



**SISSA**

---

---

**Adiabatic approaches to non-equilibrium systems:  
Topology, Optimization, and Learning  
- PhD Thesis -**

*Supervisor:* **Giuseppe E. Santoro**

*Candidate:* **Matteo M. Wauters**

September 2020

*SISSA, Trieste*

---

*“È la rievocazione storica di una disfida tra trentini e feltrini, perché ballava il tavolo”*

A wise man, 2019

## Abstract

Non-equilibrium phenomena in quantum physics are ubiquitous in current research because of their extremely rich phenomenology. Among them, adiabatic processes are particularly interesting because they connect two of the cornerstone of modern condensed matter theory: topological phases of matter and quantum information. Indeed with adiabatic periodic driving, it is possible to engineer exotic non-equilibrium phases with tunable topological properties, which are of great interest for both fundamental research and future quantum technologies. Adiabatic processes can also be used to implement optimization processes on quantum hardware and lie at the basis of hybrid quantum-classical algorithms that are extremely promising for near-term quantum devices. In this thesis, I focus on these two aspects of adiabatic dynamics in quantum systems.

In the first part, I investigate the robustness of topological phases, arising in periodically driven systems, with respect to the driving protocol and disorder. In particular, I show that quantized transport and Anderson localization coexist in one-dimensional systems displaying Thouless pumping, because of a delocalization-localization transition in the Floquet spectrum. This transition is linked to the topological nature of the adiabatic driving and disappears if the phase is trivial in the clean limit.

In the second part, I study one of the most popular hybrid optimization method, the quantum approximate optimization algorithm (QAOA). I show that QAOA can tackle the first-order phase transition arising in the infinite range p-spin model with polynomial resources, in stark contrast with adiabatic quantum computation, which requires exponentially long evolution time to reach similar performances. Finally, I present an approach to QAOA based on reinforcement learning (RL). Interestingly, the RL agent automatically adopts strategies that converge towards optimal adiabatic schedules and that can be easily transferred between systems with different sizes, even in the presence of disorder.

---

## Acknowledgements

Too many are the people I should thank, for making these four years so enjoyable and meaningful to me. Here is a list in quasi-random order: I will forget someone for sure, but I am an emissary of chaos, so they should not be surprised, neither feel offended.

First of all Giuseppe, who taught me a lot during my Ph.D., and corrected patiently all the typos I wrote.

Alberto Parola and Giuliano Benenti, my former professors at the university; without them I would not have come to SISSA.

Glen (the smart guy) and Ema (the RL guy), who followed me in my last year of Ph.D.; I really enjoyed working with them.

Nadir, who shared with me five years at the university in Como and 4 years of Phd and fencing in Trieste. It is a pity that he has been lured to the dark side, he would have been an extraordinary physicist.

Among all the friends and colleagues, a special thank for Andrea, my flatmate in the last two years, who spend with me two months of lockdown playing DBZBT2 (only for the connoisseurs), and for Giulia (pat pat) who quarreled with me on anything. Thank Matteo x2 , Lorenzo, Yusuf, Veronica, Davide, Federica who shared with me this experience.

A special place is reserved for the Cipolle Verdi, our band: Matteo x2, Nadir, Federica, Federico, Suzie, Angelo Adrianone, Bruno, Lucia. I really enjoyed playing music with them.

Uriel and all the SISSA choir. Ours is the Lamb!

The sound of swords clashing has accompanied me for all these years: Moreno, Fabio, and all the wonderful people of Trieste Scherma Storica. Hitting hard with a two-handed sword really helped me to relax.

Marco (although I've never called him so) and Erica (who still has my brain), two old friends that I know I can always count upon.

Sofia, a new friend who tried to teach me to be more spontaneous, with various degrees of success.

My family, who stood me for 24 years before I moved out, but mysteriously still want to see me.

Last but not least Gaia, who made the last two years wonderful.

## Preface

This thesis is based on the following publications and preprints:

- 1) “Quantization of the Hall conductivity in the Harper-Hofstadter model” M.M. Wauters, G.E. Santoro PRB 98 (20), 205112 (2018), discussed in chapter 3.
- 2) “Localization, Topology and Quantized Transport in Disordered Floquet Systems” M.M. Wauters, A. Russomanno, R. Citro, G.E.Santoro, L. Privitera PRL **123** (26), 266601 (2019), discussed in chapter 4.
- 3) “Polynomial scaling of QAOA for ground-state preparation of the fully-connected p-spin ferromagnet” M.M. Wauters, G.B. Mbeng, G.E. Santoro arXiv preprint arXiv:2003.07419, discussed in chapter 5.
- 4) “Reinforcement Learning for Quantum Control” M.M. Wauters, E. Panizon, G.B. Mbeng, G.E. Santoro Phys. Rev. Research 2, 033446 (2020), discussed in chapter 6.

Not discussed in the thesis is an early work on Quantum Annealing done at the beginning of my Ph.D. at SISSA:

- 5) “Direct comparison of quantum and simulated annealing on a fully-connected Ising ferromagnet” M.M. Wauters, R. Fazio, H. Nishimori, G.E. Santoro Phys. Rev. A 96 (2), 022326 (2017)



# Contents

Abstract . . . . .	iii
Acknowledgements . . . . .	iv
Preface . . . . .	v
<b>1. Introduction</b>	<b>1</b>
1.1. Adiabatic Quantum Pumping . . . . .	2
1.2. Quantum Optimization . . . . .	5
<b>I. Quantum Pumping</b>	<b>11</b>
<b>2. Quantum pumping and Floquet Theory</b>	<b>13</b>
2.1. Thouless pumping . . . . .	13
2.2. Floquet Theory . . . . .	16
2.2.1. Pumped charge and winding of quasi-energies . . . . .	18
<b>3. Non-adiabatic effects</b>	<b>21</b>
3.1. Introduction . . . . .	22
3.2. Model . . . . .	24
3.3. Robustness of Hall quantization to non adiabatic effects . . . . .	28
3.3.1. Pumping of Floquet states. . . . .	30
3.3.2. Effect of the occupation factors: sudden switch-on. . . . .	33
3.3.3. Effect of the occupation factors: continuous switch-on. . . . .	35
3.4. Conclusions . . . . .	39
<b>4. Pumping in disordered systems</b>	<b>43</b>
4.1. Introduction . . . . .	44
4.2. Model . . . . .	45
4.3. Results . . . . .	46
4.4. Winding of quasi-energies . . . . .	52
4.5. Frequency dependence of the localization length . . . . .	53
4.6. Conclusions . . . . .	53
<b>II. Quantum Optimization</b>	<b>55</b>
<b>5. QAOA on fully connected models</b>	<b>57</b>
5.1. Introduction . . . . .	58
5.2. Model and QAOA algorithm . . . . .	59
5.3. Exact ground state with $P = 1$ . . . . .	61

5.4. Numerical results for $P > 1$ . . . . .	62
5.5. Conclusions . . . . .	67
<b>6. Reinforcement Learning for quantum optimization</b>	<b>71</b>
6.1. Introduction . . . . .	72
6.2. Basis of RL . . . . .	74
6.3. RL assisted QAOA . . . . .	76
6.4. Results . . . . .	77
6.4.1. Ising model with uniform couplings . . . . .	77
6.4.2. Ising model with random couplings . . . . .	80
6.4.3. LMG model . . . . .	82
6.5. Conclusions . . . . .	84
<b>7. Conclusions</b>	<b>87</b>
<b>A. Details for chapter 3</b>	<b>93</b>
A.1. Adiabatic expansion of Floquet eigenstates . . . . .	93
A.2. Dependence on $k_x$ of Floquet quasi-energies and occupations . . . . .	95
<b>B. Details for chapter 5</b>	<b>97</b>
B.1. exact ground state preparation for $P = 1$ . . . . .	97
B.1.1. $P = 1$ : requirements on $\beta$ . . . . .	97
B.1.2. $P = 1$ and $p$ odd: requirements on $\gamma$ . . . . .	98
B.1.3. $P = 1$ and $p$ even: requirements on $\gamma$ . . . . .	99
B.1.4. Symmetries of the parameter space for general $P$ , $N$ and $p$ . . . . .	101
<b>C. Details for chapter 6</b>	<b>103</b>
C.1. Introduction to RL: Key concepts and notation . . . . .	103
C.2. Introduction to policy optimization . . . . .	109
C.3. Proximal Policy Optimization . . . . .	110
C.4. Choice of policy and training parameters . . . . .	112



# 1. Introduction

In this thesis, I discuss two important physical processes related to non-equilibrium quantum systems with an adiabatic time-dependent driving. The first one, studied in part I, deals with adiabatic quantum transport, a cornerstone of topological phases of matter arising in non-equilibrium systems. In part II, I analyze how quantum control and reinforcement learning can be applied to optimization processes which can be regarded as offsprings of the Quantum Annealing idea.

In quantum mechanics, adiabatic processes refer typically to systems described by a time-dependent Hamiltonian  $\hat{H}(t)$  which is slowly driven in time. The physics behind such phenomena is well explained by the *adiabatic theorem* [1], that can be, loosely speaking, summarized in the following way: the instantaneous time-dependent energy eigenstates  $|\phi_n(t)\rangle$  are also approximate eigenstates of the unitary evolution operator, provided the driving is *slow enough*. This means that if we prepare the system in any Hamiltonian eigenstate  $|\phi_n(t_0)\rangle$  at time  $t_0$  and evolve it according to the Schrödinger equation, the time-evolved state will remain close to the corresponding instantaneous eigenstate  $|\phi_n(t)\rangle$ :

$$\hat{U}(t, t_0) |\phi_n(t_0)\rangle \simeq |\phi_n(t)\rangle \quad (1.1)$$

where  $\hat{U}(t, t_0)$  is the evolution operator between times  $t_0$  and  $t$ . The exact meaning of “slow enough” in the theorem’s hypothesis is a crucial information, and also the bottleneck of adiabatic dynamics. Intuitively it means that the time-scale  $\tau$  at which the Hamiltonian parameters change must be much longer than the time-scale associated with the microscopic dynamics. Equivalently the rate at which energy is absorbed from the driving has to be small enough not to excite the system, assuming it is prepared in its ground state. Quantitatively, this request roughly becomes

$$\frac{\hbar}{\tau} \ll \Delta \quad (1.2)$$

where  $\Delta$  is the energy gap between the ground and the first excited state.

But why are adiabatic processes interesting? First of all, the possibility of driving a system, while keeping it close to the instantaneous eigenstate of some time-dependent Hamiltonian, is relevant for transport phenomena and quantum engines, analogously to reversible transformations in classical thermodynamics. In particular *periodic drivings*, which are interesting on their own sake also outside the adiabatic limit [2–8], can lead to out-of-equilibrium *topological* phases of matter with no static analogue [4, 9–14], of which Thouless pumping [9, 15] is a paradigmatic example. The high level of control guaranteed by laser field in optical lattices experiments [16] and photonic crystals [17] have allowed, in the last 15 years, the engineering and the experimental investigation of periodically driven systems, boosting even more the interest in these topics [18–25]. The experimental availability of adiabatic topological processes, however, raises several important questions about their robustness. Since perfect

adiabaticity is impossible, how does the finite driving frequency influence the physical observable related to the topological invariant? And how much a low dimensional non-equilibrium topological process is robust against disorder? These questions will be the main motivations behind part I of the present thesis.

A second interesting application of adiabatic processes is the problem of “ground state preparation”, which lies at the heart of many condensed matter and quantum information problems [26, 27]. The adiabatic theorem comes in as a tool to drive a quantum system from a trivial initial state (usually a product state) to the target ground state one is aiming at constructing, exploiting the time-evolution ruled by a time-dependent Hamiltonian that interpolates between these two limits. This process, known as Quantum Annealing [28, 29] or Adiabatic Quantum Computation [30], is the simplest algorithm used for quantum optimization and it has been implemented experimentally on several platforms [31–33]. Since it relies on the adiabatic theorem, it is limited by the presence of small energy gaps in the spectrum, in particular when the system crosses a phase transition, where the energy gap goes to zero in the thermodynamic limit [34]. There are many techniques developed which try to improve the annealing performance when passing through a second-order — or even worse, a first-order — quantum phase transitions. These techniques include: schedule optimization [35, 36], exploiting dissipation [37, 38] and the addition of extra driving terms, either non-stoquastic [39]<sup>1</sup>, or superadiabatic [40]. Characterizing the efficiency of such processes is a central issue among the efforts in current research towards a fully scalable and universal quantum computer [41].

Recently, hybrid algorithms [42] gained momentum. Their core idea is to integrate classical and quantum tools in order to exploit the advantages of both: a quantum hardware undergoes a unitary evolution — inaccessible to classical computers — driven by a set of gates, and a classical algorithm is used to optimize the parameters entering in the quantum hardware, usually by means of variational methods. The Variational Quantum Eigensolver (VQE) [43] and the Quantum Approximate Optimization Algorithm (QAOA) [44] are the most popular of such methods and promising candidates for near-term quantum technologies [45, 46]. In part II I investigate the performance of QAOA in fully-connected ferromagnets and propose a scheme where Reinforcement Learning (RL) is used to stabilize and increase the robustness of QAOA.

In the remaining part of this chapter, I will give a more detailed introduction to the topics discussed in this thesis.

## 1.1. Adiabatic Quantum Pumping

Topological phases of matter are one of the cornerstone of modern condensed matter theory [47–49]. The first breakthrough in this field was the discovery of Integer Quantum Hall effect (IQHE) [50–53], which describes the quantization of the conductivity in a two-dimensional electron gas subjected to a magnetic field: if the Fermi energy lies in a spectral

---

<sup>1</sup>Loosely speaking, “non-stoquastic” means that there are off-diagonal Hamiltonian terms with a sign which provokes the fermion-sign problem in quantum Monte Carlo approaches, hence potentially hard to simulate classically

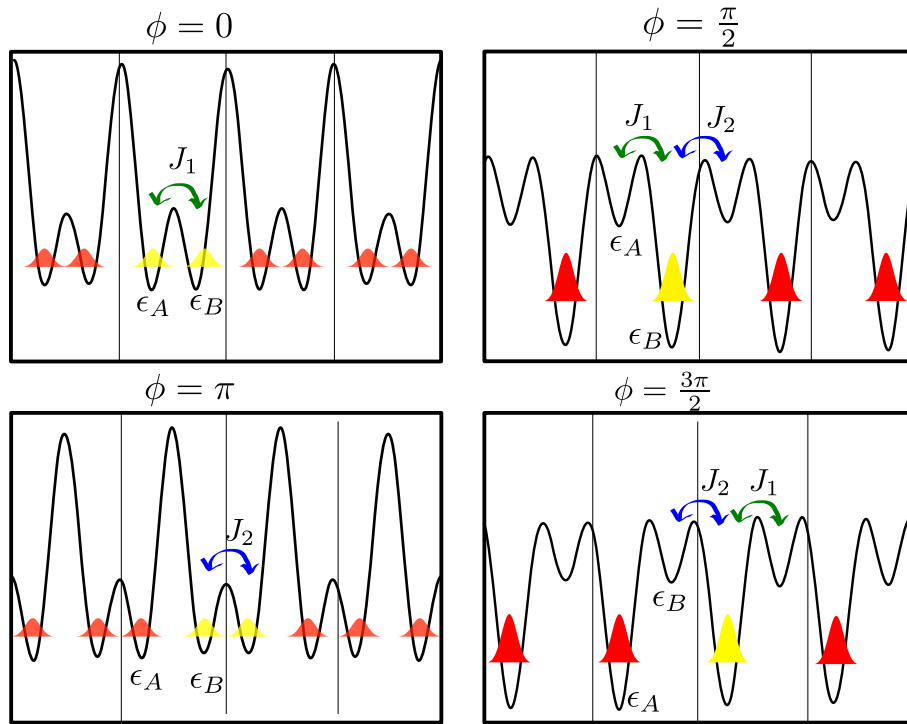


Figure 1.1.: Thouless pump scheme: atoms in a one dimensional bipartite optical lattice, generated by two laser beam with different wavelength  $\lambda_1 = 2\lambda_2$ . As the phase difference  $\phi$  between the two beams is shifted, also the on-site energies and the hopping amplitudes change, leading to a quantized motion of one particle per unit cell each period of the driving.

gap, the Hall conductivity  $\sigma_H$  is equal to an integer multiple  $n$  of a universal constant,  $\sigma_H = ne^2/h$ . The integer  $n$  is, in modern terminology, the Chern number [15, 54–56] of the occupied bands. The Chern number is a *topological invariant*, which means that it cannot change smoothly as the system's parameters are modified. Hence the topological phase characterized by a given Chern number is intrinsically robust against weak perturbations, such as disorder or interactions. Topological invariants are now widely used to classify many novel phases of matter, where the Landau-Ginzburg approach of a local order parameter fails, as in topological insulators [57–61], topological superconductor [59, 62] and Weyl semi-metals [63] among others.

Another fundamental step towards the understanding of the role of topology in physical phenomena is adiabatic quantum pumping [9, 15], or Thouless pumping for short, which realizes a non-equilibrium topological phase in a periodically driven system. It consists of the quantization of particle transport in a one-dimensional band insulator, where the Hamiltonian parameters are driven *adiabatically* in time along a closed path. If the path has a nontrivial topology, i.e. it winds around a gap-closing singularity in the spectrum (a metallicity point), an integer number of particles is pumped across the edges of the system at each period. Again, this integer number is the topological invariant Chern number of the filled bands. A sketch of the mechanism behind Thouless pumping is presented in Fig. 1.1. Since its formulation, Thouless pump has been very popular, because it connects several fundamental aspects of

topological phases of matter, from quantum Hall effect [64–66], to the modern theory of polarization [67–69] and Floquet engineering [6, 70, 71]. These non-equilibrium topological phases of matter realized through a time-periodic driving are known as *Floquet Topological Insulators* and are the subject of many recent research efforts [4, 7, 11–14, 72].

Adiabatic quantum pumping is also the subject of many recent experimental investigations [21–23, 25, 73, 74], which raises several questions concerning the robustness of the phenomenon.

First of all, any realistic implementation of a periodically driven system needs a finite frequency  $\omega$ : a perfect adiabatic limit  $\omega \rightarrow 0$  is impossible to achieve. The exact quantization of the pumped charge, however, holds only in the  $\omega \rightarrow 0$  limit. Hence, to properly interpret experimental data, one needs to understand the effects of a finite frequency. Moreover, not only the value of  $\omega$  but also the details of how the driving is turned-on influence the corrections to quantized transport. In this regard, some mathematical results are known from Refs. [64] and [65], where the IQHE is analyzed as a quantum pumping process. In Ref. [64] it is demonstrated that, under strong regularity assumptions on the driving function, the finite-frequency corrections to the Hall conductivity are non-analytic functions of the form  $e^{-\alpha/\omega}$ , suggesting an extreme robustness to non-adiabatic effects. Interesting results are also presented in Ref. [65], where it is shown that the charge pumped in one period is *necessarily* a non-analytic function of  $\omega$ , but it is not specified if the absence of a Taylor expansion close to  $\omega = 0$  is due to terms of the form  $e^{-\alpha/\omega}$  or to other forms of non-analytic behaviour. Indeed, the picture emerging from Ref. [75] suggests that increasingly fast oscillations — non-analytic in nature — are found when a driving at finite frequency is turned on *suddenly*: the fast oscillations decorate an overall envelope of the response that displays *quadratic* corrections to the quantized limit, suggesting that the average response is perturbative, but hardly “robust” to a finite  $\omega$  driving. In chapter 3 I investigate further on this issue, showing that the response of the system — being perturbative or non-analytic in  $\omega$  —, depends on the details of how the driving is turned on. Floquet theory will play a major role in this analysis because it links ground state and finite-frequency topological properties.

The second question one might ask concerns the role played by disorder in transport phenomena. If, on the one hand, topology is expected to be robust against perturbations [15, 70], insofar the ground state manifold is protected by an energy gap, on the other hand, it is well known that one-dimensional disordered systems undergo Anderson localization [76], where transport is strongly suppressed. Quantized transport could then occur only on short time scales: when averaged over many driving periods transport should vanish. Indeed, if the evolution operator is localized in space, it cannot induce transport on distances larger than the localization length, while quantum pumping should describe *ballistic* transport, where the distance covered by a particle is proportional to the number of driving periods. I will show in chapter 4 that quantum pumping survives, in one-dimension, in presence of a finite disorder  $W$  below a critical value  $W_c$ , because in the topological phase a large number of Floquet states are *delocalized*, and they undergo a true delocalization/localization transition at  $W = W_c$ . This happens only if the driving cycle is topological, while in the trivial phase both Hamiltonian eigenstates and Floquet states are localized (although the latter tend to have larger localization lengths due to the periodic driving [77, 78]).

A final comment is due on the effect of dissipation, although not directly relevant for this thesis. Adiabatic pumping is a “deeply quantum” effect, meaning that it requires the coherent unitary evolution of the wave-function for a long time. Any experimental implementation, however, suffers from some source of dissipation, being the classical noise in the control fields or the coupling of the system to a continuum of environmental modes. When looking at topological properties of a time-dependent ground state, though, dissipation effects are not necessarily detrimental. If dissipation is tuned in such a way that it helps the system to relax to the appropriate “Floquet ground state”<sup>2</sup>, disposing in such a way of part of the energy that is necessarily absorbed from the driving, the quantization of particle transport could become more robust to non-adiabatic effects, as observed in Ref. [79]. This is particularly important for *interacting* driven systems because within a strictly unitary dynamics they would tend to “heat-up” to a featureless infinite temperature steady state. In this scenario, dissipation could stabilize topological transport at large time-scales without invoking many-body localization [24], which prevents heating but also suppresses transport.

## 1.2. Quantum Optimization

Quantum optimization is a focus of current research in quantum computation (QC) [27] and information processes. At its core, quantum optimization aims at designing and developing a “machine” — depicted schematically in Fig. 1.2 — which would be able to exploit quantum mechanics to prepare the ground state of a target Hamiltonian  $\hat{H}_{\text{target}}$ , usually of a spin- $\frac{1}{2}$  model [80], provided as an input to the machine. This goal has many motivations, both from fundamental research aspects and technological ones. Concerning fundamental aspects, being able to study the low-energy properties of spin Hamiltonians is one of the building blocks of modern condensed matter physics, and quantum optimization would be able to provide information on systems inaccessible to classical computers [81]. On the technological side, classical hard optimization problems are ubiquitous and some of them might be solved more efficiently with the aid of quantum hardware and algorithms [80, 82]. Moreover, quantum optimization would be a useful tool for quantum chemistry [26, 83, 84]. Finally, the theoretical and technological efforts required to build such a wonderful machine are worth being pursued *per se*, since they would likely both increase our understanding of fundamental processes in Nature and pave the way to a whole class of new technologies based on the manipulation of individual quantum systems.

One of the first methods proposed and implemented for quantum optimization is *Quantum Annealing* (QA) [28, 29, 85–87] or equivalently *Adiabatic Quantum Computation* (AQC) [30, 88]. It relies on the adiabatic theorem to drive the system from an initial state, easy to prepare, to the ground state of the target Hamiltonian. To be more precise, let us consider a spin- $\frac{1}{2}$  Hamiltonian  $\hat{H}_{\text{target}}$ , of which we wish to find the ground state, and a second

<sup>2</sup>The concept of Floquet ground state is somewhat non-trivial: usually topological properties are associated to the ground-state (GS) of a parameterized Hamiltonian  $\hat{H}(\mathbf{R})$ , as  $\mathbf{R}$  describes some closed path. However in a real implementation,  $\mathbf{R}$  is changed in time with a finite rate  $\omega$ , hence topology is better investigated by looking at Floquet states. When  $\omega$  is small, Floquet states and Hamiltonian eigenstates are very similar, thanks to the adiabatic theorem, hence one can refer to the Floquet state with the largest projection on the GS as Floquet ground state.

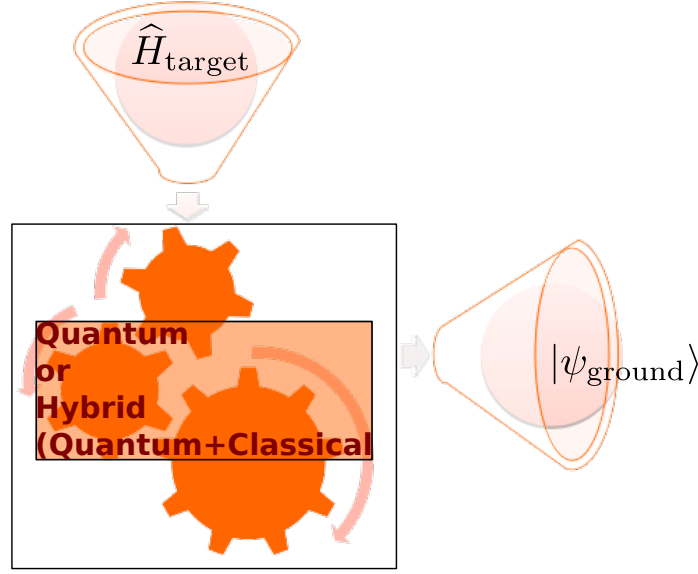


Figure 1.2.: Sketch of a “dream machine” which prepares the ground state of the Hamiltonian  $\hat{H}_{\text{target}}$  with fully quantum or hybrid quantum-classical hardware and algorithms.

Hamiltonian  $\hat{H}_{\text{drive}}$  that does not commute with  $\hat{H}_{\text{target}}$ . A common choice, relevant for both condensed matter and classical optimization problems, is

$$\hat{H}_{\text{target}} = \hat{H}_z = \sum_{i,j} J_{ij} \hat{\sigma}_i^z \hat{\sigma}_j^z, \quad \hat{H}_{\text{drive}} = \hat{H}_x = - \sum_i \hat{\sigma}_i^x. \quad (1.3)$$

In QA/AQC one considers a time-dependent interpolation of  $\hat{H}_{\text{target}}$  and  $\hat{H}_{\text{drive}}$  written as

$$\hat{H}(t) = s(t/\tau) \hat{H}_{\text{target}} + (1 - s(t/\tau)) \hat{H}_{\text{drive}}. \quad (1.4)$$

$s(t/\tau) \in [0, 1]$  is a control parameter chosen such that  $s(0) = 0$  and  $s(1) = 1$ , while  $\tau$  is the total evolution time. The system is first prepared in the ground state  $|\phi_0(0)\rangle$  of  $\hat{H}_{\text{drive}}$ , which is a trivial product state with all spins aligned in the  $\hat{\mathbf{x}}$  direction. Then, a Schrödinger evolution associated to  $\hat{H}(t)$  is implemented, and if  $\tau$  is “large enough” the adiabatic theorem ensures that the wave-function  $|\psi(t)\rangle$  describing the system will be always close to the instantaneous ground state of  $\hat{H}(t)$ . Hence at  $t = \tau$ , it will end up in the desired ground state of  $\hat{H}_{\text{target}}$ .

Unfortunately, true adiabatic evolution is quite fragile because it requires that the instantaneous ground state manifold  $|\phi_0(t)\rangle$  is always protected by an energy gap  $\Delta(t)$  from the excited states. If this is the case, the “slow enough” hypothesis of the adiabatic theorem requires that

$$\frac{\tau}{\hbar} \gg \frac{\Gamma}{\min_t \Delta^2(t)},$$

at least for a linear schedule  $s(t/\tau) = t/\tau$ . Here  $\Gamma$  is the Hamiltonian matrix element connecting the ground state and the first excited state. However, for most interesting Hamiltonians  $\hat{H}_{\text{target}}$ , the system crosses a phase transition at  $s_c < 1$  from a quantum paramagnetic phase ( $s(t) < s_c$ ) to a phase ordered in the  $\hat{\mathbf{z}}$  direction ( $s(t) > s_c$ ), whether ferromagnetic, anti-ferromagnetic or glassy. Phase transitions are always marked by an energy gap  $\Delta$  that

vanishes in the thermodynamic limit when the driving parameter reaches the critical point  $s_c$ , hence spoiling the efficiency of QA/AQC.

Two are therefore the main questions that one has to answer in QA/AQC applications. First, it is important to understand how the density of “defects” over the true desired GS scales with the number of spin variables  $N$  and with the total evolution time  $\tau$ . This allows characterizing the efficiency of the process and determines the minimum resources needed to obtain meaningful results. Second, the precise choice of the function  $s(t/\tau)$  plays a major role in the overall efficiency, hence its optimization is one of the key elements in developing a successful QA/AQC setup. Finding the optimal schedule  $s(t)$  appears at a first sight as a formidable problem, as it requires information on the target state or the instantaneous spectrum, in order to “slow down” the evolution at the appropriate point.

Despite the difficulty of optimizing the schedule  $s(t)$  and the limitations imposed by phase transitions, QA/AQC has been extremely popular since its first proposals, because it is easy to implement on noisy intermediate-scale quantum (NISQ) devices [32, 41], easier, at least, than circuit-based quantum computation. Indeed adiabatic ground state preparation requires a rather simple control setup: once the Hamiltonian is encoded in a set of physical qubits — which is a nontrivial step on its own — QA/AQC consists in the slow tuning of some external magnetic field in order to reach the target state. Therefore QA/AQC is much less sensitive to noise than circuit QC and it also suffers less from decoherence, since dissipation is not necessarily bad when looking for low-energy states of a Hamiltonian. The D-Wave processor, one of the most famous QA-based quantum computers, indeed exploits dissipation and temperature tuning during the optimization process [89].

From the theoretical point of view, one of the main disadvantages of QA/AQC is the difficulty in obtaining analytical results on the time-dependent wave-function  $|\psi(t)\rangle$ , since this is a complicated object computed as the time-ordered exponential associated to a time-dependent Hamiltonian

$$|\psi(t)\rangle = \text{Texp}\left(-\frac{i}{\hbar} \int_0^t dt' \hat{H}(s(t'))\right) |\psi(0)\rangle . \quad (1.5)$$

This evolution operator is nearly impossible to write explicitly, hence a useful approximation is to introduce a mesh of  $P$  points,  $s_1, \dots, s_P$ , substituting the schedule function  $s(t)$  with its discretized step-like form at  $s(t_m) = s_m$  in the interval  $[0, \tau]$ . Hence, the final state can be rewritten as

$$|\psi(\tau)\rangle \simeq \prod_{m=1}^{\leftarrow P} e^{-\frac{i}{\hbar} \hat{H}(s_m) \Delta t_m} |\psi(0)\rangle , \quad (1.6)$$

where  $\Delta t_m$  is the length of the  $m$ -th time interval, with  $\sum_{m=1}^P \Delta t_m = \tau$ , and the elements in the product are ordered from right to left, as  $m$  increases. This approximation might be referred to as “*step-QA*” [90]. Another simplification comes from Trotter decomposition: each term of the product can be written as the product of two simpler evolution operator

$$e^{-\frac{i}{\hbar} \hat{H}(s_m) \Delta t_m} = e^{-i \hat{H}_x \beta_m} e^{-i \hat{H}_z \gamma_m} + O(\Delta t_m^2) , \quad (1.7)$$

where we have assumed the Hamiltonian written in Eq. (1.3). The parameters  $\gamma_m$  and  $\beta_m$




Figure 1.3.: A smooth continuous-time QA schedule protocol  $s(t)$ , with the associated step-QA and digitized-QA schedules. Here  $P = 8$ . The time-intervals  $\Delta t_m$  are not identical. Figure taken from Ref. [90].

are linked to the original discretized schedule  $s_m$  in the following way

$$\begin{cases} \gamma_m = \frac{\Delta t_m}{\hbar} s_m \\ \beta_m = \frac{\Delta t_m}{\hbar} (1 - s_m) \end{cases} \Leftrightarrow \begin{cases} s_m = \frac{\gamma_m}{\gamma_m + \beta_m} \\ \Delta t_m = \hbar(\gamma_m + \beta_m) \end{cases} \quad (1.8)$$

The process associated to the parameters in Eq. (1.8) is called *digitized* Quantum Annealing (dQA), and it leads to the approximate solution

$$|\psi_P(\boldsymbol{\gamma}, \boldsymbol{\beta})\rangle = \prod_{m=1}^{\leftarrow P} e^{-i\beta_m \hat{H}_x} e^{-i\gamma_m \hat{H}_z} |+\rangle. \quad (1.9)$$

Here we write explicitly that the initially the system is prepared in the ground state of  $\hat{H}_x$  with all spins aligned in the  $\hat{\mathbf{x}}$  direction, hence  $|\psi(0)\rangle = |+\rangle$  with:

$$|+\rangle = \frac{1}{\sqrt{2^L}} \bigotimes_{i=1}^L (|\uparrow\rangle_i + |\downarrow\rangle_i), \quad (1.10)$$

while  $\boldsymbol{\gamma}$  and  $\boldsymbol{\beta}$  are real vectors of dimension  $P$ .

As advocated in Ref. [91], dQA has the advantage that each unitary in Eq. (1.9) can be implemented with universal quantum gates, with the possible inclusion of non-stoquastic terms [92] that are in general hard to simulate on classical computers [93].

Equation (1.9) is also the starting point of the main player in part II of this thesis, the Quantum Approximate Optimization Algorithm (QAOA). It is a hybrid quantum-classical variational scheme [42], first proposed by Farhi et al. in Ref. [44], which has recently gained momentum [94–97] and has been successfully realized in several experimental platforms [98, 99].

Essentially, QAOA considers the  $2P$  control parameters  $(\boldsymbol{\gamma}, \boldsymbol{\beta})$  as a set of *variational* parameters to be optimized in order to obtain the best possible approximation to the target



state. It is a hybrid quantum-classical scheme: A set of quantum gates performs the unitary evolution according to Eq. (1.9), with a given set of parameters  $(\gamma, \beta)$ , and computes the expectation value of the target Hamiltonian of the variational state

$$E_P(\gamma, \beta) = \langle \psi_P(\gamma, \beta) | \hat{H}_{\text{target}} | \psi_P(\gamma, \beta) \rangle . \quad (1.11)$$

Then, a classical minimization algorithm is used to find a minimum in the energy landscape  $E_P(\gamma, \beta)$ , leading to an approximation of the ground state  $|\psi_P(\gamma^*, \beta^*)\rangle$ , where

$$(\gamma^*, \beta^*) = \operatorname{argmin}_{(\gamma, \beta)} E_P(\gamma, \beta) . \quad (1.12)$$

QAOA could also be regarded as a control scheme that allows optimizing the digitized schedule  $s_m$  in a dQA protocol. In general, however, the minima in the energy landscape  $E_P(\gamma, \beta)$  are *not* associated to smooth protocols [36], but an iterative minimum search for increasing values of P can lead to schedules that correspond to the discretization of a continuous function, thus describing an optimized dQA [96]. Optimal dQA protocols are worth pursuing for several reasons: first, they are easier to implement experimentally, since they do not require abrupt changes in the control fields. Second, they should be more robust to noise because a small deviation from the optimal protocol leaves the system in the basin of the “smooth” local minimum, while this is not necessarily the case for highly irregular schedules. Finally, they are scalable, meaning that once the shape of the optimal schedule is known, it is possible to change the evolution time without searching for a new minimum, and they are universal, i.e. they lead to similar performances for different system sizes.

In this thesis, I focus on QAOA applied on the fully-connected Ising model, or p-spin model [100], which is a known case where standard QA/AQC fails due to the presence of a first-order phase transition and hence the need for exponentially long annealing time [100–102]. QAOA, instead, is able to prepare the ground state with resources scaling only *polynomially* with the number of variables  $N$ , therefore it outperforms QA/AQC in this model. However, this speedup comes at the price of a very complex energy landscape that makes local optimization highly dependent on the initial value of the QAOA parameters, for shallow circuits (small parameters number P), and the convergence towards local minima very slow. Indeed, we failed to find regular dQA schedules using the iterative procedures applied successfully to other problems [95, 96, 98].

In the last chapter we propose a Reinforcement Learning [103] assisted scheme for quantum optimization, where an artificial neural network (ANN) is trained to map quantum states to pairs of control parameters  $(\gamma, \beta)$  in order to minimize the variational energy in Eq. (1.11). We benchmark this method on the transverse field Ising model (TFIM) both with uniform and random couplings and on the  $p = 2$  fully connected p-spin model. Our RL scheme learns strategies converging to the optimal adiabatic solution for QAOA [96] for the translationally invariant quantum Ising chain. The same result holds in the presence of disorder, where the policy can be transferred among different disorder instances and system sizes. Indeed, our RL scheme allows the training part to be performed on small samples and transferred successfully on larger systems and different disordered realizations. Our RL-assisted QAOA is able to find regular schedules for the fully-connected Ising model, in sharp contrast with local minimization alone. It improves also on the stability of QAOA process, leading to much better results for large values of P and allowing for a better estimate of the scaling of the residual energy with P and the system size  $N$ .



**Part I.**

**Quantum Pumping**



## 2. Quantum pumping and Floquet Theory

In this chapter, I introduce the theoretical framework and the methods to understand topological phases in periodically driven quantum systems, which are the core of the first part of this thesis, for which this chapter might be regarded as a “technical introduction”. In Sec. 2.1 I give an overview of the topological interpretation of adiabatic quantum pumping, as given in Ref. [9] and generalized in Ref. [15]. In Sec. 2.2 I review the elements of Floquet theory that are going to be used in the following chapters to study the topological properties of periodically driven systems close to an adiabatic regime.

### 2.1. Thouless pumping

Let us consider a generic system made of  $N$  non-relativistic charged particles in one dimension, with periodic boundary conditions. To calculate the particle current, we include, via the minimal coupling, a vector potential  $A = \varphi/L$ , where  $\varphi$  is the magnetic flux piercing the ring of length  $L$ . On the continuum, the Hamiltonian reads

$$\hat{H}(\varphi) = \sum_{i=1}^N \left( \frac{(\hat{p}_i + \hbar\kappa)^2}{2m} + V(x_i) \right) + \frac{1}{2} \sum_{i \neq j} V_{\text{int}}(x_i - x_j), \quad (2.1)$$

where  $\hbar\kappa = \frac{e\varphi}{L}$  is the shift of momentum due to the flux  $\varphi$ ,  $V(x_i)$  the one-body potential and  $V_{\text{int}}(x_i - x_j)$  the two-body interaction. The current density operator can now be computed as [9]

$$\hat{J} = \partial_{\varphi} \hat{H}(\varphi)|_{\varphi=0} = \frac{e}{L} \sum_i \frac{\hat{p}_i}{m}. \quad (2.2)$$

Notice that the constant vector potential in Eq. (2.1) can be eliminated by a gauge transformation of the wave-function,  $\psi(x) \rightarrow \tilde{\psi}(x) = e^{-i\kappa x} \psi(x)$ , which is equivalent to *twist* the boundary condition on  $\tilde{\psi}(x)$ , i.e.,  $\tilde{\psi}(x + L) = e^{-i\theta} \tilde{\psi}(x)$  by an angle  $\theta = \kappa L = 2\pi \frac{\varphi}{\varphi_0}$ , where  $\varphi_0 = h/e$  is the quantum of magnetic flux. We can extend this derivation to a tight binding model, where the gauge transformation of the wave-function is equivalent to introducing a Peierls' phase in the hopping elements

$$\hat{c}_{l+1}^{\dagger} \hat{c}_l \rightarrow e^{i\kappa a} \hat{c}_{l+1}^{\dagger} \hat{c}_l, \quad (2.3)$$

where  $a$  is the lattice spacing. Thus, when we consider explicitly twisted boundary condition, the current is computed as

$$\hat{J} = \frac{2\pi}{\varphi_0} \partial_{\theta} \hat{H}(\theta), \quad (2.4)$$

which is equivalent to Eq. (2.2) with a non-zero magnetic flux, with  $\theta = 2\pi \frac{\varphi}{\varphi_0}$ .

If we now add a time-dependent periodic potential to the Hamiltonian, an interesting quantity to measure is the current induced by this periodic driving. Hence, let us now consider a fermionic Hamiltonian with a periodic driving,  $\hat{H}(t+\tau) = \hat{H}(t)$ , on a ring, assuming that the system is *gapped* (i.e., insulating) at any time  $t$ .  $\tau$  is the period of the driving. If the system is initially prepared in a  $N$ -particle insulating state  $|\Psi(0)\rangle = |\Phi_0(0)\rangle$  and the driving is slow enough, the adiabatic theorem guarantees that the system will approximately remain in the instantaneous ground state  $|\Phi_0(t)\rangle$  along the whole evolution. More precisely, an adiabatic expansion on the evolved states leads to [9]

$$|\Psi(t)\rangle = e^{i\gamma_0(t)} e^{-\frac{i}{\hbar} \int_0^t dt' E_0(t')} \left[ |\Phi_0(t)\rangle + i\hbar \sum_{n \neq 0} \frac{|\Phi_n(t)\rangle \langle \Phi_n(t) | \partial_t \Phi_0(t)\rangle}{E_n(t) - E_0(t)} \right] + O\left(\frac{1}{\tau^2}\right), \quad (2.5)$$

where  $\gamma_0(t)$  is the Berry phase accumulated by the ground state  $|\Phi_0(t)\rangle$ , while  $E_n(t)$  and  $|\Phi_n(t)\rangle$  are the instantaneous eigenvalues and eigenstates of  $\hat{H}(t)$ . To derive Eq. (2.5) one assumes that the Hamiltonian depends on time through a set of parameters  $\mathbf{R}(t/\tau)$ , such that the inverse of the period can be used as a small parameter for a series expansion. The sum in the second term, which runs over all the excited states ( $n \neq 0$ ), is of fundamental importance to computing transport properties.

We are now interested in the leading contribution to the average current  $\langle \Psi(t) | \hat{J}(t) | \Psi(t) \rangle$ . Since the time derivative  $\langle \Phi_n(t) | \partial_t \Phi_0(t) \rangle$  is of the order  $1/\tau \ll 1$ , being the evolution adiabatic, we will keep only terms up to first order in the expansion. A straightforward calculation leads to

$$\langle \Psi(t) | \hat{J}(t) | \Psi(t) \rangle = \langle \Phi_0(t) | \partial_\varphi \hat{H}(t) | \Phi_0(t) \rangle + i\hbar \left[ \frac{\langle \Phi_0(t) | \partial_\varphi \hat{H} | \Phi_n(t) \rangle \langle \Phi_n(t) | \partial_t \Phi_0(t) \rangle}{E_n(t) - E_0(t)} - \text{c.c.} \right]. \quad (2.6)$$

As specified in Eq. (2.2), the Hamiltonian and all eigenvalues and eigenvectors are computed for a vanishing magnetic flux  $\varphi = 0$ . To simplify Eq. (2.6), we make use of Hellmann-Feynman theorem to write the first term on the right hand side as

$$\langle \Phi_0(t) | \partial_\varphi \hat{H}(t) | \Phi_0(t) \rangle = \partial_\varphi E_0(t), \quad (2.7)$$

while the terms in the sum can be re-written as

$$\frac{\langle \Phi_0(t) | \partial_\varphi \hat{H} | \Phi_n(t) \rangle}{E_n(t) - E_0(t)} = -\langle \partial_\varphi \Phi_0(t) | \Phi_n(t) \rangle. \quad (2.8)$$

To make the interpretation of the following equations more physically transparent, we now substitute the magnetic flux  $\varphi$  with the associated momentum  $\kappa$ , first introduced in Eq. (2.1). Hence, we obtain the following expression for the average current

$$\begin{aligned} \langle \Psi(t) | \hat{J}(t) | \Psi(t) \rangle &= \frac{e}{L} \left[ \frac{1}{\hbar} \partial_\kappa E_0(t) - i \sum_{n \neq 0} (\langle \partial_\kappa \Phi_0(t) | \Phi_n(t) \rangle \langle \Phi_n(t) | \partial_t \Phi_0(t) \rangle - \text{c.c.}) \right] \\ &= \frac{e}{L} \left[ \frac{1}{\hbar} \partial_\kappa E_0(t) - i (\langle \partial_\kappa \Phi_0(t) | \partial_t \Phi_0(t) \rangle - \text{c.c.}) \right] \end{aligned} \quad (2.9)$$

To obtain the last equality we inserted the term  $n = 0$  into the sum, exploiting the fact that it gives no contribution, since  $\langle \partial_\kappa \Phi_0 | \Phi_0 \rangle \langle \Phi_0(t) | \partial_t \Phi_0(t) \rangle$  is a *real* number. Equation (2.9)

has an appealing interpretation: The first term is the group velocity of the instantaneous ground state, while the second term has the form of a *Berry curvature*, which is related to the geometrical properties of the ground state manifold. Both terms become more clear when the Hamiltonian has translational invariance and the momentum shift  $\kappa$  can be substituted by the real momentum  $k$ . In this situation, the current can be rewritten as an integral over all momenta and a sum over all occupied bands (recall that  $|\Phi_0\rangle$  describes an insulator)

$$\langle \Psi(t) | \hat{J}(t) | \Psi(t) \rangle = e \sum_v^{occ} \int_0^{2\pi/a} \frac{dk}{2\pi} \left[ \frac{1}{\hbar} \partial_k E_{v,k}(t) - i \left( \langle \partial_k u_{v,k}(t) | \partial_t u_{v,k}(t) \rangle - \text{c.c.} \right) \right], \quad (2.10)$$

where  $u_{v,k}(t)$  are the periodic part of the occupied Bloch wave-functions of the time-dependent Hamiltonian. In going from Eq. (2.9) to (2.10) we assumed implicitly the thermodynamic limit  $L \rightarrow \infty$ .

