



Trinity
College
Dublin

The University of Dublin

DOCTORAL THESIS

On the Control of Quantum Many-Body Systems

Author:

Luuk Coopmans

Supervisors:

Dr. Graham Kells

Prof. Stefano Sanvito

*A thesis submitted in fulfillment of the requirements
for the degree of Doctor of Philosophy*

in the

School of Physics

Trinity College Dublin

January 28, 2022

Declaration of Authorship

I, Luuk Coopmans, declare that this thesis titled, On the Control of Quantum Many-Body Systems, has not been submitted as an exercise for a degree at this or any other university and it is entirely my own work.

I agree to deposit this thesis in the University's open access institutional repository or allow the Library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

I consent to the examiner retaining a copy of the thesis beyond the examining period, should they so wish (EU GDPR May 2018).

Signed:

Date:

You keep on learning and learning, and pretty soon you learn something no one has learned before.

R. Feynman

Abstract

Luuk Coopmans

On the Control of Quantum Many-Body Systems

The ability to control and actively manipulate physical systems at a scale at which quantum properties manifest themselves is crucial for the development of quantum technologies. While, at present, it is possible to manufacture small-scale quantum devices, controlling quantum systems that consist of many degrees of freedom remains a formidable task. In this thesis, we look at the problem of stability and control of quantum information in complex quantum many-body systems. We explore two key aspects of this.

The first deals with the manipulation of the information encoded in the degenerate ground-state manifold of systems possessing topological quantum order (TQO). In particular, we exploit Differentiable Programming (∂P) and Natural Evolution Strategies (NES) for the optimal transport of Majorana zero modes in topological p-wave superconductors. These machine learning techniques uncover novel optimal control strategies for Majoranas that are robust with respect to disorder or interactions. Furthermore, we show, using TQO, that topological quantum memories are protected from dynamically generated phase errors caused by small, interaction-driven, energy mismatches between bulk modes. This, in turn, can be used to derive constraints on the bulk energy spectrum of a complex many-body system.

The second aspect concerns the problem of state transfer through a disordered many-body spin chain. We show that ∂P can be efficiently combined with the other quantum control techniques CRAB and shortcuts to adiabaticity. With this hybrid approach, we are able to improve the speed limit for the optimal transport of magnons in a clean Heisenberg model. In addition, in a disordered chain, perfect fidelity transport protocols can be obtained that are robust against fixed, unwanted, realizations of the noise. The fact that this setup can be implemented in a wide range of experimental platforms, makes our results relevant to real-world quantum-state transfer applications.

Publications

1. *Protocol Discovery for the Quantum Control of Majoranas with Differentiable Programming and Natural Evolution Strategies*
L. Coopmans, D. Luo, G. Kells, B. K. Clark and J. Carrasquilla (PRX Quantum 2, 2021)
2. *Dynamical Phase Error in Interacting Topological Quantum Memories*
L. Coopmans, S. Dooley, I. Jubb, K. Kavanagh and G. Kells (Phys. Rev. Research 3, 2021)
3. *Optimal Control in Disordered Quantum Systems*
L. Coopmans, A. Kiely, G. de Chiara and S. Campbell (arXiv preprint: 2201.02029)
4. *Data-Driven Time Propagation of Quantum Systems with Neural Networks*
J. Nelson, L. Coopmans, G. Kells and S. Sanvito (arXiv preprint: 2201.11647)

Presentations and Conferences

1. *Controlling Majorana Zero Modes with Machine Learning*
Poster presentation Sixth International Conference for Young Quantum information Scientists. (12-16 April 2021)
2. *Controlling Majorana Zero Modes with Machine Learning*
Online presentation Journal Club on Quantum Physics and Machine Learning (23 March 2021)
3. *Controlling Majorana Zero Modes with Machine Learning*
Online Seminar Perimeter Quantum Machine Learning Initiative, Waterloo, Canada (2 October 2020)
4. *Controlling Majorana Zero Modes with Machine Learning*
Junior Talk at the ICTP Workshop on Quantum Information Theory and Thermodynamics at the Nanoscale, El Hoceima, Morocco (2 March 2020)

Acknowledgements

Writing and doing the research for this thesis has been a tough, but worthwhile journey for me. There were times that things ran smoothly and everything about this PhD seemed enjoyable and interesting. For example, when visiting conferences and workshops abroad, or during journal clubs and seminars in the DIAS discussion room. On many other occasions, things were less pleasant and I was close to the point of quitting. Especially, I recall one time saying: 'I would rather be a plumber'. Perhaps it is not surprising to say this right after finishing my first paper, while being locked alone at home for almost 4 months. Despite this, I managed to carry through and now I am about to submit this thesis. Here, I would like to thank everybody who contributed to the completion of this journey.

First and foremost, I thank my supervisor Graham, whom I would now describe as the funny, ambitious, and optimistic chap from Laois. Does this match how you described yourself in the DCU interview? Jokes aside, from the first day I felt welcome, and I was impressed by all the knowledge you have. In particular regarding those bloody Majorana nanowires. In addition, you pushed me to greater heights and were the most critical assessor of my work. Without your help, I am sure completing this PhD would have been impossible for me.

I would like to thank my second supervisor, Stefano, for all the years of supporting me at Trinity College Dublin. From day one I have known you as a warm and open person who seems to have everything together. I admire how you are able to manage a research group of about twenty people and still have time to talk to me personally. I enjoyed our conversations about possible career options and hearing your thoughts on the field of quantum computing. You inspired me to take the next step.

This brings me to all the other people I worked with over the last four years. Firstly, my direct colleagues and friends at DIAS. After Graham, Ian you are probably the person who I worked with (and learned from) the most and I enjoyed every minute of it. Then there is Kevin, I think you are the guy with the best humor and I learned a lot from your opinions on food and movies. Shane, you are probably the most typical Irish guy I know and a great (critical) mentor. Aaron, it was great to know (and drink with) you and I hope your girlfriend (Clare) is happy with my old bike. Also Denjoe, Domenico, Werner, Marianne, Tony, Mahul, Takaki, Joost, Charles, Eucharia, Mary, Ronan, Helena, George, Colette, and Pauline, thank you all for making my time at DIAS Burlington road so enjoyable. I will never forget it.

Outside DIAS, I would like to thank some other people as well. Di, Juan, and Bryan it was great working with you on the Majorana game project. I think during this time I developed

myself enormously and I would like to thank you for that. Emma, Jonas, and Giacomo I would like to thank you for my brief random walk into the field of quantum thermodynamics. Emma, Jonas, and Patrick it was also nice hanging out with you in journal clubs and sharing drinks in pubs. Thank you, Steve, Gabriele, and Anthony for our very productive work on the disordered state transfer project. It was fun having zoom meetings with you and your cats. James, it was nice discussing and working with you on machine learning and quantum physics. Lastly, Zakaria, I had fun times with you in Morocco and I enjoyed our discussions about tensor networks and quantum information.

Then there are my office mates and fellow PhD students at Trinity College. Although I spend most of my time at DIAS, I enjoyed my brief visits to the TCD office as well. Thank you very much, Meric, Shardul, Urvesh, and Nina. It was nice sharing coffees and lunches with you. Connor and Marlon, I will never forget spending time with you in an, supposedly, Irish bar in Trieste. Also, Mario, Allesandro, and Rajarshi thank you for helping me out sometimes, and it was great knowing you.

Who else? Oh yes, I also know and met some people outside of physics. In Dublin, I would like to thank my friend Alan. I am very grateful for meeting you and I will never forget our times swimming in the sea and doing self-development challenges. I also like to thank my friends Daniel, Marise, and Patrick for our fun times hanging out in Dublin. Although I sucked at it, I enjoyed doing Hacky Sack and board games with you. Ace, you are a very wise and passionate person and it was fun talking to you, for hours, about technological developments in Stephen's Green park. Good luck with the development of your machine learning poker bot. Then Michael, Yasmeen, Maysoun, Michelle, Tyra, Chay, and Adelin thank you for all the dinners, walks and drinks we shared.

I am sure there are many more people that I need to thank. However, finally before I wrap up, I wish to thank my family and close friends in my home country, the Netherlands, who have supported me for a long time. This includes my friends Robbin, Sjoerd, Nick, Mike, Wessel, Krijn, Peter, and Jasper, who came to visit me during my stay in Dublin. We shared many drinks and good stories. Also, I am grateful for zoom calls with many other Dutch friends during the Covid '19 lockdown. This definitely helped to keep me sane during this difficult time. Oh yeah, Harri, thanks for the chicken and our trips to Eindhoven city. Then, I would like to thank my uncles, aunties, and cousins for showing interest in my physics work, despite me probably sounding gibberish. Lastly, to my parents, mam en pap, and to my sister, Aafke, I have no words left to express how grateful I am for having you in my life. Without you I would have given up on this journey a long time ago.

Contents

Declaration of Authorship	iii
Abstract	vii
Acknowledgements	xi
1 Introduction	1
1.1 Overview of the Thesis	4
2 Quantum Control and Techniques for Numerical Optimisation	7
2.1 Introduction and Formulation of the Central Problem	7
2.1.1 The Control of a Single Qubit	11
2.2 Traditional Methods for Quantum Control	13
2.2.1 Pontryagin’s Principle and Bang-Bang Protocols	14
2.2.2 Shortcuts to Adiabaticity	15
2.2.3 Numerical Techniques	18
Simulated Annealing	18
Chopped Random Basis	20
2.3 Machine Learning for Quantum Control	20
2.3.1 Reinforcement Learning	21
2.3.2 Differentiable Programming	22
Gradient Descent and Single Qubit Control	24
2.3.3 Natural Evolution Strategies	26
2.4 Summary	28
3 Topological Quantum Computation and Majorana Zero Modes	29
3.1 Quantum Computation by Anyons	30
3.2 Topological Superconductivity and Majorana Zero Modes	32
3.2.1 The Kitaev Chain	33
3.2.2 The proximity-coupled Superconductor	36
Model and Dispersion	37
The Kitaev Chain Limit and Majorana Zero Modes	38
3.2.3 Experimental Implementation	41

3.3	Moving Majorana Zero Modes	42
3.3.1	Moving Frame	44
3.3.2	Susceptibility to boundary oscillations	45
3.4	Summary	47
4	Optimising Transport of Majorana Zero Modes	49
4.1	Setup of the Optimisation Problem and the Different Majorana Control Regimes	51
4.2	Numerical Optimisation Results	53
4.2.1	Machine Learned Majorana Transport Protocols	54
4.2.2	Simulated Annealing and Benchmark Protocols	57
4.2.3	Extension to the Proximity Coupled Superconductor	59
4.3	The Jump-move-Jump Control Strategy	61
4.3.1	An Expression for the Infidelity of the Bare JMJ Control Strategy . .	62
4.3.2	The Dressed JMJ Control Strategy	65
4.4	The Effect of Disorder	66
4.5	Summary	69
5	The Effect of Repulsive Interactions on the Control of Majorana Zero Modes	71
5.1	Topological Order and Quasiadiabatic Continuation	73
5.2	Spectrum of the Interacting Kitaev Chain	78
5.2.1	Interaction Driven Bulk Energy Splittings	78
5.2.2	Interaction Driven Reduction in the Topological Energy Gap	81
5.3	Robustness of Majorana transport with respect to interactions	82
5.4	Suppression of the Interaction Driven Phase Error	86
5.4.1	Analytical Argument for <i>no Phase Error</i>	87
5.4.2	Numerical Verification	88
5.5	Summary and Further Work	90
6	Optimising Magnon Transport in Disordered Spin Chains	93
6.1	The Spin Model and Optimisation Objective	95
6.1.1	Magnons and the Optimisation Objective	97
6.2	Ansätze for the Optimal Protocols	99
6.2.1	Shortcut to Adiabaticity Inspired Protocol Ansatz	100
6.2.2	Fourier Series Ansatz	103
6.3	Optimisation Results for the Clean Spin Chain	104
6.3.1	Tightening the Heuristic Speed Limit with ∂P	107
6.4	The Effects of Disorder	108
6.4.1	Two Types of Disorder and Robustness of the STA Protocols	109
6.4.2	Optimisation in the Presence of Disorder	112

6.5	Summary	115
7	Conclusion and Future Directions	119
A	The Bogoliubov-de-Gennes formalism	123
A.1	General Solution of the Quadratic BdG Hamiltonian	123
A.2	BdG Time Evolution	124
B	Equivalence between the Interacting Kitaev Chain and XYZ Heisenberg Spin Chain	125
C	Tensor Networks	127
C.1	Matrix Product States	127
C.2	Density Matrix Renormalization Group	128
C.3	Time Evolution with MPS	129
D	Topological Order of the Non-Interacting Kitaev Chain	131
D.0.1	Groundstate Degeneracy as a Consequence of Topological Order . . .	132

Chapter 1

Introduction

Devices that harness the laws of quantum mechanics are currently being explored as a way to advance the field of information and communications technology. This technological revolution opens up new opportunities for researchers across the whole scientific community [1]. This is because quantum devices make use of properties such as the superposition principle, entanglement and quantum parallelism [2], which give them a so-called quantum advantage [3] and the potential to vastly outperform their classical counterparts. These advantages can be exploited [4] in a wide range of applications, including the simulation of complex many-body systems [5, 6], the improvement of artificial intelligence [7, 8], and also the enhancement of security systems with quantum cryptography [9, 10, 11].

One essential ingredient for the development of quantum technologies is the ability to manipulate and control physical systems at a scale at which quantum properties manifest themselves. This is also known as the area of quantum control. Standard experimental platforms for this are superconducting circuits [12], trapped ions [13, 14, 15], and nitrogen-vacancy centres in diamond [16, 17]. These platforms encode qubits, the quantum equivalent of classical bits, which can be manipulated by adjusting magnetic fields or laser pulses. One of the main questions in quantum control is how the control fields need to be tuned in time in order to realise a certain objective or task. This includes steering the qubit from one specific starting state to the desired target state or stabilising the information encoded in a quantum memory.

However, finding the specific shapes of the protocols is a huge challenge in the field of quantum control. In most cases, one not only needs to solve and have access to the full microscopic dynamics of the system but also needs to optimise the control pulses, which the dynamics implicitly depend on [18]. In addition, there are extra complications involved when one tries to model real-world devices, which cannot be perfectly isolated from their environment. Imperfections in the setup, noise sources nearby, and temperature fluctuations lead to qubit errors and the decoherence of the information encoded within them [19, 20].

Despite these complications, there is by now a large set of different techniques and optimisation methods available to explore the complicated landscape of quantum control. This includes techniques that have their roots in classical mechanics [18] and mathematical optimisation [21]. For example, the extension of the calculus of variations by Lev Pontryagin [22].

This method has been applied in a few different quantum control settings, such as quantum metrology [23] and variational quantum algorithms [24]. Another example is the now widely known Monte Carlo, or simulated annealing method [25, 26], which has found numerous applications in the optimisation of quantum systems. In addition, one can exploit techniques specifically designed for quantum control. Most notably, the shortcut to adiabaticity approaches, which have been successfully implemented in a range of different experimental platforms [27].

In this same vein, a more recently proposed approach is to control quantum systems and devices with Machine Learning (ML) techniques. In ML one trains a machine to perform a certain task without giving it an explicit set of instructions. The machine either learns from its past experiences, or from a set of examples (inputs and outputs) until it is able to do the task as best as it possibly can. This approach has gained widespread attention due to breakthroughs in natural language processing [28], computer vision [29], and genomics [30]. With respect to quantum physics [31, 32], it has been used for the representation of quantum states [33], the classification of phases of matter [34], and the enhancement of density functional theory [35], to name just a few. For quantum control, a specific ML technique known as Reinforcement Learning (RL) has been studied in great depth in a growing number of papers [36, 37, 38]. Here the control problem is formulated as a game that is played by an agent (machine) who, after much trial and error, learns to maximise a reward that encodes the control objective.

Given their success, it is natural to try to apply these techniques to control quantum systems that consist of a large number of degrees of freedom. The study of such many-body systems [39, 40] is becoming increasingly important, as the real-world quantum devices that are being built grow ever larger [41]. In particular, to reach quantum supremacy, one needs to have the ability to use devices composed of many, $\mathcal{O}(10^2)$, qubits [42]. However, on a fundamental level, many properties of quantum many-body systems are not well understood. Moreover, the simulation of them,¹ required for the application of some of the optimisation methods, is impossible with the current classical architectures. For these reasons, the control of these systems remains a challenging and relatively unexplored field.²

In this thesis, we aim to further investigate this topic and look at the problem of stability and control of quantum information in complex many-body systems. We will focus on two key aspects of this. The first concerns the control of the information encoded in the ground state of materials possessing topological quantum order (TQO). TQO is a special phase of matter that cannot be described by Landau's symmetry-breaking theory [43]. One of the main, and earliest, examples of this is the Fractional Quantum Hall State [44, 45, 46, 47], which was discovered experimentally in the early '80s. However, the name, and mathematical definition,

¹Consisting of at least $\mathcal{O}(10^2)$ qubits.

²We note that while ML has been successfully applied to some fundamental problems involving quantum many-body systems [33], in the context of quantum control, most studies are concentrated on systems involving up to a maximum of around 10 qubits.

of *topological order* was formulated a few years later, by Wen, in the context of Chiral Spin States [48].

About two decades ago, it was proposed to use these materials as a way to overcome the problem of decoherence in the quest for practical quantum devices [49, 50]. This proposal for topological quantum computation is based on the use of (non-abelian) anyons. Anyons are special quasiparticle excitations with exotic braiding statistics, which can be found in systems possessing the TQO property. Associated with the (non-abelian) anyons, these systems have a (robust) ground-state degeneracy, which can be used to encode and manipulate quantum information. The advantage of this scheme is that the anyonic quasiparticles can be moved apart, such that the information is encoded non-locally (topologically), and hence protected from local noise and error processes.

The particular topological system we will focus on is a one-dimensional p-wave superconductor [51]. This many-body system can be effectively realised in the laboratory, by proximity coupling a conventional s-wave superconductor to a spin-orbit coupled semiconducting nanowire [52, 53, 54, 55, 56].³ The specific anyonic quasiparticles in this model are known as Majorana zero modes. We will investigate the optimal control of these anyons, by utilising some of the ML optimisation methods described above. In addition, we will look at the robustness of Majorana based topological quantum memories with respect to screened electron-electron interactions. For this we make use of advanced tensor network simulation techniques [59, 60, 61, 62].

The second key aspect of the control of quantum many-body systems which we consider in this thesis is somewhat less esoteric and concerns the problem of state transfer through a many-body spin chain. For this, we make use of the XXX-Heisenberg model [63], which, via a Jordan-Wigner transformation, can also be interpreted as a fermionic many-body system. This model can be realised in a wide range of experimental setups [64, 65, 66, 67], and is often used to study magnetic properties, quantum phase transitions [68], and the spreading of (quantum) correlations [69].

Our aim is to find optimal control protocols for the transport of a localised spin excitation (Magnon) along the disordered Heisenberg spin chain. Similar to the Majorana control problem, we will apply some of the ML optimisation methods. However, this time, we will combine them with other approaches for quantum control, in particular the shortcuts to adiabaticity [27, 70]. This hybrid approach allows us to develop control protocols for quantum state transfer, which can be potentially implemented in real-world versions of this complex quantum many-body system. Our approach might be particularly relevant for the field of quantum communication [71, 72], where one studies the transmission of quantum information over large distances.

³However, we note that, while possible, the creation of these systems is extremely difficult in practice. In particular, the search for anyons has been hampered by the presence of disorder and other imperfections in these proximity-coupled setups [57, 58].

To end this introduction, in the following section, we provide a detailed outline of all the questions and problems we aim to address in this thesis.

1.1 Overview of the Thesis

Chapter 2: We will start the thesis by introducing the topic of quantum control and the (mathematical) formulation of the central problem. We will describe a simple example problem involving a single qubit, for which we introduce some of the basics of quantum information science and decoherence. Afterwards, we give an overview of the different traditional control techniques, including shortcuts to adiabaticity, simulated annealing and Pontryagin's maximum principle. In the end, we will briefly introduce machine learning, in particular reinforcement learning, and discuss in detail the Differentiable Programming (∂P) paradigm and the Natural Evolution Strategies (NES) algorithm. To make the discussion easier to follow, we will apply some of these techniques to the single-qubit example and show that they are able to come up with efficient control strategies.

Chapter 3: An introduction to the field of topological quantum computation and non-abelian anyons will be given. This is followed by a discussion of the most promising experimental material for the realisation, a topological p-wave superconductor. Afterwards, we will look at the fundamental excitations of this system, in particular the Majorana zero modes. Lastly, we will look at how to move the Majoranas by imposing and tuning an external potential profile.

Chapter 4: This chapter is based on the work published in Ref. [73], and looks at the control of Majorana Zero modes with the machine learning algorithms ∂P and NES. Here, we will first introduce four different Majorana motion control regimes, which connect with the critical velocity and resonance timescale present within the system. Then, we will apply the ∂P and NES optimisation algorithms in each of these four regimes and show that they are able to come up with novel counterintuitive Majorana transport strategies. In the end, these strategies will be analysed and their robustness with respect to disorder will be tested.

Chapter 5: We will introduce an interacting term in the p-wave model, and investigate its effects on the stability and control of Majorana zero modes. First, we will show that the interacting model inherits the TQO property from the non-interacting model, which implies that the interacting ground state is still degenerate. Then, we show that in the interacting model the ground-state degeneracy does not carry over to the excited states, and small bulk energy splittings occur. Afterwards, we will numerically simulate this system and look at the robustness of the Majorana transport strategies obtained in chapter 4. Finally, we show analytically that the broken degeneracy between the excited states does not lead to any additional phase error in interacting topological quantum memories. This result can then be used to derive constraints on the many-body spectrum. Some of these results were published in Ref. [74]

Chapter 6: Here, we will look at the optimal transport of magnons through a Heisenberg spin chain. We will start by introducing this model and the optimal control optimisation problem. We then show that a family of approximately perfect transport protocols can be obtained by exploiting an inverse engineering method based on Lewis-Riesenfeld invariants. We combine this analytical, shortcut to adiabaticity, method with ∂P to find the best protocol within this family. In addition, we use another hybrid optimisation method in which we combine ∂P with a Fourier Series ansatz. With this method, we show that we are able to improve the heuristic speed limit for the Magnons in the clean spin chain. Lastly, we perform optimisations in the disordered spin chain and show that for single fixed disorder patterns we are able to obtain perfect transport protocols.

Chapter 7: We end the thesis by giving a conclusion and putting our results in a broader context. In addition, we will give an overview of possible directions to explore for future work.

Chapter 2

Quantum Control and Techniques for Numerical Optimisation

Various strategies and methods have been developed to control and manipulate quantum information encoded in a range of different experimental platforms. This includes, for example, more traditional approaches such as the analytical optimal control method by Pontryagin [22] or the numerical optimisation algorithm named after Krotov [75]. In addition, more recently, novel methods for the control of quantum systems that make use of machine learning have been introduced. These are largely based upon algorithms for Reinforcement Learning (RL) in which an agent is trained by interacting with the problem environment such as in the Q-learning [76] or the Policy Gradient algorithms [77]. In this chapter, we will review some of these methods while providing a general introduction into the field of quantum control.

We have structured this chapter as follows. We begin by giving the necessary background of optimal control theory and the formulation of the central problem. As a simple example, we consider the control of a single qubit, which forms the fundamental building block of many quantum devices. This example serves as warm-up exercise for the control of the more complicated quantum many-body systems that we will encounter in the next chapters. Afterwards, we first introduce the traditional quantum control methods in section 2.2, before introducing the machine learning methods in section 2.3. In particular, we will discuss in detail the paradigm of Differentiable Programming and the numerical optimisation technique Natural Evolutionary Strategies, which form an important part of the rest of this thesis.

2.1 Introduction and Formulation of the Central Problem

Optimal control is a subfield of mathematical optimisation that aims to construct control strategies for various systems and dynamical problems [18]. In its most generic and abstract form, the optimal control problem is formulated via some objective function $\mathcal{L}(x, \dot{x}, t, \theta)$ that

depends on the system variable $x \equiv x(\theta, t)$, which depend on time t and the control parameters θ^1 , together with a possible set of boundary (initial) conditions $\{x_0, x_f\}$. The control problem is said to be solved when a specific optimal protocol θ_{opt} is found that minimises the objective function \mathcal{L} . We will now discuss each of the elements in this generic control problem in detail and show that quantum control problems are just a specific case of optimal control theory.

Firstly, the parameters x that describe the state of the system to be controlled and its time evolution \dot{x} are defined to satisfy some differential equation,²

$$\dot{x} = f(x, \theta), \quad (2.1)$$

that depends on the specific dynamical system under consideration. For example, when x describes the position of a simple pendulum it could be Newton's equations of motion, or, as we will see below, when we look at a quantum system, $x \equiv |\psi\rangle$ and f is the Schrödinger equation. A trajectory $x(t)$ of the system that solves this differential equation always starts in the initial configuration of the system x_0 and, provided one is given, ends in the final configuration x_f . The ability to reconstruct this trajectory from experimental measurements³ is known as the *observability* of the control problem. Specifically, the control system is said to be observable if we can reconstruct any particular initial condition x_0 from a series of measurements in a time interval $[0, t]$ [18].

The second main ingredient of the generic control problem is the collection of control parameters θ . The control parameters can be seen as the experimental control knobs, which can be tuned in time by the controller of the system. For example, an experimentalist changing a magnetic field or applying a driving force when the system is a pendulum. Importantly, in this control setup the differential equation f depends directly on the parameters θ , which therefore drastically influence the evolution of the system. The use of one particular control θ_1 can cause the system to follow a completely different trajectory $x_1(t)$ compared to the trajectory $x_2(t)$ that results from a different control $\theta_2 \neq \theta_1$.⁴ In this way the controller can influence the evolution of the system and potentially steer it to the termination state one is interested in.

An important concept related to this is the *admissibility* of the control protocol $\theta(t)$. More precisely, the controller cannot arbitrarily tune the parameters θ , since one is limited by real-world experimental constraints, such as finite rates of change or maximum amplitudes of the

¹For simplicity we write the control parameter here as a single symbol θ . In general, it can represent a single parameter, a vector, a matrix or even a tensor of control parameters. In addition, note that it is a function of time t .

²In full generality it could also be a set of (un)coupled differential equations.

³Or via any other way.

⁴Think, for example, of the difference between applying a driving force to a pendulum in the direction of its motion or applying a driving force against it.

control fields. This limits the set of evolution trajectories, $x(t)$, that the system can follow and the system is potentially not able to access any arbitrary state x . This means that in some cases a desired target state x_T can not be obtained during the evolution no matter what control is applied. In the control literature, this is known as the system not being *controllable*, as the control problem can not be solved [78]. To step ahead and make a connection with some of the concepts we will see later on in this thesis, speed limits, we note that the total evolution time τ of the system severely restricts the set of possible states the system can access.

The final element of the control problem is the cost functional \mathcal{L} in which the control objective or the control task is encoded. In general, the objective function can be written as a sum of three parts,

$$\mathcal{L} = \int_0^\tau L(x(t), \theta(t), t) dt + g(x(\tau)) + p(\theta). \quad (2.2)$$

The first term, L , corresponds to the running payoff, which is integrated over time.⁵ The second term, g , is the terminal loss, which only depends on the final state $x(\tau)$. The final term, p , includes potential penalties on using certain control parameters.⁶ An example of the running payoff is the amount of fuel used in the control problem of trying to land a spacecraft on the moon [78]. In this particular problem, the terminal loss could be, for example, the height of the spacecraft above the moon's surface and $p(\theta)$ a constraint on the maximum acceleration of the spacecraft. Similarly, for each particular control problem the specific contributions to the objective function \mathcal{L} need to be defined.⁷

The control objective is encoded in such a way that the higher the value of \mathcal{L} the further away the controller is from its goal. This means that the minimisation of \mathcal{L} with respect to θ is the mathematical formulation of the control problem. This can be written as

$$\min_{\theta \in \theta_{\text{ad}}} \mathcal{L}(\theta), \quad (2.3)$$

where θ_{ad} are all admissible control protocols. The optimal control θ_{opt} that solves this minimisation problem is the protocol that satisfies

$$\mathcal{L}(\theta_{\text{opt}}) < \mathcal{L}(\theta_{\text{nopt}}). \quad (2.4)$$

Here, θ_{nopt} are all the parameterisations of the control that are not optimal $\theta_{\text{nopt}} \neq \theta_{\text{opt}}$. Note that within this condition for optimality we can have multiple optimal protocols $\theta_{\text{opt}} \neq \theta'_{\text{opt}}$ with $\mathcal{L}(\theta_{\text{opt}}) = \mathcal{L}(\theta'_{\text{opt}}) < \mathcal{L}(\theta_{\text{nopt}})$.

⁵Note that L can depend on the derivative \dot{x} . However, the differential equation Eq. 2.1 can be used to write this a function of x .

⁶Such as excessively large fields. Note that this term is, in the literature, sometimes incorporated into the definition of L .

⁷In this thesis we will mainly concern ourselves with objective functions that only depend on the terminal state $g(x(\tau))$. However, the methods that we use and introduce below can be applied to the most generic cost function, which includes a running payoff or a penalty term as well.




FIGURE 2.1: One-dimensional example of a convex $\mathcal{L}_1(\theta)$ (a) versus a non-convex $\mathcal{L}_2(\theta)$ (b) optimisation problem. For problem (a) any line connecting two arbitrary controls θ_1 and θ_2 is always above the graph while for problem (b) there are controls θ_1 and θ_2 for which the connecting line lies under the graph as indicated by the black dashed lines. In (a) we only have one minimum which is the global minimum (black dot) whereas in (b) we have a suboptimal local minimum together with a global minimum (black dot). Applying gradient descent (GD) to problem (a) will always converge to the global minimum no matter where is started from (red arrows). Applying GD to (b), however, is completely different and might not converge to the optimum (green arrow) depending on where one starts the optimisation from.

An important aspect of this mathematical minimisation problem is the shape of the control landscape⁸ $\mathcal{L}(\theta)$. In particular, whether the control landscape is convex or not, as shown for a one-dimensional example in Fig. 2.1. We say $\mathcal{L}(\theta)$ is convex, if for all possible controls θ_1 , θ_2 , and all $\lambda \in [0, 1]$, we have

$$\mathcal{L}(\lambda\theta_1 + (1 - \lambda)\theta_2) \leq \lambda\mathcal{L}(\theta_1) + (1 - \lambda)\mathcal{L}(\theta_2) \quad (2.5)$$

Geometrically this means that for a convex landscape $\mathcal{L}(\theta)$ any line between two different controls θ_1 and θ_2 is above the graph $\mathcal{L}(\theta)$ [Fig. 2.1 (a)], while for a non-convex landscape one can pick controls connected by a line that lies under the graph [Fig. 2.1 (b)]. Another way to investigate if $\mathcal{L}(\theta)$ is convex is by checking if $\frac{d^2\mathcal{L}}{d\theta^2} \geq 0$.⁹

The convexity of $\mathcal{L}(\theta)$ is related to the type and number of extrema in the control landscape. Crucially, for a convex $\mathcal{L}(\theta)$ we have only one global minimum, which means that we can find the optimal control¹⁰ by solving the equation

$$\left. \frac{d\mathcal{L}}{d\theta} \right|_{\theta_{\text{opt}}} = 0. \quad (2.6)$$

For non-convex landscapes, we have several extrema and, hence, several solutions of this first derivative condition. However, only one of these is the global optimum. In other words, a

⁸We say landscape since θ can be multidimensional.

⁹Which generalises to the Hessian for a multidimensional θ .

¹⁰To keep the discussion simple we will focus on the case in which all possible θ are admissible.

non-convex optimisation problem has a rough control landscape $\mathcal{L}(\theta)$ with possibly many suboptimal minima, whereas a convex optimisation problem has only one optimal minimum. Identifying these convexity properties of the optimal control problem is important for determining which optimisation techniques to use.¹¹

So far the discussion has been completely general and it applies to any optimal control optimisation problem. Practical examples include minimising the amount of fuel used when trying to land a spacecraft on the moon (described above) or balancing production against consumption in a food factory. Importantly, the generality of this formulation allows us to fit quantum control problems within this framework as well. In this case, the system to be controlled is a quantum system and its state x is described by the wave function $|\psi(\theta, t)\rangle$ or the density matrix $\rho(\theta, t)$. Typically the corresponding differential equation f is the Schrödinger equation,¹²

$$i \frac{\partial}{\partial t} |\psi(\theta, t)\rangle = \mathcal{H}(\theta(t)) |\psi(\theta, t)\rangle, \quad (2.7)$$

with Hamiltonian $\mathcal{H}(\theta(t))$ for a closed quantum system¹³ or the Lindblad equation¹⁴ for an open quantum system. The control parameters θ are problem-specific, but in most cases are again the experimental fields that can be applied to the setup. The objective function \mathcal{L} has the same form as before in Eq. 2.2 and depends on the specific quantum control task at hand.

Instead of setting up the quantum control problem and objective function \mathcal{L} generally¹⁵ in the next subsection we will introduce a specific example of preparing a single qubit. This will serve as a starter for the more complicated quantum many-body control problems in later chapters. This example also allows us to introduce some of the basic concepts of quantum information.

2.1.1 The Control of a Single Qubit

Quantum processors are devices designed to perform computations, i.e. transform given inputs into desired outputs. In these devices, information is encoded within basic units called qubits, which are analogous to the logical bits in classical computers. In addition to the classical logical 0 and 1 states, qubits can also be in any linear combination of these due to the superposition principle. This can be mathematically expressed as

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle \quad (2.8)$$

¹¹For a more detailed mathematical discussion about convexity and optimisation we refer the reader to reference [79].

¹²We set $\hbar = 1$ in this thesis.

¹³Or the Dirac equation when one is looking at relativistic systems.

¹⁴This equation will be discussed below.

¹⁵We could argue that the most general formulation is the formulation already given for the general optimal control problem.

where α and β are complex numbers that satisfy the normalisation condition $|\alpha|^2 + |\beta|^2 = 1$. The information encoded within these qubits can be manipulated with special quantum gates, which each can be represented by a specific unitary matrix U that acts on the state $|\psi\rangle$ and changes the numbers α and β in the superposition. At the end of a computation the qubit state $|\psi\rangle$ is measured to determine the output.

An important (control) problem in this setup is finding the right set of unitary matrices U that bring a qubit state from a specific initial state $|\psi_0\rangle \equiv |\psi(0)\rangle$ to a target state $|\psi_\tau\rangle$ after some fixed termination time τ . A natural choice of the objective function for this task is the state infidelity function, \mathcal{I}_τ , defined as,

$$\mathcal{I}_\tau \equiv 1 - \mathcal{F}_\tau = 1 - |\langle \psi_\tau | U(\tau) | \psi_0 \rangle|^2, \quad (2.9)$$

where

$$U(\tau) = \mathcal{T}e^{-i \int_0^\tau \mathcal{H}(\theta(t)) dt} \quad (2.10)$$

is now the time-ordered unitary operator that solves the time-dependent Schrödinger Equation (2.7). \mathcal{H} is the Hamiltonian of the system, which depends on the control parameters θ . This infidelity measure is exactly zero when the target state is reached and is exactly one when it is orthogonal to it. Minimising \mathcal{I}_τ thus corresponds to maximising the overlap between the time evolved state $|\psi(\tau)\rangle$ and the desired target state $|\psi_\tau\rangle$.

The form of $\mathcal{H}(\theta)$ depends on the specific experimental platform in which the qubit is realised. There are various platforms for this, for example in the next chapter we will look at topological qubits realised in the low energies degrees of freedom of a complex many-body system. For the moment we will focus on a "simple" single spin-1/2 system in an external magnetic field as shown schematically in Fig. 2.2. The Hamiltonian for this quantum setup is given by

$$\mathcal{H}[\theta(t)] = \theta_x(t)\sigma^x + \theta_y(t)\sigma^y + \theta_z(t)\sigma^z, \quad (2.11)$$

where σ^i for $i = x, y, z$ are the Pauli spin matrices and the control field is the (three-dimensional) external magnetic field $\theta(t) \equiv [\theta_x(t), \theta_y(t), \theta_z(t)]$. The objective of the controller is to change this magnetic field $\theta(t)$ in such a way that an initial state $|\psi_0\rangle$ on the Bloch sphere (Fig. 2.2) is rotated over time to the target state $|\psi_\tau\rangle$. When we fix $\theta_z(t) = 0$ and $\theta_y(t) = 1$ this setup describes a single transverse-field Ising spin system, which is often referred to as the *two-level* system. This system is a paradigmatic example for many phenomena in condensed matter physics and its control has been extensively studied [80, 81].

In a real-world setup, there is always some additional complication in this quantum control problem due to an unavoidable coupling with the environment. Noise, such as external magnetic or electric fields from sources nearby the setup, can easily corrupt the setup, and hence the control process of the qubit. In quantum information, this is known as a decoherence process in which the information encoded in the qubit gets erased over time due to the

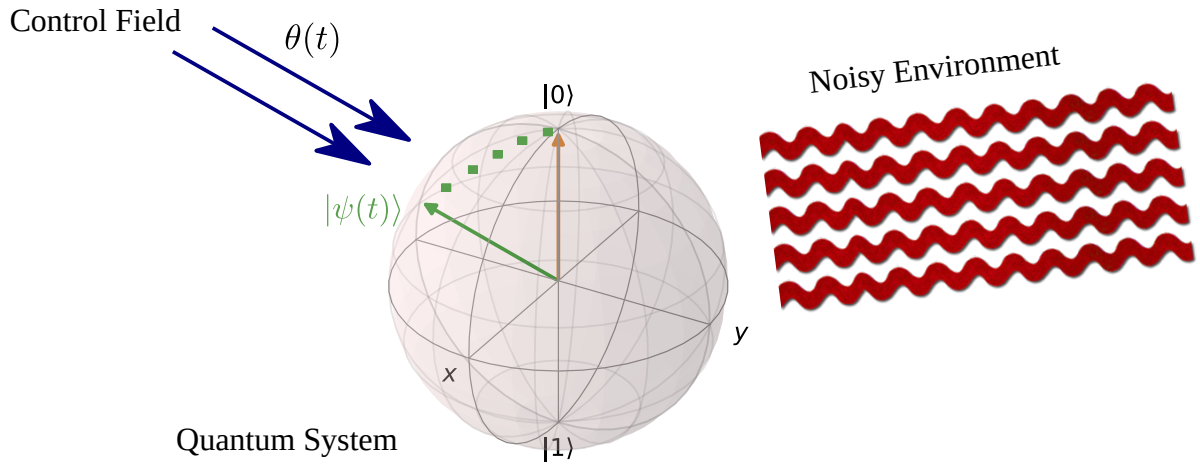


FIGURE 2.2: Basic example of quantum control. A single qubit is depicted by the Bloch sphere and manipulated by a control field $\theta(t)$ in such a way that the qubit state $|\psi(t)\rangle$ rotates from an initial state, here $|0\rangle$, to some target state (not specified here). An additional challenge in a real-world experimental setup, like an Ising spin, is that it lives in a noisy environment (red waves).

noise. One way to include such noisy effects in the dynamics is to model the entire setup as a (dissipative) Lindblad system. Here the control differential equation, f , becomes the Lindblad equation

$$\frac{d}{dt}\rho(\theta, t) = -i[H(\theta(t)), \rho(\theta, t)] + \sum_{k=1}^{N^2-1} \gamma_k \left[A_k \rho(\theta, t) A_k^\dagger - \frac{1}{2} A_k^\dagger A_k \rho(\theta, t) - \frac{1}{2} \rho(\theta, t) A_k^\dagger A_k \right], \quad (2.12)$$

where the A_k are the specific dissipation operators [20]. In this case, one can again choose a fidelity measure for the objective function \mathcal{L} if one is interested in quantum state preparation. In addition, one can also look at objective functions that include a running payoff such as the integrated heating rate, see e.g. reference [82].

2.2 Traditional Methods for Quantum Control

To solve the optimal control problem and minimise \mathcal{L} with respect to θ a range of different techniques are available. In this section, we introduce some of the traditional methods for quantum control that we will use later on in the thesis. Specifically, we focus on Pontryagin's principle, shortcuts to adiabaticity and some methods for numerical optimisation such as simulated annealing. We have termed the combined group of these methods *traditional* methods for quantum control to distinguish them from the different group of machine-learning-based approaches, but this is totally up to someone's own interpretation.¹⁶ We will apply some of the techniques to the single-qubit problem to show how they perform in practice.

¹⁶One could argue for example that the machine learning methods we use are just a form of the general methods for functional minimisation.

2.2.1 Pontryagin's Principle and Bang-Bang Protocols

The first method we discuss is Pontryagin's maximum principle [22] that can be seen as a generalisation of the calculus of variations applied to the control functional \mathcal{L} . Lev Pontryagin was one of the first to recognise that the same methods for minimising functionals in classical mechanics can be applied to control problems to obtain optimal control pulses. For this reason, Pontryagin is often regarded as one of the founding fathers of optimal control theory.¹⁷

To make the analogy with classical mechanics clear we consider for the moment the cost functional \mathcal{L} in Eq. (2.2) that only depends on the running payoff.¹⁸ We can recognise that this functional has the form of the *action* in classical physics, namely, it is an integral over some Lagrangian, $L(x, \theta(t), t)$. However, a crucial difference is that we now have the dependence on the control parameters $\theta(t)$ which vary with time, whereas in classical physics this is usually some fixed setup parameter. Furthermore, we also require the system to satisfy the differential equation $\dot{x} = f(x, \theta, t)$, which we can add in by a Lagrange multiplier $p(t)$ [18]. This results in the so-called augmented Lagrangian,

$$L_a(x, p, \theta(t), t) = L(x, \theta(t), t) + p(t) [\dot{x} - f(x, \theta(t), t)], \quad (2.13)$$

where $p(t)$ is sometimes referred to as the costate. Here, $p(t)$ can be seen as being a conjugate momentum. This means that by applying a Legendre transform $H = \sum_i \dot{q}_i \frac{\partial L_a}{\partial \dot{q}_i} - L_a$ we can form the control Hamiltonian

$$H(x, p, \theta(t), t) = p(t) [f(x, \theta(t), t)] - L(x, \dot{x}, \theta(t), t). \quad (2.14)$$

In this way, we have cast the optimal control problem in the form of a (classical) Hamiltonian. We can proceed to minimise by applying the calculus of variations,¹⁹ which results in the following Hamilton's equations for the optimal protocol $\theta^*(t) \equiv \theta_{\text{opt}}(t)$,

$$\dot{x}^* = \frac{\partial H}{\partial p}(x^*, \theta^*(t), p^*, t), \quad (2.15)$$

$$\dot{p}^* = -\frac{\partial H}{\partial x}(x^*, \theta^*(t), p^*, t), \quad (2.16)$$

$$\frac{\partial H}{\partial \theta}(x^*, \theta^*(t), p^*, t) = 0, \quad (2.17)$$

in full analogy with classical mechanics. Note that the first equation is just the differential equation of the system itself that we added in by the Lagrange multiplier. The second equation

¹⁷Fun fact: he was deriving methods to solve how to control rockets during the early days of the cold war.

¹⁸We note that we can always write the termination cost as an integral as well if it is a differentiable function of the termination time.

¹⁹The precise calculation is long but straightforward and can be found in [18].

is known as the costate equation. Both these need to be solved before the control Hamiltonian Eq. (2.17) can be minimised to find the optimum.

Importantly, these equations are only true for optimisation problems where there are no additional constraints on the possible control protocols $\theta(t)$. As discussed before, in a real-world setup there might be only a limited set of admissible protocols $\theta_{\text{ad}}(t)$. Pontryagin was able to derive a similar set of equations that specifically deals with this scenario. In this case, Eq. (2.17) gets changed into

$$H(x^*, \theta^*(t), p^*, t) = \max_{\theta_{\text{ad}}(t)} H(x^*, \theta(t), p^*, t), \quad (2.18)$$

which is nowadays regarded as Pontryagin's maximum (minimum) principle.²⁰

The maximum principle provides a necessary set of equations of motion that the optimal protocol θ_{opt} needs to fulfil. Solving these equations, however, or even writing them down is often a very hard task, strongly depending on the problem at hand. Nevertheless, examining the form of the control Hamiltonian, without solving the equations, can already give some important insights in the form of the optimal control. An important example of this is for control Hamiltonians that are linear in the control parameter $H(\theta) \sim \theta$. If the admissible protocols are bounded $\theta_{\text{ad}} \in [\theta_{\text{min}}, \theta_{\text{max}}]$, we find by applying Eq. (2.18) that the optimal protocol $\theta_{\text{opt}}(t)$ has at each instance in time t either the value $\theta_{\text{opt}}(t) = \theta_{\text{max}}$ or the value $\theta_{\text{opt}}(t) = \theta_{\text{min}}$. This special type of control protocol is called a *bang-bang* protocol and an example is shown in Fig. 2.3 (a).

While Pontryagin's principle was originally formulated for classical control problems, it also applies to quantum control problems. In particular, it was recently shown in references [24, 83] that physical quantum Hamiltonians, \mathcal{H} , linear in the control can be transformed into classical control Hamiltonians H linear in the control. This means that quantum control problems with a Hamiltonian linear in the control have bang-bang protocols as optimum. For example, for the single-qubit control problem as introduced before, which is linear in the magnetic control field $\theta(t)$, we expect optimal bang-bang protocols. This result is quite powerful, since it dramatically restricts the search space for the possible optimal protocols.²¹ This is particularly relevant for some of the numerical optimisation methods for quantum control as we will see below.

2.2.2 Shortcuts to Adiabaticity

In its original formulation by Born and Fock [84] the adiabatic theorem can be stated as follows; a quantum system remains in its instantaneous eigenstate if the change in a parameter

²⁰The proof for this is involved and we refer the reader to reference [18] for a heuristic discussion or to Pontryagin's book [22] for the full proof.

²¹The search space is still exponentially large in the time length of the individual bangs.

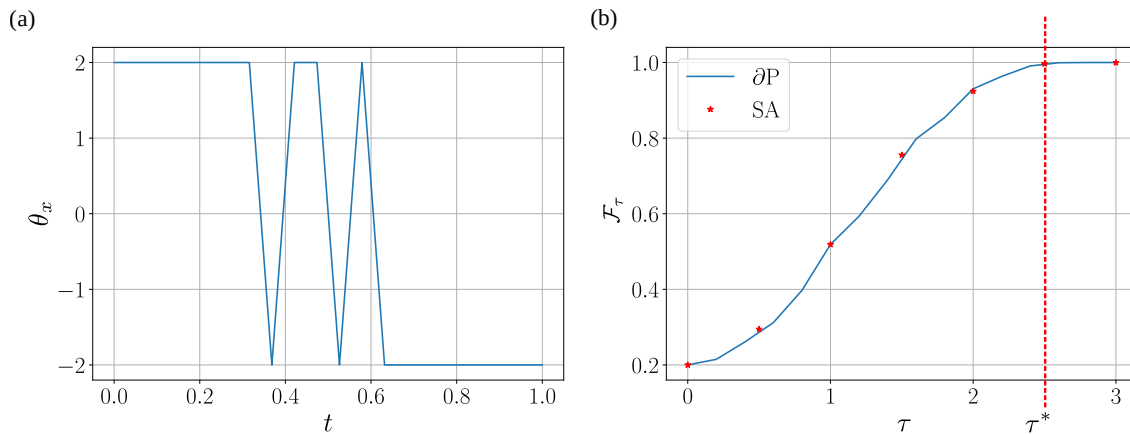


FIGURE 2.3: (a) An example of a bang-bang control protocol for the control parameter θ_x . In this type of protocol at each time θ_x only takes one of the two extremal values, i.e. $\theta_x(t) = \theta_{\max} = 2$ or $\theta_x(t) = \theta_{\min} = -2$. This particular example protocol was obtained with a SA search of the single Ising spin control problem with magnetic field θ_x and termination time $\tau = 1$. (b) Optimised target state fidelity \mathcal{F}_τ of the single Ising spin control problem as a function of the termination time τ . We compare the differentiable programming ($\partial\mathcal{P}$) method to a SA search and observe that they give the same results. For termination times $\tau > \tau^* \approx 2.5$ a perfect fidelity can be reached. This defines the speed limit for this control problem. The SA optimisations were run for 10^4 update steps and the $\partial\mathcal{P}$ method in combination with gradient descent for 200 update steps with a learning rate $\alpha = 10$.

$\theta(t)$ of the Hamiltonian is slow enough and there is a gap between the corresponding eigenvalue and the rest of the spectrum. Mathematically this can be expressed as

$$\mathcal{T} e^{-i \int_0^t \mathcal{H}(\theta(s)) ds} |\psi_0^j\rangle \approx |\psi_t^j\rangle \quad (2.19)$$

for a true adiabatic process in which $|\psi_0^j\rangle$ is the j -th eigenstate of $\mathcal{H}(0)$ and $|\psi_t^j\rangle$ the eigenstate of $\mathcal{H}(t)$ for all times $t \in [0, \tau]$. The condition under which this approximation is justified [85] is inversely proportional to the spectral gaps and hence the time of an adiabatic process will diverge when these become close to zero. Practically speaking this means that adiabatically tuning the Hamiltonian is not always a successful control strategy due to the fact that the environment now has an infinite amount of time to decohere and interact with the system.

To overcome this issue physicists started exploring other faster methods, which are now known under the name of shortcuts to adiabaticity (STA) [27]. In these methods, the specific evolution of the state at intermediate times, between the initial and termination time, is less relevant as long as the desired final instantaneous eigenstate of $\mathcal{H}(\tau)$ is reached. This means that the system at intermediate times is not necessarily following an adiabatic path and can be in a completely different state than the instantaneous eigenstate. It turns out that, in this way, strategies can be developed that can be implemented over much shorter timescales, while still reaching the target state perfectly.

One method for this is *counterdiabatic driving* [86, 87], where a specific counterdiabatic

Hamiltonian $H_{cd}(t)$ is imposed on top of the original Hamiltonian $H_0(t)$ that one tries to control. The purpose of this additional term is to suppress all unwanted nonadiabatic excitations such that the system is always in an instantaneous eigenstate of H_0 . However, this often leads to nonlocal and possibly inadmissible controls and moreover it is extremely hard to find the form of $H_{cd}(t)$ for many-body systems. This is because the construction of $H_{cd}(t)$ involves the eigenvectors of H_0 , which are nontrivial to find for many-body systems. For a more detailed review of counterdiabatic driving and a discussion of how to overcome some of these issues, we refer the reader to [88, 27].

Another shortcut method designed for the same purpose is based on Lewis-Riesenfeld invariants [89, 90, 27] and is also known as inverse engineering. A Lewis-Riesenfeld or dynamical invariant is an observable that remains invariant during an arbitrary evolution of the system $|\psi(t)\rangle$ with Hamiltonian $H(t)$. More formally, it is a hermitian operator $I(t)$ which satisfies the equation²²

$$\frac{\partial I(t)}{\partial t} + i[H(t), I(t)] = 0. \quad (2.20)$$

It can be shown that any arbitrary solution $|\psi(t)\rangle$ of the Schrödinger equation can be expanded in the eigenstates $|\phi_j(t)\rangle$ of the invariant $I(t)$ as

$$|\psi(t)\rangle = \sum_j a_j e^{i\gamma_j(t)} |\phi_j(t)\rangle, \quad (2.21)$$

in which the $\gamma_j(t)$ are the so-called Lewis-Riesenfeld phases given by

$$\gamma_j(t) = \int_0^t \langle \phi_j(s) | i \frac{\partial}{\partial s} - H(s) | \phi_j(s) \rangle ds \quad (2.22)$$

and a_j the time-independent expansion coefficients.

