time, under a suitable Hamiltonian to be determined, such that the logical qubit is spread out across the system in a way that makes it resilient to the noise on the system. This is the 'encoding' step. It is followed by the dissipative process which is the source of errors on the system. Finally, the system must evolve again, under a different Hamiltonian also to be determined, so that the logical state is mapped back onto the carrier; this is the 'decoding' step. *In situ* optimisation finds controls for the encoding and decoding step that protects the information without needing to know what the potential errors or the details of the system are.

The evolution of the carrier qubit ρ is given by the corrected dynamical map

$$M_{\text{cor}}(\rho) \equiv \text{Tr}_A \left[U_{\text{dec}}\, M_{\text{err}} \left(U_{\text{enc}}(\rho \otimes \rho_A) U_{\text{enc}}^\dagger \right) U_{\text{dec}}^\dagger \right], \tag{VI.4}$$

where ρ_A is the initial state of the ancilla qubits, taken as $|0...0\rangle\langle 0...0|$, and the other operations are as defined in Fig.VI.4. In order to negate the effect of M_{err}, M_{cor} must be as close to nothing happening as possible. The cost functional to optimise over is therefore the gate fidelity (as defined in §V.C) between the corrected error channel and the identity, $F(M_{\text{cor}}, \mathbb{I})$. Alternatively, it is possible to require a logical operation to be performed on the qubit at the same time as it is being protected from noise, in that case the gate fidelity to use is between M_{cor} and the desired unitary U.

For the examples considered, both M_{cor} and the target act on single qubits (not considering the ancillas which are added and traced out inside of M_{cor}), so measuring this gate fidelity is straight forward. It can be done by initialising the carrier qubit in one state, evolving it under M_{cor} and measuring the probability that this same state is observed at the output. Averaging these probabilities over a complete set of basis states gives the gate fidelity. This scheme can be extended to a logical space of more than one qubit. In such a case the local fidelity, discussed in depth previously, can be used to approximate the real fidelity efficiently in the size of the logical space. It is important to stress that, whichever method is used, the cost of measuring the fidelity itself does not scale with the size of the ancilla space at all as the ancilla qubits are always prepared in the same state and never measured.

An implicit assumption in the description so far is that the errors induced by the environment happen at a fixed point in the evolution of the system. This can be thought of as an approximation where U_{enc} and U_{dec} can be implemented quickly compared to the timescale on which errors occur, such that M_{err} takes place over a long period of time during which all Hamiltonians in the system are switched off. Working in such a framework has two significant advantages. Firstly, it allows errors that do not naturally arise from Lindbladians to be considered, such as limiting them to at most a single qubit flipping. This is noise that can theoretically be corrected perfectly and so provides a convenient testing ground. Secondly, it drastically reduces the cost of numerical simulations as only the closed evolution under the controls needs to be computed; having a constant Lindbladian would require the propagators to be calculated in the doubled space which is very costly. Nevertheless, we also consider some numerical examples of the more general case of errors happening continuously during the protocol. In that case, the map M_{err} is replaced by a waiting period with no control Hamiltonians and a Lindbladian acts on the system from the start of U_{enc} until the end of U_{dec}. This does not affect the way the fidelity is calculated or the steps involved in the optimisation.

The *in situ* optimisation is used to find the set of controls that give rise to the U_{enc} and U_{dec} that have the desired behaviour. The process by which it does this is very similar to the iterative optimal control method described in detail in Fig.V.2, the major change being that the fidelity used is only on the carrier and ignores the ancillas. This means that the scheme does not attempt to find predetermined encoding and decoding unitaries

separately, it only considers their joint effect on the errors. Indeed, this approach bypasses the need to determine what U_{enc} and U_{dec} should be. This is fundamental to this scheme as precalculating those would require a detailed knowledge of M_{err}, and the ability to analytically or numerically find a code to mitigate it, which is precisely what is being avoided.

There are, however, clear physical limitations as to what can be achieved by error correction, whatever approach is used. If there are too many potential errors and not enough ancilla qubits then there may not exist a code that will correct all the errors [KL96, KLV00]. This is not something *in situ* optimisation can circumvent and, in such a case, only the largest errors will be corrected, as shown in §VI.C.1. This is the best that can be expected and is very relevant to any possible physical implementation of this protocol: there are likely to be a huge number of errors that can occur and correcting them all is impossible, all that is necessary is to correct the most commonly occurring ones until the fault tolerant threshold is reached.

VI.C Numerical Investigation

In order to assess the applicability of using *in situ* optimisation to perform error correcting, a number of numerical simulations are presented here that we performed using the control libraries of QuTiP (introduced in §I.C.3) and the packages that were developed for §V.D. Additional code was also written to deal with the different fidelity used and to add CPT maps in the middle of an otherwise unitary simulation. The aim here is to show what is possible to do with this method, questions of scalability are not addressed explicitly, although the cost of finding an error correcting protocol for the different models used is given. The three error correcting protocols discussed in §VI.A.2 are the different problems considered.

VI.C.1 Three-Qubit Shor Code

The three-qubit short code, as discussed in §VI.A.2, is one of the simplest error correcting codes. The first aim is to 'rediscover' it (or at least, a code that has the same outcome) *in situ* following the method outlined in Fig.VI.4 with one carrier and two ancilla qubits. A number of different control settings are investigated. The first one is a fully controllable Ising chain where the $\sigma_z \otimes \sigma_z$ interaction between the first pair and the last pair of qubits are independently controllable, as are σ_x and σ_y on each qubit. The second scenario is a locked chain where the controls are the total interaction, $\sigma_z \otimes \sigma_z \otimes \mathbb{I} + \mathbb{I} \otimes \sigma_z \otimes \sigma_z$, and independent local controls. Lastly, there is the locally controlled chain where the interaction term is the same as for the locked chain but is now fixed and not a control parameter, and only local Hamiltonians are variable. In all cases the M_{err} is the same as

Control Type	M_{err}	Target	T_{evo}	N_{ts}	N_{upds}	p_{succ}
Full Controls	M_X	$\mathbb{I}$	1	10	78	100%
Locked Chain	M_X	$\mathbb{I}$	1	12	56	100%
Local Controls	M_X	$\mathbb{I}$	20	14	61	100%
Local Controls	M_X	X	20	14	66	90%
Local Controls	M_X	Z	20	14	64	80%

Figure VI.5 | Results of recovering the three-qubit Shor code using *in situ* optimisation. Each lines consists of 10 repeats and the standard errors in N_{upds} are all approximately ± 10. The target fidelity in all cases was 0.99; without any error correction the fidelity is 0.75. T_{evo} is the total system evolution time, N_{ts} is the number of time slots to realise both U_{enc} and U_{dec} (split evenly between them), N_{upds} is the number of iterations where the controls were updated, and p_{succ} is the probability of the scheme finding a control pulse that matches the target fidelity. A more detailed explanation of these parameters is given in §V.D.

in the standard Shor code

$$M_X(\rho) = \frac{1}{4}\left(\rho + X\mathbb{I}\mathbb{I}\rho X\mathbb{I}\mathbb{I} + \mathbb{I}X\mathbb{I}\rho\mathbb{I}X\mathbb{I} + \mathbb{I}\mathbb{I}X\rho\mathbb{I}\mathbb{I}X\right), \tag{VI.5}$$

where the tensor product between the $\mathbb{I}$ and X operators are assumed, and the target operation on the carrier qubit is $\mathcal{I}$. The results for this are shown in Fig.VI.5. In all cases, the scheme was able to find a control pulse that protected the state of the carrier qubit very well at reasonably small experimental cost. The evolution time T_{evo} and requisite time slots N_{ts} get higher as the available controls are reduced because the required pulses become more complex. It is also possible to do a gate on the carrier at the same time as protecting it from noise, as shown in the last two lines of the table. This however appears harder to do as the success rate has dropped; this is unexpected as the presence of local controls means that it ought to be trivial to do a single qubit gate at the end of U_{dec}. The fact that the optimisation procedure does not always find this solution, and that every time it failed to converge the final fidelity was 0.75, suggests that there may be a local trap or plateau in the fidelity landscape close to $\mathcal{I}$ that prevents this from happening.