The number of particles transported across any section of the ring, the *pumped charge*, is the integral over one time period of the average current, divided by the particle charge  $e$

$$Q(\tau) = \frac{1}{e} \int_0^\tau dt \langle \Psi(t) | \hat{J}(t) | \Psi(t) \rangle = -i \sum_v^{occ} \int_0^\tau dt \int_0^{2\pi/a} \frac{dk}{2\pi} \left( \langle \partial_k u_{v,k}(t) | \partial_t u_{v,k}(t) \rangle - \text{c.c.} \right), \quad (2.11)$$

where we neglected the term with  $\partial_k E_{v,k}(t)$  because it is a real periodic function, integrated over one period. This is the consequence of the insensitivity of an insulator to an external driving: hence ordinary particle transport is absent, as expected. However the pumped charge can still be non-zero. Indeed, Eq. (2.11) is the integral over a closed surface — the torus  $[0, \frac{2\pi}{a}] \times [0, \tau]$  — of the Berry curvature of the ground state:

$$\mathcal{B}_{k,t} = i \sum_v^{occ} \left( \langle \partial_k u_{v,k}(t) | \partial_t u_{v,k}(t) \rangle - \text{c.c.} \right). \quad (2.12)$$

This integral can assume only an integer value, and is now known as the *first Chern number* of the ground state: a topological invariant. Hence  $Q(\tau) \in \mathbb{Z}$  is *quantized* in the adiabatic limit  $\tau \rightarrow \infty$ : its sign shows the direction in which the average current flows. Equation (2.11), first derived by Thouless [9], is one of the cornerstones of non-equilibrium topological phases in condensed matter systems.

To conclude this section, we extend Eq. (2.11) to systems where translational invariance is broken, either by an external potential or interaction, following the work of Refs. [15, 70]. We make use of the insensitivity of insulators to the boundary conditions [104], to express the pumped charge as the average over all possible phase twists  $\theta$  at the boundary

$$Q(\tau) = \frac{1}{\hbar} \int_0^{2\pi} \frac{d\theta}{2\pi} \langle \Psi(t) | \partial_\theta \hat{H}(\theta) | \Psi(t) \rangle, \quad (2.13)$$

where we used Eq. (2.4) to write the particle current. Exploiting the same steps that led to Eqs. (2.9) and (2.11), we finally arrive at

$$Q(\tau) = -i \int_0^\tau dt \int_0^{2\pi} \frac{d\theta}{2\pi} \left( \langle \partial_\theta \Phi_0(t) | \partial_t \Phi_0(t) \rangle - \text{c.c.} \right), \quad (2.14)$$

which has again the mathematical structure of a curvature integrated on a closed surface — the torus  $[0, \tau] \times [0, 2\pi]$  — hence giving an integer Chern number.

## 2.2. Floquet Theory

In this section, I will give an overview of the Floquet theory, which is the main tool we are going to use to investigate the dynamics of periodically driven systems. In a more general framework, Floquet theory includes a series of results on systems of linear differential equations with periodic coefficients. For what regards quantum mechanics, the main field of applications is for Hamiltonians which are periodic in space, which leads to the Bloch theorem and to the band theory of crystals, or in time. Thus, let us consider a Hamiltonian which is time-periodic with a period  $\tau$ ,  $\widehat{H}(t+\tau) = \widehat{H}(t)$ , and its associated Schrödinger time-evolution operator  $\widehat{U}(t, t_0)$ . Floquet theorem, the central result of the theory, states that

$$\widehat{U}(t + n\tau, t_0) = \widehat{U}(t, t_0) \left[ \widehat{U}(t_0 + \tau, t_0) \right]^n, \quad (2.15)$$

where  $t_0 < t < \tau$  and  $n$  is an integer number. Equation (2.15) has the important implication that the evolution operator within one period of the driving  $\widehat{U}(t_0 + t, t_0)$  with  $t < \tau$  completely describes the dynamics at arbitrarily large times. Of particular interest is the evolution operator for *exactly* one period, which is often referred to as the *Floquet operator*  $\widehat{F}_{t_0}(\tau) = \widehat{U}(t_0 + \tau, t_0)$ .  $\widehat{F}_{t_0}(\tau)$  is unitary, and therefore diagonalizable. Hence, we can write a “spectral decomposition” over its basis

$$\widehat{F}_{t_0}(\tau) = \sum_{\nu} e^{-\frac{i}{\hbar} \epsilon_{\nu} \tau} |\phi_{\nu}(t_0)\rangle \langle \phi_{\nu}(t_0)|, \quad (2.16)$$

where the eigenvectors  $|\phi_{\nu}(t_0)\rangle$  are called the Floquet eigen-modes, while the  $\epsilon_{\nu}$  are the *Floquet quasi-energies*. Since the  $\epsilon_{\nu}$  enter into Eq. (2.16) as “phases”, the quasi-energies are usually taken into the first *Floquet-Brillouine zone*  $\epsilon_{\nu} \in \left[-\frac{\hbar\omega}{2}, \frac{\hbar\omega}{2}\right)$ , with  $\omega = \frac{2\pi}{\tau}$  being the driving frequency, in analogy with the quasi-momenta  $k$  entering in the Bloch theorem for translationally invariant Hamiltonians.

An important consequence of Floquet theorem concerns the time evolution of Floquet states  $|\psi_{\nu}(t)\rangle = \widehat{U}(t + t_0, t_0) |\phi_{\nu}(t_0)\rangle$ . They can be written as the product of a phase term and periodic part

$$|\psi_{\nu}(t)\rangle = \widehat{U}(t + t_0, t_0) |\phi_{\nu}(t_0)\rangle = e^{-\frac{i}{\hbar} \epsilon_{\nu} t} |\phi_{\nu}(t + t_0)\rangle, \quad (2.17)$$

where  $|\phi_{\nu}(t + \tau)\rangle = |\phi_{\nu}(t)\rangle$  is called *Floquet mode* and has the same periodicity of the Hamiltonian  $\widehat{H}(t)$ . By construction, the Floquet states solve the time-dependent Schrödinger equation:

$$i\hbar \frac{\partial}{\partial t} |\psi_{\nu}(t)\rangle = \widehat{H}(t) |\psi_{\nu}(t)\rangle. \quad (2.18)$$

Equation (2.17) is the generalization for time-periodic Hamiltonian of the Bloch theorem in one dimension, with  $\epsilon_{\nu}/\hbar$  playing the role of the quasi-momentum  $k$ .<sup>1</sup>

Equation (2.18), together with (2.17), allows us to derive directly a kind of eigenvalue equation linking Floquet modes and quasi-energies:

$$\widehat{\mathcal{K}}(t) |\phi_{\nu}(t)\rangle = \epsilon_{\nu} |\phi_{\nu}(t)\rangle, \quad (2.19)$$

<sup>1</sup>To avoid confusion, we will refer to the  $|\psi_{\nu}(t)\rangle$  in Eq. (2.17) as *Floquet states*: they are solutions of the time-dependent Schrödinger equation (2.18). The *Floquet modes*, instead, are only the periodic parts  $|\phi_{\nu}(t)\rangle$  of the Floquet states.



where the Hermitean operator

$$\widehat{\mathcal{K}}(t) = \widehat{H}(t) - i\hbar \frac{\partial}{\partial t} \quad (2.20)$$

is referred to as the *Floquet extended Hamiltonian*. It acts on the Hilbert space  $\mathcal{H}' = \mathcal{H} \otimes \mathcal{I}_\tau$ , with  $\mathcal{H}$  the original Hilbert space of  $\widehat{H}(t)$  and  $\mathcal{I}_\tau$  the space of square-integrable periodic function of period  $\tau$ . The associated inner product is the natural extension of that of  $\mathcal{H}$ :

$$\langle\langle f|g\rangle\rangle = \frac{1}{\tau} \int_0^\tau dt \langle f(t)|g(t)\rangle. \quad (2.21)$$

Armed with this machinery, we can now describe a generic wave-function evolving with a periodic  $\widehat{H}(t)$ , and prepared in the initial state  $|\psi(t_0)\rangle = |\psi_0\rangle$ . In the following, we will set  $t_0 = 0$  for simplicity. First, we write the initial wave-function on the (assumed complete) Floquet mode basis, inserting a resolution of the identity:

$$|\psi_0\rangle = \sum_\nu |\phi_\nu(0)\rangle \langle\phi_\nu(0)|\psi_0\rangle = \sum_\nu c_\nu |\phi_\nu(0)\rangle, \quad (2.22)$$

with  $c_\nu = \langle\phi_\nu(0)|\psi_0\rangle$ . Next, we use Eqs. (2.15) and (2.17) to express the state of the system at a generic time  $t' = m\tau + t$  as:

$$|\psi(m\tau + t)\rangle = \sum_\nu e^{-\frac{i}{\hbar}\epsilon_\nu(m\tau+t)} c_\nu |\phi_\nu(t)\rangle. \quad (2.23)$$

The information on the initial state is stored only in the expansion coefficients  $c_\nu = \langle\phi_\nu(0)|\psi_0\rangle$ , which are constant in time, while all the time-dependence is dictated by the phase factors (as in the ordinary time-dependent Schrödinger evolution for time-independent Hamiltonians) and in the periodic Floquet modes  $|\phi_\nu(t)\rangle$ .

Here some comments are due on the link between Floquet states and Hamiltonian eigenstates. If the evolution was truly adiabatic, the two would coincide, and the quasi-energy could be computed directly from the adiabatic theorem. Indeed, for a strictly adiabatic time-evolution of the Hamiltonian eigenstates  $|\psi_n\rangle$  we would have:

$$\widehat{U}(t,0)|\psi_n(0)\rangle = e^{i\gamma_n(t)} e^{-\frac{i}{\hbar} \int_0^t dt' E_n(t')} |\psi_n(t)\rangle. \quad (2.24)$$

Here  $|\psi_n(t)\rangle$  is the  $n$ -th instantaneous eigenstate of  $\widehat{H}(t)$ ,  $E_n(t)$  its energy and  $\gamma_n(t) = \int_0^t dt' \langle\psi_n(t')|\partial_{t'}\psi_n(t')\rangle$  the associated Berry phase. From Eq. (2.24) it is evident that  $|\psi_n(\tau)\rangle \equiv |\psi_n(0)\rangle$  is an eigenstate of the Floquet operator  $\widehat{U}(\tau,0)$ , with eigenvalue  $e^{i\gamma_n(\tau)} e^{-\frac{i}{\hbar} \int_0^\tau dt' E_n(t')}$ . Hence, in the adiabatic limit the quasi-energy reads

$$\epsilon_n = \frac{1}{\tau} \left( \hbar\gamma_n(\tau) - \int_0^\tau dt E_n(t) \right). \quad (2.25)$$

For non-equilibrium systems with a “small enough” driving frequency, this is a good estimate of the actual quasi-energies, and it is extremely useful to have some physical insight in the system dynamics, as will be further developed in chapter 3. When  $\omega > 0$  ( $\tau$  is finite), Floquet and Hamiltonian eigenstates do not coincide any more, but in general in a quasi-adiabatic scenario, Floquet states will have a dominant projection on a single instantaneous energy eigenstate, and very small overlap with the others. These overlaps can be computed with the adiabatic theorem and are analyzed in detail in appendix A.1.

### 2.2.1. Pumped charge and winding of quasi-energies

Here we link the quantization of Thouless pumping described in section 2.1 to the topological aspects of Floquet states. Let us consider a time-periodic Hamiltonian  $\hat{H}(t + \tau) = \hat{H}(t)$ , which can be used to describe adiabatic quantum pumping. In the following, we set  $e = 1$ , so that the charge current is equivalent to the particle current.

We are interested in computing the charge pumped at a generic time  $t_f$ , that we express as

$$Q(t_f) = \int_0^{t_f} dt \langle \Psi(t) | \hat{J}(t) | \Psi(t) \rangle , \quad (2.26)$$

where  $\Psi$  is the (many-body) wave-function and  $\hat{J}(t)$  the time-dependent current density operator. Let us focus on the pumped charge averaged over (infinitely) many periods

$$Q_d = \lim_{m \rightarrow \infty} \frac{1}{m} Q(m\tau) = \lim_{m \rightarrow \infty} \frac{1}{m} \int_0^{m\tau} dt \langle \Psi(t) | \hat{J}(t) | \Psi(t) \rangle . \quad (2.27)$$

Our aim is to use Floquet representation to express this quantity. First of all, we write the state, as in Eqs. (2.22) and (2.23):

$$|\Psi(0)\rangle = \sum_{\nu} c_{\nu} |\Phi_{\nu}(0)\rangle \implies |\Psi(t)\rangle = \sum_{\nu} c_{\nu} e^{-i\mathcal{E}_{\nu}t/\hbar} |\Phi_{\nu}(t)\rangle , \quad (2.28)$$

where  $|\Phi_{\nu}(t)\rangle$  is a (many-body) Floquet mode with (many-body) quasi-energy  $\mathcal{E}_{\nu}$ . Expressing the average current using such a decomposition of  $|\Psi(t)\rangle$ , we write:

$$\begin{aligned} Q(m\tau) &= \sum_{n=0}^{m-1} \sum_{\nu, \nu'} c_{\nu'}^* c_{\nu} e^{in\tau(\mathcal{E}_{\nu'} - \mathcal{E}_{\nu})/\hbar} \int_{n\tau}^{(n+1)\tau} dt e^{it(\mathcal{E}_{\nu'} - \mathcal{E}_{\nu})/\hbar} \langle \Phi_{\nu'}(t) | \partial_{\varphi} \hat{H}(\varphi, t) |_{\varphi=0} | \Phi_{\nu}(t) \rangle \\ &= \sum_{n=0}^{m-1} \sum_{\nu, \nu'} c_{\nu'}^* c_{\nu} e^{in\tau(\mathcal{E}_{\nu'} - \mathcal{E}_{\nu})/\hbar} \int_0^{\tau} dt e^{it(\mathcal{E}_{\nu'} - \mathcal{E}_{\nu})/\hbar} \langle \Phi_{\nu'}(t) | \partial_{\varphi} \hat{H}(\varphi, t) |_{\varphi=0} | \Phi_{\nu}(t) \rangle , \end{aligned} \quad (2.29)$$

where in the integral we have used the periodicity of  $\hat{H}(t)$  and of the Floquet modes. If we take the limit  $m \rightarrow \infty$  and average over the number of periods, only terms with  $\nu = \nu'$  survive, leading to

$$Q_d = \sum_{\nu} \mathcal{N}_{\nu} \int_0^{\tau} dt \langle \Phi_{\nu}(t) | \partial_{\varphi} \hat{H}(\varphi, t) |_{\varphi=0} | \Phi_{\nu}(t) \rangle , \quad (2.30)$$

where  $\mathcal{N}_{\nu} = |c_{\nu}|^2$  is the occupation number of the  $\nu$ -th many-body Floquet state, given the initial state of the system  $|\Psi(0)\rangle$ . Equation (2.30) defines the main observable of interest in part I of this thesis. Exploiting properties of the evolution operator [70], we can rewrite

$$\int_0^{\tau} dt \hat{U}^{\dagger}(t, 0) \partial_{\varphi} \hat{H}(t) \hat{U}(t, 0) = i \hat{U}^{\dagger}(\tau, 0) \partial_{\varphi} \hat{U}(\tau, 0) , \quad (2.31)$$

which, together with Eq. (2.30), leads to

$$Q_d = \tau \sum_{\nu} \mathcal{N}_{\nu} \partial_{\varphi} \mathcal{E}_{\nu}(\varphi) |_{\varphi=0} . \quad (2.32)$$

To better understand this result, it is useful to distinguish explicitly between systems possessing translational invariance and those that are not uniform in space. First of all let us assume that interactions are negligible, so that the  $N$ -particle Floquet state  $|\Phi_\nu(0)\rangle$  is just a Slater determinant with  $N$  single-particle (SP) states. If momentum is conserved, the SP Floquet states are labeled by the band index  $\alpha$  and the wave-number  $k$ , such that in second quantization we can write

$$|\Phi_\nu(0)\rangle = \prod_k f_{\alpha,k}^\dagger |0\rangle , \quad (2.33)$$

where  $f_{\alpha,k}^\dagger$  creates a particle in the SP Floquet state  $|\phi_{\alpha,k}(0)\rangle$ . An immediate consequence of Eq. (2.33) is that

$$\mathcal{E}_\nu = \sum_k \epsilon_{\alpha,k} \rightarrow \frac{L}{2\pi} \int_0^{2\pi/a} dk \epsilon_{\alpha,k} , \quad (2.34)$$

i.e., the  $N$ -particle quasi-energy is the sum of the single-particle ones associated with the states in the Slater determinant. The initial ground state  $|\Psi(0)\rangle$  can be also decomposed in single-particle orbitals  $|\psi_k(0)\rangle$ , allowing to write the occupation number as

$$\mathcal{N}_\nu = \frac{L}{2\pi} \sum_\alpha \int_0^{2\pi/a} dk |\langle \psi_k(0) | \phi_{\alpha,k}(0) \rangle|^2 . \quad (2.35)$$

Translational invariance allows to express Eq. (2.32) as a single integral over momentum and the sum over all  $N$ -particles Floquet states becomes a sum over all SP bands  $\alpha$

$$Q_d = \frac{L\tau}{2\pi} \int_0^{2\pi/a} dk n_{\alpha,k} \partial_\varphi \epsilon_{\alpha,k} . \quad (2.36)$$

Finally, we exploit  $\partial_\varphi \rightarrow \frac{e}{\hbar L} \partial_k$  to write

$$Q_d = \frac{1}{\hbar\omega} \sum_\alpha \int_0^{2\pi/a} dk n_{\alpha,k} \partial_k \epsilon_{\alpha,k} , \quad (2.37)$$

where we dropped the charge  $e$  because we are interested in *particle* current. The occupation number now is computed as  $n_{\alpha,k} = |\langle \phi_{\alpha,k}(0) | \psi_k(0) \rangle|^2$ , where  $|\phi_{\alpha,k}(0)\rangle$  and  $|\psi_k(0)\rangle$  are respectively the  $\alpha$ -th Floquet state and the initial ground-state components with momentum  $k$ . The pumped charge is written as the weighted sum of the *winding number* of the Floquet quasi-energies. Indeed, if we focus on the simple scenario  $n_{\alpha,k} = \delta_{\alpha,\alpha'}$ , meaning that only the  $\alpha'$ -th Floquet state is occupied, Eq. (2.37) reduces to

$$Q_d = \frac{1}{\hbar\omega} \int_0^{2\pi/a} dk \partial_k \epsilon_{\alpha',k} . \quad (2.38)$$

Because of lattice periodicity  $\epsilon_{\alpha',2\pi/a} = \epsilon_{\alpha',0} \bmod \hbar\omega$  — recall that the quasi-energies are the phases of the eigenvalues of the evolution operator — the integral over the momentum in Eq. 2.38 must give an integer multiple of  $\hbar\omega$ . Combining it with the prefactors, we obtain that the pumped charge is an integer number, called winding number, that relates particle transport to the topology of the Floquet states. This is schematically illustrated in Fig. 2.1, where the quasi-energies of the Rice-Mele model [75, 105] are plotted versus the momentum  $k$ . The thick line, associated with a single Floquet state, winds one time around the Floquet-Brillouin zone and the transported charge is indeed  $Q_d = 1$ . However, preparing the system




Figure 2.1.: Floquet quasi-energies vs momentum for the Rice-Mele model, with a driving frequency  $\omega = 0.2$ . The thick line is the Floquet state with the highest projection over the ground state of the Hamiltonian. The winding number can be computed counting the number of times that  $\epsilon_k$  has a jump in the Floquet-Brillouin zone, with a sign depending on the derivative:  $+1$  if  $\epsilon_k$  is increasing,  $-1$  if it is decreasing. Picture taken from Ref. [75].

in a single Floquet state is difficult, since typically the system is prepared in the ground state of the Hamiltonian, which has in general projections over many Floquet modes. Hence the pumped charge computed using Eq. (2.37) is generally not an integer number: there are corrections depending on the occupation numbers  $n_{\alpha,k}$ , which are nevertheless small when the driving is (quasi)-adiabatic.

In a disordered system, it is useful to compute the charge as its average over twisted boundary conditions. Indeed if a quantum state projector is exponentially localized, bulk properties are insensitive to a twist  $\theta$  in the boundary condition [15, 104], or equivalently they are insensitive to the presence of a flux that induces a phase shift  $\theta$  when a particle hops through the whole ring. Exploiting Eq. (2.4) we can write the charge as

$$Q_d = \frac{2\pi\tau}{\varphi_0} \sum_{\nu} \int_0^{2\pi} \frac{d\theta}{2\pi} \mathcal{N}_{\nu}(\theta) \partial_{\theta} \mathcal{E}_{\nu}(\theta). \quad (2.39)$$

Now the winding number is no longer computed as the integral over the momentum, but rather over the phase twist at the boundary, which plays the role of a “supercell pseudo-momentum”.

### 3. Non-adiabatic effects

In this chapter, I discuss a work published in Ref. [106], where we study the robustness of the quantization of the Hall conductivity towards the details of the protocol with which a longitudinal uniform driving force  $F_x(t)$  is turned on.

The integer quantum Hall effect (IQHE) is one of the cornerstones of modern condensed matter theory. It describes the current induced in a two-dimensional insulator, in a direction perpendicular to an external field  $j_y = \sigma_H F_x$ , mediated by the Hall or transverse conductivity.  $\sigma_H$  is quantized in units of  $e^2/\hbar$  in the linear response regime  $F_x \rightarrow 0$ , but experimentally it holds with stunning precision even for rather large external fields.

Here we characterize the deviation from the linear response regime in the simplest tight-binding model that displays IQHE, the Harper-Hofstadter model. Let us consider free spinless fermions on a square lattice, where each plaquette is pierced by a magnetic flux  $\phi$  and subject to a uniform driving in the  $\hat{x}$  direction  $F_x = F\hat{x}$ , as depicted in Fig. 3.1(a). Due to IQHE, current flows in the transverse direction  $\hat{y}$ . We investigate how the time-averaged (steady-




Figure 3.1.: (a) Sketch of the Harper-Hofstadter model on a square lattice with lattice spacing  $a$ , and a magnetic flux per plaquette  $\phi$ . (b) Possible deviations from quantization: perturbative (blue dashed line) or non-analytic (red line).

state) particle current density  $j_y$  deviates from the quantized value  $j_y h/F = n$  due to the finite value of  $F$  and the details of the switching-on protocol of the external field. There are two possible outcomes, summarized in Fig. 3.1: the corrections to the quantized value could be either perturbative in the field strength  $F$  or non-analytic. Regarding the possible switching on protocols of  $F_x$ , we consider a sudden quench, a linear ramp, and a smooth s-shaped ramp, the latter two taking place in tunable time  $t_0$ . The three possibilities are depicted in Fig. 3.2.




Figure 3.2.: Possible schedules for the switching-on of the uniform driving force  $F_x(t)$ . The sudden quench case  $F_x(t) = \theta(t)F$ , where  $\theta(t)$  is the Heaviside step function, is recovered for  $t_0 \rightarrow 0$ .

We tackle the problem as a quantum pumping process, using Peierls phase to introduce a complex hopping amplitude periodic in time  $J \rightarrow J e^{-i\frac{aF}{\hbar}t} = J e^{-i\omega t}$ , with a frequency proportional to the field strength  $F$ . Exploiting the time-periodicity of the Hamiltonian  $\hat{H}(t)$ , we use Floquet techniques to study this problem, as described in Sec. 3.3.

In this picture, the (Kubo) linear response  $F \rightarrow 0$  regime corresponds to the adiabatic limit for  $\hat{H}(t)$ . In the case of a sudden quench  $j_y \hbar/F$  shows  $F^2$  corrections to the perfectly quantized limit. When the switching-on is smooth, the result depends on the switch-on time  $t_0$ : for a fixed  $t_0$  we observe a crossover force  $F^*$  between a quadratic regime for  $F < F^*$  and a *non-analytic* exponential  $e^{-\gamma/|F|}$  for  $F > F^*$ . The crossover  $F^*$  decreases as  $t_0$  increases, eventually recovering the topological robustness.

### 3.1. Introduction

The quantization of the transverse conductivity  $\sigma_H$  in the Integer Quantum Hall Effect [51] (IQHE) is probably the most famous manifestation of a topological invariant, the first Chern number, in condensed matter physics [107]. Indeed the celebrated TKNN paper [54] showed that in the linear response regime, i.e. when the external electric field is small, the Hall conductivity predicted by the Kubo formula  $j_y^e = \sigma_H E_x$  is quantized and can be written as the sum of the Chern numbers of the occupied bands. Therefore it must be an integer, in units of  $e^2/h$ .

The extreme precision of the quantized Hall conductance revealed in the experiments [51] suggests a remarkable robustness of the IQH phase against many ingredients, notably the presence of impurities and interactions, and the strength of the applied electric field. Concerning the latter issue, the mathematical physics literature [64] has shown that corrections to the Kubo formula vanish in Quantum Hall systems to all orders in perturbation theory.

Quite recently, the issue of the topological robustness of a related phenomenon — Thouless pumping in one-dimensional insulators [9] — has been re-examined, showing that the details

of the preparation of the quantum non-equilibrium steady-state and of the time interval in which the pumped charge is measured deeply influence how the topological  $\tau \rightarrow \infty$  adiabatic limit is approached [75]. In particular, it was shown that the charge pumped over a *finite* number of periods shows non-analytic corrections — in the form of faster and faster oscillations as  $\omega = 2\pi/\tau \rightarrow 0$  — when the periodic driving protocol is turned-on abruptly starting from an initial uncorrelated insulating state [75]. Such a non-analytic approach of the adiabatic (topological) limit  $\omega \rightarrow 0$  was indeed predicted by Avron & Kons [65] through rigorous general arguments. What such rigorous arguments do not tell is how the limit  $\omega \rightarrow 0$  is approached when one considers the asymptotic (steady state) single-period pumped charge, where topological effects should most appropriately be looked for [66] because this involves an *infinite-time limit*. Remarkably, Ref. [75] shows that non-analytic corrections present at finite-time turn into *quadratic* corrections  $\sim \omega^2$  when the asymptotic pumped charge is considered.

Modern realizations of the IQHE physics involve artificial gauge fields in cold atomic systems [16, 19, 20, 108]. In the light of the results of Ref. [75], these experiments raise the non-trivial issue of the robustness of the quantized Hall conductance against many details, including primarily the preparation of the Quantum Hall state and the ensuing turning-on of the constant field, as well as the measurement of the transverse current. To set up and state the problem we will address, let us assume that the coherence-time [16] of these cold atomic systems is so long that it is legitimate to estimate the time-average transverse current from its infinite-time limit

$$j_y = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{t_0}^{t_0+T} dt' \langle \psi(t') | \hat{j}_y | \psi(t') \rangle . \quad (3.1)$$

Here  $\hat{j}_y$  is the space-averaged particle-current density operator, and  $|\psi(t)\rangle$  is assumed to evolve unitarily with the system Hamiltonian  $\hat{H}(t)$ , including the external uniform force field  $F_x(t)$  in the  $\hat{\mathbf{x}}$ -direction, which we represent by an extra time-dependent vector potential  $\mathcal{A}_x \hat{\mathbf{x}}$  with  $\partial_t \mathcal{A}_x(t) = F_x(t)$ . Furthermore, let us assume that the external uniform force  $F_x(t)$  is switched-on in a time  $t_0$  towards a stationary value  $F$ , i.e.,  $F_x(t) = f(t/t_0)F$ , where  $f(s)$ , with  $s = t/t_0 \in [0, 1]$ , is a switching-on function such that  $f(0) = 0$  and  $f(1) = 1$ . The Kubo formula for the IQHE implies that for small  $F$ :

$$j_y = \sigma_{yx} F = \frac{n}{h} F , \quad (3.2)$$

meaning that in the limit  $F \rightarrow 0$  the quantity  $j_y h/F$  is *exactly* an integer number  $n$ . A robust quantization against the strength of  $F$  would appear, in this context, as non-analytic corrections of the form  $j_y h/F \simeq n + Ae^{-\gamma/|F|}$ , while the presence of quadratic corrections,  $j_y h/F \simeq n + BF^2 + o(F^2)$ , would signal an ordinary perturbative response.

We investigate how the finite value of the stationary driving force  $F$  and the details of the driving protocol, encoded in  $t_0$  and in the switching-on function  $f(s = t/t_0)$ , affect the precision of the measurement of the transverse Hall response in cold atoms IQHE systems. Our investigation focuses on the Harper-Hofstadter (HH) model [109], a two-dimensional tight-binding Hamiltonian for IQHE which is particularly relevant for experimental realizations with optical lattices [19, 20, 108], providing an excellent tool to study QHE physics in a tunable and controlled system. The techniques used involve quite standard Floquet tools to study the time-periodic dynamics of the transverse current, which can be formulated as a quantum pumping problem.

We will show that the main responsible for topological robustness is the switching function  $f(t/t_0)$ . We analyze in detail three possible schedules: a sudden quench  $F_x(t) = \theta(t)F$ , a linear ramp  $F_x(t) = (t/t_0)F$  and a smoother ramp, as sketched in Fig. 3.2. We will show that in the sudden case,  $F_x(t) = \theta(t)F$ , the Hall response of the system is perturbative, and  $F^2$  corrections to  $j_y h/F$  are present. When the driving force is turned-on linearly in a time  $t_0$ ,  $F_x(t) = (t/t_0)F$ , we find two distinct regimes: for a fixed  $t_0$  we observe a crossover force  $F^*(t_0)$  between a quadratic regime for  $F < F^*$  and a *non-analytic* exponential  $e^{-\gamma/|F|}$  for  $F > F^*$ . The crossover  $F^*(t_0)$  decreases as  $t_0$  increases, eventually recovering the topological robustness. Finally, if the switching-on is smoother (with a continuous derivative),  $F_x(t) = \frac{1}{2}(1 - \cos(\pi t/t_0))F$ , we observe no qualitative differences with the linear ramp case, suggesting the main ingredient for the topological robustness seems to be the continuity of  $F_x(t)$  and a suitably long  $t_0$ .

## 3.2. Model

Our starting point is the Harper-Hofstadter Hamiltonian [109], sketched in Fig. 3.3, which describes a tight-binding system of non-interacting spinless fermions on a two-dimensional (2D) square lattice, pierced by a uniform magnetic field  $\mathbf{B} = B\hat{\mathbf{z}}$  perpendicular to the lattice plane:

$$\hat{H} = -J_0 \sum_{l,m} \left[ \hat{c}_{l+1,m}^\dagger \hat{c}_{l,m} + e^{-i2\pi\alpha l} \hat{c}_{l,m+1}^\dagger \hat{c}_{l,m} + \text{H.c.} \right]. \quad (3.3)$$

Here  $J_0$  is the bare hopping amplitude, and  $(l, m)$  are integers labelling the square lattice sites,  $\mathbf{r}_{l,m} = a(l\hat{\mathbf{x}} + m\hat{\mathbf{y}})$ , with lattice spacing  $a$ , with boundary conditions to be discussed later on. The magnetic field flux per plaquette, in units of the flux quantum  $\phi_0 = hc/e$ , is here  $\alpha = a^2 B/\phi_0$ , and results in a complex hopping amplitude through Peierls' substitution,  $J_0 e^{-i\frac{e}{hc} \int_{\mathbf{r}}^{\mathbf{r}'} \mathbf{A} \cdot d\mathbf{x}}$ , with a Landau gauge choice for the vector potential  $\mathbf{A} = Bx\hat{\mathbf{y}}$ , breaking translational invariance along the  $\hat{\mathbf{x}}$ -direction.

In a condensed matter realization of this model Hamiltonian, with charged particles in real magnetic fields, one would not be able to explore the full phase diagram of the model for  $\alpha \in [0, 1]$  since the flux per plaquette is too small, even with large laboratory fields. In modern realizations with neutral cold atoms in optical lattices [108, 110], on the contrary, synthetic gauge fields are used and all interesting values of  $\alpha$  are possible. Historically, as discovered by Hofstadter [109], the spectrum is extremely complex, with rational values of  $\alpha = p/q$  leading to  $q$  energy sub-bands with gaps in between. The crucial realization, due to Thouless and coworkers [54], is that the insulating states obtained when the Fermi energy lies inside the gaps between such sub-bands has a quantized Hall conductance

$$\sigma_H = -\frac{e^2}{h} \sum_{\nu}^{\text{occ}} \int_{\text{BZ}} \frac{d^2\mathbf{k}}{2\pi} \Omega_{\nu}(\mathbf{k}) = n \frac{e^2}{h}, \quad (3.4)$$

where  $\Omega_{\nu}(\mathbf{k}) = i \left[ \langle \partial_{k_x} u_{\nu,\mathbf{k}} | \partial_{k_y} u_{\nu,\mathbf{k}} \rangle - \langle \partial_{k_y} u_{\nu,\mathbf{k}} | \partial_{k_x} u_{\nu,\mathbf{k}} \rangle \right]$  is the Berry curvature [107] of  $\nu$ -th occupied band, and  $u_{\nu,\mathbf{k}}$  denote the periodic part of the Bloch wave-functions on the (magnetic) Brillouin Zone (BZ) of the system. This implies that a Hall current flows, for instance,






Figure 3.3.: Schematic representation of the Harper-Hofstadter model: free electrons on a square lattice of spacing  $a$ , in the presence of a uniform magnetic field  $B$  perpendicular to the lattice plane.

in the  $y$ -direction when an electric field  $E_x$  acts in the  $x$ -direction:  $j_y = \sigma_H E_x$ . The robustness of this phenomenon is remarkable: disorder and (weak) interactions do not alter the result, thus providing an exceptionally precise measurement [51] of  $e^2/h$ . Further remarkable robustness is offered by the fact that the Kubo-formula, derived from linear response theory, seems to extend its regime of validity well beyond linear response: as mathematically proven in Ref. [64], and further discussed in Ref. [65], all power-law corrections in the electric field can be shown, under suitable hypotheses, to be missing.

The availability of new experiments employing synthetic gauge fields [108, 110], directly sensitive to the time-dependent transient leading to the transverse response, calls for further scrutiny of this issue. Experimentally, the driving force  $F_x(t)$  in the  $\hat{x}$ -direction can be turned on, as a function of time, with some freedom, either abruptly or in a more or less smooth fashion. On the theory side, we can represent such a force in different gauges: quite conveniently, for a finite-length system  $L_x$  with periodic boundary conditions (PBC) in the  $\hat{x}$ -direction, we can choose a vector-potential gauge in which the force is represented by a time-dependent vector potential. The minimal coupling requires, in a tight-binding scheme, the Peierls' substitution:

$$\hat{c}_{l+1,m}^\dagger \hat{c}_{l,m} \longrightarrow e^{-ia\kappa_x(t)} \hat{c}_{l+1,m}^\dagger \hat{c}_{l,m}, \quad (3.5)$$

where  $\kappa_x(t)$  determines the force  $F_x(t)$  acting in the  $x$ -direction through  $F_x(t) = \hbar\kappa_x(t)$ , hence making the Hamiltonian time-dependent,  $\hat{H}(t)$ . More in detail, we chose  $F_x(t) = Ff(t/t_0)$ , where  $f(s = t/t_0)$  is a switch-on function interpolating between 0 and 1, i.e., such that  $f(s \leq 0) = 0$  and  $f(s \geq 1) = 1$ , and  $F$  is the stationary value of the force, attained for  $t \geq t_0$ . This choice leads to  $\kappa_x(t \leq t_0) = (t_0 F/\hbar) \int_0^{t/t_0} ds f(s)$  and  $\kappa_x(t \geq t_0) = \kappa_x(t_0) + F(t - t_0)/\hbar$ . The case of a sudden switch-on of the force is recovered by taking  $t_0 = 0$ . Since  $\kappa_x(t)$  appears in the hopping as a phase-factor, see Eq. (3.5), its linear increase for  $t \geq t_0$  implies that the

Hamiltonian becomes *time-periodic* for  $t \geq t_0$ ,  $\hat{H}(t + \tau) = \hat{H}(t)$ , with the period  $\tau$  given by

$$\tau = \frac{2\pi\hbar}{aF}, \quad (3.6)$$

which corresponds to a fundamental frequency  $\hbar\omega = aF$  entering the problem. These considerations clearly show that the question of the validity of linear response in  $F$  goes hand-in-hand with the issue of *adiabaticity* of  $\hat{H}(t)$ : Kubo linear response is essentially obtained in the fully adiabatic limit  $\omega \rightarrow 0$ .

To calculate the current, following Laughlin [52], we use PBC in the  $\hat{y}$ -direction as well, introducing a vector potential, again with a minimal-coupling Peierls' substitution:

$$\hat{c}_{l,m+1}^\dagger \hat{c}_{l,m} \longrightarrow e^{-ia\kappa_y} \hat{c}_{l,m+1}^\dagger \hat{c}_{l,m}. \quad (3.7)$$

The total current operator is obtained as a derivative of  $\hat{H}$  with respect to  $\kappa_y$ :

$$\hat{J}_y = \frac{1}{\hbar} \left. \frac{\partial \hat{H}}{\partial \kappa_y} \right|_{\kappa_y=0}. \quad (3.8)$$

The Hall response can now be seen as a non-vanishing quantum average of  $\hat{J}_y$  in presence of a force  $F_x$ , describing the transport of particles along  $\hat{y}$ -direction. We can quantify this through the linear-density of transported particles during the interval  $[t_0, t]$  (dropping the initial switching-on interval  $[0, t_0]$ ):

$$Q_y(t \geq t_0) = \int_{t_0}^t dt' \langle \psi(t') | \hat{j}_y | \psi(t') \rangle, \quad (3.9)$$

where  $|\psi(t)\rangle$  denotes the time-evolving state of the system, and  $\hat{j}_y = \hat{J}_y / (L_x L_y)$  is the space-averaged current density. Notice that  $Q_y$ , as defined, gives the number of particles per-unit-length moving along the  $\hat{y}$ -direction in the interval  $[t_0, t]$ : we will often refer to it as *pumped charge*, although the particles could be neutral.

We can ask for the charge pumped in the  $m$ -th period:

$$Q_m = Q_y(t_0 + m\tau) - Q_y(t_0 + (m-1)\tau). \quad (3.10)$$

We expect that the charge pumped in the initial periods  $Q_1, Q_2, \dots$  might be affected by *transient effects*, depending on the details of the switching-on function  $f(t/t_0)$  and time  $t_0$ . These transient effects are expected to decay for  $m \rightarrow \infty$  so that the infinite-time average

$$\bar{Q} = \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M Q_m, \quad (3.11)$$

should effectively capture the asymptotic (steady state) single-period pumped charge, where topological effects should most appropriately be looked for [66]. The Floquet theorem enormously simplifies the calculation of the infinite-time average  $\bar{Q}$ . Indeed, the state of the system at any time  $t \geq t_0$  can be expanded in terms of Floquet modes and quasi-energies [6, 111] as:

$$|\psi(t)\rangle = \sum_{\nu} e^{-i\epsilon_{\nu}(t-t_0)/\hbar} |u_{\nu}(t)\rangle \langle u_{\nu}(t_0) | \psi(t_0)\rangle \quad (3.12)$$

where  $\epsilon_\nu$  are the Floquet quasi-energies and  $|u_\nu(t)\rangle$  the associated time-periodic Floquet modes,  $|u_\nu(t+\tau)\rangle = |u_\nu(t)\rangle$ . A rather standard derivation [65, 112] shows that the infinite-time average pumped charge is dominated by the *Floquet diagonal ensemble* value:

$$\bar{Q} \equiv Q_d = \sum_\nu n_\nu \int_{t_0}^{t_0+\tau} dt' \langle u_\nu(t') | \hat{j}_y | u_\nu(t') \rangle. \quad (3.13)$$

where  $n_\nu = |\langle u_\nu(t_0) | \psi(t_0) \rangle|^2$  is the ‘‘occupation’’ of the  $\nu$ -th Floquet mode. This clearly shows that the initial preparation, with the transient loading interval  $[0, t_0]$ , is all contained in the occupation factors  $n_\nu$ .