This invariant and the knowledge of its eigenstates can be used to inverse engineer the Hamiltonian $H(t)$ such that an initial eigenstate $|\psi_0\rangle$ is driven to the target eigenstate $|\psi_\tau\rangle$ in arbitrary short times. The trick is to always remain in the instantaneous eigenstates of $I(t)$, which means we use the unitary evolution operator

$$U(t) = \sum_j e^{i\gamma_j(t)} |\phi_j(t)\rangle \langle \phi_j(0) |. \quad (2.23)$$

The Hamiltonian can then be found by inverting the equation $i\dot{U} = H(t)U$ resulting in

$$H(t) = i\dot{U}U^\dagger. \quad (2.24)$$

²²Note that this is just the time evolution equation in the Heisenberg picture with $dI/dt = 0$.

Importantly we still need to apply the boundary conditions $[I(0), H(0)]$ and $[I(\tau), H(\tau)]$ to find the $H(t)$ that brings the initial state to the target.

This method to inverse engineer the Hamiltonian and hence the control parameters $\theta(t)$ is fully dependent on there being a dynamical invariant for the quantum system one is trying to control. For the control of a single qubit the invariant and solution can be found in reference [91]. For many-body systems, in contrast, it is much harder to find a dynamical invariant. For example, for the topological system that we will look at in the next chapter it is unclear if such an invariant exists. In chapter 6 we will look, however, at a quantum many-body system that can be (approximately) mapped to a single-particle system. The known invariant of the single-particle system can then be exploited to derive approximate STA solutions for the many-body system.

2.2.3 Numerical Techniques

Pontryagin's principle and the STA methods largely depend on the ability to solve the system (eigenstates) or Hamilton's control equations. In a lot of cases, this is not possible or at least very hard, so one needs to resort to numerical methods for the minimisation of the objective function \mathcal{L} . For this, there is a whole zoo of mathematical algorithms available [92] and here we will just review two of them that can be applied to arbitrary quantum control problems.

Simulated Annealing

One of the most widely used algorithms for optimisation problems is simulated annealing (SA). This method is based on the Metropolis-Hastings Monte Carlo algorithm [25] and was originally pioneered in a few different (independent) studies [93, 26, 94]. Today it is applied to tackle a broad range of problems [95] within and also outside Physics. There are several formulations of SA but the one that we use here is based on a randomised (stochastic) search of the cost landscape, $\mathcal{L}(\theta)$. In our optimal control context this means we start from some initial configuration for the control $\vec{\theta} = (\theta_1, \dots, \theta_n)$, which we assume to be a vector of in total n individual parameters θ_i . In the next step we randomly select an individual parameter θ_i and update it according to

$$\theta_i^{j+1} = \theta_i^j \pm \Delta\theta_i \quad (2.25)$$

in which $\Delta\theta_i$ is some small finite element. We evaluate the cost, $\mathcal{L}(\theta_i^{j+1})$, and accept the update with a probability

$$\mathbb{P}_{\text{SA}} = e^{-\mathcal{L}(\theta_i^{j+1})/T_{\text{SA}}}. \quad (2.26)$$

T_{SA} is the so-called annealing temperature, which is slowly cooled down to zero while we iteratively update according to Eq. (2.25). Once the annealing temperature is zero or the cost \mathcal{L} does not change anymore the process is said to be converged and the optimisation

completed. Due to the stochastic component, the process is often restarted a few times to ensure a minimum is reached.

The idea behind this algorithm is that during the optimisation it stochastically explores a large part of the cost landscape, $\mathcal{L}(\theta)$, and finally terminates approximately near the global minimum. However, a drawback is that the search can be exhaustive due to the large family of possible control protocols $\theta(t)$ or a computationally expensive computation of the cost. In this case, it is helpful to reduce the size of the search space by restricting it to a preconditioned set of protocols, such as the bang-bang protocols. We will now exemplify this method with the single-qubit control problem.

We focus on the Ising spin case²³ in which we have a single time-dependent control field $\theta_x(t)$. We discretise this field into $n = \tau/\Delta t$ individual time steps (bins) of time Δt giving $\vec{\theta} = (\theta_1, \dots, \theta_n)$ with $\theta_i \equiv \theta_x(i\Delta t)$. We wish to start from the initial ground state $|\psi_0(\theta_x = 2)\rangle$ of $\mathcal{H}(\theta_x = 2)$ in Eq. (2.11) and end up in the ground state $|\psi_0(\theta_x = -2)\rangle$ of $\mathcal{H}(\theta_x = -2)$. The objective function is then given by the infidelity $\mathcal{L} = \mathcal{I}_\tau$ of Eq. (2.9). We can compute this numerically by diagonalising and exponentiating the Hamiltonian $\mathcal{H}(\theta_x(t))$. Since the Hamiltonian is linear in $\theta_x(t)$ we know from Pontryagin's principle that the optimal control protocols are of the bang-bang form, $\theta_x(t) = \theta_{\max}$ or $\theta_x(t) = \theta_{\min}$ for all times t . For this we assume that there is a maximum and minimum amplitude of the magnetic field, which we set to $\theta_{\max} = 2$ and $\theta_{\min} = -2$. This means that for the SA algorithm we can set $\Delta\theta_i = 4$ in Eq. (2.25), randomly select a time bin θ_i^j and update according to $\theta_i^j = \theta_{\max} \mapsto \theta_{\min}$ or $\theta_i^j = \theta_{\min} \mapsto \theta_{\max}$, i.e. switch the field. We note that while this bang-bang search space is small compared to a continuous $\theta_x(t)$ search space, it is still exponential in the number of time bins n and the search can be slow.

In Fig. 2.3 (b) we show the final target state fidelity \mathcal{F}_τ obtained with this SA search for a range of different termination times τ . We see that from approximately $\tau \approx 2.5 \equiv \tau^*$ the method is able obtain protocols that reach the target state perfectly $\mathcal{F}_\tau = 1$. This means that for times smaller than τ^* the single-qubit system is not controllable since the optimal control problem can not be solved.²⁴ This is because the system has not enough time to evolve into the target state for these small termination times no matter what protocol is used. The time τ^* , therefore, defines a speed limit for this system as is discussed in more detail in references [80, 36]. As our only aim is to show how the SA method can be used to search for optimal control protocols we will leave the discussion here and refer the interested reader to these other studies mentioned above.

²³Recall that for the Ising spin the magnetic fields in the y and z directions are fixed to be $\theta_y(t) = 1$ and $\theta_z(t) = 0$.

²⁴We verify this with a different optimisation method below and this also agrees with the results found in [36].

Chopped Random Basis

Another numerical optimisation method that can be applied to quantum control problems is the chopped random basis (CRAB) algorithm [96, 97].²⁵ In CRAB one reduces the number of control parameters θ_i , while still being able to explore a large part of the generic search space. This is done by using a complete random basis of functions as ansatz for the control protocols $\theta(t)$. The most typical example is a Fourier series ansatz,

$$\theta(t) = b(t) \sum_{i=1}^N A_i \sin \omega_i t + B_i \cos \omega_i t, \quad (2.27)$$

in which $b(t)$ imposes the boundary conditions and the Fourier parameters A_i, B_i and ω_i become the new control parameters. The number of parameters can now be reduced to be smaller than n by choosing (chopping) $N < 3n$. Moreover, in most cases, the frequencies are chosen as a fixed set $\omega_i = 2\pi i/\tau + \delta\omega$ where $\delta\omega$ is some random number. This means that one now only needs to optimise over $2N$ parameters, which is a large reduction and can lead to quicker convergence to the minimum of \mathcal{L} .

The Fourier parameters can be optimised with numerical algorithms such as the SA discussed above. Another method that is often used for this is Nelder-Mead, which is based on *simplices* of test points of the objective function [98, 99]. Both these methods do not make use of any gradient signal $\frac{d\mathcal{L}}{d\theta}$, however, which means that they can be slow compared to gradient based numerical methods for quantum control such as GRAPE [100] and Krotov [75]. To apply gradient based methods one needs to obtain a gradient signal, which typically is computed analytically or approximated numerically. In the Machine Learning community, a different efficient way of computing derivatives has been known for a long time, which is known as Differentiable Programming. We will introduce this method and some gradient based optimisation algorithms in detail in the next section.

2.3 Machine Learning for Quantum Control

In recent years machine learning (ML) techniques have been applied to solve all kinds of optimisation problems in quantum physics, see [31, 32] for extensive reviews. Similar to the traditional methods for quantum control the problems are defined by some generic cost functional \mathcal{L} that needs to be minimised. One particular ML technique for minimisation that is also explored for optimal control problems is Reinforcement Learning (RL) [101]. Here the machine is trained to perform the (control) task by reinforcing the information gained from past experiences. In this section we will briefly introduce the RL approach for quantum control and then focus on two specific optimisation methods to train the machine.

²⁵In fact it is more of a randomised ansatz for the protocols than a full algorithm.

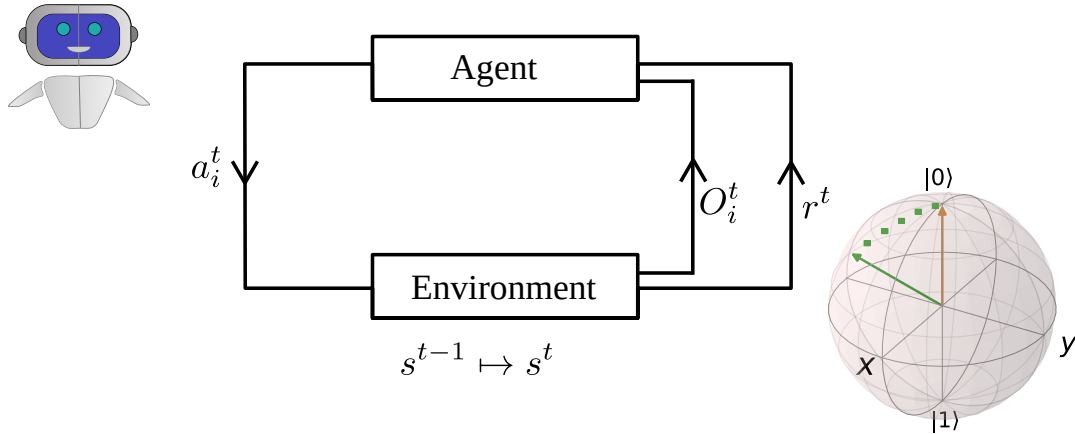


FIGURE 2.4: Schematic of the generic reinforcement learning setup. At time t in the game, an agent (robot) takes an action $a_i^t \in \mathcal{A}^t$ that changes the state $s^t \in \mathcal{S}^t$ of the environment. Afterwards, the agent receives a new observation O_i^t and possibly a reward r^t . Based on this information the agent then determines its next move $a_i^{t+1} \in \mathcal{A}^{t+1}$. This cycle is repeated until the game is finished at $t = \tau$ with total cumulative reward $R = \sum_t \gamma_t r_t$. The aim of the agent is to choose the actions such that R is maximised. In reference [36] it was shown that optimal quantum control problems, such as the single-qubit example, can be formulated within this RL game paradigm.

2.3.1 Reinforcement Learning

Reinforcement learning has achieved significant success in the context of gameplay [102, 103]. In a standard RL game setting, there is a player (agent) that needs to achieve some goal in a (game) environment and gets rewarded depending on how well it performs. During the game, the agent interacts with the environment, which means the environment changes/evolves over time. The agent is able to receive information of the state of the environment via observations, for example, the positions of individual pieces on a chessboard. Based on these observations, and the reward, the agent makes a guess for the best next move. By playing a game many times, and tracking its past experiences and rewards, the agent learns to optimise its moves such that it maximises the chance to win the game (maximises the reward).

This setup is summarised schematically in Fig. 2.4 and can be formulated mathematically as follows. The moves at times t in the game are known as *actions* a_i^t ²⁶ and the space of all possible actions is denoted as \mathcal{A}^t . We denote the *state space* of the environment by \mathcal{S}^t and *observations* of it by O_i^t . The, possibly cumulative, *reward* is given by $R = \sum_t \gamma_t r^t$ in which γ_t is a discount factor that determines the relative importance of near term versus long term rewards r^t . The optimisation task is to find a certain *policy* $\pi(a^t)$ to select the individual actions $a_i^t \in \mathcal{A}^t$ that maximises R . Specific examples of this generic RL setup can be found in [101].

The general RL problem formulation is closely related to the optimal control problem formulation. A characterisation for this analogy is given by the dictionary in Table 2.1. Here

²⁶Note that i labels the different types of actions and can be discrete (finite action space) or continuous (infinite action space).

Optimal (Quantum) Control	Reinforcement learning
Objective Function \mathcal{L}	Cumulative Reward R
Control Parameters θ	Actions a_i
Optimal Control θ_{opt}	Final Policy π_{fin}
Controller	(RL) Agent
System ($\dot{x} = f(x, \theta)$)	Environment (model)
Feedback/Measurement	Observation

TABLE 2.1: Dictionary for the similarity between optimal control language and the reinforcement learning game setting language.²⁸ In this thesis we will sometimes use the RL and control formulations interchangeably.

the objective function takes the role of the reward, the control parameters correspond to the actions and the system becomes the environment. In reference [36] it was, for example, shown that the single-qubit control problem²⁷ can be fitted within the RL game setting. The specific reward which the agent in this case needs to maximise is the target state fidelity $R = \mathcal{F}_\tau$. The single-qubit system plays the role of the environment and the actions the agent is allowed to take are manipulations of the value of the magnetic field $a_i^t \sim \theta_x(t)$. By playing this game the agent then learns how to change the magnetic field such that the desired target state is reached.

Maximising the reward R by training the agent is not very different from optimising the control functional \mathcal{L} and similar techniques can be used. However, from the RL (and more generally ML) literature, there are some additional methods available compared to the traditional control methods introduced before. These include, for example, policy gradients [77] and the Watkins Q-learning algorithm [76]. Two specific (ML-inspired) algorithms, Differentiable Programming and Natural Evolutionary strategies, will be used in the following chapters to control quantum many-body systems. We will now give a general description of these methods.

2.3.2 Differentiable Programming

Learning algorithms for the optimisation of objective functions or the training of RL agents can be divided into gradient-based and gradient-free methods. For the gradient-based methods one requires access to the gradient of the cost function with respect to the control parameters, $\frac{\partial \mathcal{L}}{\partial \theta}$. Differentiable Programming (∂P) [104, 105] is a special programming paradigm that allows to compute exact derivatives of arbitrary cost functions efficiently without requiring much implementation effort.²⁹ In some situations ∂P is computationally much more efficient than methods such as numerical differentiation or symbolic differentiation in Mathematica.

²⁷And also the control of a few coupled qubits.

²⁹Provided one makes use of a programming language that allows automatic differentiation such as JAX [106] discussed below.

This was exploited recently for the optimisation of tensor networks [107], Monte Carlo simulations [108], and also to search for quantum control protocols [109, 38].

To compute derivatives $\partial\mathcal{P}$ makes use of the computational graph of a computer programme. A computational graph is a visual representation of all the elementary functions (primitives) f^i in a computer programme that maps the input to the output. An example of this is a programme that computes the loss function $\mathcal{L}(\theta) = f^n \circ f^{n-1} \circ \dots \circ f^1(\theta)$ ³⁰ as shown in Fig. 2.5 (a). Here \circ is the composition of two functions. In such a graph the edges represent the operations (functions) and the nodes the returned data types. The most fundamental elementary operations are addition and multiplication of scalars. For practical considerations, computational graphs also allow for some granularity in which higher-order primitives like exponentiation or matrix operations are expressed as edges. Any computer programme can be decomposed in this way and represented as a computational graph.

The derivative of a generic computational graph can then be obtained by applying the chain rule. For a simple sequential graph this reads³¹

$$\frac{\partial\mathcal{L}}{\partial\theta_i} = \frac{\partial\mathcal{L}}{\partial f^{n-1}} \frac{\partial f^{n-1}}{\partial f^{n-2}} \dots \frac{\partial f^1}{\partial\theta_i} \quad (2.28)$$

for a particular input parameter θ_i that could be part of a vector $\vec{\theta} = (\theta_1, \dots, \theta_N)$ of input parameters. If all the elementary operations f^j are differentiable and have known derivatives $\frac{\partial f^j}{\partial f^i}$ with $j > i$, the total derivative can be found by recursively assembling them together.

In $\partial\mathcal{P}$ one distinguishes between two orderings of recursively evaluating the chain rule. In the so-called forward-mode differentiation one traverses the computational graph in the regular order of execution, i.e. from left to right [black arrows in Fig. 2.5 (a)]. The chain rule is then assembled via

$$\frac{\partial f^j}{\partial\theta_i} = \frac{\partial f^j}{\partial f^{j-1}} \frac{\partial f^{j-1}}{\partial\theta_i}. \quad (2.29)$$

A second method is reverse-mode differentiation³² in which one transverses the computational graph from the output all the way back to the input [orange arrows in Fig. 2.5 (a), (b)]. Here the chain rule is evaluated by recursively computing the so-called adjoint defined by

$$\bar{f}^j = \bar{f}^{j+1} \frac{\partial f^{j+1}}{\partial f^j}. \quad (2.30)$$

in which $\bar{f}^n = 1$.

Whether to use forward-mode or reverse-mode differentiation depends on the specific

³⁰Note that this equation only holds for sequential computational graphs in which all the functions are applied in series.

³¹This can be generalised to non-sequential computational graphs like the graph in Fig. 2.5 (b).

³²This generalises the notion of backpropagation for neural networks [110].

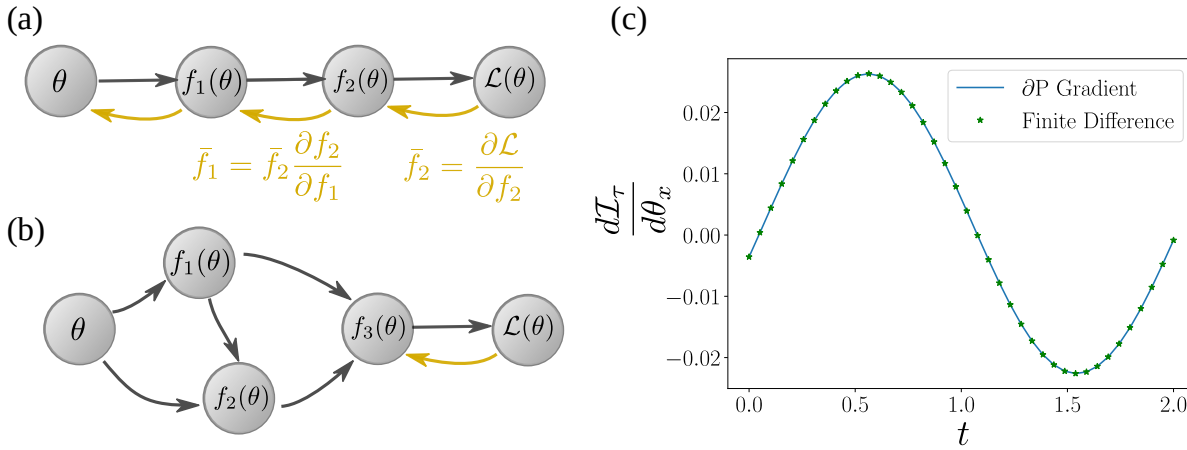


FIGURE 2.5: (a) Example of a sequential computational graph for the evaluation of some loss function \mathcal{L} . The nodes represent data structures and the edges the evaluation of primitive functions f_i . In backward-mode ∂P one evaluates the gradient of \mathcal{L} by traversing the graph from the right to the left (orange arrows). At each step, the adjoints \bar{f}_i are evaluated and assembled together recursively to find the total derivative $\frac{\partial \mathcal{L}}{\partial \theta}$. (b) An example of a non-sequential computational graph. Here the application of the primitives does not necessarily have to be done in series. The same forward and backward mode differentiation techniques can be used for the evaluation of the gradient. (c) Comparison between the gradient of the infidelity \mathcal{I}_τ with respect to the magnetic control field θ_x obtained with differentiable programming (∂P) and finite-difference. This gradient was evaluated for a constant magnetic field $\theta_x(t) = -2.0$.

problem at hand. When the number of output parameters is greater than the number of input parameters it is computationally more efficient to use the forward mode [105]. However in quantum control applications one usually has one scalar output and many, say N , input parameters. For this type of problem reverse-mode differentiation is the most efficient since it only requires evaluating the chain rule once compared to $\mathcal{O}(N)$ evaluations for forward-mode. The highest computational efficiency that can be reached for any given gradient evaluation is given by the Bauer-Strassen theorem [111]. This theorem states that the computational cost of computing a function, together with its derivative, is upper bounded by three times the computational complexity of the forward evaluation of the function. This gives a good benchmark to test the computational performance of gradient evaluations.

Gradient Descent and Single Qubit Control

The gradient $\nabla_{\vec{\theta}} \mathcal{L}$ ³³ obtained with ∂P can be used to train the RL agent and minimise the objective function \mathcal{L} . A widely used algorithm for this in the machine learning community (and also in the field of mathematical optimisation) is 'vanilla' gradient descent. In gradient descent, one steps down in the direction of the steepest descent, see e.g. the red arrows in Fig. 2.1, until a local minimum (or saddle point) $\nabla_{\vec{\theta}} \mathcal{L} = 0$ is reached. This has the following

³³For generality we assume we have a vector of control parameters $\vec{\theta} = (\theta_1, \theta_2, \dots, \theta_n)$ and hence $\nabla_{\vec{\theta}} \mathcal{L} = (\frac{\partial \mathcal{L}}{\partial \theta_1}, \dots, \frac{\partial \mathcal{L}}{\partial \theta_n})$.

iterative update rule for the control parameters

$$\theta_i^j \mapsto \theta_i^{j-1} - \alpha \frac{\partial \mathcal{L}^{j-1}}{\partial \theta_i}, \quad (2.31)$$

in which α is the so-called learning rate, which determines the size of the steps. For convex optimisation problems this update scheme, with appropriately chosen α , is expected to converge monotonically towards the global minimum [79]. For non-convex optimisation landscapes, this method could, however, get stuck in a sub-optimal local minimum. To circumvent this drawback one can restart the optimisation from different initial configurations or use other more elaborate update schemes such as the quasi-Newton method [92] or the optimiser with the acronym Adam [112].

We demonstrate the power of this gradient based optimisation approach with the single Ising spin control example. As before the objective function is the infidelity \mathcal{I}_τ and we are allowed to tune the magnetic field $\theta_x(t)$ in time. The programming code to compute \mathcal{I}_τ involves the diagonalisation and exponentiation of the Hamiltonian in Eq. (2.10) and matrix multiplications for the overlap in Eq. (2.9). Each of these operations (primitives) is differentiable [113] and we can apply $\partial\mathcal{P}$ to obtain $\frac{d\mathcal{I}_\tau}{d\theta_x}$. In practice we make use of a programming language that makes use of automatic differentiation (AD) such that we do not need to implement the derivatives of the primitives ourselves. The specific python library we use for this is JAX [106], which has as additional advantage that the computations are automatically parallelized and can be run on GPUs. The derivatives can then be obtained with a computational time close to the Bauer-Strassen bound.³⁴

To test that the code is working correctly we compare the $\partial\mathcal{P}$ derivatives with derivatives obtained from the less efficient finite difference method. In this method, the derivatives are approximated by taking a finite element $\Delta\theta$ and using

$$\left. \frac{d\mathcal{L}}{d\theta} \right|_{\theta_0} \approx \frac{\mathcal{L}(\theta_0 + \Delta\theta) - \mathcal{L}(\theta_0 - \Delta\theta)}{2\Delta\theta}. \quad (2.32)$$

This is essentially the finite approximation of the definition of the real continuous derivative, which is obtained in the limit $\Delta\theta \mapsto 0$. In Fig. 2.5 (c) we give the results for this benchmark test and we observe that the derivatives match.³⁵ However, the finite difference derivative calculation is much more demanding to compute when compared to $\partial\mathcal{P}$, since the loss \mathcal{L} is evaluated $2N$ times with N being the number of discrete time steps. This advantage of $\partial\mathcal{P}$ over the finite difference method is particularly relevant for calculating derivatives of control problems involving many-body systems.

³⁴Note however that we use a Trotter expansion for the unitary time evolution operator $U(\tau)$, which means that the memory complexity grows with the number of discrete individual time steps in the evolution.

³⁵Up to the finite precision of the finite difference method.

To then solve the single Ising spin quantum control problem we use this derivative $\frac{d\mathcal{I}_\tau}{d\theta_x}$ in combination with gradient descent on $\theta_x(t)$. Since ∂P is so efficient we can scan over a large number of termination times τ . In Fig. 2.3 we show the results of these optimisations and compare them to the infidelity values obtained earlier with the SA method. The \mathcal{I}_τ values are competitive and the same speed limit of about $\tau^* \approx 2.5$ is acquired. Although not shown here, the optimal protocols obtained with ∂P are of the bang-bang form. This means that with ∂P we are able to (independently) uncover the correct physical protocols. For ∂P this was done in only 200 gradient descent update steps, while for SA 10^4 evaluations of the objective function were required. This shows the potential of the ∂P gradient-based optimisation method for quantum control problems.

2.3.3 Natural Evolution Strategies

Aside from gradient based-optimisation algorithms, there are also gradient-free optimisation schemes to train the RL agent. These methods offer a way out in scenarios where no gradient signal can be obtained. For instance, when the exact model (Hamiltonian) of the underlying system is unknown. One of these methods is Natural Evolution Strategies (NES). Evolution strategies is a family of blackbox optimisation techniques inspired by the process of natural evolution that can be used to minimise arbitrary cost functions [114, 115]. In the context of game play it was recently shown that NES can be used as an efficient alternative optimisation approach compared to other reinforcement-learning methods [116]. For physics applications, this method is, however, rather unexplored with only a limited number of studies [117, 118, 119].

The particular version of NES that we will use in this thesis to control quantum many-body systems is shown schematically in Fig 2.6 and can be described as follows. One first draws a set of sample parameters ϕ^j from a Gaussian distribution $\mathcal{N}(\theta, \sigma^2 I)$ with mean θ and diagonal covariance matrix $\Sigma = \sigma^2 I$. Here, θ , are the control parameters that we wish to optimise and σ is a hyperparameter. The individual samples ϕ^j for $j = 1, \dots, n_{\text{pop}}$ are seen as members of a population and each individual loss, $\mathcal{L}(\phi^j)$ (their *fitness*), is evaluated. The goal of the NES algorithm is to optimise the mean θ^{36} of the Gaussian distribution such that the expectation value of the loss over the population,

$$\mathbb{E}_{\phi \sim \mathcal{N}(\theta, \sigma)} [\mathcal{L}(\phi)] = \sum_{j=1}^{n_{\text{pop}}} p_{\theta}(\phi^j) \mathcal{L}(\phi^j), \quad (2.33)$$

³⁶There are versions of the NES algorithm that also optimise the variance σ to increase the expectation value of individual members in the population [115].




FIGURE 2.6: Schematic of the Natural Evolution Strategies optimisation process. (a) In the j -th iteration of the algorithm, samples for control parameters ϕ_i^j are drawn from a (multi-variate) Gaussian distribution centred around θ_i^j with fixed standard deviation σ . (b) The mean θ_i^j of the Gaussian distribution is updated until the optimum θ_i^{opt} is reached. (c) Depending on the specific NES algorithm, the standard deviation is updated as well such that at the end the Gaussian becomes sharply peaked around the optimum. In the ideal case this would approach a delta function.

is minimised. Here $p_\theta(\phi)$ is the probability for drawing ϕ from the Gaussian distribution. This means that the NES algorithm can be seen as solving a Gaussian-blurred version of the original control problem with objective \mathcal{L} .

The minimisation is done by applying (stochastic) gradient descent

$$\theta \mapsto \theta - \alpha \nabla_\theta \mathbb{E}_{\phi \sim \mathcal{N}(\theta, \sigma^2 I)} [\mathcal{L}(\phi)] \quad (2.34)$$

to the parameter θ . The required gradient of the population is obtained by using the 'log-likelihood-trick' [115], which gives

$$\nabla_\theta \mathbb{E}_{\phi \sim \mathcal{N}(\theta, \sigma^2 I)} [\mathcal{L}(\phi)] = \nabla_\theta \sum_{j=1}^{n_{\text{pop}}} p_\theta(\phi^j) \mathcal{L}(\phi^j), \quad (2.35)$$

$$= \sum_{j=1}^{n_{\text{pop}}} p_\theta(\phi^j) \nabla_\theta \log p_\theta(\phi^j) \mathcal{L}(\phi^j), \quad (2.36)$$

$$= \sum_{j=1}^{n_{\text{pop}}} p_\theta(\phi^j) \frac{(\phi^j - \theta)}{\sigma^2} \mathcal{L}(\phi^j). \quad (2.37)$$

Then by changing variables $\phi^j = \theta + \sigma \epsilon^j$, we arrive at

$$\nabla_\theta \mathbb{E}_{\phi \sim \mathcal{N}(\theta, \sigma^2 I)} [\mathcal{L}(\phi)] = \sum_{j=1}^{n_{\text{pop}}} p_\mathbb{I}(\epsilon^j) \mathcal{L}(\theta + \sigma \epsilon^j) \epsilon^j / \sigma \quad (2.38)$$

$$= \mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)} [\mathcal{L}(\theta + \sigma \epsilon) \epsilon / \sigma]. \quad (2.39)$$

The gradient can thus be approximated by sampling the value of the objective \mathcal{L} for perturbed parameters $\theta + \sigma \epsilon$ and no exact gradient evaluations need to be done. In practice, these

evaluations can be done in parallel, which means that NES is very efficient.

2.4 Summary

In this chapter, we have introduced the general problem of quantum control and several methods to solve it. We started by giving the mathematical formulation of the optimal control optimisation problem. We saw that the most important ingredient is the control functional \mathcal{L} , which needs to be minimized to find the optimal control protocols. We then looked at some traditional methods for the minimisation of \mathcal{L} . This includes analytical approaches such as Pontryagin's principle and the STA methods, and also the numerical optimisation schemes SA and CRAB. We argue that the advantage of the numerical approaches is that they can be applied to any arbitrary control problem. The analytical methods are often harder to apply since they require specific knowledge, e.g. the eigenstates, of the system.

Afterwards, we showed how quantum control problems can be formulated in the language of reinforcement learning. This means that some additional approaches from the field of machine learning can also be applied to search for optimal control protocols. We looked at two approaches in particular, Differentiable Programming and Natural Evolution Strategies. We saw that ∂P has an advantage that it is really efficient if a gradient signal is available. NES, while possibly less efficient, can be applied to problems where one does not have access to the gradient.

As a demonstration, we successfully applied some of these methods to control a single qubit. In the next chapters, we will investigate if they perform equally well for the control of more complicated quantum many-body systems.

Chapter 3

Topological Quantum Computation and Majorana Zero Modes

To realise real-world quantum processing devices one needs to be able to physically build qubits and have an experimental procedure to control and measure the quantum information encoded within them. A whole array of different setups have been proposed for this task ranging from trapped ions [13, 15] and cold atoms [120] to superconducting circuits [12] and photonic devices [121]. One of the main challenges of these setups is that they are susceptible to noise from the environment, as we briefly discussed, for example, in the last chapter for the single Ising spin in a magnetic field. Decoherence can potentially make these devices unworkable and quantum computations prone to errors.

About two decades ago, an alternative scheme based on topological materials that host special quasiparticles called anyons [49] was proposed as a potential way to overcome this challenge. In these topological schemes quantum information can be encoded non-locally, a feature that makes it robust against local errors. A crucial property of these topological systems is that they have a spectral gap between the ground and bulk states, which protects from unwanted excitations. Despite this, finding real-world candidate materials that host (non-abelian) anyons has turned out to be very hard [122, 50]. Furthermore, it is still an open question what is the best way to control these systems.

In this chapter, we introduce some of the aspects of topological quantum many-body physics. We first explain what anyons are, and then briefly review the theoretical proposal for topological quantum computation. Afterwards, we will introduce a concrete candidate material, the proximity-coupled superconductor, which hosts Majorana zero modes. These quasiparticles are among the most promising types of (non-abelian) anyons to be realised in experiments. We will discuss how to move these Majorana modes, an operation that is required to be able to perform computations with them. This material serves as a necessary primer for later chapters, where we will look at the optimal control of these anyons and also at possible internal sources of error induced by interactions.

3.1 Quantum Computation by Anyons

Anyons are quasiparticle excitations with exotic braiding statistics that can be found in low dimensional condensed matter systems (see Fig. 3.1 a)), such as on the boundaries of proximity-coupled superconductors [56] and fractional quantum hall systems [122]. A quasiparticle is not a standard particle but can be seen as an effective particle that is obtained after dressing an ordinary particle with effects and interactions from other particles in its surrounding environment [123]. For example an electron moving in a semiconductor can be described as a free quasiparticle with a specific effective mass m^* due to the effective potential from the band structure.

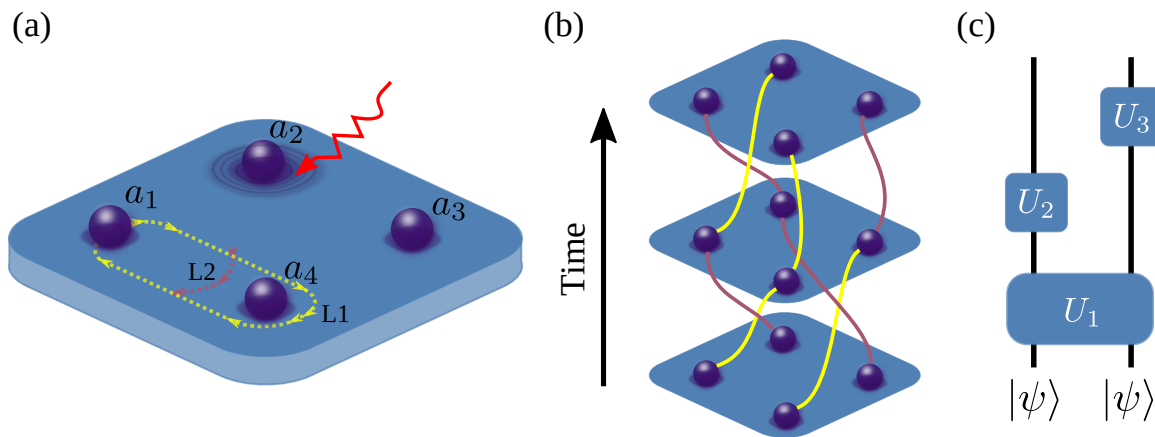


FIGURE 3.1: a) Localised anyons numbered a_{1-4} are effective quasiparticles living in two dimensional materials. Two times exchanging two anyons can be done by moving anyon a_1 via the path (L1) around anyon a_4 . Traversing the path (L2) on the other hand is equivalent to doing nothing. Local excitations caused by perturbations (red arrow) need to travel throughout the system and reach other anyons before causing an error. b) Braiding anyons in two dimensions can be used to perform quantum computation where the braids and exchanges of the anyons (red and yellow strands) are used to implement the unitary gate transformation U_i required for making the quantum circuit in c).

The low dimensionality makes that these special anyonic quasiparticles have nontrivial exchange statistics. This can be seen from the fact that paths traced out by moving anyons in two dimensions can not always be continuously deformed into one another, whereas in three dimensions they can. For instance, in 3D the loop (L1) taken by anyon a_1 in Fig. 3.1 (a) can be lifted up and transformed into the loop (L2) without cutting through the other anyon a_4 . In 2D this is not possible. This makes that twice exchanging a particle in 3D always leads to the same state, since the paths can always be connected to the trivial path (no movement). In 2D the state does not necessarily come back to itself after two exchanges of two anyons. The change of the wave function under n exchanges¹ can be expressed as $|\psi\rangle \mapsto e^{i\theta n} |\psi\rangle$. When $\theta = 0$ (π) the particles are ordinary bosons (fermions) since the wave function only changes

¹For simplicity we assume here that we only have two particles in our system described by the wave function $|\psi\rangle$.

by a sign. As such, in 3D we can only have bosons or fermions. When we have fractional phases $\theta = \pi/m$, the particles are Abelian anyons with m defining for the specific type of anyon model. This also explains the name anyon [124], which means *any* angle θ .

These particular anyons are called Abelian since the order in which they are exchanged does not matter, namely $e^{i\theta_1}e^{i\theta_2}|\psi\rangle = e^{i\theta_2}e^{i\theta_1}|\psi\rangle$. In systems with a degenerate ground-state space $|\psi\rangle = \alpha|\phi_1\rangle + \beta|\phi_2\rangle$, spanned by states $|\psi_1\rangle$ and $|\psi_2\rangle$, one can additionally have anyons with non-abelian exchange statistics. In this case, the exchange of two (non-abelian) anyons affects both the $|\phi_1\rangle$ and $|\phi_2\rangle$ states. The action of such an exchange on the wave function $|\psi\rangle \mapsto U_n|\psi\rangle$ is then represented by a unitary matrix U_n that rotates within the degenerate subspace.² As the unitary matrices U_n and $U_{n'}$, corresponding to two different exchanges, do not necessarily commute $U_n U_{n'} \neq U_{n'} U_n$ the exchange statistics are called non-abelian [50].

This property of exchanging non-abelian anyons means that they can be used as building blocks for topological quantum computers [49]. The states $|\phi_1\rangle$ and $|\phi_2\rangle$ can be seen as being the logical 0 and 1 qubit states and the physical braids (exchanges) of anyons represented by the matrices U_n act as the quantum gates, see Fig. 3.1 (c). For such a quantum computer to work one needs to be able to make all the possible quantum gates that are needed for the algorithms in universal quantum computation [125]. This means that for N qubits the braid matrices U_n need to span all representations of the group $U(2^N)$. The braid matrices of a few anyon models, such as the so-called Fibonacci anyons [126, 127], satisfy this requirement and hence can be used for universal quantum computation. The Majorana zero modes (Ising anyons) that will be discussed below do not satisfy this property but, as we will see, can be used as building blocks of a topological quantum memory.

To read the quantum information encoded in such topological qubits one can bring the anyons together and fuse them to form other (composite) quasiparticles. For example, when we fuse an abelian anyon with $\theta = \pi/m$ with another identical anyon with $\theta = \pi/m$ we get a composite particle with $\theta' = 2\pi/m$. This means that, if $m = 2$, we will obtain an ordinary fermion by combining two anyons. Computing the result of a fusion of two non-abelian anyons is more difficult since now we can have multiple different fusion outcomes resulting from the underlying non-abelian statistics.³ The specific quantum state $|\psi\rangle$ of the degenerate ground-state manifold determines the particular fusion outcome. The fusion outcomes can thus be used as a probe for the encoded quantum information. A more detailed discussion of this can be found in references [50, 128].

The use of anyons for quantum computation compared to other schemes, such as superconducting circuits and ion traps, is beneficial in the light of different sources of decoherence coming from the surrounding environment. The topological encoding of quantum information in pairs of anyons is intrinsically protected from local error and noise processes; an excitation

²Note that in the (non-degenerate) Abelian case U_n is a diagonal matrix.

³This can be seen since we are now looking for higher order representations of the non-abelian (braid) group elements [50].

occurring near one anyon needs to propagate throughout the system and reach another anyon before it can cause an error as shown in Fig. 3.1 (a). Moreover, the degenerate ground-state space is, as we will see later for a specific model, separated and protected by a robust (topological) energy gap from the excited states. In chapter 5 we will formally show that local perturbations are not able to distinguish between the different states that span this ground-state space (also known as topological order) and smooth adiabatic deformations of the states (braiding of the anyons) can be done in a finite amount of time.

This robustness of quantum information makes that anyons have been extensively studied in the last two decades for the use in (topological) quantum computers. Notably, last year the fractional exchange statistics of Abelian anyons was measured in a collider of anyons in a two dimensional electron gas [129] and also in a fractional quantum hall system [130]. Non-abelian anyons, in contrast, required for universal quantum computation, have proven to be very hard to find. One candidate are parafermions in a specific fractional quantum hall system [131], but so far they have not been realised in experiments. A more promising candidate are materials and devices based on topological superconductors. In these systems hints of the existence of non-abelian Ising anyons (Majorana zero modes) have been found in 2012 [55].⁴ We will discuss these special superconductors in the next section.

3.2 Topological Superconductivity and Majorana Zero Modes

Topological superconductors are a special type of superconductors, which have nontrivial topological bulk properties leading to the existence of non-abelian anyons at their boundaries. One specific class is made by p-wave superconductors, which have gained interest in the last years because of the possibility to realise them in real-world experimental setups. The non-abelian quasiparticles in these systems are called Majorana zero modes because, as we will see below, these modes have zero energy and their quasiparticle mode creation operators are hermitian $\gamma_j^\dagger = \gamma_j$.⁵ Associated to the zero-energy modes these systems have a degenerate ground-state manifold that can be used to make a topological qubit.

In this section, we first introduce a toy-model, the Kitaev Chain, for these topological superconductors and discuss in detail its topological properties together with the Majorana zero modes. Then we will describe how this model can be effectively realised in proximity-coupled superconductors and provide an update on its current experimental implementation.

⁴However, there is some controversy regarding some more recent experiments [57]. We will discuss some of the experimental issues in the next section.

⁵Similar to Majorana particles known from high energy physics [132].

3.2.1 The Kitaev Chain

The Kitaev chain is a lattice model introduced by Kitaev in 2001 [51], which in the continuum and low energy limits describes a spinless one-dimensional p-wave superconductor.⁶ The Hamiltonian of this one-dimensional tight-binding model is given by

$$\mathcal{H}_0 = - \sum_{x=1}^N [\mu(x) - V(x)] (c_x^\dagger c_x - 1/2) - w \sum_{x=1}^{N-1} (c_x^\dagger c_{x+1} + h.c.) + \sum_{x=1}^{N-1} (\Delta_x c_x^\dagger c_{x+1}^\dagger + h.c.) \quad (3.1)$$

in which c_x (c_x^\dagger) are the annihilation (creation) operators of electrons at the (discrete) position x in the chain. These operators satisfy the fermionic anti-commutation relation given by $\{c_i^\dagger, c_j\} = \delta_{ij}$ and $\{c_i, c_j\} = 0$. In this Hamiltonian, $\mu(x)$ describes an onsite chemical potential on which an external onsite potential $V(x)$ can be superimposed. The next term is the kinetic energy and accounts for hopping of particles between neighbouring sites with probability amplitude w . The final term is the p-wave superconducting term with superconducting gap Δ_x , which is included at a mean-field level [133, 134]. It describes a p-wave superconductor since Δ_x couples different lattice sites because the model is spinless.⁷

To investigate the topological properties of this model we first transform it to the so-called Majorana basis, given by the $2N$ Hermitian Majorana operators

$$\gamma_{2x-1} = c_x + c_x^\dagger, \quad \gamma_{2x} = \frac{c_x - c_x^\dagger}{i}, \quad (3.2)$$

which satisfy the Clifford algebra $\{\gamma_i, \gamma_j\} = 2\delta_{ij}$, $\gamma = \gamma^\dagger$. Since these operators are superpositions of both a fermionic creation and a fermionic annihilation operator, Majoranas can only be effectively realised as quasiparticles in systems that break parity conservation, e.g. superconductors. Applying this transformation to Eq. (3.1) results in the Hamiltonian

$$\mathcal{H}_0 = -\frac{i}{2} \sum_{x=1}^N (\mu(x) - V(x)) \gamma_{2x-1} \gamma_{2x} + (w + \Delta_x) \gamma_{2x} \gamma_{2x+1} + (-w + \Delta_x) \gamma_{2x-1} \gamma_{2x+2}. \quad (3.3)$$

By using this construction, Kitaev realised that for certain model parameters and open boundary conditions, the bulk of the model dimerises, while individual isolated Majoranas are residing at the edges (see Fig. 3.2). This phase of the Kitaev model is referred to as the *topological* phase. When there are no isolated Majoranas the model is in the topologically trivial phase.

In order to characterise these two distinct topological phases and the exact topological phase transition between them more precisely, we can look at the continuum Bogoliubov-de Gennes (BdG) representation (App. A) of this model and solve for the quasiparticle mode

⁶We note that this lattice model can also be obtained from a Jordan-Wigner transformation of the XYZ Heisenberg spin chain, see App. B.

⁷Two Fermionic particles are needed for Cooper pairing.




FIGURE 3.2: The Kitaev Chain is a 1D lattice model of N sites (grey disks). Each site can be occupied or unoccupied by 2 Majorana modes (purple blobs), which together make up an ordinary fermion (dashed line). In the topologically trivial phase (top row) all Majoranas pair together and no individual Majoranas can be isolated. In the topological phase (bottom row) only the Majoranas in the bulk pair together and individual unpaired Majoranas can be found at the edges (for a chain with open boundary conditions). These unpaired Majorana modes can be used to make a topological qubit (Eq. 3.7).

energies. The continuum BdG Hamiltonian in the momentum k representation (with periodic boundary conditions) is given by

$$H_k = \begin{pmatrix} \frac{k^2}{2m} - \mu & i\Delta k \\ -i\Delta k & -\frac{k^2}{2m} + \mu \end{pmatrix} \quad (3.4)$$

in which m is the mass of the electrons. The mass is related to the lattice hopping parameter via $w = 1/2ma^2$, where a is the lattice constant. Similarly, $\mu_x = \mu + 2w + V_x$ and $\Delta = \Delta_x/(2a)$ relate the lattice chemical potential and superconducting gap from Eq. (3.1) to their corresponding continuum versions in Eq. (3.4). The BdG Hamiltonian can be diagonalised to find the mode dispersion

$$\epsilon_k^\pm = \pm \sqrt{\left(\frac{k^2}{2m} - \mu\right)^2 + \Delta^2 k^2}, \quad (3.5)$$

which consists of two energy bands; one for the holes (negative energies) and one for the particle modes (positive energy). We will denote the associated quasiparticle mode creation and annihilation operators for energy ϵ_k as β_k^\dagger and β_k .

The topological phase transition can be found by solving for the parameters at which the energy gap in the quasiparticle spectrum closes, namely when $\epsilon_k^\pm = 0$ [135]. When $\Delta = 0$, this leads to $\mu = |2w|$ (taking $a = 1$). The $\Delta = 0$ point separates the two distinct topological phases, one of which, from the discussion above, has unpaired Majorana modes residing at the edges. One way to distinguish between those phases, and characterise when Majoranas form, is to use the bulk-boundary correspondence [123]. This correspondence relates properties

from the bulk dispersion (closed system) to the existence of bound states. For this, one would need to calculate a so-called winding number (in 2D it is called the Chern number) [123], which gives for the Majorana chain an integer $\nu = 0$ or $\nu = \pm 1$ that labels the two distinct phases (similar to characterising Chern classes).⁸ When this bulk index is $\nu = \pm 1$, the model is in the topological phase with Majorana modes residing at the edges. The details of this derivation can be found in [123, 135].

Here, instead, we distinguish the two phases by looking at the energy spectrum for open boundary conditions on the lattice (see Fig. 3.5 for the slightly more general proximity-coupled superconductor). Since the Majorana operators γ_i are Hermitian, this means that in this system with particle-hole symmetry the quasiparticle energies of unpaired Majorana modes should have exactly zero energy $\epsilon_i = 0$. For the Kitaev chain model this happens when $\mu \leq |2w|$.

The wave function of the Majorana zero modes can be found by solving for the eigenvectors of the continuum Hamiltonian Eq. (3.4) with eigenvalue $\epsilon_k = 0$, i.e. $[H_k][\Gamma] = 0[\Gamma]$. By taking the two-component spinor ansatz, $\Gamma = [\phi, -\phi]^T$, and imposing open boundary conditions this leads to the left boundary solution⁹

$$\phi(x) \propto e^{-x/\xi} \sin\left(\sqrt{k_F^2 - 1/\xi^2}x\right). \quad (3.6)$$

Here we have defined the Fermi momentum $k_F = \sqrt{2\mu m}$ and correlation length $\xi = 1/(\Delta m)$. From this equation, we see that the Majorana mode is localised at the left edge with a decay length equal to ξ and is Hermitian $\phi^\dagger = \phi$. In Fig. 3.5 (b) we plot this mode at the left edge of the system (the p-wave superconducting wire).

To make a topological qubit from these localised Majorana modes we can use the associated degeneracy in the many-body spectrum. For a single qubit we require in total 4 Majorana modes, which we will label by Γ_i for $i = 1, 2, 3, 4$.¹⁰ We first pair these Majoranas together into two ordinary Dirac fermionic zero modes, $\beta_0^{a/b} = \frac{1}{2}(\Gamma_{1/3} + i\Gamma_{2/4})$. We can then define

⁸This \mathbb{Z}_2 integer labelling can also be derived from the fact that the Hamiltonian in Eq. 3.4 has particle-hole and chiral symmetry which puts it in class *DIII* of the table of topological insulators and superconductors [136, 137].

⁹Note that we get two solutions, one on the right boundary of the system and one on the left.

¹⁰How one can realise 4 localised Majorana modes in a network of Kitaev chains will be described below in section 3.3.

logical zero $|\bar{0}\rangle$ and one $|\bar{1}\rangle$ qubit states from the total even parity sector¹¹ as

$$|\bar{0}\rangle = \beta_0^a \beta_0^b \prod_{i=1}^{N-2} \beta_i |vac\rangle = |0\rangle \otimes |0\rangle \otimes |n_B\rangle, \quad (3.7)$$

$$|\bar{1}\rangle = (\beta_0^a)^\dagger (\beta_0^b)^\dagger \prod_{i=1}^{N-2} \beta_i |vac\rangle = |1\rangle \otimes |1\rangle \otimes |n_B\rangle, \quad (3.8)$$

in which $|vac\rangle$ is the vacuum of the c fermions. The vector $|n_B\rangle$ corresponds to the ground state of the (real space) quasiparticle bulk modes β_i and $|0(1)\rangle$ are the unoccupied (occupied) fermionic zero modes.

Both these states have the same energy since the Dirac fermionic modes $\beta_0^{a/b}$ have both zero energy $\epsilon_0^{a/b} = 0$. This means that the $|\bar{0}\rangle$ and $|\bar{1}\rangle$ span the degenerate ground state manifold that can be manipulated by braiding the individual Majoranas. To perform quantum computation with Majoranas one thus needs to be able to move (control) them in real space. In section 3.3 we will discuss proposals for how this can be done by changing the potential profile $V(x)$. Before this, in the next subsection, we will look at how the Kitaev Chain toy-model can be realised in real world proximity-coupled superconductors.

3.2.2 The proximity-coupled Superconductor

Proximity-coupled superconductors are a class of materials in which an ordinary s-wave superconductor is placed on top of another (non-superconducting) material, such as a topological insulator, to effectively make it superconductive near the interface, see e.g. Fig 3.3 (a). This effect can be understood at a phenomenological level from the Ginzburg-Landau theory, which states that the superconducting order parameter cannot change discontinuously over distances shorter than the superconducting coherence length [133]. At a microscopic level the mechanism behind this effect is the diffusion of Cooper pairs from the parent superconductor into the coupled material that can be described by the theory of Andreev reflection [138, 139].

In 2009 Fu and Kane [54] established that the proximity effect can be used to realise an effective version of a 2D topological $(p_x + ip_y)$ superconductor.¹² Following this initial proposal, in references [53, 52] it was shown that by proximity coupling a semiconducting nanowire to an s-wave superconductor one obtains, in certain limits, the 1D Kitaev chain model. As these proximity-coupled setups can be build in the laboratory, experimentalists then started looking for the presence of Majorana zero modes in these systems. A few years later, in 2012, the first signatures for their existence were found in [55]. This was the first step towards the realisation of a real world topological quantum computer.

¹¹By parity we mean the total number of Fermionic modes in the many-body state. Note that we can not use two Majorana modes (1 Dirac mode) since the Fermion parity in a superconductor is only preserved modulo two.

¹²This is the 2D generalisation of the Kitaev chain discussed in the last section.