The *in situ* scheme can do more than just recover known codes and how to implement them in an Ising chain. It is also flexible enough to work with a range of different error models, a sample of our simulations for this is shown in Fig.VI.6. The errors are given as a set $\{p_i\, U_i\}$ which is to be understood as

$$M_{\text{err}}(\rho) = \sum_{i=1}^{N} p_i\; U_i\, \rho\, U_i^\dagger, \tag{VI.6}$$

where there is no requirements on the U_i being local, as in the second line of Fig.VI.6 where they are random. The results can be understood with the following observations. If $N \leq 4$ then the scheme works arbitrarily well, the only limitations to the final fidelity reached being the parameters used in the optimisation. This does not happen in the case

Control Type	Errors	Target	Fidelity
Full Controls	$\frac{1}{4}\mathbb{1}\mathbb{1}\mathbb{1}$, $\frac{1}{4}XXX$, $\frac{1}{4}YYY$, $\frac{1}{4}ZZZ$	$\mathbb{I}$	0.99
Full Controls	$\frac{1}{4}U_1$, $\frac{1}{4}U_2$, $\frac{1}{4}U_3$, $\frac{1}{4}U_4$	$\mathbb{I}$	0.99
Full Controls	$\frac{1}{5}\mathbb{1}\mathbb{1}\mathbb{1}$, $\frac{1}{5}X\mathbb{1}\mathbb{1}$, $\frac{1}{5}\mathbb{1}X\mathbb{1}$, $\frac{1}{5}\mathbb{1}\mathbb{1}X$, $\frac{1}{5}XXX$	$\mathbb{I}$	0.80
Full Controls	$0.22\,\mathbb{1}\mathbb{1}\mathbb{1}$, $0.22\,X\mathbb{1}\mathbb{1}$, $0.12\,\mathbb{1}X\mathbb{1}$,$0.22\,\mathbb{1}\mathbb{1}X$, $0.22\,XXX$	$\mathbb{I}$	0.88

Figure VI.6 | Result of *in situ* optimisation for different error models as defined in Eq.(VI.6). All were done with a fully controllable system where the evolution time T_{evo} is 1, the number of time slots N_{ts} is 14, and with 10 repeats. The fidelity given is, in the first two lines, the target fidelity that was reached 100% and 80% of the time respectively; in the latter two cases it was the maximum fidelity that was reached with most results clustered at or just below it. In the second line, the unitaries are randomly generated.

that $N > 4$, this being the key number here as it is the dimension of the ancilla space. For this case it is useful to only consider orthogonal errors, $\text{Tr}[U_i U_j^\dagger] = 0$ for $i \neq j$ (in order to prevent N being artificial inflated) and to order them such that $p_i \geq p_j$ for $i < j$. The *in situ* optimisation in such a situation will generate a code that protects against the largest four errors such that

$$M_{\text{cor}}(\rho) = \sum_{i=4}^{N} p_i \, U_i \, \rho \, U_i^\dagger, \tag{VI.7}$$

which gives a final fidelity of up to

$$F(M_{\text{cor}}, \mathbb{I}) = 1 - \sum_{i=4}^{N} p_i, \tag{VI.8}$$

as shown in the last line of Fig.VI.6. This is because the code space is not large enough to correct all errors on the system and the quantum capacity [BBDDS96] of the channel M_{err} is lower than one qubit, and therefore maximising the fidelity is done by removing the largest errors. Whether the largest error is a single qubit one or is otherwise correlated does not matter to the *in situ* optimiser. This is exactly the desired behaviour from the scheme as any realistic setting would potentially have any possible errors occurring, most of them with only a very small probability. Being able to automatically identify and correct the N leading errors, dependent on the number of ancillas, is the most that can be asked of the error correcting code.

VI.C.2 Five-Qubit Stabiliser Code

The same *in situ* method can be applied to recover the five-qubit stabiliser code in order to protect a logical qubit against

$$M_{X,Y,Z}(\rho) = \frac{1}{16}\left(\rho + \sum_{i=1}^{5} X_i \, \rho \, X_i^\dagger + Y_i \, \rho \, Y_i^\dagger + Z_i \, \rho \, Z_i^\dagger\right), \tag{VI.9}$$

Control Type	M_{err}	Target	T_{evo}	N_{ts}	N_{upds}	Fidelity	Run Time
Full Controls	$M_{X,Y,Z}$	$\mathbb{I}$	1	80	152	0.99	5 days
Locked Chain	$M_{X,Y,Z}$	$\mathbb{I}$	1	200	168	0.99	6 days
Local Controls	$M_{X,Y,Z}$	$\mathbb{I}$	100π	300	287	0.966	14 days

Figure VI.7 | Preliminary results of recovering the five-qubit stabiliser code using *in situ* optimisation. Due to the considerable run time involved for every optimisation, each line in the table represents a single run of the scheme on one core of a 2.9 GHz Intel Core i7 processor. The algorithm terminated when it reached the target fidelity of 0.99 or when it ran out of time. As in Fig.VI.5, T_{evo} is the total system evolution time, and N_{ts} the number of time slots, to realise both U_{enc} and U_{dec} (split evenly between them), and N_{upds} is the number of iterations where the controls where updated. For comparison, the fidelity without any error correction is 0.8125.

where X_i is an X error on qubit i and identity on the rest, and similarly for Y_i and Z_i. The same three control settings as in §VI.C.1 are investigated numerically for an Ising chain of five qubits in Fig.VI.7. The cost of doing this is, however, significantly greater than for the three-qubit code. While some of this increase is due to the evolution time, number of time slots and number of control updates required by the optimisation algorithm; the dominant factor is the exponential rise in the cost of simulating larger quantum systems classically. This led to a huge increase in the time required for the simulation (for comparison, the time per run for the three-qubit case was under one minute), which has prevented us from getting full results and running experiments with different error models like was done in Fig.VI.6. The results we do have so far, however, shows that *in situ* error correction does work for larger and more complex system. Plans to obtain statistically meaningful results by parallelising aspects of the code and running it on a high-performance cluster are underway.

VI.C.3 Collective Dephasing DFS

The results so far dealt with correcting the effects of an error channel that happened at a single point in the evolution of the system. While this is the simplest scenario to treat both analytically and numerically, our simulations also cover the case of noise acting continuously on the system, including while U_{enc} and U_{dec} are being implemented by the controls. This tests the robustness of the scheme by taking into account the possibility of errors happening during the operations that attempt to protect the logical qubit from such errors.