So far, we have made use of time-periodicity, but not of translational invariance. To proceed, we make a rational choice of the magnetic flux,  $\alpha = p/q$  with  $p$  and  $q$  co-prime integers, which leads to an enlarged ‘‘magnetic’’ unit cell of size  $qa$  in the  $x$ -direction. We now label the sites in the  $x$ -direction with a cell-index  $j = 0 \cdots N_x - 1$  and an intra-cell index  $b = 0, 1, \cdots, q - 1$ , so that  $l = qj + b$ , while  $m = 0 \cdots N_y - 1$  labels sites in the  $y$ -direction. Hence,  $L_x = N_x qa$ , and  $L_y = N_y a$ . We then define appropriate Bloch combinations of the form:

$$\begin{cases} \hat{c}_{\mathbf{k},b}^\dagger = \frac{1}{\sqrt{N}} \sum_{j=0}^{N_x-1} \sum_{m=0}^{N_y-1} e^{ia(k_x(qj+b)+k_y m)} \hat{c}_{qj+b,m}^\dagger \\ \hat{c}_{qj+b,m}^\dagger = \frac{1}{\sqrt{N}} \sum_{\mathbf{k}}^{\text{BZ}} e^{-ia(k_x(qj+b)+k_y m)} \hat{c}_{\mathbf{k},b}^\dagger \end{cases}, \quad (3.14)$$

where  $\mathbf{k} = \frac{2\pi}{a} \left( \frac{n_x}{qN_x} \hat{\mathbf{x}} + \frac{n_y}{N_y} \hat{\mathbf{y}} \right)$ , with  $n_x = 0, \cdots, N_x - 1$  and  $n_y = 0, \cdots, N_y - 1$ , define the  $N = N_x N_y$  wave-vectors inside the Brillouin Zone (BZ):  $[0, \frac{2\pi}{qa}] \times [0, \frac{2\pi}{a}]$ . The Hamiltonian for the system can then be written in the form:

$$\begin{aligned} \hat{H}(t) &= -J_0 \sum_{\mathbf{k}} \sum_{b=0}^{q-1} \left\{ 2 \cos \left( ak_y + \frac{2\pi p}{q} b \right) \hat{c}_{\mathbf{k},b}^\dagger \hat{c}_{\mathbf{k},b} \right. \\ &\quad \left. + \left[ e^{-ia(k_x + \kappa_x(t))} \hat{c}_{\mathbf{k},b+1}^\dagger \hat{c}_{\mathbf{k},b} + \text{H.c.} \right] \right\} \\ &= \sum_{\mathbf{k}}^{\text{BZ}} (\hat{c}_{\mathbf{k},0}^\dagger \cdots \hat{c}_{\mathbf{k},q-1}^\dagger) \cdot \mathbb{H}(\mathbf{k}, t) \cdot \begin{pmatrix} \hat{c}_{\mathbf{k},0} \\ \vdots \\ \hat{c}_{\mathbf{k},q-1} \end{pmatrix}, \end{aligned} \quad (3.15)$$

*i.e.*, effectively a  $q \times q$  matrix problem  $\mathbb{H}(\mathbf{k}, t)$  for every  $\mathbf{k}$ -vector in the BZ. The total current operator has a similar expression:

$$\begin{aligned} \hat{J}_y &= \frac{2aJ_0}{\hbar} \sum_{\mathbf{k}} \sum_{b=0}^{q-1} \sin \left( ak_y + \frac{2\pi p}{q} b \right) \hat{c}_{\mathbf{k},b}^\dagger \hat{c}_{\mathbf{k},b} \\ &= \sum_{\mathbf{k}}^{\text{BZ}} (\hat{c}_{\mathbf{k},0}^\dagger \cdots \hat{c}_{\mathbf{k},q-1}^\dagger) \cdot \mathbb{J}(\mathbf{k}) \cdot \begin{pmatrix} \hat{c}_{\mathbf{k},0} \\ \vdots \\ \hat{c}_{\mathbf{k},q-1} \end{pmatrix}, \end{aligned} \quad (3.16)$$

where  $\mathbb{J}(\mathbf{k}) = (1/\hbar) \partial \mathbb{H} / \partial k_y$ .

From now on, we will concentrate our study on the case  $\alpha = 1/3$ , where the Hamiltonian becomes a  $3 \times 3$  problem for every  $\mathbf{k}$ . Exploiting the  $\mathbf{k}$ -factorization of the initial state  $|\psi(0)\rangle$  and of the subsequent dynamics, using that the space-averaged current density is  $\hat{\mathbf{j}}_y = \hat{\mathbf{J}}_y/(L_x L_y)$ , and transforming the sum over  $\mathbf{k}$  into an integral on the BZ in the usual fashion for a large system, we can rewrite the asymptotic pumped charge as:

$$Q_d = \sum_{\nu} \int_{\text{BZ}} \frac{d^2\mathbf{k}}{(2\pi)^2} n_{\mathbf{k},\nu} \int_{t_0}^{t_0+\tau} dt' \langle u_{\mathbf{k},\nu}(t') | \frac{1}{\hbar} \frac{\partial \mathbb{H}}{\partial k_y} | u_{\mathbf{k},\nu}(t') \rangle, \quad (3.17)$$

where

$$n_{\mathbf{k},\nu} = |\langle u_{\mathbf{k},\nu}(t_0) | \psi_{\mathbf{k}}(t_0) \rangle|^2. \quad (3.18)$$

A generalization of the Hellman-Feynman theorem for the Floquet case [111] shows that the average current carried by a Floquet mode is easily expressed in terms of the quasi-energy velocity:

$$\int_{t_0}^{t_0+\tau} dt' \langle u_{\mathbf{k},\nu}(t') | \frac{\partial \mathbb{H}}{\partial k_y} | u_{\mathbf{k},\nu}(t') \rangle = \tau \frac{\partial \epsilon_{\mathbf{k},\nu}}{\partial k_y}. \quad (3.19)$$

Hence  $Q_d$  in Eq. (3.17) can be re-expressed as:

$$Q_d = \frac{\tau}{\hbar} \sum_{\nu} \int_{\text{BZ}} \frac{d^2\mathbf{k}}{(2\pi)^2} n_{\mathbf{k},\nu} \frac{\partial \epsilon_{\mathbf{k},\nu}}{\partial k_y}. \quad (3.20)$$

Henceforth we will refer to the quantity defined in Eq. (3.20) as the *diagonal pumped charge*. We should stress that both the occupation factors  $n_{\mathbf{k},\nu}$  and the quasi-energies  $\epsilon_{\mathbf{k},\nu}$  are dependent on the time-periodicity  $\tau$ , or the frequency  $\omega$ , a dependence that we have not explicitly indicated.

It is useful to make here the connection with the time-averaged current density mentioned in the introduction, Eq. (3.1). Using Eq. (3.11) and the fact that  $\bar{Q} \equiv Q_d$ , see Eq. (3.13), it is straightforward to derive that

$$j_y = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{t_0}^{t_0+T} dt' \langle \psi(t') | \hat{\mathbf{j}}_y | \psi(t') \rangle \equiv \frac{Q_d}{\tau}. \quad (3.21)$$

Hence, using the relationship (3.6) between the period  $\tau$  and the force  $F$ , it is simple to show that:

$$j_y = \sigma_{yx} F = \frac{a Q_d}{h} F. \quad (3.22)$$

Hence, the transverse Hall ‘‘conductance’’ is here given by  $\sigma_{yx} = (a Q_d)/h$ . Its quantization, in units of  $1/h$ , would require that  $a Q_d = n$ , an integer.

A final comment concerning transient effects. To appreciate them, the asymptotic pumped charge  $Q_d$  should be contrasted with the charge pumped in the  $m$ -th period, which would read:

$$Q_m = \frac{1}{\hbar} \int_{\text{BZ}} \frac{d^2\mathbf{k}}{(2\pi)^2} \int_{t_0+(m-1)\tau}^{t_0+m\tau} dt' \langle \psi_{\mathbf{k}}(t') | \frac{\partial \mathbb{H}}{\partial k_y} | \psi_{\mathbf{k}}(t') \rangle. \quad (3.23)$$

### 3.3. Robustness of Hall quantization to non adiabatic effects

To illustrate the previous general considerations, let us consider the case of a sudden switch-on of a constant force  $F_x(t) = \theta(t)F$ , which effectively amounts to taking  $t_0 = 0$  in the

previous expressions. The bands of the unperturbed Hamiltonian are shown in Fig. 3.4 (top), for  $N_x = 30$ , as a function of  $k_y \in (-\pi/a, \pi/a]$ : we see  $q = 3$  distinct bands, obtained by projecting the  $N_x$  different values of  $k_x$ . The initial insulating state is a Slater determinant  $|\psi(0)\rangle$  obtained by completely filling one such band, for instance the lowest one. We then calculate the charge pumped in the first period:

$$Q_1 = \frac{1}{\hbar} \int_{\text{BZ}} \frac{d^2\mathbf{k}}{(2\pi)^2} \int_0^\tau dt' \langle \psi_{\mathbf{k}}(t') | \frac{\partial \mathbb{H}}{\partial k_y} | \psi_{\mathbf{k}}(t') \rangle. \quad (3.24)$$

Fig. 3.4 (bottom) shows the value of  $Q_1$  as a function of the driving  $F$ , expressed in terms of  $\hbar\omega = aF$ . Notice that for  $F \rightarrow 0$  we recover, as expected, a pumped charge which is




Figure 3.4.: Top: Energy bands of the Harper-Hofstadter model for  $\alpha = 1/3$  vs  $k_y$  for  $N_x = 30$ . The dashed red lines represent the energy averaged over the phase  $k_x + \kappa_x(t)$ . Bottom: Charge pumped in the first period in each magnetic cell,  $3aQ_1$ , as a function of the driving field  $aF/J_0$ , where  $\hbar\omega = aF$ , after a sudden switch-on of the driving. For  $F \rightarrow 0$ ,  $3aQ_1$  is quantized to the first Chern number, respectively  $+3, -6, +3$ , of the band in which the system is initially prepared. (The simulation has been repeated by preparing the initial Slater determinant insulating state to be one of the three completely filled bands, in order to compute the Chern numbers.) The first and the third band give exactly the same response. This figure is essentially equivalent to Figure 1 of Ref. [65], where the abscissa is  $1/\omega$ .

quantized to the integer Chern numbers ( $+3, -6$  and  $+3$ ) of the three completely filled bands.

Deviations from perfect quantization are clearly visible at finite  $\omega$ : the remaining part of the chapter is precisely devoted to understanding the nature and size of these deviations.

In order to proceed with the analysis of the deviations from perfect quantization for small  $\omega$ , we shift our attention to the infinite-time average of the pumped charge, where the Floquet theory helps in elucidating the crucial ingredients. Our starting point is hence Eq. (3.20), which we rewrite below for convenience in a slightly different form:

$$Q_d = \frac{1}{\hbar\omega} \sum_{\nu} \int_{\text{BZ}} \frac{d^2\mathbf{k}}{2\pi} n_{\mathbf{k},\nu} \frac{\partial \epsilon_{\mathbf{k},\nu}}{\partial k_y}. \quad (3.25)$$

In this re-writing we have used the trivial fact that  $\tau = 2\pi/\omega$  and that the Floquet quasi-energies  $\epsilon_{\mathbf{k},\nu}$  *dependent only on  $k_y$* : to appreciate the last fact, observe that the dependence of the Hamiltonian, see Eq. (3.15), on  $k_x$  and  $t$  is all contained in the phase-factor  $e^{-ia(k_x + \kappa_x(t))} = e^{-iak_x(t_0)} e^{-i\omega(t-t_x)}$ , with  $t_x = t_0 - ak_x/\omega$ . Hence, different values of  $k_x$  effectively correspond to a shift in time  $t_0 \rightarrow t_0 - t_x$ , which in turn amounts to a unitary transformation on the Floquet operator  $\hat{U}_{\mathbf{k}}(t_0 + \tau, t_0)$ , whose eigenvector/eigenvalues are the Floquet modes/quasi-energies:

$$\hat{U}_{\mathbf{k}}(t_0 + \tau, t_0) |u_{\mathbf{k},\nu}(t_0)\rangle = e^{-i\epsilon_{\mathbf{k},\nu}\tau/\hbar} |u_{\mathbf{k},\nu}(t_0)\rangle. \quad (3.26)$$

As discussed in App. A.2, the fact that the Floquet operators at different  $k_x$  are *unitarily equivalent* implies that their eigenvalues are  $k_x$ -independent, i.e.,  $e^{-i\epsilon_{\mathbf{k},\nu}\tau/\hbar}$ . Notice that, on the contrary, the Floquet modes  $|u_{\mathbf{k},\nu}(t)\rangle$ , and hence the occupations  $n_{\mathbf{k},\nu} = |\langle u_{\mathbf{k},\nu}(t_0) | \psi_{\mathbf{k}}(t_0) \rangle|^2$ , *do depend* on  $k_x$ .

### 3.3.1. Pumping of Floquet states.

The first issue we address is what happens to the pumped charge if the initial state  $|\psi(t_0)\rangle$  is *precisely prepared to be the  $\nu$ th Floquet state*, i.e., such that  $n_{\mathbf{k},\nu'} = \delta_{\nu,\nu'}$ . Then, the corresponding value of the pumped charge is:

$$Q_{\nu}^F = \frac{1}{3a\hbar\omega} \int_0^{\frac{2\pi}{a}} dk_y \frac{\partial \epsilon_{\mathbf{k},\nu}}{\partial k_y}, \quad (3.27)$$

where we eliminated the trivial integral on  $k_x \in [0, \frac{2\pi}{3a}]$ . To better understand the physical implications of this formula, let us start from the extreme adiabatic limit  $\omega \rightarrow 0$ , where the predictions of the adiabatic theorem give us a 0th-order expression for the quasi-energies, in an extended Floquet BZ scheme [6], in the form:

$$\begin{aligned} \epsilon_{\mathbf{k},\nu}^{(0)} &= \frac{1}{\tau} \int_0^{\tau} dt [E_{\mathbf{k},\nu}(t) - i\hbar \langle \phi_{\mathbf{k},\nu}(t) | \partial_t \phi_{\mathbf{k},\nu}(t) \rangle] \\ &= \epsilon_{\mathbf{k},\nu}^d + \epsilon_{\mathbf{k},\nu}^g. \end{aligned} \quad (3.28)$$

Here  $\phi_{\mathbf{k},\nu}(t)$  and  $E_{\mathbf{k},\nu}(t)$  are the instantaneous eigenstates/eigenvalues of  $\hat{H}(t)$ , while  $\epsilon_{\mathbf{k},\nu}^d$  and  $\epsilon_{\mathbf{k},\nu}^g$  denote dynamical and geometric contributions [107]. See section 2.2, in particular Eqs. (2.24) and (2.25), for a derivation of Eq. (3.28). These contributions are in turn expressed

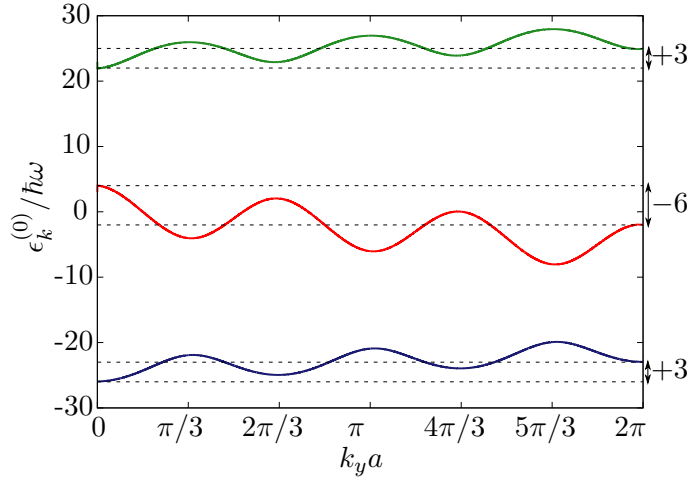


Figure 3.5.: Adiabatic quasi-energies  $\epsilon_{k_y, \nu}^{(0)}$  in units of  $\hbar\omega$  for the three bands of the Harper-Hofstadter model. The frequency is  $\hbar\omega = 0.1J_0$ .  $\epsilon_{k_y, \nu}^{(0)}$  is the sum of the  $k_x$ -averaged band, dashed line in Fig. 3.4 (upper panel), plus the geometric contribution, giving rise to the loss of periodicity for  $\epsilon_{k_y, \nu}^{(0)}$  in the BZ:  $\epsilon_{\frac{2\pi}{a}, \nu}^{(0)} = \epsilon_{0, \nu}^{(0)} + \hbar\omega C_\nu$ , where  $C_\nu$  is the Chern number of the  $\nu$ -th band.

as:

$$\begin{aligned}\epsilon_{k_y, \nu}^d &= a \int_0^{\frac{2\pi}{a}} \frac{dk_x}{2\pi} E_{\mathbf{k}, \nu}(0) \\ \epsilon_{k_y, \nu}^g &= -\hbar\omega \int_0^{\frac{2\pi}{a}} \frac{dk_x}{2\pi} \mathcal{A}_x^{(\nu)}(\mathbf{k}).\end{aligned}\quad (3.29)$$

where  $\mathcal{A}_x^{(\nu)}(\mathbf{k}) = i\langle\phi_{\mathbf{k}, \nu}(0)|\partial_{k_x}\phi_{\mathbf{k}, \nu}(0)\rangle$  is the Berry connection of the  $\nu$ -th band. In both terms, the time integral has been transformed into a  $k_x$ -integral using the fact that the dependence on  $t$  is through the variable  $ak_x + \omega t$ . As a consequence, both terms are functions of  $k_y$  only. The dynamical contribution is the  $k_x$ -averaged Bloch band, and is strictly periodic in  $k_y$  of the BZ, see dashed red line in Fig. 3.4. On the contrary, the geometric term *winds* over the BZ, ending up acquiring an overall integer equal to the Chern number of the corresponding band:

$$\epsilon_{\frac{2\pi}{a}, \nu}^g - \epsilon_{0, \nu}^g = \hbar\omega C_\nu. \quad (3.30)$$

This immediately leads to the expected integer quantization

$$3aQ_\nu^{(0)} = \frac{1}{\hbar\omega} \int_0^{\frac{2\pi}{a}} dk_y \frac{\partial \epsilon_{k_y, \nu}^{(0)}}{\partial k_y} = C_\nu. \quad (3.31)$$

We now consider finite- $\omega$  effects beyond the adiabatic limit. Fig. 3.6 shows the Floquet quasi-energy bands for  $\hbar\omega/J_0 = 0.1$ , plotted versus  $k_y$  in the region  $[0, 2\pi/(3a)]$ , due to a periodicity  $\epsilon_{k_y + \frac{2\pi}{3a}, \nu} = \epsilon_{k_y, \nu}$  discussed in App. A.2. Notice that the quasi-energies are here naturally represented in the *Floquet Brillouin Zone* [6]  $[-\hbar\omega/2, \hbar\omega/2]$ , as they are obtained

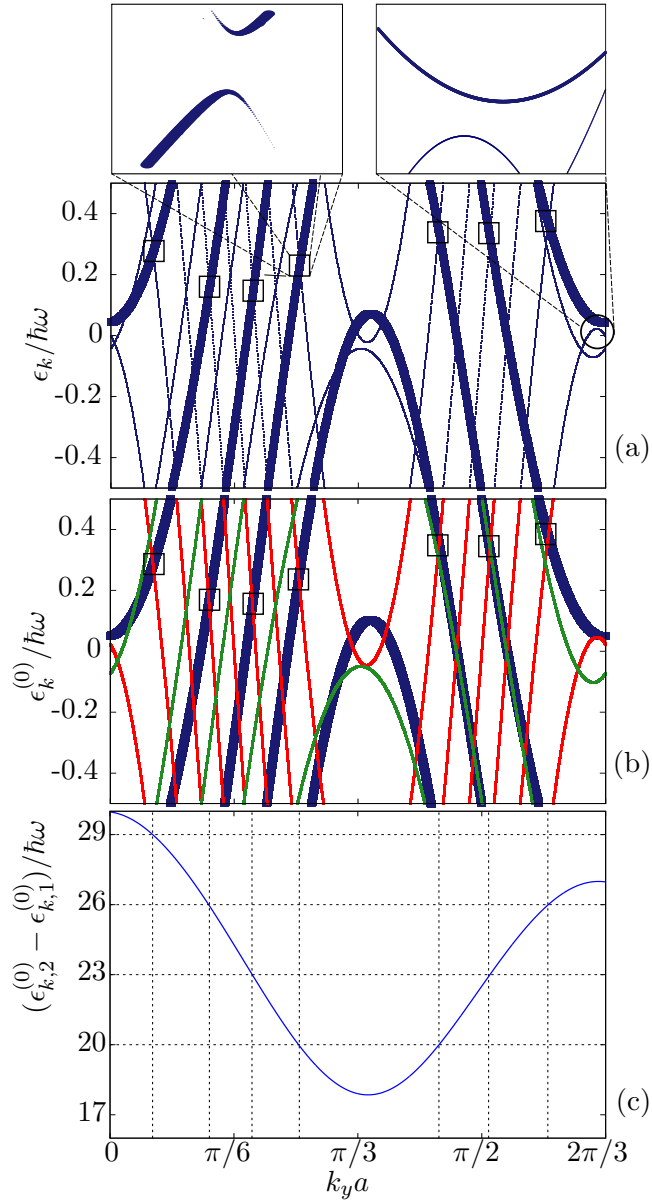


Figure 3.6.: (a) Floquet spectrum as a function of  $k_y$ , for  $\hbar\omega/J_0 = 0.1$ . The squares signal the Floquet resonance avoided crossings, the circle an ordinary avoided crossing. Both are magnified in the top insets, where the size of the points is proportional to the  $k_x$ -averaged occupation  $n_{k_y,\nu}$ , see Eq. (3.33). (b) Floquet adiabatic quasi-energies, Eq. (3.28), folded in the Floquet BZ. c)  $(\epsilon_{k,2}^{(0)} - \epsilon_{k,1}^{(0)}) / \hbar\omega$ , the energy difference between the two lowest adiabatic bands in the extended-zone scheme, for  $\hbar\omega = 0.1J_0$ . The vertical lines indicate the *Floquet resonances*,  $\epsilon_{k,2}^{(0)} - \epsilon_{k,1}^{(0)} = m\hbar\omega$  with  $m = 29, 26, 23, 20$ , giving rise to the avoided crossing gaps of panel (a).

by a numerical diagonalization of the Floquet operator. The thick line represents the quasi-energy band emerging from the low-energy band of Fig. 3.5. We observe two conspicuous features:



- i) an apparent *winding* over the Floquet Brillouin Zone, as a quasi-energy crossing  $\hbar\omega/2$  re-enters at  $-\hbar\omega/2$  (and vice-versa). This is the winding expected from the geometric contribution to the adiabatic quasi-energies shown in Fig. 3.5. It would lead to:

$$\frac{1}{\hbar\omega} \int_0^{\frac{2\pi}{a}} dk_y \frac{\partial \epsilon_{k_y, \nu}}{\partial k_y} = C_\nu, \quad (3.32)$$

where  $C_\nu$  is the Chern number of the  $\nu$ -th band ( $C_\nu = +3$ , for the thick band shown in Fig. 3.6).

- ii) an apparent *crossing* of quasi-energies belonging to different Floquet bands.

The crossings between different Floquet bands can develop very small *anti-crossing gaps* [65, 113], as indeed we find at the points signaled by a square (see inset of Fig 3.6). To better understand the nature of such anti-crossing gaps, we reconsider again the adiabatic bands. The central panel of Fig 3.6 shows a plot of the adiabatic bands  $\epsilon_{k_y, \nu}^{(0)}$  folded back into the Floquet BZ: quite evidently, they are a good approximation to the true quasi-energies for such a value of  $\omega$ . Notice, however, that here all the band crossings are genuine ones. The bottom panel of Fig 3.6, finally, shows  $\epsilon_{k_y, 2}^{(0)} - \epsilon_{k_y, 1}^{(0)}$ , the energy difference between the two lowest adiabatic bands, which clearly suggests that the anti-crossing points — signalled by vertical dashed lines — are associated to *Floquet resonances* when  $\epsilon_{k_y, 2}^{(0)} - \epsilon_{k_y, 1}^{(0)} = m\hbar\omega$ . Surprisingly, not all possible resonances actually lead to the opening of an anti-crossing gap, but only a sequence of them, here with  $m = 29, 26, 23, 20$ . The periodicity of  $\Delta m = 3$  is likely associated with our choice of flux  $\alpha = 1/3$ , but the precise location of the resonance is not fully understood. One thing that we can say, however, is that the resonances open up gaps [114] in the quasi-energy spectrum, which are *exponentially small in  $1/\omega$* . This makes such gaps quite difficult to pinpoint precisely, but our numerical evidence is reasonably robust on that issue. Fig. 3.7 shows the deviation from integer quantization,  $3 - 3aQ_\nu^F$ , — calculated assuming  $n_{\mathbf{k}, \nu} = 1$  and using Eq. (3.17), which, as opposed to Eq. (3.27), avoids derivatives of numerically determined quasi-energies — as a function of  $J_0/\hbar\omega$ : in the  $\omega$ -region we plot, an overall exponential decay is clearly visible for  $3 - 3aQ_\nu^F \sim e^{-\gamma J_0/(\hbar\omega)}$ , with  $\gamma \sim 0.5$ , superimposed on a saw-tooth behaviour due to the sudden formation of larger gaps when two nearby gaps coalesce together upon decreasing  $\omega$ . Summarizing, if the quasi-energy avoided-crossing gaps opening were the main responsible for finite-frequency/field corrections to the quantized pumped charge, such deviations would be exponentially small in  $1/\omega \propto 1/F$ : therefore non-analytic in the field strength [65] and exceedingly small for most practical purposes: for instance, in an experiment in which  $\hbar\omega = 10^{-2}J_0$ , we would estimate  $3 - 3aQ_\nu^F \approx 10^{-22}$ .

### 3.3.2. Effect of the occupation factors: sudden switch-on.

The second source of deviations from perfect quantization arises from the fact that the prepared state  $|\psi(t_0)\rangle$  is *not precisely a Floquet state*, i.e., that Floquet occupation factors deviate from  $n_{\mathbf{k}, \nu} = \delta_{\nu, \nu'}$ . The inset of Fig. 3.6, where the size of the dots is proportional to the Floquet occupation, shows that sizeable deviations occur whenever  $\omega > 0$ , even if small, at the quasi-energy avoided level crossing. Indeed, for a quasi-adiabatic evolution, the Floquet modes will be “close” to the eigenstates of the Hamiltonian, to which they reduce for  $\omega \rightarrow 0$ .




Figure 3.7.: The deviation of  $3aQ_\nu^F$  from integer quantization,  $3 - 3aQ_\nu^F$ , for the lowest Floquet band, assuming  $n_{\mathbf{k},\nu'} = \delta_{\nu,\nu'}$ , showing that the exponentially small gaps in the quasi-energy spectrum at finite  $\omega$  lead to exponentially small deviations.

If we initialize the system in an insulating phase by filling the lowest-energy band, one of the Floquet occupation number  $n_{\mathbf{k},\nu}$  will be close to 1 and much higher than the others: the corresponding Floquet mode will be the main one responsible for charge transport. In the following, we will refer to such a state as *adiabatic* or *lowest-energy* Floquet state: it is indeed the Floquet state which has the largest overlap with the instantaneous Hamiltonian ground state. This is highlighted in the Fig. 3.6, where the Floquet spectrum is plotted vs  $k_y$  with thickness proportional to  $k_x$ -averaged occupation factor

$$n_{k_y,\nu} = \frac{3a}{2\pi} \int_0^{\frac{2\pi}{3a}} dk_x n_{\mathbf{k},\nu}. \quad (3.33)$$

Let us now focus on the occupation of such “adiabatic” Floquet state. If the driving field is suddenly turned on from  $F_x(t \leq 0) = 0$  to  $F_x(t > 0) = F$ ,  $|\psi(t_0 = 0)\rangle$  coincides with a Slater determinant Bloch eigenstate of  $\hat{H}(0)$ , and  $n_{\mathbf{k},\nu}$  is given by the overlap of such a state with the adiabatic Floquet state:  $n_{\mathbf{k},\nu} = |\langle u_{\mathbf{k},\nu}(0) | \phi_{\mathbf{k},\nu} \rangle|^2$ . When  $\omega$  is small, we can combine adiabatic perturbation theory [115] (APT) to obtain an approximate expression for the Floquet modes  $|u_{\mathbf{k},\nu}(0)\rangle$ , see App. A.1 for details. Following this approach  $n_{\mathbf{k},\nu}$  can be calculated to be:

$$n_{\mathbf{k},\nu} = 1 - \left( \frac{\hbar\omega}{2\pi} \right)^2 \sum_{\mu \neq \nu} \left| \frac{M_{\mu,\nu}^{(\mathbf{k})}}{\Delta_{\mu,\nu}^{(\mathbf{k})}} \right|^2 + O(\omega^3). \quad (3.34)$$

Here  $M_{\mu,\nu}^{(\mathbf{k})}$  and  $\Delta_{\mu,\nu}^{(\mathbf{k})}$  are calculated from instantaneous Hamiltonian eigenvalues/eigenstates,  $\hat{H}_{\mathbf{k}}(s) |\phi_{\mathbf{k},\nu}(s)\rangle = E_{\mathbf{k},\nu}(s) |\phi_{\mathbf{k},\nu}(s)\rangle$  where  $s = \omega t$  is the rescaled time, as:

$$\begin{aligned} \Delta_{\mu,\nu}^{(\mathbf{k})}(s) &= E_{\mathbf{k},\mu}(s) - E_{\mathbf{k},\nu}(s), \\ M_{\mu,\nu}^{(\mathbf{k})}(s) &= \frac{\langle \phi_{\mathbf{k},\mu}(s) | \partial_s \hat{H}_{\mathbf{k}}(s) | \phi_{\mathbf{k},\nu}(s) \rangle}{\Delta_{\nu,\mu}^{(\mathbf{k})}}. \end{aligned} \quad (3.35)$$

In Eq. (3.34) all quantities are evaluated at  $s = 2\pi$ , corresponding to  $t = \tau$ , a full period. Therefore, if the matrix elements  $M_{\mu,\nu}$  are not all equal to zero, which in general they are




Figure 3.8.: Correction to the  $\mathbf{k}$ -averaged adiabatic Floquet mode occupation  $n_\nu$ , Eq. (3.36), vs  $1/\omega$ , showing the good agreement between the numerical data and the perturbation theory prediction from Eq. (3.34).

not, we expect to see power-law corrections to the occupation number of the Floquet states, leading to a similar behavior for the pumped charge. Fig. 3.8 shows the  $\mathbf{k}$ -averaged occupation

$$n_\nu = 3a^2 \int_{\text{BZ}} \frac{d^2\mathbf{k}}{(2\pi)^2} n_{\mathbf{k},\nu}, \quad (3.36)$$

calculated numerically, compared to the perturbation theory estimate in Eq. (3.34), as a function of  $\omega$ : the  $\omega^2$  deviation is quite clearly visible. This quadratic correction to the occupation factors reflects itself into the pumped charge, both the single-period charge  $Q_1$ , Eq. (3.24), as well as the infinite-time average  $Q_d$ , Eq. (3.25), as seen from Fig. 3.9. The faster-and-faster oscillations seen in  $Q_1$  for  $\omega \rightarrow 0$  originate from the essential singularity in  $\omega = 0$  of the expectation value of the current operator [65]; the oscillations are smeared in  $Q_d$ , due to the infinite-time average. This behavior is very similar to that reported in Ref. [75] for adiabatic quantum pumping in the Rice-Mele model.

### 3.3.3. Effect of the occupation factors: continuous switch-on.

The picture becomes richer if we switch-on the driving in a continuous fashion, taking  $F_x(t) = Ff(t/t_0)$  with a suitably smooth function  $f(s = t/t_0)$ . The first obvious choice is a linear switch-on,  $f(s) = s$ , with a fixed switch-on time  $t_0$ . As shown in Fig. 3.10, we now observe two regimes: a first one, for relatively large  $\omega$ , where the corrections to the occupation  $n_\nu$  of the adiabatic Floquet band appear to be exponentially small in  $1/\omega$ , and a second regime, for small  $\omega$ , where the corrections are  $\propto \omega^2$ :

$$1 - n_\nu \sim \begin{cases} A e^{-\frac{\gamma J_0}{\hbar\omega}} & \text{for } \omega > \omega^* \\ B \frac{\hbar^4 \omega^2}{J_0^4 t_0^2} & \text{for } \omega < \omega^* \end{cases}. \quad (3.37)$$

The two regimes have markedly different behaviors. The non-analytic exponential observed at higher  $\omega$  is *universal* — with  $\gamma \simeq 0.5$  from our data, and at most a very mild dependence




Figure 3.9.: Pumped charge vs  $\omega$ , both for one period (solid line)  $3aQ_1$ , Eq. (3.24), and in the diagonal ensemble (dashed line)  $3aQ_d$ , Eq. (3.25). Notice the oscillations in  $Q_1$ , signaling an essential singularity in  $\omega = 0$ . The inset shows the deviation from the quantized value  $3aQ_1(\omega \rightarrow 0) = 3$  vs  $1/\omega$ . For small frequency this deviation is quadratic in  $\omega$ .

of  $A$  on  $t_0$  — and, as we will argue, it is directly related to the width of the resonances of the Floquet spectrum. The power-law regime is non-universal, with an amplitude decreasing as  $1/t_0^2$ : hence the crossover frequency  $\omega^*$  between these two regimes, which is approximately given by:

$$\frac{J_0}{\hbar\omega^*} \simeq \log \left( \sqrt{\frac{A}{B}} \frac{J_0 t_0}{\hbar} \right), \quad (3.38)$$

is shifted towards smaller  $\omega$  as  $t_0$  increases. Notice that the crossover  $\omega^*$  exists only if  $\frac{J_0 t_0}{\hbar} \geq \sqrt{\frac{B}{A}} \frac{e}{4}$ ; indeed if  $t_0$  is too small, only the power law regime survives, leading to the ordinary “perturbative response” observed for the sudden quench case.

It is interesting to ask why the continuity in time of the force field  $F_x(t)$  is so important. As explained in Sec. 3.3.1 the topological properties at finite frequency are related to the Floquet states, while the system is initially prepared in a state  $|\psi(0)\rangle$  which coincides with the Hamiltonian ground state. By switching on the driving force in a continuous manner,  $F_x(t) = Ff(t/t_0)$ , the initial state is *continuously* deformed into a state which is “closer” to the “lowest-energy” Floquet state at the final frequency  $\omega$ . Fig. 3.11 helps to illustrate what happens as we turn on the driving frequency: as the instantaneous  $\omega(t) = aF_x(t)/\hbar$  increases, each Floquet mode winds around the expanding Floquet-Brillouin zone (FBZ) and encounters a series of (avoided) level crossings in the quasi-energy spectrum, with exponentially small gaps  $\Delta$ . Since the gaps  $\Delta$  are exponentially small, however, a finite value of  $t_0$  will lead the system to cross them *adiabatically*. The final Floquet state will show an occupation which can be interpreted [116] as the excitation probability after many Landau-Zener [117, 118] events. Following Ref. [116], at each avoided crossing we obtain a transition probability

$$P_{\text{ex}}(\omega, t_0) = e^{-\frac{\Delta^2 t_0}{4\hbar^2 \zeta \omega}}, \quad (3.39)$$

where we used the fact that the speed at which the gap is crossed can be estimated as  $\partial_t(\epsilon_2 - \epsilon_1) \simeq \zeta \hbar \omega / t_0$ ,  $\zeta$  being the difference in slope between the two quasi-energy bands

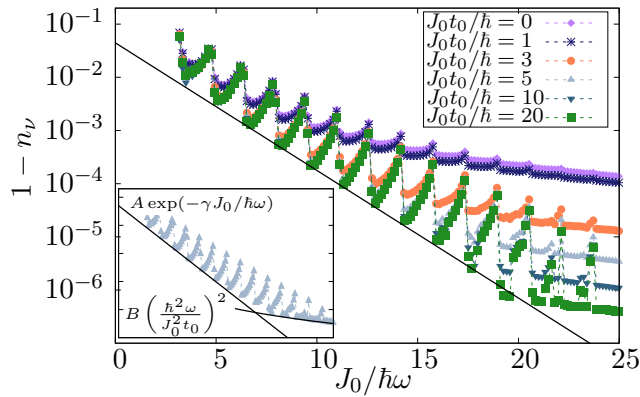


Figure 3.10.: Correction to the occupation of the lowest energy Floquet band for different switch-on times  $t_0$ . The inset highlights the crossover from an exponential regime  $e^{-1/\omega}$  to the quadratic one  $\omega^2$  for finite ramp time  $t_0$ . The solid lines correspond to the functions  $0.05e^{-\frac{0.5J_0}{\hbar\omega}}$  and  $0.002(\hbar\omega/J_0)^2$ .




Figure 3.11.: Floquet quasi-energies in  $\mathbf{k} = 0$  as a function of the frequency. The width of the line is proportional to the occupation of the state when the system is in the ground state with a filling factor  $1/3$ . The inset zooms on a level crossing to highlight the presence of gaps, showing also that the adiabatic band is the “excited” state after the avoided crossing. The solid black lines are the boundary of the first Floquet-Brillouine zone.




Figure 3.12.: Correction to the occupation of the  $\nu$ -th Floquet mode compared with the prediction of Eq. (3.42) (solid black lines).

as they wind around the Floquet BZ. Since the gaps  $\Delta(\omega) \sim J_0 e^{-\gamma J_0/\hbar\omega}$  are the smallest quantities, it is legitimate to expand the exponential in Eq. (3.39) to lowest order in  $\Delta^2$ . A further simplification is because the dominant contribution to the sequence of LZ processes comes from the largest gap encountered, which corresponds to the end of the ramp when the frequency is maximum. Hence we obtain the following estimate for the corrections to  $n_\nu$

$$1 - n_\nu \sim \frac{\Delta^2 t_0}{4\hbar^2 \zeta \omega} \sim \frac{J_0^2 t_0}{4\hbar^2 \zeta \omega} e^{-2\gamma J_0/\hbar\omega}. \quad (3.40)$$

This rather crude estimate gives a hint on the physical mechanism behind the non-perturbative corrections to the integer occupation of the Floquet mode observed when the electric field is turned on at a finite rate  $1/t_0$ . Incidentally, Eq. (3.40) also suggests that increasing the ramp time  $t_0$  would lead to *larger corrections* to both  $n_\nu$  and  $Q_d$ , although our numerical data do not show this, possibly because of the limited range of  $t_0$  explored. Indeed, at the level crossing, a larger  $t_0$  would increase the adiabaticity of the process, therefore decreasing the occupation of the lowest energy Floquet mode, which corresponds to the “excited” state in the quasi-energy spectrum, as shown in Fig. 3.11. An alternative explanation is given in Ref. [119], where it is suggested that the increasing deviations from the adiabatic preparation of Floquet states for very large ramp times  $t_0$  is related to the absorption of energy from the external field, leading to heating of the system.

This picture breaks down for small  $\omega$ , where the crossover with the quadratic regime occurs. The observed  $\omega^2$  scaling suggests that a Floquet adiabatic perturbation theory (FAPT) [2, 119, 120] might be appropriate here. Unfortunately, the standard framework of application of such a theory is when the slowly changed parameters  $\lambda(t)$  do not involve the crossing of Floquet resonances [119], which is certainly not the case for  $\omega \rightarrow 0$ . So, we construct here a simplified version of FAPT which should capture the  $\omega \rightarrow 0$  regime. To do so, we start from an expansion of the state  $|\psi_{\mathbf{k}}(t)\rangle$  in terms of instantaneous Floquet modes  $|u_{\mathbf{k},\mu}(\omega(t), t)\rangle$  corresponding to a frequency  $\omega(t)$  (which is slowly evolving in time), with associated phase

factor given by the *adiabatic* Floquet quasi-energy  $\epsilon_{\mathbf{k},\mu}^{(0)}(\omega(t))$ :

$$|\psi_{\mathbf{k}}(t)\rangle = \sum_{\mu} c_{\mathbf{k},\mu}(t) e^{-\frac{i}{\hbar} \int_0^t \epsilon_{\mathbf{k},\mu}^{(0)}(\omega(t'))} |u_{\mathbf{k},\mu}(\omega(t), t)\rangle. \quad (3.41)$$

Proceeding as in the standard APT, assuming that at  $t = 0$  we have  $c_{\mathbf{k},\mu}(0) = \delta_{\mu,\nu}$  and keeping only the lowest-order terms we end-up writing:

$$c_{\mathbf{k},\mu \neq \nu}(t_0) \approx - \int_0^{\omega} d\omega' \langle u_{\mathbf{k},\mu} | \partial_{\omega'} u_{\mathbf{k},\nu} \rangle e^{-\frac{it_0}{\hbar\omega} \int_0^{\omega'} (\epsilon_{\mathbf{k},\nu}^{(0)} - \epsilon_{\mathbf{k},\mu}^{(0)})}$$

where we assumed a linear adiabatic switch-on,  $\omega(t) = (t/t_0)\omega$ , and changed variable to an integral over frequency. Here  $|u_{\mathbf{k},\mu}\rangle$  stands for  $|u_{\mathbf{k},\mu}(\omega', t(\omega'))\rangle$ , where  $t(\omega') = t_0\omega'/\omega$ . Noticing now that the adiabatic quasi-energy differences  $(\epsilon_{\mathbf{k},\nu}^{(0)} - \epsilon_{\mathbf{k},\mu}^{(0)})$  are large compared to  $\omega'$ , we integrate by part, as in standard APT, ending up with:

$$c_{\mathbf{k},\mu \neq \nu}(t_0) \approx \frac{i\hbar\omega}{t_0} \frac{\langle u_{\mathbf{k},\mu}(\omega', t(\omega')) | \partial_{\omega'} u_{\mathbf{k},\nu}(\omega', t(\omega')) \rangle}{\epsilon_{\mathbf{k},\nu}^{(0)}(\omega') - \epsilon_{\mathbf{k},\mu}^{(0)}(\omega')} e^{-\frac{it_0}{\hbar\omega} \int_0^{\omega'} (\epsilon_{\mathbf{k},\nu}^{(0)} - \epsilon_{\mathbf{k},\mu}^{(0)})} \Big|_{\omega'=0}^{\omega'=\omega}. \quad (3.42)$$

Finally we compute the scalar products  $\langle u_{\mathbf{k},\mu} | \partial_{\omega'} u_{\mathbf{k},\nu} \rangle$  by using the expansion derived in App. A.1, in particular Eq. A.8, which allows us to write:

$$\langle u_{\mathbf{k},\mu} | \partial_{\omega'} u_{\mathbf{k},\nu} \rangle = \frac{M_{\mu,\nu}^{(\mathbf{k})}}{\Delta_{\mu,\nu}^{(\mathbf{k})}} + O(\omega'). \quad (3.43)$$

Substituting back into Eq. (3.42), we get an expression that can be computed numerically. Once the projections  $c_{\mathbf{k},\mu \neq \nu}$  have been computed, the correction to the occupation number of the ‘‘adiabatic’’ Floquet state reads

$$1 - n_{\nu} = 3a^2 \sum_{\mu \neq \nu} \int_{BZ} \frac{d^2\mathbf{k}}{(2\pi)^2} |c_{\mathbf{k},\mu}|^2. \quad (3.44)$$

As shown in Fig. 3.12 this simplified FAPT describes quite well the quadratic regime and its scaling with  $t_0$ . We observe that the accuracy of the approximation seems to decrease as  $t_0$  grows, probably because non-adiabatic corrections to the time-evolved eigenstates need to be taken into account in computing Eq. (3.43).

As a final check, we have considered whether imposing continuity also on the first derivative of  $\omega(t)$  makes any difference or not. Fig. 3.13 shows the occupation of the ‘‘adiabatic’’ Floquet state when the frequency is increased smoothly from 0 to its final value  $\omega$  with a switching function  $f(s) = \frac{1}{2}(1 - \cos(\pi s))$ . Beside some small numerical difference, the situation is qualitatively similar to the one obtained with the linear ramp (Fig. 3.10), with a crossover between an exponential regime for  $\omega > \omega^*(t_0)$  and a power law tail for  $\omega < \omega^*(t_0)$ . This suggests that while a *necessary* condition — albeit not sufficient — to obtain non-perturbative corrections is indeed the *continuity* of the force field  $F_x(t)$ , its differentiability seems not to be required.

## 3.4. Conclusions

In this chapter, we discussed the robustness of the quantization of the Hall conductivity beyond the validity range of linear response theory (adiabatic limit), in the Harper-Hofstadter




Figure 3.13.: Correction to the occupation number of the adiabatic Floquet state, when the driving force is smoothly turned on with the switching function  $f(s = t/t_0) = \frac{1}{2}(1 - \cos(\pi s))$ . Eq. (3.42), with a slight modification due to the different driving schedule, still gives a good estimate of the quadratic regime for small  $\omega$ .

model. Our work was mainly motivated by the possibility of realizing simple tight-binding Hamiltonians, such as the HH one, in cold atoms experiments with synthetic gauge fields, where the model parameters can be easily fine-tuned. By employing Floquet theory for time-periodic systems, we showed that the quantization of the transverse pumped charge  $Q_d$  depends mainly on the occupation factor  $n_\nu$  of the lowest energy Floquet state. In particular, we found that a continuous and sufficiently slow switching-on of the driving force is necessary to obtain corrections to the Kubo formula which are non-analytic in the force amplitude  $F$ , scaling as  $e^{-\gamma/|F|}$ . If the switching time  $t_0$  is too small, or the force is turned on abruptly, corrections of the order  $O(F^2)$  are always recovered when  $F \rightarrow 0$ . A crossover force amplitude  $F^*(t_0)$  between the quadratic and the exponential regimes is clearly shown by our numerical analysis for any finite switching time  $t_0$ , and it would be interesting to see if this crossover can be indeed be observed in experimental realizations of IQHE or quantum pumping in optical lattices experiments.