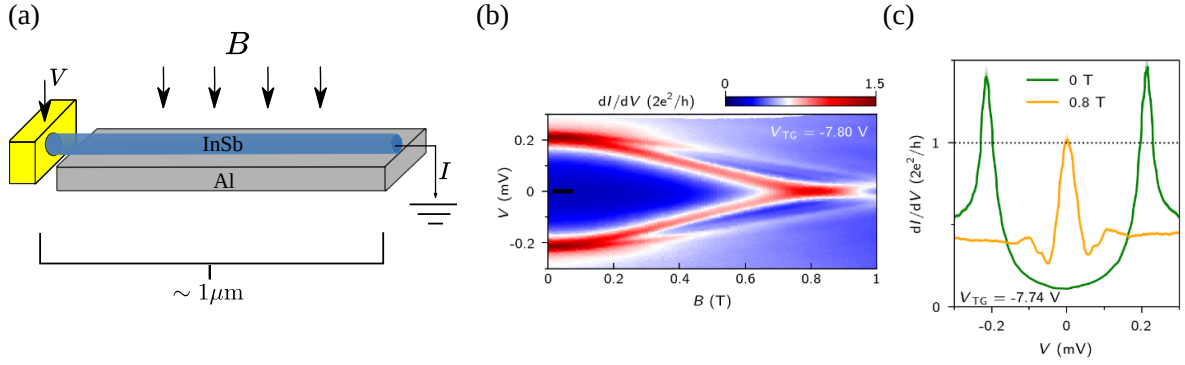


FIGURE 3.3: (a) The proximity-coupled superconductor used in the experiments for detecting Majorana zero modes. In this setup a spin-orbit coupled semiconducting nanowire (InSb) is placed on top of a s-wave superconductor (Al). The nanowire becomes an effective p-wave superconductor when an external magnetic field B is applied. For the tunnelling transport experiments a voltage, V , is applied to the gate electrode (yellow) and the current I is measured at the end of the wire. (b) Experimental results for the differential conductance, dI/dV , as a function of the magnetic field. (c) A zero-bias peak in the experimental differential conductance occurs when the magnetic field is increased to $B = 0.8$, a signature of the presence of Majorana zero modes. Figures (b) and (c) are taken from [140].

We will now introduce the microscopic model for the specific proximity-coupled superconductor used in this experiment. Afterwards we then show how to obtain the Kitaev chain from this model and provide a discussion on the latest experimental developments.

Model and Dispersion

In the experiments reported in references [55, 140] a semiconducting nanowire (NW) with strong spin-orbit coupling was proximity coupled to a s-wave superconductor and placed in an external magnetic field, as shown in Fig. 3.3 (a). The continuum model for this setup [53, 52] is given by the Hamiltonian

$$\mathcal{H}^{\text{NW}} = \frac{1}{2} \int dk \Psi_k^\dagger H_{\text{BdG}}^{\text{NW}}(k) \Psi_k \quad (3.9)$$

with (Nambu) spinors $\Psi_k^\dagger = [c_{k\uparrow}^\dagger, c_{k\downarrow}^\dagger, c_{-k\uparrow}, c_{-k\downarrow}]$ and the BdG (App. A) Hamiltonian

$$H_{\text{BdG}}^{\text{NW}}(k) = \left(\frac{k^2}{2m} - \mu + V + B\sigma_x + \alpha k\sigma_y \right) \tau_z + \Delta\sigma_y\tau_x. \quad (3.10)$$

Here the c_{ks} (c_{ks}^\dagger) are fermionic annihilation (creation) operators for particle modes of momentum k and hole modes of momentum $-k$, which have either up spin $s = \uparrow$ or down spin $s = \downarrow$. These operators satisfy the fermionic anti-commutation relations $\{c_{ks}^\dagger, c_{k's'}\} = \delta_{kk'}\delta_{ss'}$ and $\{c_{ks}, c_{k's'}\} = 0$. The σ_i and τ_i for $i = \{x, y, z\}$ represent the Pauli spin matrices, which act on the spin-degrees of freedom (σ_i) and on the particle-hole space (τ_i).

The first term of the BdG Hamiltonian is a standard kinetic energy term for free electrons (and holes) of mass m in an external potential V and chemical potential μ . The term proportional to $B \equiv \frac{\mu_B g B_x}{2}$ describes the effect of the external magnetic field, which is coupled to the spin of the particles. α is the strength of the (Rashba) spin-orbit coupling [141] in the semi-conducting nanowire. Lastly, the s-wave superconductivity is included at a mean-field level with superconducting gap Δ . We note that Δ is the superconducting gap within the nanowire induced by the proximity-coupled parent superconductor, which itself has a different gap.

The energies of the quasiparticles (excitations) of the system can be found by diagonalising the BdG Hamiltonian $H_{\text{BdG}}^{\text{NW}}(k)$ in Eq. (3.10). This results in the energy dispersion

$$\epsilon_{\pm}^2(k) = \xi_k^2 + B^2 + \alpha^2 k^2 + |\Delta|^2 \pm 2\sqrt{B^2|\Delta|^2 + \xi_k^2 B^2 + \alpha^2 k^2 \xi_k^2}, \quad (3.11)$$

in which we have defined $\xi_k = \frac{k^2}{2m} + V - \mu$. This dispersion relation describes four energy bands, which are plotted for various choices of the model parameters in Fig. 3.4. In the absence of the external field B , spin-orbit coupling α and the proximity effect Δ we have the familiar quadratic dispersion for free fermions in which the two different spin components are degenerate. The spin degeneracy is broken when B is turned on (panel b) due to the Zeeman effect [85]. The main effect of the spin-orbit coupling α , in addition to breaking the conservation of spin, is a shift of the Fermi momentum k_F . Here, k_F is defined to be the momentum k where $|\epsilon_k|$ is minimised. Lastly, Δ opens up a small energy gap between the conduction and valence bands (panel d). This gives rise to the characteristic dispersion for a p-wave superconductor for the middle two bands.

The Kitaev Chain Limit and Majorana Zero Modes

The correspondence between the middle energy bands of the proximity-coupled superconducting nanowire setup, Eq. (3.10), and the dispersion, Eq. (3.5), of a p-wave superconductor can be made explicit by considering certain limits of the model parameters. The first limit that achieves this consists of making the magnetic field very large, $B \gg \alpha, \Delta$, such that the bands corresponding to the different spin components can be effectively decoupled [142, 143]. Formally this can be seen by expressing $H_{\text{BdG}}^{\text{NW}}(k)$ in the eigenbasis of the Hamiltonian $H_{\text{BdG}}^{\text{NW}}(k)|_{\Delta=0}$. This results in

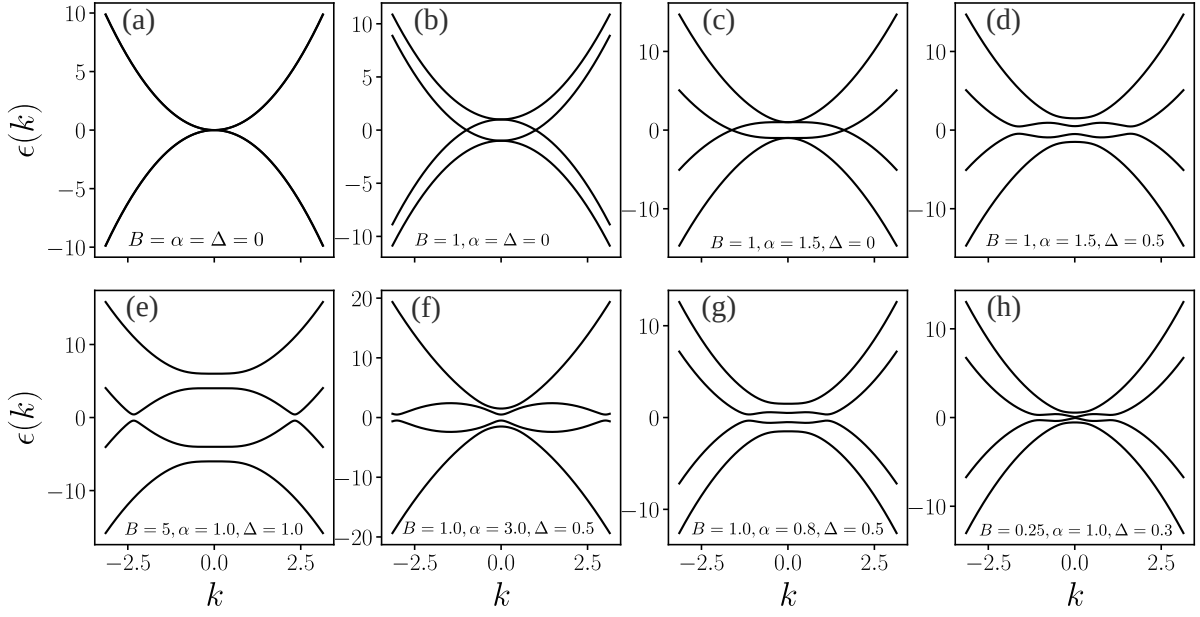


FIGURE 3.4: Energy dispersion for different parameters of the Hamiltonian Eq. 3.10 of the proximity-coupled semiconducting nanowire in a Zeeman field. The top row (a-d) shows the contributions from the various terms in the Hamiltonian. The two bottom left plots (e, f) are for parameters that recover the Kitaev chain limit for which we can estimate the energy gap. The bottom right plots (g, h) are for parameters outside this limit and we need to resolve the size of the gap numerically.

$$H^{\text{NW}}(k) = \begin{pmatrix} -\sqrt{B^2 + k^2\alpha^2} - \xi_k & 0 & \frac{-iB\Delta}{\sqrt{B^2 + k^2\alpha^2}} & \frac{k\alpha\Delta}{\sqrt{B^2 + k^2\alpha^2}} \\ 0 & \sqrt{B^2 + k^2\alpha^2} - \xi_k & \frac{-k\alpha\Delta}{\sqrt{B^2 + k^2\alpha^2}} & \frac{iB\Delta}{\sqrt{B^2 + k^2\alpha^2}} \\ \frac{iB\Delta}{\sqrt{B^2 + k^2\alpha^2}} & \frac{-k\alpha\Delta}{\sqrt{B^2 + k^2\alpha^2}} & -\sqrt{B^2 - k^2\alpha^2} + \xi_k & 0 \\ \frac{k\alpha\Delta}{\sqrt{B^2 + k^2\alpha^2}} & \frac{-iB\Delta}{\sqrt{B^2 + k^2\alpha^2}} & 0 & \sqrt{B^2 + k^2\alpha^2} + \xi_k \end{pmatrix}. \quad (3.12)$$

In the limit $B \gg \alpha, \Delta$, the terms proportional to $\propto iB\Delta$ that couple the different spin components become ineffective¹³ and we can project onto the middle two bands to obtain¹⁴

$$H_{\text{eff}}^{\text{NW}}(k) \approx \begin{pmatrix} \sqrt{B^2 + k^2\alpha^2} - \xi_k & \frac{-k\alpha\Delta}{\sqrt{B^2 + k^2\alpha^2}} \\ \frac{-k\alpha\Delta}{\sqrt{B^2 + k^2\alpha^2}} & -\sqrt{B^2 - k^2\alpha^2} + \xi_k \end{pmatrix}. \quad (3.13)$$

This effective Hamiltonian, $H_{\text{eff}}^{\text{NW}}(k)$, matches the BdG Hamiltonian, Eq. (3.4), for the Kitaev Chain with effective p-wave superconducting gap $\Delta_{\text{Kit}} \approx -\frac{\alpha\Delta}{B}$. To confirm this result, in panel (e) of Fig. 3.4 we plot the energy dispersion of the proximity-coupled superconductor in the high magnetic field limit. The different spin components clearly decouple and the middle

¹³Due to the large energy separation of the bands that these terms couple, see e.g. Fig. 3.4 (e).

¹⁴Note that projecting on the two other bands gives the same result.

bands show a 'w'-like p-wave dispersion with the minimum energy gap appearing at k_F .

Another limit that results in a p-wave superconductor is the strong spin-orbit coupling limit $\alpha \gg B, \Delta$. The derivation for this limit is similar to the high magnetic-field limit and is discussed in reference [144]. For this limit the effective Kitaev chain Hamiltonian has a p-wave superconducting gap $\Delta_{\text{Kit}} \approx -\frac{\Delta}{m\alpha}$. In panel (f) of Fig. 3.4 we show the energy dispersion in this limit. We observe that the spins do not decouple for small k , but near k_F there is a large energy separation between the different spin bands. The middle two bands have again the characteristic p-wave dispersion.

In both of these limits Majorana zero modes can appear at the domain boundaries between topological and non-topological regions in the proximity-coupled nanowire. However, these are not the only parameters for which there are localised Majorana modes [52]. For example, two other parameter regimes for which the model is in a topological phase are shown in panels (g) and (h) of Fig. 3.4. Crucially, all terms in the nanowire model in Eq. (3.9) need to be there for Majorana modes to exist. In order to check if for certain choices of the model parameters we have Majorana zero modes we can diagonalise the BdG Hamiltonian on the lattice with open boundary conditions and, like for the Kitaev chain, look for zero energy solutions. An example of this is shown in Fig. 3.5 (b).

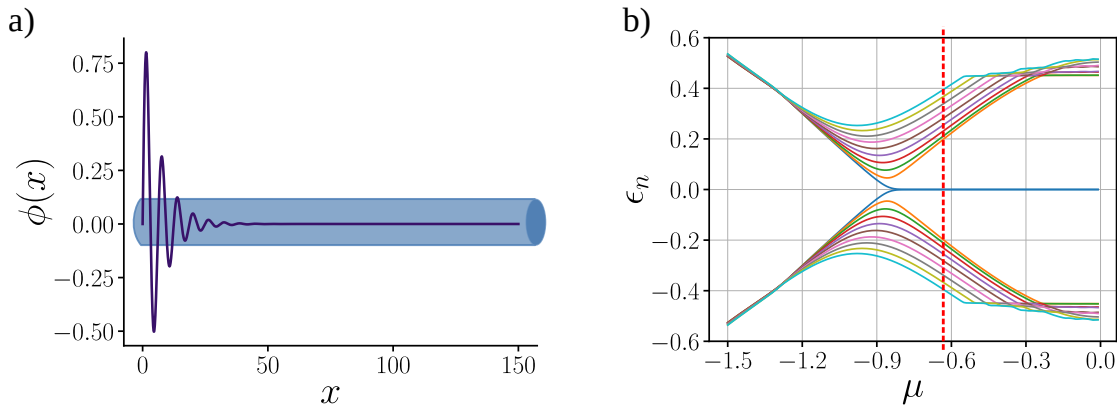


FIGURE 3.5: (a) Localisation of the Majorana zero mode (Eq. 3.6) at the left edge of the wire. Note that we have not plotted the Majorana mode localised at the right edge of the wire. (b) Excitation spectrum as a function of the chemical potential μ for the lowest energies of the BdG Hamiltonian (Eq. 3.9) for the proximity-coupled nanowire model on the lattice. For approximately $-0.8 \leq \mu \leq 0.8$ we get two degenerate zero energy Majorana modes (blue line). When μ is close to the bottom of the band, $\mu \approx -0.9$, we get also low energy Andreev bound states (orange lines). These Andreev states could cause the appearance of non-topological zero-bias peaks in experiments [58]. The mode energies were obtained by numerically diagonalising the Hamiltonian with parameters outside the Kitaev Chain limits $B = \alpha = 1.0$, $\Delta = 0.5$.

3.2.3 Experimental Implementation

The proximity-coupled nanowire model can be used to look for experimental signatures of Majorana zero modes [55, 56]. There are various types of measurements that one can perform, but so far finding definite proof for the existence of a topological Majorana zero modes has been very challenging. One signature is that in the presence of Majorana zero modes there is a 4π (fractional) Josephson effect [145]. As a consequence, observable quantities like the current through a Josephson junction $I_{\text{jos}}(\phi)$ made of two p-wave superconducting nanowires are 4π periodic as a function of the flux (superconducting phase difference) ϕ . One then needs to measure the current, while changing the applied flux ϕ slowly.

However, a huge challenge is that when this is done too slowly, an effect known as quasi-particle poisoning [146, 147] due the coupling to the environment (or even from the parent superconductor) kicks in. This effect destroys the 4π periodicity of observable quantities back to the normal 2π Josephson periodicity. In contrast, doing it very fast is also not an option because changing the phase too quickly could lead to Landau-Zener tunnelling [148] through small energy gaps. This, incorrectly, can give a 4π effect, while the material under study really has a 2π Josephson effect. Despite these challenges, recently evidence for the 4π Josephson effect was reported in an experiment [149].

Another, arguably more robust, signature is a zero-bias peak in the measured tunnelling conductance G through the Majorana wire [140, 150]. This type of measurement allows one to observe the presence of subgap states at low energies, such as the Majorana zero modes. To probe this one applies a bias voltage V to a lead connected to one end the nanowire and measures the current I at the other end, see Fig. 3.3 (a). The tunnel conductance is then given by $G = dI/dV$. For Majoranas G is quantised to $2e/h$ when $V = 0$ due to a perfect Andreev reflection at the lead-nanowire interface [151].¹⁵ These measurements were performed in several experimental studies [55, 140] and, as shown in Fig. 3.3 (b) and (c), indeed show a conductance peak for zero bias when the external magnetic field B is large enough.

Although these observations are a strong indication that Majoranas are present in this system, there are also other mechanisms that could explain such zero-bias peaks. In reference [58] several different effects were theoretically investigated and the authors identified three different regimes, the "good, bad and ugly", for the detection of Majorana zero modes. The good regime is where there is an actual topological Majorana, which is the theoretical scenario we have discussed above in the proximity-coupled superconductor. A smooth confining potential [152] or the presence of unintended quantum dots due to imperfect manufacturing of the device could, however, also lead to a conductance peak, not always exactly quantised at $2e/h$, caused by low energy Andreev bound states. This possibility can be circumvented by increasing the external magnetic field B and is therefore dubbed the bad regime. An even worse scenario is when there is inhomogeneous disorder in either the chemical potential or

¹⁵This exact quantisation is a signature of the nontrivial topology.

the magnetic field. This is called the ugly regime since disorder cannot be removed easily. After an extensive comparison between the experimental results and numerical simulations of the setup with disorder it was concluded [58] that likely most, if not all, observations of a conductance peak could be explained by the presence of disorder and hence are in the ugly regime.

Many improvements in the setup or measurements thus need to be made before the existence of Majorana modes can be proven. However in this thesis we take a positive (theoretical) perspective and assume that we have access to localised Majorana zero modes. We highlight that the quantum control methods we use in Chapter 4 and the conclusions for error processes we draw in Chapter 5 are generic and can be applied to more complex setups, which have the discussed sources of error from real-world devices.

3.3 Moving Majorana Zero Modes

To perform quantum computation with Majorana zero modes in superconducting nanowires one needs to move (braid) the individual Majoranas around each other, while remaining within the protected ground-state manifold. A way to achieve this is by attaching a so-called 'keyboard' of gate electrodes to the Majorana wire [153] as shown schematically in Fig. 3.6 (a). A potential V_i can be applied to each individual gate electrode to control the chemical potential (electron density) in the wire. These gate potentials can be used to create domain boundaries between topological ($\mu + V_i < |2w|$) and non-topological ($\mu + V_i > |2w|$) regions. By tuning the individual V_i in time it is possible to move such domain walls and hence the Majoranas.




FIGURE 3.6: (a) Majoranas γ_L and γ_R localised at domain boundaries between topological (darker blue) and non-topological (lighter blue) regions in the nanowire can be moved by tuning a 'keyboard' of gate potentials applied to the wire. (b) Instead of modelling the individual discrete gates we use a smooth potential, Eq. (3.14), that can be tuned to move the Majoranas in our simulations. (c) The wires (and gate potentials) can be deployed in 2D network architectures to exchanges Majoranas. In the specific example shown three moves are required to exchange the middle two Majoranas γ_2 and γ_3 .

Modelling analytically or numerically the individual gate electrodes separately is rather cumbersome.¹⁶ Here, instead, we will utilise the continuous inhomogeneous potential profile given by

$$V(x) = V_{\text{height}}[f(x - x_L) + f(x_R - x)], \quad (3.14)$$

in which $f(x) = 1/(1 + \exp\{x/\sigma\})$ are sigmoid (step) functions as shown in Fig. 3.6 (b). V_{height} is the height of the potential and is chosen to lie far outside the topological phase $V_{\text{height}} \gg |2w|$. This means that x_L and x_R represent the positions of two domain boundaries. The idea is that, with a fine enough 'keyboard' of gate electrodes, it should be possible to approximate this potential in experiments. We note that for numerical simulations this potential $V(x)$ will go into the lattice Hamiltonian in Eq. (3.1).

However, using the potential to control the position of Majoranas in a single one-dimensional nanowire is not enough to be able to braid Majoranas in real space since they cannot be moved through each other. This issue can be resolved by deploying the individual nanowires in 2D wire networks, such as the example shown in Fig. 3.6 (c). In these architectures it is possible to move the Majoranas for a small distance in the extra dimension (vertical wires) so to effectively exchange them. Another option is to use the two dimensional generalisation of the Kitaev model, the $p + ip$ superconductor [135]. In here topological defects such as flux vortices can host Majoranas. These vortices can potentially be moved more easily around each other.

For fault-tolerant quantum computation the system should be in the ground state at the end of the braiding process. Moving the Majoranas too fast, however, can result in unwanted excitations outside the ground-state manifold and a loss or corruption of the encoded quantum information. A quantitative measure for how much quantum information is excited at the end of the motion is the infidelity given by

$$\mathcal{I}_\tau \equiv 1 - \mathcal{F}_\tau = 1 - |\langle \psi_\tau | \mathcal{T} e^{-i \int_0^\tau \mathcal{H}(t) dt} | \psi_0 \rangle|^2. \quad (3.15)$$

We saw this measure before in the single qubit control problem in section 2.1.1. This time, the Hamiltonian $\mathcal{H}(t)$ used for the evolution can either be the Kitaev chain of Eq. (3.1) or the proximity-coupled nanowire of Eq. (3.10). The Hamiltonian is time dependent because we are changing the potential profile $V(x)$ in time. The initial state $|\psi_0\rangle$ is the ground state of the system $\mathcal{H}(0)$ and $|\psi_\tau\rangle$ the target ground state of $\mathcal{H}(\tau)$ after the movement. \mathcal{I}_τ is zero when the target state is reached perfectly and is equal to one when it is orthogonal to it.

One way to exactly reach the target state is by moving the Majoranas adiabatically, which means the system remains in the instantaneous ground state throughout the entire motion. However, in real-world experimental setups this is not a desirable strategy because of noise

¹⁶It can be done, see for example [154].

from the environment. For long movement times τ processes like quasiparticle poisoning and thermal noise are able to decohere the encoded quantum information. Trying to move faster than adiabatically on the other hand is tricky and generically increases the infidelity when not the right strategy is chosen. Nevertheless it turns out to be possible to move faster in this topological system by a concept known as superadiabaticity [155, 144]. Superadiabaticity is a notion of adiabaticity in moving frames. We will discuss and investigate this further in the next two subsections.

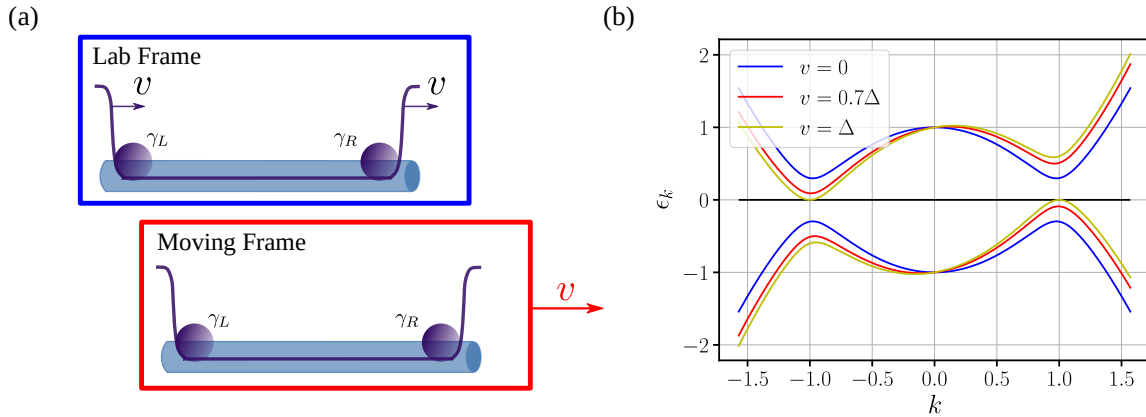


FIGURE 3.7: (a) The left and right domain boundary in a single wire are being moved at a constant velocity v . In the lab frame (blue box) it looks like the localised Majoranas are moving with velocity v and the wire itself is stationary. In the moving frame (red box) we move along at the same speed as the motion of the domain boundaries and it looks like everything is stationary. (b) The energy dispersion ϵ_k of the Kitaev chain in the moving frame for three different velocities. When the velocity approaches the critical velocity $v_{\text{crit}} = \Delta$ the energy gap closes and we have a topological phase transition in the moving frame.

3.3.1 Moving Frame

Moving at a constant velocity is intimately tied with the concept of a moving reference frame. In this setting one distinguishes between a so-called *lab frame* and a reference frame moving at some constant velocity v . In our context, when we are moving both domain walls in the nanowire with a constant velocity v , this means that in the lab frame it looks like we are moving the Majoranas, whereas in the moving frame everything seems static, see Fig. 3.7 (a). Switching between such reference frames can be useful for evaluating the effect of the motion on the Majoranas. For example, from this we can find that above a certain critical velocity, v_{crit} , the spectral gap of the Hamiltonian in the moving frame closes and the topological phase is destroyed. In the following we will derive this critical point formally.

In order to bring the system in the moving frame with velocity v (in non-relativistic quantum mechanics) we can apply the time-dependent (Galilean) translation operator $U(t) = e^{-ik \int_0^t v(t') dt'}$ to the system Hamiltonian. This results, for the Kitaev chain, in the moving-frame Hamiltonian

$$H_v(t) = U^\dagger(t)H_kU(t) + i\frac{dU^\dagger}{dt}U(t) = \begin{pmatrix} \frac{k^2}{2m} - \mu + v(t)k & i\Delta k \\ -i\Delta k & -\frac{k^2}{2m} + \mu + v(t)k \end{pmatrix}. \quad (3.16)$$

We see that compared to the lab-frame Kitaev chain Hamiltonian of Eq. (3.1) there is an extra contribution $v(t)$ on the diagonal. This leads to a modified mode dispersion

$$\epsilon_k = \pm \sqrt{\left(\frac{k^2}{2m} - \mu\right)^2 + \Delta^2 k^2} + vk, \quad (3.17)$$

which is plotted in Fig. 3.7 (b). When the velocity approaches the critical velocity $v = v_{\text{crit}} \approx \Delta$ the gap in the excitation spectrum closes indicating a topological phase transition.

At this transition point the Majoranas delocalise in the moving frame. This can be seen by solving for the zero energy quasiparticle modes of the moving-frame Hamiltonian of Eq. (3.16), which gives

$$\phi(x) \propto e^{-x/(\gamma\xi)} \sin\left(\sqrt{k_F^2 + 1/(\gamma\xi)^2}x\right). \quad (3.18)$$

Note that this equation matches the lab-frame Majorana localisation in Eq. (3.6) when we set $v = 0$. The effective Majorana coherence length is now given by $\xi_v = \gamma\xi = \gamma/\Delta m$, which is dilated by a factor $\gamma = 1/\sqrt{1 - \frac{v^2}{\Delta^2}}$. When we approach the critical velocity, Δ , the localisation length will diverge $\xi_v \mapsto \infty$ and the local character of the Majoranas will be lost.¹⁷ Superadiabaticity can now be explained as the ability to move the Majoranas at finite velocities below v_{crit} , which means it is adiabatic in the moving frame.

3.3.2 Susceptibility to boundary oscillations

The ability to move at finite velocities does not mean that it is possible to instantly start moving the domain walls at a large velocity¹⁸ without losing some fidelity. For superadiabaticity to work it is important to first reach the moving-frame ground state, which can take some finite amount of time. In the next chapter we will discuss several strategies for this in more detail. Here we briefly review some important related concept, which we call resonant Majorana motion.

This concept can be understood by considering the scenario in which we are shuttling the position of the left domain wall $x_L(t)$ in Eq. (3.14) with a certain fixed frequency ω according

¹⁷Note that this can also be shown for the more experimentally relevant proximity-coupled semiconducting nanowire model [144]. Here within the high magnetic field Kitaev chain limit the critical velocity is given by $v_{\text{crit}} \approx \frac{\alpha\Delta}{B}$. And in the strong spin-orbit coupling limit the critical velocity is $v_{\text{crit}} \approx \frac{\Delta}{m\alpha}$. Outside these limits v_{crit} needs to be determined numerically.

¹⁸But below critical v_{crit} .

to the velocity protocol given by

$$v_L(t) = v_{\max} \sin \omega t. \quad (3.19)$$

When the maximum velocity, v_{\max} , of the motion is small enough, such that the amplitude¹⁹ is tiny compared to the total length of the system N , we can consider this motion as a perturbation on top of the static Hamiltonian \mathcal{H}_0 with fixed domain wall positions $x_L(0)$ and x_R . Applying perturbation theory, specifically Fermi's golden rule [156, 85], then leads to the following rate of fidelity loss

$$\lim_{\tau \rightarrow \infty} \frac{\mathcal{I}_\tau}{\tau} = \lim_{\tau \rightarrow \infty} \frac{1}{\tau} [1 - |\langle \psi_0 | \exp \left\{ -i \int_0^\tau \mathcal{H}(t) dt \right\} | \psi_0 \rangle|^2], \quad (3.20)$$

$$= \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \sum_{i \neq 0} |\langle \psi_i | \exp \left\{ -i \int_0^\tau \mathcal{H}(t) dt \right\} | \psi_0 \rangle|^2, \quad (3.21)$$

$$= 2\pi \sum_{i \neq 0} (\delta(E_i - E_0 + \omega) + \delta(E_i - E_0 - \omega)) |\langle \psi_i | \delta V | \psi_0 \rangle|^2. \quad (3.22)$$

Here, we have expanded in terms of the eigenstates $|\psi_i\rangle$ of \mathcal{H}_0 , and δV is the amplitude of the perturbation. From this equation we can see that when ω is exactly equal to the energy difference between the ground state, E_0 , and a specific excited state, E_i , there is a resonance. As we explained above, within the topological phase, the system is gapped with gap size $E_{\text{gap}} \equiv E_1 - E_0 = \Delta k_F$ for the Kitaev chain. Thus, when oscillating the Majorana with $\omega_{\text{res}} = \Delta k_F$, we obtain a resonance in the infidelity.

Below the resonance frequency $\omega < \omega_{\text{res}}$ no resonance appears because these frequencies are below the size of the energy gap. This low frequency regime is desired when one tries to accelerate the system from a static ground state into the moving-frame ground state at a finite velocity. The other limit is the high frequency limit $\omega \mapsto \infty$, which seems useful since there is no fidelity loss for back and forth shuttling of the Majorana. However, this can be explained because these frequencies are so high that the system does not have time to respond to the motion and, hence always remains in its initial state.

As in the following chapters we are interested in forward motion of Majoranas, we verify that in this case the same resonance occurs. For this we move the left domain wall with velocity

$$v_L(t) = v_{\max} \frac{1 - \cos \omega t}{2}, \quad (3.23)$$

and compute \mathcal{I}_τ numerically. The results of these simulations for a wide range of frequencies and maximum velocities are shown in Fig. 3.8. We clearly observe a resonance in the infidelity when the frequency is equal to ω_{res} . Furthermore, for small enough $\omega < \omega_{\text{res}}$ and $v_{\max} \ll v_{\text{crit}}$ we obtain near perfect fidelities for the Majorana motion.

¹⁹Maximum distance from the starting position of the wall $x_L(0)$.

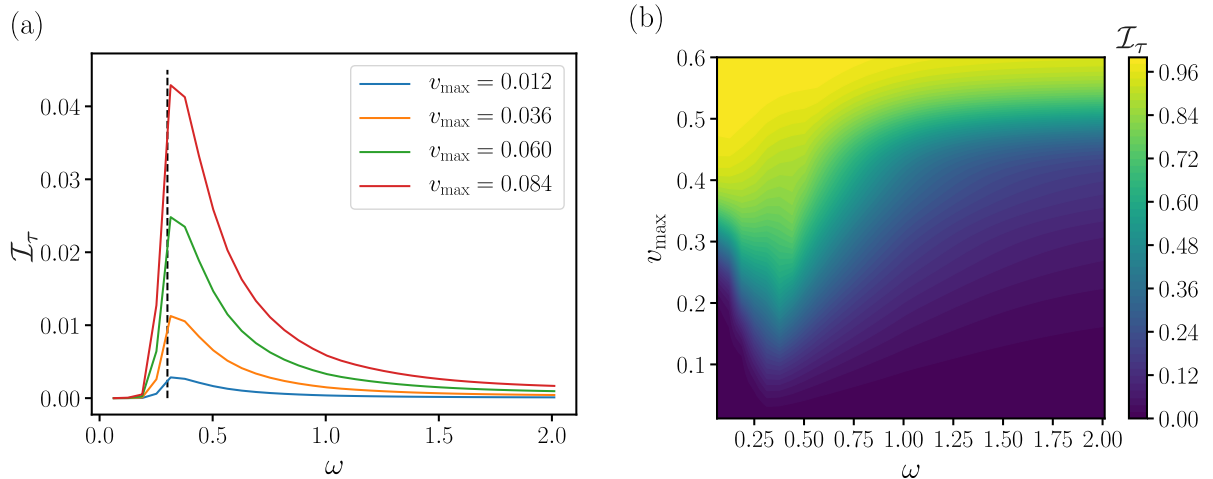


FIGURE 3.8: (a) Shuttling of the left Majorana in a single nanowire leads to a resonance in the infidelity when the frequency is approximately equal the size of the energy gap $\omega_{\text{res}} = \Delta k_F$ (black dashed line). (b) This resonance remains visible when the maximum velocity v_{\max} is increased up to $v_{\text{crit}} = \Delta = 0.3$. Above v_{crit} there are no slow adiabatic frequencies (with respect to the gap) left to oscillate at. The parameters used in these simulations are $N = 140$, $\mu = 1$, $w = 1$, $\Delta = 0.3$ and $\sigma = 1$.

3.4 Summary

In this chapter we have introduced quantum computation with non-abelian anyons. After a general introduction we have focused on Majorana zero modes in topological p-wave superconductors. We saw that these Majoranas need to be moved (braided) to make quantum gates. For fault-tolerant quantum computation the Majoranas need to be transported with perfect fidelity. However, adiabatic or superadiabatic transport might not be the best strategy due to various sources of decoherence. In the next chapter we will look at how the traditional quantum control and machine-learning algorithms introduced in Chapter 2 can be exploited to come up with alternative strategies for moving Majoranas with high fidelities.

Chapter 4

Optimising Transport of Majorana Zero Modes

The ability to control Majorana zero modes and more generally non-abelian anyons is, as we saw in the last chapter, of critical importance for the development of real-world topological quantum devices. The Majorana modes, from which we construct a topological qubit, need to be moved around each other or fused together to be able to perform computations. By definition, adiabatic transport of Majoranas in topological superconductors by tuning the external gate potential $V(x)$ is a slow process, and is therefore susceptible to noise such as quasiparticle poisoning [146]. This motivates the search for alternative control strategies that are faster and more efficient.

A different approach is to use schemes where the Majoranas do not need to be transported over excessively long distances. This includes protocols where one aims to braid the Majoranas by controlled tunnelling [157] or by changing the interaction strength [158, 159, 160]. In addition, there are the so-called measurement-only braiding schemes in which one applies topological charge measurements to implement the logical quantum gates [161, 160]. However, all these methods require the ability to precisely tune specific parameters in experiments, namely the tunnelling rates or the interaction and measurement strengths. Thus, in each case the specific form of the optimal control pulses needs to be determined. This means that these schemes are possibly still limited by adiabatic time constraints and in some cases shortcuts to adiabaticity need to be applied [162]. Furthermore, for the measurement-only schemes an additional drawback is that the state manipulation is inherently probabilistic [163].

Instead of looking for completely different braiding schemes one could also try to see if the (adiabatic) time scales for the movement of Majoranas can be shortened by using some of the techniques for quantum control introduced in Chapter 2. Such an optimal control approach for the physical movement of Majorana zero modes was taken, for example, by Karzig et al. in reference [83]. The authors applied Pontryagin's principle to a specific Majorana edge setup and showed that the optimal transport protocol is of the bang-bang form. The specific bang-bang protocol was then obtained by applying a similar simulated annealing (SA) search as we did for the single-qubit problem presented in section 2.2.3. Importantly, we note that

for this argument to work one requires the Hamiltonian to be linear in the control parameter.

In this chapter, we build upon such approach, using numerical optimal control techniques to design fast and high-fidelity transport protocols for the movement of Majoranas.¹ Specifically, we exploit Differentiable Programming (∂P) and the Natural Evolutionary Strategies (NES) optimisation algorithm for this task. With these methods we find alternative strategies that, for the range of parameters and searches we did, perform better than standard parameterised bang-bang protocols. For short transport times, τ , and relatively long transport distances, l , we find a novel counter-intuitive strategy, dubbed *jump-move-jump* (JMJ) and shown in Fig. 4.6, which makes use of pulse-like jumps at the beginning and end of the protocol. In between the jumps there is a period of motion at a nearly constant velocity. For long (super)adiabatic time scales the methods recover the expected smooth adiabatic protocols from previous studies [144, 164].

To capture (all) these different behaviours for the control strategies as a function of the transport time τ we propose a new characterisation in terms of different Majorana motion control regimes based on the resonance times scale T_{res} and critical velocity v_{crit} . As these quantities only depend on the size of the topological energy gap E_{gap} we argue that this categorisation applies both to the Kitaev chain as well as the proximity-coupled semiconducting nanowire model. We provide an understanding of the (physical) working mechanisms behind the various strategies in these different regimes. We find that the small jumps in the position of the Majorana in the JMJ strategy do not lead to a large decrease in the ground-state fidelity and that during the *move* part the system is stable when evaluated in a moving frame. Based on this, we derive an approximate equation for the ground state fidelity of the JMJ strategy that allows us to predict the value and its performance for a wide range of parameters (l, τ).

We have structured this chapter as follows. In section 4.1, we introduce the specific Majorana quantum-control setup and show that the average velocity constraint leads to distinct control regimes that connects (on a physical level) with the moving-frame picture as described in the last chapter. We give the specific mathematical objective function to which the numerical optimisation algorithms can be applied and also explain the problem from a reinforcement learning (RL) perspective. In section 4.2 we then show the optimisation results obtained with ∂P and NES and benchmark them with Majorana transport protocols obtained with the traditional (numerical) quantum control-methods, specifically SA. Afterwards in section 4.3, we discuss and explain the novel JMJ strategy that the optimisation algorithms came up with by evaluating it in a moving frame and applying a sudden approximation. Here we also look at the robustness of the protocol with respect to an inhomogeneous (disordered) chemical potential in the Majorana wire. Finally, in section 4.5 we summarise and give an outlook for further work.

¹The results presented in this chapter have been published in [73] coauthored by the author of this thesis.

4.1 Setup of the Optimisation Problem and the Different Majorana Control Regimes

We start by introducing the mathematical formulation of the Majorana control setup as shown schematically in Fig. 4.1. Our objective is to find the optimal control $x_L(t)$ to move the left Majorana in a single nanowire from a position $x_L(0) = x_A$ to a position $x_L(\tau) = x_B \equiv x_A + l$ in a time τ by tuning the external potential $V(x, t)$, see Eq. (3.14). As discussed in the last chapter, for fault-tolerant quantum computation we want to end up in the ground state of the system. A natural measure for the quality of the Majorana transport is thus the infidelity,²

$$\mathcal{I}_\tau = 1 - |\langle \psi_B | U(\tau) | \psi_A \rangle|^2, \quad (4.1)$$

in which $|\psi_A\rangle$ and $|\psi_B\rangle$ are the ground states of the wire with the Majorana at the positions x_A and x_B respectively. $U(\tau)$ is again the time-ordered unitary evolution operator that describes the time evolution of the system. In the following we take \mathcal{I}_τ as our cost (objective) function, which we wish to minimise. We consider both the Kitaev chain model of Eq. (3.1) and separately the proximity-coupled superconductor of Eq. (3.10) as the underlying system Hamiltonian $\mathcal{H}(t)$ to compute this infidelity.

For the application of the NES and ∂P methods we recall from Chapter 2 that we can also formulate the optimisation problem in RL language. In the RL formulation the Majorana nanowire, together with the potential profile, corresponds to the *environment*, which is manipulated by an *agent* (controller). This agent is able to perform *actions* and receive feedback in the form of *rewards* and *observations* of the environment. In our specific Majorana setting the allowed actions of the agent are a continuous modification of the position of the left domain boundary, $x_L(t)$, in the potential profile, $V(x, t)$. The only feedback the agent receives is the final infidelity \mathcal{I}_τ for the NES algorithm, and the gradient with respect to the control $\frac{d\mathcal{I}_\tau}{dx_L}$ for ∂P .³ This feedback is then used to update the control $x_L(t)$ and gradually minimise the infidelity \mathcal{I}_τ over a large number of iterations. This can be interpreted as the agent learning from past experiences by trying many different paths for the Majorana as shown in Fig. 4.1 (b).

An important constraint in this optimisation problem is the average velocity $v_{\text{avg}} = l/\tau$ of the Majoranas, which we set externally when we fix the termination time τ and movement length $l = x_B - x_A$. After some initial optimisation trials we found that this constraint dictates the character of the optimal protocols obtained with the optimisation algorithms. In particular, in relation to the critical velocity v_{crit} and the resonance timescale $T_{\text{res}} = \frac{2\pi}{\omega_{\text{res}}}$, which are present

²We recall that we have seen this infidelity measure before in Eq. (3.15) for the single qubit and Eq. (3.15) for the shuttling of Majoranas. For our current specific optimisation problem we have however slightly different initial and target states so we have defined it here again.

³It is possible to provide more feedback signals to the agent, such as observations retrieved from measurements [165, 166, 167] or for example a sum of instantaneous infidelity values at different times during the evolution $\mathcal{L} = \sum_i \mathcal{I}(t_i)$.




FIGURE 4.1: (a) For the Majorana control problem we give a ML agent (robot) the task to move the left Majorana γ_L from a position x_A to a position x_B in a total time τ . The aim of the agent is to find the right strategy to change the external potential $V(x)$ in time such that the system ends up in the final ground state $|\psi_B\rangle$. (b) During the learning (training) process the agent attempts different paths $x_L(t)$ for the Majorana. In this way the agent can use its past experiences (NES) together with the gradient signal (∂P) to minimise the ground state infidelity.

in this topological system, we find that it is useful to set and interpret the magnitude of the average velocity constraint v_{avg} with respect to these model parameters. Motivated by this we divide the Majorana control up into the four distinct Majorana motion regimes shown in Fig. 4.2. These regimes hold for both the proximity coupled superconductor as well as the Kitaev chain and are defined⁴ as follows:

- Regime I corresponds to the critical regime in which the Majorana must move on average above the critical velocity $v_{\text{avg}} > v_{\text{crit}}$. In this regime the ground-state fidelity is expected to rapidly decrease to zero.
- Regime II is the sub-critical regime for which the velocity is on average close to but nonetheless below v_{crit} . This regime is open ended in both the termination time τ and the movement length l . The key feature distinguishing this regime from regime IV below is the character of the found optimal protocols.
- Regime III is again a sub-critical regime, defined for times shorter than the resonance time T_{res} , but also with low velocity. This region cannot be used to efficiently move Majorana states over long distances, but we expect it to be relevant for braiding protocols based on small relative movements that change the effective couplings between Majoranas.
- Finally regime IV is the adiabatic regime in which we are above T_{res} and we have sufficient time to expect that the slow superadiabatic ramp-up/ramp-down protocols [144, 164] from earlier studies are optimal. Ideally, one would always like to be in this regime,

⁴Note that these definitions are taken from [73] coauthored by the author of this thesis.




FIGURE 4.2: The different Majorana motion control regimes labelled (I-IV). The red line corresponds to the critical velocity. Note that the blue line is not a hard boundary but determined by the specific protocols the algorithms come up with. This specific control 'phase' diagram is not limited to the Kitaev chain but can be generalised to include the proximity coupled superconductor and we speculate it is also relevant to braiding schemes where the coupling between different Majoranas is altered [160].

however, as discussed earlier, a gradual build up of noise and decoherence may make it necessary to get things done more quickly.

In the next section we aim to find the optimal controls $x_L(t)$ for the transport of a Majorana in each of these four regimes.

4.2 Numerical Optimisation Results

In order to find the optimal⁵ strategies we aim to minimise the infidelity \mathcal{I}_τ with some of the numerical optimisation methods introduced in Chapter 2 for each regime separately. Specifically we apply the NES and ∂P methods and benchmark them with SA and some parameterised bang-bang and superadiabatic protocols. In this section we will first focus on the results obtained with NES and ∂P in which the Kitaev chain model in Eq. (3.1) is used to compute the infidelity. For this, we find a novel counter-intuitive jump-like strategy in regimes I, II and III and recover the expected smooth superadiabatic protocols in regime IV (see Fig. 4.3). We then show that these ML protocols outperform our parameterised benchmark protocols (see Table 4.1). Lastly, we will briefly discuss that the obtained Majorana transport strategies

⁵Note that for the numerical optimisation methods for a potentially non-convex problem like ours it is hard to prove that we have obtained the "optimal" protocol. By optimal we therefore mean the protocol with the lowest cost value that we encounter and is obtained after the optimisation methods are converged.

also work for the proximity coupled superconductor by showing some optimisation results for this slightly more complicated model.

4.2.1 Machine Learned Majorana Transport Protocols

To study the distinct Majorana motion regimes we pick four representative points, $(l, \tau) = \{(4.32, 12), (4.95, 22), (0.48, 8), (2.4, 40)\}$, each of which belongs to a different regime. For the simulations that we perform l is chosen to be relatively small, only a few lattice sites. This is needed to avoid a large τ for some of the regimes, which would result in excessively long computation times.⁶ We choose a system size of $N = 110$, which is large enough to remove finite size effects and the Majoranas are localised with a small, but finite, energy splitting of $\delta E_0 \sim \mathcal{O}(10^{-8})$. To the best of our knowledge the NES and ∂P optimisation methods we use here have not been previously applied to (free) dynamical many-body systems of this size.

The next important practical aspect of the optimisation is the definition of the set of protocols (ansatz) over which to optimise, i.e. how to parameterise the control $x_L(t)$. We saw in section 2.2.3 that restricting to bang-bang protocols or a Fourier series (CRAB method) can dramatically shrink the size of the search space. However, since the Hamiltonian, through $V(x, t)$, is nonlinear in the control, we do not necessarily expect optimal bang-bang protocols. For this reason, we do not *a priori* fix the form of the control and revert to an unrestricted time-bin optimisation. This means that, like for the ∂P single-qubit control example, we discretise the control parameter into $[x_L(t_1), x_L(t_2), \dots, x_L(t_M)]$. Here, the individual time bins $x_L(t_i)$ with $t_i = i\Delta t$ become the in total $M = \tau/\Delta t$ new control parameters that can take continuous values. Importantly, we take the discretisation of the control Δt to be larger than the discretisation of the underlying simulation $\Delta t \geq 10\delta t$ to avoid potentially discontinuous cuts in the Majorana motion. Aside from the domain wall position it is also possible to parameterise by the velocity $v_L(t_i)$, from which the position can be obtained via $x_L(t) = \int_0^t v_L(t)dt$. Another option is to represent the control as the output of a neural network $x_L(t) = \text{NN}_\eta(t)$ where η are now the parameters of the neural network to optimise over.

After extensive tests and comparisons, see reference [73], we find a combination of these different parameterisations to work the best in practice. For the simulations presented in Fig. 4.3 and analysed below we use for ∂P a combination of a neural network⁷ together with a subsequent position optimisation. For NES we parameterise by position in regimes I and II

⁶The maximum simulation times we test are on the same order of magnitude as previous studies that looked at Majorana motion within in this system [83, 164]. It remains interesting however to investigate the optimal protocols for very large termination times and long transport distances. We will come back to this matter in the concluding remarks.

⁷The standard feedforward neural network that we use consists of three layers with rectified linear unit (ReLU) activation functions combined with one single sigmoid output neuron. The input to the network is the time t and the output the position $x_L(t)$ of the Majorana. This means we view the neural network as a special type of function approximator [168].

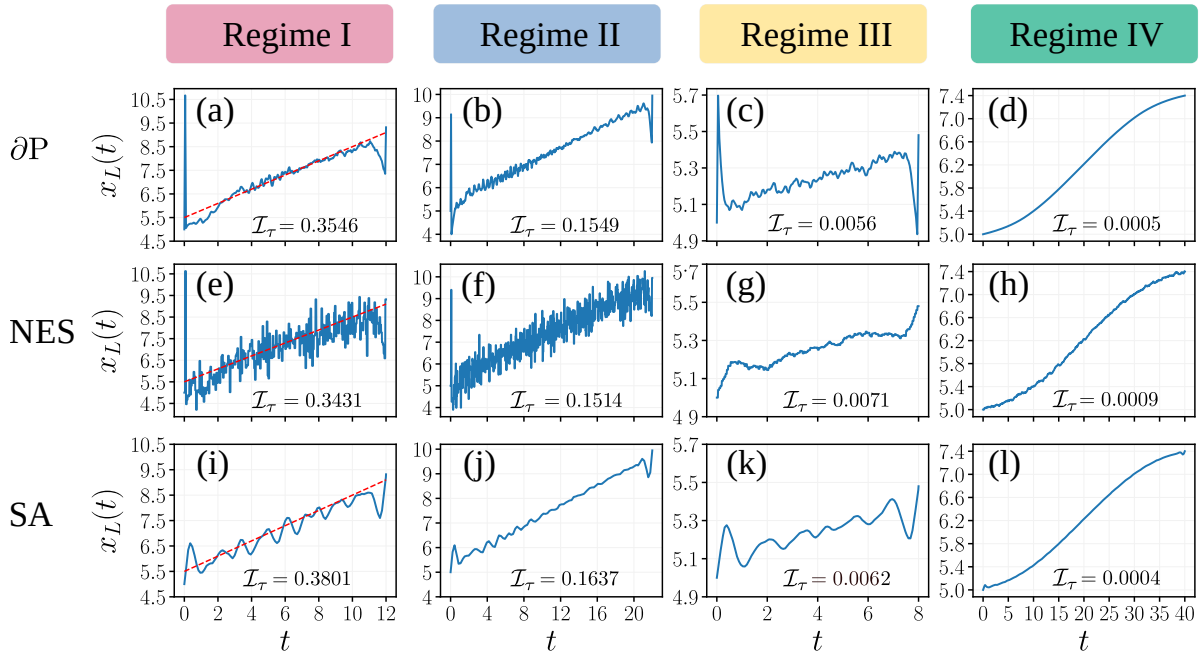


FIGURE 4.3: Control strategies $x_L(t)$ for Majorana transport obtained after minimisation of the infidelity \mathcal{I}_τ with Differentiable Programming (top row), Natural Evolution Strategies (middle row) and simulated annealing (bottom row). Four optimisations in the different control regimes I, II, III, and IV are performed (columns) for which we choose the external constraint parameters $(l, \tau) = \{(4.32, 12), (4.95, 22), (0.48, 8), 2.4, 40\}$. For regime I we have included the red dashed line to compare to linear motion at the critical velocity v_{crit} . In regime IV the algorithms reproduce the superadiabatic ramp up-down protocols whereas in the other regimes we find a new type of strategy which we have called jump-move-jump. In here the simulations for the infidelity are based on a Kitaev Chain with parameters $N = 110$, $\mu = 1$, $w = 1$, $\Delta = v_{\text{crit}} = 0.3$, $V_{\text{height}} = 30.1$, and $\sigma = 1$.

and velocity in regimes III and IV. We find that optimising with a neural network parameterisation gives suboptimal smooth protocols that can be further refined with either the position or velocity optimisations respectively. In addition, we find that optimising the velocity with NES seems to give smoother protocols compared to a direct position optimisation. This explains why the velocity optimisation is preferable in the low velocity regimes III and IV. Lastly, we note that for the neural-network parameterisations we perform updates with the optimisation algorithm Adam, while for all the other parameterisations we use vanilla gradient descent. We determine the hyperparameters for these optimisations empirically by scanning over a range of different values and we also restart the optimisations a few times from different initial conditions to avoid getting stuck in a local minimum.