The model used to study this is two-qubit collective dephasing which has a simple DFS to protect a single logical qubit, as discussed in §VI.A.2. The noise in this case is not given by a single map M_{err} as in the gate model of Fig.VI.4, but by the Lindbladian

$$L(\rho) = -\gamma\,[s_z,\,[s_z,\,\rho]] \qquad \text{where} \quad s_z = \sigma_z \otimes \mathbb{I} + \mathbb{I} \otimes \sigma_z, \tag{VI.10}$$

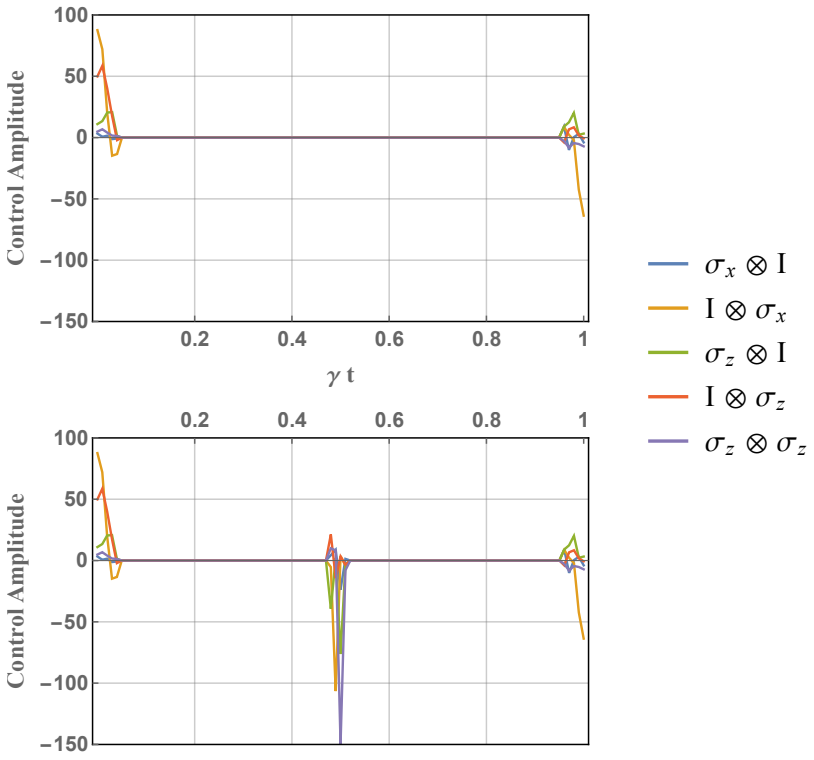
γT_{enc}	γT_{wait}	γT_{dec}	F_{targ}	N_{upds}	p_{succ}
0.05	0.9	0.05	0.90	33 ± 2	98%
0.05	0.9	0.05	0.95	90 ± 7	98%
0.05	0.9	0.05	0.98	260 ± 24	90%
0.05	1.9	0.05	0.95	87 ± 12	92%
0.01	0.98	0.01	0.98	65 ± 2	96%
0.01	0.98	0.01	0.99	170 ± 14	100%
0.01	0.98	0.01	0.995	320 ± 21	82%
0.01	1.98	0.01	0.99	160 ± 12	96%

Figure VI.8 | Parameters in finding a DFS from continuous collective dephasing noise, Eq.(VI.10), *in situ*. The simulations had $N_{\text{ts}} = 10$ time slots split evenly between T_{enc} and T_{dec}, with all Hamiltonians switched off in the waiting period. The control Hamiltonians were σ_x and σ_z on each qubit as well as the $\sigma_z \otimes \sigma_z$ interaction between them. As previously, N_{upds} is the number of iterations where the controls were updated and F_{targ} is the target gate fidelity to the identity map on the carrier qubit. Each row corresponds to 50 repeats and the errors given for N_{upds} are the standard errors. For comparison, the fidelity without any error correction for total $\gamma T = 1$ is 0.68, and for $\gamma T = 2$ is 0.57.

which bears close resemblance to the single qubit dephasing Lindbladian of §I.B.6. As before, the evolution can be divided into three parts. The first, encoding, and the last, decoding, are the same as before except that the collective dephasing Lindbladian is present while the controls are acting. In the middle step all the controls are switched off and the noise acts on its own. The results investigating how the time spent in each of those steps T_{enc}, T_{wait}, and T_{dec} affected the fidelity that could be reached are shown in Fig.VI.8. There are three main observations to be made. Firstly, the high fidelities obtained show that the simulation is able to protect the initial state of the carrier qubit from the noise. Secondly, increasing the speed with which the system is encoded and decoded into the DFS increases the maximum fidelity that can be reached. Lastly, the fidelities that can be attained are unchanged by increasing the waiting time. This is consistent with the *in situ* scheme properly mapping the system in and out of the DFS: the noise has negligible impact during the waiting period and only causes loss of quantum information while the logical qubit is not completely inside of it. The limiting factor is therefore the maximum frequency and amplitude of the control Hamiltonians that can be implemented.

It is also possible to perform a gate on the logical qubit while it is encoded. This is done by taking one of the pulses found by the method described above, plotted in Fig.VI.9a, that has a fidelity of 0.95. This is used as an initial pulse for a second control problem where the target operation is a Hadamard gate, rather than the identity, and the controls are only allowed to vary for a short period of time in the middle of the waiting period. This control problem asks the *in situ* optimisation to perform the gate while the logical qubit is encoded across both physical qubits, as opposed to §VI.C.1 where there was nothing preventing the algorithm from implementing the gate at the end of U_{dec}. The number of time slots available to perform the logical gate is too small to allow the qubit

(a) Control pulse to reach identity



(b) Control pulse to reach Hadamard gate

Figure VI.9 | Plot of control pulse found with *in situ* optimisation that realises a Hadamard gate for a logical qubit encoded in a DFS. The first pulse had $N_{\text{ts}} = 5$ time slots lasting a total of $\gamma T = 0.05$ for both the encoding and decoding steps, a waiting time of $\gamma T = 0.9$, and had a fidelity with the identity map of 0.95. The second pulse was initialised with the first and the controls for the encoding and decoding steps kept fixed, but controls were optimised over $N_{\text{ts}} = 4$ time slots lasting a total of $\gamma T = 0.04$ in the middle waiting time. The target there was a Hadamard gate and had a total fidelity of 0.925.

to be decoded, be acted on by a Hadamard, and re-encoded. This ensures that the pulse shown in Fig.VI.9b does indeed perform a Hadamard gate on a logical qubit inside the DFS, and does so with a total fidelity of 0.925. This drop in fidelity is only half of that caused by the pulse which takes the qubit in and out of the DFS. This shows that *in situ* optimisation, in addition to performing gates on physical qubits and being able to protect them from noise, can also be used to directly find control pulses that realise gates inside the protected space.

VI.D Conclusion

The results shown in this chapter demonstrate that *in situ* optimisation can be readily applied to perform different variants of error correction. Even though there is dissipation

in the system, the only thing being optimised are the control Hamiltonians, no feedback based mechanisms are used; this is advantageous because it keeps the controls completely coherent. The scheme is therefore similar to that used in the previous chapters: it allows gradient based optimisation to be used, and does not add any extra requirement to an experimental implementation. It also means that scaling up the code space in order to correct for more errors does not increase the complexity of evaluating the gate fidelity. However, as a comparison between the three-qubit Shor code and the five-qubit stabiliser code shows, the complexity of the required controls can increase as more ancillas are added. This is a direct consequence of needing to create a specific highly entangled state between many qubits.