Another important element to complete the understanding of the physics of quantized pumping is how the robustness of the topological phase and the crossover with a perturbative regime are affected by the presence of disorder or dissipation. Concerning disorder, it is well known that in solid state realizations of IQHE a certain amount of impurities, with associated localized states, are crucial to the robustness of the Hall plateaus. The robustness of the topological state against disorder [121] or absence of translational invariance [122] has also been tested in simple tight-binding models: the crucial question, for what concerns our story, is if disorder tends to increase the “robustness” of the time response, as we have formulated it, by increasing the extent of the region in which non-analytic corrections to the Kubo formula dominate. We observe that the dimensionality might play a role: while for clean samples a two-dimensional (2D) lattice model with a constant drift is essentially equivalent to a one-dimensional (1D) chain with a time-periodic driving, such as the Rice-Mele model [75], the disorder could affect 1D and 2D system in different ways. The interplay between disorder and topology in 1D Floquet insulators is indeed the main focus of chapter 4.



Similar questions can be formulated concerning the role of dissipation: while the linear response regime is quite well understood [123, 124], the interplay between non-adiabatic effects and the coupling with a thermal bath still requires a precise characterization. Recent results on the effect of dissipation in the periodically driven Rice-Mele model [79] show that dissipation towards a low-temperature bath can be beneficial in increasing the occupation of the lowest-energy Floquet states, thus making the pumped charge closer to the Thouless adiabatic limit.



## 4. Pumping in disordered systems

In this chapter, I present the results published in Ref. [125], where we investigate the effects of disorder in Thouless pumping.

Given a one-dimensional insulator subjected to a periodic driving, Thouless (or adiabatic) pumping consists in the quantization of the number of particles transported each period across any section of the system. The value of the pumped charge depends on the topological properties of the Hamiltonian eigenstates. A typical tight-binding model that displays quantized pumping is the Rice-Mele (RM) model, sketched in Fig. 4.1. In the RM model the




Figure 4.1.: Sketch of the clean Rice-Mele model. The unit cell contains two sites because of the staggered potential  $\pm\Delta(t)$  and the different intra-cell and inter-cell hopping amplitudes  $J_1(t)$  and  $J_2(t)$ .

pumped charge is  $Q = 1$  if the system is driven along a close path enclosing the origin of the parameter space  $(\Delta, J_1 - J_2)$ , as summarized in Fig. 4.2.<sup>1</sup>

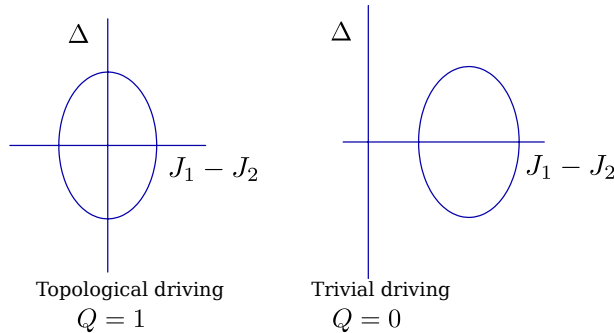


Figure 4.2.: topological driving cycle on the left, with quantized pumping  $Q = 1$ , and trivial driving cycle on the right, with  $Q = 0$ .

<sup>1</sup>A half-filled state with  $J_1 = J_2$  and  $\Delta = 0$  is clearly metallic, hence a closed path around it encloses a singularity of the Berry curvature. By transforming the problem to  $k$ -space, it is simple to argue that the two loops shown in Fig. 4.2 are topologically distinct.

When disorder is added to the sample, quantized pumping survives, because the nontrivial topology is protected by the energy gap between the ground-state manifold and the excited states. This, however, is in contrast with the expectation that any one-dimensional system undergoes Anderson localization for arbitrarily small disorder. Transport should be therefore suppressed since localized states cannot carry a steady-state current.

The solution of this puzzle is remarkable: while instantaneous energy eigenstates are Anderson localized, the periodic driving plays a fundamental role in *delocalizing Floquet states* over the whole system, henceforth allowing for a steady-state quantized current. This is linked to a localization/delocalization transition of the Floquet states at strong disorder, which occurs for periodic driving corresponding to a non-trivial loop in the parameter space.

The localization/delocalization transition is reflected in a crossover of the winding number of quasi-energies and in the spectral properties of Floquet states. Indeed a transition occurs between a continuous spectrum where the Floquet states are extended and quantized pumping holds, to a pure point spectrum in the localized phase, where transport is strongly suppressed.

## 4.1. Introduction

Thouless pumping [9, 15] provides one of the simplest manifestations of topology in quantum systems, and has attracted a lot of recent interest, both theoretically [12, 14, 70, 75, 79, 126–129] and experimentally [21–23, 25, 73, 74]. Thouless pumping is also the first example of a topological phase emerging in a periodically driven system with no static analog. Such phases have been the subject of many recent proposals [7, 12–14, 70, 130–132]. In this respect, understanding the role of disorder has a twofold purpose: on one hand, it is important to understand the robustness to disorder of the topology of driven systems [74, 133] *per se*; on the other hand, localization properties in the topological phase are relevant for the possibility of stabilizing topological pumping in *interacting* systems [134, 135] by means of many-body localization [24, 136, 137].

Since the seminal works on quantum pumping by Thouless & Niu in Ref. [15, 54] it has been claimed that the quantization of the charge transport is robust against disorder, and in general against any substrate potential, as long as the system stays in an insulating ground-state. In the original formulation, this was guaranteed by the presence of an energy gap that allows the system to evolve adiabatically for small enough driving frequency. Thouless & Niu’s idea is to define the Berry connection through the derivative of the instantaneous eigenstates with respect to an external flux piercing the system (with PBC) or equivalently with respect to a twist  $\theta$  of the boundary condition. Since observables must be independent of  $\theta$  in the thermodynamic limit, one can integrate over  $\theta$  and define a two-dimensional torus spanned by the time  $t \in [0, \tau)$ ,  $\tau$  being the period of the driving, and  $\theta \in [0, 2\pi)$ . While Thouless & Niu set up the theoretical framework to understand quantized charge transport in presence of disorder or interaction, they did not investigate the details of the transition between a topological phase and a trivial one due to increasing disorder strength.

More recently Kitagawa *et al.* [70] claimed that even in the presence of disorder, integer quantum pumping can be explained through the winding number of the many-body Floquet

quasi-energy around a Floquet-Brillouin zone  $[-\omega/2, \omega/2) \times [0, 2\pi)$ , where again one averages the system response over all possible values of the twist angle  $\theta$  at the boundary. However, they did not provide an analysis to study the actual robustness of the topological phase.

In this chapter we tackle the problem from a genuine time-periodic perspective, hence exploiting Floquet theory. Restricting ourselves to non-interacting systems, we focus on the localization properties of the Floquet states, which play a major role in explaining the persistence of Thouless pumping at finite disorder strength. Indeed, while quantized transport over a single period of the driving is expected at small disorder [15], its robustness over many driving cycles is not trivial, since it would imply the existence of *extended* Floquet states. But in the adiabatic limit, where the charge is strictly quantized, Floquet states for a generic driving should coincide with the Hamiltonian eigenstates, which are Anderson localized in one-dimension (1d). So, how can Thouless pumping in Anderson insulators be stable in the long-time limit? It has already been shown that a strong driving can dramatically change the localization properties of an Anderson insulator, see Flach [77] or Agarwal [78], leading to a diffusive transport of particles up to the localization length of Floquet states, which can be much larger of that of the energy eigenstates, even divergent for  $2\pi/\tau = \omega \rightarrow 0$ . This suggests that, at finite frequency, transport will be always suppressed over a certain distance, leading to a trivial steady-state with vanishing current across the system. Moreover, no evidence of truly extended states nor a clear link between the localization properties and topology have been found.

Focusing on the finite-size scaling of the localization length of Floquet states, the long-time dynamics and the winding of Floquet quasi-energies, we show that when the clean system is in topological phase with quantized currents, the intrinsic chirality of the model leads to a survival of extended Floquet states even at finite disorder. Remarkably, as disorder increases these states undergo a true delocalization/localization transition at a critical disorder strength  $W_c$ , which reflects itself in the breakdown of quantized transport. Crucially, topology plays a fundamental role in the existence of such extended states and on the character of the phase transition, as we illustrate by explicit comparison with the case of a trivial adiabatic driving protocol.

## 4.2. Model

We consider a disordered version of the driven Rice-Mele model [105], sketched in Fig. 4.1. For a system of spinless fermions on a chain of  $L = 2N$  sites, with  $\hat{c}_{\alpha,j}^\dagger$  creating a fermion on the  $\alpha$  site of the  $j$ -th cell, the Hamiltonian reads

$$\begin{aligned} \hat{H}(t) = & - \sum_{j=1}^N \left( J_1(t) \hat{c}_{A,j}^\dagger \hat{c}_{B,j} + J_2(t) \hat{c}_{B,j}^\dagger \hat{c}_{A,j+1} + \text{H.c.} \right) + \Delta(t) \sum_{j=1}^N \left( \hat{c}_{A,j}^\dagger \hat{c}_{A,j} - \hat{c}_{B,j}^\dagger \hat{c}_{B,j} \right) \\ & + W \sum_{\alpha=A,B} \sum_{j=1}^N \zeta_{\alpha,j} \hat{c}_{\alpha,j}^\dagger \hat{c}_{\alpha,j}. \end{aligned} \quad (4.1)$$

Here  $J_{1(2)}(t)$  and  $\Delta(t)$  describe hopping amplitudes and on-site energies for the clean model, while  $W\zeta_{\alpha,j}$  describes the on-site disorder of strength  $W$ , with  $\zeta_{\alpha,j} \in [-\frac{1}{2}, \frac{1}{2}]$  uniformly

distributed random numbers. We assume periodic boundary conditions (PBC).

In absence of disorder,  $W = 0$ , and for generic  $J_{1(2)}$  and  $\Delta$ , the instantaneous spectrum is split in two bands, separated by a gap. Thus, at *half-filling* — i.e., when number of particles is equal to  $N$  —, the charge pumped in one period is equal in the adiabatic limit to the Chern number of the occupied band [9]. This integer is different from 0 when the driving is topologically non-trivial, i.e., when the path in the space  $(J_1 - J_2, \Delta)$  encloses the gapless point  $(0, 0)$ .

To characterize the topological phase we compute the average number of particles pumped over an infinite number of periods [65, 75]

$$\bar{Q} = \lim_{M \rightarrow \infty} \frac{1}{M} \int_0^{M\tau} dt \langle \Psi(t) | \hat{J}(t) | \Psi(t) \rangle, \quad (4.2)$$

from the quantum-average of the current density operator  $\hat{J}(t)$  [9]. Here,  $\tau = 2\pi/\omega$  is the driving period, and the system is initially prepared in the  $N$ -particle ground-state  $|\Psi_0\rangle$  of  $\hat{H}(t=0)$ .

Since the Hamiltonian is time periodic, we can exploit the Floquet representation [111] of the evolution operator  $\hat{U}(t, 0) = \sum_{\nu} e^{-i\mathcal{E}_{\nu}t/\hbar} |\Phi_{\nu}(t)\rangle \langle \Phi_{\nu}(0)|$ , where  $|\Phi_{\nu}(t)\rangle = |\Phi_{\nu}(t + \tau)\rangle$  are  $N$ -particle Floquet modes and  $\mathcal{E}_{\nu}$  are the many-body quasi-energies.  $\bar{Q}$  can be computed directly in the Floquet diagonal ensemble [75, 112]

$$\bar{Q} = Q_d = \sum_{\nu} \mathcal{N}_{\nu} \int_0^{\tau} dt \langle \Phi_{\nu}(t) | \hat{J}(t) | \Phi_{\nu}(t) \rangle, \quad (4.3)$$

where  $\mathcal{N}_{\nu} = |\langle \Psi_0 | \Phi_{\nu}(0) \rangle|^2$  is the occupation number of the  $\nu$ -th Floquet state. For non-interacting fermions, it suffices to know the single-particle (SP) Floquet states  $|\phi_{\alpha}(t)\rangle$  and their occupation number  $n_{\alpha}$  to explicitly calculate the diagonal pumped charge [65, 75].

Details about the derivation of Eqs. (4.2) and (4.3) can be found in chapter 2.

### 4.3. Results

Let us start from the phenomenological aspect of the problem, and compute directly the pumped charge with Eq. (4.3) as a function of the disorder strength  $W$ . Figure 4.3 (a) shows the average over many instances  $[Q_d]_{av}$ , pumped in a topological driving cycle with  $J_{1(2)}(t) = J_0 \pm \delta_0 \cos(\omega t)$ , and  $\Delta(t) = \Delta_0 \sin(\omega t)$ . The system is initially prepared in the ground-state of  $\hat{H}(t=0)$  and then evolved directly with a finite frequency  $\hbar\omega = 0.01J_0$ . We can clearly distinguish between three regimes. For sufficiently small  $W \lesssim 3J_0$  topological pumping persists and the charge is still quantized, up to the non-adiabatic corrections studied in Ref. [75] for the clean version of this model (see also the discussion in chapter 3). The regime of large  $W \gtrsim 8J_0$  is also rather clear: disorder is so strong that the particles are localized on a single site during the whole driving cycle, hence  $Q_d = 0$ . The intermediate region  $W/J_0 \approx 4$  shows large sample-to-sample fluctuations and corresponds to a transition in the transport properties. Our aim in this chapter is to characterize precisely this transition in terms of localization properties of Floquet states.

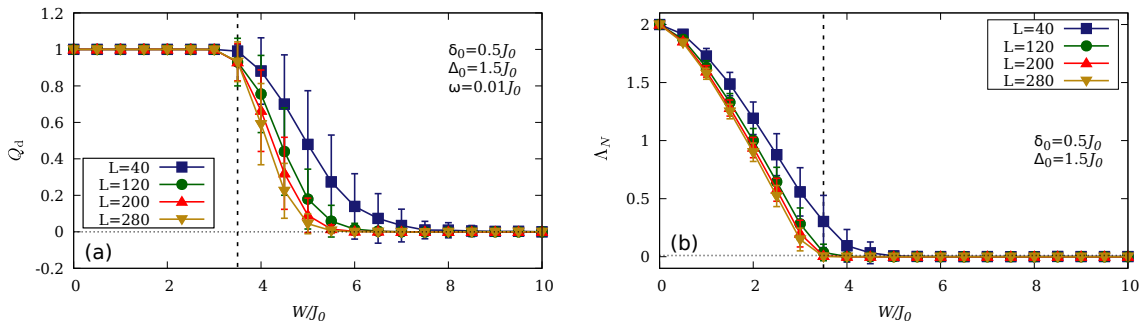


Figure 4.3.: (a) Disorder average of the diagonal pumped charge plotted against disorder strength. The transition between the quantized charge regime and the trivial one  $Q_d = 0$  is linked to the closing of the minimum energy gap due to the disorder, as highlighted by the vertical dashed line. (b) Many-body energy gap  $\Lambda_N$  in Eq. (4.4) versus disorder strength  $W$ . All data in both panels (a) and (b) have been averaged over 200 disordered configurations.

Figure 4.3(b) shows that the drop of  $[Q_d]_{av}$  starts when disorder closes the instantaneous gap

$$\Lambda_N \equiv \min_{t \in [0, \tau]} [E_{N+1}(t) - E_N(t)], \quad (4.4)$$

where  $E_N(t)$  is the  $N$ -particle ground-state energy at time  $t$ . Apparently, quantized pump has a very simple requirement: disorder must not close the gap at any point during the driving, to allow for an adiabatic evolution of the ground-state. Non-adiabatic corrections do not increase as the gap shrinks due to rising disorder, until  $\Lambda_N < \hbar\omega$  and the driving excites states in the conduction band, gradually destroying quantized pumping. Notice however that in the transient regime  $W/J_0 \approx 4$  transport is still present and the current can be rather large, depending on the particular disorder instance under investigation.

Tackling the problem from this perspective, one might ask whether the choice of the Hamiltonian parameters, and hence the driving cycle in the space  $(J_1 - J_2, \Delta)$ , plays any role in determining the disorder strength  $W_c$  at which charge quantization breaks down. Of course it does, and it depends on the *gap structure* induced by disorder in the parameter space. To illustrate this, we show in Fig. 4.4 the logarithm of  $\Lambda_N$  as a function of  $J_1 - J_2$  and  $\Delta$ , and compare the position of the gapless points with the driving cycle we are considering (black ellipse). When disorder is absent,  $W = 0$ , there is a single gapless point in  $(0, 0)$  and the driving cycle circles around it, leading to topological transport [9, 75]. As disorder increases, the number of gapless points increases and they are connected by a net of lines where  $\Lambda_N$  is very small. As the net size increases, gapless points are crossed during the driving, thus breaking the quantization of  $Q_d$ . However at this point — see for instance  $W = 4J_0$  — one could enlarge the cycle, i.e. making the variation of the Hamiltonian parameters larger, and again enclose inside the path all the gapless points, restoring quantized transport. When disorder increase again, however, the gapless lines cross necessarily any realistic path, leading to a “frozen” localized system with no transport at all.

We turn now our attention to the localization properties of Floquet states, in order to




Figure 4.4.: Gap-closing structure in the  $(J_1 - J_2, \Delta)$  parameter space for increasing disorder in a single random instance. The solid black line is the driving cycle we consider throughout the chapter. The gap we plot is  $\Lambda_N$  in Eq. (4.4).

understand why quantized transport survives in spite of the instantaneous eigenstates being Anderson localized [138]. We will show that, crucially, a significant fraction of the SP Floquet states remain *delocalized* even for very low frequency, due to a driving-induced mixing of localized states [77, 78]. We analyzed localization/delocalization of states through the real-space inverse participation ratio (IPR) [139] of the single-particle Floquet modes  $|\phi_\alpha(0)\rangle$

$$\text{IPR}_\alpha = \sum_l |\langle l | \phi_\alpha(0) \rangle|^4, \quad (4.5)$$

with  $|l\rangle = \hat{c}_l^\dagger |0\rangle$  being a particle localized at site  $l$ . For a finite system,  $\text{IPR}_\alpha \in [L^{-1}, 1]$ , where  $\text{IPR}_\alpha \sim L^{-1}$  signals a completely delocalized (plane-wave-like) state, while  $\text{IPR}_\alpha = 1$  corresponds to a perfect localization on a single site.

Figure 4.5(a) shows the distribution of IPRs of Floquet states for three values of the disorder strength  $W$ . Notice the presence of a very sharp peak in the IPR distribution, signalling the presence of a bunch of states with *large localization length*. By analyzing the peak position as a function of the system size  $L$ , we find that it scales as  $\text{IPR}_\alpha \sim L^{-1}$  for  $W/J_0 = 2$  and 4, see Fig. 4.5(b), suggesting that the *mode*<sup>2</sup> of the IPR distribution corresponds to extended states. For stronger disorder instead, the mode reaches a finite value, as expected for a localized system. Incidentally, the *average* IPR suggests always a localized scenario, as  $\overline{\text{IPR}} = \frac{1}{L} \sum_\alpha \text{IPR}_\alpha$  saturates to a finite value for  $L \rightarrow \infty$ . However, the shape of the distribution also suggests that the average alone is not enough to study the

<sup>2</sup>The “mode” of a distribution is the value that appears more often, i.e., the maximum.



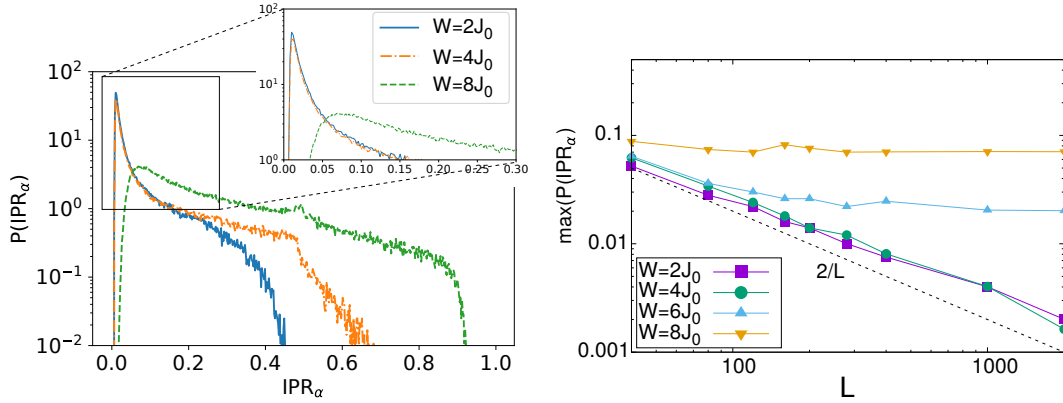


Figure 4.5.: (a): Disorder averaged IPR distribution of SP Floquet states for several values of  $W$  and  $L = 280$ . The inset shows the sharp peaks almost superimposed at small IPR for  $W = 2J_0$  and  $W = 4J_0$ . (b): Scaling of the mode of the IPR distribution with the system size  $L$ ; the dashed line highlights the delocalized behavior  $\sim L^{-1}$ .

localization properties of the model, hence we believe that the mode of the distribution is a better indicator.

Intriguingly, a similar distributions (not shown) emerges — for the same disorder strength — also when the driving protocol is topologically *trivial*. To better analyze the size-dependence of the IPR peak, and its correlation with the topology of the driving, we estimate a characteristic localization length for a chain of size  $L$  from the inverse of the peak's position in the IPR distribution (inverse of the mode)

$$\bar{\xi}_L(W) = \frac{1}{\operatorname{argmax}(P(\operatorname{IPR}_\alpha))}. \quad (4.6)$$

Figures 4.6(a) and (b) show the size-scaling of  $\bar{\xi}_L(W)$  for a trivial and topological driving, respectively. When the driving is trivial, our data suggest that  $\bar{\xi}_L(W)$  scales as  $W^{-\beta}$  with  $\beta \simeq 2.5$  for large  $W$ , see Fig. 4.6(a), while it saturates to the system size  $\sim L$  when  $W$  is small. Hence we can extract a crossover disorder strength  $W^* \sim L^{-1/\beta}$  separating these two regimes, vanishing in the thermodynamic limit: here, truly extended Floquet states appear only at zero disorder. By rescaling the data, we see a very good collapse of  $\bar{\xi}_L(W)/L$  versus  $L^{1/\beta}W$ , see inset of Fig. 4.6(a). On the other hand, when the driving is topological the same phenomenology holds with a *finite critical disorder strength*  $W_c$ , see Fig. 4.6 (b): For  $W > W_c \simeq 3.5J_0$  we observe that  $\bar{\xi}_L(W) \sim (W - W_c)^{-\beta}$ , with  $\beta \simeq 2$ , while, again, the localization length saturates to  $L$  when  $W < W_c$ , thus indicating the presence of a true localization/delocalization phase transition. The critical value  $W_c$  extracted by our scaling analysis is compatible with the breaking of quantization in Fig. 4.3.

To understand the mechanism behind the delocalization/localization transition, we study the relation between the time-averaged energy of the SP Floquet states

$$\langle E \rangle_\alpha = \frac{1}{\tau} \int_0^\tau dt \langle \phi_\alpha(t) | \hat{H}(t) | \phi_\alpha(t) \rangle, \quad (4.7)$$

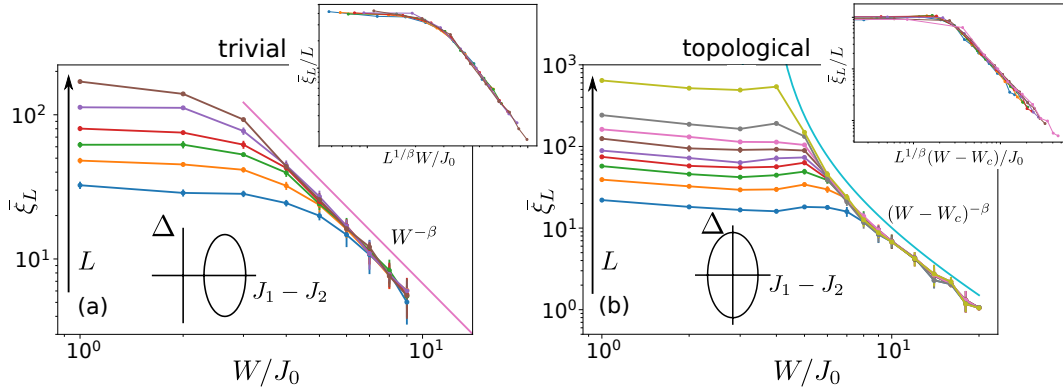


Figure 4.6.: (a): Characteristic localization length  $\bar{\xi}_L(W)$  as a function of disorder, for a trivial driving cycle ( $J_1 = 4J_0 + \delta_0 \sin(\omega t)$ ). The inset shows the collapse obtained by  $\bar{\xi}_L \rightarrow \bar{\xi}_L/L$  and  $W \rightarrow WL^{1/\beta}$ , with  $\beta \simeq 2.5$ . (b): Characteristic localization length when the system exhibits topological transport ( $J_1 = J_0 + \delta_0 \sin(\omega t)$ ). The inset shows the collapse  $\bar{\xi}_L \rightarrow \bar{\xi}_L/L$  and  $W \rightarrow (W - W_c)L^{1/\beta}$ , with  $W_c \simeq 3.5J_0$  and  $\beta \simeq 2$ . The parameters used in the simulation are  $\delta_0 = 0.5J_0$ ,  $\Delta_0 = 1.5J_0$ ,  $\hbar\omega = 0.01J_0$ .

and the corresponding  $\text{IPR}_\alpha$  (Fig. 4.7). For weak disorder, extended states carrying charge in the positive (negative) direction lay in the middle of the lower (higher) band, while localized ones stay closer to the edges. Floquet states from different bands are separated in energy by a gap — closely related to the instantaneous energy gap — and by a mobility edge. The presence of localized states at the band edges suggests that the driving-induced mixing occurs mainly at the band center, as long as  $W \lesssim W_c$  (see panel (b) of Fig. 4.7, with  $W = 3.5J_0$ ). This implies an additional robustness against non-adiabatic effects — indeed we observe that our results do not depend on the precise value of  $\omega$ , as further examined in section 4.5—, even when the gap is almost completely closed. When  $W \gg W_c$  the two bands merge into a single one where extended states transporting opposite charges hybridize into localized states, see panel (c) with  $W = 8J_0$ , and the current stops flowing.

This phenomenology is similar to what happens in integer quantum Hall effect (IQHE) in 2D systems, where there *must be* spectral regions of extended states [53, 140, 141], in order to have a non-zero quantized transverse conductivity. Also the exponent  $\beta \simeq 2$  found for  $W > W_c$ , is in good agreement with a similar scaling analysis performed on the density of extended states in IQHE in a disordered sample [142]. Even though the parallelism between the physics of clean 1D topological charge pumping and 2D integer quantum Hall effect is well established [15, 66, 143], this connection is not obvious for disordered samples. Indeed a localization/delocalization transition clearly associated with topological properties, as it happens in IQHE [140, 141], has not been investigated previously in a one dimensional driven Anderson insulator.




Figure 4.7.:  $IPR_\alpha$  vs the time-averaged energy of the corresponding Floquet state  $\langle E \rangle_\alpha$ . The dashed lines indicates the value  $1/L$  associated with extended states. The data refer to several realizations of a chain with 200 sites and increasing disorder strengths:  $W = 2J_0$  in panel (a),  $W = 3.5J_0$  in panel (b), and  $W = 8J_0$  in panel (c).  $Q_d = \pm 1$  is the charge transported when a single band is completely filled.

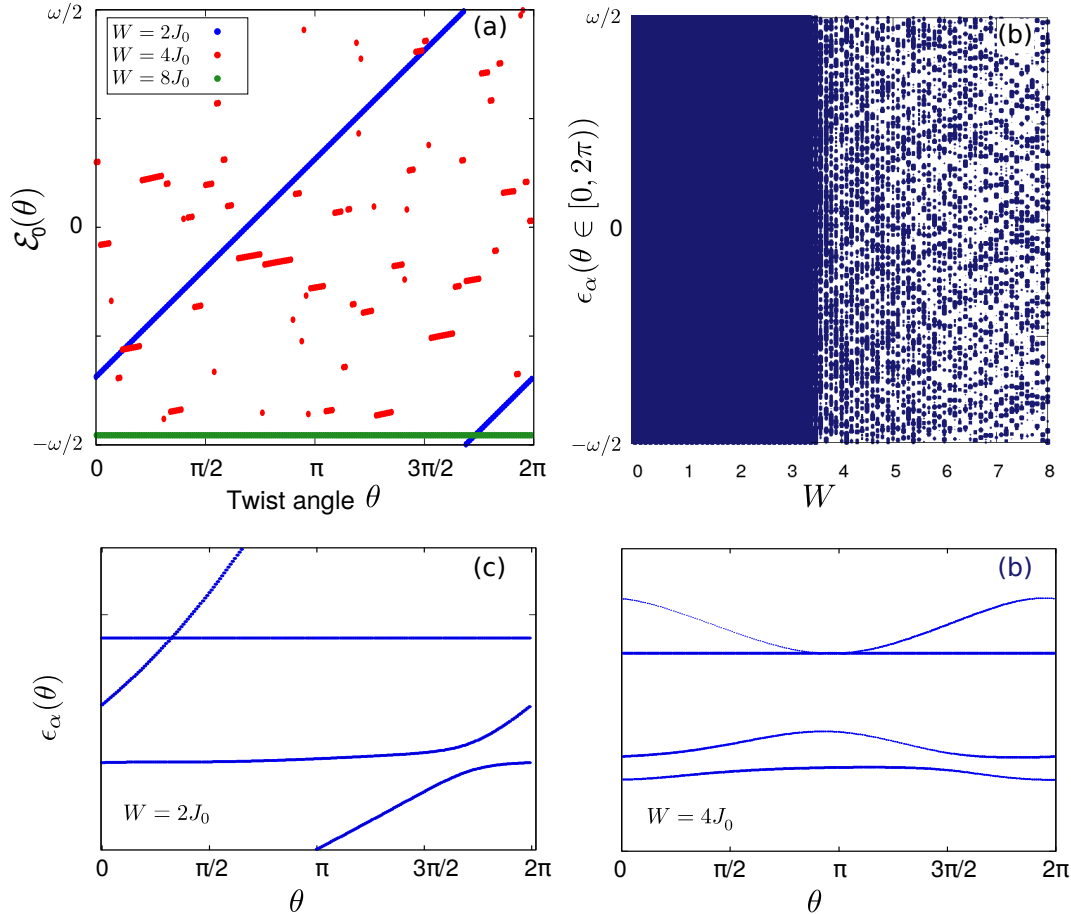


Figure 4.8.: (a) Quasi-energy  $\mathcal{E}_0(\theta)$  of the many-body Floquet state with lowest energy. The winding is well defined only for  $W = 2J_0$  and  $W = 8J_0$ , when  $|\Psi_0\rangle$  has a non vanishing projection on a single MB Floquet state. (b): SP quasi-energy spectrum for all possible angles  $\theta$  as a function of the disorder. (c) and (d): details of the dependence of SP quasi-energy  $\epsilon_\alpha$  on the phase twist  $\theta$  for  $W = 2J_0$  and  $W = 4J_0$ . Only a portion of the quasi-energy spectrum is reported, with the linewidth proportional to the occupation of the corresponding Floquet state.

## 4.4. Winding of quasi-energies

In a clean system, quantized pumping corresponds to a non-trivial winding of the quasi-energy of the occupied Floquet bands in  $k$ -space, as the momentum spans the first Brillouin zone. A detailed derivation is provided in section 2.2. When translational invariance is broken, a common procedure is to introduce a phase twist  $\theta \in [0, 2\pi)$  between site 1 and site  $L$  and then take the average of  $Q_d$  over  $\theta$  [15, 70]. This operation is justified because when the state projector is exponentially localized, the dependence of observables on the twisted boundary decays exponentially with  $L$  [104]. Recalling the relation between pumped charge and winding number in section 2.2, Eq. (4.3) can then be equivalently written as

$$Q_d = \int_0^{2\pi} \frac{d\theta}{2\pi} Q_d(\theta) = \frac{\tau}{\hbar} \sum_{\nu} \int_0^{2\pi} \frac{d\theta}{2\pi} \mathcal{N}_{\nu}(\theta) \partial_{\theta} \mathcal{E}_{\nu}(\theta). \quad (4.8)$$

Here  $\mathcal{E}_{\nu} = \epsilon_{\alpha_1} + \dots + \epsilon_{\alpha_N}$  is the  $N$ -particle quasi-energy associated with the Floquet state  $|\Phi_{\nu}\rangle$  given by a Slater determinant of the SP states  $|\phi_{\alpha_1}\rangle, \dots, |\phi_{\alpha_N}\rangle$ ;  $\mathcal{N}_{\nu} = |\langle \Psi_0 | \Phi_{\nu} \rangle|^2$  is the occupation number. In this context, the winding number is the number of times that  $\mathcal{E}_{\nu}(\theta)$  wraps around the first Floquet-Brillouin zone as  $\theta$  goes from 0 to  $2\pi$ . Besides non-adiabatic corrections that depend on the initial state  $|\Psi_0\rangle$  [75, 106] discussed in chapter 3,  $Q_d$  is quantized when a single many-body Floquet state is occupied, e.g.  $\mathcal{N}_{\nu} \simeq \delta_{0,\nu}$  independently of  $\theta$ , and that state has a nontrivial winding number. Henceforth we focus on the Floquet state with the lowest initial energy  $|\Phi_0(\theta)\rangle$ , computed as the Slater determinant of the  $N$  SP Floquet states with the highest projection on the ground-state in which the state is initially prepared.

We report in Fig. 4.8(a)  $\mathcal{E}_0(\theta)$  in the first Floquet-Brillouin zone for three different disorder strengths  $W/J_0 = 2, 4$  and  $8$ , which are respectively below, close, and above the transition point. In Fig. 4.8(b) the SP quasi-energy spectrum is plotted with respect to  $W$  as  $\theta$  spans the interval  $\in [0, 2\pi)$ . A localized state is characterized by a quasi-energy  $\epsilon_{\alpha}$  *periodic* in  $\theta$ , while extended ones with positive winding satisfy the relation  $\epsilon_{\alpha}(2\pi) = \epsilon_{\alpha+1}(0)$ . Hence, we distinguish between three situations.

- 1)  $W < W_c$ :  $|\Phi_0\rangle$  coincides essentially with  $|\Psi_0\rangle$  ( $\mathcal{N}_{\nu} \simeq \delta_{0,\nu}$ ), because adiabaticity is preserved at many-body level, and has winding number equal to 1, blue line in Fig. 4.8(a). The SP quasi-energy spectrum is continuous in  $\theta$  and there are no gaps in the Floquet-Brillouin zone, Fig. 4.8(b). Most of the SP states feel the twist at the boundary, obey  $\epsilon_{\alpha}(2\pi) = \epsilon_{\alpha+1}(0)$  and contribute to the winding of  $\mathcal{E}(\theta)$ , as shown in Fig. 4.8(c).
- 2)  $W \gtrsim W_c$ : the initial ground-state  $|\Psi_0\rangle$  has relevant projections over many MB Floquet states and the SP occupation numbers  $n_{\alpha}$  depend non-trivially on  $\theta$ . Hence the set  $\alpha_1, \dots, \alpha_N$  of SP Floquet states in the Slater determinant  $|\Phi_0\rangle$ , changes with  $\theta$ , thus making  $\mathcal{E}_0(\theta)$  discontinuous, see red dots in Fig. 4.8(a). Gaps start to appear in the SP quasi-energy spectrum, see Fig. 4.8(b), and the occupation number itself depends non-trivially on  $\theta$ . The SP Floquet states with opposite transported charge start to be mixed in pairs of localized states, with quasi-energies periodic in  $\theta$ , see Fig. 4.8(d).
- 3)  $W > W_c$ : both SP Floquet states and Hamiltonian eigenstates are strongly localized and there is no current. Again  $\mathcal{N}_{\nu} = \delta_{0,\nu}$ , but the winding number is trivial, see

the green horizontal line in Fig. 4.8(a), because SP quasi-energy spectrum has only a pure-point contribution and localization makes the system insensitive to the boundary twist.

## 4.5. Frequency dependence of the localization length

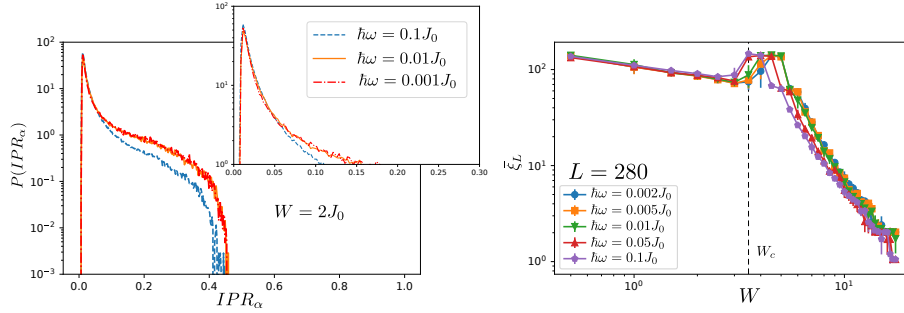


Figure 4.9.: Left: SP Floquet states IPR distribution for frequencies  $\hbar\omega = 0.1J_0$ ,  $0.01J_0$  and  $0.001J_0$ ,  $W = 2J_0$  and  $L = 280$ . Notice the almost perfect overlap around the peak of extended states. Right: localization length versus disorder strength for a chain of 280 sites for different values of the driving frequency.

When considering the frequency dependence of the localization/delocalization transition discussed in section 4.3, two main issues can be raised: *i*) how does the “fraction of extended states” and *ii*) how the “mode” of the IPR distribution change with  $\omega$ ?

The fraction of extended states is related to the area under the peak in the IPR distribution. In the left panel of Fig. 4.9 we compare  $P(IPR)$  for three different frequencies spanning two orders of magnitude. This shows that the shape of the peak remains almost identical, suggesting that the fraction of delocalized states does not depend much on  $\omega$ .

The localization length  $\bar{\xi}_L$ , instead, shows a mild dependence on the frequency, namely it increases as  $\omega$  decreases when  $W > W_c$ . This is in agreement with general results on driven Anderson insulators [77, 78]. However, the curves for  $\hbar\omega \leq 0.01J_0$  are almost superimposed, suggesting that the critical disorder strength converges rapidly with decreasing  $\omega$  to the value at which the minimum energy gap closes  $W_c \simeq 3.5J_0$ . When the charge is quantized ( $W < W_c$ ) the localization length, instead, does not show any dependence on the frequency.

## 4.6. Conclusions

In this chapter we analyzed the steady-state current flowing in a one-dimensional Floquet-Anderson insulator: the topological periodic driving mixes localized Hamiltonian eigenstates, giving rise to extended Floquet modes. The dynamics is adiabatic only at many-body level, but not at the SP one, where driving-induced mixing of localized states occurs even at very low frequencies. Delocalization makes quantized pumping robust until extended Floquet states with opposite winding coalesce for large disorder. Even though the physics of quantum

pumping in clean systems is the same as 2D IQHE, this analogy is not trivial in the presence of disorder, since the 1D periodically driven chain would be mapped in an *extremely anisotropic* disordered 2D model.

A subtle point emerges in the adiabatic limit  $\omega \rightarrow 0$ . In a truly adiabatic evolution the SP Floquet states would coincide with the Hamiltonian eigenstates, thus being localized, at least when the disorder-induced SP level crossings are actually *avoided crossings*. This happens generically when one takes the adiabatic limit in a finite system, where the energy levels are protected by finite gaps, although exponentially small in  $L$ . Quantized pumping still works because disorder induces resonances in the spectrum, allowing for large distance tunnelling [144, 145]. Instead, if disorder induces true level crossings in the occupied band of the SP spectrum, as in the *control freak limit* [146], a fine tuned driving is still able to mix localized energy eigenstates into extended Floquet states, even at vanishing frequency. When the thermodynamic limit  $L \rightarrow \infty$  is taken first, the spectrum becomes *dense* and the driving mixes SP levels at arbitrarily small frequency, leading again to a quantized pumping. These arguments suggest that quantized pumping is obtained regardless of which of the two limits ( $\omega \rightarrow 0$  or  $L \rightarrow \infty$ ) is taken first, although the physical mechanisms appear somewhat different.

**Part II.**

# **Quantum Optimization**





## 5. QAOA on fully connected models

In this chapter I discuss our work on the quantum approximate optimization algorithm, available on the arXiv [147].

As discussed in the introduction in section 1.2, one of the central problems in quantum computing is “ground state preparation”, which is a relevant aspect for many different fields, from condensed matter problems and material design to classical optimization problems and information theory.

A recent proposal to find the ground state of a spin Hamiltonian  $\hat{H}_z$  is the quantum approximate optimization algorithm (QAOA), sketched in Fig. 5.1. It consists in the optimization of a variational wave-function  $|\psi_P(\gamma, \beta)\rangle$  obtained by applying a set of P unitary operators

$$\hat{U}(\gamma_m, \beta_m) = e^{-i\beta_m \hat{H}_x} e^{-i\gamma_m \hat{H}_z}, \quad (5.1)$$

to evolve the initial state. Here  $\hat{H}_x$  is a driving Hamiltonian that induces fluctuations in the basis of  $\hat{H}_z$ .

In this chapter, we test QAOA on the fully connected p-spin Ising ferromagnet

$$\hat{H}_z = - \left( \sum_j \hat{\sigma}_j^z \right)^p, \quad (5.2)$$

a problem that notoriously poses severe difficulties to a Quantum Annealing (QA) approach, due to the exponentially small gaps encountered at first-order phase transition for  $p \geq 3$ .

We demonstrate that QAOA can construct with polynomially scaling resources the ground state of the p-spin ferromagnet with very shallow circuits  $P = 1$  or  $P = 2$  depending on

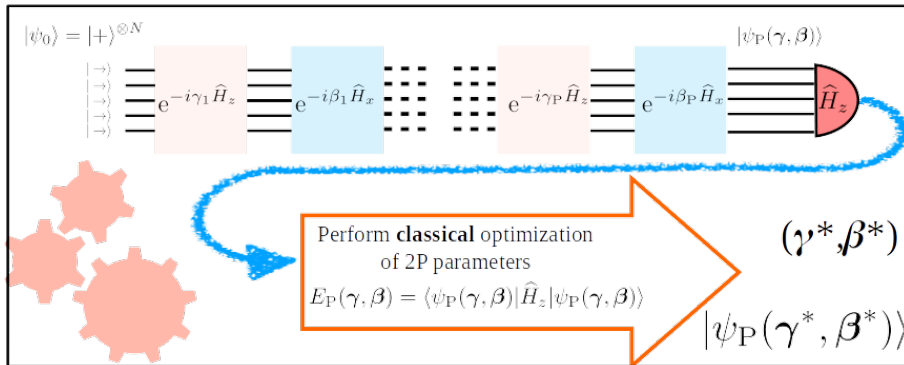


Figure 5.1.: Schematic representation of QAOA protocol.

the parity of the number of variables  $N$ . This optimal minimum, however, is hidden in an extremely rugged energy landscape, which makes minimization attempts prone to remain trapped in bad local minima. Indeed for a generic target state, we show that an appropriate QAOA parameter initialization is necessary to achieve a good performance of the algorithm when the number of variational parameters  $2P$  is much smaller than the system size  $N$ , because of the large number of sub-optimal local minima. When  $P$  exceeds a critical value  $P_N^* \propto N$  instead, the structure of the parameter space simplifies and all minima become degenerate. This allows us to achieve the ground state with perfect fidelity with a number of parameters scaling extensively with  $N$ , and with resources scaling polynomially with  $N$ .