In Fig. 4.3 we plot the resulting Majorana position protocols $x_L(t)$ obtained after the learning process has been completed. We first remark that both NES and ∂P converge to qualitatively similar-looking final protocols in each of the four regimes. In the slower regimes III and IV, ∂P seems to perform slightly better, in terms of final infidelity value. In contrast, NES performs better in the short-time regimes I and II. The differences are, however, small $\mathcal{O}(10^{-3})$ and could be partially attributed to the fact that different parameterisations for the

control are used. Furthermore, we note that the results for the ∂P optimisation are obtained after 440 update steps and for NES after in total 50 update steps (with 100 infidelity evaluations), which could also explain the small difference between the two methods.

We observe that the resulting protocols in regimes I, II, and III are similar and each consist of three main parts (features). Right at the start of the protocol the Majorana is kicked abruptly. This kick consists of an instantaneous forward jump followed by a rapid backward motion over a short time.⁸ Afterwards, in the middle part of the protocol, the Majorana moves on average with a nearly constant velocity motion that is dressed to various degrees with some rapid oscillations. At the end of the protocol a kick similar to the first one is applied, but now with first backward motion and then a large instantaneous forward step in position. We note that the constant-velocity middle parts are always at a velocity lower or right at the critical velocity, even in regime I.⁹ Another observation is that the controls move the domain wall position sometimes outside the target range, i.e. before the starting position $x_L(t) < x_A$ and after the final position $x_L(t) > x_B$. This seems counter-intuitive, but it is important to keep in mind that we are looking here at the driving protocol in $\mathcal{H}(t)$. The instantaneous ground state of this does not exactly match up with the time evolved state $|\psi(t)\rangle$ for quick non-adiabatic changes (the jumps and oscillations).

For the (super)adiabatic time scales, i.e. regime IV, we observe that the optimisation algorithms are able to recover the smooth superadiabatic ramp up-down protocols that were discussed earlier in literature [144, 169, 164]. The rationale behind this strategy is that the velocity of the Majorana is slowly building up to some finite velocity below v_{crit} and then decelerating back to zero again, in a symmetric fashion.¹⁰ During the entire protocol, when done perfectly, the ground state follows the moving-frame ground state and the motion is called superadiabatic, as we discussed in Chapter 3. However, a distinction with the previously studied protocols is that the ML protocols start with a small nonzero velocity $v_L(t=0)$, which does not lead to a huge increase in infidelity. This indicates that there is still some freedom within the family of superadiabatic protocols.

The infidelity values of these protocols are reported as labels in Fig. 4.3. It can be seen that, as expected, the performance is insufficient in the critical regime I. In this regime the strategies have infidelity values as bad as $\mathcal{I}_\tau \approx 0.35$. This confirms that trying to move the Majorana with a velocity which is on average greater than the critical velocity is a bad strategy. It could still be a relevant regime (and strategy) for near term Majorana demonstration experiments, where high-fidelity values are not yet required. In all the other regimes we find significantly lower infidelity values. In particular in regimes III and IV, where the obtained strategies result in infidelities values of about $\mathcal{I}_\tau \approx \mathcal{O}(10^{-3})$ and lower. For regime IV, however, this requires

⁸This can be best observed in panel 4.3 (c).

⁹We explain below in the jump-move-jump section why this is needed for this type of strategy to work.

¹⁰We will give an exact equation for this type of protocol below.

long transport times, which potentially leads to other problems.¹¹ Furthermore, in regime III the Majoranas are only moved over very short distances. Therefore, regime II would likely be the most relevant regime. In this regime our protocols give an infidelity of about $\mathcal{I}_\tau \approx 0.15$. However, to better interpret these results, and to argue that they are relevant for experiments, we need to make sure that there are no other control protocols, such as bang-bang protocols, with lower infidelities. We will look into this matter in the next subsection.

4.2.2 Simulated Annealing and Benchmark Protocols

In order to benchmark the performance of the ML protocols we compare them to parameterised bang-bang protocols, ramp up-down protocols, and protocols obtained with an unrestricted SA search. In Table 4.1 we provide an overview of the lowest numerical values of the infidelity obtained with these different methods. We will now discuss them one by one and compare the resulting protocols.

Regime	linear	bang-bang	ramp up-down	SA	DP	NES
I	0.4738	0.3923	0.4817	0.3801	0.3546	0.3431
II	0.2236	0.1713	0.2350	0.1637	0.1549	0.1514
III	0.0120	0.0067	0.0125	0.0062	0.0056	0.0071
IV	0.0077	0.0004	0.0045	0.0004	0.0005	0.0009

TABLE 4.1: Infidelity values of different types of parameterised protocols versus the machine learned protocols for optimal Majorana transport. We compare optimal bang-bang protocols, ramp up-down protocols and protocols obtained with an unrestricted SA search to the optimal ∂P and NES protocols. For completeness we have also included the values for a naive linear ramp $x_L(t) = v_{\text{avg}}t$. The simulations for these values were done with a Kitaev chain of system size $N = 110$ and the same parameters as in Fig. 4.3 above.

For the unrestricted¹² SA search we use a variant of the method described in section 2.2.3 that was also applied to optimise the Majorana edge setup in reference [83]. Here, we parameterise by the wall velocity $v_L(t)$ and iteratively update it by choosing two random intervals of time Δt of which one interval is increased by Δv and the other is decreased by Δv . The new infidelity \mathcal{I}_i is calculated for the updated velocity profile and the move is accepted with a probability $e^{-\delta\mathcal{I}/T_{SA}}$, where $\delta\mathcal{I} = \mathcal{I}_i - \mathcal{I}_{i-1}$ is the difference in the infidelity with respect to previous iteration. As before, the annealing temperature T_{SA} is slowly cooled down to zero. The results of this benchmark approach in each of the four regimes are shown in Fig. 4.3(i-l) and are qualitatively similar to the results obtained with NES and ∂P [Fig. 4.3(a-h)]. We note that for these SA optimisations we need about $\mathcal{O}(10^5)$ infidelity evaluations before convergence is reached. Thus, in this specific example the SA approach is far more computationally demanding than the NES and ∂P methods.

¹¹Such as disorder, see the discussion in section 4.4.

¹²By unrestricted we mean that we put no constraints on the family of protocols over which we optimise.

Next, we benchmark against the 'optimal' bang-bang protocols obtained by performing an SA search with bang-bang ansatz, like we did for the single qubit control in Chapter 2. This time we consider bang-bang protocols in the velocity of the Majorana $v_L(t_i)$, where at each time step t_i the velocity is either $v_L(t_i) = v_{\max}$ or $v_L(t_i) = v_{\min}$. Importantly, the width of the bangs $t_{i+1} - t_i$ is at least $10\delta t$, where δt is the discretisation of the Trotter evolution of the numerical integration method. We fix $v_{\min} = 0$, like in reference [83], and optimise for several different $v_{\max} \leq 8v_{\text{avg}}$. We note that this sometimes means that $v_{\max} > v_{\text{crit}}$, which is different from [83]. The high velocity jumps in the ML protocols inspired us to follow such strategy.

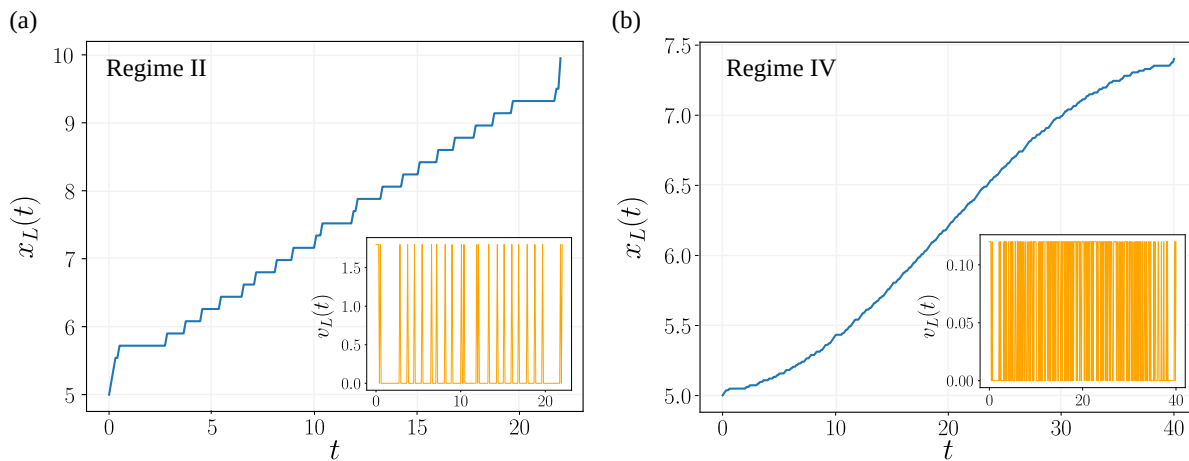


FIGURE 4.4: Control protocols for the domain wall position $x_L(t)$ (main panels) and corresponding velocities (insets) obtained with a simulated annealing search that was restricted to bang-bang protocols. We plot both the results for the high velocity regime I (panel a) and slow adiabatic regime IV (panel b). We observe that in both cases the optimal bang-bang protocols start to approximate the protocols found with the ML methods (c.f. Fig. 4.3). For these simulations we have the same Kitaev chain parameters as in Fig. 4.3.

The resulting infidelity values of these optimisations are reported in Table 4.1 and the protocols for regimes II and IV given in Fig. 4.4. We see that for large enough maximum velocity v_{\max} and small enough bang widths $t_{i+1} - t_i$ the bang-bang protocols start to approximate the jump-like ML protocols in regimes I, II, and III, and the superadiabatic protocols in regime IV. The infidelity performance of these protocols is slightly worse compared to the NES and ∂P protocols in the non-adiabatic regimes. Therefore, we conclude that the bang-bang protocols, as far as they can be still called bang-bang, only approximate the ML and superadiabatic protocols and are thus not optimal in this setting.¹³ However, it remains interesting to investigate if the optimal protocols become of bang-bang form when the slope of the domain wall, σ in $V(x)$, is increased, such that the control becomes more linear, or when we look at a disordered system.

¹³An important remark here is that one needs to take into account that the bang-bang protocols are restricted by a maximum and minimum velocity, whereas for the ∂P and NES unrestricted methods we did not need to impose this limitation.

The final protocol we compare our ML results to is part of the family of superadiabatic protocols studied in reference [144, 164]. This is based on the idea of slowly ramping up and down the velocity of the Majorana, while following the moving-frame ground state. This can be expressed mathematically by the following domain wall velocity

$$v_L(t) = \begin{cases} v_{\max} \frac{1 - \cos \omega t}{2}, & 0 \leq t \leq \frac{\pi}{\omega} \\ v_{\max}, & \frac{\pi}{\omega} \leq t \leq \tau - \frac{\pi}{\omega} \\ v_{\max} \frac{1 - \cos(\omega t - \omega \tau)}{2}, & \tau - \frac{\pi}{\omega} \leq t \leq \tau \\ 0, & \text{otherwise} \end{cases} \quad (4.2)$$

in which ω is the parameter that determines how quickly the velocity is accelerated to v_{\max} . For true superadiabaticity one needs to choose $\omega \ll \omega_{\text{res}}$. In order to determine the best ramp up-down benchmark protocol, in each of the four regimes separately, we scan over a range of v_{\max} and ω , and pick the protocol with the lowest infidelity value. In regimes I, II and III this leads to a very high ω with $v_{\max} \approx v_{\text{avg}}$, while in regime IV, by definition, there is enough time for this protocol to work. This can be seen in Table 4.1 where only in regime IV it performs well. However even in this regime it performs slightly worse than the ML protocols. This is likely due to the fact that NES and ∂P have more freedom to make small refined changes compared to the parameterised ramp up-down protocols. For example we saw that the NES and ∂P protocols in regime IV start with a small, but nonzero, velocity $v_L(0)$ while the parameterised protocols always need to start at zero velocity.

4.2.3 Extension to the Proximity Coupled Superconductor

To end this numerical results section, we show here that for the proximity-coupled superconductor model of Eq. (3.10) we obtain qualitatively similar-looking protocols in all four regimes. For this we apply the same ∂P and NES algorithms and use the same parameterisations for the control of the domain wall as we did for the optimisations of the Kitaev chain. In practice, this means we only change the Hamiltonian¹⁴ in the underlying numerical simulations for the time evolution and keep all other the components of the numerical optimisation algorithms the same. However, as a consequence of the extra spin degree of freedom, the proximity-coupled superconductor Hamiltonian [Eq. (3.10)], for a fixed lattice size N , is twice as big as the Kitaev chain Hamiltonian. This means that the simulations become computationally more expensive and the optimisation more challenging. Nevertheless, we find that ∂P and NES still converge relatively quickly ($\mathcal{O}(10^2)$ update steps) to the optimal protocols.

The results for these optimisations are shown in Fig. 4.5 for two different sets of parameters of the proximity-coupled semiconducting nanowire model. One set of parameters (in panels a-h) is within the Kitaev chain limit (high B -field). The other set of parameters is

¹⁴Note that the control is still encoded in the external potential profile $V(x, t)$ in Eq. (3.14).

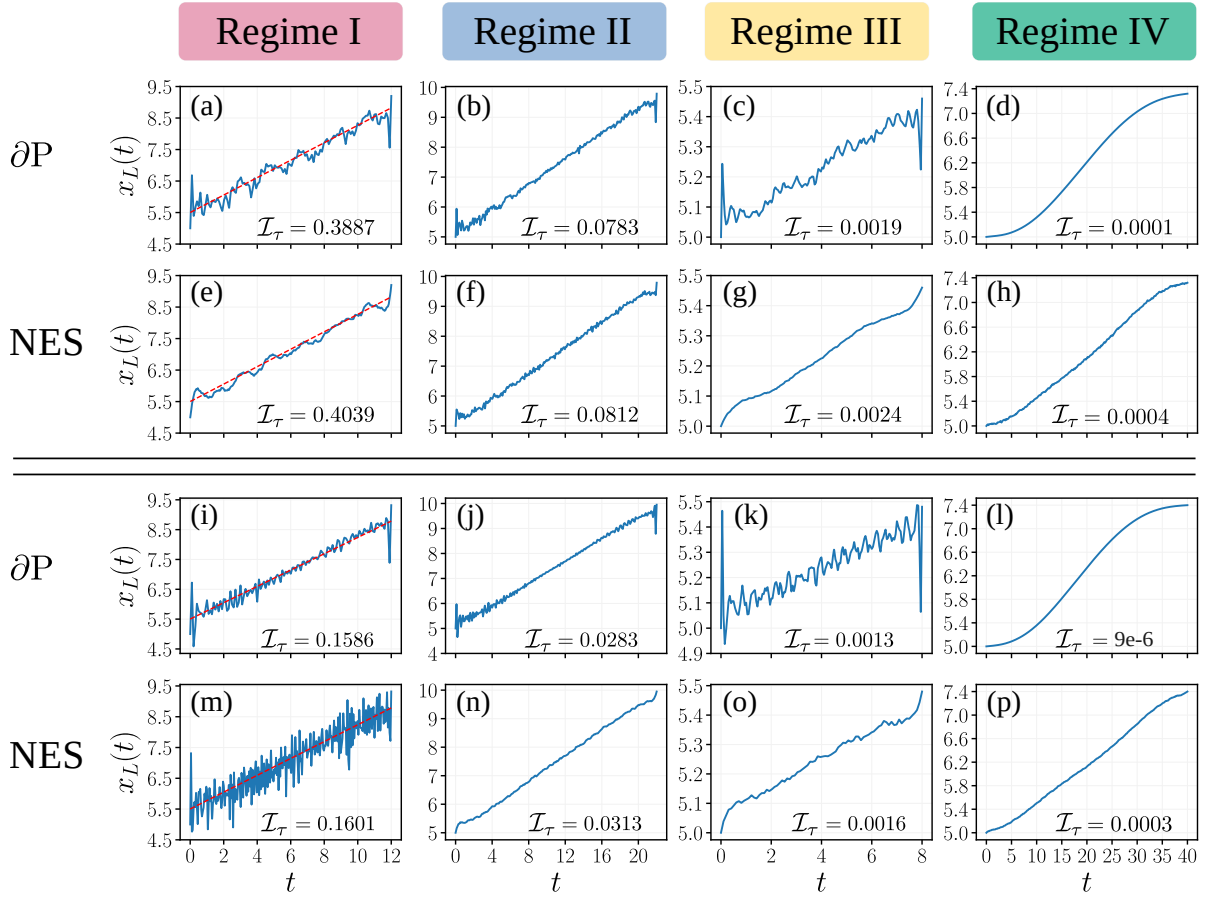


FIGURE 4.5: Control strategies $x_L(t)$ for the domain wall of the Majorana transport in the proximity coupled semiconducting nanowire obtained after minimisation of the infidelity \mathcal{I}_τ with Differentiable Programming (top rows, panels a-d, i-l) and Natural Evolution Strategies (bottom rows, panels e-h, m-p) in the four different regimes. The optimisations were performed for two sets of parameters. The top two rows (panels a-h) correspond to simulations with the parameters inside the Kitaev chain limit $N = 100$, $\mu = 0.0$, $w = 1$, $\Delta = 0.8$, $B = 2$, $\alpha = 0.8$, $V_{\text{height}} = 30.1$, and $\sigma = 1$. For the bottom two rows (panels i-p) we have the parameters $N = 100$, $\mu = -0.55$, $w = 1$, $\Delta = 0.5$, $B = 1$, $\alpha = 1.0$, $V_{\text{height}} = 30.1$, and $\sigma = 1$, which are outside the Kitaev chain limit. For both of these the external constraint parameters (l, τ) for the different regimes in chronological order are given by $\{(4.32, 12), (4.95, 22), (0.48, 8), (2.4, 40)\}$. The resulting protocols are qualitatively similar to the protocols obtained for the Kitaev chain in Fig. 4.3.

outside the Kitaev limit, but in a phase where there are still localised Majoranas. We choose the parameters in both cases such that v_{crit} is approximately the same as for the simulations of the Kitaev chain. This means that a direct comparison with the previous results for the Kitaev chain in Fig. 4.3 is possible for the same external constraint parameters (l, τ) . We observe in all panels that the optimal protocols for the movement of the Majoranas $x_L(t)$ are qualitatively similar-looking to the protocols obtained for the Kitaev chain. That is, in regimes I, II and III we obtain a protocol comprising two pulses at the boundaries that consist of forward jumps and backward motions over short periods of time. In between the jumps this protocol has a period of on average constant motion with a velocity below v_{crit} . In regime IV, we again obtain the smooth adiabatic ramp up-down protocols with a low ground state infidelity.

One difference is that in this case the boundary jumps in the protocols in the non-adiabatic regimes are of smaller magnitude compared to the protocols in Fig. 4.3. Furthermore, for the parameters outside the Kitaev chain limit the initial and final pulses consist of three parts; a forward jump followed by a backward jump and again a forward jump. The performance of all these protocols as quantified by the infidelity is similar however, and even slightly better in regime II, compared to the Kitaev chain. This means that despite those small differences in protocols the performance is still robust and our results from the Kitaev chain generalise to the more complicated proximity coupled nanowire model.

4.3 The Jump-move-Jump Control Strategy

Thus far, we have obtained control strategies for Majorana zero modes with ∂P and NES and shown that they outperform the previously studied bang-bang and superadiabatic protocols. In particular, in the short time regimes I, II, and III, we found an intriguing jump-like strategy with low infidelities for relatively long movement distances. In this section we aim to uncover the underlying (physical) working mechanisms for the success of this strategy.

To this end, and to make (semi-)analytical analysis possible, we simplify the ML strategies and focus on a simple model strategy. The simple model that we consider, which we call the jump-move-jump (JMJ) strategy, is shown in Fig. 4.6 and encapsulates the main features of the ML protocols. At the start of this strategy, there is a pulse¹⁵ in which the Majorana is kicked instantaneously forward over a distance $\Delta x_{\text{forward}}$ and then pulled back over a distance Δx_{back} in a time Δt_{back} . This initial pulse is followed by a motion at a constant velocity v . At the end of the protocol, a pulse similar to the first one is performed but in the reverse order; a movement backward over a distance Δx_{back} followed by a final instantaneous kick in the position of size $\Delta x_{\text{forward}}$. Note that, while we restrict here to exactly the same initial and final jumps, it is straightforward to generalise to unequal jump sizes.

In the following analysis¹⁶ we will show that we can derive an approximate expression for the infidelity \mathcal{I}_τ of this protocol based on some of the system parameters. We will evaluate the JMJ strategy in the moving frame and treat the instantaneous jumps by invoking the sudden approximation. We will first stick to protocols in which $\Delta x_{\text{back}} = 0$ and refer to this as the bare JMJ strategy. The scenario when $\Delta x_{\text{back}} \neq 0$ we refer to as the dressed JMJ strategy and will be treated numerically.

¹⁵Note that by *jumps* in jump-move-jump we refer to these pulses. When backward motion is included these can be seen as dressed jumps.

¹⁶This follows the derivation presented in [73] coauthored by the main author of this thesis.

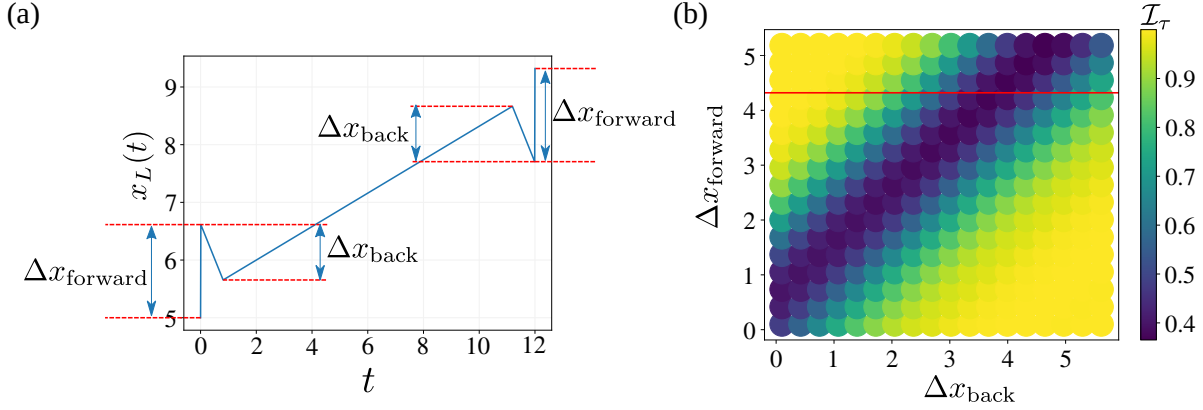


FIGURE 4.6: (a) The JMJ strategy $x_L(t)$ inspired by the protocols obtained with ∂P and NES in regimes I, II, and III. This model strategy consists of three core parts; two dressed jumps consisting of forward and backward motions at the start and end of the protocol and a period of nearly constant motion between the jumps. (b) Infidelity surface of the JMJ strategy as a function of the forward $\Delta x_{\text{forward}}$ and backward Δx_{back} jump sizes in regime I. The red line indicates where $\Delta x_{\text{forward}}$ is equal to the total movement length $l = 4.32$. Larger jump sizes thus end up outside the target range $x_L(t) > x_B$. We set $\Delta t_{\text{back}} = 0.07$ and the parameters for the Kitaev chain simulations of the infidelity are the same as in Fig. 4.3.

4.3.1 An Expression for the Infidelity of the Bare JMJ Control Strategy

When $\Delta x_{\text{back}} = 0$ the sizes of the initial and final forward jumps of the JMJ strategy can be expressed as $\Delta x_{\text{forward}} = (l - v\tau)/2 \equiv \delta$ with v the constant velocity of the middle part of the protocol. This means that we can just focus on v and the other parameters are fixed by the external constraints (l, τ) . To evaluate this strategy in the moving frame we recall from section 3.3.1 that the entire system (Majorana wire) needs to move at the same velocity. For this reason we will focus on the case in which both the left and right Majorana are moved with the JMJ strategy. That is, the right domain wall in $V(x, t)$, Eq. (3.14), becomes time dependent as well via $x_R(t) := x_L(t) + C_x$, where C_x is the fixed distance between the two Majoranas. During the middle part of the protocol the whole system thus moves at a velocity v .

The effect of the initial forward boundary jump $x_L \mapsto x_L = x_A + \delta$ is evaluated by expanding the infidelity in terms of the instantaneous eigenbasis $|\psi_{A+\delta}^i\rangle$ of the Hamiltonian right after the first jump, which gives

$$\mathcal{I}_\tau = 1 - \left| \sum_i \langle \psi_B | U_\tau | \psi_{A+\delta}^i \rangle \langle \psi_{A+\delta}^i | \psi_A \rangle \right|^2, \quad (4.3)$$

$$= 1 - \left| \langle \psi_B | U_\tau | \psi_{A+\delta}^0 \rangle \langle \psi_{A+\delta}^0 | \psi_A \rangle + \sum_{i>0} \langle \psi_B | U_\tau | \psi_{A+\delta}^i \rangle \langle \psi_{A+\delta}^i | \psi_A \rangle \right|^2. \quad (4.4)$$

Here $U_\tau = \mathcal{T} e^{-i \int_0^\tau \mathcal{H}(t) dt}$ is the time-evolution operator that follows the time dependence of the control strategy. Note that $\mathcal{H}(t)$ can correspond to the Kitaev chain Hamiltonian in Eq. (3.1)

or the proximity coupled nanowire model in Eq. (3.10) and that our analysis does not depend on it. We assume that after this initial kick we are still close to the ground state $|\psi_{A+\delta}^0\rangle$ and in the following ignore¹⁷ the comparatively small contributions coming from the excited states ($i > 0$) in the sum in Eq. (4.4).

During the move part of the protocol the system is static in the moving frame and we insert projections onto the ground state $|\psi_v^0\rangle$ of the moving frame Hamiltonian H_v in Eq. (3.16). This results in the approximate¹⁸ infidelity

$$\mathcal{I}_\tau \approx 1 - \left| \langle \psi_B | \psi_v^0 \rangle \langle \psi_v^0 | U_\tau | \psi_v^0 \rangle \langle \psi_v^0 | \psi_{A+\delta}^0 \rangle \langle \psi_{A+\delta}^0 | \psi_A \rangle \right|^2. \quad (4.5)$$

Importantly, the effect of the unitary operator U_τ in the moving frame is trivial (up to some phase) and we have $|\langle \psi_v^0 | U_\tau | \psi_v^0 \rangle|^2 = 1$. Then, we account for the final instantaneous jump to x_B by projecting onto the instantaneous ground state $|\psi_{B-\delta}^0\rangle$ of \mathcal{H} with the domain wall at position $x_L = x_B - \delta$, which gives

$$\mathcal{I}_\tau \approx 1 - \left| \langle \psi_B | \psi_{B-\delta}^0 \rangle \langle \psi_{B-\delta}^0 | \psi_v^0 \rangle \langle \psi_v^0 | \psi_{A+\delta}^0 \rangle \langle \psi_{A+\delta}^0 | \psi_A \rangle \right|^2. \quad (4.6)$$

With these projections we have reduced the expression of the target state infidelity for the bare JMJ strategy to a product of four different overlaps. The amplitudes $\langle \psi_{A(B)} | \psi_{A(B)\pm\delta}^0 \rangle$ come from the boundary jumps in position and can be computed numerically for our model \mathcal{H} . We find that they can be characterised by a Gaussian¹⁹

$$O_\delta \equiv \left| \langle \psi_{A(B)} | \psi_{A(B)\pm\delta}^0 \rangle \right|^2 \sim e^{-\left(\frac{\delta}{\epsilon + \alpha\lambda_F}\right)^2}, \quad (4.7)$$

in which $\lambda_F = 2\pi/k_F$ is the Fermi wavelength and (α, ϵ) are two fitting parameters that depend on the model. For the Kitaev chain with $\Delta = 0.3$ with two-wall movements (α, ϵ) are $(-0.12, 0.3)$ and for one-wall movement $(-0.33, 0.44)$.

The other two overlaps, $\langle \psi_{A(B)\pm\delta}^0 | \psi_v^0 \rangle$, correspond to boosts from the static-frame ground states to the moving frame. To estimate these contributions we use the results from reference [164], which showed that there is a (relativistic) kinetic energy correction $E \sim \gamma M v^2$ to the moving-frame ground state with effective mass $M \propto k_F/\Delta$ and $\gamma = \sqrt{1 - v^2/v_{\text{crit}}^2}$. From this we approximate the ground state overlaps as

$$O_v = \left| \langle \psi_{A(B)\pm\delta}^0 | \psi_v^0 \rangle \right|^2 \sim 1 - \beta \left(\frac{v}{\Delta}\right)^2 \gamma, \quad (4.8)$$

in which the parameter β needs to be determined from a to our model. For the Kitaev chain

¹⁷This assumption possibly breaks down when δ becomes very large.

¹⁸Approximate because we do not include the excited state contributions in Eq. (4.4). The projection into the moving frame is exact.

¹⁹Note that in this overlap the specific starting and target positions x_A and x_B do not matter.

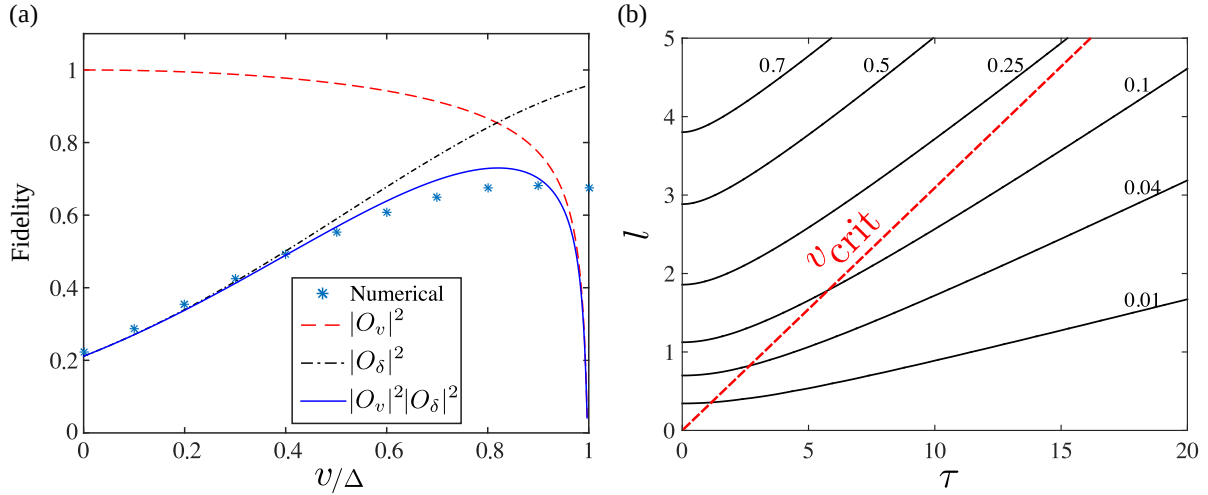


FIGURE 4.7: (a) Fidelity \mathcal{F}_τ and overlaps of the various contributions to the expression Eq. 4.9 for the JMJ strategy compared to a numerically obtained fidelity from simulations of the Kitaev chain. In here we set $(l, \tau) = (4.32, 12)$ (regime I) and move only the left domain wall (i.e. $\beta = 0.065$). The other parameters are the same as in Fig. 4.3. Our approximate analysis (solid blue line) captures the behaviour of the JMJ strategy (blue stars) well for low velocities v but breaks down near $v_{\text{crit}} = \Delta$. (b) Infidelity surface obtained after maximisation of Eq. 4.9 for a range of parameters (l, τ) . The red dashed line shows where $v_{\text{avg}} = v_{\text{crit}}$ and can be used to distinguish the critical regime I from the other control regimes. For these results we have used the same Kitaev chain parameters as in panel (a).

we find $\beta \approx 0.13$ for motion with two walls and $\beta \approx 0.065$ for motion with one wall. It is straightforward to generalise this equation and the fitting parameters to the proximity-coupled semiconducting nanowire model.

By combining these contributions we thus arrive at the following expression for the infidelity

$$\mathcal{I}_\tau = 1 - |O_\delta O_v|^2 = 1 - \exp\left\{\left(-\frac{(l - v\tau)^2}{2(\epsilon + \alpha\lambda_F)^2}\right)\right\} \times \left[1 - \beta \left(\frac{v}{\Delta}\right)^2 \gamma\right]^2, \quad (4.9)$$

in which we have substituted $\delta = (l - v\tau)/2$ in O_δ . This equation can be used to make estimates for the value of the infidelity of the JMJ strategy based on the fitting parameters and the external constraint parameters. In Fig. 4.7 (a) we show the estimated values for the fidelity $\mathcal{F}_\tau = 1 - \mathcal{I}_\tau$ in regime I and compare it to numerical simulation data of the Kitaev chain. We observe that up to velocities that are very near $v_{\text{crit}} = \Delta$ the predicted values match closely with the simulation data. The break down near v_{crit} is expected as we have put a harsh penalty on moving at the critical velocity by our approximation for O_v in Eq. (4.8).

In the same figure we also show the individual contributions O_δ and O_v to the infidelity. We observe that at the specific velocity $v \approx 0.85v_{\text{crit}}$ the product of these two contributions and hence \mathcal{F}_τ is maximised. This means that to find the optimal bare JMJ strategy one carefully needs to choose v such that the product of the two contributions is maximised. This can be interpreted as a trade-off between the amount of fidelity lost by the jumps ($1 - O_\delta$) and the

fidelity lost due the constant velocity motion ($1-O_v$). The maximisation can be done efficiently for a wide range of parameters (l, τ) for which we show results in Fig. 4.7 (b). Here we see that the bare JMJ strategy in the critical regime I (above the red dashed line) gives infidelity values $\mathcal{I}_\tau > 0.25$ while in regimes II and III (below the red dashed line) it performs well with infidelities of $\mathcal{O}(10^{-2})$ and smaller.

4.3.2 The Dressed JMJ Control Strategy

In addition to the instantaneous forward jumps, a key feature of all the ML protocols in regimes I, II and III are rapid backward motions for short periods of time. This motivates us to dress the bare JMJ strategy with additional backward motions over distances $\Delta x_{\text{back}} \neq 0$ as shown in Fig. 4.6 (a). In this section we investigate the effect of these dressings and show numerically that they allow to better target the moving frame ground state during the *move* part of the (dressed) JMJ control strategy.

In Fig. 4.6 (b) we show the target state infidelity \mathcal{I}_τ of the dressed JMJ protocol in regime I as a function of $\Delta x_{\text{forward}}$ and Δx_{back} . We observe a clear diagonal structure in the surface plot, which shows that there is an optimal combined jump size $\Delta x_{\text{opt}} = \Delta x_{\text{forward}} - \Delta x_{\text{back}}$ where the infidelity of the JMJ protocols is minimised. Once Δx_{opt} is fixed the magnitudes of $\Delta x_{\text{forward}}$ and Δx_{back} themselves only slightly change the infidelity value, i.e. Δx_{opt} is within the valley of the infidelity surface $\mathcal{I}_\tau(\Delta x_{\text{forward}}, \Delta x_{\text{back}})$. This indicates that the rationale behind the dressed JMJ strategy is the same as for the bare JMJ strategy and the purpose of the dressed jumps is to bring²⁰ the system in a state moving at a particular constant velocity v . Note that when $\Delta x_{\text{back}} = 0$ the optimal 'combined' jump size Δx_{opt} reduces to the optimal forward jump, $\Delta x_{\text{forward}}$, that minimises the infidelity of the bare JMJ strategy in Eq. (4.9).

In order to understand the benefits of the backward motions for targeting a moving-frame ground state we compute the instantaneous infidelity

$$I(t) = 1 - |\langle \psi(t) | \psi_{\text{ins}}^0(t) \rangle|^2, \quad (4.10)$$

with $|\psi(t)\rangle = \mathcal{T} e^{-i \int_0^t H(t) dt} |\psi_A\rangle$ for various dressed JMJ protocols in the lab and moving frames of the system. For the lab frame we use the standard Kitaev chain Hamiltonian $H(t) = \mathcal{H}(t)$ (Eq. 3.1) with $|\psi_{\text{ins}}^0(t)\rangle$ its instantaneous ground state at time t . The lab frame results are shown in Fig. 4.8 (a). In here we see that while the bare JMJ strategy (in red) initially has a lower $I(t)$ over time the dressed JMJ protocols start to perform better. At the final time $t = \tau$ the dressed protocols have a lower target state infidelity $\mathcal{I}_\tau = I(\tau)$.

For the moving-frame analysis we use the moving-frame Hamiltonian $H(t) = H_v(t)$, Eq. (3.16), with v equal to the velocity of the middle part of the protocols. Importantly, to make a fair comparison we scan over different dressed JMJ protocols that all have the same

²⁰Or get the system out in case of the final dressed jump.

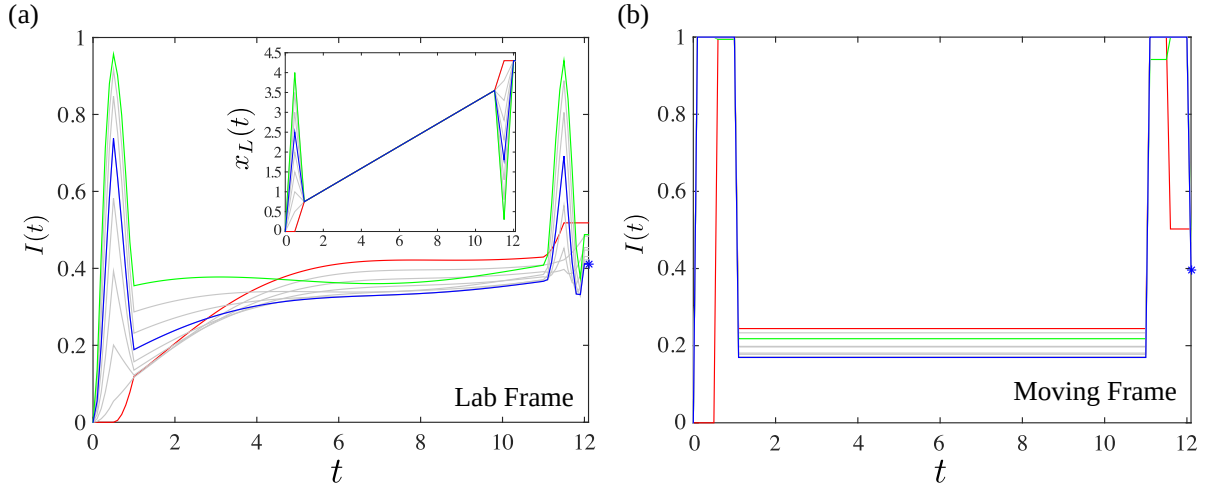


FIGURE 4.8: Instantaneous infidelity $I(t)$ of the dressed jump-move-jump strategy in (a) the lab frame and (b) the moving frame at a velocity v . The inset shows the control protocols $x_L(t)$ that were used for the simulations in the main panels. The best protocol measured by the target state infidelity \mathcal{I}_τ is given in blue and the bare JMJ protocol in red. The green curves correspond to the protocols with the largest backward motion Δx_{back} . The main observation is that the dressed protocols are able to better target the moving frame ground state.

v , see the inset in Fig. 4.8 (a). Here $|\psi_{\text{ins}}^0(t)\rangle$ is the instantaneous ground state of $H_v(t)$. The results for this exercise are plotted in Fig. 4.8 (b), where we see that all the protocols dressed with backward motions have a lower instantaneous infidelity $I(t)$ than the bare protocol. This is interesting and shows that the dressings $\Delta x_{\text{back}} \neq 0$ can be used to better target the moving-frame ground state.

Then, to find the optimal dressed JMJ protocols that minimise the target state infidelity in the different Majorana motion regimes, we scan over the parameters $\{\Delta x_{\text{forward}}, \Delta x_{\text{back}}, \Delta t_{\text{back}}\}$ and select the values corresponding the lowest infidelity value that we encounter. In Fig. 4.9 we show the JMJ protocols obtained in this way in regimes I, II and III. In regimes I and III we observe a large dressing effect with jumps that go outside the target ranges, while in regime II this effect is less pronounced. The infidelity values of these protocols are similar to the values of the optimal ML protocols, with only small differences of $\mathcal{O}(10^{-2})$. This shows that the optimal parameterised dressed JMJ protocols that we analysed are competitive with the protocols obtained with our numerical optimisation methods.

4.4 The Effect of Disorder

In section 3.2.3 we discussed that the real-world implementation of Majorana based topological quantum devices is hampered by the presence of disorder and other imperfections in the experimental setup. For this reason, in this section, we will look at the robustness of our optimal transport protocols with respect to disorder in the Majorana nanowires. We will show that the jumps of the JMJ strategy are robust, whereas during the move part of the protocol the




FIGURE 4.9: The optimal dressed JMJ strategies $x_L(t)$ in regimes I, II and III. The parameters $(\Delta x_{\text{forward}}, \Delta x_{\text{back}}, \Delta t_{\text{back}})$ for these protocols are $\{(7.992, 7.506, 0.05), (1.909, 1.550, 0.31), (0.548, 0.457, 0.33)\}$ and the parameters of the Kitaev chain simulations to compute \mathcal{I}_τ are the same as in Fig. 4.3. We see that the protocols in regimes I and III have a large forward jump that goes outside the target range $x_L(t) > x_B$ while the jumps in regime II stay within the target range. The performance \mathcal{I}_τ of these protocols is competitive with the ML protocols in Fig. 4.3.

infidelity gradually increases over time. The optimal superadiabatic protocols in regime IV, in contrast, are not robust and, for sufficiently strong disorder strength, start to perform worse than the naive linear benchmark protocol²¹.

The effect of disorder in the wires [170, 171, 172] is modelled by consider the Kitaev chain Hamiltonian $\mathcal{H}(t)$ with a disordered chemical potential

$$\mu_{\text{dis}}(x) = \mu(x) + \tilde{\mu}(x). \quad (4.11)$$

Here, we have added the Gaussian noise term $\tilde{\mu}(x)$ with mean 0 and standard deviation λ to the clean $\mu(x)$. As a consequence of this noise, the effective coherence length ξ_{eff} of the Majoranas increases and is given by

$$\frac{1}{\xi_{\text{eff}}} = \frac{1}{\xi} - \frac{1}{2l_{\text{dis}}} \quad (4.12)$$

in the continuum limit of the model. The length $l_{\text{dis}} = v_F^2/\lambda^2$ is the characteristic disorder length scale. This means that in the presence of disorder the Majoranas effectively delocalise. When the disorder strength is such that $2l_{\text{dis}} = \xi$ the effective Majorana coherence length diverges $\xi_{\text{dis}} \mapsto \infty$ and the topological phase is destroyed.

To assess the robustness of our optimal protocols obtained in the clean system we simulate them in this disordered Kitaev chain and compute the disorder averaged infidelity measures $\langle I(t) \rangle$ and $\langle \mathcal{I}_\tau \rangle$. In Fig. 4.10 (a) we show the instantaneous infidelity $\langle I(t) \rangle$ for various disorder strengths for the optimal JMJ protocol in regime I. We observe that the value of $\langle I(t) \rangle$ right after the first jump does not change (visibly) with increasing disorder strength. This is

²¹Note that here again we follow the results and discussion published in [73] coauthored by the main author of this thesis.

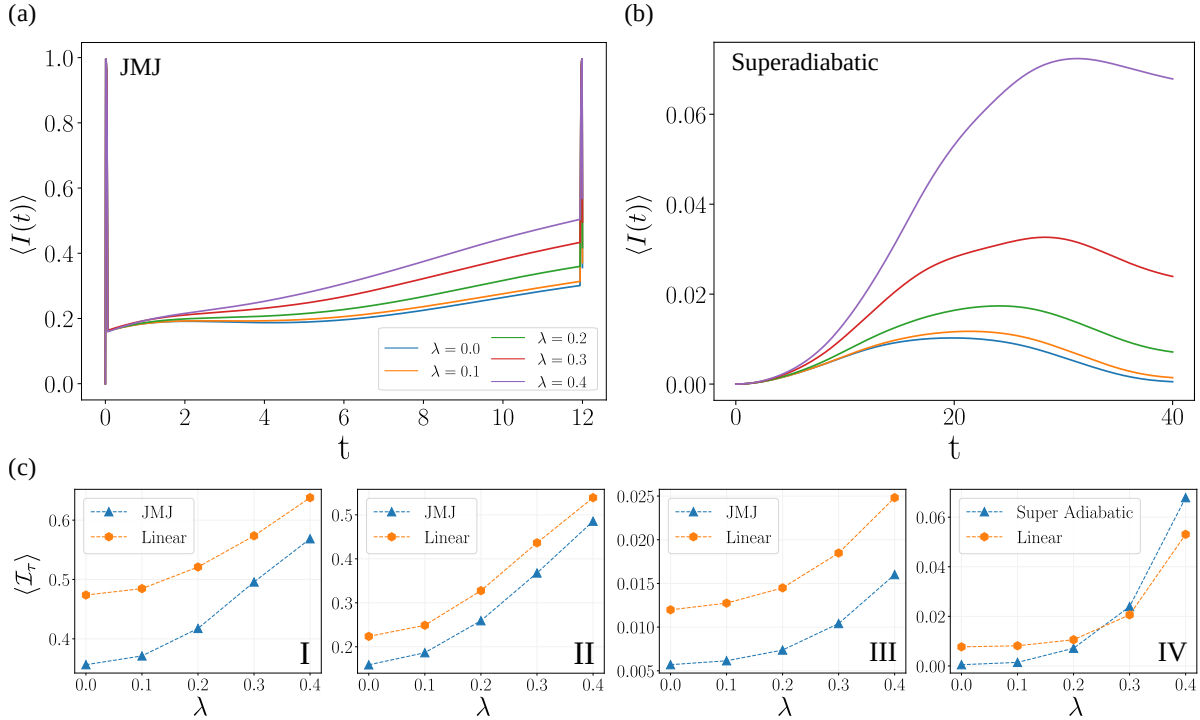


FIGURE 4.10: Robustness of the optimal JMJ and superadiabatic protocols with respect to disorder in the chemical potential $\tilde{\mu}(x)$. In panels (a) and (b) we plot the disorder averaged instantaneous infidelity $\langle I(t) \rangle$ of the protocols in regimes I and IV respectively. We observe that the jumps of the JMJ strategy are robust whereas during linear motion and also superadiabatic accelerating motion $\langle I(t) \rangle$ increases with disorder strength λ . In (c) we show the disorder averaged target state infidelity $\langle \mathcal{I}_\tau \rangle$ versus disorder strength in all four regimes. In here we benchmark against a linear protocol (orange) and see that while the optimal JMJ protocols outperform the linear protocol at all λ , the superadiabatic strategy becomes worse than linear motion near $\lambda \approx 0.2$. For all these simulations we used 500 noise realisations and the parameters of the Kitaev chain were the same as in Fig. 4.3.

also the case for the final jump, although it is a bit harder to observe. During the move part of the protocol, $\langle I(t) \rangle$ increases slightly with disorder strength. This increase could be explained by the lower critical velocity [173] in the disordered system resulting from the fact that we are pushing the Majorana through a rough potential landscape. Another possible reason could be that the longer ξ_{eff} implies that the Majoranas are less robust to moving at a specific constant velocity v .

As a consequence of this increase, the final averaged infidelity $\langle \mathcal{I}_\tau \rangle$ increases approximately quadratically²² with the disorder strength as shown in Fig. 4.10 (c). For low disorder strengths $\lambda \leq 0.1$ the increase is still small with a difference of about $\mathcal{O}(10^{-2})$. In contrast, for strengths $\lambda > 0.2$ the infidelity shoots up significantly to values of $\langle \mathcal{I}_\tau \rangle > 0.4$. Despite this the optimal JMJ protocol still outperforms a naive linear benchmark protocol at all disorder strengths. This indicates that lowering the velocity $v < v_{\text{avg}}$ of the middle part of the protocol by dressed initial and final jumps is still beneficial compared to just constant motion at v_{avg} .

²²Note that this is observed for all regimes but we are still focusing our discussion in this paragraph on regime I.

For the other regimes we find a similar quadratically increasing behaviour of $\langle \mathcal{I}_\tau \rangle$ as a function of the disorder strength. In regimes II and III, the optimal MJJ protocols seem to be still robust since they outperform the linear protocol at all λ . In regime IV, where in the clean case we have an optimal superadiabatic protocol, we find that this quadratic increase, at an disorder strength of about $\lambda \approx 0.25$, leads to an infidelity value that is worse than the value of the linear benchmark protocol. From the $\langle I(t) \rangle$ plot in Fig. 4.10 (b) we see that this break down of the performance of the protocol in regime IV comes from a gradual build²³ up of fidelity loss over time that does not find its way back to the target ground state. This possibly means that the superadiabatic ramp up-down protocols are not able to slowly bring the disordered system in a moving frame. It could be that the slow ramping up of the velocity needs to be done over a much longer time frame. Another possible explanation is that the maximum velocity of the protocol is higher than the new effective (disordered) v_{crit} .

Although these results look promising for the MJJ control strategy²⁴, we cannot claim that it is also the optimal strategy in the disordered system. Moreover, it is still open question if there is a better strategy than linear motion in regime IV. To answer these questions, we need to perform numerical optimisations with the ML or SA methods directly in the presence of disorder. This can be achieved by using the disorder averaged target state infidelity $\langle \mathcal{I}_\tau \rangle$ as the cost functional. However, we remark that simulating the averaged infidelity is computationally expensive, which means that such an optimisation becomes expensive and challenging.

4.5 Summary

In this chapter we looked at the optimisation of the transport of Majorana zero modes in superconducting nanowires. We formulated the transport problem in the form of a mathematical optimisation problem that can be tackled by the ML optimisation techniques ∂P and NES. We defined four different Majorana motion control regimes depending on the movement length and total time of the transport. With the ML methods we found in the non-adiabatic regimes I, II and III a novel counter-intuitive strategy that makes use of dressed jumps at the start and end of the protocol and in between the jumps has a period of nearly constant velocity motion. In the long-time regime IV we recovered the optimal superadiabatic ramp up and down protocols that were studied in previous works. In addition, we showed that the ML protocols retain their high efficiency in the proximity-coupled semiconducting nanowire model and that they outperform parameterised bang-bang protocols. Interestingly, from an SA search restricted to bang-bang protocols, we obtained optimal bang bang protocols that mimic the behaviour of the optimal ML protocols, but with a lower performance in regimes I, II and III.

²³For which the rate depends on the disorder strength.

²⁴And not so good for the superadiabatic protocols.

On a theoretical level we analysed a parameterised JMJ model strategy for the ML protocols in the non-adiabatic regimes. We found that without dressing the jumps with backward motions an approximate expression for the infidelity can be derived by evaluating the protocol in the moving frame. Thus, we learned that a possible explanation for the strategy is that the jumps try to target the steady moving-frame ground state during the middle part of the protocol. The additional backward motions can then be used to further maximise this overlap. By optimising these jumps we additionally showed that we can obtain JMJ strategies that perform competitively to the ML protocols and are robust to weak disorder in the wire.

For future work it will be interesting to investigate if the JMJ strategy has a connection with other protocols studied in the optimal control literature. The bare JMJ strategy could possibly be seen as a special form of a bang-off-bang protocol [174, 80]. Another example of a possibly related family of protocols are the so-called bang-anneal-bang protocols that were discovered in [175] for a quantum annealing problem. Analysing the performance of these protocols in the Majorana nanowire or another many-body system will likely tell if they are related in some way. Furthermore, our theoretical analysis for the JMJ strategy can potentially shed light on the physical working mechanisms behind the success of these protocols.