The numerical results showed the range of applications of *in situ* error correction. Well known codes were recovered, and automatically adapted by the optimiser to deal with different types of errors on the system. This was without the optimiser knowing beforehand what the errors are, and therefore without any need to characterise the errors in an experimental setting; it simply learnt what controls were needed to prevent those errors from being detrimental. This worked equally well when the noise was on single qubits or acted collectively, thereby removing the need to make such assumptions about the type of dissipation present. That this could be done in different control settings, including with the interaction terms being fixed, show that the approach outlined above may be of real practical use for quantum simulators, and complete the method to upgrade them to quantum computers proposed in §V.

At the time of writing, we are in discussions about an experimental implementation of a DFS in a transmon experiment. The experiment is composed of a cavity which mediates the interaction between two different, off-resonant, transmon qubits [KYGHS07]. The available controls are the microwave field fed into the cavity and the frequency of one the qubits. The noise in the system is partly the finite lifetime of the cavity but is not completely characterised, providing a excellent platform on which to test the ideas of §VI.C.3. There are some difficulties in doing numerical simulations of this process centred around the non-linearity of the controls and the modelling of the cavity, but there is no expectation that those will be insurmountable.

There are also a number of directions in which to expand the ideas put forward here. Among them is pushing to higher dimensional systems such that several logical qubits can be considered at the same time. A particularly interesting case would be to have two separate systems each with a single logical qubit where U_{enc} and U_{dec} are found for each using the methods presented above, and to then use *in situ* optimisation to find an entangling gate between them. Independently, a way to simplify the optimisation may also be to reduce the number of control parameters by using the close relation between U_{enc} and U_{dec}; it may be possible to lower the number of independent control parameters by constraining the two operations.

More broadly, the question of whether this scheme is a scalable and fault tolerant way of dealing with errors needs to be investigated. The two primary points to address are the ability to concatenate the error correcting code, and whether logical gates can be done transversely [NC00, DMN13]. The first point is a way of arbitrarily reducing the errors by nesting the code, specifically, the logical qubit protected by one layer of the code is used as a 'physical' qubit for the next layer of the code. The difficulty of doing this for *in situ* error correction is that each layer will have a different error structure (different M_{err}), and therefore needs to be recalculated by a new optimisation. The number of different layers will scale only logarithmically with the total number of physical qubits used, so the cost of those new optimisations is not expected to be high, but this needs to be assessed rigorously. Transversal gates are another highly desired property of error correcting code as they limit the rate at which errors can spread. A gate between two logical qubits is said to be transversal if it acts identically on every underlying physical qubit. In the *in situ* context described here, the concept can be mapped to having transversal Hamiltonians. The key question is whether a universal gate set can be realised with a given drift Hamiltonian using controls that are transversal between pairs of physical qubits. Lie algebraic methods should go a long way towards answering this question, but numerical simulations will be needed to verify that the optimisation process works when highly symmetric Hamiltonians are used.

CHAPTER VII

CONCLUSION

VII.A Summary

This thesis has covered research we have done on a range of topics on the control of open quantum systems. The two main questions of control theory, "What can we do?" and "How it can be done?", have been applied to noisy quantum dynamics in many different ways, with a common thread.

The story started with §II by identifying which quantum channels could not be reached by a system with a fixed generator of dissipation using any possible control Hamiltonians. As well as being of interest for purely theoretical reasons, it is of use for simulating open quantum systems using an existing source of noise. An alternative way of achieving the same simulations is to use a closed system with purely unitary dynamics, as well as the ability to add and remove ancillary systems. This was studied in §III where it was shown that any analytic family of quantum channels can be replicated with a time-dependent Hamiltonian and a finite ancillary space. In addition to being a novel way to achieve time-continuous simulations, it is also a lift of Stinespring's dilation theorem from the level of propagators to generators as the required conditions for such a Hamiltonian to be continuous and bounded were derived. It was also proved that the dilation can always be well approximated to have those properties. Furthermore, it provides a testing ground to see how adding control Hamiltonians on the microscopic system-bath level can modify the generator for the reduced system. This question, which was a key assumption in §II, was specifically addressed in §IV in the case where the controls and the Lindbladian on the reduced system commute. It was conjectured that in such a case the Hamiltonian does not modify the noise no matter what the underlying system-bath interaction is. Although a definitive proof either way was not forthcoming, it was nevertheless showed that there are some interesting constraints on the dynamics, and that the conjecture does not hold in a specific example of non-Markovian noise.

This need to know how control Hamiltonians affect an open quantum system is a major difficulty when designing control pulses, and one that was circumvented in the last

two chapters. The idea behind this is to use the system for which the controls need to be found to calculate, in tandem with a classical computer, those controls *in situ*. As well as removing the very difficult task of performing accurate process tomography on the quantum system such that it can be modelled, it also removes the exponential computational cost of solving this model numerically. In §V this idea is applied to a quantum simulator and numerical evidence was presented showing that the experimental cost of turning it into a universal quantum computer scales polynomially with the number of qubits. A key step in reaching this scaling is the introduction of the local estimator to the fidelity. This is a bound on the gate fidelity between a general quantum channel and a target unitary which has a tensor product structure; the number of different input-measurement combinations needed to reconstruct it scales only with the dimension of the largest subspace and not with the total system size. This provides an exponential speed up in estimating fidelities when the target is a logical gate. The same fundamental idea was also at work in §VI, where the aim was to protect some quantum information in a many-body system. Numerical simulations of the *in situ* scheme showed that it could 'discover' well known error correcting schemes, as well as new ones tailored to the specific errors on that system. In addition it could also perform logical operations on a qubit's worth of information that was encoded across several physical qubits without taking it out of the protected code space.

VII.B Outlook

As is inevitably the case when doing research, every result that we found raised more new questions than it answered, and those are detailed at the end of each chapter. Taken as whole, however, they point towards one major avenue of future research. Calculating optimal control pulses classically is hard for closed quantum dynamics as the cost of doing numerical simulations scales exponentially with the system size. Doing the same for large open systems that work in a physical realisation seems so much more difficult to be unfeasible. All the problems that exist in the unitary case are still present, and in addition the noise acting on the system has to be characterised. This can already be challenging but it is not enough because, as has been demonstrated in §III and §IV, adding Hamiltonian controls generally will affect the dissipative part of the dynamics. Therefore not only does the noise need to be characterised, so does its response to the possible controls that will be added to it. While existing techniques do work for small systems with relatively few degrees of freedom, it is unlikely that they would generalise well to substantially larger problems. It seems completely natural to us that the way forward is to use quantum systems (such as simulators or, eventually, computers) to do as much of the calculation involved in optimal control as possible. This removes the need to have a classical computer more powerful than the quantum system being optimised; it also, when done *in situ*, bypasses having to acquire an accurate model of the system.

This approach can also be applied more generally than to optimal control. In the same way as classical computers have played an integral part in research for decades, by running models and generally number-crunching, it makes sense that future quantum computers will play a similarly important role. Being able to use such machines to solve problems in quantum information or to design improvements in quantum technology can only be beneficial to the field; the *in situ* approach detailed in this thesis is an example of such a scheme. After all, no one would expect a modern microchip to be designed by hand. It it is therefore natural that developing tools which make use of the special capabilities of quantum technology, in conjunction with traditional classical computers and analytic pen-and-paper work, will provide unique opportunities for future growth and improvement in this area of research.