## 5.1. Introduction

Optimization and ground state preparation are two of the building blocks in the current research on quantum technologies [45, 148]. Optimization is a long-standing problem in physics and computer science and lies at the roots of the efforts to show a possible “quantum supremacy” [149–151] over classical algorithms. A robust state preparation strategy, in turn, would be a crucial tool for quantum technologies, and would also allow to “solve”, using quantum hardware, many long-standing problems in condensed matter theory or quantum chemistry [26, 83, 84]. The two are intimately connected, as many optimization tasks can be reformulated in terms of finding the classical ground state of an appropriate spin-glass Hamiltonian [80].

A traditional tool in this field has been Quantum Annealing [28, 29, 85, 86] (QA) — *alias* Adiabatic Quantum Computation [30, 88] —, which relies on the adiabatic theorem to find the ground state of a target Hamiltonian, starting from a trivial initial state. Although QA appeared to be more efficient than its classical counterpart for certain problems [86, 152–155], it is limited by the smallest gap encountered during the evolution, which vanishes, in the thermodynamic limit, when the system crosses a phase transition. In this context, the fully-connected  $p$ -spin Ising ferromagnet is a simple but useful benchmark for optimization, because QA fails due to the exponentially small gap at the first-order phase transition encountered for  $p \geq 3$  [100–102]. Several techniques have been advocated to overcome the slowness induced by such an exponentially small gap, including the introduction of non-stoquastic terms [39, 156], pausing [157], dissipative effects [38, 158], or approximated counterdiabatic driving [159]. Their successful application, however, often depends on the knowledge of the spectrum or on the phase diagram of the model, thus making these techniques highly problem-specific.

Recent alternative ground state preparation approaches [44, 160, 161] rely on hybrid quantum-classical variational techniques [42] to tackle such problems, avoiding the limitations imposed by a QA adiabatic evolution. Here we focus on one such scheme, the Quantum Approximate Optimization Algorithm (QAOA) [44, 95].

The core idea of QAOA is to write a trial wave-function as a product of many unitary operators, each depending on a classical variational parameter, applied to a state simple to construct, usually a product state with spins aligned in the  $x$ -direction. A quantum hardware performs the discrete quantum dynamics and measures the expectation value of the target

Hamiltonian, which is then minimized by an external classical algorithm, as a real function in a high dimensional parameter space.

Although QAOA is a universal computational scheme [94], its performance strongly depends on the details of the target Hamiltonian. QAOA seems to perform rather well on Max-Cut problems [95] and on short-range spin systems [96, 162]. The Grover search problem has also been studied within QAOA, showing that it leads to the optimal square root speed-up with respect to classical algorithms [163]. For generic long-ranged Hamiltonians, however, many open questions remain. The questions concern, in particular, the efficiency of the algorithm when a large number of unitaries are employed, or the ability to overcome large energy barriers in presence of first-order phase transitions, or the presence of “smooth” sets of optimal parameters [90, 95, 98].

In this chapter, I will show that QAOA can construct with polynomially scaling resources the ground state of the fully-connected  $p$ -spin Ising ferromagnet for all  $p \geq 2$ , hence including the case where a first-order phase transition occurs. For a generic target state, I find that an appropriate QAOA parameter initialization is necessary to achieve a good performance of the algorithm when the number of variational parameters  $2P$  is much smaller than the system size  $N$ , because of the large number of sub-optimal local minima. Finally, I will show that when  $P > P_N^* \propto N$ , the structure of the parameter space simplifies, and all minima become degenerate. This allows us to achieve the ground state with perfect fidelity with a number of parameters scaling extensively with  $N$ , and with resources scaling polynomially with  $N$ .

## 5.2. Model and QAOA algorithm

As a benchmark for QAOA on long-range models we focus on the ferromagnetic fully-connected  $p$ -spin model [100–102, 164]:

$$\hat{H}_{\text{target}} = -\frac{1}{N^{p-1}} \left( \sum_{j=1}^N \hat{\sigma}_j^z \right)^p - h \left( \sum_{j=1}^N \hat{\sigma}_j^x \right), \quad (5.3)$$

where  $\hat{\sigma}_j^{x,z}$  are Pauli matrices at site  $j$ ,  $N$  is the total number of sites, and  $h$  a transverse field. This model displays, for  $p = 2$ , a second-order quantum phase transition, at a critical transverse field  $h_c = 2$ , from a paramagnetic ( $h > h_c$ ) to a symmetry-broken ferromagnetic phase ( $h < h_c$ ). The transition becomes first-order for  $p > 2$ , and  $h_c$  decreases for increasing  $p$ , with  $h_c \rightarrow 1$  for  $p \rightarrow \infty$  [100].

The QAOA algorithm [44] is a variational method to find the ground state of a target Hamiltonian  $\hat{H}_{\text{target}}$ . Starting from an initial spin state polarized along the  $\hat{x}$  direction

$$|+\rangle = \left( \frac{|\uparrow\rangle + |\downarrow\rangle}{\sqrt{2}} \right)^{\otimes N}, \quad (5.4)$$

QAOA writes the following variational *Ansatz*

$$|\psi_P(\boldsymbol{\gamma}, \boldsymbol{\beta})\rangle = e^{-i\beta_P \hat{H}_x} e^{-i\gamma_P \hat{H}_z} \dots e^{-i\beta_1 \hat{H}_x} e^{-i\gamma_1 \hat{H}_z} |+\rangle \quad (5.5)$$

in terms of  $2P$  variational parameters  $\boldsymbol{\gamma} = (\gamma_1 \cdots \gamma_P)$  and  $\boldsymbol{\beta} = (\beta_1 \cdots \beta_P)$ , where  $\widehat{H}_z$  and  $\widehat{H}_x$  are non-commuting Hamiltonians depending on the problem we wish to solve. Here we take  $\widehat{H}_x = -\sum_j \hat{\sigma}_j^x$ , the standard transverse field term, and an interaction term  $\widehat{H}_z$

$$\widehat{H}_z = \left( \sum_{j=1}^N \hat{\sigma}_j^z \right)^P, \quad (5.6)$$

chosen for convenience to have a super-extensive form with an integer spectrum. These choices allow us to restrict the parameter space for  $\gamma_m$  and  $\beta_m$  to the interval  $[0, \pi]$ . See appendix B.1.4 for a detailed discussion on the symmetries in the variational parameter space. In each QAOA run the variational energy cost function

$$E_P(\boldsymbol{\gamma}, \boldsymbol{\beta}) = \langle \psi_P(\boldsymbol{\gamma}, \boldsymbol{\beta}) | \widehat{H}_{\text{target}} | \psi_P(\boldsymbol{\gamma}, \boldsymbol{\beta}) \rangle, \quad (5.7)$$

is minimized, until convergence to a local minimum  $(\boldsymbol{\gamma}^*, \boldsymbol{\beta}^*)$  is obtained. The quality of the variational solution is gauged by computing the residual energy density [96]

$$\epsilon_P^{\text{res}}(\boldsymbol{\gamma}^*, \boldsymbol{\beta}^*) = \frac{E_P(\boldsymbol{\gamma}^*, \boldsymbol{\beta}^*) - E_{\min}}{E_{\max} - E_{\min}}, \quad (5.8)$$

where  $E_{\min}$  and  $E_{\max}$  are the lowest and largest eigenvalues, respectively, of the target Hamiltonian.

The connection with a QA approach is interesting [96]. In QA one would write an interpolating Hamiltonian [30]  $\widehat{H}(s) = s\widehat{H}_{\text{target}} + (1-s)\widehat{H}_x$ , with  $s(t)$  driven from  $s(0) = 0$  to  $s(\tau) = 1$  in a sufficiently large annealing time  $\tau$ . A lowest-order Trotter decomposition of the corresponding step-discretized evolution operator — with  $s_{m=1\dots P}$  constant for a time-interval  $\Delta t_{m=1\dots P}$  — would then result in a state of the form of Eq. (5.5) with:

$$\begin{cases} \gamma_m = \frac{s_m \Delta t_m}{\hbar} \frac{1}{N^{P-1}} \\ \beta_m = \frac{\Delta t_m}{\hbar} (1 - s_m(1-h)) \end{cases} \quad (5.9)$$

where the total evolution time would be given by:

$$\frac{\tau}{\hbar} = \sum_{m=1}^P \frac{\Delta t_m}{\hbar} = \sum_{m=1}^P (\beta_m + (1-h)\gamma_m N^{P-1}). \quad (5.10)$$

While an optimization of the parameters  $s_m$  and  $\Delta t_m$  is in principle possible, the standard linear schedule  $s(t) = t/\tau$  would result in a digitized-QA scheme where  $s_m = m/P$  and  $\Delta t_m = \Delta t = \tau/P$  [91, 165]. With these choices, a convenient starting point for the QAOA optimization algorithm would be to take

$$\begin{cases} \gamma_m^0 = \frac{\Delta t}{\hbar} \frac{m}{P} \frac{1}{N^{P-1}} \\ \beta_m^0 = \frac{\Delta t}{\hbar} \left(1 - \frac{m}{P}(1-h)\right) \end{cases}, \quad (5.11)$$

with possible addition of a small noise term. Alternatively, we might choose a completely random initial point with  $\gamma_m^0, \beta_m^0 \in [0, \pi]$ . These two alternative choices will be henceforth referred to as l-init and r-init.

### 5.3. Exact ground state with $P = 1$

Ref. [166] has shown that the target ground state of the  $p = 2$  fully connected Ising ferromagnet with  $h = 0$ , the so-called Lipkin-Meshov-Glick [167] model, can be perfectly constructed, with unit fidelity, with the smallest QAOA circuit,  $P = 1$ . Ref. [168] has recently shown that a whole class of spin-glass models can be constructed where QAOA shows such a property.

Here I show that the general  $p$ -spin model in Eq. (5.3) belongs to the class of  $P = 1$  QAOA-solvable problems, provided that  $h = 0$  and that the number of spin variable  $N$  is odd, when the system is initially prepared in the fully  $x$ -polarized state in Eq. (5.4). This result is already noteworthy, as it shows that one can construct the exact  $h = 0$  classical ground state with an algorithm whose equivalent computational time, see Eq. (5.10), scales as  $N^{p-1}$ . On the contrary, for any finite  $N$ , a QA algorithm would need to cope with a minimum spectral gap at the transition point [100–102, 164] which scales as  $\Delta \sim N^{-1/3}$  if  $p = 2$  and  $\Delta \sim e^{-\alpha_p N}$  if  $p \geq 3$ : with a linear schedule annealing, this would imply a total annealing time  $\tau \propto \Delta^{-2}$ , hence  $\tau \sim N^{2/3}$  for  $p = 2$  and  $\tau \sim e^{2\alpha_p N}$  for  $p > 2$ . Therefore, QAOA shows an *exponential speed-up* with respect to a linear-schedule QA for  $p > 2$ , without exploiting any knowledge on the spectrum or on the phase diagram.

In the following, I sketch the analysis of the optimal solution for  $P = 1$  and  $N$  odd, while a detailed proof is provided in appendix. B.1. The core idea of the proof starts from observing that for  $P = 1$  the fidelity reads:

$$\begin{aligned} \mathcal{F}(\gamma, \beta) &= \left| \langle \psi_{\text{targ}} | e^{-i\beta \hat{H}_x} e^{-i\gamma \hat{H}_z} | + \rangle \right|^2 \\ &= \left| \frac{1}{\sqrt{2^N}} \sum_l e^{-i\gamma E_l} \langle \psi_{\text{targ}} | e^{-i\beta \hat{H}_x} | l \rangle \right|^2, \end{aligned} \quad (5.12)$$

where  $|\psi_{\text{targ}}\rangle$  is the  $h = 0$  target ground state, and the sum in the second line runs over the  $2^N$  basis states  $|l\rangle$  of the computational basis, with  $\hat{H}_z |l\rangle = E_l |l\rangle$ . Eq. (5.12) shows that  $\mathcal{F}$  is the scalar product of two  $2^N$ -dimensional unit vectors of components

$$\mathcal{F} = |\mathbf{v}^\dagger \cdot \mathbf{u}|^2, \quad (5.13)$$

where

$$(\mathbf{v}(\gamma))_l = \frac{1}{\sqrt{2^N}} e^{i\gamma E_l} \quad \text{and} \quad (\mathbf{u}(\beta))_l = \langle \psi_{\text{targ}} | e^{-i\beta \hat{H}_x} | l \rangle. \quad (5.14)$$

To ensure  $\mathcal{F} = 1$ , the Cauchy-Schwarz inequality requires  $\mathbf{v}(\gamma)$  and  $\mathbf{u}(\beta)$  to be parallel up to an overall phase factor  $\mathbf{u}(\beta) = e^{i\theta} \mathbf{v}(\gamma)$ . By imposing  $|(\mathbf{u}(\beta))_l|^2 = |(\mathbf{v}(\gamma))_l|^2$ , one obtains  $\beta = \frac{\pi}{4} \pmod{\pi}$ , as discussed in detail in appendix B.1. This result is rather general and holds both for a classical ferromagnetic target state  $|\psi_{\text{targ}}\rangle = |\uparrow\rangle^1$  and a quantum superposition  $|\psi_{\text{targ}}\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle + |\downarrow\rangle)$ , that corresponds respectively to  $p$  odd and  $p$  even in the  $p$ -spin model.

Furthermore, a unit fidelity further imposes [168] that all terms appearing in the sum in Eq. (5.12) are pure phase factors, which have to be identical for all states  $|l\rangle$ , modulo  $2\pi$ . For  $p$  odd, the energy is the odd power of an odd number  $E_l = -(M_l)^p$ ,  $M_l$  being the

<sup>1</sup> $|\uparrow\rangle = |\uparrow \dots \uparrow\rangle$  is the state with all spins pointing up.

magnetization eigenvalue,  $\widehat{S}_z |l\rangle = M_l |l\rangle$ , and the key observation is the following property of arithmetic congruence:

$$M_l^{p-1} = 1 \pmod{8} \implies M_l^p = M_l \pmod{8} \quad \text{if } p \text{ is odd.} \quad (5.15)$$

Hence, we can write

$$\begin{aligned} \mathcal{F}\left(\frac{\pi}{4}, \frac{\pi}{4}\right) &= \left| \frac{1}{2^N} \sum_l \exp\left(i \frac{2\pi}{8} \left[ (M_l^p + 2N_l^\downarrow) \pmod{8} \right] \right) \right|^2 \\ &= \left| \frac{1}{2^N} \sum_l \exp\left(i \frac{2\pi}{8} \left[ (M_l + 2N_l^\downarrow) \pmod{8} \right] \right) \right|^2 \\ &= \left| \frac{1}{2^N} \sum_l \exp\left(i \frac{2\pi}{8} [N \pmod{8}] \right) \right|^2 = 1, \end{aligned} \quad (5.16)$$

where we have inserted the contribution for the  $\mathbf{u}(\beta = \frac{\pi}{4})$  vector

$$\langle \psi_{\text{targ}} | e^{-i \frac{\pi}{4} \widehat{H}_x} |l\rangle = \frac{1}{\sqrt{2^N}} e^{i \frac{\pi}{2} N_l^\downarrow}. \quad (5.17)$$

$N_l^\downarrow = (N - M_l)/2$  is the number of spins pointing downward in the  $|l\rangle$  basis vector. Equation (5.16) demonstrates that  $(\gamma, \beta) = (\frac{\pi}{4}, \frac{\pi}{4})$  allows targeting exactly the ferromagnetic ground state  $|\uparrow\rangle$ , when both  $p$  and  $N$  are odd. A more detailed proof is given in appendix B.1, where we perform all the calculations explicitly, for both  $p$  odd and even. In the latter case, the proof is more complicated and we find that the precise values of  $\gamma$  leading to perfect fidelity depend on  $p$ . As a final remark, notice that in the theoretical proof we use for convenience the fidelity, instead of the residual energy. In general, however, we prefer the latter as a figure of merit, since it is directly linked to the variational minimization of the expectation value of the target Hamiltonian. Moreover computing the fidelity requires the full knowledge of the target ground state, which in general is not available for large systems. The energy instead is computed more easily and it is accessible also in experimental implementations of QAOA [98], without performing full tomography of the variational state.

## 5.4. Numerical results for $P > 1$

The possibility of finding exactly the ground state with  $P = 1$  holds only when the target Hamiltonian is classical (the transverse field is  $h = 0$ ) and the system has an odd number of spin variables  $N$ . When  $N$  is even, we have numerical evidence that  $P = 2$  is enough to target the ferromagnetic ground state for  $h = 0$ , but already at this level the efficiency of local optimization is severely limited by the presence of many local minima. Indeed we find that the energy landscape  $E_P(\boldsymbol{\gamma}, \boldsymbol{\beta})$  is extremely rugged for  $P \geq 2$ , as shown in Fig. 5.2, making local optimizations highly dependent on the initial set of parameters  $(\boldsymbol{\gamma}^0, \boldsymbol{\beta}^0)$ . Specifically, we use the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm [169]. We observe a very different behavior if the minimization is initialized with parameters  $\gamma_m^0$  and  $\beta_m^0$  chosen randomly in  $[0, \pi]$  (r-init), or rather with an initial guess based on a linear schedule,  $\gamma_m^0 = \frac{\Delta t}{h} \frac{m}{P} \frac{1}{N^{P-1}}$  and  $\beta_m^0 = \frac{\Delta t}{h} (1 - \frac{m}{P} (1 - h))$  (l-init).




Figure 5.2.: Section of the residual energy landscape  $\epsilon_{\text{res}}^P(\boldsymbol{\gamma}, \boldsymbol{\beta})$  along two random orthogonal directions, for  $N = 32$ ,  $p = 2$  and  $P = 2$ . The picture is centered on one minimum corresponding to the ferromagnetic ground state.

The results for the random initialization for  $p = 2$  are summarized in Fig. 5.3, where we show the normalized residual energy, Eq. (5.8), versus the number of QAOA steps  $P$  for  $h = \frac{\sqrt{5}-1}{2} < h_c$ , whose target state lies in the ferromagnetic phase for any value of  $p$ . Data for different system sizes  $N$  collapse perfectly after rescaling  $P \rightarrow (P - 2)/N$  (see inset of Fig. 5.3) and drop below machine precision at  $P = P_N^* = \frac{N}{2} + 2$ . Correspondingly, the variance of the residual energy distribution, which is rather large for  $P < P_N^*$  as witnessed by the error bars, drops to 0 at  $P_N^*$ , implying that all local minima become degenerate. The colored area around each curve shows the range between the lowest and the highest residual energy obtained for each value of  $P$  and  $N$ . The distribution of individual optimizations is symmetric around the average, with the exception of small values of  $P$  where r-init QAOA occasionally finds a local minimum with very small residual energy. For better readability, in the following figures we report only error-bars corresponding to the standard deviation of our data. This behavior holds for any value of the transverse field  $h$ , if the QAOA minimization is initialized with random parameters. In general, we find that the residual energy follows:

$$\epsilon_P^{\text{res}} = \begin{cases} \left(1 - \frac{P}{P_N^*}\right)^b & \text{if } P < P_N^* \\ 0 & \text{if } P \geq P_N^* \end{cases}, \quad (5.18)$$

with  $b \simeq 3$ . Remarkably, this scaling holds also for  $p > 2$ , see Fig. 5.4, with similar values of  $b$ , with the only difference that  $P_N^* = N + 1$  for  $p$  odd, because of the lack of the  $\mathbb{Z}_2$  symmetry. This implies that for *finite*  $N$  one can attain a perfect control of the state with a circuit depth  $P = P_N^* \propto N$ , physically corresponding to a total evolution time that scales as a power-law with  $N$ . Once again, this is at variance with a standard linear-schedule QA, where the total evolution time has to scale *exponentially* with  $N$  when the transition is first-order, i.e., for  $p > 2$ .

We have shown that a QAOA circuit with  $P = P_N^* \propto N$  is sufficient to prepare the exact

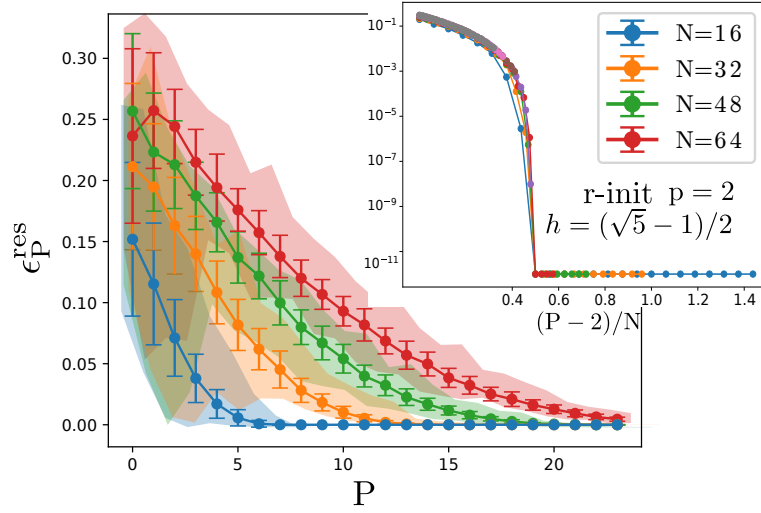


Figure 5.3.: Results of local optimizations with  $(\gamma^0, \beta^0)$  initialized randomly in  $[0, \pi]$  (r-init) averaged over 200 different realizations, for several values of  $N$  and  $h = \frac{\sqrt{5}-1}{2}$ , for  $p = 2$ . The shaded areas show the range between the best and the worst result obtained for each set of data. The inset shows the collapsed data (in log scale) after rescaling  $P \rightarrow (P - 2)/N$ .

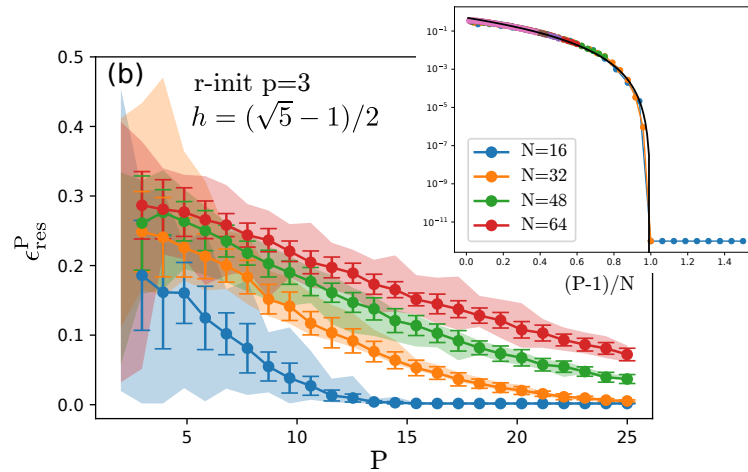


Figure 5.4.: Results of local optimizations with  $(\gamma^0, \beta^0)$  with averaged over 200 random initializations, for several values of  $N$  and  $h = \frac{\sqrt{5}-1}{2}$ , for  $p = 3$ . The inset shows the collapsed data (in log scale) after rescaling  $P \rightarrow (P - 1)/N$ . The solid black line is the scaling described in Eq. (5.18).






Figure 5.5.: Number of iterations required for BFGS algorithm to converge, averaged over 20 optimizations with random initializations of QAOA parameters. The corresponding value of  $P$  is  $N/2 + 2$  for  $p$  even, and  $N + 1$  for  $p$  odd, which is sufficient to obtain a residual energy below the numerical error.

ground state of the  $p$ -spin model for an arbitrary target  $h$ . However, to estimate the total computational complexity of running the QAOA algorithm to solve the  $p$ -spin model, we must include the computational cost of finding the QAOA variational parameters (with BFGS). Indeed, during the optimization process, the quantum device is used  $N_{\text{iter}}$  times, to sample the optimization landscape associated with QAOA circuits of  $P = P_N^*$ . In Fig. 5.5 we show the number of iterations  $N_{\text{iter}}$ , that the BFGS required for convergence as a function of  $N$ .  $N_{\text{iter}}$  appears to increase linearly with  $N$ , with a slope that only depends on the parity of  $p$ . Hence, the total computational time needed for converging to the exact ground state, at arbitrary transverse field, is at most polynomial in  $N$ , since it requires an order  $O(N)$  of iterations and a similar number of variational parameters, all in the range  $[0, \pi]$ .

A linear initialization of QAOA parameters, with a small noise (see caption of Fig. 5.6 for details), improves drastically the QAOA performance. This is illustrated in Fig. 5.6 where the results of the two competing schemes, random (r-init) versus linear (l-init) initialization, are shown for a system with  $N = 64$  for both  $p = 2$  (main plot) and  $p = 3$  (inset), and three fixed values of  $P = 5, 15, 25$ . Notice how the linear initialization is able to “detect” the quantum paramagnetic phase, for  $h > h_c$ , as being “easy”, with the QAOA minima found having vanishingly small residual energy, almost to machine precision, even if  $P < P_N^*$ . This occurs not only in the second-order transition case with  $p = 2$ , but also in the more “difficult” first-order case with  $p = 3$ . At variance with that, a random initialization performs on average quite independently of the target transverse field  $h$ , and knows nothing about the location of the critical field. Interestingly, this suggests that QAOA for small  $P$  is sensitive to the phase diagram of the target Hamiltonian: choosing a good Ansatz for the initial parameter set  $(\gamma^0, \beta^0)$  is fundamental to initialize the variational wave-function in a “good” basin of attraction, where the minimization leads to small values of  $\epsilon_p^{\text{res}}$ . Whether this feature is unique to infinite range models or is a common property of long range Hamiltonians is an interesting issue to pursue in future works.

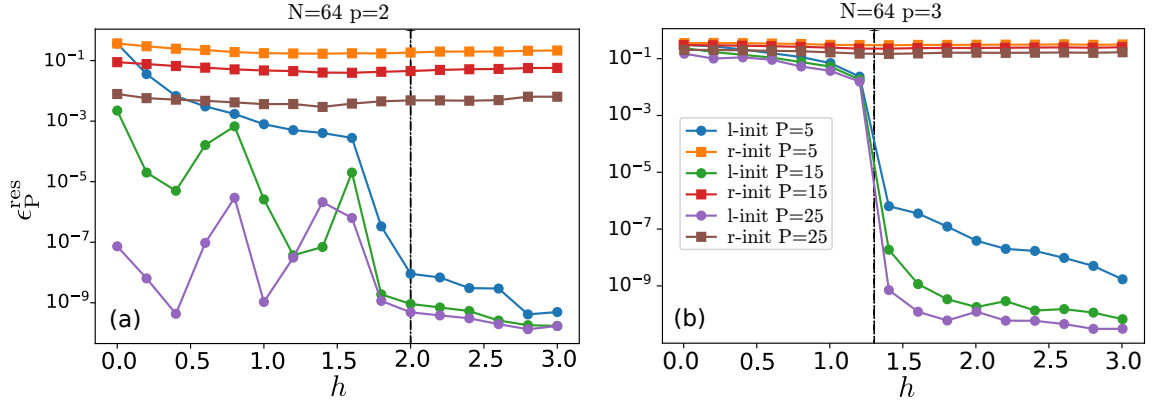


Figure 5.6.: QAOA residual energy versus the transverse field  $h$  for a system with  $N = 64$ , and three values of  $P = 5, 15, 25$ , for random (r-init) and linear initialization (l-init) of the QAOA parameters. Notice that  $P_N^* = N/2 + 2 = 34$  for  $p = 2$  and  $P_N^* = N + 1 = 65$  for  $p = 3$ . The vertical dashed lines show the critical transverse field:  $h_c = 2$  for  $p = 2$  in panel (a) and  $h_c \simeq 1.3$  for  $p = 3$  in panel (b). The linear initialization (l-init) corresponds to Eq. (5.9) multiplied element-wise by a noise factor  $(1 + r)$ , with  $r \in [-0.05, 0.05]^{2P}$  a vector of uniformly distributed random numbers. Data are averaged over 100 different instances of  $r$ .

The linear initialization displays better efficiency, compared to the random one, also when the target state belongs to the ferromagnetic phase ( $h < h_c$ ), and  $P < P_N^*$ . This is illustrated in Fig. 5.7 for  $p = 2$  (a) and  $p = 3$  (b). Here, however, the improvement is only quantitative —  $\epsilon_P^{\text{res}}$  decreases faster and scales better with system size — since the actual change in the landscape, with degenerate global minima, occurs only at  $P_N^*$ . Moreover, the system displays a large roughness of the variational energy landscape, which makes the task of obtaining good variational minima extremely demanding, especially for  $p \geq 3$ , hence justifying the poorer improvement of l-init over r-init observed in Fig. 5.7(b).

A smooth change of the control parameters is required, or at least useful, for experimental implementations of QAOA algorithms [98]. Finding local minima  $(\gamma^*, \beta^*)$  which can be seen as the discretization of some continuous function, proves however to be a difficult task for this model. In contrast with Refs. [95, 96], an iterative procedure that initializes  $(\gamma^0, \beta^0)$  from an interpolation of a smooth set obtained for a smaller parameter space does not seem to work straightforwardly. Our failed attempts do not exclude that smart smooth choices for  $(\gamma^0, \beta^0)$  can be constructed: they only signal that finding them is a non-trivial task, due to the extreme roughness of the energy landscape. The linear initialization we have adopted is able to find reasonably smooth  $(\gamma^*, \beta^*)$  only for small values of  $P$ , as reported in Fig. 5.8. In the upper panels, it is shown the expectation value of the target Hamiltonian during the digital evolution  $E_m = \langle \psi_m | \hat{H}_{\text{target}} | \psi_m \rangle$ , with  $|\psi_m\rangle$  being the wave-function evolved with only the first  $m$  out of  $P$  parameters. In the lower panel, it is shown the interpolation parameter

$$s_m = \frac{N^{P-1} \gamma_m}{(1-h)N^{P-1} \gamma_m + \beta_m}. \quad (5.19)$$

As the dimensionality of the parameter space increases, and so does the roughness and the number of local minima, the optimal parameters obtained starting from a linear initialization

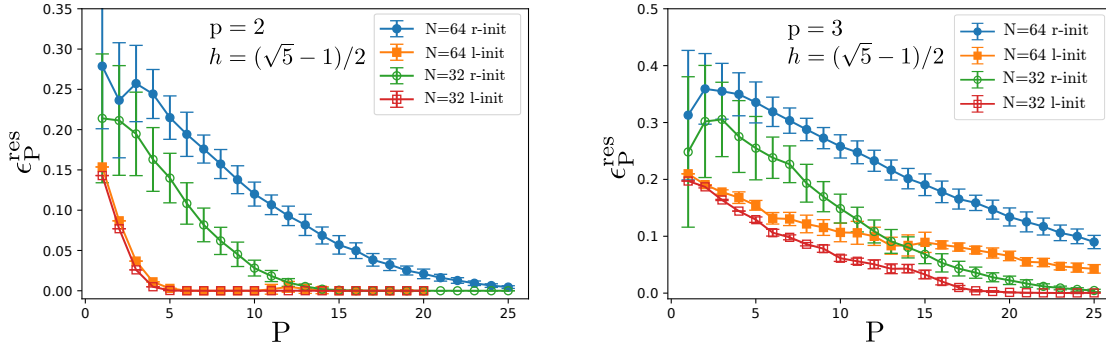


Figure 5.7.: Comparison between the optimized residual energy obtained from a linear initial guess plus small noise (l-init) and from random initialization (r-init), for two system sizes  $N = 32$  and  $N = 64$ . In (a)  $p = 2$ , in (b)  $p = 3$ .

scheme appear to be increasingly irregular, and so does the expectation value of the energy.

## 5.5. Conclusions

We analyzed the performance of QAOA on the fully connected  $p$ -spin model, showing that it is able to find exactly the ferromagnetic ground state with polynomial resources, even when the system encounters a first-order phase transition. In particular, the algorithm prepares the ground state of  $\hat{H}_z$  with only  $P = 1$  (if  $N$  is odd) or  $P = 2$  (if  $N$  is even) steps, with a corresponding evolution time that scales as  $N^{P-1}$ , while QA would require an exponentially long annealing time. This exact minimum however exists only for zero transverse field,  $h = 0$ . Interestingly, the exact minimum, which clearly survives for  $P \geq 2$ , is very hard to find with gradient-based optimization schemes due to the extreme roughness of the energy landscape, especially for  $p > 2$ . The “hardness” of the problem for  $p > 2$  is thus reflected in the difficulty in finding the correct absolute minimum, rather than in the resources (i.e. the computational time) needed.

The performance of the optimization itself strongly depends on the initialization of the variational parameters  $(\gamma^0, \beta^0)$ . For a random initialization, the residual energy drops below machine precision as  $(P_N^* - P)^b$ , with  $b \sim 3$  and  $P_N^*$  growing linearly with  $N$ . This behavior is independent of the target transverse field  $h$  and from  $p$ , with the only difference that  $P_N^* = N/2 + 2$  for  $p$  even and  $P_N^* = N + 1$  for  $p$  odd. With a linear initialization, the algorithm performs much better and is able to detect the presence of a phase transition, although the improvement deteriorates rapidly as  $P$  increases, because of the growing number of “bad” local minima.

For future developments, it would be interesting to understand whether infinite or long-range Hamiltonians can be used to boost QAOA performance on short-range models. The idea is to add a further unitary  $e^{-i\epsilon\hat{H}'_z}$  in Eq. (5.5), generated by a long-range Hamiltonian  $\hat{H}'_z$  unrelated to the problem to solve. This enlarges the portion of Hilbert space approximated with a QAOA Ansatz, at a fixed number of Trotter steps  $P$ , at the price however of increasing




Figure 5.8.: Comparison between different values of  $P$ , of the expectation value of the target Hamiltonian at the  $m$ -th step  $E_m$  and the interpolation parameter  $s_m$ . The thick line is the average over 50 l-init optimizations, while the faint gray lines are single instances. For  $P = 5$  all optimization give the same set of parameters, hence only the average is visible. The transverse field is  $h = \frac{\sqrt{5}-1}{2}$ .

the number of variational parameters.

At variance with Refs. [95, 96, 98], we are unable to construct minima in the energy landscape associated with smooth parameters  $(\gamma^*, \beta^*)$ . These regular parameter choices are often desired since they might be linked to adiabatic schedules, that can be used to infer the optimal protocol  $s(t)$  in a continuous annealing scenario, and allow for a faster minimum search in the  $2P$ -dimensional parameter space, once a solution for  $P' < P$  is known [96, 98]. Preliminary results [170] with reinforcement learning [103] methods applied to the QAOA evolution suggest however that smooth choices of  $(\gamma^*, \beta^*)$  do indeed exist, but they are hard to find with local optimizations. This topic will be investigated in the next chapter. Whether *global* minima are related to smooth values of  $(\gamma^*, \beta^*)$  remains an open and interesting question.



## 6. Reinforcement Learning for quantum optimization

I discuss here a reinforcement learning (RL) scheme for feedback quantum control within the quantum approximate optimization algorithm (QAOA), which we have recently published [170].

The main idea presented in this chapter is to rephrase the minimum search in the energy landscape of QAOA, discussed in the previous chapter, as a reinforcement learning process. An RL agent learns a strategy (called policy) to associate a particular action choice — in the QAOA framework, a pair of parameters  $\gamma$  and  $\beta$  — to any state  $|\psi\rangle$  that belongs to the region of the Hilbert space explored during the training. The information on the current state of the system is provided by measuring a set of observables.

The difference between the standard QAOA process and our proposal is summarized in Fig. 6.1. While QAOA treats the digitized evolution as a whole and looks for a minimum of the variational energy  $E(\gamma, \beta)$  in the  $2P$  dimensional parameter space, an RL-assisted optimization gets partial information on the wave-function at each discrete time step  $t$  and then evolves the system, according to a learned strategy, with the goal of minimizing the final energy. This feedback control mechanism is repeated until  $P$  steps are covered and the expectation value of the target Hamiltonian is computed.

I will show that our RL scheme finds a policy converging to the optimal adiabatic solution for QAOA, found by Mbeng *et al.* [96] for the translationally invariant quantum Ising chain. The same result holds in the presence of disorder, where the policy can be transferred among different disorder instances and system sizes. Indeed, our RL scheme allows the training part to be performed on small samples and transferred successfully on larger systems and different disordered realizations.

RL-assisted QAOA is able to find smooth regular schedules for the  $p$ -spin model (with  $p = 2$ ) described in chapter 5, in sharp contrast with local minimization alone. It improves also the stability of the QAOA process, leading to much better results for large values of  $P$ , and allowing for a better estimate of the scaling of the residual energy with  $P$  and the system size  $N$ .



**Disclaimer** In this chapter, there is a slight change in notation with respect to the previous one. Here we will indicate with  $t$  the discrete time-step in a digitized evolution, which previously was denoted with  $m$ . For instance, the QAOA parameters will be labelled as  $\gamma_t$  and  $\beta_t$  instead of  $\gamma_m$  and  $\beta_m$ . This is to keep our notation consistent with that standard in Reinforcement Learning.

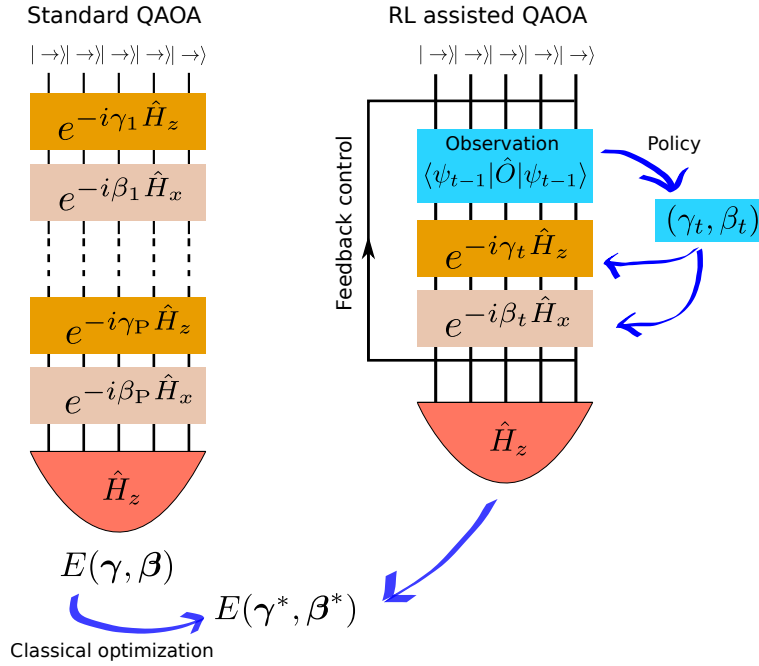


Figure 6.1.: Pictorial representation of the difference between standard QAOA and Reinforcement Learning assisted QAOA.

## 6.1. Introduction

Reinforcement Learning [103] (RL) is a class of machine learning algorithms where an *agent* faces the task of maximizing a reward function, by acting on an environment. The agent has no prior knowledge of which are good strategies to follow, hence it learns, by trials and errors, how its actions modify the environment and influence the reward it receives. This “learning from experience alone” makes RL suitable for several control processes, such as automated driving [171–173] and robot controls [174, 175]. Moreover, RL has been successfully applied in complex strategic games such as Go [176, 177], where future planning is essential to any winning strategy.

Our idea here is to apply RL as a feedback control mechanism on quantum optimization, specifically for a digitized quantum evolution such as the *Ansatz* of the Quantum Approximate Optimization Algorithm (QAOA) [42, 44], which we used already in Chap. 5.

Before entering in the details of our proposals, let us recap briefly the methods used for quantum optimization that are relevant for our discussion.



**Info:** As commonly done in optimization techniques for physical problems, we restrict ourselves to spin-1/2 Hamiltonians.

In Quantum Annealing (QA) [28, 29, 85–87] *alias* Adiabatic Quantum Computation (AQC) [30, 88] (QA/AQC) one constructs an interpolating Hamiltonian  $\hat{H}(s) = s\hat{H}_z + (1-s)\hat{H}_x$ , where  $\hat{H}_z$  is the problem Hamiltonian whose ground state (GS) we are searching [178], while



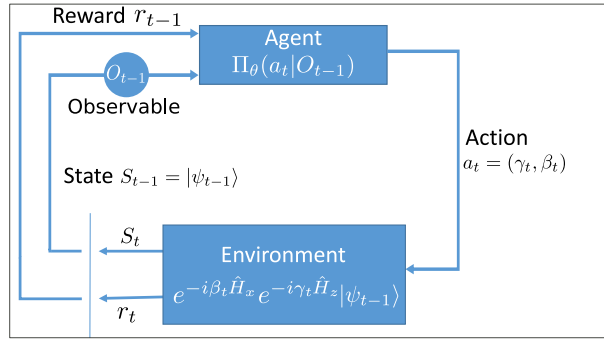


Figure 6.2.: Scheme of a single step of Reinforcement Learning for QAOA.

$\hat{H}_x = -h \sum_j \hat{\sigma}_j^x$  is a transverse field term. An adiabatic dynamics is then attempted by slowly increasing  $s(t)$  from  $s(0) = 0$  to  $s(\tau) = 1$  in a large annealing time  $\tau$ , starting from some easy-to-prepare initial state  $|+\rangle$ , the GS of  $\hat{H}_x$ . The difficulty is usually associated with the growing annealing time  $\tau$  necessary when the system crosses a transition point, especially of first order [179].

QAOA, instead, uses a variational *Ansatz* of the form

$$|\psi_P(\boldsymbol{\gamma}, \boldsymbol{\beta})\rangle = \left( \prod_t^{P \leftarrow 1} e^{-i\beta_t \hat{H}_x} e^{-i\gamma_t \hat{H}_z} \right) |+\rangle, \quad (6.1)$$

where  $\boldsymbol{\gamma} = \gamma_1, \dots, \gamma_P$  and  $\boldsymbol{\beta} = \beta_1, \dots, \beta_P$  are  $2P$  real parameters. The product is a time-ordered product with  $t = 1 \dots P$  increasing from right to left. The standard QAOA approach consists of a classical minimum search in such a  $2P$ -dimensional energy landscape, which is, in general, not a trivial task [180]. Indeed, there are in general very many local minima in the QAOA-landscape, and local optimizations with random starting points produce irregular parameter sets  $(\boldsymbol{\gamma}^*, \boldsymbol{\beta}^*)$ , hard to implement and sensitive to noise. To obtain stable and regular solutions  $(\boldsymbol{\gamma}^*, \boldsymbol{\beta}^*)$  that can be easily generalized to different values of  $P$  and implemented experimentally, it is necessary to employ iterative procedures during the minimum search [96, 98, 181]. Interestingly, as discovered in Ref. [96] for quantum Ising chains, smooth regular optimal schedules for  $\gamma_t$  and  $\beta_t$  can be found, which are *adiabatic* in a digitized-QA/AQC [91] context.