Another direction to explore is to apply the optimisation methods to different setups for the braiding of Majoranas and compare the performance of the obtained control protocols. One could, for example, think of the earlier discussed measurement-only or controlled coupling schemes. To apply NES to these schemes one needs to define a specific cost measure that encodes the control objective for which the value can be determined in some way, either numerically or experimentally. In addition, for ∂P one needs to have a numerical model of the cost measure that one can differentiate through. The advantage of ∂P and NES is that any potential additional constraints of these setups or a combination of different performance measures can be easily incorporated within the cost functional.

Lastly, from an optimisation perspective it will be interesting to compare the (computational) efficiency of different RL optimisation algorithms or numerical optimal control methods to the efficiency of the NES and ∂P methods discussed here. We saw for example that, compared to the SA method, the ML methods ∂P and NES require much less numerical evaluations of the infidelity. It remains to be seen how many evaluations other ML algorithms such as the Watkins Q-learning algorithm [76] and the Policy Gradients [77] need, and if they can do better. To make definite statements about which optimisation methods has the highest efficiency, however, one needs to have a proper understanding of the convexity of the Majorana-control cost landscape. Investigating this landscape and how it changes in the presence of disorder would be another interesting avenue for future studies.

Chapter 5

The Effect of Repulsive Interactions on the Control of Majorana Zero Modes

In the last two chapters we looked at controlling Majorana zero modes in superconducting nanowires. For this we made use of a mean-field description for topological p-wave superconductivity via the Kitaev chain and proximity-coupled superconductor lattice models. Here the system can be described effectively by a set of non-interacting quasiparticles because the Hamiltonian is quadratic in the creation and annihilation operators $c_x^{(\dagger)}$ for the electrons. This has several advantages since the mean-field model can be solved and simulated numerically without requiring access over the full many-body Hilbert space. This means that we can efficiently optimise the ground-state infidelity and obtain control strategies for the transport of the Majoranas.

To make contact with real-world topological quantum devices, however, it is important to investigate possible additional effects caused by interactions [176, 177, 178, 179, 180, 181]. It remains to be seen, for instance, which aspects of the jump-move-jump (JMJ) control strategy for Majorana transport are robust and if the target state fidelity does not rapidly decrease to zero in the presence of interactions. Furthermore, one might also reasonably worry about the possibility for completely new types of interaction driven qubit errors [19]. Such errors could be seriously troublesome if they lead to undetectable corruption or decoherence of the topologically protected quantum information. One example of such an error is a dynamically generated *phase error* between the even and odd parity sectors of an interacting topological system. These phase mismatches can occur due to interaction driven energy splittings [180, 182] between the bulk energy modes as explained in Fig. 5.1.

In this chapter, we investigate these possible additional interaction driven effects, and show that, in fact, a (large) part of the intuition gained for the mean-field models can be carried over to an interacting model for topological superconductivity. Specifically, we focus on the effects of a screened (coulomb) interaction [183, 184, 39] between two neighbouring electrons on the




FIGURE 5.1: As a consequence of the presence of interactions [$u_x \neq 0$ in Eq. (5.1)] splittings occur between the excited states of different topological sectors, here labelled e for the even parity sector and o for the odd parity sector. This opens up an apparent phase error loophole if a local process $\delta_L(t)$ is able to couple the protected ground states with the bulk energy modes (yellow arrow). During such a process relative dynamical phases $e^{-i\epsilon_j t}$ between the sectors are picked up in the bulk which when returned back to the ground-state manifold $|0\rangle_e$ and $|0\rangle_o$ potentially leads to an undetectable phase error in the topological qubit. In this chapter we show based on TQO that this error does not happen for times shorter than a time proportional to the system size L .

lattice. We model this by adding the term

$$\mathcal{H}_I = \sum_{x=1}^{N-1} u_x c_x^\dagger c_x c_{x+1}^\dagger c_{x+1} \quad (5.1)$$

to the non-interacting Kitaev chain Hamiltonian \mathcal{H}_0 of Eq. (3.1). u_x is the interaction strength, which for positive values is repulsive. It can be seen that this interaction term is not quadratic in the creation (annihilation) operators, which means that the simulation of the interacting system is a nontrivial and computationally expensive task.

By using tensor network simulation techniques for this interacting system we show numerically that the *jumps* in the jump-move-jump strategy are robust, whereas during the *move* part the ground-state fidelity decreases with increasing interaction strength u_x ¹. Moreover, we show,² by using an argument based on the topological quantum order (TQO) property [48, 185, 186], that the potential additional phase error described above is suppressed with the size of the system L , such that by making the system longer it can be effectively removed. This means that the phase error occurring in this interacting system is not any different from the same error in the non-interacting (mean-field) system. This rather peculiar result can be used to derive constraints on the bulk energy splittings between excited states in a complicated interacting many-body system [74].

We have structured this chapter as follows. In section 5.1 we show that the interacting

¹These results were published in [73], coauthored by the main author of this thesis.

²This will be an extended discussion of the results presented in [74], coauthored by the main author of this thesis.

Kitaev chain model inherits the TQO property from the non-interacting model, which means that it retains some of the useful properties of topological quantum computation. To make this argument, we will formally define TQO and discuss a notion called quasiadiabatic continuation. In section 5.2 we show the consequences of the interacting term on the energy spectrum, including the bulk energy splittings between different topological sectors. Afterwards, in section 5.3, we present numerical results of the various protocols for Majorana transport, such as the JMJ strategy that we discussed in the last chapters, but now in the presence of interactions. In section 5.4 we then show by using the TQO property that the bulk energy splittings do not lead to any additional phase error. Finally in section 5.5 we comment on the consequences of this result for the bulk energy spectrum and conclude.

5.1 Topological Order and Quasiadiabatic Continuation

The ground-state degeneracy and the protection against local perturbations that make some topological systems so interesting for the application in quantum computation are a consequence of a deeper, mathematical, property known as topological quantum order (TQO) [48, 185]. In this section we formally define the TQO property and show that the non-interacting Kitaev Chain Hamiltonian, \mathcal{H}_0 , possesses it. Then we show how the interacting model $\mathcal{H} = \mathcal{H}_0 + \mathcal{H}_I$ inherits approximate TQO from the free model, utilising quasiadiabatic continuation. This means that some of the (nice) features of the non-interacting model extend to the more complicated interacting model, and also that we can use it to argue that the interaction driven phase error is suppressed.

Before defining TQO we first need some preliminary notions and definitions. For our arguments we consider gapped lattice Hamiltonians which are a sum of Hermitian operators H_i that act only on nearest neighbour³ sites

$$H = \sum_i H_i. \quad (5.2)$$

Furthermore, we assume that we are looking at an anyonic model, which means that H has fractional excitations. The (interacting) Kitaev chain model \mathcal{H} is a specific example of such a Hamiltonian.

Since we are looking at anyon models, we can define different topological sectors corresponding to the different anyon flavours in the model. For simplicity we restrict here to two topological sectors (even and odd), for which we label the ground states $|e\rangle \equiv |0\rangle_e$ and $|o\rangle \equiv |0\rangle_o$, as is the case for the Kitaev chain. However, the arguments can be generalised to other anyon models with more than two distinct topological sectors.

³Most Hamiltonian models in condensed matter physics fall into this category. This includes both fermionic and bosonic models. In addition, we note that we not restrict to a specific dimension d of the lattice.

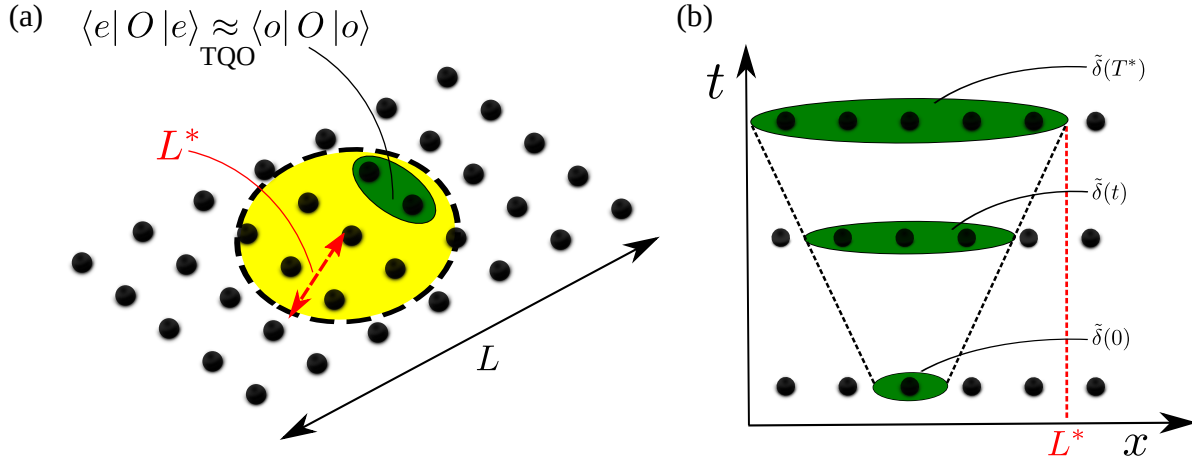


FIGURE 5.2: (a) Schematic explanation of the topological quantum order property. Local operators, such as the 2-local operator O (green region), that act within a region of size $\max L^* = cL$ (yellow region) have approximately the same expectation values, see Eq. 5.4, between the different parity sectors $|e\rangle$ and $|o\rangle$. Since L^* scales with L , we can always increase the system length L to include larger operators O . (b) We are specifically interested in initially local time dependent perturbations $\tilde{\delta}(t)$ whose support in space (green region) grows, by the Lieb-Robinson result [187], in time. After a time T^* the support becomes larger than L^* (non-local) and differences in the expectation values between the sectors can start occurring. Thus, by making the system larger we can increase T^* .

Lastly, for the definition of topological order, we also need to define what we mean by *local*. We define a generic operator O to be k -local if it acts in a region of size k on the lattice. That is, O can be written as

$$O = \sum_{A \in S(k)} O_A \quad (5.3)$$

where A is a subset of the ball $S(k)$ with radius k on the lattice, see Fig. 5.2 for an example. The operators O_A in the subset A can be connected, as in a connected product of neighbouring operators (a string), or disconnected, where O_A constitutes a sum of disconnected strings of size most k . For example, the nearest neighbour operators H_i in the generic Hamiltonian H (Eq. 5.2) are 2-local operators.

The definition for topological quantum order is then given as follows.

Definition (TQO): Assume that there exists some length L^* that scales with the total size (radius) L of the system via some constant $c > 0$ as $L^* > cL$. Then the ground states $|e\rangle$ and $|o\rangle$ satisfy topological quantum order⁴ if for every local operator, O , supported in a region $S(L^*)$ of diameter at most L^* , we have

$$\langle e|O|e\rangle = \langle o|O|o\rangle + \mathcal{O}(e^{-L/\xi}) \quad , \quad (5.4)$$

⁴See Fig. 5.2 for a schematic explanation of this definition.

for some constant $\xi > 0$. That is, local operators cannot be used to distinguish the two sectors, up to exponential corrections in system size L .

To make this formal (abstract) definition more concrete we can show that the ground states of the non-interacting Kitaev chain, \mathcal{H}_0 in Eq. (3.1), satisfy the TQO property. We recall that within the topological phase ($\mu \leq |2t|$) the two degenerate ground states for \mathcal{H}_0 can be defined as, see section 3.2.1,

$$|e\rangle \equiv \prod_{i=1}^{L/2} \beta_i^\dagger |\text{c-vac}\rangle, \quad |o\rangle \equiv \beta_0^\dagger |o\rangle. \quad (5.5)$$

Here $\beta_0^\dagger = \frac{1}{2}(\Gamma_L - i\Gamma_R)$ is the creation operator of a Dirac fermionic zero mode, and $\Gamma_{L/R}$ are the hermitian Majorana bound state operators acting on the left and right edges of the one-dimensional wire. We can write these Majorana operators as

$$\Gamma_L = i \sum_{x=1}^L (c_x^\dagger - c_x) u_L(x), \quad (5.6)$$

$$\Gamma_R = \sum_{x=1}^L (c_x^\dagger + c_x) u_R(x), \quad (5.7)$$

where $u_L(x)$ and $u_R(x)$ are the exponentially decaying Majorana functions at the left and right of the system. For simplicity we consider $\Delta = w$ and $\mu = 0$ for which we have perfectly localised Majoranas, i.e. $u_L(x) = \delta_{x,1}$ and $u_R(x) = \delta_{x,L}$. To show TQO we can now write

$$\langle e|O|e\rangle = \langle o|\beta_0 O \beta_0^\dagger|o\rangle \quad (5.8)$$

$$= \langle o|\beta_0 \beta_0^\dagger O|o\rangle + \mathcal{O}(e^{-L/\xi}) \quad (5.9)$$

$$= \langle o|O|o\rangle + \mathcal{O}(e^{-L/\xi}), \quad (5.10)$$

which holds when $[O, \beta_0^\dagger] = \mathcal{O}(e^{-L/\xi})$. Note that this is trivially true when O is *not* parity preserving, in which case both $\langle e|O|e\rangle$ and $\langle o|O|o\rangle$ are exactly zero. When O is parity preserving, one can show that $\langle e/o|\beta_0 [O, \beta_0^\dagger]|e/o\rangle = \mathcal{O}(e^{-L/\xi})$ provided that O is supported in a region smaller than L^* , see App. D.

The result that \mathcal{H}_0 has TQO shows that the topological encoding of quantum information in the ground states of the non-interacting system is robust against local error processes. This is because local perturbations O are not able to distinguish between the two distinct topological sectors $|e\rangle$ and $|o\rangle$. Only when an excitation is able to travel from one edge of the system to another, and effectively becomes non-local, can we start observing differences between the two sectors. Furthermore, another consequence of the TQO property is that the ground states of \mathcal{H}_0 are exponentially degenerate⁵. This can be seen from the fact that the non-interacting

⁵This is true for any system possessing the TQO property. So any topological system possessing TQO has a robust ground-state degeneracy that can potentially be used to encode quantum information.

Kitaev chain Hamiltonian is a sum of 2-local operators, which, by virtue of the TQO property, are not able to distinguish between the two sectors. In total there are $\mathcal{O}(L)$ terms in the sum which are all exponentially suppressed with L , which means that the sum is also exponentially suppressed. For completeness, we include this calculation in detail in App. D.

We cannot use the same construction to show TQO for the interacting model \mathcal{H} , since for this model we technically do not have a zero energy mode.⁶ Nor can we obtain expressions for the many-body ground states. However, we can show that \mathcal{H} possesses the TQO property by locally adiabatically⁷ connecting it to the non-interacting model. This construction is known as quasiadiabatic continuation [186]. In the following will review this construction in detail and apply it to the Kitaev chain. For formal (technical) proofs we refer the reader to [186].

In quasiadiabatic continuation we define a set of parameterised Hamiltonians given by

$$H(s) = \mathcal{H}_0 + sH_\delta, \quad (5.11)$$

where H_δ is some local perturbation away from the original unperturbed Hamiltonian⁸ \mathcal{H}_0 (the non-interacting Kitaev chain). We tune s adiabatically from the unperturbed $H(0) = \mathcal{H}_0$ to the perturbed Hamiltonian $H(1) = \mathcal{H}_0 + H_\delta$. We denote the instantaneous eigenstates along the process by $H(s)|\psi_s^i\rangle = E_s^i|\psi_s^i\rangle$, and require that the energy gap between the ground-state manifold and first excited states $E_{\text{gap}}(s) = E^1(s) - E^0(s)$ remains open for every s .⁹ Importantly, we note that we do not require or assume the states in the ground-state manifold to be degenerate or satisfy TQO. This manifold only constitutes a set of states that is separated by a large energy gap from the bulk states.

We can then define a special unitary *continuation* operator $V(s) = \sum_i |\psi_s^i\rangle \langle \psi_0^i|$ that takes us from the unperturbed eigenstates $|\psi_0^i\rangle$ to the perturbed eigenstates $|\psi_s^i\rangle$ of $H(s)$. This operator takes local operators O acting on the perturbed system into "dressed" operators

$$O_{\text{dres}}(s) = V(s)^\dagger O V(s), \quad (5.12)$$

which act on the unperturbed system. This means that we can relate expectation values¹⁰ in both pictures via

$$\langle \psi_s^i | O | \psi_s^i \rangle = \langle \psi_0^i | V(s)^\dagger O V(s) | \psi_0^i \rangle. \quad (5.13)$$

However, this does not mean that the TQO property in Eq. (5.4) of the unperturbed system \mathcal{H}_0 , directly implies TQO for the perturbed system H . This is because the dressed operators

⁶Here we mean that we cannot have non-interacting quasiparticles modes in this interacting system, and hence not a zero mode. However, we can formally find zero modes defined via the commutator [188, 189].

⁷Without closing the gap.

⁸For simplicity of the argument we restrict here to the non-interacting Hamiltonian, but in full generality quasiadiabatic continuation can be applied to any gapped system.

⁹We use the label $i = 1$ for the first excited state, with energy $E^1(s)$, above the possibly many degenerate ground states with energy $E^0(s)$.

¹⁰Here we have not restricted to the ground-state manifold yet and i runs over all the states.

$O_{\text{dres}}(s)$ are by no means local, as the unitary continuation operators $V(s)$ involve sums over possibly non-local eigenstates $|\psi_s^i\rangle$.

To resolve this, in [186] it is proven that $O_{\text{dres}}(s)$ can, in fact, be approximated by a local operator acting within a distance of maximum size l from O , if one allows for a small error, on the order of $\mathcal{O}(e^{-l/\xi})$, between the expectation values of states in the ground-state manifold. For this a special local *quasiadiabatic* unitary continuation operator $\tilde{V}(s)$ can be defined, which replaces the former adiabatic operator $V(s)$. Similar to before, we can define dressed operators as

$$\tilde{O}_{\text{dress}} = \tilde{V}(s)^\dagger O \tilde{V}(s) \quad (5.14)$$

and prove¹¹ that the expectation values of the states in the ground-state manifold are related by

$$\langle \psi_s^0 | O | \psi_s^0 \rangle = \langle \psi_0^0 | Q_0(s) \tilde{O}_{\text{dress}} Q_0^\dagger(s) | \psi_0^0 \rangle + \mathcal{O}(e^{-l/\xi}). \quad (5.15)$$

Here $Q_0(s)$ is some operator that only acts within the ground-state manifold, see [186]. The important difference with Eq. 5.13 is that \tilde{O}_{dress} is an operator with a support up to length l around the support of the original undressed operator O .

With this we can show that the perturbed system \mathcal{H} satisfies an approximate (with error $\mathcal{O}(e^{-l/\xi})$) TQO condition. By quasiadiabatic continuation it follows from Eqs. 5.15 and 5.4 that, for local operators O and provided that $2l$ is smaller than L^* , we have

$$\langle e |_s O | e \rangle_s = \langle o |_s O | o \rangle_s + \mathcal{O}(e^{-l/\xi}). \quad (5.16)$$

This construction is completely general¹² and only depends on the gap remaining open throughout the whole quasiadiabatic tuning process, i.e. for every s . In particular it does not depend on the specific form of the perturbation H_δ , as long as it is a sum of local terms. This means that we can quasiadiabatically continue the interacting Kitaev chain Hamiltonian to the non-interacting Hamiltonian, since $H_\delta \equiv \mathcal{H}_I$ is a sum of local nearest neighbour terms. The exact topological order (with error $\mathcal{O}(e^{-L/\xi})$) from the non-interacting system can thus be connected to approximate TQO in the interacting system.

As another example of this general quasiadiabatic continuation approach, we can show that the moving-frame Hamiltonian for the Kitaev chain in Eq. (3.16) possesses TQO for velocities lower than the critical velocity $v < v_{\text{crit}}$. We again start from the unperturbed Hamiltonian \mathcal{H}_0 of the non-interacting Kitaev chain. Then we define the perturbation to be equal to the term resulting from the unitary transformation to the moving frame¹³ with velocity

¹¹The proof is quite involved so we will not repeat it here and refer the interested reader to [186].

¹²In [186] a few examples are shown of how the degenerate ground states of the Ising model and some fractional Quantum Hall systems are robust against small perturbations.

¹³This term can be obtained by Fourier transforming the diagonal term $\sim vk$ in Eq. (3.16) back to real (and discrete) space.

v , which gives

$$H_\delta = \sum_{x=1}^{N-1} \frac{v}{2i} (c_x^\dagger c_{x+1} - c_x^\dagger c_{x-1}). \quad (5.17)$$

Now, since H_δ consists only of local 2-body terms, and since the gap remains open for velocities $v < v_{\text{crit}}$, we can quasiadiabatically connect it to \mathcal{H}_0 , and show that the ground states of moving-frame Hamiltonian satisfy an approximate TQO condition.

Lastly, we remark that, although we have been focusing on the Kitaev chain model, this construction can, in theory, also be applied to the proximity-coupled semiconducting nanowire model. Here one again needs to start from some set of parameters of the Hamiltonian for which one can show exact TQO. Then one can quasiadiabatically connect it, via a path in parameter space along which the energy gap remains open, to the interacting Hamiltonian. However, in this case, the type of interaction [183] can be different since it involves spin degrees of freedom. Nevertheless, as long as the interaction term is local, the quasiadiabatic continuation can be applied and an approximate TQO condition can be derived.

5.2 Spectrum of the Interacting Kitaev Chain

While the TQO property ensures that the interacting system has (up to finite size corrections) an exact ground-state degeneracy,¹⁴ in this section we show that this degeneracy does not carry over to the excited states above the gap, and that interactions lead to small energy splittings between the even and odd parity sectors. We will also briefly look at the spectral gap itself and show numerically that when the interaction strength u_x becomes too strong it closes, see Fig. 5.4. The closing of the gap is important because it means that the topological order is destroyed, and that the quasiadiabatic continuation construction from the last section fails.

5.2.1 Interaction Driven Bulk Energy Splittings

To show how the interacting term in Eq. (5.1) leads to small energy splittings between bulk states in the even and odd parity sectors we review here the arguments given in [180] and [182]. Unlike for the non-interacting system, we cannot solve the spectrum exactly by finding the quasiparticle excitations. To make quantitative statements about the excited states of the interacting system, one therefore requires other methods, such as the Bethe ansatz [190] or exact diagonalisation [191]. However, these methods are generally hard to apply or limited to small system sizes. Another approach, which we adopt here, is based on perturbation theory in which the interaction term \mathcal{H}_I is treated as the perturbation. Although this approach is

¹⁴By the same argument as for the non-interacting Kitaev chain, i.e. the fact that the Hamiltonian consists only of 2-local terms, see App. D.




FIGURE 5.3: In the presence of interactions u_x there are small energy gaps (avoided crosses) between energy bands with different zero mode occupancy number. These splittings result from the matrix elements $B_{nm} \sim L_{nm} \pm R_{nm}$ that contain one zero mode operator β_0 in Eq. (5.19). $l \equiv L_{nm}$ comes from the contribution of the left Majorana zero mode and $r \equiv R_{nm}$ from the right Majorana zero mode. This figure was, with permission from the main author, taken from [182].

perturbative, it allows us to show the origin of the splittings, which one can verify numerically (see Fig. 5.3).

To give the argument for the splittings, we first expand the interaction term \mathcal{H}_I in the eigenstates of the non-interacting system \mathcal{H}_0 . We assume we have diagonalised the non-interacting system $u_x = 0$ in terms of the mode creation operators β_i^\dagger , i.e. $\mathcal{H}_0 = \sum_i (\beta_i^\dagger \beta_i - 1/2)$. In addition, we assume we are in the topological phase and have one exact zero mode β_0^\dagger in the system. We can then employ the inverse BdG quasiparticle transformation¹⁵

$$c_x^\dagger = \sum_n U_{xn}^* \beta_n^\dagger + V_{xn} \beta_n \quad (5.18)$$

to write the interaction \mathcal{H}_I in terms of the non-interacting mode operators $\beta_i^{(\dagger)}$. This gives

$$\begin{aligned} \mathcal{H}_I = \sum_{x=1}^{L-1} u_x c_x^\dagger c_x c_{x+1}^\dagger c_{x+1} &= \sum_{x=1}^{L-1} U_x \left[\sum_k U_{xk}^* \beta_k^\dagger + V_{xk} \beta_k \right] \times \\ &\left[\sum_l U_{xl} \beta_l + V_{xl}^* \beta_l^\dagger \right] \times \\ &\left[\sum_m U_{x+1m}^* \beta_m^\dagger + V_{x+1m} \beta_m \right] \times \\ &\left[\sum_n U_{x+1n} \beta_n + V_{x+1n}^* \beta_n^\dagger \right]. \end{aligned} \quad (5.19)$$

Expanding out the four products in the equation above is tedious and leads to a long expression, see [40, 180]. We recognise however that we get a large sum over terms that

¹⁵The details of the BdG formalism are discussed in App. A

are products of a combination of four mode creation or annihilation operators, i.e. terms $\sim \beta_i^{(\dagger)} \beta_j^{(\dagger)} \beta_k^{(\dagger)} \beta_l^{(\dagger)}$. We note that sometimes these products can be reduced to 2-body terms, but we will focus here on the most general 4-body terms. In this large sum we can then distinguish, depending on the zero modes, between three different types of terms:

1. Terms that do not contain the zero mode operators β_0 or β_0^\dagger , which we call (A).
2. Terms that include a single zero mode operator, which we call (B).
3. Terms that include an even multiple (two or four) of the zero mode operators, which we call (C).

We evaluate the effect of each of these terms in the eigenbasis of the diagonalised non-interacting system

$$|0n\rangle_e, \quad |1n\rangle_e, \quad |0n\rangle_o, \quad |1n\rangle_o, \quad (5.20)$$

in which n corresponds to the n -th many-body eigenstate (see Fig. 5.1). The first entry in the ket corresponds to the occupancy of the zero mode. The labels e/o still distinguish the even and odd parity sectors. We are interested in the matrix elements $\langle \mathcal{H}_I \rangle_{e/o}$, which we can evaluate by using the expansion in Eq. (5.19). Specifically, for the splittings, we want to know the differences between the even and odd subspace Hamiltonians

$$H_e = \langle \mathcal{H} \rangle_e = E_e + A_e + B_e + C_e, \quad (5.21)$$

$$H_o = \langle \mathcal{H} \rangle_o = E_o + A_o + B_o + C_o. \quad (5.22)$$

The $A_{e/o}$, $B_{e/o}$, $C_{e/o}$ matrices correspond to the evaluation of the terms defined above in the even and odd sectors respectively. $E_{e/o}$ are the contributions coming from the non-interacting Hamiltonian $\langle \mathcal{H}_0 \rangle_{e/o}$, which, by the presence of the exact zero mode, are equal in both sectors $E_e = E_o$.

The contributions $A_{e/o}$, from the terms without zero mode operators, are also equal in both sectors. This is because for these terms \mathcal{H}_I does not change the zero mode occupancy, and for every odd parity state there is a corresponding even parity state with the opposite zero mode occupancy. In a similar way we can argue that the terms $C_{e/o}$ are the same in both sectors because the overall occupancy of the Majorana zero mode is not changed. The terms $B_{e/o}$ are different in both sectors, since they switch the occupancy of the Majorana mode. This means that we can show that the corresponding matrix elements can be written as $B_{nm} \sim L_{nm} \pm R_{nm}$, where L_{nm} results from the contribution of the left Majorana mode and R_{nm} from the right mode.¹⁶ Importantly, the sign of the R_{nm} contribution depends on the sector, such that a difference between the two occurs, see Fig. 5.3 (a) for a schematic example.

¹⁶This can be most easily seen from looking at a small example of a few bulk modes and the zero mode.

This observation allows one to argue on a perturbative level for topologically protected degeneracies between even and odd eigenstates in cases where there are no overlaps between bands with different quasiparticle number [180]. Such scenarios arise when one considers the model around the flat-band limit. However, away from this limit, when there are overlapping energy bands with different zero mode occupancy, the even and odd parity subspace Hamiltonians are different. This means that the eigenvalues are also different and we get small energy splittings (avoided crossings) between the two sectors. The zero mode in this interacting system is therefore called a 'weak' zero mode instead of a 'strong' zero mode (fully degenerate spectrum) [192, 193, 194, 195]. Despite this, we will show below, in section 5.4, that these splittings do not lead to any additional phase error.

5.2.2 Interaction Driven Reduction in the Topological Energy Gap

Aside from inducing bulk energy splittings, the interaction term \mathcal{H}_I also reduces the size of the energy gap E_{gap} . Intuitively this can be understood from the idea that repulsive interactions compete with the superconducting pairing interaction Δ , and hence lead to a smaller effective energy gap (weaker superconductor). This reduction can be found for example by again utilising mean-field theory and running some Hartree-Fock-Bogoliubov minimisation to find the effective quasiparticles [40].

Here we make use of tensor network techniques to find the many-body spectral gap numerically, see App. C. First we use the density-matrix renormalization group (DMRG) [61, 62] to find the even and odd $|e/o\rangle$ ground states of the interacting system \mathcal{H} in terms of matrix product states (MPS).¹⁷ Then we evolve these MPS ground states in time by using the time-dependent variational principle (TDVP) [59, 60], and compute the single particle Green's functions defined by

$$G^{e/o}(x, x', t) = e^{iE_0 t} \langle e/o | U(x, x', t) | e/o \rangle. \quad (5.23)$$

Here the unitary operator is given by

$$U(x, x', t) = (c_{x'}^\dagger + c_{x'}) e^{-i\mathcal{H}t} (c_x^\dagger + c_x). \quad (5.24)$$

We fix $x = 1$ and we let x' run over all L lattice sites. By Fourier transforming these Green's functions $G^{e/o}$ in space and in time we can find the spectral correlation functions¹⁸

$$A^{e/o}(k, \omega) = \int_{-\infty}^{\infty} dt e^{-i\omega t} \int_{-\infty}^{\infty} dx' e^{-ikx'} G^{e/o}(1, x', t). \quad (5.25)$$

¹⁷DMRG can be applied (computationally) efficiently since the ground states satisfy an area law for the Von-neumann entanglement entropy, see [196] and App. C.

¹⁸Note that in our simulations everything is discrete, so in practise we use a discrete Fourier transforms.

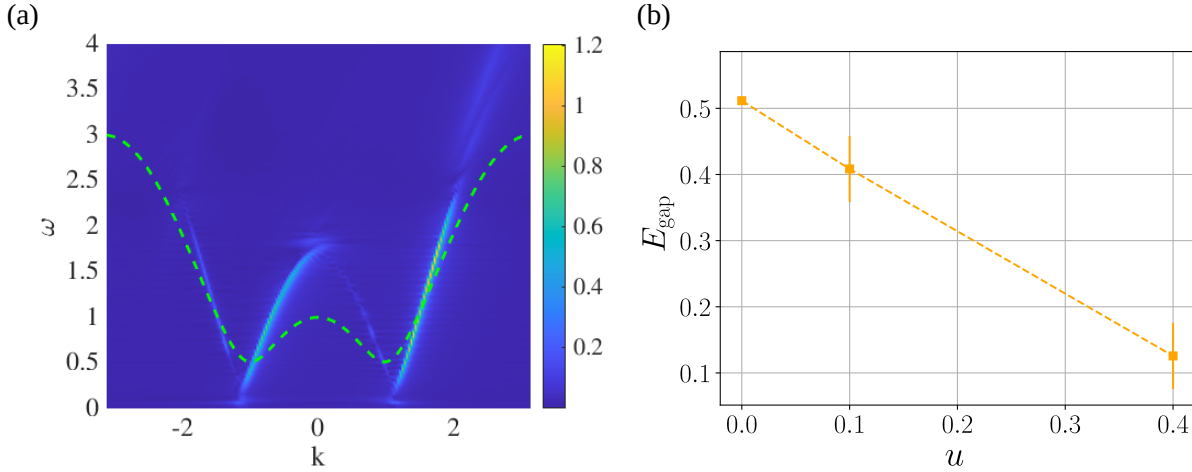


FIGURE 5.4: (a) Spectral function $A^e(\omega, t)$ as defined in Eq. (5.25) with interaction strength $u_x = 0.4$ and other parameters $\mu = -1$, $\Delta = 0.3$, $w = 1$, and $N = 200$. The Fourier transform was cutoff in time to smoothen the resolution in energy ω . We obtain a 'w'-like dispersion with a smaller gap and slightly larger k_F compared to the non-interacting single particle dispersion (dashed line). (b) The topological energy gap $E_{\text{gap}} = E^1 - E^0$ as a function of interaction strength u . The gap was resolved by computing the spectral function A^e with the same parameters, aside from u_x , as in panel (a). We find that the gap reduces approximately linearly with the interaction strength. Note that we include error bars on the finite interaction strength results because of the resolution of the spectral correlation function and the manual reading of the gap.

These spectral functions $A^{e/o}$ are related to the single excitation energies [183] and can be used to determine the spectral gap as shown in Fig. 5.4. In panel 5.4 (a) we see that, even in the presence of strong interactions $u_x = 0.4$, we get a roughly 'w'-like looking dispersion for the excitation energies. We recall from section 3.2.1 that this dispersion is characteristic for a p-wave superconductor. However, this time the energy gap is significantly reduced compared to the free non-interacting dispersion (dashed line). As we show in panel 5.4 (b), this reduction seems to scale linearly with the interaction strength u_x . As a consequence, the gap is expected to close for sufficiently strong interactions. We note that these results are in agreement with the results found in [176] which showed, with a different (DMRG only) technique, that the topological gap reduces with increasing interactions.

5.3 Robustness of Majorana transport with respect to interactions

The fact that in the presence of interactions $u_x \neq 0$ the gap in the spectrum of the Kitaev chain in topological phase gets reduced and also that the spectrum is not exactly two-fold degenerate anymore makes asking whether the protocols for Majorana transport that we investigated in chapter 4 remain robust a valid question. In this section we answer this question numerically for the Jump-move-Jump (JMJ), Superadiabatic and oscillating boundary protocols by




FIGURE 5.5: Response of the system to boundary wall oscillations for four different interactions strengths. We see that the resonance peak distinguishing the adiabatic from non-adiabatic regimes shifts to lower values when increasing the uniform interaction strength u . For these simulations the left wall is oscillated with frequency ω and max velocity $v_{\max} = 0.1$ and we set the other parameters $N = 50$, $\mu = -1$, $w = 1$ and $\Delta = 0.5$.

encoding them in the time-dependent potential profile $V(x, t)$ ¹⁹ in the interacting Hamiltonian \mathcal{H} . For these simulations we use the same TDVP-MPS time-evolution approach as described for the spectral function. In here we time evolve the initial even and odd degenerate ground states $|e\rangle$ and $|o\rangle$ with the time-dependent Hamiltonian $\mathcal{H}(t)$ that follows the Majorana transport protocol. We then use these time evolved states to compute the target state infidelity²⁰ \mathcal{I}_τ by contracting the time evolved MPSs, see App. C. We will now discuss the results of these simulations for the different Majorana control strategies.

First in figure 5.5 we show the response of the system with respect to boundary wall oscillations with frequency ω for various interaction strengths. For this we oscillate the left domain wall position $x(t)$ encoded in $V(x, t)$ with a velocity $v(t) = \dot{x}(t) = v_{\max} \cos \omega t$. The specific infidelity measure we use is the qubit-loss (quantum information lost from the degenerate ground states) as defined by

$$P_{\text{loss}}(t) = 1 - |\langle e|U(t)|e\rangle|^2 - |\langle o|U(t)|o\rangle|^2 \quad (5.26)$$

with $U(t)$ the unitary evolution operator of the interacting system. Similarly to the non-interacting case discussed in section 3.3,²¹ we see that a clear resonance occurs that distinguishes between a regime in which there is almost no loss of quantum information (low frequency) and a regime (high frequency) where the loss decreases with increasing frequency.

¹⁹Recall Eq. (3.14), and see Fig. 5.7 for a two-wire Majorana-based topological qubit.

²⁰Recall that $\mathcal{I}_\tau \equiv 1 - \mathcal{F}_\tau = |\langle \psi_\tau | U(\tau) | \psi(0) \rangle|^2$ with $U(\tau)$ the time evolution operator is a quantitative measure for the robustness of the Majorana transport protocols as described in section 4.1.

²¹Compare Fig. 3.8 (a).

The Fermi-golden rule argument also directly applies here, which means that the low frequencies $\omega < \omega_{\text{res}}$ are not able to couple the ground states to the excited states since they are smaller than the energy gap E_{gap} . The resonance frequency is found to decrease with interaction strength, which can be directly attributed to the interaction driven reduction of the energy gap. These results can possibly be argued for by using mean-field theory, which results in an effectively smaller superconducting pairing Δ . However, we note that contrary to simulations with a smaller Δ , see Fig. 2 (b) in [164]²², we observe that interactions increase the qubit loss at all frequencies.

Next we compute the robustness of the optimal JMJ protocols for the domain wall position $x_L(t)$ in the non-adiabatic Majorana transport regimes²³ with respect to interactions. For this we take the specific JMJ protocols obtained for zero interaction strength with the exact same jump sizes and keep all the model parameters the same except turning on the interacting term \mathcal{H}_I . These protocols move the left Majorana from a position x_A to a position x_B in the interacting system. Importantly we note that for these protocols we have not performed any additional optimisation in the presence of interactions.

In Fig. 5.6 a) we show the instantaneous infidelity

$$I(t) = 1 - |\langle \psi_t | U(t) | \psi_0 \rangle|^2 \quad (5.27)$$

of the optimal JMJ in regime I for a few different interaction strengths. In here $|\psi_t\rangle$ is the instantaneous ground state of the interacting system $\mathcal{H}(t)$ which means that $I(t)$ gives insight in how far away the evolved state is from the adiabatic path. We observe that the initial and final jumps (at times $t = 0$ and $t = \tau - dt$) remain robust with increasing interactions since the change in $I(t)$ at these times does not (noticeably) change for different interaction strengths. During the *move* part of the protocol, i.e. between the jumps, however the instantaneous infidelity becomes gradually bigger over time and increases with interactions. This increase means that the final target state infidelity $\mathcal{I}_\tau \equiv I(\tau)$ becomes bigger with interacting strength. This can also be observed from Fig. 5.6 (c) where we have plotted \mathcal{I}_τ as function of u . As a comparison we have included data of a naive linear protocol as well which shows that the target state infidelity of the JMJ is superior to linear motion in regimes I, II and III at all interaction strengths.

These results can again possibly be explained on the basis of mean-field theory with a lower energy gap and potentially lower critical velocity v_{crit} .²⁴ A lower gap and critical velocity imply that moving at the same constant velocity will give a larger fidelity loss. This is especially important in regime I if the velocity of the move part becomes above the new

²²Here it was found that when Δ is reduced a topological qubit can become more resilient to high frequency noise $\omega \gg \omega_{\text{res}}$.

²³See sections 4.3 and 4.1 for the discussion of these protocols and the Majorana motion regimes respectively.

²⁴It remains to be seen however if there is also a moving frame in the interacting system in the same way as there is in the free system.

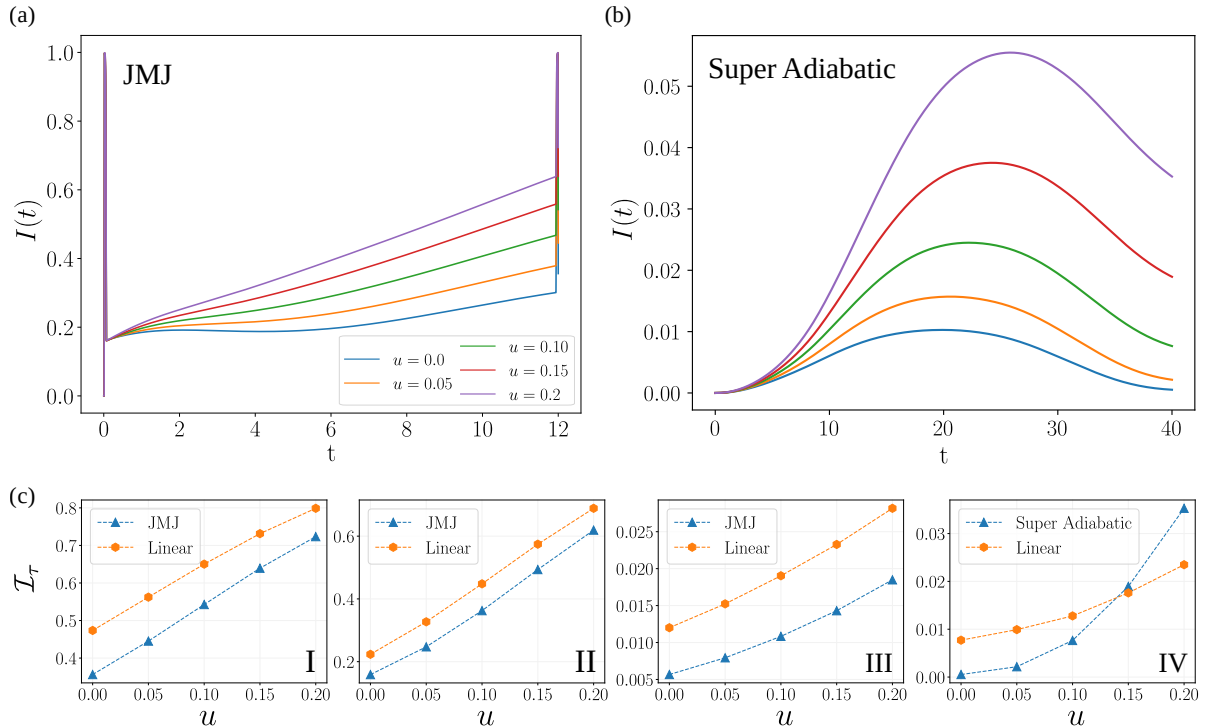


FIGURE 5.6: Overview of the robustness results of the Jump-move-Jump (JMJ) and superadiabatic Majorana control protocols with respect to uniform interactions u . In panel (a) the instantaneous infidelity $I(t)$ of the JMJ protocol in regime I is plotted. We observe that the jumps are robust with respect to interactions while during the move part the infidelity increases with interactions. Panel (b) shows $I(t)$ as well but then for the superadiabatic ramp up-down protocol in regime IV for the same interaction strengths as in (a). We see that for increasing interactions the infidelity increases significantly at all times. In panel (c) the final target state infidelity \mathcal{I}_τ of these protocols is plotted (blue) compared to a naive linear benchmark protocol (orange). Here for completeness we have also included regimes II and III. We see that the JMJ strategy outperforms the linear protocol in regimes I, II and III at all interaction strengths while the infidelity of superadiabatic protocol in regime IV becomes bigger than the linear protocol for sufficiently strong interactions. These results were obtained with TDVP-MPS simulations of the interacting system \mathcal{H} with parameters $N = 110$, $\mu = 1$, $w = 1$, $\Delta = 0.3$, $V_{\text{height}} = 30.1$, and $\sigma = 1$.

effective v_{crit} . For example for the biggest interaction strengths that we investigated the gap is expected to be almost fully closed so that any motion would result in a significant fidelity loss. Furthermore, on a mean-field level one generically finds that the couplings in the model become nonuniform and thereby effectively induce disorder in the wire. Disorder as we stated above in section 4.4 results in a strongly decreased v_{crit} which can also explain the increase in infidelity with interaction strength.

In addition we also look at the robustness of the superadiabatic ramp-up and down protocols that we found to be optimal in the long transport time regime IV after optimization in the non-interacting system. In Fig. 5.6 (b) we show the instantaneous infidelity $I(t)$ of this protocol as a function of interaction strength. We clearly observe that the superadiabatic protocol seems to break down since it is now not able to bring the infidelity fully down to zero during the ramp-down part, i.e. towards the end of the protocol. This means that the final infidelity

values \mathcal{I}_τ increase drastically with interaction strength as shown in Fig. 5.6 (c). Importantly for an interaction of about $u_x \approx 0.15$ the simple linear protocol obtains better values hence being a better strategy.

Two reasons similar to the ones discussed before can be given for this observed behaviour in regime IV. Firstly, we can argue again on the basis of mean field theory that it is caused by a lower v_{crit} but this would not directly explain why linear motion suddenly becomes better. The second possible reason is based on the fact that a lower gap means that the resonance frequency is lower as we showed for the oscillating boundary protocol in Fig. 5.5. This implies that the timescale τ_{res} above which regime IV is defined gets longer which could mean that by not changing the external constraint we effectively end up in one of the other regimes. If this is the case and we end up in one of the other regimes it is expected that the MJM strategy is optimal again. To test this one needs to run more simulations to determine the exact energy gap as a function of the interactions and also optimise in a corresponding free (mean-field) system with the same energy gap. This is quite involved and we leave this for exploration in future studies.

We thus find that the MJM strategy seems to be robust with respect to interactions whereas the superadiabatic protocols lose their low target state infidelity. We note however that since we have not performed the optimisations of the Majorana transport directly in the presence of interactions we cannot argue that the optimal MJM protocol from the non-interacting system is also the optimal protocol in the interacting system. It might be that there is another movement strategy that is able to outperform the MJM strategy in regimes I, II, and III in the presence of interactions that is totally different. Moreover, it remains to be seen what the optimal strategy is in regime IV²⁵ in the presence of interactions.

To investigate these questions one needs to embed the MPS-TDVP simulation method for the interacting system with the ML optimisation techniques NES and ∂P used in chapter 4. For the NES algorithm this poses no additional challenges other than a potential blow up of the computational complexity since many MPS-TDVP simulations need to be performed. Amending the ∂P requires more effort because one needs to code up the MPS-TDVP algorithm in a way such that each individual primitive is differentiable. In this regard we note that a step towards the realisation of this has already been made in [107] which showed that the DMRG and tensor network contraction algorithms are differentiable.

5.4 Suppression of the Interaction Driven Phase Error

In addition to the effect of interactions on the qubit-loss (or infidelity) error we look at the potential *phase error* in the topological ground-state manifold caused by interaction driven

²⁵Both where we define regime IV with respect to the non-interacting gap or with respect to the effective mean-field gap induced by interactions.




FIGURE 5.7: The 2-wire setup of the interacting topological quantum memory we consider to investigate the phase error numerically. Four Majorana Zero Modes form at the domain boundaries between topological and non-topological regions encoded by $V(x)$ that can be combined to make a topological qubit. We consider noise processes $\delta_x(t)$ that happen on the left of the system, i.e. near the left domain boundary, and in our arguments we effectively forget about the right wire.

splittings between bulk energy modes of different topological (parity) sectors (see Fig.5.1). Although the *phase error* is possibly more violent²⁶ and can be caused by a local perturbation here we show that it is suppressed with the length of the system L . For this we make use of an analytical argument based on the TQO property which we verify numerically with MPS-TDVP simulations of a two wire interacting topological memory as shown in Fig. 5.7. We follow here the discussion and results presented in [74] coauthored by the main author of this thesis.

5.4.1 Analytical Argument for *no Phase Error*

The *phase error* that we consider is a dynamically accumulated phase difference between the even and odd ground states $|e\rangle$ and $|o\rangle$ of the interacting system \mathcal{H} ²⁷. Specifically we say there is a phase error if the expectation values of the two ground states

$$\langle e|U(t)|e\rangle \neq \langle o|U(t)|o\rangle \quad (5.28)$$

differ after some time t in which $U(t)$ is the unitary time evolution operator that describes the evolution of the system.

Since we are interested in dephasing caused by the little energy splittings in the bulk modes we include a local time-dependent perturbation $\delta(t)$ that over time excites parts of the ground states to the bulk modes²⁸. The time-dependent unitary is then defined as

$$U(t) = \mathcal{T}\{e^{-i\int_0^t dt' \mathcal{H} + \delta(t')}\}. \quad (5.29)$$

²⁶As it is undetectable.

²⁷The two wire form of the potential profile $V(x, t)$ does not matter at this point as long as we have two degenerate topological groundstates.

²⁸We note that without a perturbation only a phase difference can be picked up by the relative energy splitting between the two ground states. This splitting is however exponentially small with the system size such that it always can be removed. Also this ground-state mechanism for phase error is not unique to the interacting system.

in which \mathcal{T} ensures the time ordering. At this point we do not put any restrictions on the form of this perturbation and only require that it is local, for example an oscillating left domain boundary of the Majorana wire.

To now show that there is no phase error we need to show that this unitary acts similarly in both topological sectors, i.e. $\langle e|U(t)|e\rangle \approx \langle o|U(t)|o\rangle$. We can achieve this by first transforming $U(t)$ to the interaction picture which gives

$$U(t) = e^{-i\mathcal{H}t} \times \mathcal{T}\{e^{-i\int_0^t dt' \tilde{\delta}(t')}\}, \quad (5.30)$$

in which $\tilde{\delta}(t) := e^{i\mathcal{H}t}\delta(t)e^{-i\mathcal{H}t}$. We can then write the expectation values for the even ($a = e$) and odd ($a = o$) sectors as

$$\begin{aligned} \langle a|U(t)|a\rangle &= \langle a|e^{-i\mathcal{H}t} \times \mathcal{T}\{e^{-i\int_0^t dt' \tilde{\delta}(t')}\}|a\rangle \\ &= e^{-iE_0t} \langle a|\mathcal{T}\{e^{-i\int_0^t dt' \tilde{\delta}(t')}\}|a\rangle, \end{aligned} \quad (5.31)$$

in which E_0 is the ground-state energy which is the same in both sectors up to corrections of order $\mathcal{O}(e^{-L}/\xi)^{29}$.

In this way we have removed the contribution from the unperturbed Hamiltonian \mathcal{H} and only need to look at the difference in expectation values of the perturbation $\tilde{\delta}(t)$ in the interaction picture. As this perturbation is now dressed it is not necessarily local and we cannot directly apply the TQO condition in Eq. (5.4) to argue that both expectation values are the same. However, by the Lieb-Robinson bounds [187] the support of the operator $\tilde{\delta}(t)$, and hence the time-ordered unitary $U(t)$, grows approximately linear in time with the Lieb-Robinson velocity v_l as shown schematically in Fig. 5.2 (b). This means that when the support remains smaller than the length L^* we can apply the TQO condition and argue that the expectation values $\langle a|U(t)|a\rangle$ are the same up to corrections of order $\mathcal{O}(e^{-L}/\xi)$. We therefore have no phase error for times $|t| < T^* \sim L^*/v$. Importantly the phase error can thus always be removed by increasing the system length L since $L^* = cL$.

We remark that this argument for no phase error can be extended to two separate initially local time-dependent perturbations and also to time-independent local unitary kicks, see [74]. Moreover since our argument only makes use of the TQO condition in combination with the spreading of operators by the Lieb-Robinson velocity it also applies to other gapped local Hamiltonians with topological order.

5.4.2 Numerical Verification

To verify these analytical results we perform numerical simulations of the interacting system \mathcal{H} with a local time-dependent perturbation $\delta_x(t)$. To model a real-world topological memory,

²⁹Recall that the ground-state degeneracy for the interacting system directly follows from the TQO condition.