BIBLIOGRAPHY

[AAAAA13] J. Aasi *et al.* "Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light" *Nat. Photonics* **7**, 613 (2013).

[ABFH17] C. Arenz, D. Burgarth, P. Facchi & R. Hillier. "Dynamical Decoupling of Unbounded Hamiltonians" *arXiv:1704.06143* (2017).

[AGB14] C. Arenz, G. Gualdi & D. Burgarth. "Control of open quantum systems: Case study of the central spin model" *New J. Phys.* **16**, 065023 (2014).

[AHFB15] C. Arenz, R. Hillier, M. Fraas & D. Burgarth. "Distinguishing decoherence from alternative quantum theories by dynamical decoupling" *Phys. Rev. A* **92**, 022102 (2015).

[AKN98] D. Aharonov, A. Kitaev & N. Nisan. "Quantum Circuits with Mixed States" in *Proc. Thirtieth Annu. ACM Symp. Theory Comput.* (Dalas, 1998).

[Ali95] R. Alicki. "Comment on 'reduced dynamics need not be completely positive'" *Phys. Rev. Lett.* **75**, 3020 (1995).

[BBBG09] D. Burgarth, S. Bose, C. Bruder & V. Giovannetti. "Local controllability of quantum networks" *Phys. Rev. A* **79**, 060305(R) (2009).

[BBCFH16] K. Bongs *et al.* "The UK National Quantum Technologies Hub in sensors and metrology (Keynote Paper)". Tech. Rep., University of Nottingham (2016).

[BBDDS96] C. H. Bennett *et al.* "Mixed-state entanglement and quantum error correction" *Phys. Rev. A* **54**, 3824 (1996).

[BCOR09] S. Blanes, F. Casas, J. A. Oteo & J. Ros. "The Magnus expansion and some of its applications" *Phys. Rep.* **470**, 151 (2009).

[BD16] F. Buscemi & N. Datta. "Equivalence between divisibility and monotonic decrease of information in classical and quantum stochastic processes" *Phys. Rev. A* **93**, 012101 (2016).

[BDN12] I. Bloch, J. Dalibard & S. Nascimbène. "Quantum simulations with ultracold quantum gases" *Nat. Phys.* **8**, 267 (2012).

[BF17] D. Burgarth & P. Facchi. "Positive Hamiltonians can give purely exponential decay" *Phys. Rev. A* **96**, 010103(R) (2017).

[Bha97] R. Bhatia. *Matrix Analysis* (Springer, 1997).

[BL57] P. Bocchieri & A. Loinger. "Quantum Recurrence Theorem" *Phys. Rev.* **107**, 337 (1957).

[BLP09] H.-P. Breuer, E.-M. Laine & J. Piilo. "Measure for the Degree of Non-Markovian Behavior of Quantum Processes in Open Systems" *Phys. Rev. Lett.* **103**, 210401 (2009).

[BLPV15] H.-P. Breuer, E.-M. Laine, J. Piilo & B. Vacchini. "Non-Markovian dynamics in open quantum systems" *Rev. Mod. Phys.* **88**, 021002 (2015).

[BM13] B. Bartels & F. Mintert. "Smooth optimal control with Floquet theory" *Phys. Rev. A* **88**, 052315 (2013).

[BP02] H.-P. Breuer & F. Petruccione. *The theory of open quantum systems* (Oxford University Press, 2002).

[BR12] R. Blatt & C. F. Roos. "Quantum simulations with trapped ions" *Nat. Phys.* **8**, 277 (2012).

[BWPRW17] J. D. Biamonte *et al.* "Quantum Machine Learning" *Nature* **549**, 195 (2017).

[BWSH16] V. Bergholm, F. K. Wilhelm & T. Schulte-Herbrüggen. "Arbitrary n-Qubit State Transfer Using Coherent Control and Simplest Switchable Local Noise" *arXiv:1605.06473* (2016).

[BZ06] I. Bengtsson & K. Zyczkowski. *Geometry of Quantum States* (Cambridge University Press, 2006).

[CBK10] J. Clausen, G. Bensky & G. Kurizki. "Bath-optimized minimal-energy protection of quantum operations from decoherence" *Phys. Rev. Lett.* **104**, 040401 (2010).

[CBK12] J. Clausen, G. Bensky & G. Kurizki. "Task-dependent control of open quantum systems" *Phys. Rev. A* **85**, 052105 (2012).

[CCM11] T. Caneva, T. Calarco & S. Montangero. "Chopped random-basis quantum optimization" *Phys. Rev. A* **84**, 022326 (2011).

[CD11] B. Coecke & R. Duncan. "Interacting quantum observables: categorical algebra and diagrammatics" *New J. Phys.* **13**, 043016 (2011).

[CKS17] J. Combes, J. Kerckhoff & M. Sarovar. "The SLH framework for modeling quantum input-output networks" *Adv. Phys. X* **2:3**, 784 (2017).

[CPFSG10] M. Cramer *et al.* "Efficient quantum state tomography" *Nat. Commun.* **1**, 149 (2010).

[CRJP13] J. Cai, A. Retzker, F. Jelezko & M. B. Plenio. "A large-scale quantum simulator on a diamond surface at room temperature" *Nat. Phys.* **9**, 168 (2013).

[Cyw14] L. Cywinski. "Dynamical-decoupling noise spectroscopy at an optimal working point of a qubit" *Phys. Rev. A* **90**, 042307 (2014).

[CZ12] J. I. Cirac & P. Zoller. "Goals and opportunities in quantum simulation" *Nat. Phys.* **8**, 264 (2012).

[D'A08] D. D'Alessandro. *Introduction to Quantum Control and Dynamics* (Taylor and Francis, 2008).

[DBM16] B. Dive, D. Burgarth & F. Mintert. "Isotropy and control of dissipative quantum dynamics" *Phys. Rev. A* **94**, 012119 (2016).

[DCM11] P. Doria, T. Calarco & S. Montangero. "Optimal control technique for many-body quantum dynamics" *Phys. Rev. Lett.* **106**, 190501 (2011).

[DE98] D. Deutsch & A. Ekert. "Quantum Computation" *Phys. World* **11**, 47 (1998).

[dFSGK11] de P. Fouquières, S. G. Schirmer, S. J. Glaser & I. Kuprov. "Second order gradient ascent pulse engineering" *J. Magn. Reson.* **212**, 412 (2011).

[DHJMT00] A. C. Doherty, S. Habib, K. Jacobs, H. Mabuchi & S. M. Tan. "Quantum feedback control and classical control theory" *Phys. Rev. A* **62**, 012105 (2000).

[DHKSH09] G. Dirr, U. Helmke, I. Kurniawan & T. Schulte-Herbrüggen. "Lie-semigroup structures for reachability and control of open quantum systems: kossakowski-lindblad generators form lie wedge to markovian channels" *Reports Math. Phys.* **64**, 93 (2009).

[DiB02] E. DiBenedetto. *Real Analysis* (Springer, 2002).

[DiV00] D. P. DiVincenzo. "The Physical Implementation of Quantum Computation" *arXiv:quant-ph/0002077* (2000).

[Div14] B. Dive. *Quantum Control of Open System.* Mres report, Imperial College London (2014).