One might indeed reformulate the QAOA minimization as an optimal control process [182] in which one acts sequentially on the system in order to maximize a final reward. This reformulation seems particularly suited for Reinforcement Learning (RL) [103, 176, 183, 184]. As schematically represented in Fig. 6.2, at each discrete time step  $t$  an “agent” is given some information, typically through measuring some observables  $O_{t-1}$  on the state  $S_{t-1} = |\psi_{t-1}\rangle$  of the system on which it acts (the “environment”). The agent then performs an action  $a_t$  — here choosing the appropriate  $(\gamma_t, \beta_t)$  and applying the corresponding unitaries to the state — obtaining a new state  $S_t = |\psi_t\rangle$  and receiving a “reward”  $r_t$ , measuring the quality of the variational state constructed.

In our work we address several questions that have not been investigated in the literature that flourished recently on RL applied to quantum optimization [185–190]:

- i) is such RL-assisted QAOA able to “learn” *optimal* schedules?
- ii) Are the schedules found *smooth* in  $t$ ?
- iii) How to dwell with the fact that getting information from  $|\psi_t\rangle$  involves quantum measurements which *destroy* the state?
- iv) Are the strategies learned easily *transferable* to larger systems?

In this chapter I show, on the paradigmatic example of the transverse field Ising chain, that optimal strategies — well known in that case, see Ref. [96] — can be effectively learned by an agent trained with Proximal Policy Optimization (PPO) algorithm [191] employing very small neural networks (NN). We show that RL automatically learns *smooth* control parameters, hence realizing an optimal controlled digitized-QA algorithm [96, 192]. By working with disordered quantum Ising chains we show that strategies “learned” on small samples can be successfully transferred to larger systems, hence alleviating the “measurement problem”: one can learn a strategy on a small problem which can be simulated on a computer, and implement it on a larger experimental setup [193]. Finally, we show that RL stabilizes the performance of QAOA applied on the  $p = 2$  p-spin model, investigated in chapter 5, and learns smooth schedules, which we did not find using the standard QAOA approach.

## 6.2. Basis of RL

Here we provide a very concise introduction to Reinforcement Learning, discussed in more detail in appendix C.

The term “Reinforcement Learning” refers to an entire class of algorithms and methods developed to optimize sequences of consecutive actions in order to obtain a well defined target objective. While several key concepts of RL were developed already in past decades, the field is currently seeing a new surge of interest, thanks also to its success in tackling complex strategic tasks such as those necessary in real-time strategic videogames.

At their core, RL algorithms deal with an *agent* which interacts with an outside *environment*, i.e. changing its *state* ( $s$ ) through some *actions* ( $a$ ). The agent receives information on the consequences of its actions in the form of the new *state* of the environment and the instantaneous *reward* ( $r$ ) relative to its objective. RL algorithms are particularly useful when a long-term strategy is fundamental to the solution of the problem, i.e. in cases where, quoting Ref. [103], “*actions influence not just immediate rewards, but also subsequent situations, or states, and through those future rewards*”. The schematics of an agent/environment system are represented in Fig. 6.2, where all the elements are already specified in the QAOA framework. The goal of any RL algorithm is to find the optimal *policy*, i.e., a mapping from each possible *state* to the distribution probability of *actions* to take, aiming at maximizing a future *reward*.

Two key concepts in RL are the *policy* and the *state value function*. The policy  $\Pi(a|s)$  is the strategy followed by the RL agent and, in a stochastic process language, it describes the conditional probability of taking the action  $a$ , knowing that the environment is in the state  $s$ .

The state value function  $V_{\Pi}(s)$  describes how well the policy  $\Pi$  performs, when evolving the environment from the initial state  $s$ . It is a well defined (but generally unknown) quantity:

$$V^{\Pi}(s) = \mathbf{E}_{\Pi}[\sum_t \lambda^t r_t | s_0 = s] .$$

It represents the expected cumulative future reward that one will obtain starting from  $s$ , always choosing actions following a given policy  $a \sim \Pi$ . The discount-factor  $\lambda \in [0, 1]$  is used to modulate the relative importance of long-term and short-term rewards<sup>1</sup>

The basic structure of an RL algorithm is to produce trajectories  $\mathbf{T}$  selecting actions that follow a *policy*  $\Pi_k$  and recording at each time  $t$  the step  $(s_t, a_t, r_t, s_{t+1})$ . Using the experienced data, a new updated policy  $\Pi_{k+1}$  is constructed such that actions that “performed best” have increased probability to happen. The period between the policy updates is called an *epoch*. New trajectories are produced following  $\Pi_{k+1}$  and the process is repeated.

RL-algorithms differ mainly in *i)* how the performance of the actions is evaluated *ii)* how the value of states is approximated and *iii)* how policy is updated. What specific flavor of RL algorithms is best suited to solve a task depends therefore on the specific features of the problem at hand, e.g., whether the state and action spaces are discrete or continuous or whether the agent has access to a complete description of the state of the environment or only to a partial observation.

Here, we utilize a Proximal-Policy Optimization (PPO) algorithm [191], which belongs to the family of actor-critic algorithms. The detailed description of the implementation of the PPO algorithm can be found here, and is well outside the scope of this work: In the following, we describe only the basic ideas involved, while some extra details are given in appendix C.3. In PPO two Neural Networks (NNs) are used as parametric approximations of the functions describing policy  $\Pi(a|s)$  (the actor) and the value  $V^{\Pi}(s)$  (critic).

**Value function** The value function is an expectation over a potentially infinite number of trajectories, hence the experience gained from those sampled in a single epoch in general does not contain sufficient information to faithfully approximate it. The rewards obtained in the sampled trajectories can, however, be used to extrapolate better estimations of the value functions with different methods. At the end of each epoch, the NN is updated to minimize the error between the parametrized value function  $V_{\theta'}^{\Pi}(s)$  and the new one extrapolated from experience  $V_{\mathbf{T}}^{\Pi}(s)$  for all visited states.

**Policy** To update the policy, the outcome of the experienced actions has to be evaluated. The “goodness” of having chosen at time  $t$  a particular action  $a_t$ , is measured with the Advantage function

$$A^{\Pi}(s_t, a_t) = r_t + \lambda V^{\Pi}(s_{t+1}) - V^{\Pi}(s_t) \approx r_t + V_{\theta'}^{\Pi}(s_{t+1}) - V_{\theta'}^{\Pi}(s_t). \quad (6.2)$$

This seemingly complex expression has a conceptually simple meaning: The expectation value of being in the state  $s_t$ ,  $V_{\theta'}^{\Pi}(s_t)$ , is measured against the known fact that it has experimentally

<sup>1</sup>Usually in RL notation, the discount factor is denoted with  $\gamma$ . Here I use  $\lambda$  to avoid confusion with the  $\gamma$  parameter of QAOA.

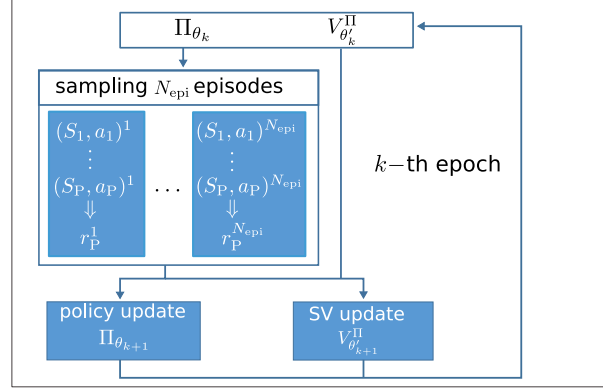


Figure 6.3.: Scheme of the “episode” loop in each  $k$ -th training “epoch”, with the “policy” and “state-value” neural networks  $\Pi_{\theta_k}$  and  $V_{\theta_k}^{\Pi}$ .

produced a reward  $r_t$  and it has changed the state of the environment to  $s_{t+1}$  (with its expectation value  $V_{\theta}^{\Pi}(s_{t+1})$ ). The policy is then updated such that the probability of taking actions with a positive score of  $A^{\Pi}$  is increased, and conversely that of actions with negative  $A^{\Pi}$  is decreased. A sketch of the training process of the PPO algorithm is depicted in Fig. 6.3. The defining feature of PPO is that a clipping function is used such that the changes between old policy and new policy cannot be larger than a fixed amount. This simple additional feature adds remarkable stability to the learning process.

### 6.3. RL assisted QAOA

The key ingredients of the RL-assisted algorithm, as schematized in Fig. 6.2, are as follows.

**State)** The *state*  $S_t$  at time step  $t = 1, \dots, P$  is encoded by the wave-function  $|\psi_t\rangle$ , defined iteratively as  $|\psi_t\rangle = e^{-i\beta_t \hat{H}_x} e^{-i\gamma_t \hat{H}_z} |\psi_{t-1}\rangle$ , with  $|\psi_0\rangle = |+\rangle = \frac{1}{\sqrt{2^N}} \otimes_i (|\uparrow\rangle_i + |\downarrow\rangle_i)$ . Due to the symmetry of both  $\hat{H}_x$  and  $\hat{H}_z$ ,  $|\psi_t\rangle$  is always  $\mathbb{Z}_2$  symmetric. The agent has partial information through a number of *observables*  $O_{t-1}$  measured on  $|\psi_{t-1}\rangle$ . Our choice (with  $t-1 \rightarrow t$ ) is

$$O_t = \{ \langle \psi_t | \hat{\sigma}_j^z \hat{\sigma}_{j+1}^z | \psi_t \rangle, \langle \psi_t | \hat{\sigma}_j^x | \psi_t \rangle \}, \quad (6.3)$$

where a single value of  $j$  is enough when translational invariance is respected. Interestingly, the agent seems to achieve comparable results even with a single observable  $O_t = \langle \psi_t | \hat{\sigma}_j^z \hat{\sigma}_{j+1}^z | \psi_t \rangle$ , see appendix C.4 for further information.

**Action)** The action  $a_t$  at time  $t$  corresponds to choosing  $(\gamma_t, \beta_t)$ . The conditional probability of  $a_t$  given the observables  $O_{t-1}$ , the policy, is denoted by  $\Pi_{\theta}(a_t | O_{t-1})$ , where  $\theta$  are the parameters of a Neural Network (NN) encoding. Our policy is stochastic, to help exploration:  $\Pi_{\theta}(a | O)$  is chosen as a Gaussian distribution, whose mean and standard deviation are computed by the NN. From this,  $a_t = (\gamma_t, \beta_t)$  is extracted.<sup>2</sup>

<sup>2</sup>At variance with section 6.2, here I specify that the policy depends on the state  $s$  through the observable

**Reward)** A reward  $r_t$  is calculated at time  $t$ . In our present implementation,  $r_{t=1,\dots,P-1} = 0$  and only  $r_P > 0$ . The final reward  $r_P = R(E_P)$  is associated to minimizing the final expectation value  $E_P = \langle \psi_P | \hat{H}_{\text{target}} | \psi_P \rangle$ . Here  $R(E_P)$  is monotonically increasing when  $E_P$  decreases. Specifically, we take  $R(E_P) = -E_P$ , but different non-linear choices have been tested.

**Training)** The training process consists of a number  $N_{\text{epo}}$  of “epochs”, as sketched in Fig. 6.2(b). During each epoch the RL agent explores, with a *fixed* policy, the state-action trajectories for a certain number  $N_{\text{epi}}$  of “episodes”, each episode involving  $P$  steps  $t = 1, \dots, P$ . At the end of each epoch, the policy is updated to favor trajectories with higher rewards. The particular RL algorithm we used is the Proximal Policy Optimization (PPO) algorithm [191], from the OpenAI SpinningUp library [194], which is well suited for continuous state-action pairs. In our numerical simulations, we used NNs with two fully-connected hidden layers of 32, 16 neurons, and linear-rectification (ReLU) activation function, which is defined as

$$f(x) = \begin{cases} 0 & \text{if } x < 0, \\ x & \text{if } x \geq 0. \end{cases} \quad (6.4)$$

It has been chosen because it has always a positive gradient when the signal is positive ( $x > 0$ ), leading to a faster policy optimization compared to a sigmoid<sup>3</sup> activation function [103].

## 6.4. Results

We test our scheme on three different spin models, namely the transverse field Ising model (TFIM), the random Ising model (rTFIM) and the Lipkin-Meshov-Glick (LMG) model. The latter is equivalent to the  $p$ -spin model studied in chapter 5 with  $p = 2$ . All these models are solvable with exact numerical techniques, which allows for a full characterization of the dynamics: the one-dimensional TFIM can be diagonalized with the Jordan-Wigner transformation [195], while the  $p$ -spin is restricted to the maximally polarized subsector of the Hilbert space. Moreover, both the TFIM and the  $p$ -spin model have been studied in depth concerning QAOA processes [96, 147], so they represent perfect benchmarks for RL assisted quantum optimization.

### 6.4.1. Ising model with uniform couplings

Let us first consider the transverse field Ising model (TFIM) in one dimension, with periodic boundary conditions, where detailed QAOA results are already known [96]. Specifically, we

---

$O(s)$ , although in RL literature the notation  $\Pi(a|s)$  is more common even when the agent has only a partial knowledge on the state.

<sup>3</sup>A sigmoid is a function which interpolates smoothly between 0 and 1, such as  $(1 + e^{-x})^{-1}$ . It has zero gradient when  $|x| \gg 1$  but it is differentiable in  $x = 0$ , at variance with a ReLU function. The best choice depends on the specific problem at hand.

define the target Hamiltonian  $\hat{H}_{\text{target}} = \hat{H}_z + h\hat{H}_x$  with

$$\hat{H}_z = - \sum_{j=1}^N J_j \hat{\sigma}_j^z \hat{\sigma}_{j+1}^z, \quad \hat{H}_x = - \sum_j \hat{\sigma}_j^x. \quad (6.5)$$

We start considering the uniform TFIM, where  $J_j = J$ . The model has a paramagnetic ( $h > J$ ) and a ferromagnetic ( $h < J$ ) phase, separated by a 2<sup>nd</sup>-order transition at  $h = J$ . The performance of QAOA on the uniform TFIM chain has been studied in detail in Refs. [96,162]. Given a set of QAOA parameters  $(\gamma, \beta)$ , we gauge the quality of the resulting state from the residual energy density

$$\epsilon_P^{\text{res}}(\gamma, \beta) = \frac{E_P(\gamma, \beta) - E_{\min}}{E_{\max} - E_{\min}}, \quad (6.6)$$

where  $E_P(\gamma, \beta) = \langle \psi_P(\gamma, \beta) | \hat{H}_{\text{target}} | \psi_P(\gamma, \beta) \rangle$  is the variational energy, and  $E_{\max}$  and  $E_{\min}$  are the highest and lowest eigenvalues of the target Hamiltonian. Specifically, the results presented below will concern targeting the classical state for  $h = 0$ , although the approach can be easily extended to the case with  $h > 0$ . At  $h = 0$  the residual energy is bounded by the inequality [96]

$$\epsilon_P^{\text{res}}(\gamma, \beta) \geq \begin{cases} \frac{1}{2P+2} & \text{if } 2P < N \\ 0 & \text{if } 2P \geq N \end{cases}, \quad (6.7)$$

which becomes an equality if and only if  $(\gamma, \beta)$  are optimal QAOA parameters.

In the RL training, the system is initially prepared in the state  $|\psi_0\rangle = |+\rangle$ , while the NNs for the policy and the state-value function are both initialized with random parameters. The agent is then trained for  $N_{\text{epo}} = 1024$  epochs, each comprising  $N_{\text{epi}} = 100$  episodes of  $P$  steps each. After training, we test the RL algorithm with  $\sim 50$  runs.




Figure 6.4.: Residual energy density  $\epsilon_P^{\text{res}}$ , Eq. (6.6), vs  $P$ . The target state is ferromagnetic with  $h = 0$ . Full symbols: results from RL only; empty symbols: a local optimization (LO) supplements the RL actions (RL+LO); data are averaged over 50 test runs. The black dashed line is the lower bound of Eq. (6.7)

Fig. 6.4(a) shows the results obtained by the RL-trained policy. For  $P \leq 6$ , the trained RL agent finds optimal QAOA parameters, saturating the bound for  $\epsilon_P^{\text{res}}$  in Eq. (6.7). In particular, for small system sizes  $N$ , when  $P > N/2$ , the agent finds the exact target ground state, and  $\epsilon_P^{\text{res}} = 0$ . For longer episodes ( $P > 6$ ), the residual energy deviates from the lower bound due to two factors: *i*) the longer the episode, the more difficult it is to learn the policy, as a larger number of training epochs are necessary to reach convergence; *ii*) since we are using a stochastic policy, the error due to the finite width of the action distributions is accumulated during an episode, leading to larger relative errors for longer trajectories. To cure this fact, we adopted the following strategy: we supplement the RL-trained policy with a final *local optimization* (LO) of the parameters  $(\gamma, \beta)$ , employing the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm [169]. This last step is computationally cheap since the RL training brings the agent already close to a local minimum, provided  $N_{\text{epo}}$  is large enough. The residual energy data obtained in this way, denoted by RL+LO in Fig. 6.4, falls on top of the optimal curve  $\epsilon_P^{\text{res}} = \frac{1}{2P+2}$ .




Figure 6.5.: The schedule  $s_t = \gamma_t / (\gamma_t + \beta_t)$ . Full blue lines denote  $s_t$  learned after  $N_{\text{epo}} = 1024$  epochs on a chain of  $N = 128$  sites; Dashed red lines, the RL+LO results; Black empty squares, the iterative LO smooth solution [96]. The RL actions are in the basin of the same optimal minimum. Inset: same data for  $N_{\text{epo}} = 128$  training epochs, where not all the LO optimized actions sets fall onto the iterative LO solution.

To visualize the action choices, we translate  $\gamma_t$  and  $\beta_t$  into the corresponding interpolation parameter  $s_t$  which a Trotter-digitised QA/AQC would show, which for  $h = 0$  is given by: [96]

$$s_t = \frac{\gamma_t}{\gamma_t + \beta_t}. \quad (6.8)$$

Fig. 6.5 shows the interpolation parameter  $s_t$  during an episode  $t = 1, \dots, P$ , for a chain of  $N = 128$  spins and  $P = 8$ . Different curves are obtained by repeating a test run of the same stochastic policy, trained for  $N_{\text{epo}} = 1024$  epochs. The parameters obtained through the RL policy are smooth, and different tests result in similar s-shaped profiles for  $s_t$ . When a final local minimization is added, the curves for  $s_t$  coalesce and coincides with the smooth optimal schedule obtained in Ref. [96] through an independent iterative local optimization




Figure 6.6.: Residual energy, Eq. (6.6), vs  $P$  for a single instance of the random TFIM: comparison between bare RL and RL followed by local optimization (LO) results (RL+LO).

strategy. When the training is at an early stage, i.e., the number of epochs is small, see inset of Fig. 6.4(b), the profiles  $s_t$  are more irregular and do not fall all in the same smooth minimum upon performing the LO (see the three dashed red lines in the inset).

#### 6.4.2. Ising model with random couplings

Next, we turn to the random TFIM case, described by the Hamiltonian

$$\hat{H}_z = - \sum_j J_j \hat{\sigma}_j^z \hat{\sigma}_{j+1}^z, \quad \hat{H}_x = - \sum_j \hat{\sigma}_j^x. \quad (6.9)$$

The couplings  $J_j \in [0, 1]$  are uniformly distributed random numbers. Here, for each chain length  $N$  we fix a given disorder instance  $\{J_j\}_{j=1, \dots, N}$ , both for the training and the test of the RL policy. Since translational invariance is now lost, one would naively imagine that the relevant observables  $O_t$  in Eq. (6.3) would involve a list of  $2N$  measurements. However, our experience has taught us that we can efficiently go on with a reduced list comprising only the two Hamiltonian terms,  $O_t = \{\langle \psi_t | \hat{H}_z | \psi_t \rangle, \langle \psi_t | \hat{H}_x | \psi_t \rangle\}$ , hence chain-averaged quantities. All the parameters involved in training the NNs are fixed as in the uniform TFIM case.

Fig. 6.6 shows the residual energy  $\epsilon_P^{\text{res}}$  vs  $P$  obtained from the bare RL (full symbols) and from RL followed by a local optimization (RL+LO, empty symbols). The local optimization significantly improves the quality for large  $P \geq 10$ . A detailed study of the behaviour of  $\epsilon_P^{\text{res}}$  for large  $P$  and a comparison with the results obtained [164] by a linear-QA/AQC scheme, with  $s(t) = t/\tau$ , is left to a future study.

Fig. 6.7 shows the optimal parameter  $s_t = \gamma_t / (\gamma_t + \beta_t)$  found by the RL+LO method, compared to the  $s_t$  constructed with the iterative optimization strategy described in Ref. [96]: the agreement between the two is remarkable, showing that the RL-assisted QAOA effectively “learns” smooth action trajectories.






Figure 6.7.: The optimized  $s_t$  obtained with different procedures. Empty squares: the iterative LO process of Ref. [96]; Blue circles: RL+LO performed directly on a  $N = 128$  chain; Gray lines:  $RL_{N=8}+LO_{N=128}$ , i.e., training of a  $N = 8$  chain used as *Ansatz* for LO of the  $N = 128$  chain.

The most remarkable fact, however, is shown by the series of grey lines present in Fig. 6.7. These are obtained by training the RL agent on a much smaller instance with  $N = 8$  sites, and then use the suggested actions to initialize local optimizations for a larger system with  $N = 128$ . Each line corresponds to a different disorder instance of the longer “test” chain. These results show that a single trained policy can be easily adapted to any realization of the rTFIM model, changing both the couplings and the chain length, without decreasing the quality of the approximate ground state  $|\psi_P(\gamma^*, \beta^*)\rangle$  compared to a policy trained directly on the target system. It also reduces drastically the resource needed for RL assisted optimization, because training the policy over a large system is computationally much heavier than training on a smaller one. This approach does not even change the final residual energy after the local optimization on top of the RL actions. In Fig. 6.8 we report the residual energies obtained on 10 random instances of a system of  $N_{\text{test}} = 128$  spins, with policies trained on chains of different length  $N_{\text{tr}}$ . The dashed black line refers to the residual energy of the coupling choice where the largest system ( $N_{\text{tr}} = 128$ ) has been trained; after a local optimization this value is totally independent of the actual training instance and system size, showing that the RL can be perfectly transferred. Indeed the dispersion of the data of each training size is due only on the dispersion of the residual energies of different disordered instances.

The actions  $(\gamma_t, \beta_t)$  can be even be transferred directly from the  $N = 8$  chain to the  $N = 128$  one without a final local minimization, obtaining residual energies comparable with an RL agent trained on the larger system, again without the action optimization at the end. We have to remark however that only the actions and not the policy itself can be applied on different random instances; indeed  $\Pi(a|O)$  learns only a specific choice of couplings. It is possible however that a cleverer design of the learning process would allow for a direct transfer of the policy, without using the original system as a reference to choose the actions to apply afterward on different ones.

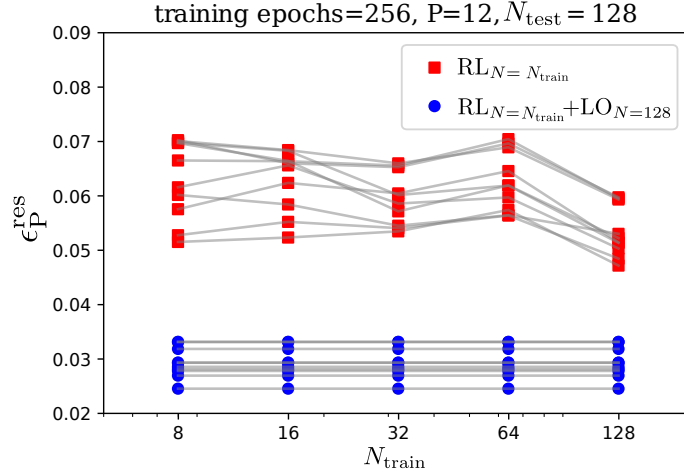


Figure 6.8.: Comparison of the residual energy obtained by training the policy on a single disordered instance with  $N_{tr}$  sites and tested on 10 instances of length  $N_{test} = 128$ . When a local optimization for the test instance is added on top of the RL actions (blue circles), the final residual energy is *independent* from the training instance and system size.

Finally, the action transferability suggests the following way-out from the “measurement problem” involved in the construction of the state observables  $O_t$ . Indeed, in an experimental implementation of RL-assisted QAOA, the RL agent could observe a small system, efficiently simulated on a classical hardware, and then use the learned actions to evolve the larger experimental system. This reduces drastically the number of measurements to be performed and allows to test RL-assisted QAOA on physical quantum platforms.

### 6.4.3. LMG model

Here we consider the Lipkin-Meshkov-Glick (LMG) model, corresponding to the  $p = 2$  p-spin model we analyzed in chapter 5, described by the Hamiltonian

$$\hat{H} = -\frac{1}{N} \left( \sum_i \hat{\sigma}_i^z \right)^2 - h \sum_i \hat{\sigma}_i^x. \quad (6.10)$$

This model provides another useful benchmark for our method because it displays some peculiarities within QAOA framework. It has a very rugged energy landscape [147] which makes local optimizations unstable: it is hard to find good minima when  $P < N/2$  and in a previous work [147] we failed in finding smooth parameter sets. In particular, the iterative optimization used in Refs. [96, 181] does not work.

Here I show that RL makes the local optimization more stable than QAOA alone. We consider a target Hamiltonian corresponding to a nonzero transverse field  $\hat{H}_{target} = \hat{H}_z + h\hat{H}_x$  and focus on reaching the ferromagnetic phase when the system is initially prepared in the paramagnetic state  $|+\rangle$ . In the upper panel of Fig. 6.9 O present the data obtained with the different protocols for a chain of  $N = 64$  spins and target transverse field  $h = \frac{\sqrt{5}-1}{2}$ .




Figure 6.9.: Upper panel: Comparison between the residual energy curves obtained with different protocols: QAOA with random initialization (blue circles), QAOA with linear initialization (green squares), RL (red triangles), RL+LO (black triangles). The data refers to a chain of  $N = 64$  spins. Lower panel: data collapse of the residual energy curves after rescaling  $P \rightarrow P/\log_2 N$ .

While the quality of RL alone deteriorates rapidly when  $P$  increases, even if still better than QAOA with random initializations, it makes the results of a subsequent local optimization much more stable. A data collapse shows that results for different chain length collapse nicely when  $P$  is rescaled with the logarithm of the system size  $N$ . Thus, among the QAOA variational *Ansatz* for the LMG model, there exist a class of minima that allow reaching very small residual energy with an evolution time increasing only *logarithmically* with the system size  $N$ . However these minima are very hard to find with local optimization; indeed only RL-assisted QAOA is able to address them correctly, among all the techniques we tested (random initialization, linear initialization, iterative local optimization).

Another nice feature of RL-assisted QAOA is the (partial) smoothness of the interpolation parameter  $s_t$ , which was absent in the standard QAOA approach. Since the transverse field

of the target Hamiltonian is non-zero,  $s_t$  takes the form

$$s_t = \frac{\gamma_t}{(1-h)\gamma_t + \beta_t}. \quad (6.11)$$

In Fig. 6.10(a) I report the interpolation parameter  $s_t$  (blue lines) for ten different test runs of a RL policy trained on a system of  $N = 64$  spins with  $P = 10$  and target transverse field  $h = \frac{\sqrt{5}-1}{2}$ . Alongside, I plot the  $s_t$  relative to the local optimization (dashed red lines) on top of the actions chosen by the RL agent. Comparing with the results obtained from l-init QAOA in Fig. 6.10(b), the difference in smoothness is striking. Even after the final local optimization, the RL actions are much more regular and the different trajectories clearly suggest the presence of a common basin linked to some continuous schedule  $s(t)$ . An interesting feature of the schedule learned by the algorithm is that it is not the discretization of an annealing protocol that interpolates between  $\hat{H}_{\text{drive}}$  and  $\hat{H}_{\text{target}}$ . Indeed  $s_t$  does not start close to 0 at the beginning of the episode, but as  $t \rightarrow 0$   $s_t$  reaches a finite value close to 0.5, indicating that the associated continuous schedule  $s(t)$  is (likely) not adiabatic with respect to the instantaneous Hamiltonian.

Finally, let us discuss the transferability of the policy in the LMG model. I report in Fig. 6.11 the residual energies I obtained by training the NN on a system of  $N_{\text{tr}}$  spins for 1024 epochs and then testing the policy on a larger set with  $N_{\text{test}} = 128$ , both with and without a local optimization on top (blue circles and red squares, respectively). At difference with the random TFIM model presented in the previous section, now the transferred policy displays a clear change in performance depending on the ratio between the training and the test system sizes. The final local optimization tends to smear out this difference in the residual energies, but clearly the performance is better when  $N_{\text{tr}} = N_{\text{test}}$ , as indicated also by the smaller error-bar.

## 6.5. Conclusions

I have shown that the optimal QAOA strategies well known for the TFIM [96] can be effectively learned with a simple PPO-algorithm [191] employing rather small NNs. The observables measured on a state, referring to the two competing terms in the Hamiltonian and providing information to the “agent”, seem to be effective in the learning process. I have shown that RL learns *smooth* control parameters, hence realizing an RL-assisted feedback Quantum Control for the schedule  $s(t)$  of a digitized QA/AQC algorithm [96], in absence of any spectral information. By working with disordered quantum Ising chains, I showed that strategies “learned” on small samples can be successfully transferred to larger systems, hence alleviating the “measurement problem”: one can learn a strategy on a small problem simulated on a computer, and implement it on a larger experimental setup.

A discussion of previous RL-work on quantum systems is here appropriate. RL as a tool for quantum control and quantum-error-correction has been investigated in Refs. [185, 186]. Regarding applications to QAOA, Refs. [187, 188, 190] have all formulated RL strategies to learn optimal variational parameters  $(\gamma, \beta)$ . While sharing similar RL tools, their approach is markedly different from ours: they identify the RL “state” with the whole set of QAOA parameters. The agent has no access to the internal quantum state, and no information

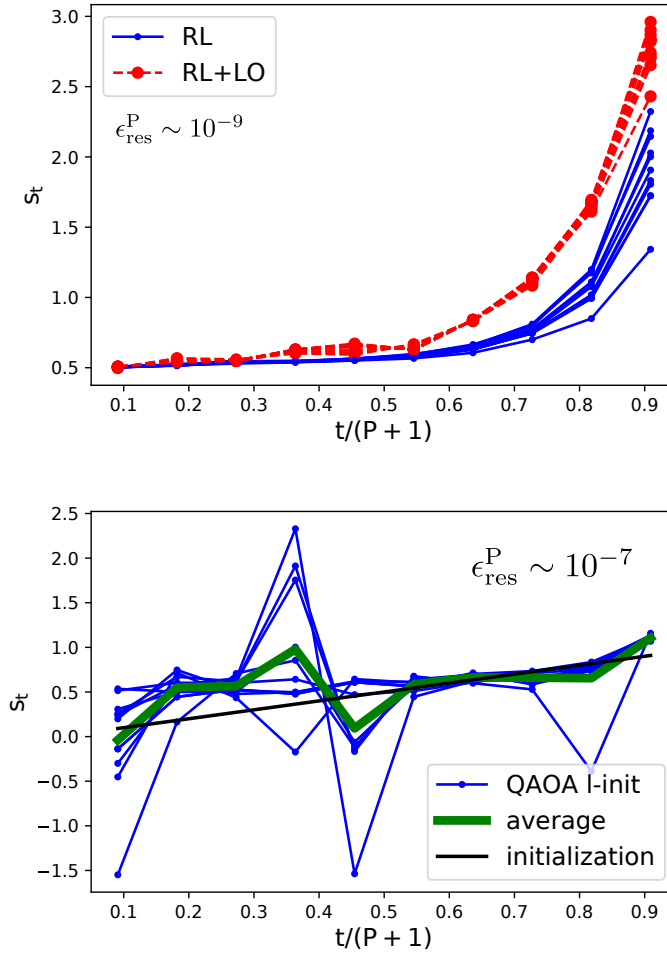


Figure 6.10.: (a) Learned actions after 1024 training epochs on a chain of  $N = 64$  (blue lines), and their local optimization (red dashed lines). (b) optimal parameter sets from QAOA with local optimization (blue lines), their average (thick green line), and the linear *Ansatz* (black line). In both panels  $P = 10$  and the target transverse field is  $h = \frac{\sqrt{5}-1}{2}$ . In both panels, we indicate the typical residual energy found for those parameter sets.




Figure 6.11.: Residual energy obtained by training the policy on a system of  $N_{tr}$  spins and then applying the actions on a test system of  $N_{test} = 128$  spins. Red squares refer to the direct transfer of the RL actions, while blue circles have a local optimization on top. The inset shows the same data in log scale, where it is possible to appreciate the difference in the average performance depending on the value of  $N_{tr}$ .

on the evolution process can be exploited in the optimization. In this way, the issue of measuring the intermediate quantum state is bypassed. This choice, however, reduces RL to a heuristic optimization which forfeits one of the most relevant features of the RL framework: The possibility to drive the process with a step-by-step evolution. An alternative proposal, closer to ours in methods but tackling different physical questions, has recently appeared in Ref. [189].

Concerning future developments, we mention possible improvements to the “measurement problem”. One possibility is to introduce ancillary bits to provide intermediate information to the RL agent without destroying the state of the system, in a way similar to Ref. [186]. A possible alternative is to perform weak measurements [196]. A second issue is the sensitivity to noise: preliminary results show that noise in the initial state preparation does not harm the ability to learn the correct strategies. Finally, the application to other models is worth pursuing, to better investigate the universality of our approach.

## 7. Conclusions

In the present Thesis I have discussed how adiabatic dynamics and related processes generate interesting non-equilibrium physical phenomena and provide useful tools for developing quantum technologies. Two are the aspects that are analyzed: topological phases arising in periodically driven systems and optimization techniques rooted in adiabatic quantum computation.

Periodically driven systems are the focus of part I, where I investigate the robustness of quantum pumping with respect to finite frequency effects and substrate disorder. Floquet engineering and non-equilibrium phases are indeed the objects of flourishing research activities, and understanding the response of such systems to possible perturbations is necessary to properly interpret experimental data.

As a first contribution, in chapter 3 I extend the results presented in Ref. [75], showing that the relative time scales of the periodic driving and the ramp that turns on the driving itself determine whether the system displays a perturbative or a non-analytic response to the finite frequency. A detailed analysis of the quasi-energy spectrum at small frequency allows connecting the topological properties of the instantaneous ground state of the time-period Hamiltonian to those of the Floquet states relative to the actual frequency of the driving. A topologically robust response of the form  $e^{-\alpha\omega}$  is obtained only if the Floquet state with the same Chern number as the ground state is prepared through a continuous deformation of the latter, accelerating the driving up to the final frequency. The acceleration, however, needs not to be too small, because this process must be *adiabatic* with respect to the small gap due to resonances in the Floquet spectrum. This leads to a crossover between a non-analytic and a perturbative response depending on the ratio between the final driving frequency and the switching-on time.

As a second contribution, in chapter 4 I demonstrated that the robustness of quantized transport to static on-site disorder, which also breaks the symmetry protecting the topological phase, is due to a remarkable delocalization/localization transition of the Floquet states. The pumped charge averaged over an infinite number of driving periods  $Q_d$  displays a smooth crossover between a topological plateau, where  $Q_d$  is independent of the disorder strength  $W$ , and a trivial regime where disorder completely suppresses transport. The two regimes are connected by a region where pumping is not quantized and the precise value of  $Q_d$  strongly depends on the actual disorder instance. Hence no sharp transition is observed. However, the inverse participation ratio (IPR) distribution of the single-particle Floquet states shows a marked peak corresponding to extended states, that disappear when transport is suppressed. An analysis of the localization length and spectral properties associated with this peak suggests the presence of an actual delocalization/localization transition of Floquet states driven by the disorder, which is a remarkable fact in a one-dimensional system with an

Anderson localized energy spectrum. This transition is linked to the topological nature of the driving in the clean system and disappears in the trivial phase, where no particle transport occurs in the insulating ground state.

A key future development will be understanding the interplay between interaction, disorder, and topology. On the one hand, it is known that Thouless pumping can only survive as a metastable state in clean interacting Floquet systems [135], before reaching a featureless infinite temperature state. On the other hand, interaction and disorder can induce many-body localization, where a periodically driven state can be stabilized even at low frequencies [24]. Quantized transport, however, requires extended Floquet states, in contrast with the localized energy spectrum. The key question, therefore, is whether a topological periodic driving in an interacting disordered system can induce extended Floquet states, while at the same time the localized Hamiltonian eigenstates prevent the system to absorb energy from the driving. Another interesting perspective is understanding if disorder makes quantum pumping more robust to some of the detail of its implementation, such as the chemical potential  $\mu$ . Indeed, this is well known in the integer quantum Hall effect (IQHE), where disorder creates localized states in the energy gap between different Landau levels (LL) and makes the quantization of the conductivity possible as long as  $\mu$  lies within the energy of localized states. Although quantum pumping is a dimensional reduction of the IQHE, in the clean limit, when disorder is present this parallelism is not so clear anymore. When mapping a one-dimensional driven system to a two-dimensional one [197], one obtains a disordered pattern organized in stripes, hence preserving translational invariance in the “time” direction, which is very different to actual uniform disorder in 2D.

For what regards quantum optimization, I investigated in chapter 5 the performance of the Quantum Approximate Optimization Algorithm on the fully-connected p-spin model, where QAOA is able to prepare the ground state with polynomial resources. This is in sharp contrast with continuous adiabatic evolution, where the total time required scales exponentially with the system size  $N$ , due to the presence of a first order phase transition when  $p \geq 3$ . However, the presence of many local minima poses severe difficulties in finding a good approximation to the ground state for small values of  $P$ , where the target state can not be addressed exactly. Our results suggest that infinite-range Hamiltonians can be useful tools to allow collective transformations of the variational wavefunction and thus explore more efficiently the energy landscape. This might be used in QAOA optimizations for generic spin models, adding extra generators to the unitary transformations, beside the target and the driving Hamiltonians. However, this approach does not guarantee the convergence of QAOA to the true ground state in the  $P \rightarrow \infty$  limit, where it should be able to recover an analogic adiabatic evolution. With more kinds of unitary operators, the variational Ansatz describes a larger class of functions, but convergence to the correct minimum might become extremely hard. Further investigations on this topic are due to uncovering the full power of such hybrid variational optimization techniques,

Finally, I proposed in chapter 6 a reinforcement learning (RL) assisted QAOA, where an RL agent learns good approximation for optimal smooth parameter schedules for the transverse field Ising model (TFIM). Indeed, without any constraint on the action choices, the RL agent learns a policy which lies in the same basin of attraction of the optimal digitized-QA schedule found with iterative optimization in Ref. [96]. Hence when a local optimization is added on top



---

of the RL actions, the optimal schedule is recovered. The most interesting feature is that the learned policy for the TFIM with random couplings can be easily transferred among systems with different sizes and different disordered instances. After the final local optimization step, the residual energy is independent of the particular system where the agent has been trained, provided the number of learning epochs is large enough to bring the policy close to the correct minimum. Our method works well also on the p-spin ( $p = 2$ ) model, making the local optimization much more stable with respect to standard QAOA, both regarding the residual energy and the optimal parameters.

Several open questions remain: First of all the measurement problem, i.e. the fact that RL requires constantly to monitor the state of the system, in order to choose the proper action pair. Besides exploiting the transferability of the policy, an amenable solution would be reducing or changing the kind of information given to the agent. For instance, one could try to provide single-shot outcomes, instead of the average of operators, which are much less expensive in experimental realization. Another possibility is to couple the system to auxiliary qubits and performing projective or weak measurements only on those [186]. It would introduce some decoherence in the dynamics, but for optimization purposes it might not be necessarily a drawback, as some degree of dissipation is exploited in many quantum annealing processes [198–200]. Also, a detailed study of the learning process itself might be interesting, to characterize the collapse of the learned policies to the optimal strategy after a local optimization. This indeed might help to improve the method, for instance finding the minimum number of training epochs required to reach the basin of attraction of an optimal smooth minimum, depending on the size of the system and the number of Trotter steps  $P$ . Finally, a fundamental issue is the universality of the method, and primarily of the transferability of the policy. Future works will for sure deal with more complicated, non-integrable models, such as the XXZ chain, more complex Max-Cut problems, such as those studied in Ref. [181], or frustrated spin-glasses, as the Sharrington-Kirkpatrick model studied in Ref. [193].



# Appendix



# A. Details for chapter 3

## A.1. Adiabatic expansion of Floquet eigenstates

In this appendix we derive an expansion in powers of  $1/\tau$  for the Floquet modes, looking in particular for their overlap with the Hamiltonian eigenstates. To obtain this expansion, we first exploit adiabatic perturbation theory [115] (APT) to compute the Floquet operator and then we use ordinary perturbation theory (PT) to calculate the corrections to the Floquet modes.