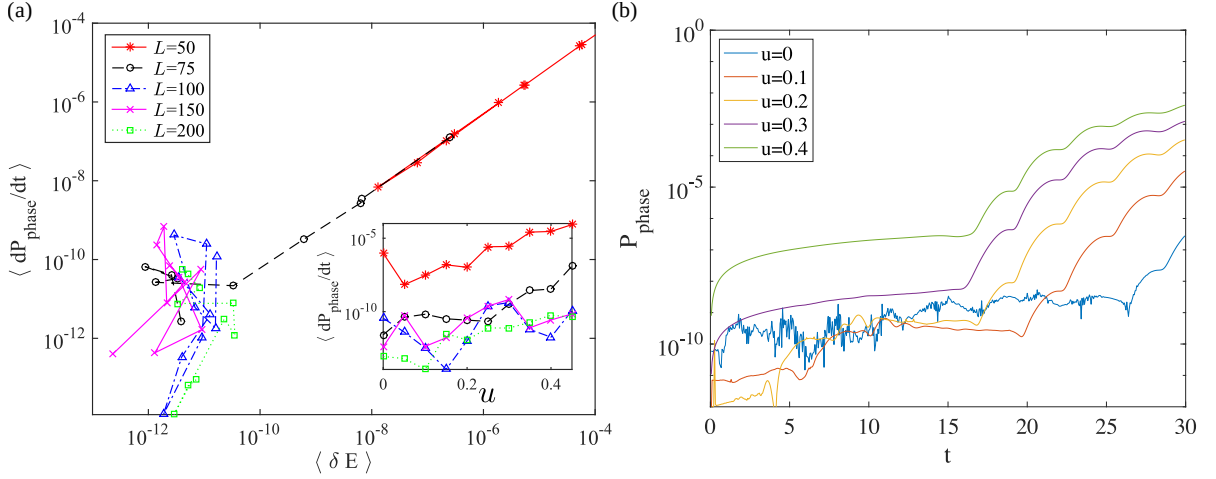


FIGURE 5.8: (a) Time averaged rate of phase error as a function of the ground-state energy splitting $\langle \delta E \rangle$ (main panel) and the uniform interaction strength u (inset) for a few different system sizes L . From about $\langle \delta E \rangle = 10^{-10}$ we start detecting a phase error that results from the ground-state splitting. We however see no additional error arising that could come from the mismatches bulk levels. The other parameters that we use are $\mu = -1.5$, $\Delta = 0.7$, $\omega = 1.0$ and $v_{\max} = 0.1$. (b) Phase error as a function of time for two time dependent perturbations $x_1(t)$ and $x_2(t)$ at the different edges of the left wire. After a time $T^* = L^*/v$ proportional to the system size the phase error suddenly kicks in. The propagation velocity v increases with the interaction strength u . For these simulations we use $L = 70$, $w = 1$, $\mu = -1.2$, $\Delta = 0.8$, $\omega = 1$ and $v_{\max} = 0.1$.

for which 4 Majorana zero modes are required, we encode the potential profile as

$$V_x = V_{\text{outer}}[f(x_1 - x) + f(x - x_4)] + V_{\text{inner}}[f(x - x_2) - f_3(x - x_3)], \quad (5.32)$$

in which we have four domain walls at the positions x_i for $i = 1, 2, 3, 4$ and $f(x) = (1 + \exp(-x/\sigma))^{-1}$ is a Sigmoid function. We have shown this setup schematically in Fig. 5.7 and note that we only focus on the two parity sectors $|e\rangle$ and $|o\rangle$ corresponding to the occupancy of the Dirac zero mode of the left wire $\beta_0 = \frac{1}{2}(\gamma_1 + i\gamma_2)$. To this end we choose the inner potential barrier V_{inner} and the separation $|x_3 - x_2|$ large enough such that the right wire is effectively disconnected. We quantify the phase error as

$$P_{\text{phase}}(t) = \frac{1}{4} |\langle e|U(t)|e\rangle - \langle o|U(t)|o\rangle|^2 \quad (5.33)$$

and note that this measure is zero for no phase error.

We first focus on a single oscillating perturbation $\delta_x(t)$ that acts on the left of the system. Specifically this perturbation corresponds to small oscillations of the left domain wall position $x_1(t) = (v_{\max}/\omega) \sin(\omega t)$ in the potential Eq. 5.32. As we want to couple to the bulk we

choose the frequency larger than E_{gap} (adiabatic). In Fig. 5.8 (a) we show the results of MPS-TDVP simulations for this perturbation in the interacting system as a function of the time-averaged ground-state energy splitting $\langle \delta E \rangle$ (main panel) and interaction strength u (inset). We observe that with our numerical methodology we are able to detect phase errors due to a ground-state splitting up to orders as small as $\langle \delta E \rangle \sim 10^{-10}$. When making the system size large enough such that the ground-state splitting $\langle \delta E \rangle \sim \mathcal{O}(e^{-L/\xi})$ becomes negligible, we do not detect any additional increase in P_{phase} .

We also consider the scenario in which we add a second time-dependent perturbation that oscillates the right domain wall position $x_2(t)$ with the same frequency and velocity as the left wall $x_1(t)$. In Fig. 5.8 (b) we show numerical simulation results for this case and we observe that while P_{phase} first slowly increases according to the ground-state splitting after some time, which we call T^* , the error shoots up. This is as expected from the argument above since the support of both of the two perturbations spreads out in time and at the time $T^* \sim L/(2v)$ when they start overlapping we expect a phase error³⁰. By making the Majorana wires longer we could delay and eventually even remove this effect. Moreover we note that exactly the same type of effect occurs in the non-interacting system which means that the observed phase error at T^* is completely understood on a mean-field level and is not necessarily attributed to the splittings between the bulk energies.

5.5 Summary and Further Work

In this chapter we have looked at the effects repulsive interactions on the control and decoherence of Majorana-based topological quantum devices. We saw that the topological quantum order property that ensures the stability of topologically encoded quantum information against local perturbations can be (quasiadiabatically) continued from the non-interacting to the interacting Kitaev chain model for p-wave superconductivity. This is important since it implies that the ground states of the interacting system remain degenerate and hence could still be used to make a topologically protected qubit.

Despite this we also showed that as a consequence of interactions the topological energy gap is reduced and, crucially, that the exact degeneracy between the excited states of the different topological sectors above the gap is broken. Although this seemed detrimental for the robustness of Majorana based topological memories, we showed, using an argument based on TQO, that this does not lead to any additional phase error other than the one coming from mean-field or finite size effects. We verified this numerically with tensor network simulation techniques. This also allowed us to show that the optimal JMJ Majorana control strategy

³⁰Note that the propagation velocity v depends on the specific perturbations and the parameters of the system. We can determine it numerically from the numerics by reading of the time T^* when the error kicks up.

from the free system still outperforms the naive linear benchmark protocol. In contrast, the robustness of the superadiabatic protocols gets destroyed in the presence of interactions.

This leaves us with a few different directions to explore for future work. Firstly, regarding the robustness of the Majorana control strategies, we could incorporate the tensor network simulation method within the ∂P and NES optimisation algorithms to search for optimal controls directly in the presence of interactions. While in theory this might be possible to do, it is questionable if this will lead to completely new insights or control strategies other than the ones already encountered in the free system. Thus, it might be better to do a proper mean-field analysis of the interacting term, and then optimise within the effective mean-field system. Such an analysis could also possibly provide an explanation for the breakdown of the robustness of the superadiabatic protocols.

Another, arguably more interesting, direction to explore are the implications of the no phase error results. The fact that there is no phase error can be used to derive a large set of constraints on the splittings between the excited states. This is partly explored in reference [74], coauthored by the author of this thesis, where a finite time cutoff of the Fourier transform in the spectral correlation function $A^{e/0}(\omega, t)$ is used to argue that the splittings need to be smaller than some finite resolution. It remains to be seen if more quantitative and exact bounds can be derived. We note that making such statements is compelling since it means that TQO, normally a ground-state property, has implications on the excited states as well.

Lastly, it will be interesting to look at the combined effect of disorder and interactions. An open question still is if disorder has consequences for the time scale of the phase error. In this respect, possibly a connection with the works [197, 194, 198] on prethermalisation can be made.

Chapter 6

Optimising Magnon Transport in Disordered Spin Chains

An important aspect of quantum information processing and the quest for practical quantum devices is the ability to transfer quantum states over large distances. For example, in the field of quantum communication one is interested in using quantum channels to transmit information from one sender to a receiver [71, 72]. Moreover, for quantum computing applications one may need to transport qubit states in networks of quantum computers or between independent qubits via some form of quantum bus, see e.g. [199]. For such transport of quantum information over long distances a very well studied and implemented system is the optical fibre [200, 201]. In these fibres photons carry the quantum information. However, for transport over shorter distances, e.g. in quantum networks, other schemes that can be combined with solid state architectures are more desirable.

In 2003 Bose [202] proposed to make use of spin chains with short range Heisenberg interactions to transmit quantum information. In this proposal, the information is encoded in a single quantum spin at one edge of the system (open boundaries) and then transported to the final spin at the other edge, see e.g. Fig. 6.1. The main advantage of using such a spin chain is that a wide array of (potential) experimental platforms are available to realise it in the laboratory. This includes systems consisting of an array of Josephson junctions [65], coupled quantum dots [67], optical lattices [64] and also atomic nuclei [203]. In particular, for cold atom setups it has been recently shown that it is possible to transport spins [66]. Aside from this transport application, Bose also showed that the Heisenberg chain can be used to create distributed entanglement over large distances. This represents another useful resource for quantum communication.

A crucial aspect for all these architectures is the ability to control the encoded quantum information in the spin chain and find optimal protocols for the transport. Various studies have looked into this, starting from the seminal work by Bose himself. An important contribution was made by reference [204], which proposed to use an external magnetic field in the form of a harmonic trap, as shown in Fig. 6.1, and move it adiabatically from the first spin to the last. This has as advantage that the transport is relatively robust against static disorder and




FIGURE 6.1: Setup of the disordered Magnon transport problem. The objective of the agent is to move the harmonic (parabolic) trap B_x in such a way that a localised wave packet (Magnon excitation) is transported from the first spin site to the last spin site (note that in the real simulations we define a different starting and final site). Outside the trap the spins are aligned (blue arrows), whereas inside the trap the spin is anti-aligned (red arrow). The disorder is uniform across the spin chain (yellow wiggly line) and stays constant in time.

also it is not required to have full control over the individual spin couplings, as was studied earlier in reference [205]. However, a drawback is that, as we know from previous chapters, adiabatic protocols tend to be slow and hence susceptible to noise from the environment. Other faster, non-adiabatic, protocols have been studied numerically with the Krotov method in reference [206] and an analytic shortcut to adiabaticity approach in reference [70].

In this chapter, we build upon these studies and aim to look for fast control protocols for transporting a magnon¹ along an isotropic Heisenberg spin chain. For the optimisation we use the numerical optimisation method Differentiable Programming (∂P) that we have introduced and applied before in the context of Majorana transport. An important distinction² is, however, that here we are looking at state transfer and not at the transport of a single quasiparticle mode. What is more, here we are interested in the transmission of information instead of the manipulation of quantum information, which was the aim for the movement of Majoranas. To model the real-world experimental scenario closely we include onsite non-uniform disorder along the chain. We apply ∂P directly in this noisy setup and compare the results to the clean, non-disordered, setup. As before, the main advantage of ∂P is that it allows us to efficiently obtain the required gradients, which makes gradient based optimisation of this quantum many-body system computationally manageable.

¹A localised wave packet (spin excitation) as in Fig. 6.1.

²Aside from the system itself.

With this ∂P method we obtain transport protocols that have high fidelity even in the presence of fixed, unwanted, disorder realisations. As part of this we show that we can efficiently combine ∂P with several other traditional quantum control methods, specifically the shortcut to adiabaticity protocol method (STA) in reference [70] and also the CRAB fourier protocol ansatz described in Chapter 2. With this combined approach we are able to improve the heuristic speed limit for the clean system found in reference [70], which did not make use of additional numerical optimisation. Lastly, upon inspection of the optimal transport strategies in the clean system, we find a connection between the frequency of the applied harmonic trap and the frequencies of the protocols. This indicates a potentially deeper physical origin for their high fidelity³.

We have structured this chapter as follows. In section 6.1, we introduce the Heisenberg spin chain model and define the optimisation objective. Then, in section 6.2, we will look at the two different protocol ansätze obtained with the traditional methods, which we aim to optimise further, i.e. combine, with ∂P . In section 6.3, we discuss the results for the optimisation in the clean spin chain and show that we can improve the heuristic speed limit. Afterwards in section 6.4, we look at the effects and optimisation of disorder before concluding in section 6.5.

6.1 The Spin Model and Optimisation Objective

The model we consider is the one-dimensional Heisenberg model [63] of N interacting spin-1/2 particles with the Hamiltonian given by

$$\mathcal{H} = -\frac{J}{2} \sum_{i=1}^{N-1} \vec{\sigma}_i \cdot \vec{\sigma}_{i+1} + \sum_{i=1}^N B_i \sigma_i^z, \quad (6.1)$$

in which $\vec{\sigma}_i \equiv (\sigma_i^x, \sigma_i^y, \sigma_i^z)$ are the usual Pauli matrices and we have used open boundary conditions. We have an isotropic ($J \equiv J^x = J^y = J^z$) nearest-neighbour coupling⁴ with constant J . Note that, in other literature, this model is sometimes also referred to as the XXX -model in a local external magnetic field B_i . After a Jordan-Wigner transformation, see App. B, this model maps to the Hubbard model [209, 39] for interacting fermions hopping on a one-dimensional lattice.

Diagonalising this type⁵ of Hamiltonian and looking for the eigenstates and eigenenergies is generally a hard task since the model cannot be mapped to a free fermion model. To find the full spectrum and set of eigenvectors of the 1D Heisenberg chain one therefore needs to make use of techniques like the Bethe Ansatz [190] or exact diagonalisation. However, in the

³Some of the results in this chapter are, after the submission of this thesis, published in Ref. [207]

⁴This describes (magnetic) exchange interactions [208] between the neighbouring spins, see for a review [85]

⁵Heisenberg-like Hamiltonians including the XXX , XYZ , etc. models.

following, we will see that for our single spin transport problem we do not need access to the full many-body Hilbert space and that we can restrict to a specific subspace. Moreover, by simple inspection of the terms in the Hamiltonian Eq. 6.1 one can already find, without full diagonalisation, that in the absence of a magnetic field the ground state for $J > 0$ is ferromagnetic, i.e. all the spins are aligned⁶. This is a direct consequence of the exchange-interaction that aligns the spins.

In order to show that we do not need to work in the full many-body space, we first observe that the total spin excitation number, $\hat{S} = \sum_{i=1}^N \sigma_i^z$, is preserved by the Heisenberg Hamiltonian, since we have

$$[\mathcal{H}, \hat{S}] = 0. \quad (6.2)$$

This means that we can divide up the problem into separate (non-connecting) sectors with each a fixed total number of spin excitations $S = \langle \hat{S} \rangle$. The simplest sector is the single spin-excitation sector in which in total one spin is up, denoted by $|\uparrow\rangle$, whereas all the other spins are down $|\downarrow\rangle$. In total we have N of such 1-spin excitation states which we label by $|n\rangle = |\downarrow\rangle_1 \otimes \cdots \otimes |\uparrow\rangle_n \otimes \cdots \otimes |\downarrow\rangle_N$ in which only the n -th spin is up.

The Hamiltonian Eq. (6.1) projected into this subspace then becomes

$$H = \sum_{n=1}^N [B_n - 2J] |n\rangle \langle n| + J |1\rangle \langle 1| + J |N\rangle \langle N| \quad (6.3)$$

$$+ J \sum_{n=1}^{N-1} |n\rangle \langle n+1| + |n+1\rangle \langle n|. \quad (6.4)$$

Now since for our optimisation problem we are interested in transporting a single spin and as the Hamiltonian preserves the total spin, for our simulations it is sufficient to only consider the single spin excitation subspace Hamiltonian H . This can be efficiently implemented as a $N \times N$ matrix on a computer, which is an exponential reduction in memory compared to the $2^N \times 2^N$ many-body Hamiltonian of the full system.

We note that the single excitation Hamiltonian Eq. (6.4) is similar to a standard tight-binding Hamiltonian with hopping $t \sim J$ and chemical potential $\mu_i \sim B_i - 2J$ (in the bulk). To be precise, up to some edge effects⁷ this model can be mapped to a discretised version of a single particle in a potential trap. In the continuum limit this has a Hamiltonian

$$H_c = \frac{k^2}{2m} + V(x), \quad (6.5)$$

in which we associate the potential $V(x)$ with the local (discrete) magnetic field B_i and the

⁶This means it is twofold degenerate, i.e. all spins up or all spins down.

⁷The two diagonal terms proportional to J in Eq. (6.4).

inverse mass in the kinetic energy term $1/2m$ with the coupling J . This approximate correspondence is useful for solving analytical STA control protocols, as we will do in section 6.2. Importantly, the physics in both scenarios is different. In one case we look at a localised spin excitation in a discrete system, and in the other case at a single particle in a continuous trapping potential.

6.1.1 Magnons and the Optimisation Objective

For our optimisation problem we aim to transport magnons from one position in the chain to another, see Fig. 6.1. Magnons are spin wave excitations travelling through the chain, which can be written as a (generic) series expansion

$$|\psi_k\rangle = \sum_n c_{n,k} |n\rangle \quad (6.6)$$

of the single spin excitation states $|n\rangle$. When $c_{n,k} = e^{ink}$, we have plane (Fourier) waves and the magnon excitations $|\psi_k\rangle$ are the eigenstates of the single excitation Hamiltonian $H = \sum_k E_k |\psi_k\rangle \langle \psi_k|$, for $B_n = 0$ and closed boundary conditions. The energy dispersion for these modes is then given by

$$E_k = -2J(1 - \cos k), \quad (6.7)$$

where k is the momentum. This means that the group velocity of the spin waves is given by

$$v_g \equiv \frac{dE_k}{dk} = -2J \sin k. \quad (6.8)$$

The group velocity v_g limits how fast the magnons can maximally travel and, hence, puts a lower bound on the heuristic speed limit of our control problem.

Freely propagating magnons, without a confining potential or magnetic field B_i , disperse over time throughout the chain. This effect of propagating magnons has been investigated in detail in reference [210] and is a direct result of the nonlinear dispersion relation Eq. (6.7). In order to counteract this spreading it is necessary to guide the magnon transport [211, 204]. This can be done by imposing the external magnetic field $B_n(t)$ to have the form of a trap around the wave packet and then slowly change the position of the trapping minimum over time. Various kinds of magnetic traps with different spatial profiles can be used for this, such as the Pöschl-Teller [212] potential, a square well [210] potential, or a harmonic trap [204].

For our control problem we focus on a (harmonic) parabolic trap, which has the form

$$B_n(t) = -\frac{\omega(t)^2}{4J} \left[\frac{n - X_0(t)}{\Delta x} \right]^2 |n\rangle \langle n|, \quad (6.9)$$

where $\omega(t)$ and $X_0(t)$ are the now time dependent trapping frequency and minimum of the trap. These variables, $\omega(t)$ and $X_0(t)$, are our control parameters. We set the lattice spacing, Δx , to $\Delta x = 1$ in what follows. The motivation for this specific form of trap stems from the correspondence with the single particle Hamiltonian Eq. (6.4) for which the potential becomes a harmonic oscillator (well) potential $V(x) = \frac{1}{2}m\omega(t)^2 [x - X_0(t)]$. This means that we can exploit known analytic optimal control pulses derived with STA methods for the harmonic oscillator, see e.g. [213, 70] and section 6.2. We remark that, due to this specific nonlinear encoding of the control, we do not expect the optimal protocols to be of the bang-bang form studied in section 2.2.

The magnetic trap $B_n(t)$ has the form of a parabola and can be used to control where the magnon excitation is localised by tuning the minimum $X_0(t)$. Moreover, the trapping frequency $\omega(t)$ can be used to control the spatial extent of the wave packet, which becomes very narrow for very large values of ω . For our magnon transport problem we initialise these parameters (at $t = 0$) such that we have a nicely confined (Gaussian) wave packet

$$|\psi_A\rangle = \frac{1}{\sqrt{\sum_n c_n^2(x_A)}} \sum_n c_n(x_A) |n\rangle \quad (6.10)$$

at a position x_A in the chain. Here, we set $c_n(x) = \exp\{-(n-x)^2/2\sigma\}$, which means the magnon has the shape of a Gaussian. We fix $\omega(t) \equiv \omega_0$ to be constant in time and aim to change $X_0(t)$ in such a way that we reach the target state

$$|\psi_B\rangle = \frac{1}{\sqrt{\sum_n c_n^2(x_B)}} \sum_n c_n(x_B) |n\rangle \quad (6.11)$$

after a total time τ . This target state has the same magnon profile but now localised at a position $x_B = x_A + d$, where d is the transport distance. Importantly, we choose x_A and x_B far enough from the boundaries to avoid any finite system size effects.

The quality of a generic magnon transport protocol parameterised by $X_0(t)$ can be measured by the standard infidelity measure that we saw before (section 2.1.1). However, this time we define the starting and target states to be $|\psi_0\rangle \equiv |\psi_A\rangle$ [Eq. (6.10)] and $|\psi_\tau\rangle \equiv |\psi_B\rangle$ [Eq. (6.11)], which gives

$$\mathcal{I}_\tau = 1 - |\langle \psi_B | U(\tau) | \psi_A \rangle|^2. \quad (6.12)$$

Here, $U(\tau) = \mathcal{T} \exp\{-i \int_0^\tau H(t) dt\}$ is the time-ordered unitary that solves the time-dependent Schrödinger equation with the single spin excitation Hamiltonian $H(t)$. This infidelity \mathcal{I}_τ forms the objective function, which we minimise with respect to $X_0(t)$ to obtain the optimal magnon control pulses. We note that this objective function slightly changes if we include disorder⁸ in the chain, as we will discuss in section 6.4.

⁸It becomes an average of the infidelity values of a batch of disorder realisations.

As a final remark, we note that this optimisation problem for transporting localised wave packets with an infidelity objective function feels very similar to the transport of localised Majorana modes in Chapter 4. In fact, on a mathematical level, and from an abstract optimisation perspective, there are no significant differences and the same techniques can be used. However, the underlying physical systems and control goals are different. For example, locality in the spin-picture does not directly transfer to locality in the fermion picture for Majoranas. Moreover, in one case we have a gapped topological system whereas in the other we have a continuous spectrum.⁹ This means that we do not expect the same type of protocols or intuition gained from the Majorana problem, i.e. the superadiabatic and Jump-move-Jump protocols, to apply to the magnon transport as well.

6.2 Ansätze for the Optimal Protocols

To search for the optimal magnon transport protocols $X_0^{\text{opt}}(t)$ that minimise the infidelity Eq. (6.12) we use a few different optimisation approaches and protocol ansätze. Specifically, we compare the performance of protocols obtained with two hybrid ∂P - traditional quantum control approaches to protocols obtained with a standard ∂P time-bin optimisation approach and a naive linear benchmark protocol. In this section, we will first briefly introduce the latter two approaches and then discuss the hybrid ∂P methods in detail. In the following sections we will then show the optimisation results.

The linear benchmark protocol is given by

$$X_0(t) = x_A + \frac{x_B - x_A}{\tau}t \quad (6.13)$$

and was first studied in reference [204]. The advantage of this protocol is that it is relatively simple to implement experimentally. However, it has as main disadvantage that it requires long (adiabatic) transport times τ before the infidelity \mathcal{I}_τ becomes zero. This time scale becomes even longer in the presence of disorder. This means that, in general, this approach is not useful for fast and robust magnon transport. As such, we use the linear protocol simply to benchmark the increase in performance of our optimised control protocols.

The second approach is the standard ∂P time-bin optimisation method that we have used before in section 2.3.2 for the single-qubit control problem and in chapter 4 for the optimal Majorana transport. Here, we discretise the control into M individual time bins $X_0(t) \mapsto [X_0(t_1), X_0(t_2) \cdots X_0(t_M)]$ of time width $\Delta t = \tau/M$ and we define $t_n = n\Delta t$. The optimisation task becomes finding the values for each $X_0(t_n)$ that minimise \mathcal{I}_τ . The advantage of this method is that we do not *a priori* put any restrictions on the form of the protocols $X_0(t)$. However, this means that the size of the search space scales with the number of (discrete)

⁹Which is discretised on the lattice.

time steps in the evolution and that the optimisation problem can become computationally expensive¹⁰.

This scaling is one of the motivations for using the two hybrid $\partial\mathcal{P}$ -traditional quantum control methods. In these hybrid approaches we first specify a parameterised ansatz for the form of the optimal protocols and then optimise the parameters of the ansatz with $\partial\mathcal{P}$. The hybrid nature comes about because we use the traditional quantum control methods STA and CRAB to provide us with two specific ansätze. We will now describe (and derive for STA) these ansätze and discuss how to optimise them with $\partial\mathcal{P}$.

6.2.1 Shortcut to Adiabaticity Inspired Protocol Ansatz

For the STA protocol ansatz we derive optimal control protocols for a particle in a harmonic trap [see Eq. (6.5)] with the inverse engineering method based on Lewis-Riesenfeld (LR) invariants described in section 2.2. The approximate correspondence with the single spin excitation Hamiltonian H [Eq. (6.5)] means that these protocols are also expected to work for the magnon transport in a Heisenberg chain, as was shown in reference [70]. Our derivation follows [70], which also uses some of the results from the works [89, 90, 214, 213] for generic (quantum) harmonic oscillators.

To apply the inverse engineering STA method we recall from section 2.2 that we need to find a suitable dynamical invariant $I(t)$. For the single particle in a harmonic trap with Hamiltonian¹¹

$$H_c(t) = \frac{k^2}{2m} + \frac{1}{2}m\omega(t)^2 [x - X_0(t)]^2 \quad (6.14)$$

a quadratic in momentum LR invariant [89, 213] is given by

$$I(t) = \frac{1}{2m} [\eta(k - m\dot{\alpha}) - m\dot{\eta}(x - \alpha)]^2 + \frac{1}{2}m\omega_0^2 \left(\frac{x - \alpha}{\eta} \right)^2, \quad (6.15)$$

where $\omega(0) = \omega_0$. Here, the time dependent functions $\eta \equiv \eta(t)$ and $\alpha \equiv \alpha(t)$ need to satisfy the equations

$$\ddot{\eta} + \omega(t)^2\eta = \frac{\omega_0^2}{\eta^3}, \quad (6.16)$$

$$\ddot{\alpha} + \omega(t)^2(\alpha - X_0(t)) = 0, \quad (6.17)$$

so that $\frac{\partial I(t)}{\partial t} + i[H_c(t), I(t)] = 0$. These equations (6.16), (6.17) put constraints on the control functions $X_0(t)$ and $\omega(t)$, such that $I(t)$ is an exact dynamical invariant.

¹⁰This is most evident for non-gradient based optimisation methods such as simulated annealing but also applies (to a lesser extent) to gradient based optimisation methods.

¹¹Note that this is just Eq. (6.5) with our specific parabolic trapping potential, (Eq. 6.9), substituted in.

We recognise that $I(t)$ has the form of a harmonic oscillator, which means its eigendecomposition $I(t) |\phi_n(t)\rangle = \lambda_n |\phi_n(t)\rangle$ can be written in terms of harmonic oscillator eigenfunctions and eigenvalues. The eigenvalues of $I(t)$ are therefore given by $\lambda_n = \omega_0(n + 1/2)$ and the eigenmodes (in real space) can be found [214] to be

$$\phi_n(x, t) = e^{im[\dot{\eta}x^2/2\eta + (\dot{\alpha}\eta - \alpha\dot{\eta})/\eta]} \eta^{-1/2} f_n\left(\frac{x - \alpha}{\eta}\right), \quad (6.18)$$

in which the $f_n(x)$ are the standard quantum harmonic oscillator eigenfunctions. Any particular solution of the Schrödinger equation for $H_c(t)$ can be expanded in terms of these eigenstates via $|\psi(t)\rangle = \sum_n a_n e^{i\gamma_n(t)} |\phi_n(t)\rangle$ [recall Eq. (2.21)]. Here, the LR-phases have the form [89, 214]

$$\gamma_n(t) = - \int_0^t dt' \left[\frac{\lambda_n}{\eta^2} + \frac{m(\dot{\alpha}\eta - \alpha\dot{\eta})^2}{2\eta} \right]. \quad (6.19)$$

In order to find optimal transport protocols for the harmonic trap we need to ensure that the initial and final (at $t = \tau$) eigenstates of $I(t)$ match with the initial and target states, $|\psi_A\rangle$ and $|\psi_B\rangle$, of our control problem.¹² This can be achieved by imposing $[I(0), H(0)] = [I(\tau), H(\tau)] = 0$, which results in the following additional boundary conditions

$$\begin{aligned} \alpha(0) = x_A, \quad \alpha(\tau) = x_B, \quad \left. \frac{d^n \alpha}{dt^n} \right|_{t=0, \tau} = 0, \\ \eta(0) = 1, \quad \eta(\tau) = \sqrt{\frac{\omega_0}{\omega(\tau)}}, \quad \left. \frac{d^n \eta}{dt^n} \right|_{t=0, \tau} = 0 \end{aligned} \quad (6.20)$$

for $n = 1, 2$. The final shortcut protocol solutions for $X_0(t)$ and $\omega(t)$ are then obtained by solving the invariant equations (6.16), (6.17) with arbitrary functions $\alpha(t)$ and $\eta(t)$ that satisfy the boundary conditions Eq. (6.20). Since we are only interested in $X_0(t)$ and set $\omega(t) = \omega_0$ we only need to use $\alpha(t)$ in the following.

In practice we can pick any function $\alpha(t)$, and $\eta(t)$, but we restrict to generic polynomials $\alpha(t) = \sum_{n=1}^7 b_n t^n$ that satisfy the boundary conditions Eq. (6.20). This gives a whole family of STA protocols for $X_0(t)$ in which each member corresponds to a specific configuration of the expansion coefficients b_n . In the case of the single particle in the harmonic trap each of these protocols will yield perfect target state fidelity. However, since the correspondence with the magnon transport is not exact (discretisation and finite size), and also because of potential anharmonicities, we do not expect all of these protocols to work equally well for the magnon transport.¹³ Furthermore, real-world experimental constraints make that possibly

¹²In this derivation $|\psi_A\rangle$ and $|\psi_B\rangle$ are the eigenstates of H_c in Eq. (6.14) with the harmonic trap at positions $x(0) = x_A$ and $x(\tau) = x_B$. By the correspondence with the single spin excitation Hamiltonian H we can analogously think of them as the initial and target states of the magnon excitation in the Heisenberg chain.

¹³We will show below that in fact the set of STA protocols with perfect fidelity for the magnon transport grows with the operation time τ .

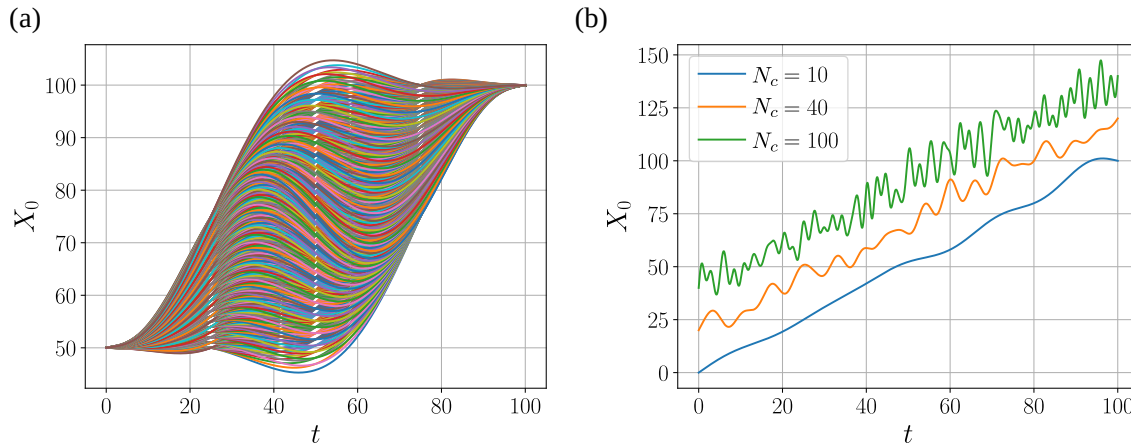


FIGURE 6.2: (a) Set of example protocols that fulfil the boundary conditions of the Lewis-Riesenfeld shortcut to adiabaticity approach in Eqs. (6.17) and (6.20). These protocols were used for the simulations in Fig. 6.3. The parameters for these protocols are $\omega_0 = 0.5$, $x_A = 50$, $x_B = 100$, $C_1 \in [50, 75]$ and $C_2 \in [75, 100]$. (b) Fourier Ansatz example protocols for three different total number N_c of frequency components $\omega_n = \frac{2\pi n}{\tau}$. The Fourier coefficients A_n were chosen randomly from a uniform distribution $A_n \in [-1.5, 1.5]$ and $d = x_B - x_A = 90$. The higher the frequencies involved, i.e. the bigger N_c , the more rapid fluctuations and details the protocols can have.

some of these protocols are not admissible. This means that additional optimisation within this family of STA protocols is desirable.

To be able to optimise within this family of protocols we need to parameterise it in such a way that the minimisation algorithms can be applied. One possible way is to define the coefficients b_n as the new control parameters, but we have found that it is nontrivial to apply ∂P in this case. This is because for each set b_n the boundary conditions, Eq. (6.20), need to be imposed and the STA equations, (6.17), inverted. In theory this can be done in a fully differentiable way by, for example, using Lagrange multipliers. However, this can become complicated and possibly unstable in practice. Here, instead, we parameterise (a subset) of this family by defining two¹⁴ new control parameters $C_1 = X_0(\tau/4)$ and $C_2 = X_0(3\tau/4)$, which are the positions of the trap at times $\tau/4$ and $3\tau/4$. These can be used as control parameters by imposing them as additional constraints when solving the polynomial for $\alpha(t)$. The advantage of this is that the constraints only need to be imposed once. In addition, the resulting fixed polynomial of C_1 and C_2 can be implemented as a differentiable function for ∂P . The downside is that, since we only have two parameters, we do not cover the full family and only look at a subset.

A sample set of these STA protocols $X_0(t)$ as a function of C_1 and C_2 is shown in Fig. 6.2 (a). It can be seen that all of these have an initially slow acceleration (second derivative), which, to various extents, initially increases and finally slows back down to zero at the final time. Some of the protocols go beyond the initial or target state positions $x_{A/B}$, which

¹⁴For practicality we restrict to two parameters, this approach can however be generalised to 4,6,...,N control parameters without any (significant) additional computational cost for ∂P .

might not be feasible to implement in experiments if they go outside the chain. Moreover, as shown in Fig. 6.3, not all these protocols work equally well, as not all of them have a perfect fidelity \mathcal{F} . In fact, we find that the subset of protocols of this parameterised group that has near perfect fidelity shrinks with the transport time τ (shrinking white-yellow island in Fig. 6.3). This observation is in agreement with the results in reference [70], which found for one specific unparameterised STA protocol a heuristic speed limit of $\tau = \tau^* \approx J$ below which the magnon cannot be perfectly transported. With our hybrid ∂P approach we can, however, optimise the parameters C_1 and C_2 with gradient descent and find which member of the (sub)-family of STA protocols performs the best.




FIGURE 6.3: Fidelity surfaces (for three different transport times τ) of the family of shortcut to adiabaticity (STA) protocols parameterised by C_1 and C_2 and shown in Fig. 6.2 (a). As C_1 and C_2 are the control parameters these surfaces are also the control landscapes. The landscapes look convex and the set of STA protocols with near perfect target fidelity (coloured island) shrinks with the total transport time.

6.2.2 Fourier Series Ansatz

The final ansatz for the optimal protocols that we use and aim to optimise with ∂P is based on a Fourier basis and inspired by the CRAB method [96, 97] discussed in section 2.2.3. Specifically, to fulfil the boundary conditions $X_0(0) = x_A$ and $X_0(\tau) = x_B$ we define the Fourier ansatz to be the family of protocols parameterised as

$$X_0(t) = x_A + (x_B - x_A) \frac{t}{\tau} + \sum_{n=1}^{N_c} A_n \sin \omega_n t. \quad (6.21)$$

The control parameters we wish to optimise are now the Fourier coefficients A_n , as we fix the frequencies $\omega_n = \frac{n\pi}{\tau}$ to be the first $n < N_c$ harmonics. A few examples of these protocols are shown in Fig. 6.2 (b).

The set of protocols defined by Eq. (6.21) is again a constrained set of protocols if we fix N_c to be finite, which means we are not optimising over the full space of available protocols. This could mean that outside this set there are potentially better protocols with lower infidelity values. However, the main advantages of this ansatz are that the boundary conditions are automatically satisfied and also that we can put a cutoff on the maximum frequency ω_{N_c} . This means we can choose a specific N_c based on knowledge about the system or by considering the maximal rates of change that can be reached in real experimental setups. To go back to the Majorana transport optimisation problem, we could have for example restricted to frequencies lower than the resonance frequency (superadiabatic frequencies) if we had used this Fourier ansatz.

In order to find the optimal protocols in this restricted family we use $\partial\mathcal{P}$ to compute the derivatives $\frac{d\mathcal{I}}{dA_n}$. These derivatives can then be employed in standard gradient based optimisation algorithms such as vanilla gradient descent. This means that we can scale to thousands of fourier coefficients A_n without much additional computational cost, whereas previous studies [215, 36] that used CRAB in combination with non-gradient methods were restricted to $\mathcal{O}(10^1)$ optimisation parameters.

6.3 Optimisation Results for the Clean Spin Chain

Our aim is to minimise the infidelity of the magnon transport \mathcal{I}_τ in Eq. (6.12) and find the optimal control protocols for the trapping centre $X_0(t)$. For this task we apply $\partial\mathcal{P}$ in combination with the various ansätze and protocols described in the last section. To define the external constraints of the control problem we fix the total transport distance $d = 50$ and choose four different total transport times $\tau = 40, 60, 80, 100$. For the specific spin chain parameters that we consider we have a heuristic speed limit velocity $v_{\tau^*} = d/\tau^* \approx J = 1$ following [70]. However, we will show that we can improve this with the $\partial\mathcal{P}$ method and bring the transport velocity closer to the maximum magnon group velocity of $v_g = 2J = 2$. Finally, in order to obtain the optimal protocols we minimise with the vanilla gradient descent algorithm (section 2.3.2) for a maximum of 200 update steps¹⁵ with an empirically determined learning rate.¹⁶ The use of this (basic) gradient descent algorithm is motivated by the shape of the control landscapes for the STA protocols in Fig. 6.3, which look convex.

The results for these optimisations are given in Fig. 6.4 and the corresponding infidelity values are reported in Table 6.1. Firstly, we note that above the speed limit time $\tau^* = 50$ we are able to obtain protocols that have infidelities between the orders of $\mathcal{O}(10^{-3})$ to 10^{-4} with all three optimisation methods. In contrast, the linear reference protocol has significantly

¹⁵The minimisation is stopped when an infidelity lower than $\mathcal{O}(10^{-3})$ is reached.

¹⁶Note that for the various different ansätze we use different derivatives such as $\frac{\partial\mathcal{I}_\tau}{\partial C_i}$ for the STA protocol ansatz and $\frac{\partial\mathcal{I}_\tau}{\partial A_n}$ for the Fourier ansatz.

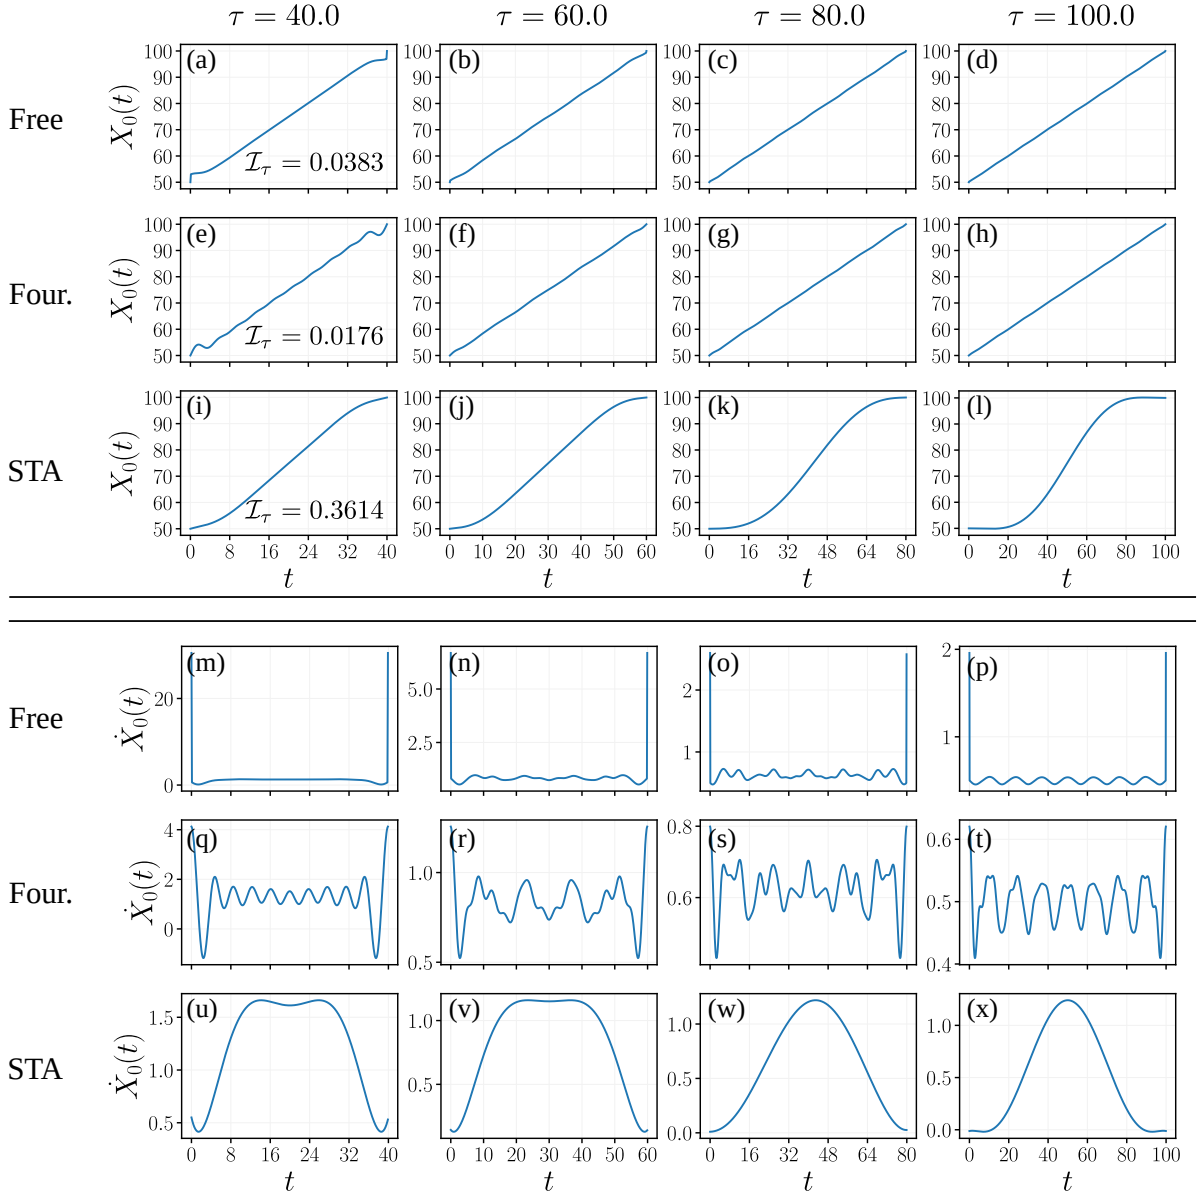


FIGURE 6.4: Clean (disorder-free) optimisation results obtained with ∂P in combination with no constraints on the protocols (rows labelled 'Free'), a Fourier ansatz with $N_c = \text{int}[\tau/2]$ (rows labelled 'Four.') and the STA protocol ansatz parameterised by C_1 and C_2 (rows labelled 'STA'). In panels (a-l) we plot the position of the trapping minimum $X_0(t)$ as a function of time and in panels (m-x) the corresponding velocity protocols $\dot{X}_0(t)$. The simulations were done for 4 different total times τ (columns) of which the first column is below the heuristic speed limit time $\tau^* = 50$ and the others above. The infidelity values below the speed limit are reported as a label in the subplots (a, e, i) and the others are reported in Table. 6.1 but are 'perfect' $\mathcal{O}(10^{-3} - 10^{-4})$. We obtain different protocols for the three different methods (rows) and note that the STA method is not able to reach a low infidelity below the speed limit. The system parameters are chosen to be $N = 251$, $d = x_B - x_A = 50$, $\omega_0 = 0.5$, and $J = 1.0$.

bigger infidelity values, $\mathcal{O}(10^{-1})$. Below the speed limit the ∂P Free and Fourier approaches are still able to achieve low infidelities $\mathcal{O}(10^{-2})$, whereas ∂P in combination with STA does not perform as well but still better than the linear protocol. This is likely due to the fact that we

have a restricted search space and that the speed limit is a hard limit for the STA protocols. In addition, it might be that increasing the number of control parameters C_i of the STA protocols will lead to a better performance of the ∂P STA optimisation combination below the speed limit.

τ	∂P Free	∂P STA	∂P Fourier	Linear Ramp (Eq. 6.13)
40	0.03832	0.36139	0.01755	0.98450
60	0.00146	0.000126	0.00516	0.67802
80	0.00086	0.000229	0.00034	0.07908
100	0.00241	0.00023	0.00173	0.44136

TABLE 6.1: Infidelity \mathcal{I}_τ results for optimisation in the clean system $N = 251$, $J = 1$, ω_0 and $d = 50$ and the linear benchmark protocol. The corresponding control protocols are shown in Fig. 6.4.

The shape of the optimal protocols $X_0(t)$ themselves, see Fig. 6.4, depends, not surprisingly, on the different methods and also if we are above or below the speed limit. We observe that the protocols for the Free and Fourier methods are more or less similar above the speed limit. These protocols start and end with a small quench in position and in the middle oscillate with an on average constant velocity (see panels m-f). The size of the quenches in position grows with decreasing transport time, while the oscillations tend to become smoother with increasing time (most notably for ∂P Free in panel p). The STA protocols are drastically different and all look like a form of ramp-up and down protocol in velocity. These protocols slowly accelerate (ramp) to a finite velocity and then symmetrically decelerate again to zero velocity (panels u-x). This behaviour is reminiscent of the superadiabatic protocols that we used for the Majorana transport in Chapter 4.

In order to analyse the frequencies ω_p of the oscillations of the ∂P Free protocols further we show in Fig. 6.5 (a) velocity protocols $\dot{X}_0(t)$ for different trapping frequencies ω_0 . In the main panel we observe that the protocol frequency ω_p tends to increase with increasing ω_0 . From the inset we can then see that the relationship is linear with a coefficient very close to one, i.e. $\omega_p \approx \omega_0$. This indicates that there is a connection between the setup parameter ω_0 and the optimal protocols obtained with ∂P Free. Moreover, it provides a good starting point for the optimisation of systems with a different ω_0 . For this potentially a parameterised protocol can be extrapolated. We leave the investigation of this and a possible deeper physical explanation for future studies.

Another feature of the optimal protocols that we briefly report about is the maximum velocity v_{\max} . Specifically for the STA protocols, since they have the shape of a superadiabatic protocol. We determine v_{\max} for a range of different protocols obtained after grid scans of the parameters c_1 and c_2 . These results are plotted in Fig. 6.5 (b) in which we observe that at the heuristic speed limit time τ^* we get a maximum velocity $v^* \approx 1.38 \approx \sqrt{2J}$.¹⁷ For shorter

¹⁷We note that this value did not seem to depend on d .

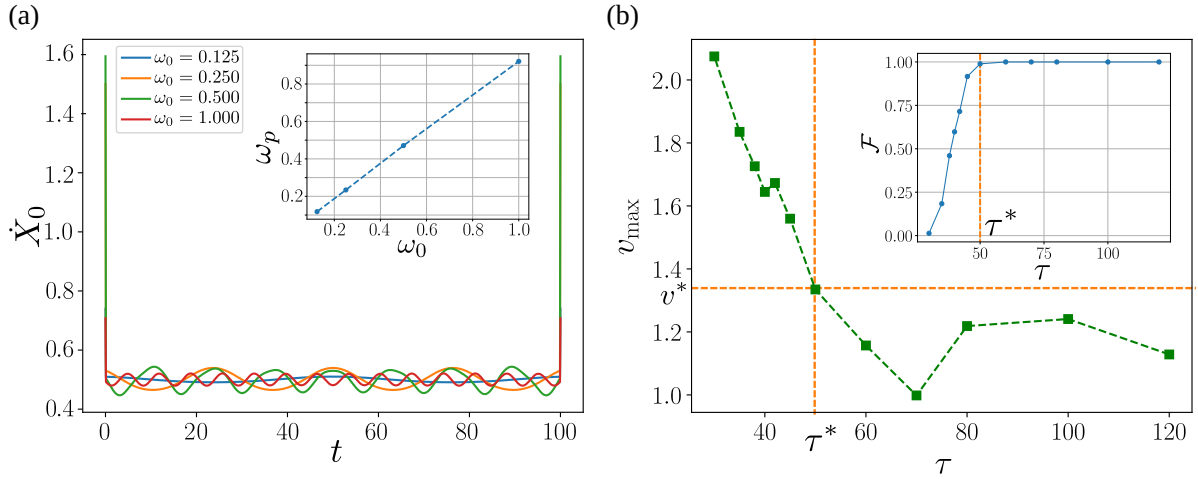


FIGURE 6.5: (a) Main panel: velocities of the Magnon transport protocols $\dot{X}(t)$ obtained with the ∂P Free approach for 4 different trapping frequencies ω_0 . Inset: the characteristic frequency ω_p of the protocols in the main panel versus ω_0 . A linear relationship $\omega_p \approx \omega_0$ is observed. For these simulations $\tau = 100$, $N = 251$, $J = 1$ and $d = 50$. (b) Main panel: maximum velocity of the STA protocols for various different total transport times τ . Inset: the corresponding target state fidelity \mathcal{F}_τ of these protocols. At the critical (speed limit) time τ^* the maximum velocity of the magnon is v^* . For shorter times, $\tau < \tau^*$, the fidelity starts to decrease. The parameters for these simulations are the same as in Fig. 6.4.

times we get higher maximum velocities but the fidelity starts decreasing (inset). This v^* is the maximum velocity of the trap and hence the magnon for the STA protocols.¹⁸ Although the determined v^* is slightly higher than the heuristic speed limit on the total transport velocity v_{τ^*} , it is still lower than the maximum group velocity v_g . We note that this is just a feature of the STA protocols and, in fact, in the next subsection we will show that it is possible to get closer to v_g with ∂P in combination with the Fourier ansatz.

6.3.1 Tightening the Heuristic Speed Limit with ∂P

To show that we can tighten the heuristic speed limit time $\tau^* \approx d/J$ we first show how it can be determined from the original (unparameterised) STA protocol for magnon transport derived in reference [70]. This protocol is just one specific example of our parameterised family of STA protocols derived in section 6.2.1 and is given by

$$X_0(t) = x_A + d \left[6t^5 - 15t^4 + 10t^3 + \frac{60t}{\omega_0^2 \tau^2} (1 - 3t + 2t^2) \right]. \quad (6.22)$$

This protocol is a function of the total transport time τ , the trapping frequency ω_0 and the initial and target positions x_A and x_B . This means that it can be determined (extrapolated) for a range of different values of τ and d .

¹⁸Since the system is not topological or gapped however we cannot directly make the analogy with the critical velocity of the superadiabatic Majorana motion. Nevertheless, it is still possible that the STA protocols are following a superadiabatic path and the magnon is actually in the ground state of a moving frame Hamiltonian.

In Fig. 6.6 (a) we show the target state Fidelity \mathcal{F}_τ of this protocol for a specific range of values of d and τ . We observe a clear lightcone-like surface with a velocity of $v_{\tau^*} \approx d/J$. Outside this lightcone a fidelity of at least $\mathcal{F}_\tau > 0.5$ can be obtained, whereas inside it the fidelity drops quickly to zero. The speed v_{τ^*} now defines the heuristic speed limit for this protocol. We see that this speed limit is still far away from the lightcone of the group velocity $v_g = 2J$, and even further from the Lieb-Robinson velocity $v_l = 6J$ lightcone found in [216]. The Lieb-Robinson speed defines (mathematically) the fastest speed of information spreading in the spin chain.