[DJ99] A. C. Doherty & K. Jacobs. "Feedback control of quantum systems using continuous state estimation" *Phys. Rev. A* **60**, 2700 (1999).

[DJR14] D. D'Alessandro, E. Jonckheere & R. Romano. "On the Control of Open Quantum Systems in the Weak Coupling Limit" in *21st Int. Symp. Math. Theory Networks Syst.* (Groningen, 2014).

[DL05] D. P. DiVincenzo & D. Loss. "Rigorous Born approximation and beyond for the spin-boson model" *Phys. Rev. B* **71**, 035318 (2005).

[DLCP11] M. P. Da Silva, O. Landon-Cardinal & D. Poulin. "Practical characterization of quantum devices without tomography" *Phys. Rev. Lett.* **107**, 210404 (2011).

[DMB15] B. Dive, F. Mintert & D. Burgarth. "Quantum simulations of dissipative dynamics: Time dependence instead of size" *Phys. Rev. A* **92**, 032111 (2015).

[DMN13] S. J. Devitt, W. J. Munro & K. Nemoto. "Quantum Error Correction for Beginners" *Reports Prog. Phys.* **76**, 076001 (2013).

[DNBT02] J. L. Dodd, M. A. Nielsen, M. J. Bremner & R. T. Thew. "Universal quantum computation and simulation using any entangling Hamiltonian and local unitaries" *Phys. Rev. A* **65**, 040301 (2002).

[DPMB17] B. Dive, A. Pitchford, F. Mintert & D. Burgarth. "In situ upgrade of quantum simulators to universal computers" *arXiv:1701.01723v2* (2017).

[DS78] E. B. Davies & H. Spohn. "Open quantum systems with time-dependent Hamiltonians and their linear response" *J. Stat. Phys.* **19**, 511 (1978).

[DZWKH08] R. Doll, D. Zueco, M. Wubs, S. Kohler & P. Hanggi. "On the conundrum of deriving exact solutions from approximate master equations" *Chem. Phys.* **347**, 243 (2008).

[Ell09] D. L. Elliott. *Bilinear Control System* (Springer, 2009).

[EPJ14] EPJ. *Quantum Technology* (Springer, 2014).

[EPS14] EPSRC. "UK Quantum Technology Hubs" (2014).

[Eur16] European Commission. "Quantum Technologies Flagship" (2016).

[EW14] D. J. Egger & F. K. Wilhelm. "Adaptive hybrid optimal quantum control for imprecisely characterized systems" *Phys. Rev. Lett.* **112**, 240503 (2014).

[Fac16] P. Facchi. *Notes on the Friedrichs-Lee model* (Private Correspondence, 2016).

[FB07] J. Fischer & H.-P. Breuer. "Correlated projection operator approach to non-Markovian dynamics in spin baths" *Phys. Rev. A* **76**, 052119 (2007).

[Fey82] R. P. Feynman. "Simulating Physics with Computers" *Int. J. Theor. Phys.* **21**, 467 (1982).

[FM15] C. Ferrie & O. Moussa. "Robust and efficient in situ quantum control" *Phys. Rev. A* **91**, 052306 (2015).

[Fri85] A. Frigerio. "Covariant Markov Dilations of Quantum Dynamical Semigroups" *Publ. RIMS, Kyoto Uni.* **21**, 657 (1985).

[Fun04] T. C. Fung. "Computation of the matrix exponential and its derivatives by scaling and squaring" *Int. J. Numer. Methods Eng.* **59**, 1273 (2004).

[GAN14] I. M. Georgescu, S. Ashhab & F. Nori. "Quantum simulation" *Rev. Mod. Phys.* **86**, 153 (2014).

[GBCKK15] S. J. Glaser *et al.* "Training Schrodinger's cat: Quantum optimal control: Strategic report on current status, visions and goals for research in Europe" *Eur. Phys. J. D* **69** (2015).

[GKS76] V. Gorini, A. Kossakowski & E. C. G. Sudarshan. "Completely positive dynamical semigroups of N-level systems" *J. Math. Phys.* **17**, 821 (1976).

[GLN05] A. Gilchrist, N. K. Langford & M. A. Nielsen. "Distance measures to compare real and ideal quantum processes" *Phys. Rev. A* **71**, 062310 (2005).

[Gol89] D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning* (Addison-Wesley Longman Publishing, 1989).

[Got97] D. Gottesman. *Stabilizer codes and quantum error correction.* Ph.D. thesis, California Institute of Technology (1997).

[Got07] D. Gottesman. "Fault-Tolerant Quantum Computation" *arXiv:quant-ph/0701112* (2007).

[Gre01] M. Gregoratti. "The Hamiltonian operator associated with some quantum stochastic evolutions" *Commun. Math. Phys.* **222**, 181 (2001).

[GSKB12] M. G. Genoni, A. Serafini, M. S. Kim & D. Burgarth. "Dynamical Recurrence and the Quantum Control of Coupled Oscillators" *Phys. Rev. Lett.* **108**, 150501 (2012).

[HBDRK15] B. Hensen *et al.* "Experimental loophole-free violation of a Bell inequality using entangled electron spins separated by 1.3 km" *Nature* **526**, 682 (2015).

[HDR09] T.-S. Ho, J. Dominy & H. Rabitz. "Landscape of unitary transformations in controlled quantum dynamics" *Phys. Rev. A* **79**, 013422 (2009).

[HP84a] R. L. Hudson & K. R. Parthasarathy. "Communications in Quantum Ito's Formula and Stochastic Evolutions" *Commun. Math. Phys.* **93**, 301 (1984).

[HP84b] R. L. Hudson & K. R. Parthasarathy. "Construction of quantum diffusions" in *Quantum Probab. Appl. to Quantum Theory Irreversible Process.* (Villa Mondragone, 1984).

[HTK12] A. A. Houck, H. E. Türeci & J. Koch. "On-chip quantum simulation with superconducting circuits" *Nat. Phys.* **8**, 292 (2012).

[IBM17] IBM. "IBM Quantum Computing" (2017). URL `http://researchweb.watson.ibm.com/ibm-q/`.

[IOP15] IOP. *Quantum Science and Technology* (IOP Science, 2015).

[Jac14] K. Jacobs. *Quantum Measurement Theory and its Applications* (Cambridge University Press, 2014).

[JCJ14] T. H. Johnson, S. R. Clark & D. Jaksch. "What is a quantum simulator?" *EPJ Quantum Technol.* **1**, 1 (2014).

[JNN13] J. R. Johansson, P. D. Nation & F. Nori. "QuTiP 2: A Python framework for the dynamics of open quantum systems" *Comput. Phys. Commun.* **184**, 1234 (2013).

[JNPGG16] R. Johansson *et al.* "QuTiP-4.0.0" (2016). URL `https://doi.org/10.5281/zenodo.220867`.

[JR92] R. S. Judson & H. Rabitz. "Teaching lasers to control molecules" *Phys. Rev. Lett.* **68**, 1500 (1992).

[JVW09] M. Johanning, A. F. Varón & C. Wunderlich. "Quantum simulations with cold trapped ions" *J. Phys. B* **42**, 154009 (2009).

[Kat80] T. Kato. *Perturbation Theory for Linear Operators* (Springer, 1980).

[KBA05] M. Koniorczyk, V. Bužek & P. Adam. "Simulation of generators of Markovian dynamics on programmable quantum processors" *Eur. Phys. J. D* **37**, 275 (2005).