Given a gapped periodic Hamiltonian  $\hat{H}(t + \tau) = \hat{H}(t)$  with  $\hat{H}(t)|\phi_\nu(t)\rangle = E_\nu(t)|\phi_\nu(t)\rangle$  denoting instantaneous eigenstates/eigenvalues, if the evolution is slow enough, the adiabatic theorem states that we can write the *time evolved* state  $|\psi_\nu(\tau)\rangle$  originating from  $|\psi_\nu(t = 0)\rangle \equiv |\phi_\nu(0)\rangle$  up to the 0-th order in  $1/\tau$ , as:

$$|\psi_\nu(\tau)\rangle \approx |\psi_\nu^{(0)}(\tau)\rangle = e^{-i\epsilon_\nu\tau/\hbar}|\phi_\nu(0)\rangle, \quad (\text{A.1})$$

where

$$\epsilon_\nu\tau = \int_0^\tau dt (E_\nu(t) - i\hbar\langle\phi_\nu(t)|\partial_t\phi_\nu(t)\rangle).$$

Hence, writing the Floquet operator as

$$\hat{F}(\tau) \equiv \hat{U}(\tau, 0) = \sum_\nu |\psi_\nu(\tau)\rangle\langle\phi_\nu(0)|, \quad (\text{A.2})$$

the adiabatic theorem tells us that, when  $\omega = \frac{2\pi}{\tau} \rightarrow 0$ , it reduces to the expression

$$\hat{F}^{(0)}(\tau) = \sum_\nu e^{-i\epsilon_\nu\tau/\hbar}|\phi_\nu(0)\rangle\langle\phi_\nu(0)|. \quad (\text{A.3})$$

Notice that  $|\phi_\nu(\tau)\rangle = |\phi_\nu(0)\rangle$  due to the time periodicity of the Hamiltonian. Eq. (A.3) means that in this limit the Floquet modes  $|\psi_\nu(0)\rangle$  and the instantaneous eigenstates  $|\phi_\nu(0)\rangle$  at the end of each period coincide. To obtain finite frequency corrections, we need to write  $|\psi_\nu(\tau)\rangle$  through an adiabatic perturbation series [115]

$$|\psi_\nu(\tau)\rangle = \sum_{p=0}^{\infty} \left(\frac{\hbar}{\tau}\right)^p |\psi^{(p)}(\tau)\rangle, \quad (\text{A.4})$$

which leads to a similar expression for the Floquet operator

$$\hat{F}(\tau) = \sum_{p=0}^{\infty} \left(\frac{\hbar}{\tau}\right)^p \hat{F}^{(p)}(\tau). \quad (\text{A.5})$$

Now we assume that at  $t = 0$  the system is prepared in the  $\nu$ -th eigenstate  $|\phi_\nu(0)\rangle$  of the Hamiltonian. From Eq. (A.3) we expect a single Floquet state to have a large overlap with  $|\phi_\nu(0)\rangle$  if the period  $\tau$  is large, and therefore its occupation number  $n_\nu$  should be close to one. We wish to exploit perturbation theory to compute the lowest order corrections in  $\frac{1}{\tau}$  to

$$n_\nu = |\langle \phi_\nu(0) | u_\nu(0) \rangle|^2 = |\langle \phi_\nu(\tau) | u_\nu(\tau) \rangle|^2 . \quad (\text{A.6})$$

As we will show in the following, the lowest order terms are quadratic in  $\frac{1}{\tau}$ —or equivalently in  $\omega$ — and they originate from second-order corrections due to  $\hat{F}^{(1)}(\tau)$ . Indeed first-order corrections to a given eigenstate in perturbation theory are always orthogonal to the unperturbed one and therefore  $\hat{F}^{(2)}(\tau)$  can only give contribution of order  $O(\frac{1}{\tau^4})$ . Thus we just need to compute  $\hat{F}^{(1)}(\tau)$ . Before proceeding we define the following quantities depending on a rescaled time  $s = t/\tau$

$$\begin{aligned} \tau\epsilon_\nu &= \tau \int_0^1 ds E_\nu(s) - i\hbar \int_0^1 ds \langle \phi_\nu(s) | \partial_s \phi_\nu(s) \rangle , \\ \Delta_{\nu,\mu}(s) &= E_\nu(s) - E_\mu(s) , \end{aligned}$$

where  $\{E_\nu(s)\}$  is the set of instantaneous eigenvalues of  $\hat{H}(s)$  and  $\{|\phi_\nu(s)\rangle\}$  the corresponding eigenvectors. The adiabatic expansion of the time evolved state  $\psi_\nu(t)$  will be written in terms of

$$\begin{aligned} M_{\mu,\nu}(s) &= \frac{\langle \phi_\mu(s) | \partial_s \hat{H}(s) | \phi_\nu(s) \rangle}{\Delta_{\nu,\mu}} , \\ J_{\mu,\nu} &= \int_0^1 ds \frac{|M_{\mu,\nu}(s)|^2}{\Delta_{\mu,\nu}(s)} , \end{aligned}$$

which again depend only on the instantaneous spectrum of the Hamiltonian. Following Ref. [115] we can write the first-order correction to the evolved eigenstate  $|\psi_\nu(\tau)\rangle$  as:

$$\begin{aligned} |\psi_\nu^{(1)}\rangle &= i \sum_{\mu \neq \nu} e^{-i\tau\epsilon_\nu/\hbar} J_{\mu,\nu} |\phi_\nu\rangle + \\ & i \sum_{\mu \neq \nu} \frac{M_{\mu,\nu}}{\Delta_{\mu,\nu}} \left( e^{-i\tau\epsilon_\nu/\hbar} - e^{-i\tau\epsilon_\mu/\hbar} \right) |\phi_\mu\rangle , \end{aligned} \quad (\text{A.7})$$

where the  $s$  dependence is omitted since all quantities are computed in  $s = 1$ . The ‘‘perturbation’’ of order  $1/\tau$  to the Floquet operator consists of a diagonal part (first term of the RHS in Eq. (A.7)) and an off-diagonal part (second term). The former acts only as a renormalization of the eigenvalues (the Floquet quasi-energies) of the operator but does not change the eigenvector, since it is diagonal in the original basis. We can now apply perturbation theory for linear operators to obtain the correction to the Floquet modes. The first-order term reads

$$\begin{aligned} |u_\nu^{(1)}\rangle &= \frac{\hbar}{\tau} \sum_{\mu \neq \nu} |\phi_\mu\rangle \frac{\langle \phi_\mu | \hat{F}^{(1)}(\tau) | \phi_\nu \rangle}{e^{-i\tau\epsilon_\nu/\hbar} - e^{-i\tau\epsilon_\mu/\hbar}} \\ &= i \frac{\hbar}{\tau} \sum_{\mu \neq \nu} \frac{M_{\mu,\nu}}{\Delta_{\mu,\nu}} |\phi_\mu\rangle , \end{aligned} \quad (\text{A.8})$$

where the off diagonal elements of  $F^{(1)}(\tau)$  are obtained by combining Eq. (A.2) and Eq. (A.7), leading to

$$F_{\mu,\nu}^{(1)}(\tau) = i \frac{M_{\mu,\nu}}{\Delta_{\mu,\nu}} \left( e^{-i\tau\epsilon_\nu/\hbar} - e^{-i\tau\epsilon_\mu/\hbar} \right) |\phi_\mu\rangle \langle \phi_\nu| . \quad (\text{A.9})$$

Since we are interested in computing the projection  $\langle \phi_\nu | u_\nu \rangle$ , only the terms proportional to  $|\phi_\nu\rangle$  are needed. Clearly Eq. (A.8) gives no contribution — all terms are orthogonal to  $|\phi_\nu\rangle$ — but it can be used to obtain the next order by imposing the normalization condition  $\langle u_\nu | u_\nu \rangle = 1$

$$|u_\nu^{(2)}\rangle = -\frac{\hbar^2}{2\tau^2} |\phi_\nu\rangle \sum_{\mu \neq \nu} \left| \frac{M_{\mu,\nu}}{\Delta_{\mu,\nu}} \right|^2 + \text{terms orthogonal to } |\phi_\nu\rangle . \quad (\text{A.10})$$

Hence the occupation at finite frequency of the targeted Floquet mode reads

$$\begin{aligned} n_\nu &= \left| 1 - \frac{\hbar^2}{2\tau^2} \sum_{\mu \neq \nu} \left| \frac{M_{\mu,\nu}}{\Delta_{\mu,\nu}} \right|^2 \right|^2 + o(1/\tau^2) \\ &= 1 - \frac{\hbar^2}{\tau^2} \sum_{\mu \neq \nu} \left| \frac{M_{\mu,\nu}}{\Delta_{\mu,\nu}} \right|^2 + o(1/\tau^2) . \end{aligned} \quad (\text{A.11})$$

Therefore, if the matrix elements  $M_{\mu,\nu}$  are not all equal to zero, we expect to see power law correction to the occupation number of Floquet modes, when the system is prepared in the  $\nu$ -th state  $|\phi_\nu(0)\rangle$  of  $\hat{H}(t=0)$ .

## A.2. Dependence on $k_x$ of Floquet quasi-energies and occupations

Here we discuss the dependency of Floquet modes and quasi-energies from  $k_x$  and how the system can be effectively described in only 1+1 dimensions (space + time). The starting point is the block diagonal Hamiltonian in momentum space, which reads

$$\begin{aligned} \hat{H}_{\mathbf{k}}(t) = J_0 \sum_{b=0}^{q-1} \left\{ 2 \cos \left( ak_y + \frac{2\pi p}{q} b \right) \hat{c}_{\mathbf{k},b}^\dagger \hat{c}_{\mathbf{k},b} + \right. \\ \left. \left[ e^{-ia(k_x + \kappa_x(t))} \hat{c}_{\mathbf{k},b+1}^\dagger \hat{c}_{\mathbf{k},b} + \text{H.c.} \right] \right\} , \end{aligned} \quad (\text{A.12})$$

where  $\hat{c}_{\mathbf{k},q} = \hat{c}_{\mathbf{k},0}$  and  $a\kappa_x(t) = \omega(t - t_0)$  when the force field  $F = \frac{\hbar\omega}{a}$  is stationary.  $t_0$  is the initial time for which the system is prepared with a non periodic driving. Notice that  $\hat{H}_{\mathbf{k}}(t)$  depends on  $k_x$  and time only through the phase  $ak_x + \omega(t - t_0)$ . Hence we can define  $t_x = t_0 - k_x/\omega$ , so that the evolution operator over one period (the Floquet operator) for a given  $t_0$  can be written as

$$\hat{F}_{\mathbf{k}}(\tau) = \hat{U}_{\mathbf{k}}(\tau + t_0, t_0) = \hat{U}_{k_y}(\tau + t_x, t_x) , \quad (\text{A.13})$$

Here a subscript  $k_y$  indicates that the associated quantity is evaluated in  $\mathbf{k} = (0, k_y)$ . By applying the composition property of evolution operator and exploiting Floquet theorem in the form

$$\hat{U}(t_0 + t + \tau, t_0) = \hat{U}(t + t_0, t_0) \hat{U}(\tau + t_0, t_0) ,$$

one obtains

$$\hat{F}_{\mathbf{k}}(\tau) = \hat{U}_{k_y}(t_x, t_0) \hat{F}_{k_y}(\tau) \hat{U}_{k_y}^\dagger(t_x, t_0) , \quad (\text{A.14})$$

which can be written explicitly as

$$\hat{F}_{\mathbf{k}}(\tau) = \sum_{\nu} e^{-i\epsilon_{k_y,\nu}\tau/\hbar} |u_{k_y,\nu}(t_x)\rangle \langle u_{k_y,\nu}(t_x)|, \quad (\text{A.15})$$

since the phase factors arising from the action of  $\hat{U}_{k_y}(t_x, t_0)$  and  $\hat{U}_{k_y}^{\dagger}(t_x, t_0)$  exactly cancel each other. Therefore the Floquet modes shifted along  $k_x$  are

$$\begin{aligned} |u_{\mathbf{k},\nu}(t_0)\rangle &= |u_{k_y,\nu}(t_0 - \frac{ak_x}{\omega})\rangle \\ &= e^{i\epsilon_{k_y,\nu}(t_x-t_0)/\hbar} \hat{U}_{k_y}(t_x, t_0) |u_{k_y,\nu}(t_0)\rangle, \end{aligned} \quad (\text{A.16})$$

i.e. the periodic part of the  $\nu$ -th Floquet state in  $\mathbf{k} = (0, k_x)$  evolved for a time  $t_x - t_0 < \tau$ . Thus the Floquet operator at any point in the  $k$ -space with  $k_x \neq 0$  can be obtained by a *unitary* transformation applied on  $\hat{F}_{k_y}$ . The most important implication is that the quasi-energies  $\epsilon_{\mathbf{k},\nu} = \epsilon_{k_y,\nu}$  are independent from  $k_x$ . The Floquet modes instead still depend on  $k_x$ , because of Eq. (A.16). Hence when computing the infinite time average pumped charge

$$Q_{\mathbf{d}} = \frac{\tau}{\hbar} \sum_{\nu} \int_{\text{BZ}} \frac{d^2\mathbf{k}}{(2\pi)^2} n_{\mathbf{k},\nu} \frac{\partial \epsilon_{k_y,\nu}}{\partial k_y}, \quad (\text{A.17})$$

the only remaining dependence on  $k_x$  is in the occupation number  $n_{\mathbf{k},\nu} = |\langle \psi_{\mathbf{k}} | u_{k_y,\nu}(t_x) \rangle|^2$ .

Another interesting property of the Hamiltonian as written in Eq. (A.12), is that the spectrum is invariant for a discrete shift of the momentum in the  $\hat{\mathbf{y}}$  direction  $k_y \rightarrow k_y + \frac{2\pi p}{qa}$ . Indeed this transformation is equivalent to a shift of  $a$  in real space of the magnetic unit cell, leading to a simple relabelling of the internal index  $b \rightarrow b + 1$ . For the case investigated in chapter 3 ( $p = 1, q = 3$ ), this property is clearly shown in Fig. 3.4, where the invariance for  $k_y \rightarrow k_y + \frac{2\pi}{3a}$  is evident. This symmetry in the Hamiltonian is inherited also by the quasi-energy spectrum, which is also repeated three times inside the Brillouin zone. This symmetry is nothing else than gauge invariance: the spectrum must depend on the same way by  $k_x$  and  $k_y$ , because the breaking of translational invariance along the  $\hat{\mathbf{x}}$  direction is only due to the gauge choice, which can not influence any observable. If we had chosen  $\mathbf{A} = -By\hat{\mathbf{x}}$ , the magnetic unit cell would have consisted of three sites along the  $\hat{\mathbf{y}}$  direction and thus the first Brillouin zone would have been  $[0, \frac{2\pi}{a}] \times [0, \frac{2\pi}{3a}]$ , leading to a periodicity of  $\frac{2\pi}{3a}$  in  $k_y$ .



## B. Details for chapter 5

In this appendix we will show that one can get the exact target ground state of the p-spin model with a single QAOA step,  $P = 1$ , starting from the fully x-polarized state  $|+\rangle = \frac{1}{\sqrt{2^N}} \otimes_{j=1}^N (|\uparrow\rangle_j + |\downarrow\rangle_j)$ , provided the number of sites  $N$  is *odd*. This holds true for all possible values of  $p$ , and generalizes the result of Ref. [166] to  $p > 2$ . We will then analyze some general symmetry properties of the QAOA landscape for the problem for general  $P$ .

### B.1. exact ground state preparation for $P = 1$

#### B.1.1. $P = 1$ : requirements on $\beta$

For  $P = 1$  the QAOA state has only two parameters, which we will denote by  $\gamma$  and  $\beta$ , without an index. Let  $|\psi_{\text{targ}}\rangle$  denote the (target) ground state of the model, and define the fidelity:

$$\mathcal{F}(\gamma, \beta) = \left| \langle \psi_{\text{targ}} | \psi_{P=1}(\gamma, \beta) \rangle \right|^2 = \left| \langle \psi_{\text{targ}} | e^{-i\beta \hat{H}_x} e^{-i\gamma \hat{H}_z} |+\rangle \right|^2 = \left| \frac{1}{\sqrt{2^N}} \sum_l e^{-i\gamma E_l} \langle \psi_{\text{targ}} | e^{-i\beta \hat{H}_x} |l\rangle \right|^2, \quad (\text{B.1})$$

where we have expanded the initial state  $|+\rangle = \frac{1}{\sqrt{2^N}} \sum_l |l\rangle$  as an equal superposition of all possible  $2^N$  classical  $\mathbf{z}$ -basis configurations  $|l\rangle$ , and we used that  $\hat{H}_z |l\rangle = E_l |l\rangle$ , where  $E_l$  is the energy of the configuration  $|l\rangle$ . Let us now define the following two  $2^N$  dimensional complex vectors:

$$(\mathbf{v}(\gamma))_l = \frac{1}{\sqrt{2^N}} e^{i\gamma E_l} \quad \text{and} \quad (\mathbf{u}(\beta))_l = \langle \psi_{\text{targ}} | e^{-i\beta \hat{H}_x} |l\rangle. \quad (\text{B.2})$$

Simple algebra shows that they have unit norm,  $\|\mathbf{v}(\gamma)\| = 1$  and  $\|\mathbf{u}(\beta)\| = 1$ , and that the fidelity can be expressed as a scalar product of them:  $\mathcal{F}(\gamma, \beta) = |\mathbf{v}^\dagger(\gamma) \cdot \mathbf{u}(\beta)|^2$ . Hence, by the Cauchy-Schwarz inequality:

$$1 = \mathcal{F}(\gamma, \beta) = |\mathbf{v}^\dagger(\gamma) \cdot \mathbf{u}(\beta)|^2 \iff \exists \theta \in \mathbb{R} \quad \text{such that} \quad \mathbf{u}(\beta) = e^{i\theta} \mathbf{v}(\gamma), \quad (\text{B.3})$$

i.e., the two vectors coincide, up to an overall phase factor. Since  $|(\mathbf{v}(\gamma))_l|^2 = \frac{1}{2^N}$ , this in turn implies that we must have

$$\left| \langle \psi_{\text{targ}} | e^{-i\beta \hat{H}_x} |l\rangle \right|^2 = |(\mathbf{u}(\beta))_l|^2 = \frac{1}{2^N} \quad \forall l. \quad (\text{B.4})$$

So far, our arguments have been rather general. We now specialize in our discussion to the case where  $|\psi_{\text{targ}}\rangle$  is the ground state of the classical p-spin ferromagnet.

For  $p$  *odd*, we have  $|\psi_{\text{targ}}\rangle = |\uparrow \cdots \uparrow\rangle$ , and a simple calculation shows that:

$$\langle \psi_{\text{targ}} | e^{-i\beta \hat{H}_x} | l \rangle = \prod_{j=1}^N \langle \uparrow | \cos \beta \hat{1}_j + i \sin \beta \hat{\sigma}_j^x | l_j \rangle = (\cos \beta)^{N_l^\uparrow} (i \sin \beta)^{N_l^\downarrow}, \quad (\text{B.5})$$

where  $N_l^\uparrow$  and  $N_l^\downarrow$  denote the number of  $\uparrow$  and  $\downarrow$  spins in the configuration  $|l\rangle$ . Hence the requirement given by Eq. (B.4) is satisfied only if

$$\cos^2 \beta = \sin^2 \beta = \frac{1}{2} \quad \implies \beta = \frac{\pi}{4}, \frac{3\pi}{4}, \frac{5\pi}{4}, \frac{7\pi}{4}. \quad (\text{B.6})$$

Similar arguments have been used, see Ref. [168], for the more general case in which  $|\psi_{\text{targ}}\rangle$  is the classical ground state of a generic spin-glass Hamiltonian.

For  $p$  *even* the calculation is slightly more involved, since the target state is now a non-classical superposition of the two degenerate ferromagnetic states

$$|\psi_{\text{targ}}\rangle = \frac{1}{\sqrt{2}} (|\uparrow \cdots \uparrow\rangle + |\downarrow \cdots \downarrow\rangle). \quad (\text{B.7})$$

Hence:

$$\left| \langle \psi_{\text{targ}} | e^{-i\beta \hat{H}_x} | l \rangle \right|^2 = \frac{1}{2} \left| (\cos \beta)^{N_l^\uparrow} (i \sin \beta)^{N_l^\downarrow} + (\cos \beta)^{N_l^\downarrow} (i \sin \beta)^{N_l^\uparrow} \right|^2. \quad (\text{B.8})$$

Once again, one easily verifies that  $\beta = \frac{\pi}{4}$  satisfies the requirement (B.4), *provided N is odd*, so that  $N_l^\uparrow$  and  $N_l^\downarrow$  have opposite parity and therefore  $|i^{N_l^\uparrow} + i^{N_l^\downarrow}|^2 = 2$ .

From now on we will therefore restrict our choice of  $\beta$  to  $\beta = \frac{\pi}{4}$ , a necessary condition for unit fidelity, and study the conditions that  $\gamma$  has to verify. Essentially, the value of  $\gamma$  will have to be chosen in such a way that the various phase factors interfere constructively in a way that is independent of  $l$ . To this goal, we notice that the energy  $E_l$  of the configuration  $|l\rangle$  can be expressed as:

$$E_l = -\langle l | \left( \sum_j \hat{\sigma}_j^z \right)^P | l \rangle = -(N_l^\uparrow - N_l^\downarrow)^P = -M_l^P, \quad (\text{B.9})$$

where  $M_l = N_l^\uparrow - N_l^\downarrow$  is the total magnetization of the configuration.

### B.1.2. $P = 1$ and $p$ *odd*: requirements on $\gamma$

Here we prove that for odd values of  $p$  and  $N$ , the  $p$ -spin QAOA circuit of depth  $P = 1$  and parameters  $(\gamma = \frac{\pi}{4}, \beta = \frac{\pi}{4})$  is sufficient to prepare the ferromagnetic target state  $|\psi_{\text{targ}}\rangle = |\uparrow \cdots \uparrow\rangle$ . Substituting  $\beta = \frac{\pi}{4}$  and  $i = e^{i\frac{\pi}{2}}$  in Eq. (B.5), and using  $E_l = -M_l^P$  in Eq. (B.1) we get:

$$\mathcal{F}(\gamma, \frac{\pi}{4}) = \left| \frac{1}{2^N} \sum_l e^{i\gamma M_l^P} e^{i\frac{\pi}{2} N_l^\downarrow} \right|^2. \quad (\text{B.10})$$

Taking  $\gamma = \frac{\pi}{4}$  gives

$$\begin{aligned}
\mathcal{F}\left(\frac{\pi}{4}, \frac{\pi}{4}\right) &= \left| \frac{1}{2^N} \sum_l \exp\left(i\frac{2\pi}{8} \left[ \left(M_l^p + 2N_l^\downarrow\right) \bmod 8 \right] \right) \right|^2 \\
&= \left| \frac{1}{2^N} \sum_l \exp\left(i\frac{2\pi}{8} \left[ \left(M_l + 2N_l^\downarrow\right) \bmod 8 \right] \right) \right|^2 \\
&= \left| \frac{1}{2^N} \sum_l \exp\left(i\frac{2\pi}{8} \left[ N \bmod 8 \right] \right) \right|^2 = 1, \tag{B.11}
\end{aligned}$$

where we have used the fact that for  $N$  odd,  $M_l = N - 2N_l^\downarrow$  is also odd and the following property of arithmetic congruences holds:

$$M_l^{p-1} = 1 \pmod{8} \implies M_l^p = M_l \pmod{8} \quad \text{if } p \text{ is odd.} \tag{B.12}$$

Eq. (B.11) proves our initial claim, that the QAOA protocol ( $\gamma = \frac{\pi}{4}, \beta = \frac{\pi}{4}$ ) prepares the target ground state of  $\widehat{H}_z$  for the  $p$ -spin ferromagnet with unit fidelity, provided  $N$  and  $p$  are both odd.

### B.1.3. $P = 1$ and $p$ even: requirements on $\gamma$

For even values of  $p$ , the system is  $\mathbb{Z}_2$  symmetric. The  $p$ -spin QAOA circuit preserves such symmetry. Therefore, the targeted ground state of  $\widehat{H}_z$  is  $|\psi_{\text{target}}\rangle$  in Eq. (B.7). As we did in the previous section, we compute the fidelity between the output  $|\psi_{P=1}(\gamma, \beta = \frac{\pi}{4})\rangle$  of the QAOA circuit and the (non-classical) target state  $|\psi_{\text{target}}\rangle$  in Eq. (B.7):

$$\mathcal{F}(\gamma, \beta = \frac{\pi}{4}) = \left| \frac{1}{2^N} \sum_l e^{i\gamma M_l^p} \left( \frac{e^{i\frac{\pi}{2}N_l^\downarrow} + e^{i\frac{\pi}{2}N_l^\uparrow}}{\sqrt{2}} \right) \right|^2. \tag{B.13}$$

We observe that for  $N = N_l^\uparrow + N_l^\downarrow$  odd,  $N_l^\uparrow$  and  $N_l^\downarrow$  must have *opposite parity*, and the term inside parenthesis is a pure phase factor, which can be expressed as:

$$\left( \frac{e^{i\frac{\pi}{2}N_l^\downarrow} + e^{i\frac{\pi}{2}N_l^\uparrow}}{\sqrt{2}} \right) = e^{i\frac{\pi}{4}N} e^{-i\pi f(M_l)}, \tag{B.14}$$

where

$$f(M) = \begin{cases} 0 & \text{for } M \bmod 8 = \pm 1 \\ 1 & \text{for } M \bmod 8 = \pm 3 \end{cases}. \tag{B.15}$$

Hence, omitting the irrelevant  $l$ -independent common factor  $e^{i\frac{\pi}{4}N}$ , we can rewrite the fidelity as:

$$\mathcal{F}(\gamma, \frac{\pi}{4}) = \left| \frac{1}{2^N} \sum_l e^{i(\gamma M_l^p - \pi f(M_l))} \right|^2. \tag{B.16}$$

The arithmetics to prove that the various phase factors can be made  $l$ -independent for a judicious choice of  $\gamma$  are now slightly more involved, for even  $p$ . By experimenting with this

expression for  $p \leq 10$ , we have come out with the following unconventional parameterization of an even value of  $p$ : for every even  $p$ , two natural numbers  $n$  and  $k$  can be found such that:

$$p = 2^{k+1} + n2^k. \quad (\text{B.17})$$

Correspondingly, given the value of  $k$  in Eq. (B.17), we will set the value of  $\gamma$  to:

$$\gamma_k = \frac{2\pi}{2^{k+4}}. \quad (\text{B.18})$$

The crucial arithmetic identity which we will use — see Sec. B.1.3 for a proof — is the following:

$$m^{2^{k+1}+n2^k} \bmod 2^{k+4} = f(m)2^{k+3} + 1 \quad \forall m \in \mathbb{Z} \text{ with } m \text{ odd}, \quad (\text{B.19})$$

where  $f(m)$  is the function given in Eq. (B.15).

With these definitions, it is immediate to verify that:

$$\begin{aligned} \mathcal{F}(\gamma_k, \frac{\pi}{4}) &= \left| \frac{1}{2^N} \sum_l e^{-i\pi f(M_l)} \exp\left(i \frac{2\pi}{2^{k+4}} (M_l)^{2^{k+1}+n2^k}\right) \right|^2 \\ &= \left| \frac{1}{2^N} \sum_l e^{-i\pi f(M_l)} \exp\left(i \frac{2\pi}{2^{k+4}} \left[ (M_l)^{2^{k+1}+n2^k} \bmod 2^{k+4} \right]\right) \right|^2 \\ &= \left| \frac{1}{2^N} \sum_l e^{-i\pi f(M_l)} \exp\left(i \frac{2\pi}{2^{k+4}} \left( f(M_l) 2^{k+3} + 1 \right)\right) \right|^2 \\ &= \left| \frac{1}{2^N} \sum_l e^{-i\pi f(M_l)} e^{i\pi f(M_l) + i \frac{2\pi}{2^{k+4}}} \right|^2 = 1. \end{aligned} \quad (\text{B.20})$$

**Proof of identity in Eq. (B.19)** For completeness, we also present a proof of the arithmetic identity Eq. (B.19). To prove Eq. (B.19), it is sufficient to show that

$$\forall k \in \mathbb{N}, m \in \mathbb{Z}, m \text{ odd}: \quad (m^{2^{k+1}} - 1) \bmod 2^{k+4} = f(m)2^{k+3}. \quad (\text{B.21})$$

We prove Eq. (B.21) by induction over  $k$ :

- (i) We show that Eq. (B.21) holds for  $k = 0$ .

For  $k = 0$ , a direct computation, for odd  $m$ , gives:

$$\begin{aligned} (m^{2^{0+1}} - 1) \bmod 2^{0+4} &= (m^2 - 1) \bmod 16 \\ &= (m - 1)(m + 1) \bmod 16 \\ &= f(m)2^3 \end{aligned} \quad (\text{B.22})$$

- (ii) We show that if Eq. (B.21) holds for a given  $k \in \mathbb{N}$  and for all odd  $m \in \mathbb{N}$ , then it holds also for  $k + 1$ .

Using Eq. (B.21), we write

$$m^{2^{k+1}} = a_m 2^{k+4} + f(m)2^{k+3} + 1, \quad (\text{B.23})$$

with  $a_m \in \mathbb{Z}$ . Then, we have

$$\begin{aligned} (m^{2^{(k+1)+1}} - 1) &= (m^{2^{k+1}} - 1)(m^{2^{k+1}} + 1) \\ &= (a_m 2^{k+4} + f(m) 2^{k+3})(a_m 2^{k+4} + f(m) 2^{k+3} + 2) \\ &= (a_m 2^{k+5} + f(m) 2^{k+4})(a_m 2^{k+3} + f(m) 2^{k+2} + 1). \end{aligned} \quad (\text{B.24})$$

From this, we derive

$$\begin{aligned} (m^{2^{(k+1)+1}} - 1) \bmod 2^{(k+1)+4} &= f(m) 2^{k+4} (a_m 2^{k+3} + f(m) 2^{k+2} + 1) \bmod 2^{(k+1)+4} \\ &= f(m) 2^{k+4}, \end{aligned} \quad (\text{B.25})$$

where we have used that  $f(m) = 0, 1$  for all odd  $m \in \mathbb{Z}$ . This indeed implies that for all  $k \in \mathbb{N}$ :

$$(m^{2^{k+1}} - 1) \bmod 2^{k+4} = f(m) 2^{k+3} \implies (m^{2^{(k+1)+1}} - 1) \bmod 2^{(k+1)+4} = f(m) 2^{(k+1)+3} \quad (\text{B.27})$$

This concludes the proof by induction of Eq. (B.21).

Incidentally, as an immediate consequence of Eq. (B.21) we get that, for any  $n \in \mathbb{N}$ :

$$m^{2^{k+1}} \bmod 2^{k+4} = f(m) 2^{k+3} + 1 \quad (\text{B.28})$$

$$m^{2^{k+2}} \bmod 2^{k+4} = 1 \quad (\text{B.29})$$

$$m^{2^{k+1}+n2^k} \bmod 2^{k+4} = f(m) 2^{k+3} + 1. \quad (\text{B.30})$$

Notice that Eq. (B.29) also follows from the properties of the multiplicative group of integers modulo  $2^k$  discussed in Refs. [201, 202] (eg.  $(\mathbb{Z}/2^{k+4}\mathbb{Z})^\times \cong C_2 \times C_{2^{k+2}}$ ).

#### B.1.4. Symmetries of the parameter space for general $P$ , $N$ and $p$

We discuss here the symmetries in the parameter space of the function

$$E_P(\boldsymbol{\gamma}, \boldsymbol{\beta}) = \langle \psi_P(\boldsymbol{\gamma}, \boldsymbol{\beta}) | \hat{H}_{\text{target}} | \psi_P(\boldsymbol{\gamma}, \boldsymbol{\beta}) \rangle. \quad (\text{B.31})$$

A first trivial operation that leaves the energy unaltered is the inversion  $(\boldsymbol{\gamma}, \boldsymbol{\beta}) \rightarrow -(\boldsymbol{\gamma}, \boldsymbol{\beta})$ , which corresponds to the complex conjugate of Eq. (B.31). Indeed it is immediate to see that

$$|\psi_P(-\boldsymbol{\gamma}, -\boldsymbol{\beta})\rangle = \prod_{m=1}^P e^{i\beta_m \hat{H}_x} e^{i\gamma_m \hat{H}_z} |\psi_0\rangle = |\psi_P(\boldsymbol{\gamma}, \boldsymbol{\beta})\rangle^*, \quad (\text{B.32})$$

given that  $|\psi_0\rangle = |+\rangle$  is a real wavefunction in the basis of  $\hat{S}_z$ .

The symmetries on the  $\boldsymbol{\beta}$  parameters are shared by all QAOA wavefunctions where quantum fluctuations are induced by a magnetic field transverse to the computational basis. We can write a single evolution operator  $e^{-i\beta_m \hat{H}_x}$  as a set of rotation on each individual spin

$$e^{-i\beta_m \hat{H}_x} = e^{i\beta_m \sum_{j=1}^N \hat{\sigma}_j^x} = \bigotimes_{j=1}^N (\cos \beta_m + \hat{\sigma}_j^x \sin \beta_m). \quad (\text{B.33})$$

$\forall N, p$	$E(-\boldsymbol{\gamma}, -\boldsymbol{\beta}) = E(\boldsymbol{\gamma}, \boldsymbol{\beta})$
$p$ odd	$E(\boldsymbol{\gamma}, \boldsymbol{\beta} + \boldsymbol{\pi}) = E(\boldsymbol{\gamma}, \boldsymbol{\beta})$
$p$ even	$E(\boldsymbol{\gamma}, \boldsymbol{\beta} + \frac{\boldsymbol{\pi}}{2}) = E(\boldsymbol{\gamma}, \boldsymbol{\beta})$
$N$ odd	$E(\boldsymbol{\gamma} + \boldsymbol{\pi}, \boldsymbol{\beta}) = E(\boldsymbol{\gamma}, \boldsymbol{\beta})$
$N$ even	$E(\boldsymbol{\gamma} + \frac{\boldsymbol{\pi}}{2^{p-1}}, \boldsymbol{\beta}) = E(\boldsymbol{\gamma}, \boldsymbol{\beta})$

Table B.1.: Symmetry operations for the QAOA process of the  $p$ -spin model. It is understood that any component of  $\boldsymbol{\gamma}$  or  $\boldsymbol{\beta}$  can be modified.

A shift  $\beta_m \rightarrow \beta_m + \pi$  changes the sign of each term in the product, leading to

$$e^{-i(\beta_m + \pi)\hat{H}_x} = \bigotimes_{j=1}^N (-\cos \beta_m - \hat{\sigma}_j^x \sin \beta_m) = (-1)^N \bigotimes_{j=1}^N (\cos \beta_m + \hat{\sigma}_j^x \sin \beta_m) , \quad (\text{B.34})$$

which is a trivial global phase that does not change the energy in Eq. (B.31). Moreover, if  $p$  is even, the target Hamiltonian is  $\mathbb{Z}_2$  symmetric. Recall that  $\hat{H}_x = -\hat{S}_x$  (twice the total spin), which implies that:

$$e^{i\frac{\pi}{2}\hat{S}_x} \hat{H}_{\text{target}} e^{-i\frac{\pi}{2}\hat{S}_x} = \hat{H}_{\text{target}} , \quad (\text{B.35})$$

because  $e^{-i\frac{\pi}{2}\hat{S}_x}$  is a  $\pi$ -rotation around the  $\mathbf{x}$ -direction which gives a global spin flip  $\hat{\sigma}_j^z \rightarrow -\hat{\sigma}_j^z$ , leading to  $E_P(\boldsymbol{\gamma}, \boldsymbol{\beta} + \frac{\boldsymbol{\pi}}{2}) = E_P(\boldsymbol{\gamma}, \boldsymbol{\beta})$ .

The symmetry for  $\boldsymbol{\gamma}$  is subtler and is model-specific. Notice first that  $\hat{S}_z = \sum_j \hat{\sigma}_j^z$  has integer eigenvalues, even or odd depending on  $N$ , and so does  $\hat{H}_z = -\hat{S}_z^p$ . Following the same notation introduced previously, we write a single QAOA evolution operator as

$$e^{-i\gamma_m \hat{H}_z} = \sum_l e^{i\gamma_m M_l^p} |l\rangle \langle l| . \quad (\text{B.36})$$

If  $N$  is odd the eigenvalues  $M_l^p$  of  $\hat{S}_z^p$  are also odd, and the periodicity of  $\gamma_m$  is  $\pi$ , because

$$(\gamma_m + \pi)M_l^p = \gamma_m M_l^p + \pi \pmod{2\pi} . \quad (\text{B.37})$$

Hence the shift  $\gamma_m \rightarrow \gamma_m + \pi$  introduces a global phase  $e^{-i(\gamma_m + \pi)\hat{H}_z} = -e^{-i\gamma_m \hat{H}_z}$ , which is irrelevant in the expectation value of the energy. If  $N$  is even, the eigenvalues  $M_l^p$  of  $\hat{S}_z^p$  are multiples of  $2^p$ , hence

$$(\gamma_m + \frac{\pi}{2^{p-1}})M_l^p = \gamma_m M_l^p \pmod{2\pi} , \quad (\text{B.38})$$

which means that  $e^{i(\gamma_m + \frac{\pi}{2^{p-1}})\hat{S}_z^p} = e^{i\gamma_m \hat{S}_z^p}$ . In table B.1 we summarize the symmetries we have discussed.

## C. Details for chapter 6

In this appendix, we provide a brief introduction to reinforcement learning mathematical formulation and describe with some detail the algorithm we used in chapter 6. We follow quite closely the introduction to reinforcement learning given in the documentation of the OpenAI Spinningup library.

First, let us give an intuitive picture of what can be called reinforcement learning. In a nutshell, RL studies how an artificial intelligent (AI) agent acts on an environment and how it can learn by trial and error. It gives a formal structure to the idea that rewards and punishments help to repeat or forego a certain behavior in the future. RL methods have been applied successfully to a great variety of problems, from teaching computers to control robots to self-driving vehicles. It has also been proved able to learn to play complex strategy games, most notably Go and Dota, where RL agents outperform human players, or to teach computers to play Atari games from pixel inputs.

### C.1. Introduction to RL: Key concepts and notation

The two main elements of RL are the agent and the environment. The environment is the world where the agent lives and that changes as a consequence of the agent's actions. At every step of the agent-environment interaction, the agent gets partial information on the state of the world through an observation and then decides which action to perform. In the simplest formulation of RL, the state of the environment changes when the agent acts on it, even if in general it could also change on its own.

The agent receives a reward from the environment, a real number that measures how good or bad the present state of the world is. The agent's goal is to maximize the cumulative reward along all the actions it takes, called return. Reinforcement learning describes a set of strategies that the agent can learn to achieve its goal.

To enter in deeper details on RL formalism, we need to introduce some further concepts:

- states and observations,
- action spaces,
- policies,
- trajectories,
- different formulations of return,

- the RL optimization problem,
- value functions.

**States and Observations** A state  $s$  is a complete description of the environment and contains all possible knowledge on the world. An observation  $o(s)$  is only a partial description of the state  $s$ , and in general loses some information. In most RL applications both states and observations are represented with real vectors, matrices, or higher order tensors. When the agent obtains the complete information on the state of the environment, we say the environment is fully observed. When only partial information is known, we say that the environment is partially observed.



**Info:** In RL notation the symbol for the state,  $s$ , is often used instead of the one for observation,  $o$ , even if the latter would be more appropriate, for instance when considering the action choice made by the agent. However, it should be clear from the context which of the two is meant.

**Action Spaces** The proper actions that can be performed depend on the problem, and hence on the environment, that the agent is facing. The action space is the set of all possible valid actions. It can be discrete, when a finite number of actions can be taken, such as in Atari games and Go, or continuous, as in the optimization scheme presented in chapter 6.

**Policies** A policy is a function of the environment state and determines the action chosen by the agent. Here we consider only stochastic policies, denoted by  $\pi$ :

$$a_t \sim \pi(\cdot | s_t). \quad (\text{C.1})$$

The last equation means that the action  $a_t$  is extracted from the conditional probability distribution described by the policy, given the system is in state  $s_t$ . In deep RL, the policies are usually parametric functions, where the output depends on a set of parameters (eg the weights and biases of a neural network) which can be tuned to modify the agent behavior through some optimization algorithms. This set of parameters is denoted by  $\theta$ , and the corresponding parametrized policy is  $\pi_\theta(\cdot | s_t)$ .

Now we discuss briefly two of the most common stochastic policies: categorical policies, for discrete action spaces, and diagonal Gaussian policies, for continuous action spaces.



**Info:** In chapter 6 the policy is denoted with  $\Pi(a|s)$ , to avoid confusion with the number  $\pi$ . Here we choose to follow more strictly common RL notation and call the policy  $\pi(a|s)$ .

**Categorical Policies** In machine learning language, a categorical policy is a classifier over the discrete action space. If the current action  $a$  has to be chosen among  $k$  different options,



and the environment is in state  $s$ , the output of the policy will be a  $k$  dimensional vector  $P_\theta(s)$  that represents the probability of taking each action.  $a$  is thus extracted according to this probability distribution. Of particular interest for policy optimization, which we will describe further on, is the logarithm of the probability distribution, which reads

$$\log \Pi_\theta(a|s) = \log [P_\theta(s)]_a. \quad (\text{C.2})$$

In the last equation we used the action  $a$  as the index for the vector describing the probability distribution  $P_\theta(s)$ .

**Diagonal Gaussian Policies** A multivariate Gaussian distribution is identified by vector  $\mu$  describing the mean of each variable, and a covariance matrix,  $\sigma$ . A diagonal Gaussian distribution is a particular subclass where the covariance matrix has only diagonal entries and can, therefore, be described by a vector too.

In RL, a diagonal Gaussian policy always has a neural network that computes the mean actions given the state (or the observation)  $\mu_\theta(s)$ . Regarding the covariance matrix instead, it can be represented in two ways. In the first, there is a single vector of logarithms of standard deviations,  $\log \sigma$ , which is independent of the state. The Proximal Policy Optimization (PPO) algorithm we used in chapter 6 is implemented in this way. As an alternative, the log standard deviations may be computed by a neural network  $\log \sigma_\theta(s)$ , which could share some layers with the one describing the averages  $\mu_\theta(s)$ .

In both cases, it is convenient to consider the log standard deviations instead of standard deviations directly. This is because  $\log \sigma$  can take any value in  $(-\infty, \infty)$ , while  $\sigma$  must be non-negative. In general, it is easier to train parameters without enforcing such constraint, and the actual standard deviation can easily be recovered by exponentiating its logarithm, without losing information.

The log-likelihood of a  $k$ -dimensional action  $a$  is easily computed using logarithm properties

$$\log \pi_\theta(a|s) = -\frac{1}{2} \left( \sum_{i=1}^k \left( \frac{(a_i - \mu_i)^2}{\sigma_i^2} + 2 \log \sigma_i \right) + k \log 2\pi \right). \quad (\text{C.3})$$

**Trajectories** A trajectory  $\mathbf{T}$  is a sequence of states of the environment and actions performed by the agent,

$$\mathbf{T} = (s_0, a_0, s_1, a_1, \dots). \quad (\text{C.4})$$

When all trajectories have a well defined length or they have a terminal state, they are also called episodes.

State transitions that describe the evolution from  $s_t$  to  $s_{t+1}$ , depend on the dynamics of the environment and on the last action performed *only*. They can be either deterministic,

$$s_{t+1} = f(s_t, a_t) \quad (\text{C.5})$$

or stochastic,

$$s_{t+1} \sim P(\cdot | s_t, a_t). \quad (\text{C.6})$$

The actions are chosen by the agent following a certain policy.

**Reward and Return** The reward function  $R$  is crucial in reinforcement learning, because it describes to the agent the task it is facing. It depends on the state of the environment at time  $t$ , the corresponding action, and the next state  $s_{t+1}$ :

$$r_t = R(s_t, a_t, s_{t+1}) , \quad (\text{C.7})$$

even though it is often simplified to a dependence on the current state only ,  $r_t = R(s_t)$ , or on the state-action pair  $r_t = R(s_t, a_t)$ .

The agent aims to maximize a cumulative reward over a trajectory, denoted with  $R(\mathbf{T})$ . The exact definition of  $R(\mathbf{T})$  depends on the dynamics involved in the environment and on the length of the trajectory.

The simplest definition is just the sum of all the rewards obtained along a trajectory of finite length  $P$

$$R(\mathbf{T}) = \sum_{t=0}^P r_t. \quad (\text{C.8})$$

This is called *finite-horizon undiscounted return*.

Another common choice is the *infinite-horizon discounted return*, where the sum runs over all rewards ever got by the agent, multiplied by a discount factor  $\gamma \in (0, 1)$  that becomes smaller the farther away in the future the reward is obtained.

$$R(\mathbf{T}) = \sum_{t=0}^{\infty} \gamma^t r_t. \quad (\text{C.9})$$

This formulation is particularly useful to derive analytical results because it is easy to ensure convergence thanks to the geometric series. Moreover, it encodes the intuitive idea that a reward now is better than a reward later.  $\gamma > 0$  is fundamental however to learn good strategies for future rewards.

**The RL Problem** The goal of reinforcement learning is to find a policy that maximizes the expected return, when the agent follows it. To treat formally the expected return, it is helpful to consider first the probability distribution of trajectories, in a situation where both the policy and the environment transition are stochastic. The probability of a  $P$ -step trajectory reads

$$P(\mathbf{T}|\pi) = \rho_0(s_0) \prod_{t=0}^{P-1} P(s_{t+1}|s_t, a_t) \pi(a_t|s_t). \quad (\text{C.10})$$

Here  $\rho_0(s_0)$  is the distribution over the initial state. In chapter 6 we considered the simple case where all trajectories start from the same state.

Then we can write the expected return  $J(\pi)$  as

$$J(\pi) = \int_{\mathbf{T}} P(\mathbf{T}|\pi) R(\mathbf{T}) = \mathbf{E}_{\mathbf{T} \sim \pi} R(\mathbf{T}) , \quad (\text{C.11})$$

which is the average return over all trajectories generated by the policy  $\pi$ , each weighted with its probability distribution. The key optimization problem in RL is then summarized by the search of the optimal policy  $\pi^*$ , which satisfies

$$\pi^* = \arg \max_{\pi} J(\pi). \quad (\text{C.12})$$

**Value Functions** Another key element of RL is evaluating how good or bad is to be in a certain state, or to pick an action from a given state, and then evolve afterward according to the policy  $\pi$ . The *value* of a state is represented by its expected reward. We will introduce four functions that formally encode this information.