In order to show that we can improve this speed limit with $\partial\mathcal{P}$ we focus on one specific distance slice $d = 50$ and minimise \mathcal{I}_τ for a range of different total times τ near the original speed limit time $\tau^* = 50$. We use the Fourier ansatz with $N_c = 50$ frequency components. To push the limits of our optimisation methodology we run 500 update steps¹⁹ with the Adam [112] update scheme. The resulting fidelity values as function of τ compared to the original STA protocol are shown in Fig. 6.6 (b). Compellingly, we see that we can push the heuristic speed limit closer to the group velocity time $\tau_g = d/v_g$. When we take the $\mathcal{F}_\tau > 0.5$ measure again, the speed limit is about $\tau_{\partial\mathcal{P}} \approx 29.0$, which approaches $\tau_g = 25.0$. Below τ_g the fidelity quickly drops to zero and no improvement seems to be able to be made with the optimisation. This is in accordance with the group velocity being the maximum speed of the magnons [211]. Moreover, it agrees with the speed limit found in reference [206] with the numerical Krotov method. The authors there used the same spin system but with a different form of spin excitation.²⁰

6.4 The Effects of Disorder

To model real-world experimental setups more closely in this section we look at the effects of disorder in the Heisenberg spin chain [217, 218, 204, 210, 70]. We will first introduce two different types of disorder and investigate its effects on the performance of the STA protocols. Afterwards, we will apply our $\partial\mathcal{P}$ optimisation method in combination with the Fourier ansatz directly in the presence of disorder. We will show that with this method it is still possible to find an optimal control that transports the Magnon with perfect fidelity for a single fixed realisation of the noise. However, for an average over a batch of disorder realisations we find that the (averaged) infidelity for the obtained optimal protocols increases approximately quadratically with the disorder strength. While $\partial\mathcal{P}$ is able to slightly suppress this increase for the simulations we perform, it is not able to reach an averaged infidelity of zero.²¹

¹⁹Again we stop the optimisation when an infidelity of $\mathcal{O}(10^{-5})$ is reached.

²⁰We note that this study incorporated in addition a time dependent trapping frequency $\omega(t)$ which might have gotten the limit closer to v_g than without it.

²¹Note that the results presented here have not yet been published and are still being expanded.

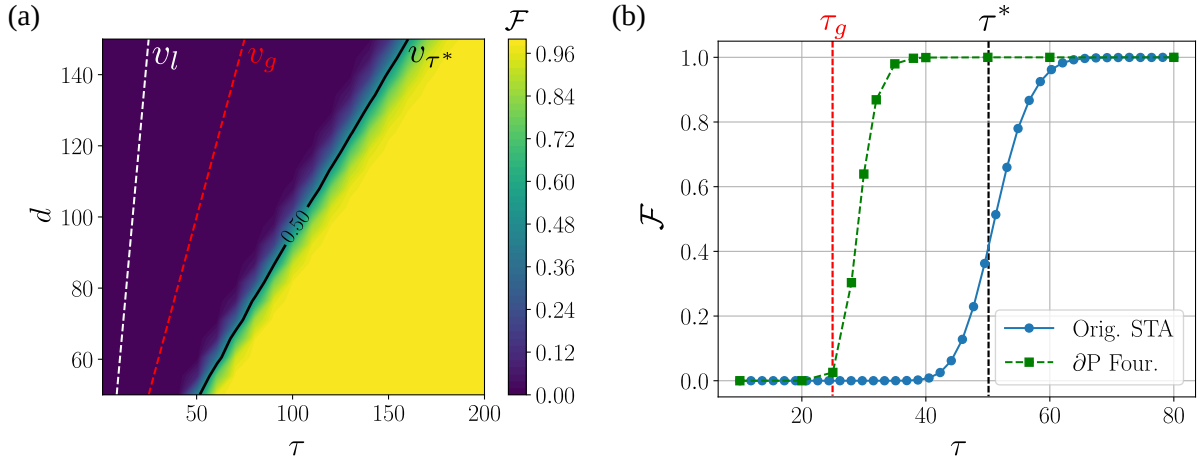


FIGURE 6.6: (a) Target state fidelity surface of the standard non-parameterized STA protocol defined in Eq. 6.22 as function of the transport distance d and transport time τ . The black line gives the $\mathcal{F}_\tau = 0.5$ contour for which the heuristic speed limit $v_{\tau^*} = 1$ is defined, the red dashed line the group velocity $v_g = 2$ and the white dashed line the Lieb-Robinson velocity $v_l = 6$. (b) Target state fidelity as a function of the transport time τ for the original STA protocol (blue solid line) versus the optimised ∂P Fourier protocols (green line). We see that the optimised ∂P Four. protocols keep the fidelity at one for smaller transport times τ than the original STA protocol. The group velocity bound time is given by the red dashed line. The parameters for these simulations were $N = 251$, $\omega_0 = 0.5$, $J = 1$ and in (b) $d = 50$.

6.4.1 Two Types of Disorder and Robustness of the STA Protocols

We consider two different types of static disorder, motivated by potential experimental imperfections. The first is a randomised exchange interaction between the neighbouring spins which can arise from spatial disorder in a realistic spin chain. To model this [204] we change the originally homogeneous spin coupling $J_n \equiv J$ to

$$J_n = J(1 + \epsilon_n) \quad (6.23)$$

in the Hamiltonian in Eq. (6.1). Here, ϵ_n is a random number drawn from a uniform distribution $\epsilon_n \in [-\Delta, \Delta]$ with noise strength Δ . As a consequence of this type of (coupling) disorder, the energy spectrum of the Hamiltonian fluctuates when the potential trap is moved along the chain as shown in [210]. This means that the transport infidelity is expected to increase due to additionally created excitations and dispersion of the magnon.

The second type of disorder that we consider is an inhomogeneous magnetic trapping field

$$B_n(t) \mapsto B_n(t) + \epsilon_n. \quad (6.24)$$

Here, ϵ_n is again a random number from a uniform distribution with disorder strength Δ . Like

²²Note that, while the trapping field is time dependent for the control of the position of the magnon, the disorder itself is static.

for the coupling disorder, this scenario can arise due to imperfections in realistic implementations of the spin chain. This time the noise is local (acting on each spin site individually), however, and we will sometimes refer to this case as "onsite" disorder to distinguish it from the "coupling" disorder.

The main effect of the disordered magnetic field is the localisation of the single spin excitation wave functions. We have shown this in Fig. 6.7 for the lowest energy eigenstate of the disordered Hamiltonian H for several disorder strengths. We observe that the localisation length ξ decreases approximately with the square root of the disorder strength, $\xi \sim 1/\sqrt{\Delta}$. This localisation behaviour can be understood from the approximate mapping to the single particle model, which results for the onsite disordered case in the Anderson model for Anderson localisation [219].²³ The wave function localisation means that the transport infidelity is expected to increase when the magnon is moved for a distance of at least a few disorder length scales, ξ .

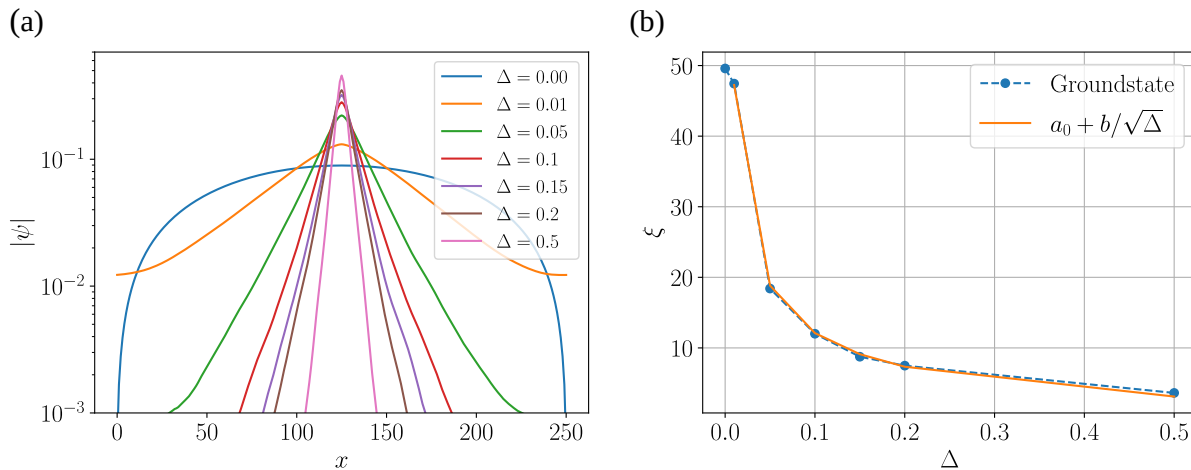


FIGURE 6.7: (a) Disorder averaged amplitude $|\psi|$ of the lowest energy eigenstate of the single spin excitation Heisenberg Hamiltonian in Eq. 6.4 with a disordered magnetic field $B_n = \epsilon_n$. ϵ_n is a random number from a uniform distribution $[-\Delta, \Delta]$. For increasing Δ the amplitude $|\psi|$ becomes spatially more confined (localised) to a smaller number lattice sites x . For these simulations we diagonalised H for in total 1000 disorder realisations and shifted the maximum of $|\psi|$ to the middle of the chain at $x = 125$. The other parameters we used were $N = 251$, $\omega_0 = 0.5$ and $J = 1$. (b) Localisation length ξ of the amplitudes in (a) obtained after fitting a decaying exponential, $a e^{-(x-x_0)/\xi}$, to $|\psi|$. ξ (blue dashed line) falls off approximately as a square root of Δ as shown by the fit (orange solid line).

In order to test the effects of both of these types of disorder on the robustness of the magnon transport we perform numerical simulations of the family of STA protocols in the disordered systems. As a figure of merit for the performance we now look at the disorder averaged fidelity $\langle \mathcal{F}_\tau \rangle = 1 - \langle \mathcal{I}_\tau \rangle$. In Fig. 6.8 we show the disorder averaged fidelity landscapes for both types of disorder with disorder strength $\Delta = 0.1$. The localisation length for the onsite disorder is approximately $\xi \approx 12$ in this case and, hence, we look at relatively strong

²³We note that the coupling disorder model also has localised wave functions. However, this type of disorder corresponds to a disordered hopping (kinetic energy) term after the mapping to the single particle Hamiltonian.

disorder. In these simulations we move the magnon for a total distance of $d = 50$ sites, which means that the magnon is moved for about four disorder length scales ξ . We set the total transport time $\tau = 100$, such that we are far above the (disorder free) heuristic speed limit time, $\tau^* = 50$, where the members of the family of analytic STA protocols [shown in Fig. 6.2 (a)] have almost all perfect fidelity.

For the onsite disorder case, we observe in the surface plots 6.8 (a) and (b) a rich structure with some peaks and valleys.²⁴ There are two areas in the (C_1, C_2) parameter space of the STA protocols that perform the best. These are around $(C_1, C_2) = (58, 78)$ and $(C_1, C_2) = (73, 93)$, for which the average fidelity is about $\langle \mathcal{F}_\tau \rangle \approx 0.89$. This is still a relatively good fidelity given that we have transported the magnon for four localisation lengths. In between these peaks there is a valley and the fidelity drops to about $\langle \mathcal{F}_\tau \rangle \approx 0.82$ near $(C_1, C_2) = (67, 85)$. Although this seems bad, these STA protocols still perform significantly better than a naive linear ramp, which has a fidelity of $\langle \mathcal{F}_\tau \rangle = 0.52$ averaged over 500 disorder realisations. Note that the linear ramp does not have an infidelity close to one because the simulation time, τ , is quite large compared to τ^* .

For the coupling disorder case, shown in Figs. 6.8 (c) and (d), we get a very different fidelity surface. Firstly, the surface is smooth and has only one broad rounded peak (island) with a large fidelity.²⁵ The island is centred at around $(C_1, C_2) = (67, 85)$ ²⁶ and has a maximum averaged fidelity of about $\langle \mathcal{F}_\tau \rangle \approx 0.98$. This is nearly perfect and the STA protocols on this island are robust with respect to disorder. The maximum average fidelity is also better than the maximum of the onsite disorder case.

From an optimisation perspective the difference between the two different disorder surfaces is important. As we discussed in Chapter 2, the convexity properties of the control landscape determine the difficulty of the optimisation problem. The coupling disorder landscape in Fig. 6.8 (c) is convex, since it is smooth and has only one global minimum. The onsite disorder landscape in Fig. 6.8 (a) is non-convex and seems to have two maxima (peaks) with approximately the same size (fidelity). Thus, optimising the STA protocols is potentially harder for the onsite disorder case compared to the coupling disorder case. This does not tell us anything, however, about the complexity or shape of control landscapes in which other ansätze for the protocols are used. It remains open if the coupling disorder problem in general (for any type of ansatz) is convex. Similarly, if the onsite disorder optimisation problem is always non-convex.

²⁴Compare also to the clean landscapes in Fig. 6.3, which have a much smoother structure with only one area (island) where the fidelity is high.

²⁵It looks similar to a bar stool that somebody has carefully scrubbed near the edges with some sandpaper.

²⁶Remarkably, this point in the (C_1, C_2) parameter space is exactly in the low fidelity valley for the onsite disorder case.

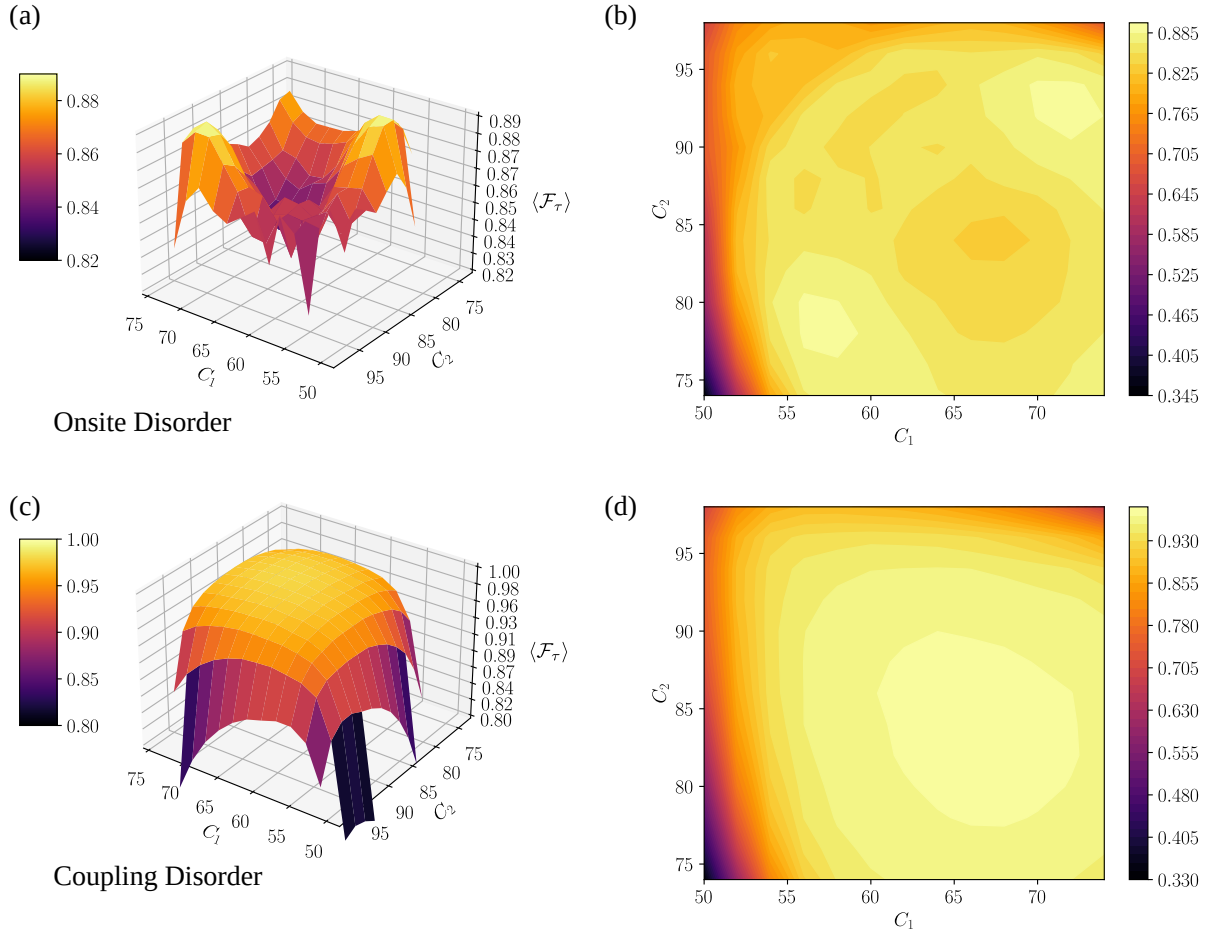


FIGURE 6.8: Disorder averaged fidelity surfaces $\langle \mathcal{F}_\tau \rangle$ for onsite disorder $B_n(t)$ [Eq. (6.24)] shown in panels (a) and (b) and coupling disorder J_n [Eq. (6.23)] shown in panels (c) and (d). Panels (b) and (d) are top views of the 3d surfaces in (a) and (c). The disorder strength for these simulations is $\Delta = 0.1$ and the other parameters are the same as in 6.7. The C_1 and C_2 parameterise the family of STA protocols as in Fig. 6.2 (a). We observe a rough (non-convex) landscape for the onsite disorder case in which $\langle \mathcal{F}_\tau \rangle$ varies significantly with (C_1, C_2) . For the coupling disorder case the surface is smooth and convex with one high fidelity island around $(C_1, C_2) = (67, 85)$.

6.4.2 Optimisation in the Presence of Disorder

In order to see if we can improve the robustness of the magnon transport with respect to disorder, in this section we optimise directly in the presence of disorder. We focus here on the onsite disorder case, which by the non-convexity of the fidelity landscape is possibly a harder optimisation problem. We use ∂P in combination with the Fourier ansatz in Eq. (6.21) and aim to find the optimal protocol(s) in two different scenarios. In the first scenario we focus on one specific disorder realisation and in the second we optimise over a batch of different disorder realisations. We will now discuss the specifics and results of these scenarios one by one.

For the first scenario we assume that the exact disorder pattern of the experimental setup is known. This means that we first pick one particular fixed disorder realisation $\vec{e}^p = (\epsilon_1^p, \dots, \epsilon_N^p)$

(one fixed set of random numbers ϵ_n), which we label by p .²⁷ For this pattern p we then aim to minimise the infidelity \mathcal{I}_τ^p . This infidelity measure is the same as the standard target state infidelity \mathcal{I}_τ from before, but computed with the fixed disorder pattern p added to the magnetic field $B_n(t) + \epsilon_n^p$ in the Hamiltonian. For the minimisation we use \mathcal{I}_τ^p as the cost functional and we compute its derivatives with respect to the control parameters with ∂P . These derivatives are then used in combination with Adam to update the controls A_n of the Fourier ansatz. We repeat the complete optimisation process for a few different disorder patterns p and finally average the result $\langle \mathcal{I}_\tau \rangle_s = \frac{1}{p_{\max}} \sum_p \mathcal{I}_\tau^p$. Here, the sublabel 's' indicates that the infidelity was optimised for each single disorder pattern individually.

In Fig. 6.9 (a) we show the results for these optimisations as a function of the disorder strength Δ . We also show benchmark results of the optimal protocols obtained with the ∂P STA and ∂P Fourier methods in the clean system. We observe that the disorder ∂P optimisation method (green line) is able to keep the infidelity $\langle \mathcal{I}_\tau \rangle_s$ close to zero (on the order of 10^{-3} and smaller). This is remarkable because the infidelity $\langle \mathcal{I}_\tau \rangle_s$ of the benchmark protocols (blue and orange lines) for the same fixed disorder patterns is found to be increasing approximately quadratically with disorder strength. This shows that our ∂P optimisation method is still able to obtain high fidelity control protocols in the presence of single fixed disorder patterns. This means that disorder in the setup is not immediately an issue as long as the exact form of the pattern is known.

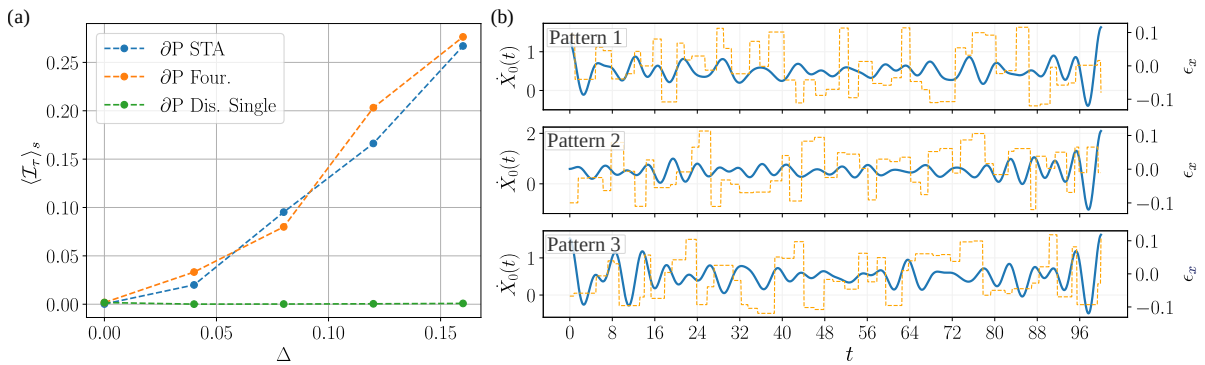


FIGURE 6.9: (a) Disorder averaged infidelity $\langle \mathcal{I}_\tau \rangle_s$ for 5 single disorder realisations versus disorder strength Δ . For the ∂P Dis. Single optimisation method (green dashed line) we optimised the \mathcal{I}_τ^p for each pattern p individually with ∂P before computing the average $\langle \mathcal{I}_\tau \rangle_s$. This optimisation method is able to keep the infidelity near zero while the infidelity of the benchmark protocols obtained for optimisation in the clean system increases quadratically. For these simulations we used Adam with 500 update steps on a Fourier ansatz of 50 frequency components $n = 50$ with external constraints $(\tau, d) = (50, 100)$. All the other parameters are the same as in Fig. 6.7. (b) The obtained velocity protocols (left y-axis) with the $\langle \mathcal{I}_\tau \rangle_s$ method for 3 specific disorder patterns at $\Delta = 0.12$. The disorder value ϵ_x where $x = \text{int}[X_0(t)]$ is plotted in orange with the scale on the right y-axis. For each different pattern we obtain a different optimal transport protocol.

²⁷Note that this is just a N -dimensional random vector from the uniform distribution $[-\Delta, \Delta]$.

To try to understand how the optimisation method is able to achieve these low infidelity values in the presence of disorder we show in Fig. 6.9 (b) the optimal velocity protocols $\dot{X}_0(t)$ for three specific disorder patterns. The obtained optimal protocols $\dot{X}_0(t)$ are different for each different disorder pattern. This means that during the optimisation process the machine learns about the specific disorder realisation and finds a way to correct for it. This is likely due to the fact that the information of a fixed pattern p is directly encoded in the derivatives $\frac{\partial \mathcal{I}_\tau^p}{\partial A_n}$ that the optimiser gets. However, the difference in the strategies also means that there is possibly no universal strategy that obtains a low infidelity value for any disorder realisation. This can be seen because a simple averaging of the strategies will give a different strategy that does not perform well. We leave a more formal mathematical exploration of this for future studies.

The second scenario we consider is the case in which we do not know the exact disorder realisations. In other words, we are interested in figuring out the optimal control strategy that performs on average the best for all possible disorder realisations. For this we take the disorder averaged infidelity $\langle \mathcal{I}_\tau \rangle$ as the figure of merit and minimise it. To minimise it with ∂P we use batch gradient descent. Here, we first fix a set of 200 different disorder realisations and divide it up into 20 individual batches of 10 realisations. For each batch we compute the disorder averaged gradients $\frac{\partial \langle \mathcal{I}_\tau \rangle}{\partial A_n}$ and use them to update the control protocol with Adam. We do this cyclically, which means we start from the first batch do an update and then move on to the second batch. When all batches have had one update we have completed one learning 'episode' and repeat the cycle again. During this process, the value of $\langle \mathcal{I}_\tau \rangle$ averaged over all the 200 disorder patterns is gradually reduced until convergence is reached. We note that due to large number of disorder realisations the optimisation becomes computationally demanding. The gradient evaluations of ∂P can be done in parallel, however, which makes the optimisation task manageable.

In Fig. 6.10 (a) we show the resulting averaged infidelity $\langle \mathcal{I}_\tau \rangle$ values for the optimal protocols obtained with batch gradient descent for 49.5 episodes.²⁹ For these $\langle \mathcal{I}_\tau \rangle$ values we have used a new set of disorder realisations, which were not used during the optimisation process. This gets rid of any optimisation bias to particular disorder realisations and we can fairly evaluate the performance. This time, we see that the optimisation method is not able to keep the infidelity close to zero for increasing disorder strength Δ . Instead, it is increasing quadratically, similarly to the benchmark protocols. Although the batch optimisation method seems to slightly decrease the quadratic increase compared to the other protocols, the difference with the optimal STA protocol is almost negligible. This is remarkable and seems to indicate that the optimisation in the presence of disorder made barely any difference. It is therefore reasonable to ask whether there exist useful strategies in the presence of disorder. To investigate this,

²⁹Note that we first optimised for 24.5 episodes on one fixed set of 200 disorder realisations and then 25 episodes on a different set of 200 disorder realisations.

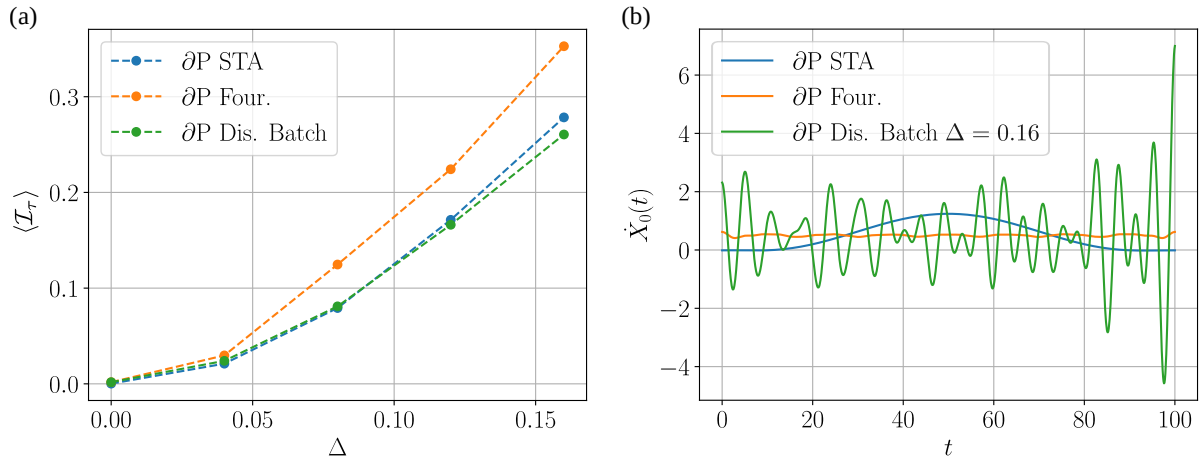


FIGURE 6.10: (a) Disorder averaged infidelity $\langle \mathcal{I}_\tau \rangle$ versus disorder strength for protocols obtained with the ∂P batch optimisation method (green dashed line) compared to the optimal clean protocols.²⁸ The disorder averaging was done over 500 disorder realisations and all the model parameters are the same as in Fig. 6.7. The external constraints are set to $(\tau, d) = (100, 50)$. All the protocols increase approximately quadratically with disorder strength and the batch optimisation slightly outperforms the benchmark protocols. (b) The optimal velocity protocol $\dot{X}_0(t)$ (green line) obtained with the batch optimisation method at $\Delta = 0.16$. We have also shown the optimal protocols for the disorder free system. The batch optimisation protocol has much more drastic/rapid changes in velocity compared to the optimal clean protocols.

more (computationally expensive) optimisations need to be performed for different constraint parameters (d, τ) . In particular, near the heuristic speed limit for the clean system where the performance of STA protocols breaks down.

As a final remark we note that although the batch optimisation method does not give an enormous improvement in infidelity compared to the STA protocols the form of the protocols is completely different. We show in Fig. 6.10 (b) the optimal velocity protocol obtained at a disorder strength of $\Delta = 0.16$ compared to optimal clean Fourier and STA protocols. We observe that the batch optimisation protocol has much more rapid changes in the velocity and also obtains higher speeds. This might result from the presence of disorder but more optimisations are needed to confirm this.

6.5 Summary

In this chapter, we have looked at the optimal transport of magnons in clean and disordered Heisenberg spin chains. For this optimisation task we have exploited a hybrid numerical-analytical approach. We took two different types of analytical ansätze for the optimal magnon transport protocols and optimised them with ∂P . The first ansatz we used was a family of parameterised protocols obtained with an analytic STA approach following [70]. The second ansatz was a Fourier series ansatz and was inspired by the CRAB method for quantum control [97].

With these methods we found in the clean system that, for transport times that are longer than the heuristic speed limit time τ^* , we could efficiently find magnon transport protocols with perfect fidelity. In particular, we found a new type of strategy which comprises two small quenches in position at the beginning and end of the protocol. In between the quenches this protocol has a period of perfectly oscillatory motion. The frequency of this oscillatory motion scales linearly with the harmonic frequency of the magnetic trap. In addition, with the ∂P Fourier method, we were able to push the heuristic speed limit closer to the maximum group velocity of the magnons.

We also looked at the effects and optimisation in the presence of disorder. Here we found a structural difference in the control landscapes of the STA protocols between disordered couplings and a disordered external magnetic field. The coupling disorder has a smooth convex landscape, whereas the onsite disorder case has a rough non-convex landscape. The maximum infidelity values are found to be relatively robust with respect to both types of disorder. With optimisation we were able to find optimal magnon transport strategies with near perfect fidelity in the presence of a fixed onsite disorder pattern. However, for a batch of onsite disorder realisations additional optimisation for a long transport time τ did not seem to make any significant difference. In this case, the disorder averaged infidelity was found to increase approximately quadratically with disorder strength.

For future work a few different directions can be explored. Firstly, it would be interesting to investigate if it is possible to show that there is no universal strategy that has a perfectly fidelity in the presence of any random disorder pattern. This could explain why the batch optimisation with ∂P is not able to make any significant improvement in performance compared to the performance of the clean optimal protocols. Related to this is the question if there is an explanation for the quadratic growth of the fidelity in the presence of disorder. In particular, if the rate of this quadratic increase depends on the total transport time τ .

Another direction to explore is to see if the heuristic speed limit can be even further improved by using more (and different) control parameters. For example, one could allow a time-dependent trapping frequency $\omega(t)$ or the ability to change the couplings in time. It might be possible to transport the magnons faster than the group velocity in this way. In this respect, it will also be interesting to see what role disorder plays in the heuristic speed limit. In references [220, 221] it was found that the Lieb-Robinson velocity is changed in disordered spin chains (a logarithmic lightcone instead of a linear lightcone appears). It will be interesting to see if this has some consequences for the heuristic speed limit as well.

From an optimisation perspective several additional interesting avenues could be investigated. The effect of disorder on the control landscapes for any arbitrary type of protocol ansatz is one of these. A few other related works [222, 223], for slightly different systems, have looked at the shape of control landscapes but did not incorporate the effect of disorder. It might be that some of the smooth convex landscapes in these problems get destroyed by

disorder which could affect the optimisation and optimal protocols.

To explore these kind of questions about the effects of disorder often many computationally expensive numerical simulations need to be performed. To resolve this, and bring the computational complexity down, a final interesting question would be to see if one can train a neural network to predict the value of the cost function, like was done in [222]. The input of the neural network could, for example, be the parameters of the protocol ansatz and the output the value of the infidelity. Applying the neural network would potentially be computationally cheaper than computing the full unitary evolution numerically. A step towards this is already made in [224]. In that work it is shown that neural networks are able to correctly predict the reduced dynamics of the (many-body) XYZ Heisenberg spin model.

Chapter 7

Conclusion and Future Directions

In this thesis we have investigated the problem of stability and control of quantum information in complex many-body systems. We began by introducing the problem of optimal quantum control and several numerical and analytical techniques to solve it. In order to make the abstract formulation of the problem more accessible, we discussed a simple single-qubit example and applied some of the methods for quantum control to its initialisation. In particular, we focused on the machine learning (ML) optimisation techniques Differentiable Programming (∂P) and Natural Evolution Strategies (NES). Our motivation for using ML techniques was that they have proven to be very successful in a wide range of applications in physics [31, 32], with a few in the specific area of quantum control [36, 37, 38].

In the following chapters, we exploited some of the introduced techniques to control quantum many-body systems. The first aspect we looked at concerned topological quantum order (TQO) and the control of information encoded in the ground-state degeneracy associated with non-abelian anyons. In Chapter 3, we introduced this topic and explained that non-abelian anyons need to be braided to make topological quantum gates. We focused on one specific type of anyon, the Majorana zero mode, which potentially can be realised in real-world hybrid superconducting devices. Although we showed that these Majoranas can be moved by tuning an external gate potential, (super)adiabatically changing this potential is a slow process and therefore susceptible to noise and decoherence.

In Chapter 4, we aimed to address this issue and formulated the optimal transport of Majoranas in terms of a mathematical optimisation problem. In this guise, we took the target state infidelity as the figure of merit, which we minimised with the ML techniques. From this minimisation we obtained different types of Majorana control strategies depending on the total transport time and movement distance. To capture all the different strategies we proposed a categorisation in terms of 4 distinct transport regimes, defined by the critical velocity and the resonance time scale. In the non-adiabatic regimes we found a counterintuitive strategy that we dubbed the jump-move-jump control strategy. In the adiabatic regime the ML optimisation methods correctly recovered the smooth superadiabatic protocols. These results shed new light into this area where previously, for a related model [83], bang-bang protocols were proven to be optimal. This shows that ∂P and NES can be efficiently used to discover control

protocols for complex quantum many-body systems.

In Chapter 5 we then added an interacting term to the superconducting p-wave model and investigated its effects on the stability and control of the topologically encoded quantum information. This allowed us to make closer contact with real-world topological quantum devices and also to address some fundamental questions with respect to the spectrum of this complex many-body system. Firstly, we found that the jumps of the jump-move-jump control strategy are robust with respect to interactions, whereas during the move part of the protocol the infidelity gradually increases. We provided some potential explanations for this behaviour on the basis of mean-field theory and the critical velocity. However, a more detailed mathematical analysis is still needed to obtain a better understanding.

The second main result obtained in Chapter 5 was that the interaction-driven splittings between bulk energy modes of different topological sectors do not cause an additional phase error in the ground-state space. For this, we used an argument based on the topological quantum order (TQO) property in combination with Lieb-Robinson bounds. This result indicates that one can use TQO to derive constraints on the energy eigenvalue statistics of an interacting many-body system. Investigating this question, and also the interplay between disorder and interactions, would be an interesting avenue for further work. There would be a clear connection with works on strong zero modes [225, 180, 226] and prethermalisation [197, 194, 198].

In Chapter 6 we switched gears to a different, but related, quantum many-body system and looked at the problem of quantum-state transfer. Specifically, we focused on the XXX-Heisenberg spin chain and aimed to transfer a localised single-spin excitation, a magnon, from one end of a spin chain to the other. The Heisenberg model has as advantage that it can be implemented in a wide range of experimental platforms, which makes our results and analysis relevant for real-world quantum-state transfer applications.

For the optimisation of the magnon transport we showed that ∂P can be efficiently combined with an analytical control ansatz derived from a shortcut to adiabaticity (STA) method and also a Fourier series ansatz. With this hybrid approach, we were able to improve the heuristic speed limit for the magnons found with a fully analytical STA approach in Ref. [70]. In addition, we included disorder along the spin chain and saw that our method can find perfect state transfer protocols for fixed disorder realisations. For a batch of disorder patterns we found that the optimised averaged infidelity increases quadratically with the disorder strength. This compelling result might indicate that no universal strategy exists for perfectly transporting magnons in disordered media.

For future studies it will be interesting to see if the methods for quantum control developed and used in this thesis can be applied to different problems in quantum information. One of these is the optimisation of control pulses for open quantum systems, namely those governed by the Lindblad equation. For example, recently in Ref. [227] it was shown that the performance of quantum thermal machines can be enhanced by controlling the system-bath

coupling. Another example is the minimisation of the dissipated heat during the erasure of a quantum bit of information, as studied in references [228, 229]. Applying ∂P or NES to such quantum thermodynamics problems can potentially lead to higher efficiencies and novel control protocols. In this respect, we note that for ∂P this was already demonstrated for the optimisation of a quantum Otto engine [230]. The application of NES will also be interesting, since it does not require a model of the system and can be based on real experimental data.

Another aspect of quantum information that is worth studying deeper is the influence of projective measurements and/or continuous monitoring of quantum systems. Aside from more fundamental motivations, such as understanding measurement-induced entanglement phase transitions [231, 232], this is directly relevant for quantum control. One can exploit the information obtained with such measurements to improve the training of the ML control agent. This was done, for example, in references [233, 234, 235] and connects with the field of (quantum) feedback control [236]. It will be interesting to see, if the ML techniques combined with measurement data, or some techniques for feedback control, can improve the ability to control quantum many-body systems.

In relation to the particular many-body systems discussed in this thesis, there are also some future directions that can be investigated. For the transport of Majoranas one could look at the application of the optimisation methods directly in the presence of disorder or interactions. In addition, the methods could be used to study the control of different setups for Majorana braiding, such as the measurement-only schemes. For the magnon transport, one could look at the performance of the Natural Evolution Strategies optimisation algorithm. It will be compelling if this algorithm is also able to obtain high fidelity protocols in the presence of fixed disorder patterns. One could then, for example, try to apply NES directly to an experimental implementation of the Heisenberg chain.

From an algorithmic perspective it would be interesting to further investigate the shape of the control landscapes of these optimisation problems. Specifically, what the effects are of interactions, disorder or a coupling to the environment. For example, for the magnon transport we saw that disorder significantly affects the shape of the cost landscape. Studying these landscapes gives insights both into the control problem itself as well as which optimisation technique is best to use. For instance, the roughness of this landscape determines the complexity of the optimisation problem and the chances of successfully finding an optimal control protocol. Thus, this knowledge is relevant for the realisation of controllable quantum devices.

Lastly, to end this thesis with a more broader and futuristic outlook, it will be fascinating to explore the possibility for a fully autonomous ML quantum-control agent. This ML agent could, for example, be composed of two different (sub)machines. The first machine tries to understand and learn the behaviour of the quantum system and gives as output an estimate for the value of the control functional \mathcal{L} , for instance the infidelity. This machine could learn to

predict \mathcal{L} from the feedback by experimental measurements or by training on a mathematical model. After the training is completed an estimate for \mathcal{L} could then be obtained with a much lower computational complexity than the complexity of the full many-body quantum evolution. A similar idea was exploited in references [222, 224].

The second machine can then use the prediction of the cost by the first machine to devise efficient control strategies. This ML scheme for quantum control might be related to generative adversarial networks (GANs) [237, 238], where two ML algorithms play a zero-sum game. In order to make this connection one needs to be able to formulate the quantum control problem in a way that GANs can be applied. It will be intriguing to see if this is possible and if the development of such an autonomous ML control agent can lead to stable and controllable large scale quantum devices.

Appendix A

The Bogoliubov-de-Gennes formalism

Here we shortly introduce the Bogoliubov-de-Gennes (BdG) formalism for (free) fermions hopping on a lattice [40, 135]. We will start from the generic quadratic Hamiltonian

$$\mathcal{H} = \frac{1}{2} \sum_{jk} \xi_{jk} c_j^\dagger c_k - [\xi_{jk}]^T c_j c_k^\dagger + \Delta_{jk} c_j^\dagger c_k^\dagger + [\Delta_{jk}]^\dagger c_j c_k \quad (\text{A.1})$$

in which ξ_{jk} describes an onsite chemical potential when $j = k$, and hopping between different lattice sites j and k otherwise. The inclusion of Δ_{ij} gives the possibility for superconductivity (electron pair creation/annihilation). In the first section below, we will describe how to solve this Hamiltonian. In the second section we review quantum time evolution within the BdG formalism.

A.1 General Solution of the Quadratic BdG Hamiltonian

First one writes Hamiltonian (A.1) in BdG form,

$$\mathcal{H}_0 = \frac{1}{2} \begin{bmatrix} c_{\leftrightarrow}^\dagger & c_{\leftrightarrow} \end{bmatrix} \begin{bmatrix} \xi & \Delta \\ \Delta^\dagger & -\xi^T \end{bmatrix} \begin{bmatrix} c_{\uparrow} \\ c_{\downarrow}^\dagger \end{bmatrix} = \frac{1}{2} [C^\dagger] [H_{BdG}] [C] \quad (\text{A.2})$$

with $C^\dagger \equiv [c_{\leftrightarrow}^\dagger \quad c_{\leftrightarrow}] \equiv [c_1^\dagger \dots c_i^\dagger \dots c_N^\dagger \quad c_1 \dots c_i \dots c_N]$, and diagonalises

$$[H_{BdG}][W] = [W][D] \quad (\text{A.3})$$

with eigenvectors $W = \begin{bmatrix} U & V^* \\ V & U^* \end{bmatrix}$ and eigenmode energies $[D]_{ii} = \epsilon_i$. The U, V matrices correspond to the coefficients of the BdG quasiparticle transformation β_n , given in Eq. 5.18 in the main text. Applying this formalism to the Kitaev chain model in the momentum representation (Eq. 3.4), results in the energy dispersion ϵ_k in Eq. 3.5.

A.2 BdG Time Evolution

In this section we show how to time evolve many-body states with ordinary Schrödinger time evolution within the BdG formalism. We also show how to calculate overlaps of states in the many-body picture by utilising the BdG quasiparticle mode picture.

To time evolve an initial state in the many-body picture $|\psi(0)\rangle$ up to time T one needs to apply the unitary evolution operator

$$u(T) = \mathcal{T} e^{-i \int_0^T H(t') dt'}. \quad (\text{A.4})$$

Here $H(t)$ is the time dependent many-body Hamiltonian of the system, and \mathcal{T} the time ordering operator. By discretising time, in time steps δt , and exploiting the Trotter decomposition (equation C.4) when $[H(t), H(t + \delta t)] \neq 0$, $u(T)$ can be approximated by

$$u(T) \approx \prod_{t=0}^{T/\delta t} e^{-iH(t)\delta t} + \mathcal{O}(t^2). \quad (\text{A.5})$$

This can be transformed from the many-body picture to the BdG picture with a time dependent BdG transformation

$$\beta_n(t) = \sum_j V_{jn}^*(t) c_j^\dagger + U_{jn}^*(t) c_j. \quad (\text{A.6})$$

$V(t)$ and $U(t)$ diagonalise $H_{\text{BdG}}(t)$ at time t ,

$$[H_{\text{BdG}}(t)][W(t)] = [E(t)][W(t)], \quad (\text{A.7})$$

with corresponding eigenvalue matrix $E(t)$. Then the evolution operator in the BdG picture is given by

$$\mathcal{U}(t) = W(t) e^{iE(t)\delta t} W^\dagger(t). \quad (\text{A.8})$$

The BdG quasiparticle modes can now be evolved forward in time by $W(t + \delta t) = \mathcal{U}(t)W(t)$.

To go back to the many-body picture and calculate overlaps, one can make use of the Onishi formula

$$|\langle \psi(0) | \psi(T) \rangle|^2 = \det [U(0)^* U(T) + V(0)^* V(T)], \quad (\text{A.9})$$

which can be derived with help of the Thouless theorem [40].

Appendix B

Equivalence between the Interacting Kitaev Chain and XYZ Heisenberg Spin Chain

In this appendix we show the equivalence/duality between the (interacting) Kitaev chain [51] and the XY(Z)-Heisenberg [63] spin-1/2 model. We will use the 1928 Jordan-Wigner [239] result that qubits (spin-1/2) particles are dual to fermions obeying the fermionic commutation relations $\{c_i, c_j^\dagger\} = \delta_{ij}$ and $\{c_i^{(\dagger)}, c_j^{(\dagger)}\} = 0$. This isomorphism between spins and fermions has been extended to more than one dimension [240, 241], but we focus here on the one-dimensional version.

This Jordan-Wigner transformation is given by a linear map which maps the Pauli spin operators $\vec{\sigma}_i$ into the spinless fermionic creation and annihilation operators

$$\sigma_i^x = \prod_{i < j} (1 - 2c_j^\dagger c_j) (c_i + c_i^\dagger), \quad (\text{B.1})$$

$$\sigma_i^y = i \prod_{i < j} (1 - 2c_j^\dagger c_j) (c_i - c_i^\dagger), \quad (\text{B.2})$$

$$\sigma_i^z = (1 - 2c_i^\dagger c_i). \quad (\text{B.3})$$

Applying this map to the XYZ Heisenberg N -site spin chain

$$H = - \sum_{i=1}^N (J_x \sigma_i^x \sigma_{i+1}^x + J_y \sigma_i^y \sigma_{i+1}^y + J_z \sigma_i^z \sigma_{i+1}^z + h_z \sigma_i^z) \quad (\text{B.4})$$

with coupling constants \vec{J} and external field h_z , we obtain

$$H = - \sum_{i=1}^N [(-2h_z - 2J_z - 2J_z|_{i>2}) c_i^\dagger c_i + (J_x + J_y) (c_i^\dagger c_{i+1} + h.c.) + (J_x - J_z) (c_i^\dagger c_{i+1}^\dagger + h.c.) + 4J_z c_i^\dagger c_i c_{i+1}^\dagger c_{i+1} + h_z + J_z]. \quad (\text{B.5})$$

It can be seen that H resembles the interacting Kitaev chain model given in Eqs. 3.1 & 5.1 in the main text. The coupling constants are linked to the Kitaev chain parameters by

$$\mu - V = -2h_z - 2J_z, \quad t = J_x + J_y, \quad \Delta = J_x - J_y, \quad U = -J_z. \quad (\text{B.6})$$

Note that we did not include some boundary effect term for the chemical potential resulting from the Z -interaction. However, the interpretation is different in both pictures (spins vs. fermions), this makes the Jordan-Wigner transformation useful if one wants to gain more insight in another physical picture. For example, [242] solved the one-dimensional antiferromagnetic XY chain with a Jordan-Wigner transformation to the fermion picture.

Appendix C

Tensor Networks

In this appendix we introduce the topic of tensor networks and the density matrix renormalisation group (DMRG), which has been a widely used numerical method to study quantum systems of low-entanglement ever since the initial invention by Steven White in 1992 [61]. Because we are interested in interactions, which makes the interacting Kitaev chain not exactly solvable with methods like BdG, the use of MPS and DMRG, which are able to capture interactions, is very useful. This method will now be shortly introduced, starting from the definition of matrix product states. For more extensive reviews we refer the reader to references [62, 243].

C.1 Matrix Product States

Before giving the definition of a matrix product state (MPS), we first recall the definition of the Schmidt decomposition of a quantum state $|\Psi\rangle$. $|\Psi\rangle$ is living in the 2^N dimensional Hilbert space \mathcal{H} and can be expanded into a tensor product of two partitions A and B , with $2^{N/2}$ dimensional orthonormal bases $|i\rangle_A$ and $|i\rangle_B$, giving

$$|\Psi\rangle_{A+B} = \sum_i \lambda_i |i\rangle_A |i\rangle_B. \quad (\text{C.1})$$

Here, the λ_i are the square roots of the eigenvalues of the density matrix, which can be obtained from the singular values of the singular value decomposition (SVD) of $|\Psi\rangle_{A+B}$.

We restrict to quantum systems in one-spatial dimension consisting of N lattice sites, the extension of MPS to higher dimensions is called projected entangled pair states (PEPS) and can be found in [244]. By partitioning the system in N subsystems, one for each site, by performing $N - 1$ Schmidt decompositions (Singular Value Decompositions) we obtain the MPS

$$|\Psi\rangle_N = \sum_{i_1, i_2, \dots, i_N} \vec{A}[i_1] \vec{A}[i_2] \dots \vec{A}[i_N] |i_1\rangle |i_2\rangle \dots |i_N\rangle. \quad (\text{C.2})$$

The \vec{A}_i for each site are the two-index tensors, matrices in this case, with local onsite physical basis states $|i_n\rangle$ (spin up/down for a spin-1/2 chain). The matrices are bonded together in terms

of a tensor product, which is contracting the shared singular values of the expansion between two sites. This is an example of a tensor network, a quantum state with Hilbert space of dimension 2^N is expanded into smaller local tensors connected by bonds with bond dimension χ . We can do a similar expansion for operators such as Hamiltonians, called matrix product operators MPO, as can be found in references [62, 245]. In the MPO the only difference is that there are two physical indices i_n and j_n , one for the physical kets and one for the physical bras, as in $|j_n\rangle\langle i_n|$.

Since writing down all the tensors with indices is quite tedious, there is a graphical notation for MPS/MPO as shown in figure C.1 for a $N = 10$ system. The tensors are represented by boxes, the physical basis by vertical bonds and the tensors are connected by horizontal (virtual) bonds. Expectation values and overlaps can be easily calculated in this notation by contracting over the indices.

$$\Psi^{i_1 \dots i_{10}} = \begin{array}{cccccccccccc} \text{---} & \boxed{A} & \text{---} & \boxed{A} & \text{---} & \boxed{A} & \text{---} & \boxed{A} & \text{---} & \boxed{A} & \text{---} & \boxed{A} & \text{---} & \boxed{A} & \text{---} & \boxed{A} & \text{---} & \boxed{A} & \text{---} & \boxed{A} & \text{---} \\ & | & & | & & | & & | & & | & & | & & | & & | & & | & & | & & | \\ & i_1 & & i_2 & & i_3 & & i_4 & & i_5 & & i_6 & & i_7 & & i_8 & & i_9 & & i_{10} & & \end{array}$$

$$\mathcal{M}_{i_1 \dots i_{10}}^{j_1 \dots j_{10}} = \begin{array}{cccccccccccc} & i_1 & & i_2 & & i_3 & & i_4 & & i_5 & & i_6 & & i_7 & & i_8 & & i_9 & & i_{10} & & \\ & | & & | & & | & & | & & | & & | & & | & & | & & | & & | & & | & & \\ \text{---} & \boxed{M} & \text{---} & \boxed{M} & \text{---} & \boxed{M} & \text{---} & \boxed{M} & \text{---} & \boxed{M} & \text{---} & \boxed{M} & \text{---} & \boxed{M} & \text{---} & \boxed{M} & \text{---} & \boxed{M} & \text{---} & \boxed{M} & \text{---} & \boxed{M} & \text{---} \\ & | & & | & & | & & | & & | & & | & & | & & | & & | & & | & & | & & \\ & j_1 & & j_2 & & j_3 & & j_4 & & j_5 & & j_6 & & j_7 & & j_8 & & j_9 & & j_{10} & & \end{array}$$

FIGURE C.1: Graphical representation for tensor networks. The top expansion is a MPS of the form $|\Psi\rangle = \sum_i^{10} \Psi^{i_1 \dots i_{10}} |i_1\rangle \dots |i_{10}\rangle$. The bottom expansion is a MPO of the form $M = \sum_i^{10} M_{i_1 \dots i_{10}}^{j_1 \dots j_{10}} |j_1\rangle \dots |j_{10}\rangle \langle i_1| \dots \langle i_{10}|$.