[KBAL05] M. Koniorczyk, V. Bužek, P. Adam & A. Laszlo. "Simulation of Markovian quantum dynamics on quantum logic networks" *arXiv:quant-ph/0205008* (2005).

[KBCCC14] J. Kelly *et al.* "Optimal quantum control using randomized benchmarking" *Phys. Rev. Lett.* **112**, 240504 (2014).

[KBFMJ16] J. Kelly *et al.* "Scalable in-situ qubit calibration during repetitive error detection" *Phys. Rev. A* **94**, 032321 (2016).

[KBGKE11] M. Kliesch, T. Barthel, C. Gogolin, M. Kastoryano & J. Eisert. "Dissipative quantum Church-Turing theorem" *Phys. Rev. Lett.* **107**, 120501 (2011).

[KL96] E. Knill & R. Laflamme. "Theory of quantum error-correcting codes" *Phys. Rev. A* **55**, 900 (1996).

[KLV00] E. Knill, R. Laflamme & L. Viola. "Theory of Quantum Error Correction for General Noise" *Phys. Rev. Lett.* **84**, 2525 (2000).

[Kos72] A. Kossakowski. "On Quantum Statistical Mechanics of Non-Hamiltonian Systems" *Reports Math. Phys.* **3**, 247 (1972).

[KRKSHG05] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen & S. J. Glaser. "Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms" *J. Magn. Reson.* **172**, 296 (2005).

[KSW08] D. Kretschmann, D. Schlingemann & R. F. Werner. "A Continuity Theorem for Stinespring's Dilation" *J. Funct. Anal.* **255**, 1889 (2008).

[KYGHS07] J. Koch *et al.* "Charge-insensitive qubit design derived from the Cooper pair box" *Phys. Rev. A* **76**, 042319 (2007).

[LB13] D. A. Lidar & T. A. Brun. *Quantum Error Correction* (Cambridge University Press, 2013).

[LBH15] Y. LeCun, Y. Bengio & G. Hinton. "Deep learning" *Nature* **521**, 436 (2015).

[LBRdLM16] H. Labuhn *et al.* "Realizing quantum Ising models in tunable two-dimensional arrays of single Rydberg atoms" *Nature* **534**, 667 (2016).

[Len87] K. Lendi. "Evolution matrix in a coherence vector formulation for quantum Markovian master equations of N-level systems" *J. Phys. A* **20**, 15 (1987).

[LG16] D. Leiner & S. J. Glaser. *Experimental Details* (Private Correspondence, 2016).

[LHNMG11] B. P. Lanyon *et al.* "Universal Digital Quantum SImulation with Trapped Ions" *Science* **334**, 57 (2011).

[Lid12] D. A. Lidar. "Review of Decoherence Free Subspaces, Noiseless Subsystems, and Dynamical Decoupling" *arXiv:1208.5791* (2012).

[Lin76] G. Lindblad. "On the generators of quantum dynamical semigroups" *Commun. Math. Phys.* **48**, 119 (1976).

[LLLKP17] D. Lu *et al.* "Towards quantum supremacy: enhancing quantum control by bootstrapping a quantum processor" *arxiv:1701.01198* (2017).

[Llo95] S. Lloyd. "Almost Any Quantum Logic Gate is Universal" *Phys. Rev. Lett.* **75**, 346 (1995).

[LM13] F. Levi & F. Mintert. "Hierarchies of multipartite entanglement" *Phys. Rev. Lett.* **110**, 150402 (2013).

[LMPZ96] R. Laflamme, C. Miquel, J. P. Paz & W. H. Zurek. "Perfect Quantum Error Correcting Code" *Phys. Rev. Lett.* **77**, 198 (1996).

[LPP13] S. Lorenzo, F. Plastina & M. Paternostro. "Geometrical characterization of non-Markovianity" *Phys. Rev. A* **88**, 020102 (2013).

[LYLSL17] S.-K. Liao *et al.* "Long-distance free-space quantum key distribution in daylight towards inter-satellite communication" *Nat. Photonics* **11**, 509 (2017).

[LYPS17] J. Li, X. Yang, X. H. Peng & C.-P. Sun. "Physical Realization of a Quantum Oracle Machine" *Phys. Rev. Lett.* **118**, 150503 (2017).

[Maa85] H. Maassen. "Quantum Markov processes on Fock space described by integral kernels" in *Quantum Probab. Appl. II Proc. a Work.* (Heidelberg, 1985).

[Mic17] Microsoft. "Station Q" (2017). URL `https://stationq.microsoft.com/`.

[MMHHL16] P. L. McMahon *et al.* "A fully programmable 100-spin coherent Ising machine with all-to-all connections" *Science* **354**, 614 (2016).

[MP14] S. Machnes & M. B. Plenio. "Surprising Interactions of Markovian noise and Coherent Driving" *arXiv:1408.3056* (2014).

[MRL08] M. Mohseni, A. T. Rezakhani & D. A. Lidar. "Quantum-process tomography: Resource analysis of different strategies" *Phys. Rev. A* **77**, 032322 (2008).

[MRNBD17] M. Mohseni *et al.* "Commercialize early quantum technologies" *Nature* **543**, 171 (2017).

[MSGdFG11] S. Machnes *et al.* "Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework" *Phys. Rev. A* **84**, 022305 (2011).

[MSK15] A. S. L. Malabarba, A. J. Short & P. Kammerlander. "Clock-driven quantum thermal engines" *New J. Phys.* **17**, 045027 (2015).

[MTWA15] S. Machnes, D. J. Tannor, F. K. Wilhelm & E. Assémat. "Gradient optimization of analytic controls: the route to high accuracy quantum optimal control" *arXiv:1507.04261* (2015).

[Nak03] M. Nakahara. *Geometry, Topology and Physics* (IOP Publishing, 2003).

[NC00] M. A. Nielsen & I. L. Chuang. *Quantum Computation and Quantum Information* (Cambridge University Press, 2000).

[NDGD06] M. A. Nielsen, M. R. Dowling, M. Gu & A. C. Doherty. "Optimal control, geometry, and quantum computing" *Phys. Rev. A* **73**, 062323 (2006).

[NM65] J. A. Nelder & R. Mead. "A Simplex Method for Function Minimization" *Comput. J.* **7**, 308 (1965).

[NPJ15] NPJ. *Quantum Information* (Nature, 2015).

[OBKRM16] P. J. J. O'Malley *et al.* "Scalable Quantum Simulation of Molecular Energies" *Phys. Rev. X* **6**, 031007 (2016).

[ODSH12] C. O'Meara, G. Dirr & T. Schulte-Herbrüggen. "Illustrating the Geometry of Coherently Controlled Quantum Channels" *IEEE Trans. Automat. Contr.* **57**, 2050 (2012).

[Ope17] Open Journal. "Quantum" (2017). URL `http://quantum-journal.org/`.

[Par92] K. R. Parthasarathy. *An Introduction to Quantum Stochastic Calculus* (Springer, 1992).

[Pas14] S. Pascazio. "All You Ever Wanted to Know About the Quantum Zeno Effect in 70 Minutes" *Open Syst. Inf. Dyn.* **21**, 1440007 (2014).

[PCZ97] J. F. Poyatos, J. I. Cirac & P. Zoller. "Complete Characterization of a Quantum Process: The Two-Bit Quantum Gate" *Phys. Rev. Lett.* **78**, 390 (1997).