C.2 Density Matrix Renormalization Group

In principle any 1D quantum state can be represented by a MPS if one does not put a bound on the bond dimension χ . However, this does not give much advantage in computational speed over, for example, exact diagonalisation methods. Steven White introduced an ansatz [61], now known as the density matrix renormalisation group (DMRG), in which there is a bound $\chi' \leq \chi$ on the amount of entanglement, i.e. the number of singular values kept from the SVD for the Schmidt decomposition [Fig. C.2 (c)]. In this way, many low-entangled ground states of systems with local Hamiltonians can be efficiently described by an MPS. The downside is

that this bound results in a truncation error given by

$$\| |\Psi\rangle - |\Psi\rangle_{trun} \|^2 \leq 2 \sum_i^N (\epsilon_i(\chi) - \epsilon_i(\chi')), \quad (\text{C.3})$$

where $\epsilon_i(x)$ is the sum of the x biggest singular values of the SVD at site i . This will give problems for high entangled phases, for example near phase transition points, however, in regions far away from these points it will not give problems. The main advantage of DMRG is that for a Hamiltonian given as MPO, the ground state can be efficiently found variationally for large system sizes, much larger than with exact diagonalisation is possible. This variational ground state search can be done by solving a generalized eigenvalue problem as explained in [62].




FIGURE C.2: Diagram representation of time evolution in DMRG. (a) Local exponentiated term, two-site MPO. (b) Application of the local MPOs first to the even bonds and then to the odd bonds. (c) The new time evolved MPS sites are reduced by performing an SVD with truncation D on the number of singular values.

C.3 Time Evolution with MPS

Time evolution is nontrivial in DMRG because one expects the entanglement to grow with time as the system is evolved. Therefore, the required bond dimension χ' to keep the truncation error small will need to grow, resulting in a higher computational cost. Time evolution is done by making the evolution operator as a MPO and applying it to the MPS to move it forward in time. This new evolved MPS is reduced by bringing the bond dimension to manageable sizes as shown in Fig. C.2. There are several schemes available to do this [62], but one common way is via Trotter time evolution (TEBD). The time evolution operator is

approximated via a Trotter decomposition

$$e^{-\tau\hat{H}} \approx \prod_{i=1}^{N/2} e^{-\tau\hat{h}_i} \prod_{i=1}^{N/2} e^{-\tau\hat{h}_{i+1}} + \mathcal{O}(\tau^2), \quad (\text{C.4})$$

where one applies local exponentiated MPO terms $e^{-\tau\hat{h}_i}$ (ordinary matrix exponential) first to the odd and then to the even bonds (figure C.2). Another more stable scheme, without Trotter error, to do this is the time-dependent-variational principle (TDVP). This method is solving a minimisation problem for the Schrödinger equation to find an approximation for the new time-evolved MPS as discussed in reference [59].

Appendix D

Topological Order of the Non-Interacting Kitaev Chain

In this appendix we show that

$$\langle e/o | \beta_0 [O, \beta_0^\dagger] | e/o \rangle = e^{-L/\xi} \quad (\text{D.1})$$

which is required to show the TQO property the non-interacting Kitaev chain \mathcal{H}_0 by the argument in section 5.1 in the main text. In here O is some sufficiently local parity preserving operator and $\beta_0^\dagger = \frac{1}{2}(\Gamma_l - i\Gamma_r)$ is the creation operator for the Dirac fermionic zero mode. Since we have assumed perfectly localized Majoranas ($\Delta = w, \mu = 0$)¹ we have $\Gamma_l = \gamma_1$ and $\Gamma_r = \gamma_L$ where $x = 1, L$ are the first and last lattice sites of the chain. The γ_1 and γ_L are thus the uncoupled Majorana bound state operators whereas the bulk Majorana operators γ_i with $i \neq 1, L$ are the dimerized bulk Majorana operators, see section 3.2.1 in the main text.

To show Eq. D.1 we first assume a simple 2-body parity preserving form

$$O = \sum_{i,j} O_{ij} \gamma_i \gamma_j. \quad (\text{D.2})$$

To make sure this operator is sufficiently local we impose $|i - j| < L$, i.e. they do not cross the entire system. The commutator in Eq. D.1 can then be written as

$$\frac{1}{2} \sum_{j,k} O_{jk} [\gamma_j \gamma_k, \gamma_1 - i\gamma_L] = \frac{1}{2} \sum_{j,k \neq L} O_{ij} [\gamma_j \gamma_k, \gamma_1] - \frac{i}{2} \sum_{j,k \neq 1} O_{ij} [\gamma_j \gamma_k, \gamma_L]. \quad (\text{D.3})$$

Focusing for the moment only on the first term we note that this commutator is only non-zero when either $j = 1, k \neq 1$ or $j \neq 1, k = 1$. When $j = 1$ we find $[\gamma_1 \gamma_k, \gamma_1] = -2\gamma_k$ and when $k = 1$ we find $[\gamma_j \gamma_1, \gamma_1] = 2\gamma_j$. However the expectation value in Eq. D.1 for these commutators

$$\langle e/o | \beta_0 [O, \beta_0^\dagger] | e/o \rangle = \langle e/o | (\gamma_1 + \gamma_L) \gamma_j | e/o \rangle = 0 \quad (\text{D.4})$$

¹We can always use quasiadiabatic continuation to argue TQO for a \mathcal{H}_0 with some different parameters for which the gap is still open.

because $j \neq 1, L$. The same argument, with permuted indices, can be made for the second term in Eq. D.3 to show it also has zero contribution.

This shows that for the particular 2-body operator in Eq. D.2 we have an exact TQO condition as long as it has a maximum support of size $L^* = L - 1$. The exactness comes from the fact that the RHS in Eq. D.1 in this case is exactly zero which results from the perfectly localized Majoranas. For different parameters for which the Majoranas have some finite decay length $u_l(x) \neq \delta_{x,1}$ one gets the approximate TQO condition with a correction on the order of $\mathcal{O}(e^{-L/\xi})$. Furthermore, it is straightforward to generalize the argument above to N-body operators O which have support only on a region of maximum size $L^* = L - 1$. This can be seen by again considering the commutation properties with the localized Majorana operators γ_1 and γ_L .

D.0.1 Groundstate Degeneracy as a Consequence of Topological Order

For completeness we now also show the mathematical derivation behind the argument that we use in the main text for the groundstate degeneracy of some systems² that possess TQO. We consider Hamiltonians of the form

$$H = \sum_k H_k \quad (\text{D.5})$$

where the H_k are maximally L^* -local. Now by the TQO condition Eq. 5.4 we have

$$E_e^0 - E_o^0 = \sum_k \langle e | H_k | e \rangle - \langle o | H_k | o \rangle \quad (\text{D.6})$$

$$= k \times \mathcal{O}(e^{-L/\xi}). \quad (\text{D.7})$$

This difference is exponentially suppressed because the number of terms in the sum k grows at most linearly with the system size L . From this we conclude that the groundstates of the even and odd sectors are degenerate $E_e^0 \approx E_o^0$ for long enough system sizes L .

²Systems that can be described by local Hamiltonians

Bibliography

- [1] Jonathan P Dowling and Gerard J Milburn. Quantum technology: The second quantum revolution. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, 361:1655–74, 09 2003.
- [2] David Deutsch and Roger Penrose. Quantum theory, the church’s turing principle and the universal quantum computer. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 400(1818):97–117, 1985.
- [3] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018.
- [4] Stephen Jordan. Quantum algorithm zoo.
- [5] Christof Zalka. Simulating quantum systems on a quantum computer. *Proceedings: Mathematical, Physical and Engineering Sciences*, 454(1969):313–322, 1998.
- [6] Adam Smith, M. S. Kim, Frank Pollmann, and Johannes Knolle. Simulating quantum many-body dynamics on a current digital quantum computer. *npj Quantum Information*, 5(1):106, Nov 2019.
- [7] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549(7671):195–202, Sep 2017.
- [8] M. Schuld and F. Petruccione. *Supervised Learning with Quantum Computers*. Quantum Science and Technology. Springer International Publishing, 2018.
- [9] Charles H. Bennett, François Bessette, Gilles Brassard, Louis Salvail, and John Smolin. Experimental quantum cryptography. *Journal of Cryptology*, 5(1):3–28, Jan 1992.
- [10] Stephen Wiesner. Conjugate coding. *SIGACT News*, 15(1):78–88, January 1983.
- [11] S. Pirandola, U. L. Andersen, L. Banchi, M. Berta, D. Bunandar, R. Colbeck, D. Englund, T. Gehring, C. Lupo, C. Ottaviani, J. L. Pereira, M. Razavi, J. Shamsul Shaari, M. Tomamichel, V. C. Usenko, G. Vallone, P. Villoresi, and P. Wallden. Advances in quantum cryptography. *Adv. Opt. Photon.*, 12(4):1012–1236, Dec 2020.

- [12] Jay M. Gambetta, Jerry M. Chow, and Matthias Steffen. Building logical qubits in a superconducting quantum computing system. *npj Quantum Information*, 3(1):2, Jan 2017.
- [13] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Phys. Rev. Lett.*, 74:4091–4094, May 1995.
- [14] Justin G. Bohnet, Brian C. Sawyer, Joseph W. Britton, Michael L. Wall, Ana Maria Rey, Michael Foss-Feig, and John J. Bollinger. Quantum spin dynamics and entanglement generation with hundreds of trapped ions. *Science*, 352(6291):1297–1301, 2016.
- [15] Colin D. Bruzewicz, John Chiaverini, Robert McConnell, and Jeremy M. Sage. Trapped-ion quantum computing: Progress and challenges. *Applied Physics Reviews*, 6(2):021314, 2019.
- [16] Brian B. Zhou, Alexandre Baksic, Hugo Ribeiro, Christopher G. Yale, F. Joseph Heremans, Paul C. Jerger, Adrian Auer, Guido Burkard, Aashish A. Clerk, and David D. Awschalom. Accelerated quantum control using superadiabatic dynamics in a solid-state lambda system. *Nature Physics*, 13(4):330–334, Apr 2017.
- [17] Francesco Casola, Toeno van der Sar, and Amir Yacoby. Probing condensed matter physics with magnetometry based on nitrogen-vacancy centres in diamond. *Nature Reviews Materials*, 3(1):17088, Jan 2018.
- [18] D.E. Kirk. *Optimal Control Theory: An Introduction*. Dover Books on Electrical Engineering. Dover Publications, 2012.
- [19] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, New York, NY, USA, 10th edition, 2011.
- [20] H.P. Breuer, P.I.H.P. Breuer, F. Petruccione, and S.P.A.P.F. Petruccione. *The Theory of Open Quantum Systems*. Oxford University Press, 2002.
- [21] J.A. Snyman and D.N. Wilke. *Practical Mathematical Optimization: Basic Optimization Theory and Gradient-Based Algorithms*. Springer Optimization and Its Applications. Springer International Publishing, 2018.
- [22] L.S. Pontryagin. *Mathematical Theory of Optimal Processes*. CRC Press, 2018.
- [23] Chungwei Lin, Yanting Ma, and Dries Sels. Optimal control for quantum metrology via pontryagin’s principle. *Phys. Rev. A*, 103:052607, May 2021.

- [24] Zhi-Cheng Yang, Armin Rahmani, Alireza Shabani, Hartmut Neven, and Claudio Chamon. Optimizing variational quantum algorithms using pontryagin’s minimum principle. *Phys. Rev. X*, 7:021027, May 2017.
- [25] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- [26] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.
- [27] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martínez-Garaot, and J. G. Muga. Shortcuts to adiabaticity: Concepts, methods, and applications. *Rev. Mod. Phys.*, 91:045001, Oct 2019.
- [28] Tom Young, Devamanyu Hazarika, Soujanya Poria, and Erik Cambria. Recent Trends in Deep Learning Based Natural Language Processing. *arXiv:1708.02709 [cs]*, November 2018.
- [29] Athanasios Voulodimos, Nikolaos Doulamis, Anastasios Doulamis, and Eftychios Protopapadakis. Deep Learning for Computer Vision: A Brief Review. *Computational Intelligence and Neuroscience*, 2018:7068349, 2018.
- [30] Gökçen Eraslan, Žiga Avsec, Julien Gagneur, and Fabian J. Theis. Deep learning: New computational modelling techniques for genomics. *Nature Reviews Genetics*, 20(7):389–403, July 2019.
- [31] Giuseppe Carleo, Ignacio Cirac, Kyle Cranmer, Laurent Daudet, Maria Schuld, Naftali Tishby, Leslie Vogt-Maranto, and Lenka Zdeborová. Machine learning and the physical sciences. *Rev. Mod. Phys.*, 91:045002, Dec 2019.
- [32] Juan Carrasquilla. Machine learning for quantum matter. *Advances in Physics: X*, 5(1):1797528, Jan 2020.
- [33] Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, February 2017.
- [34] Juan Carrasquilla and Roger G. Melko. Machine learning phases of matter. *Nature Physics*, 13(5):431–434, May 2017.
- [35] John C. Snyder, Matthias Rupp, Katja Hansen, Klaus-Robert Müller, and Kieron Burke. Finding density functionals with machine learning. *Phys. Rev. Lett.*, 108:253002, Jun 2012.

- [36] Marin Bukov, Alexandre G. R. Day, Dries Sels, Phillip Weinberg, Anatoli Polkovnikov, and Pankaj Mehta. Reinforcement learning in different phases of quantum control. *Phys. Rev. X*, 8:031086, Sep 2018.
- [37] Xiao-Ming Zhang, Zezhu Wei, Raza Asad, Xu-Chen Yang, and Xin Wang. When does reinforcement learning stand out in quantum control? A comparative study on state preparation. *npj Quantum Information*, 5(1):1–7, October 2019.
- [38] Frank Schäfer, Michal Kloc, Christoph Bruder, and Niels Lörch. A differentiable programming method for quantum control. *Machine Learning: Science and Technology*, May 2020.
- [39] A. Altland, P.T.C.M.P.A. Altland, and B. Simons. *Condensed Matter Field Theory*. Cambridge University Press, 2006.
- [40] P. Ring and P. Schuck. *The Nuclear Many-Body Problem*. Springer-Verlag, New York, 1980.
- [41] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, Oct 2019.
- [42] Alexander M. Dalzell, Aram W. Harrow, Dax Enshan Koh, and Rolando L. La Placa. How many qubits are needed for quantum computational supremacy? *Quantum*, 4:264, May 2020.
- [43] Xiao-Gang Wen. Colloquium : Zoo of quantum-topological phases of matter. *Reviews of Modern Physics*, 89(4), Dec 2017.

- [44] K. v. Klitzing, G. Dorda, and M. Pepper. New method for high-accuracy determination of the fine-structure constant based on quantized hall resistance. *Phys. Rev. Lett.*, 45:494–497, Aug 1980.
- [45] D. C. Tsui, H. L. Stormer, and A. C. Gossard. Two-dimensional magnetotransport in the extreme quantum limit. *Phys. Rev. Lett.*, 48:1559–1562, May 1982.
- [46] R. B. Laughlin. Anomalous quantum hall effect: An incompressible quantum fluid with fractionally charged excitations. *Phys. Rev. Lett.*, 50:1395–1398, May 1983.
- [47] Xiao-Gang Wen. Topological orders and edge excitations in fractional quantum hall states. *Advances in Physics*, 44(5):405–473, 1995.
- [48] X. G. WEN. Topological orders in rigid states. *International Journal of Modern Physics B*, 04(02):239–271, 1990.
- [49] A.Yu. Kitaev. Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303(1):2 – 30, 2003.
- [50] Chetan Nayak, Steven H. Simon, Ady Stern, Michael Freedman, and Sankar Das Sarma. Non-abelian anyons and topological quantum computation. *Rev. Mod. Phys.*, 80:1083–1159, Sep 2008.
- [51] A Yu Kitaev. Unpaired majorana fermions in quantum wires. *Physics-Usppekhi*, 44(10S):131–136, oct 2001.
- [52] Yuval Oreg, Gil Refael, and Felix von Oppen. Helical liquids and majorana bound states in quantum wires. *Phys. Rev. Lett.*, 105:177002, Oct 2010.
- [53] Roman M. Lutchyn, Jay D. Sau, and S. Das Sarma. Majorana fermions and a topological phase transition in semiconductor-superconductor heterostructures. *Phys. Rev. Lett.*, 105:077001, Aug 2010.
- [54] Liang Fu and C. L. Kane. Superconducting proximity effect and majorana fermions at the surface of a topological insulator. *Phys. Rev. Lett.*, 100:096407, Mar 2008.
- [55] V. Mourik, K. Zuo, S. M. Frolov, S. R. Plissard, E. P. A. M. Bakkers, and L. P. Kouwenhoven. Signatures of majorana fermions in hybrid superconductor-semiconductor nanowire devices. *Science*, 336(6084):1003–1007, 2012.
- [56] R. M. Lutchyn, E. P. A. M. Bakkers, L. P. Kouwenhoven, P. Krogstrup, C. M. Marcus, and Y. Oreg. Majorana zero modes in superconductor–semiconductor heterostructures. *Nature Reviews Materials*, 3(5):52–68, May 2018.

- [57] Piet Brouwer, Klaus Ensslin, David Goldhaber-Gordon, and Patrick Lee. Nature paper "Quantized Majorana conductance", report from independent experts, March 2021.
- [58] Haining Pan and S. Das Sarma. Physical mechanisms for zero-bias conductance peaks in majorana nanowires. *Physical Review Research*, 2(1), Mar 2020.
- [59] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pižorn, H. Verschelde, and F. Verstraete. Time-dependent variational principle for quantum lattices. *Phys. Rev. Lett.*, 107:070601, Aug 2011.
- [60] S. Paeckel, T. Köhler, A. Swoboda, S.R. Manmana, U. Schollwöck, and C. Hubig. Time-evolution methods for matrix-product states. *Annals of Physics*, 411:167998, 2019.
- [61] Steven R. White. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.*, 69:2863–2866, Nov 1992.
- [62] Ulrich Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics*, 326(1):96 – 192, 2011. January 2011 Special Issue.
- [63] W. Heisenberg. *Zur Theorie des Ferromagnetismus*, pages 580–597. Springer Berlin Heidelberg, Berlin, Heidelberg, 1985.
- [64] L.-M. Duan, E. Demler, and M. D. Lukin. Controlling spin exchange interactions of ultracold atoms in optical lattices. *Phys. Rev. Lett.*, 91:090402, Aug 2003.
- [65] Alessandro Romito, Rosario Fazio, and C. Bruder. Solid-state quantum communication with josephson arrays. *Phys. Rev. B*, 71:100501, Mar 2005.
- [66] Sebastian Hild, Takeshi Fukuhara, Peter Schauß, Johannes Zeiher, Michael Knap, Eugene Demler, Immanuel Bloch, and Christian Gross. Far-from-equilibrium spin transport in heisenberg quantum magnets. *Phys. Rev. Lett.*, 113:147205, Oct 2014.
- [67] Haifeng Qiao, Yadav P. Kandel, Kuangyin Deng, Saeed Fallahi, Geoffrey C. Gardner, Michael J. Manfra, Edwin Barnes, and John M. Nichol. Coherent multispin exchange coupling in a quantum-dot spin chain. *Phys. Rev. X*, 10:031006, Jul 2020.
- [68] S. Sachdev. *Quantum Phase Transitions*. Cambridge University Press, 2011.
- [69] Gabriele De Chiara, Simone Montangero, Pasquale Calabrese, and Rosario Fazio. Entanglement entropy dynamics of heisenberg chains. *Journal of Statistical Mechanics: Theory and Experiment*, 2006(03):P03001–P03001, mar 2006.
- [70] Anthony Kiely and Steve Campbell. Fast and robust magnon transport in a spin chain. mar 2021.

- [71] P. Zoller, Th. Beth, D. Binosi, R. Blatt, H. Briegel, D. Bruss, T. Calarco, J. I. Cirac, D. Deutsch, J. Eisert, A. Ekert, C. Fabre, N. Gisin, P. Grangiere, M. Grassl, S. Haroche, A. Imamoglu, A. Karlson, J. Kempe, L. Kouwenhoven, S. Kröll, G. Leuchs, M. Lewenstein, D. Loss, N. Lütkenhaus, S. Massar, J. E. Mooij, M. B. Plenio, E. Polzik, S. Popescu, G. Rempe, A. Sergienko, D. Suter, J. Twamley, G. Wendin, R. Werner, A. Winter, J. Wrachtrup, and A. Zeilinger. Quantum information processing and communication. *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, 36(2):203–228, Nov 2005.
- [72] H. J. Kimble. The quantum internet. *Nature*, 453(7198):1023–1030, Jun 2008.
- [73] Luuk Coopmans, Di Luo, Graham Kells, Bryan K. Clark, and Juan Carrasquilla. Protocol discovery for the quantum control of majoranas by differentiable programming and natural evolution strategies. *PRX Quantum*, 2:020332, Jun 2021.
- [74] L. Coopmans, S. Dooley, I. Jubb, K. Kavanagh, and G. Kells. Dynamical phase error in interacting topological quantum memories. *Physical Review Research*, 3(3), Jul 2021.
- [75] V. F. Krotov. Global Methods in Optimal Control Theory. In Alexander B. Kurzhanski, editor, *Advances in Nonlinear Dynamics and Control: A Report from Russia*, Progress in Systems and Control Theory, pages 74–121. Birkhäuser, Boston, MA, 1993.
- [76] Christopher J. C. H. Watkins and Peter Dayan. Q-learning. *Machine Learning*, 8(3):279–292, May 1992.
- [77] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. A Bradford Book, Cambridge, MA, USA, 2018.
- [78] Lawrence Evans. An introduction to mathematical optimal control theory version 0.2. 02 2013.
- [79] Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Applied Optimization. Springer US, 2003.
- [80] Gerhard C. Hegerfeldt. Driving at the quantum speed limit: Optimal control of a two-level system. *Phys. Rev. Lett.*, 111:260501, Dec 2013.
- [81] Marin Bukov, Alexandre G. R. Day, Dries Sels, Phillip Weinberg, Anatoli Polkovnikov, and Pankaj Mehta. Reinforcement learning in different phases of quantum control. *Phys. Rev. X*, 8:031086, Sep 2018.
- [82] Sebastian Deffner. Optimal control of a qubit in an optical cavity. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 47(14):145502, jul 2014.

- [83] Torsten Karzig, Armin Rahmani, Felix von Oppen, and Gil Refael. Optimal control of majorana zero modes. *Phys. Rev. B*, 91:201404, May 2015.
- [84] M. Born and V. Fock. Beweis des adiabatensatzes. *Zeitschrift für Physik*, 51(3):165–180, Mar 1928.
- [85] D.J. Griffiths and D.F. Schroeter. *Introduction to Quantum Mechanics*. Cambridge University Press, 2018.
- [86] Mustafa Demirplak and Stuart A. Rice. On the consistency, extremal, and global properties of counterdiabatic fields. *The Journal of Chemical Physics*, 129(15):154111, 2008.
- [87] M V Berry. Transitionless quantum driving. *Journal of Physics A: Mathematical and Theoretical*, 42(36):365303, aug 2009.
- [88] Dries Sels and Anatoli Polkovnikov. Minimizing irreversible losses in quantum systems by local counterdiabatic driving. *Proceedings of the National Academy of Sciences*, 114(20):E3909–E3916, May 2017.
- [89] H. R. Lewis and W. B. Riesenfeld. An exact quantum theory of the time-dependent harmonic oscillator and of a charged particle in a time-dependent electromagnetic field. *Journal of Mathematical Physics*, 10(8):1458–1473, 1969.
- [90] H. Ralph Lewis and P. G. L. Leach. A direct approach to finding exact invariants for one-dimensional time-dependent classical hamiltonians. *Journal of Mathematical Physics*, 23(12):2371–2374, 1982.
- [91] A Ruschhaupt, Xi Chen, D Alonso, and J G Muga. Optimally robust shortcuts to population inversion in two-level quantum systems. *New Journal of Physics*, 14(9):093040, sep 2012.
- [92] J. Nocedal and S. Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer New York, 2006.
- [93] Martin Pincus. Letter to the editor—a monte carlo method for the approximate solution of certain types of constrained optimization problems. *Operations Research*, 18(6):1225–1228, 1970.
- [94] A. Khachaturyan, S. Semenovsovskaya, and B. Vainshtein. The thermodynamic approach to the structure analysis of crystals. *Acta Crystallographica Section A*, 37(5):742–754, Sep 1981.

- [95] C. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer Texts in Statistics. Springer New York, 2013.
- [96] Patrick Doria, Tommaso Calarco, and Simone Montangero. Optimal control technique for many-body quantum dynamics. *Physical Review Letters*, 106(19), May 2011.
- [97] Tommaso Caneva, Tommaso Calarco, and Simone Montangero. Chopped random-basis quantum optimization. *Physical Review A*, 84(2), Aug 2011.
- [98] J. A. Nelder and R. Mead. A Simplex Method for Function Minimization. *The Computer Journal*, 7(4):308–313, 01 1965.
- [99] Jeffrey Lagarias, James Reeds, Margaret Wright, and Paul Wright. Convergence properties of the nelder–mead simplex method in low dimensions. *SIAM Journal on Optimization*, 9:112–147, 12 1998.
- [100] Navin Khaneja, Timo Reiss, Cindie Kehlet, Thomas Schulte-Herbrüggen, and Stefan J. Glaser. Optimal control of coupled spin dynamics: design of nmr pulse sequences by gradient ascent algorithms. *Journal of Magnetic Resonance*, 172(2):296–305, 2005.
- [101] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018.
- [102] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dhharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, February 2015.
- [103] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, January 2016.
- [104] R. E. Wengert. A simple automatic derivative evaluation program. *Commun. ACM*, 7(8):463–464, August 1964.
- [105] Atılım Günes Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. Automatic differentiation in machine learning: A survey. *J. Mach. Learn. Res.*, 18(1):5595–5637, January 2017.

- [106] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, and Skye Wanderman-Milne. JAX: composable transformations of Python+NumPy programs, 2018.
- [107] Hai-Jun Liao, Jin-Guo Liu, Lei Wang, and Tao Xiang. Differentiable programming tensor networks. *Phys. Rev. X*, 9:031041, Sep 2019.
- [108] Sandro Sorella and Luca Capriotti. Algorithmic differentiation and the calculation of forces by quantum monte carlo. *The Journal of Chemical Physics*, 133(23):234111, 2010.
- [109] Nelson Leung, Mohamed Abdelhafez, Jens Koch, and David Schuster. Speedup for quantum optimal control from automatic differentiation based on graphics processing units. *Phys. Rev. A*, 95:042318, Apr 2017.
- [110] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, Oct 1986.
- [111] Walter Baur and Volker Strassen. The complexity of partial derivatives. *Theoretical Computer Science*, 22(3):317–330, 1983.
- [112] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [113] Mike B. Giles. Collected matrix derivative results for forward and reverse mode algorithmic differentiation. In Christian H. Bischof, H. Martin Bücker, Paul Hovland, Uwe Naumann, and Jean Utke, editors, *Advances in Automatic Differentiation*, pages 35–44, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.
- [114] Ingo Rechenberg. *Evolutionstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. Stuttgart : Frommann-Holzboog, 1973.
- [115] Daan Wierstra, Tom Schaul, Tobias Glasmachers, Yi Sun, Jan Peters, and Jürgen Schmidhuber. Natural evolution strategies. *Journal of Machine Learning Research*, 15(27):949–980, 2014.
- [116] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. Evolution Strategies as a Scalable Alternative to Reinforcement Learning. *arXiv:1703.03864 [cs, stat]*, September 2017.

- [117] Chris Beeler, Uladzimir Yahorau, Rory Coles, Kyle Mills, Stephen Whitlam, and Isaac Tamblyn. Optimizing thermodynamic trajectories using evolutionary and gradient-based reinforcement learning, 2021.
- [118] Tianchen Zhao, Giuseppe Carleo, James Stokes, and Shravan Veerapaneni. Natural evolution strategies and variational monte carlo. *Machine Learning: Science and Technology*, 2(2):02LT01, Jan 2021.
- [119] Abhinav Anand, Matthias Degroote, and Alán Aspuru-Guzik. Natural evolutionary strategies for variational quantum computation. *Machine Learning: Science and Technology*, 2(4):045012, Jul 2021.
- [120] Christian Gross and Immanuel Bloch. Quantum simulations with ultracold atoms in optical lattices. *Science*, 357(6355):995–1001, 2017.
- [121] E. Knill, R. Laflamme, and G. J. Milburn. A scheme for efficient quantum computation with linear optics. *Nature*, 409(6816):46–52, Jan 2001.
- [122] Ady Stern. Non-abelian states of matter. *Nature*, 464(7286):187–193, Mar 2010.
- [123] Tudor Stanescu. *Introduction to Topological Quantum Matter Quantum Computation*. 12 2016.
- [124] Frank Wilczek. Quantum mechanics of fractional-spin particles. *Phys. Rev. Lett.*, 49:957–959, Oct 1982.
- [125] Adriano Barenco, Charles H. Bennett, Richard Cleve, David P. DiVincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John A. Smolin, and Harald Weinfurter. Elementary gates for quantum computation. *Phys. Rev. A*, 52:3457–3467, Nov 1995.
- [126] Michael H. Freedman, Michael Larsen, and Zhenghan Wang. A modular functor which is universal for quantum computation. *Communications in Mathematical Physics*, 227(3):605–622, Jun 2002.
- [127] Simon Trebst, Matthias Troyer, Zhenghan Wang, and Andreas W. W. Ludwig. A Short Introduction to Fibonacci Anyon Models. *Progress of Theoretical Physics Supplement*, 176:384–407, 06 2008.
- [128] J.K. Pachos. *Introduction to Topological Quantum Computation*. Cambridge University Press, 2012.
- [129] H. Bartolomei, M. Kumar, R. Bisognin, A. Marguerite, J.-M. Berroir, E. Bocquillon, B. Plaçais, A. Cavanna, Q. Dong, U. Gennser, Y. Jin, and G. Fève. Fractional statistics in anyon collisions. *Science*, 368(6487):173–177, 2020.

- [130] J. Nakamura, S. Liang, G. C. Gardner, and M. J. Manfra. Direct observation of anyonic braiding statistics. *Nature Physics*, 16(9):931–936, Sep 2020.
- [131] David J. Clarke, Jason Alicea, and Kirill Shtengel. Exotic non-abelian anyons from conventional fractional quantum hall states. *Nature Communications*, 4(1):1348, Jan 2013.
- [132] G.F. Bassani. *Ettore majorana scientific papers: On occasion of the centenary of his birth*. 01 2006.
- [133] M. Tinkham. *Introduction to Superconductivity*. International series in pure and applied physics. McGraw Hill, 1996.
- [134] A.L. Fetter and J.D. Walecka. *Quantum Theory of Many-particle Systems*. Dover Books on Physics. Dover Publications, 2003.
- [135] B.A. Bernevig and T.L. Hughes. *Topological Insulators and Topological Superconductors*. Princeton University Press, 2013.
- [136] Alexander Altland and Martin R. Zirnbauer. Nonstandard symmetry classes in mesoscopic normal-superconducting hybrid structures. *Phys. Rev. B*, 55:1142–1161, Jan 1997.
- [137] Alexei Kitaev. Periodic table for topological insulators and superconductors. *AIP Conference Proceedings*, 1134(1):22–30, 2009.
- [138] A F Andreev. Thermal conductivity of the intermediate state of superconductors. *Zh. Eksperim. i Teor. Fiz.*
- [139] Johan Nilsson, A. R. Akhmerov, and C. W. J. Beenakker. Splitting of a cooper pair by a pair of majorana bound states. *Physical Review Letters*, 101(12), Sep 2008.
- [140] Hao Zhang, Michiel W. A. de Moor, Jouri D. S. Bommer, Di Xu, Guanzhong Wang, Nick van Loo, Chun-Xiao Liu, Sasa Gazibegovic, John A. Logan, Diana Car, Roy L. M. Op het Veld, Petrus J. van Veldhoven, Sebastian Koelling, Marcel A. Verheijen, Mihir Pendharkar, Daniel J. Pennachio, Borzoyeh Shojaei, Joon Sue Lee, Chris J. Palmstrøm, Erik P. A. M. Bakkers, S. Das Sarma, and Leo P. Kouwenhoven. Large zero-bias peaks in insb-al hybrid semiconductor-superconductor nanowire devices, 2021.
- [141] Y. Bychkov and E. Rashba. Oscillatory effects and the magnetic susceptibility of carriers in inversion layers. *Journal of Physics C: Solid State Physics*, 17:6039–6045, 1984.

- [142] Jason Alicea. Majorana fermions in a tunable semiconductor device. *Phys. Rev. B*, 81:125318, Mar 2010.
- [143] M.-T. Rieder, G. Kells, M. Duckheim, D. Meidan, and P. W. Brouwer. Endstates in multichannel spinless p -wave superconducting wires. *Phys. Rev. B*, 86:125423, Sep 2012.
- [144] M. S. Scheurer and A. Shnirman. Nonadiabatic processes in majorana qubit systems. *Phys. Rev. B*, 88:064515, Aug 2013.
- [145] H.-J. Kwon, K. Sengupta, and V. M. Yakovenko. Fractional ac josephson effect in p - and d -wave superconductors. *The European Physical Journal B - Condensed Matter*, 37(3):349–361, Feb 2003.
- [146] Diego Rainis and Daniel Loss. Majorana qubit decoherence by quasiparticle poisoning. *Phys. Rev. B*, 85:174533, May 2012.
- [147] Daniel Frombach and Patrik Recher. Quasiparticle poisoning effects on the dynamics of topological josephson junctions. *Phys. Rev. B*, 101:115304, Mar 2020.
- [148] D. I. Pikulin and Yuli V. Nazarov. Phenomenology and dynamics of a majorana josephson junction. *Physical Review B*, 86(14), Oct 2012.
- [149] Dominique Laroche, Daniël Bouman, David J. van Woerkom, Alex Proutski, Chaitanya Murthy, Dmitry I. Pikulin, Chetan Nayak, Ruben J. J. van Gulik, Jesper Nygård, Peter Krogstrup, Leo P. Kouwenhoven, and Attila Geresdi. Observation of the 4π -periodic josephson effect in indium arsenide nanowires. *Nature Communications*, 10(1):245, Jan 2019.
- [150] Chun-Xiao Liu, Jay D. Sau, and S. Das Sarma. Distinguishing topological majorana bound states from trivial andreev bound states: Proposed tests through differential tunneling conductance spectroscopy. *Physical Review B*, 97(21), Jun 2018.
- [151] K. T. Law, Patrick A. Lee, and T. K. Ng. Majorana fermion induced resonant andreev reflection. *Phys. Rev. Lett.*, 103:237001, Dec 2009.
- [152] G. Kells, D. Meidan, and P. W. Brouwer. Near-zero-energy end states in topologically trivial spin-orbit coupled superconducting nanowires with a smooth confinement. *Phys. Rev. B*, 86:100503, Sep 2012.
- [153] Jason Alicea, Yuval Oreg, Gil Refael, Felix von Oppen, and Matthew P. A. Fisher. Non-abelian statistics and topological quantum information processing in 1d wire networks. *Nature Physics*, 7(5):412–417, Feb 2011.

- [154] Bela Bauer, Torsten Karzig, Ryan Mishmash, Andrey Antipov, and Jason Alicea. Dynamics of majorana-based qubits operated with an array of tunable gates. *SciPost Physics*, 5(1), Jul 2018.
- [155] Michael Victor Berry. Quantum phase corrections from adiabatic iteration. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 414(1846):31–46, 1987.
- [156] Jun John Sakurai. *Modern quantum mechanics; rev. ed.* Addison-Wesley, Reading, MA, 1994.
- [157] Jay D. Sau, David J. Clarke, and Sumanta Tewari. Controlling non-abelian statistics of majorana fermions in semiconductor nanowires. *Phys. Rev. B*, 84:094505, Sep 2011.
- [158] B van Heck, A R Akhmerov, F Hassler, M Burrello, and C W J Beenakker. Coulomb-assisted braiding of majorana fermions in a josephson junction array. *New Journal of Physics*, 14(3):035019, mar 2012.
- [159] M. Burrello, B. van Heck, and A. R. Akhmerov. Braiding of non-abelian anyons using pairwise interactions. *Phys. Rev. A*, 87:022343, Feb 2013.
- [160] Parsa Bonderson. Measurement-only topological quantum computation via tunable interactions. *Phys. Rev. B*, 87:035113, Jan 2013.
- [161] Parsa Bonderson, Michael Freedman, and Chetan Nayak. Measurement-only topological quantum computation. *Phys. Rev. Lett.*, 101:010501, Jun 2008.
- [162] Torsten Karzig, Falko Pientka, Gil Refael, and Felix von Oppen. Shortcuts to non-abelian braiding. *Phys. Rev. B*, 91:201102, May 2015.
- [163] Chuanchang Zeng, Girish Sharma, Tudor D. Stanescu, and Sumanta Tewari. Feasibility of measurement-based braiding in the quasi-majorana regime of semiconductor-superconductor heterostructures. *Phys. Rev. B*, 102:205101, Nov 2020.
- [164] A. Conlon, D. Pellegrino, J. K. Slingerland, S. Dooley, and G. Kells. Error generation and propagation in majorana-based topological qubits. *Phys. Rev. B*, 100:134307, Oct 2019.
- [165] Thomas Fösel, Petru Tighineanu, Talitha Weiss, and Florian Marquardt. Reinforcement learning with neural networks for quantum feedback. *Phys. Rev. X*, 8:031084, Sep 2018.
- [166] Riccardo Porotti, Antoine Essig, Benjamin Huard, and Florian Marquardt. Deep reinforcement learning for quantum state preparation with weak nonlinear measurements, 2021.

- [167] Pierpaolo Sgroi, G. Massimo Palma, and Mauro Paternostro. Reinforcement learning approach to nonequilibrium quantum thermodynamics. *Phys. Rev. Lett.*, 126:020601, Jan 2021.
- [168] Shiyu Liang and R. Srikant. Why deep neural networks for function approximation? In *ICLR*, 2017.
- [169] Christina Knapp, Michael Zaletel, Dong E. Liu, Meng Cheng, Parsa Bonderson, and Chetan Nayak. The nature and correction of diabatic errors in anyon braiding. *Phys. Rev. X*, 6:041003, Oct 2016.
- [170] Piet W. Brouwer, Mathias Duckheim, Alessandro Romito, and Felix von Oppen. Probability distribution of majorana end-state energies in disordered wires. *Phys. Rev. Lett.*, 107:196804, Nov 2011.
- [171] Piet W. Brouwer, Mathias Duckheim, Alessandro Romito, and Felix von Oppen. Topological superconducting phases in disordered quantum wires with strong spin-orbit coupling. *Phys. Rev. B*, 84:144526, Oct 2011.
- [172] Wade DeGottardi, Diptiman Sen, and Smitha Vishveshwara. Majorana fermions in superconducting 1d systems having periodic, quasiperiodic, and disordered potentials. *Phys. Rev. Lett.*, 110:146404, Apr 2013.
- [173] Torsten Karzig, Gil Refael, and Felix von Oppen. Boosting majorana zero modes. *Phys. Rev. X*, 3:041017, Nov 2013.
- [174] Xi Chen, E. Torrontegui, Dionisis Stefanatos, Jr-Shin Li, and J. G. Muga. Optimal trajectories for efficient atomic transport without final excitation. *Physical Review A*, 84(4), Oct 2011.
- [175] Lucas T. Brady, Christopher L. Baldwin, Aniruddha Bapat, Yaroslav Kharkov, and Alexey V. Gorshkov. Optimal protocols in quantum annealing and quantum approximate optimization algorithm problems. *Physical Review Letters*, 126(7), Feb 2021.
- [176] E. M. Stoudenmire, Jason Alicea, Oleg A. Starykh, and Matthew P.A. Fisher. Interaction effects in topological superconducting wires supporting majorana fermions. *Phys. Rev. B*, 84:014503, Jul 2011.
- [177] S. Gangadharaiah, B. Braunecker, P. Simon, and D. Loss. Majorana edge states in interacting one-dimensional systems. *Phys. Rev. Lett.*, 107:036801, Jul 2011.
- [178] G. Yang and D. E. Feldman. Exact zero modes and decoherence in systems of interacting majorana fermions. *Phys. Rev. B*, 89:035136, Jan 2014.

- [179] H. T. Ng. Decoherence of interacting majorana modes. *Scientific Reports*, 5:12530 EP–, 07 2015.
- [180] G. Kells. Many-body majorana operators and the equivalence of parity sectors. *Phys. Rev. B*, 92:081401, Aug 2015.
- [181] M. McGinley, J. Knolle, and A. Nunnenkamp. Robustness of majorana edge modes and topological order: Exact results for the symmetric interacting kitaev chain with disorder. *Phys. Rev. B*, 96:241113, Dec 2017.
- [182] G. Kells, N. Moran, and D. Meidan. Localization enhanced and degraded topological order in interacting p -wave wires. *Phys. Rev. B*, 97:085425, Feb 2018.
- [183] Henrik Bruus and Karsten Flensberg. *Many-body quantum field theory in condensed matter physics: an introduction*. Oxford University Press, United Kingdom, 2003.
- [184] R.D. Mattuck. *A Guide to Feynman Diagrams in the Many-Body Problem: Second Edition*. Dover Books on Physics. Dover Publications, 2012.
- [185] S. Bravyi, M. B. Hastings, and S. Michalakis. Topological quantum order: Stability under local perturbations. *Journal of Mathematical Physics*, 51(9):093512, 2010.
- [186] M. B. Hastings and Xiao-Gang Wen. Quasiadiabatic continuation of quantum states: The stability of topological ground-state degeneracy and emergent gauge invariance. *Phys. Rev. B*, 72:045141, Jul 2005.
- [187] E.H. Lieb and D.W. Robinson. The finite group velocity of quantum spin systems. *Commun. Math. Phys.*, 28:251–257, 1972.
- [188] G. Kells. Multiparticle content of majorana zero modes in the interacting p -wave wire. *Phys. Rev. B*, 92:155434, Oct 2015.
- [189] D. Pellegrino. Strong zero modes in generalized spin chains and topological quantum computation. *Maynooth University, National University of Ireland Maynooth*, PhD Thesis, 2019.
- [190] H. Bethe. Zur theorie der metalle. *Zeitschrift für Physik*, 71(3):205–226, Mar 1931.
- [191] Alexander Weiß and Holger Fehske. *Exact Diagonalization Techniques*, pages 529–544. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [192] Iman Mahyaeh and Eddy Ardonne. Study of the phase diagram of the kitaev-hubbard chain. *Phys. Rev. B*, 101:085125, Feb 2020.

- [193] Paul Fendley. Strong zero modes and eigenstate phase transitions in the XYZ/interacting majorana chain. *Journal of Physics A: Mathematical and Theoretical*, 49(30):30LT01, jun 2016.
- [194] Dominic V. Else, Paul Fendley, Jack Kemp, and Chetan Nayak. Prethermal strong zero modes and topological qubits. *Phys. Rev. X*, 7:041062, Dec 2017.
- [195] N. Moran, D. Pellegrino, J. K. Slingerland, and G. Kells. Parafermionic clock models and quantum resonance. *Physical Review B*, 95(23), Jun 2017.
- [196] J. Eisert, M. Cramer, and M. B. Plenio. Colloquium: Area laws for the entanglement entropy. *Reviews of Modern Physics*, 82(1):277–306, Feb 2010.
- [197] David A. Huse, Rahul Nandkishore, Vadim Oganesyan, Arijeet Pal, and S. L. Sondhi. Localization-protected quantum order. *Phys. Rev. B*, 88:014206, Jul 2013.
- [198] S. Nuly, J. Vala, D. Meidan, and G. Kells. Constrained thermalization and topological superconductivity. *Phys. Rev. B*, 102:054508, Aug 2020.
- [199] J. Majer, J. M. Chow, J. M. Gambetta, Jens Koch, B. R. Johnson, J. A. Schreier, L. Frunzio, D. I. Schuster, A. A. Houck, A. Wallraff, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Coupling superconducting qubits via a cavity bus. *Nature*, 449(7161):443–447, Sep 2007.
- [200] Govind P Agrawal. *Fiber-optic communication systems*, volume 222. John Wiley & Sons, 2012.
- [201] Guilherme B. Xavier and Gustavo Lima. Quantum information processing with space-division multiplexing optical fibres. *Communications Physics*, 3(1):9, Jan 2020.
- [202] Sougato Bose. Quantum communication through an unmodulated spin chain. *Phys. Rev. Lett.*, 91:207901, Nov 2003.
- [203] P. Cappellaro, C. Ramanathan, and D. G. Cory. Dynamics and control of a quasi-one-dimensional spin system. *Phys. Rev. A*, 76:032317, Sep 2007.
- [204] Vinitha Balachandran and Jiangbin Gong. Adiabatic quantum transport in a spin chain with a moving potential. *Phys. Rev. A*, 77:012303, Jan 2008.
- [205] K Eckert, O Romero-Isart, and A Sanpera. Efficient quantum state transfer in spin chains via adiabatic passage. *New Journal of Physics*, 9(5):155–155, may 2007.
- [206] Michael Murphy, Simone Montangero, Vittorio Giovannetti, and Tommaso Calarco. Communication at the quantum speed limit along a spin chain. *Phys. Rev. A*, 82:022318, Aug 2010.

- [207] Luuk Coopmans, Anthony Kiely, Gabriele De Chiara, and Steve Campbell. Optimal control in disordered quantum systems, 2022.
- [208] W. Heisenberg. Mehrkörperproblem und resonanz in der quantenmechanik. *Zeitschrift für Physik*, 38(6):411–426, Jun 1926.
- [209] J. Hubbard. Electron correlations in narrow energy bands. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 276:238 – 257, 1963.
- [210] M. Ahmed. Confined magnon transport in low dimensional ferromagnetic structures. *Ph.D Thesis, RMIT University, Melbourne, Australia*, 2017.
- [211] M. H. Ahmed and A. D. Greentree. Guided magnon transport in spin chains: Transport speed and correcting for disorder. *Phys. Rev. A*, 91:022306, Feb 2015.
- [212] Melissa I. Makin, Jared H. Cole, Charles D. Hill, and Andrew D. Greentree. Spin guides and spin splitters: Waveguide analogies in one-dimensional spin chains. *Phys. Rev. Lett.*, 108:017207, Jan 2012.
- [213] E. Torrontegui, S. Ibáñez, Xi Chen, A. Ruschhaupt, D. Guéry-Odelin, and J. G. Muga. Fast atomic transport without vibrational heating. *Phys. Rev. A*, 83:013415, Jan 2011.
- [214] A K Dhara and S V Lawande. Feynman propagator for time-dependent lagrangians possessing an invariant quadratic in momentum. *Journal of Physics A: Mathematical and General*, 17(12):2423–2431, aug 1984.
- [215] Luca Innocenti, Gabriele De Chiara, Mauro Paternostro, and Ricardo Puebla. Ultrafast critical ground state preparation via bang–bang protocols. *New Journal of Physics*, 22(9):093050, Sep 2020.
- [216] Jeffrey M. Epstein and K. Birgitta Whaley. Quantum speed limits for quantum-information-processing tasks. *Physical Review A*, 95(4), Apr 2017.
- [217] Daniel Burgarth and Sougato Bose. Perfect quantum state transfer with randomly coupled quantum chains. *New Journal of Physics*, 7:135–135, may 2005.
- [218] Gabriele De Chiara, Davide Rossini, Simone Montangero, and Rosario Fazio. From perfect to fractal transmission in spin chains. *Phys. Rev. A*, 72:012323, Jul 2005.
- [219] P. W. Anderson. Absence of diffusion in certain random lattices. *Phys. Rev.*, 109:1492–1505, Mar 1958.
- [220] Christian K. Burrell and Tobias J. Osborne. Bounds on the speed of information propagation in disordered quantum spin chains. *Phys. Rev. Lett.*, 99:167201, Oct 2007.

- [221] Christian K. Burrell, Jens Eisert, and Tobias J. Osborne. Information propagation through quantum chains with fluctuating disorder. *Phys. Rev. A*, 80:052319, Nov 2009.
- [222] Mogens Dalgaard, Felix Motzoi, and Jacob Sherson. Predicting quantum dynamical cost landscapes with deep learning, 2021.
- [223] Alexandre G.R. Day, Marin Bukov, Phillip Weinberg, Pankaj Mehta, and Dries Sels. Glassy phase of optimal quantum control. *Physical Review Letters*, 122(2), Jan 2019.
- [224] James Nelson, Luuk Coopmans, Graham Kells, and Stefano Sanvito. Data-driven time propagation of quantum systems with neural networks, 2022.
- [225] P. Fendley. Strong zero modes and eigenstate phase transitions in the XYZ/interacting majorana chain. *Journal of Physics A: Mathematical and Theoretical*, 49(30):30LT01, jun 2016.
- [226] D Pellegrino, G Kells, N Moran, and J K Slingerland. Constructing edge zero modes through domain wall angle conservation. *Journal of Physics A: Mathematical and Theoretical*, 53(9):095006, feb 2020.
- [227] Nicola Pancotti, Matteo Scandi, Mark T. Mitchison, and Martí Perarnau-Llobet. Speed-ups to isothermality: Enhanced quantum thermal machines through control of the system-bath coupling. *Phys. Rev. X*, 10:031015, Jul 2020.
- [228] Harry J.D. Miller, Giacomo Guarnieri, Mark T. Mitchison, and John Goold. Quantum fluctuations hinder finite-time information erasure near the landauer limit. *Physical Review Letters*, 125(16), Oct 2020.
- [229] Tan Van Vu and Keiji Saito. Finite-time quantum landauer principle and quantum coherence, 2021.
- [230] Ilia Khait, Juan Carrasquilla, and Dvira Segal. Optimal control of quantum thermal machines using machine learning, 2021.
- [231] Brian Skinner, Jonathan Ruhman, and Adam Nahum. Measurement-induced phase transitions in the dynamics of entanglement. *Physical Review X*, 9(3), Jul 2019.
- [232] Xiangyu Cao, Antoine Tilloy, and Andrea De Luca. Entanglement in a fermion chain under continuous monitoring. *SciPost Physics*, 7(2), Aug 2019.
- [233] Thomas Fösel, Petru Tighineanu, Talitha Weiss, and Florian Marquardt. Reinforcement learning with neural networks for quantum feedback. *Physical Review X*, 8(3), Sep 2018.

- [234] Frank Schäfer, Pavel Sekatski, Martin Koppenhöfer, Christoph Bruder, and Michal Kloc. Control of stochastic quantum dynamics by differentiable programming. *Machine Learning: Science and Technology*, 2(3):035004, Apr 2021.
- [235] Pierpaolo Sgroi, G. Massimo Palma, and Mauro Paternostro. Reinforcement learning approach to nonequilibrium quantum thermodynamics. *Physical Review Letters*, 126(2), Jan 2021.
- [236] D. Dong and I.R. Petersen. Quantum control theory and applications: a survey. *IET Control Theory Applications*, 4(12):2651–2671, Dec 2010.
- [237] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014.
- [238] Lan Lan, Lei You, Zeyang Zhang, Zhiwei Fan, Weiling Zhao, Nianyin Zeng, Yidong Chen, and Xiaobo Zhou. Generative adversarial networks and its applications in biomedical informatics. *Frontiers in Public Health*, 8:164, 2020.
- [239] P. Jordan and E. Wigner. Über das paulische äquivalenzverbot. *Zeitschrift für Physik*, 47(9):631–651, Sep 1928.
- [240] Eduardo Fradkin. Jordan-wigner transformation for quantum-spin systems in two dimensions and fractional statistics. *Phys. Rev. Lett.*, 63:322–325, Jul 1989.
- [241] Alexei Kitaev. Anyons in an exactly solved model and beyond. *Annals of Physics*, 321(1):2 – 111, 2006. January Special Issue.
- [242] Elliott Lieb, Theodore Schultz, and Daniel Mattis. Two soluble models of an antiferromagnetic chain. *Annals of Physics*, 16(3):407 – 466, 1961.
- [243] F. Verstraete, V. Murg, and J.I. Cirac. Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems. *Advances in Physics*, 57(2):143–224, 2008.
- [244] F Verstraete and J I. Cirac. Renormalization algorithms for quantum-many body systems in two and higher dimensions. 08 2004.
- [245] Gregory M. Crosswhite and Dave Bacon. Finite automata for caching in matrix product algorithms. *Phys. Rev. A*, 78:012356, Jul 2008.