[Pec94] P. Pechukas. "Reduced dynamics need not be completely positive" *Phys. Rev. Lett.* **73**, 1060 (1994).

[Pec95] P. Pechukas. "Reply to comment on 'reduced dynamics need not be completely positive'" *Phys. Rev. Lett.* **75**, 3021 (1995).

[PS10] X. H. Peng & D. Suter. "Spin qubits for quantum simulations" *Front. Phys. China* **5**, 1 (2010).

[PS11] A. Pernice & W. T. Strunz. "Decoherence and the nature of system-environment correlations" *Phys. Rev. A* **84**, 062121 (2011).

[PT11] A. N. Pechen & D. J. Tannor. "Are there traps in quantum control landscapes?" *Phys. Rev. Lett.* **106**, 120402 (2011).

[PWZVS16] P. Palittapongarnpim, P. Wittek, E. Zahedinejad, S. Vedaie & B. C. Sanders. "Learning in quantum control: High-dimensional global optimization for noisy quantum dynamics" *Neurocomputing* **0**, 1 (2016).

[Rag02] M. Raginsky. "Entropy production rates of bistochastic strictly contractive quantum channels on a matrix algebra" *J. Phys. A* **35**, L585 (2002).

[RB01] R. Raussendorf & H. J. Briegel. "A One-Way Quantum Computer" *Phys. Rev. Lett.* **86**, 5188 (2001).

[RdVRMK16] H. Rabitz, de R. Vivie-Riedle, M. Motzkus & K. Kompa. "Whither the Future of Controlling Quantum Phenomena?" *Science* **288**, 824 (2016).

[RFZB12] T. Rybár, S. N. Filippov, M. Ziman & V. Bužek. "Simulation of indivisible qubit channels in collision models" *J. Phys. B* **45**, 154006 (2012).

[RH11] Á. Rivas & S. F. Huelga. *Open Quantum Systems. An Introduction* (Springer, 2011).

[RHP14] Á. Rivas, S. F. Huelga & M. B. Plenio. "Quantum Non-Markovianity: Characterization, Quantification and Detection" *Reports Prog. Phys.* **77**, 094001 (2014).

[RKK15] D. M. Reich, N. Katz & C. P. Koch. "Exploiting Non-Markovianity of the Environment for Quantum Control" *Sci. Rep.* **5**, 12430 (2015).

[RRMAG10] C. A. Rodríguez-Rosario, K. Modi & A. Aspuru-Guzik. "Linear assignment maps for correlated system-environment states" *Phys. Rev. A* **81**, 012313 (2010).

[RRW16] B. Russell, H. Rabitz & R.-B. Wu. "Quantum Control Landscape Are Almost Always Trap Free" *arxiv:1608.06198* (2016).

[RSPL16] P. Rebentrost, M. Schuld, F. Petruccione & S. Lloyd. "Quantum gradient descent and Newton's method for constrained polynomial optimization" *arxiv:1612.01789* (2016).

[RSSHW09] P. Rebentrost, I. Serban, T. Schulte-Herbrüggen & F. K. Wilhelm. "Optimal control of a qubit coupled to a non-Markovian environment." *Phys. Rev. Lett.* **102**, 090401 (2009).

[Sch12] K. Schmudgen. *Unbounded Self-adjoint Operators on Hilbert Space* (Springer, 2012).

[SGB14] S. Sauer, C. Gneiting & A. Buchleitner. "Stabilizing entanglement in the presence of local decay processes" *Phys. Rev. A* **89**, 022327 (2014).

[SGGRZ02] D. Stucki, N. Gisin, O. Guinnard, G. Ribordy & H. Zbinden. "Quantum Key Distribution over 67 km with a plug & play system" *Quantum Phys.* **4**, 41 (2002).

[SHMFKL12] T. Schulte-Herbrüggen *et al.* "Control aspects of quantum computing using pure and mixed states." *Philos. Trans. R. Soc. A* **370**, 4651 (2012).

[Sho95] P. W. Shor. "Scheme for reducing decoherence in quantum computer memory" *Phys. Rev. A* **52**, R2493 (1995).

[SKHDD16] A. Smirne, J. Kołodyński, S. F. Huelga & R. Demkowicz-Dobrzański. "Ultimate Precision Limits for Noisy Frequency Estimation" *Phys. Rev. Lett.* **116**, 120801 (2016).

[SS04] S. G. Schirmer & A. I. Solomon. "Constraints on relaxation rates for N-level quantum systems" *Phys. Rev. A* **70**, 022107 (2004).

[SSBCM16] I. A. Silva *et al.* "Observation of time-invariant coherence in a nuclear magnetic resonance quantum simulator" *Phys. Rev. Lett.* **117**, 160402 (2016).

[SSP14] R. Sweke, I. Sinayskiy & F. Petruccione. "Simulation of single-qubit open quantum systems" *Phys. Rev. A* **90**, 022331 (2014).

[Ste96] A. M. Steane. "Simple Quantum Error Correcting Codes" *Phys. Rev. A* **54**, 4741 (1996).

[VKL99] L. Viola, E. Knill & S. Lloyd. "Dynamical Decoupling of Open Quantum Systems" *Phys. Rev. Lett.* **82**, 2417 (1999).

[WBC15] C. J. Wood, J. D. Biamonte & D. G. Cory. "Tensor networks and graphical calculus for open quantum systems Choi-Matrix" *Quant. Inf. Comp.* **15**, 0579 (2015).

[WC08] M. M. M. Wolf & J. I. Cirac. "Dividing quantum channels" *Commun. Math. Phys.* **279**, 147 (2008).

[WD05] H. M. Wiseman & A. C. Doherty. "Optimal unravellings for feedback control in linear quantum systems" *Phys. Rev. Lett.* **94**, 070405 (2005).

[WDKGK17] K. H. Wan, O. Dahlsten, H. Kristjánsson, R. Gardner & M. S. Kim. "Quantum generalisation of feedforward neural networks" *Quantum Inf.* **3** (2017).

[WECC08] M. M. M. Wolf, J. Eisert, T. S. Cubitt & J. I. Cirac. "Assessing Non-Markovian Quantum Dynamics" *Phys. Rev. Lett.* **101**, 150402 (2008).

[XCLMY05] R. X. Xu, P. Cui, X. Q. Li, Y. Mo & Y. Yan. "Exact quantum master equation via the calculus on path integrals" *J. Chem. Phys.* **122**, 041103 (2005).

[XXCLY07] R. X. Xu, J. Xu, P. Cui, B. Li & Y. Yan. "Quantum Dissipation: From Path Integral to Hierarchical Equations of Motion and to Continued Fraction Formalisms" in *Proc. Dyn. Open Quantum Syst. Collab. Comput. Proj. 6* (Daresbury, 2007).

[Yua11] H. Yuan. "Simulation of open quantum dynamics in Markovian environment" *Proc. IEEE Conf. Decis. Control* 5565 (2011).

[ZÁK16] A. Zwick, G. A. Álvarez & G. Kurizki. "Maximizing Information on the Environment by Dynamically Controlled Qubit Probes" *Phys. Rev. Appl.* **5**, 014007 (2016).

[ZŠB05] M. Ziman, P. Štelmachovič & V. Bužek. "Description of Quantum Dynamics of Open Systems Based on Collision-Like Models" *Open Syst. Inf. Dyn.* **12**, 81 (2005).