The *On-Policy Value Function*, or *State Value Function*,  $V^\pi(s)$  gives the expected return if the environment starts initially in state  $s$  and always evolve according to policy  $\pi$ :

$$V^\pi(s) = \mathbf{E}_{\mathbf{T} \sim \pi} R(\mathbf{T}) | s_0 = s \quad (\text{C.13})$$

The On-Policy Action-Value Function  $Q^\pi(s, a)$ , describes the expected return if the environment starts in state  $s$ , any action  $a$  is chosen, even different from the one suggested by the policy, and then the agent forever after acts according to policy  $\pi$ :

$$Q^\pi(s, a) = \mathbf{E}_{\mathbf{T} \sim \pi} R(\mathbf{T}) | s_0 = s, a_0 = a \quad (\text{C.14})$$

It can also be written using the on-policy value function

$$Q^\pi(s, a) = R(s, a, s') + \gamma V^\pi(s'), \quad (\text{C.15})$$

where the first term of the right-hand side is the reward got when going from state  $s$  to state  $s'$  due to the action  $a$ .  $\gamma$  is the discount factor.

The Optimal Value Function  $V^*(s)$  is the value function when the agent follows the optimal policy starting from state  $s$

$$V^*(s) = \max_{\pi} \mathbf{E}_{\mathbf{T} \sim \pi} R(\mathbf{T}) | s_0 = s \quad (\text{C.16})$$

The Optimal Action-Value Function  $Q^*(s, a)$  is equivalent to eq. (C.14), but with the agent following the optimal policy, after action  $a$  is taken

$$Q^*(s, a) = \max_{\pi} \mathbf{E}_{\mathbf{T} \sim \pi} R(\mathbf{T}) | s_0 = s, a_0 = a \quad (\text{C.17})$$

**i**

**Info:** When value functions do not have a time-dependence, they refer to the expected infinite-horizon discounted return. Value functions for finite-horizon undiscounted return need to time as an extra argument.

**i**

**Info:** There are two important relations between value functions and state-action value functions that follow directly from their definitions:

$$V^\pi(s) = \mathbf{E}_{a \sim \pi} Q^\pi(s, a), \quad (\text{C.18})$$

and

$$V^*(s) = \max_a Q^*(s, a). \quad (\text{C.19})$$

An important connection exists between the optimal action-value function  $Q^*(s, a)$  and the action selected by the optimal policy. Following from eq. (C.17) and the definition of the optimal policy  $\Pi^*$ , if we have  $Q^*$ , we can directly obtain the optimal action  $a^*(s)$  via

$$a^*(s) = \arg \max_a Q^*(s, a). \quad (\text{C.20})$$

The optimal action may not be unique, but there might be multiple choices that maximize  $Q^*(s, a)$ , in which case the optimal policy may randomly select any of them.

**Bellman Equations** All the value functions we introduced obey special self-consistency equations called Bellman equations. The Bellman equations are a central element in RL theory to implement policy optimization. The Bellman equations for the on-policy value functions are

$$\begin{aligned} V^\pi(s) &= \mathbf{E}_{a \sim \pi, s' \sim P^r}(s, a) + \gamma V^\pi(s'), \\ Q^\pi(s, a) &= \mathbf{E}_{s' \sim P^r}(s, a) + \gamma \mathbf{E}_{a' \sim \pi} Q^\pi(s', a') \end{aligned} \quad (\text{C.21})$$

where  $s' \sim P$  means that the next state  $s'$  is sampled from the environment's dynamics;  $a \sim \pi$  and  $a' \sim \pi$  indicate that the actions are chosen following the policy  $\pi$  when the state is  $s$  and  $s'$  respectively.

The Bellman equations for the optimal value functions are

$$\begin{aligned} V^*(s) &= \max_a \mathbf{E}_{s' \sim P^r}(s, a) + \gamma V^*(s'), \\ Q^*(s, a) &= \mathbf{E}_{s' \sim P^r}(s, a) + \gamma \max_{a'} Q^*(s', a'). \end{aligned} \quad (\text{C.22})$$

Notice the presence of the max over the actions in eq.(C.22), which is the main difference between the Bellman equations for the On-Policy value functions and the optimal ones. It reflects the fact that whenever the optimal action always leads to the maximum value function.



**Info:** The right-hand side of Bellman equations is called "Bellman backup" in the RL literature.

**Advantage Functions** The advantage function describes how much a specific action  $a$  is better than your average expectation, instead of telling how much it is good in absolute. Formally, the advantage function  $A^\pi(s, a)$  corresponding to a policy  $\pi$  is defined by

$$A^\pi(s, a) = Q^\pi(s, a) - V^\pi(s) = (r + \gamma V^\pi(s')) - V^\pi(s). \quad (\text{C.23})$$

Here  $s'$  is the state reached from  $s$  after action  $a$ . The advantage function is crucially important to policy gradient methods such as PPO, as we will discuss later on.

## C.2. Introduction to policy optimization

In this section, we present the basic mathematical formulation to understand and write policy optimization gradient-based algorithms. RL methods in this class represent a policy explicitly as  $\pi_\theta(a|s)$ . They aim to maximize the expected return  $J(\pi_\theta)$ , by optimizing the parameters  $\theta$  either directly by gradient ascent objective function, or indirectly, by maximizing local approximations of  $J(\pi_\theta)$ . We focus on model-free on-policy methods; model-free means that the agent does not know the dynamic laws of the environment, but has to figure them through the exploration. On-policy means that the data used to estimate the performance, gradients and value functions are all generated following the current policy.

We make use of three fundamental tools in the theory of policy gradients: the simplest equation for the gradient with respect to policy parameters of the policy performance itself, a way to eliminate useless term in that expression, and a rule to add useful terms to that expression. Tying those results together we describe the advantage-based update that lies at the core of most common policy gradient methods.

**Deriving the Simplest Policy Gradient** Let us consider consider a stochastic, parameterized policy,  $\pi_\theta$ . The goal is to maximize the expected return  $J(\pi_\theta) = \mathbf{E}_{\mathbf{T} \sim \pi_\theta} R(\mathbf{T})$ . In the following, we take  $R(\mathbf{T})$  to be the undiscounted finite-horizon return, defined in eq. (C.8). The mathematical derivation for the infinite-horizon discounted return, eq. (C.9), is almost identical.

We would like to maximize the policy performance by steepest ascent

$$\theta_{k+1} = \theta_k + \alpha \nabla_\theta J(\pi_\theta)|_{\theta_k}. \quad (\text{C.24})$$

The gradient of the expected return,  $\nabla_\theta J(\pi_\theta)$ , is called the *policy gradient*, and algorithms that exploit it to optimize the policy belong to the class of policy gradient algorithms. Proximal Policy Optimization, the algorithm we used in chapter 6 is one of those<sup>1</sup>.

The first thing we need is an expression for the policy gradient that can be numerically computed. We will show in the following that the analytical form of the gradient is an expectation value over trajectories, which will be evaluated from the finite number of data as a monte-carlo average. Here we derive the simplest form for these expressions.

Before starting, we highlight some definition and properties which will be useful afterwards.

1. *Probability of a Trajectory.* The probability of a trajectory  $\mathbf{T} = (s_0, a_0, \dots, s_{P+1})$  with actions chosen from the policy  $\pi_\theta$  is

$$P(\mathbf{T}|\theta) = \rho_0(s_0) \prod_{t=0}^P P(s_{t+1}|s_t, a_t) \pi_\theta(a_t|s_t). \quad (\text{C.25})$$

2. *The Log-Derivative Trick.* The log-derivative trick follows directly from the derivative

<sup>1</sup>rigorously is slightly inappropriate to call PPO a policy gradient algorithm, but these technical details are beyond the scope of this appendix

of logarithms

$$\nabla_{\theta} P(\mathbf{T}|\theta) = P(\mathbf{T}|\theta) \nabla_{\theta} \log P(\mathbf{T}|\theta). \quad (\text{C.26})$$

3. *Log-Probability of a Trajectory.* The log-prob of a trajectory  $\mathbf{T}$  is simply

$$\log P(\mathbf{T}|\theta) = \log \rho_0(s_0) + \sum_{t=0}^P \left( \log P(s_{t+1}|s_t, a_t) + \log \pi_{\theta}(a_t|s_t) \right). \quad (\text{C.27})$$

4. *Gradients of Environment Functions.* All functions of the environment only has no dependence on  $\theta$ , so gradients of  $\rho_0(s_0)$ ,  $P(s_{t+1}|s_t, a_t)$ , and  $R(\mathbf{T})$  are zero.

5. *Grad-Log-Prob of a Trajectory.* The gradient of the log-prob of a trajectory is therefore

$$\begin{aligned} \nabla_{\theta} \log P(\mathbf{T}|\theta) &= \nabla_{\theta} \log \rho_0(s_0) + \sum_{t=0}^P \left( \nabla_{\theta} \log P(s_{t+1}|s_t, a_t) + \nabla_{\theta} \log \pi_{\theta}(a_t|s_t) \right) \\ &= \sum_{t=0}^P \nabla_{\theta} \log \pi_{\theta}(a_t|s_t). \end{aligned} \quad (\text{C.28})$$

Using all these tools, we are able to derive the basic policy gradient expression

$$\nabla_{\theta} J(\pi_{\theta}) = \nabla_{\theta} \mathbf{E}_{\mathbf{T} \sim \pi_{\theta}} R(\mathbf{T}) \quad (\text{C.29})$$

$$= \nabla_{\theta} \int_{\mathbf{T}} P(\mathbf{T}|\theta) R(\mathbf{T}) \quad \text{Expand expectation} \quad (\text{C.30})$$

$$= \int_{\mathbf{T}} \nabla_{\theta} P(\mathbf{T}|\theta) R(\mathbf{T}) \quad \text{Bring gradient under integral} \quad (\text{C.31})$$

$$= \int_{\mathbf{T}} P(\mathbf{T}|\theta) \nabla_{\theta} \log P(\mathbf{T}|\theta) R(\mathbf{T}) \quad \text{Log-derivative trick} \quad (\text{C.32})$$

$$= \mathbf{E}_{\mathbf{T} \sim \pi_{\theta}} \nabla_{\theta} \log P(\mathbf{T}|\theta) R(\mathbf{T}) \quad \text{Return to expectation form} \quad (\text{C.33})$$

$$\therefore \nabla_{\theta} J(\pi_{\theta}) = \mathbf{E}_{\mathbf{T} \sim \pi_{\theta}} \sum_{t=0}^T \nabla_{\theta} \log \pi_{\theta}(a_t|s_t) R(\mathbf{T}) \quad \text{Expression for grad-log-prob} \quad (\text{C.34})$$

This is an average over the trajectories distribution, which in practice is estimated numerically with a sample mean. If we collect a set of trajectories  $\mathcal{N} = \{\mathbf{T}_i\}_{i=1, \dots, N}$  acting according the policy  $\pi_{\theta}$ , the policy gradient can be estimated with

$$\hat{g} = \frac{1}{|\mathcal{D}|} \sum_{\mathbf{T} \in \mathcal{D}} \sum_{t=0}^P \nabla_{\theta} \log \pi_{\theta}(a_t|s_t) R(\mathbf{T}), \quad (\text{C.35})$$

where  $|\mathcal{N}| = N$  is the number of sampled trajectories.

Assuming that the policy is written in way which allows to compute  $\nabla_{\theta} \log \pi_{\theta}(a|s)$ , eq. (C.35) makes it possible to estimate the policy gradient from a set of trajectories run on the environment, and then update the policy parameters  $\theta$ .

### C.3. Proximal Policy Optimization

PPO is an on-policy algorithm that can be used both with discrete and continuous action spaces. It is motivated by the following question: how can we make the biggest possible

improvement step when updating a policy, without stepping so far that the performance collapses?

There are several methods used in PPO algorithms to keep the new policy close to the old one. Here we describe PPO-Clip, which is the one implemented in the OpenAI SpinningUp library. In a nutshell, PPO-Clip relies on specialized clipping in the objective function to the maximum allowed change in the policy. More precisely, the policy parameters are updated following

$$\theta_{k+1} = \arg \max_{\theta} \mathbb{E}_{s, a \sim \pi_{\theta_k}} [L(s, a, \theta_k, \theta)]. \quad (\text{C.36})$$

Usually, multiple steps of stochastic gradient descent (SDG) are taken to maximize the objective function. Here  $L$  is given by

$$L(s, a, \theta_k, \theta) = \min \left( \frac{\pi_{\theta}(a|s)}{\pi_{\theta_k}(a|s)} A^{\pi_{\theta_k}}(s, a), g(\epsilon, A^{\pi_{\theta_k}}(s, a)) \right), \quad (\text{C.37})$$

where  $A^{\pi_{\theta_k}}(s, a)$  is the parameterized advantage function defined in eq. (C.23), and

$$g(\epsilon, A) = \begin{cases} (1 + \epsilon)A & A \geq 0 \\ (1 - \epsilon)A & A < 0. \end{cases} \quad (\text{C.38})$$

$\epsilon$  is a small hyperparameter that controls how far away the new policy and the old one can be.

To understand how this update works, it is useful to focus on a single state action pair  $(s, a)$ , where two cases may happen.

Advantage is positive: if  $A^{\pi_{\theta_k}}(s, a) > 0$  its contribution to the objective function reduces to

$$L(s, a, \theta_k, \theta) = \min \left( \frac{\pi_{\theta}(a|s)}{\pi_{\theta_k}(a|s)}, (1 + \epsilon) \right) A^{\pi_{\theta_k}}(s, a). \quad (\text{C.39})$$

Since the advantage is positive, the objective will increase if  $\pi_{\theta}(a|s)$  becomes larger, meaning that the action is chosen more likely. But the min in eq. (C.39) limits how much the objective can increase. Once  $\pi_{\theta}(a|s) > (1 + \epsilon)\pi_{\theta_k}(a|s)$ , the this term is clipped by the min to  $(1 + \epsilon)A^{\pi_{\theta_k}}(s, a)$ , preventing the new policy from taking a benefit in going too far away from the old one.

Advantage is negative: if  $A^{\pi_{\theta_k}}(s, a) < 0$ , its contribution to the objective function reads

$$L(s, a, \theta_k, \theta) = \max \left( \frac{\pi_{\theta}(a|s)}{\pi_{\theta_k}(a|s)}, (1 - \epsilon) \right) A^{\pi_{\theta_k}}(s, a). \quad (\text{C.40})$$

Because the advantage is negative, the objective will increase if  $\pi_{\theta}(a|s)$  decreases, that is if the action is chosen less likely. Again the max here limit how much the objective function can increase. When  $\pi_{\theta}(a|s) < (1 - \epsilon)\pi_{\theta_k}(a|s)$ , it is clipped to  $(1 - \epsilon)A^{\pi_{\theta_k}}(s, a)$  by the max in eq. (C.40), and again the new policy does not benefit from changing too much.

The hyperparameter  $\epsilon$  thus works as a regularizer that prevents dramatic changes in the policy, which may lead to large drops in the performance.

PPO is an on-policy method to train a stochastic policy. This means that the exploration is carried out by sampling actions following only the current version of the stochastic policy.

The randomness in the chosen actions is influenced both by the initial condition and by the training process. During the training, the action selection becomes less and less random, as the update encourages the agent to exploit choices already explored. However, the drawback of encouraging exploitation over exploration at the end of the training is that the policy may get trapped in local maxima.

### PPO Pseudocode

- 1: initialize policy parameters  $\theta_0$  and value function parameters  $\phi_0$
- 2: for  $k = 0, 1, 2, \dots$  do
- 3: sample batch of trajectories  $\mathcal{N}_k = \mathbf{T}_i$  following policy  $\pi_k = \pi_{\theta_k}$ .
- 4: compute rewards  $\hat{R}_t$ .
- 5: Estimate advantages  $\hat{A}_t$  using current value function  $V_{\phi_k}$ .
- 6: update policy via eq. (C.36).
- 7: Fit value function and update  $\phi_{k+1}$ .

## C.4. Choice of policy and training parameters

The results presented in chapter 6 are obtained by training a diagonal Gaussian policy with PPO algorithm, for 1024 epochs of 100 episodes each. The reward function is the simplest possible  $r_t = -\delta_{t,P} \langle \psi(\boldsymbol{\gamma}, \boldsymbol{\beta}) | \hat{H}_{\text{target}} | \psi(\boldsymbol{\gamma}, \boldsymbol{\beta}) \rangle$ , and the RL agent receives it only at the end of each episode ( $t = P$ ). Here we discuss briefly our choices of training parameters, the *hyperparameters* in RL language.

The PPO algorithm has been chosen because it is one of the most advanced RL methods suited for problems with a continuous action space [191, 194], such as QAOA. This algorithm is implemented in the OpenAI SpinningUp [194] library with a stochastic diagonal Gaussian policy. This means that at each step the two parameters which constitute the action  $a_t = (\gamma_t, \beta_t)$  are extracted from independent Gaussian distributions, with the averages given by the output of the Neural Network that parametrizes the policy. The logarithm of the variance of the two Gaussian distributions are also parameters learned during the training process. The code for the Quantum environment is publicly available on GitHub <https://github.com/mwauters92/QuantumRL>.

The reward function must measure the variational quality of the final wavefunction  $|\psi(\boldsymbol{\gamma}, \boldsymbol{\beta})\rangle_P$ , hence it must be a monotonic increasing function of *minus* the final energy  $E_P(\boldsymbol{\gamma}, \boldsymbol{\beta})$ . We tested two choices:

$$R(E_P(\boldsymbol{\gamma}, \boldsymbol{\beta})) = -E_P(\boldsymbol{\gamma}, \boldsymbol{\beta}) \quad (\text{C.41})$$

and

$$R(E_P(\boldsymbol{\gamma}, \boldsymbol{\beta})) = e^{-4E_P(\boldsymbol{\gamma}, \boldsymbol{\beta})/N} , \quad (\text{C.42})$$






Figure C.1.: Difference between the residual energy during the training and QAOA optimal value  $\frac{1}{2(P+2)}$ , for a uniform Ising chain of  $N = 32$  spins and episode length  $P = 10$ . We compare two choices of the reward function given by Eq. (C.41) and Eq. (C.42), blue and red lines respectively. The black solid line is a heuristic upper bound for the convergence speed of the residual energy obtained from RL protocol to the optimal value.

where the system size  $N$  is used to prevent the reward to diverge. The factor 4 increases the steepness of the reward function towards the optimal control value  $E_P(\boldsymbol{\gamma}, \boldsymbol{\beta}) = -N$  (for the TFIM model). A possible advantage of the exponential choice over the linear one is indeed the higher derivative towards the maximum possible reward, which should improve policy optimization when the agent has already reached a good strategy. However, we do not see an appreciable difference between the two choices, as reported in fig. C.1, where we show the convergence of the corresponding residual energy to the QAOA optimal value  $\epsilon_P^{\text{res}} = (2P + 2)^{-1}$ . There is little difference between the two choices of the reward function. Moreover one can see that the residual energy decrease very slowly with the number of epochs, making it inconvenient to try reaching the optimal value by increasing  $N_{\text{epo}}$  instead of adding a local optimization on top of the Reinforcement Learning process.

To choose the number of episodes per training epoch, i.e. the number of trajectories used to evaluate and update the policy, we tested several values  $N_{\text{epo}} = 50, 100, 150, 200$ , on a single TFIM model with  $N = 32$  and  $P = 10$ . We found that  $N_{\text{epi}} = 100$  gives the lowest average residual energy, as reported in Fig. C.2.

Regarding the observables provided to the agent, there are several available choices. The first is full tomography of the wavefunction, which has the huge disadvantage of requiring an exponentially large number of measurements to provide reliable information. Moreover, it is a redundant description of the state, and the neural network needs first to learn how to compress it and extract the relevant information, before optimizing the policy. The performance of the method with this choice turns out to be rather poor and the convergence towards an optimal strategy very slow. Furthermore, the number of nodes in the NN has to scale with the system size in order to be able to extract the information, worsening efficiency at larger sizes, and hampering transferability.

When focusing on observables, one of the most intuitive choices is providing the exp-




Figure C.2.: Average residual energy at the end of the training versus the number of episodes per training epoch. The data are obtained for a uniform Ising chain of  $N = 32$  spins and episode length  $P = 10$ .

tation values of the problem and the driving Hamiltonians,  $\hat{H}_z$  and  $\hat{H}_x$ . This is what we used throughout this Letter. This choice has the advantage of being easily accessible from the Jordan-Wigner representation of the Ising chain and allows us to visualize the policy as a vector function of two real variables. Moreover, it is an efficient description of the state for what regards the optimization task, since  $\langle \hat{H}_z \rangle$  is directly linked to the reward function. This set of observable has been enlarged to include correlation functions at longer distances  $\langle \hat{\sigma}_i^z \hat{\sigma}_{i+k}^z \rangle$ . The test we performed did not indicate any increase in the performance. Preliminary results instead suggest that  $\langle \hat{H}_z \rangle$  alone is sufficient to learn optimal smooth schedules.

An alternative choice would be computing the expectation value of the three components of the magnetization  $\langle \hat{S}^\alpha \rangle$ , with  $\alpha = x, y, z$ , as done in Ref. [189]. This, however, does not work when the state preserves  $\mathbb{Z}_2$  symmetry, as in the TFIM case.

## Bibliography

- [1] A. Messiah. *Quantum mechanics*, volume 2. North-Holland, Amsterdam, 1962.
- [2] K. Drese and M. Holthaus. Floquet theory for short laser pulses. *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, 5(1):119–134, Jan 1999.
- [3] Zhenghao Gu, HA Fertig, Daniel P Arovas, and Assa Auerbach. Floquet spectrum and transport through an irradiated graphene ribbon. *Physical review letters*, 107(21):216601, 2011.
- [4] Jérôme Cayssol, Balázs Dóra, Ferenc Simon, and Roderich Moessner. Floquet topological insulators. *Physica status solidi (RRL)-Rapid Research Letters*, 7(1-2):101–108, 2013.
- [5] N. Fläschner, B. S. Rem, M. Tarnowski, D. Vogel, D.-S. Lühmann, K. Sengstock, and C. Weitenberg. Experimental reconstruction of the berry curvature in a floquet bloch band. *Science*, 352(6289):1091–1094, 2016.
- [6] Martin Holthaus. Floquet engineering with quasienergy bands of periodically driven optical lattices. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 49(1):013001, 2016.
- [7] Jan Carl Budich, Ying Hu, and Peter Zoller. Helical floquet channels in 1d lattices. *Phys. Rev. Lett.*, 118(10):105302, 2017.
- [8] Fenner Harper, Rahul Roy, Mark S. Rudner, and S.L. Sondhi. Topology and broken symmetry in floquet systems. *Annual Review of Condensed Matter Physics*, 11(1):345–368, 2020.
- [9] David J. Thouless. Quantization of particle transport. *Phys. Rev. B*, 27(10):6083, 1983.
- [10] Yakir Aharonov and J Anandan. Phase change during a cyclic quantum evolution. *Phys. Rev. Lett.*, 58(16):1593, 1987.
- [11] Lorenzo Privitera and Giuseppe E. Santoro. Quantum annealing and nonequilibrium dynamics of floquet chern insulators. *Phys. Rev. B*, 93(24):241406, 2016.
- [12] Paraj Titum, Erez Berg, Mark S. Rudner, Gil Refael, and Netanel H. Lindner. Anomalous floquet-anderson insulator as a nonadiabatic quantized charge pump. *Phys. Rev. X*, 6:021013, May 2016.
- [13] Iliya Esin, Mark S. Rudner, Gil Refael, and Netanel H. Lindner. Quantized transport and steady states of floquet topological insulators. *Phys. Rev. B*, 97:245401, Jun 2018.

- [14] Michael H. Kolodrubetz, Frederik Nathan, Snir Gazit, Takahiro Morimoto, and Joel E. Moore. Topological floquet-thouless energy pump. *Phys. Rev. Lett.*, 120:150601, Apr 2018.
- [15] Qian Niu and DJ Thouless. Quantised adiabatic charge transport in the presence of substrate disorder and many-body interaction. *J. Phys. A-Math. Gen.*, 17(12):2453, 1984.
- [16] Immanuel Bloch, Jean Dalibard, and Wilhelm Zwerger. Many-body physics with ultracold gases. *Rev. Mod. Phys.*, 80:885–964, Jul 2008.
- [17] Tomoki Ozawa, Hannah M. Price, Nathan Goldman, Oded Zilberberg, and Iacopo Carusotto. Synthetic dimensions in integrated photonics: From optical isolation to four-dimensional quantum hall physics. *Phys. Rev. A*, 93:043827, Apr 2016.
- [18] Alexander B Khanikaev, S Hossein Mousavi, Wang-Kong Tse, Mehdi Kargarian, Allan H MacDonald, and Gennady Shvets. Photonic topological insulators. *Nature materials*, 12(3):233–239, 2013.
- [19] Hirokazu Miyake, Georgios A. Siviloglou, Colin J. Kennedy, William Cody Burton, and Wolfgang Ketterle. Realizing the harper hamiltonian with laser-assisted tunneling in optical lattices. *Phys. Rev. Lett.*, 111:185302, Oct 2013.
- [20] M. Aidelsburger, M. Atala, M. Lohse, J. T. Barreiro, B. Paredes, and I. Bloch. Realization of the hofstadter hamiltonian with ultracold atoms in optical lattices. *Phys. Rev. Lett.*, 111:185301, Oct 2013.
- [21] Ling Lu, John D. Joannopoulos, and Marin Soljavcic. Topological photonics. *Nature Photonics*, 8:821, Oct 2014.
- [22] Michael Lohse, Christian Schweizer, Oded Zilberberg, Monika Aidelsburger, and Immanuel Bloch. A thouless quantum pump with ultracold bosonic atoms in an optical superlattice. *Nat. Phys.*, 12(4):350–354, 2016.
- [23] Shuta Nakajima, Takafumi Tomita, Shintaro Taie, Tomohiro Ichinose, Hideki Ozawa, Lei Wang, Matthias Troyer, and Yoshiro Takahashi. Topological thouless pumping of ultracold fermions. *Nat. Phys.*, 12:296–300, 2016.
- [24] Pranjal Bordia, Henrik Lüschen, Ulrich Schneider, Michael Knap, and Immanuel Bloch. Periodically driving a many-body localized quantum system. *Nature Physics*, 13:460, Jan 2017.
- [25] Michael Lohse, Christian Schweizer, Hannah M. Price, Oded Zilberberg, and Immanuel Bloch. Exploring 4d quantum hall physics with a 2d topological charge pump. *Nature*, 553:55, march 2018.
- [26] Seth Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, 1996.
- [27] M. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

- 
- [28] A. B. Finnila, M. A. Gomez, C. Sebenik, C. Stenson, and J. D. Doll. Quantum annealing: A new method for minimizing multidimensional functions. *Chem. Phys. Lett.*, 219:343, 1994.
- [29] T. Kadowaki and H. Nishimori. Quantum annealing in the transverse Ising model. *Phys. Rev. E*, 58:5355, 1998.
- [30] Tameem Albash and Daniel A. Lidar. Adiabatic quantum computation. *Rev. Mod. Phys.*, 90:015002, 2018.
- [31] R. Harris et al. Experimental investigation of an eight qubit unit cell in a superconducting optimization processor. *Phys. Rev. B*, 82:024511, 2010.
- [32] M. W. Johnson et al. Quantum annealing with manufactured spins. *Nature*, 473:194–198, 2011.
- [33] S. Boixo, T. Albash, F. M. Spedalieri, N. Chancellor, and D. A. Lidar. Experimental signature of programmable quantum annealing. *Nat. Commun.*, 4:2067, 2013.
- [34] Subir Sachdev. *Quantum Phase Transitions (Second Edition)*. Cambridge University Press, 2011.
- [35] J. E. Avron, M. Fraas, G. M. Graf, and P. Grech. Optimal time schedule for adiabatic evolution. *Phys. Rev. A*, 82:040304, Oct 2010.
- [36] 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.
- [37] Maximilian Keck, Simone Montangero, Giuseppe E. Santoro, Rosario Fazio, and Davide Rossini. Dissipation in adiabatic quantum computers: lessons from an exactly solvable model. *New J. Phys.*, 19(11):113029, 2017.
- [38] G. Passarelli, G. De Filippis, V. Cataudella, and P. Lucignano. Dissipative environment may improve the quantum annealing performances of the ferromagnetic  $p$ -spin model. *Phys. Rev. A*, 97:022319, Feb 2018.
- [39] Hidetoshi Nishimori and Kabuki Takada. Exponential enhancement of the efficiency of quantum annealing by non-stoquastic hamiltonians. *Frontiers in ICT*, 4:2, 2017.
- [40] Adolfo del Campo and Kihwan Kim. Focus on shortcuts to adiabaticity. *New Journal of Physics*, 21(5):050201, may 2019.
- [41] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018.
- [42] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2):023023, feb 2016.
- [43] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M. Chow, and Jay M. Gambetta. Hardware-efficient variational quantum eigen-

- solver for small molecules and quantum magnets. *Nature*, 549(7671):242–246, Sep 2017.
- [44] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A Quantum Approximate Optimization Algorithm. *arXiv e-prints*, page arXiv:1411.4028, Nov 2014.
- [45] Gershon Kurizki, Patrice Bertet, Yuimaru Kubo, Klaus Mølmer, David Petrosyan, Peter Rabl, and Jörg Schmiedmayer. Quantum technologies with hybrid systems. *Proc. Natl. Acad. Sci. USA*, 112(13):3866–3873, 2015.
- [46] Nikolaj Moll, Panagiotis Barkoutsos, Lev S Bishop, Jerry M Chow, Andrew Cross, Daniel J Egger, Stefan Filipp, Andreas Fuhrer, Jay M Gambetta, Marc Ganzhorn, Abhinav Kandala, Antonio Mezzacapo, Peter MÅ $\frac{1}{4}$ ller, Walter Riess, Gian Salis, John Smolin, Ivano Tavernelli, and Kristan Temme. Quantum optimization using variational algorithms on near-term quantum devices. *Quantum Science and Technology*, 3(3):030503, jun 2018.
- [47] N. D. Mermin. The topological theory of defects in ordered media. *Rev. Mod. Phys.*, 51:591–648, Jul 1979.
- [48] David J Thouless. *Topological Quantum Numbers in Nonrelativistic Physics*. WORLD SCIENTIFIC, 1998.
- [49] Alexander Altland and Ben D. Simons. *Condensed Matter Field Theory*. Cambridge University Press, 2 edition, 2010.
- [50] 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.
- [51] Klaus von Klitzing. The quantized Hall effect. *Rev. Mod. Phys.*, 58:519, 1986.
- [52] Robert B Laughlin. Quantized hall conductivity in two dimensions. *Phys. Rev. B*, 23(10):5632, 1981.
- [53] Steven M. Girvin Richard E. Prange. *The Quantum Hall Effect*. Graduate Texts in Contemporary Physics. Springer-Verlag New York, 2 edition, 1990.
- [54] DJ Thouless, M Kohmoto, MP Nightingale, and M Den Nijs. Quantized Hall conductance in a two-dimensional periodic potential. *Phys. Rev. Lett.*, 49(6):405, 1982.
- [55] Qian Niu, D J Thouless, and Yong-Shi Wu. Quantized hall conductance as a topological invariant. *Phys. Rev. B*, 31(6):3372, 1985.
- [56] Mahito Kohmoto. Topological invariant and the quantization of the hall conductance. *Annals of Physics*, 160(2):343–354, 1985.
- [57] Liang Fu, Charles L Kane, and Eugene J Mele. Topological insulators in three dimensions. *Phys. Rev. Lett.*, 98(10):106803, 2007.
- [58] YL Chen, JG Analytis, J-H Chu, ZK Liu, S-K Mo, Xiao-Liang Qi, HJ Zhang, DH Lu, Xi Dai, Zhong Fang, et al. Experimental realization of a three-dimensional topological insulator, bi<sub>2</sub>te<sub>3</sub>. *Science*, 325(5937):178–181, 2009.

- 
- [59] Alexei Kitaev. Periodic table for topological insulators and superconductors. *arXiv preprint arXiv:0901.2686*, 2009.
- [60] Charles L. Kane. Topological band theory and the  $z_2$  invariant. In Marcel Franz and Laurens W. Molenkamp, editors, *Topological Insulators*, volume 6, page 3. Elsevier, 2013.
- [61] Yoichi Ando. Topological insulator materials. *Journal of the Physical Society of Japan*, 82(10):102001, 2013.
- [62] Xiao-Liang Qi and Shou-Cheng Zhang. Topological insulators and superconductors. *Rev. Mod. Phys.*, 83(4):1057, 2011.
- [63] Binghai Yan and Claudia Felser. Topological materials: Weyl semimetals. *Annual Review of Condensed Matter Physics*, 8(1):337–354, 2017.
- [64] M Klein and Ruedi Seiler. Power-law corrections to the kubo formula vanish in quantum hall systems. *Comm. Math. Phys.*, 128(1):141–160, 1990.
- [65] Joseph E Avron and Zvi Kons. Quantum response at finite fields and breakdown of chern numbers. *J. Phys. A-Math. Gen.*, 32(33):6097, 1999.
- [66] Joseph E Avron, Daniel Osadchy, and Ruedi Seiler. A topological look at the quantum hall effect. *Phys. Today*, 56(8):38–42, 2003.
- [67] RD King-Smith and David Vanderbilt. Theory of polarization of crystalline solids. *Phys. Rev. B*, 47(3):1651, 1993.
- [68] Gerardo Ortiz and Richard M Martin. Macroscopic polarization as a geometric quantum phase: Many-body formulation. *Phys. Rev. B*, 49(20):14202, 1994.
- [69] Raffaele Resta. Macroscopic polarization in crystalline dielectrics: the geometric phase approach. *Rev. Mod. Phys.*, 66:899–915, Jul 1994.
- [70] Takuya Kitagawa, Erez Berg, Mark Rudner, and Eugene Demler. Topological characterization of periodically driven quantum systems. *Phys. Rev. B*, 82(23):235114, 2010.
- [71] Netanel H Lindner, Gil Refael, and Victor Galitski. Floquet topological insulator in semiconductor quantum wells. *Nature Physics*, 7(6):490–495, 2011.
- [72] Paraj Titum, Netanel H. Lindner, Mikael C. Rechtsman, and Gil Refael. Disorder-induced floquet topological insulators. *Phys. Rev. Lett.*, 114:056801, Feb 2015.
- [73] Wenchao Ma, Longwen Zhou, Qi Zhang, Min Li, Chunyang Cheng, Jianpei Geng, Xing Rong, Fazhan Shi, Jiangbin Gong, and Jiangfeng Du. Experimental observation of a generalized thouless pump with a single spin. *Phys. Rev. Lett.*, 120:120501, Mar 2018.
- [74] Inbar Hotzen Grinberg, Mao Lin, Cameron Harris, Wladimir A. Benalcazar, Christopher W. Peterson, Taylor L. Hughes, and Gaurav Bahl. Robust temporal pumping in a magneto-mechanical topological insulator. *arXiv e-prints*, page arXiv:1905.02778, May 2019.

- [75] Lorenzo Privitera, Angelo Russomanno, Roberta Citro, and Giuseppe E. Santoro. Nonadiabatic breaking of topological pumping. *Phys. Rev. Lett.*, 120:106601, Mar 2018.
- [76] P. W. Anderson. Absence of diffusion in certain random lattices. *Phys. Rev.*, 109:1492–1505, Mar 1958.
- [77] H. Hatami, C. Danieli, J. D. Bodyfelt, and S. Flach. Quasiperiodic driving of anderson localized waves in one dimension. *Physical Review E*, 93(6), jun 2016.
- [78] Kartiek Agarwal, Sriram Ganeshan, and R. N. Bhatt. Localization and transport in a strongly driven anderson insulator. *Phys. Rev. B*, 96:014201, Jul 2017.
- [79] Luca Arceci, Lucas Kohn, Angelo Russomanno, and Giuseppe E Santoro. Dissipation assisted thouless pumping in the rice–mele model. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(4):043101, apr 2020.
- [80] Andrew Lucas. Ising formulations of many np problems. *Frontiers in Physics*, 2:5, 2014.
- [81] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6):467–488, Jun 1982.
- [82] Endre Boros and Peter L. Hammer. Pseudo-boolean optimization. *Discrete Applied Mathematics*, 123(1):155 – 225, 2002.
- [83] Iulia Buluta and Franco Nori. Quantum simulators. *Science*, 326(5949):108–111, 2009.
- [84] I. M. Georgescu, S. Ashhab, and Franco Nori. Quantum simulation. *Rev. Mod. Phys.*, 86:153–185, Mar 2014.
- [85] J. Brooke, D. Bitko, T. F. Rosenbaum, and G. Aeppli. Quantum annealing of a disordered magnet. *Science*, 284:779, 1999.
- [86] Giuseppe E. Santoro, R. Martovnák, E. Tosatti, and R. Car. Theory of quantum annealing of an Ising spin glass. *Science*, 295:2427, 2002.
- [87] Giuseppe E. Santoro and Erio Tosatti. Optimization using quantum mechanics: Quantum annealing through adiabatic evolution. *J. Phys. A: Math. Gen.*, 39:R393–R431, 2006.
- [88] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda. A quantum adiabatic evolution algorithm applied to random instances of an NP-Complete problem. *Science*, 292:472, 2001.
- [89] Sergio Boixo, Troels F. Rønnow, Sergei V. Isakov, Zhihui Wang, David Wecker, Daniel A. Lidar, John M. Martinis, and Matthias Troyer. Evidence for quantum annealing with more than one hundred qubits. *Nat. Phys.*, 10(3):218–224, 03 2014.
- [90] Glen Bigan Mbeng, Rosario Fazio, and Giuseppe E. Santoro. Optimal quantum control with digitized quantum annealing, 2019.
- [91] R. Barends, A. Shabani, L. Lamata, J. Kelly, A. Mezzacapo, U. Las Heras, R. Babbush,



- A. G. Fowler, B. Campbell, Yu Chen, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, E. Solano, H. Neven, and John M. Martinis. Digitized adiabatic quantum computing with a superconducting circuit. *Nature*, 534:222 EP –, 06 2016.
- [92] S. Bravyi, D. P. DiVincenzo, R. I. Oliveira, and B. M. Terhal. The complexity of stoquastic local hamiltonian problems. *Quantum Inf. Comput.*, 8:361–385, 2008.
- [93] Matthias Troyer and Uwe-Jens Wiese. Computational complexity and fundamental limitations to fermionic quantum monte carlo simulations. *Phys. Rev. Lett.*, 94:170201, May 2005.
- [94] Seth Lloyd. Quantum approximate optimization is computationally universal. *arXiv e-prints*, page arXiv:1812.11075, Dec 2018.
- [95] Leo Zhou, Sheng-Tao Wang, Soonwon Choi, Hannes Pichler, and Mikhail D. Lukin. Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices. *Phys. Rev. X*, 10:021067, Jun 2020.
- [96] Mbeng, Glen B. and Fazio, Rosario and Santoro, Giuseppe E. Quantum Annealing: a journey through Digitalization, Control, and hybrid Quantum Variational schemes. *arXiv e-prints*, page arXiv:1906.08948, Jun 2019.
- [97] Mauro E. S. Morales, Jacob Biamonte, and Zoltán Zimborás. On the universality of the quantum approximate optimization algorithm. 2019.
- [98] G. Pagano, A. Bapat, P. Becker, K. S. Collins, A. De, P. W. Hess, H. B. Kaplan, A. Kyprianidis, W. L. Tan, C. Baldwin, L. T. Brady, A. Deshpande, F. Liu, S. Jordan, A. V. Gorshkov, and C. Monroe. Quantum Approximate Optimization with a Trapped-Ion Quantum Simulator. *arXiv e-prints*, page arXiv:1906.02700, Jun 2019.
- [99] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Sergio Boixo, Michael Broughton, Bob B. Buckley, David A. Buell, Brian Burkett, Nicholas Bushnell, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Sean Demura, Andrew Dunsworth, Edward Farhi, Austin Fowler, Brooks Foxen, Craig Gidney, Marissa Giustina, Rob Graff, Steve Habegger, Matthew P. Harrigan, Alan Ho, Sabrina Hong, Trent Huang, L. B. Ioffe, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Cody Jones, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Seon Kim, Paul V. Klimov, Alexander N. Korotkov, Fedor Kostritsa, David Landhuis, Pavel Laptev, Mike Lindmark, Martin Leib, Erik Lucero, Orion Martin, John M. Martinis, Jarrod R. McClean, Matt McEwen, Anthony Megrant, Xiao Mi, Masoud Mohseni, Wojciech Mruczkiewicz, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Florian Neukart, Hartmut Neven, Murphy Yuezheng Niu, Thomas E. O'Brien, Bryan O'Gorman, Eric Ostby, Andre Petukhov, Harald Putterman, Chris Quintana, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Andrea Skolik, Vadim Smelyanskiy, Doug Strain, Michael Streif, Kevin J. Sung, Marco Szalay, Amit Vainsencher, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, and Leo Zhou. Quantum approximate optimization of non-planar graph problems on a planar superconducting

- processor. 2020.
- [100] Victor Bapst and Guilhem Semerjian. On quantum mean field models and their quantum annealing. *JSTAT*, page P06007, 2012.
  - [101] T. JÄPrg, F. Krzakala, J. Kurchan, A. C. Maggs, and J. Pujos. Energy gaps in quantum first-order mean-field-like transitions: The problems that quantum annealing cannot solve. *EPL*, 89(4):40004, 2010.
  - [102] Matteo M. Wauters, Rosario Fazio, Hidetoshi Nishimori, and Giuseppe E. Santoro. Direct comparison of quantum and simulated annealing on a fully connected ising ferromagnet. *Phys. Rev. A*, 96:022326, Aug 2017.
  - [103] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning, An Introduction*. The MIT Press, 2 edition, 2018.
  - [104] Haruki Watanabe. Insensitivity of bulk properties to the twisted boundary condition. *Phys. Rev. B*, 98:155137, Oct 2018.
  - [105] MJ Rice and EJ Mele. Elementary excitations of a linearly conjugated diatomic polymer. *Phys. Rev. Lett.*, 49(19):1455, 1982.
  - [106] Matteo M. Wauters and Giuseppe E. Santoro. Quantization of the hall conductivity in the harper-hofstadter model. *Phys. Rev. B*, 98:205112, Nov 2018.
  - [107] Di Xiao, Ming-Che Chang, and Qian Niu. Berry phase effects on electronic properties. *Rev. Mod. Phys.*, 82:1959–2007, Jul 2010.
  - [108] M. Aidelsburger, M. Lohse, C. Schweizer, M Atala, J. T. Barreriro, S. NascimbÄ“ne, N. R. Cooper, I. Bloch, and N. Goldman. Measuring the chern number of hofstadter bands with ultracold bosonic atoms. *Nat. Phys.*, 11:162, 2015.
  - [109] Douglas R. Hofstadter. Energy levels and wave functions of bloch electrons in rational and irrational magnetic fields. *Phys. Rev. B*, 14:2239–2249, Sep 1976.
  - [110] Jean Dalibard, Fabrice Gerbier, Gediminas Juzeliūnas, and Patrik Öhberg. Colloquium: Artificial gauge potentials for neutral atoms. *Rev. Mod. Phys.*, 83:1523–1543, Nov 2011.
  - [111] Milena Grifoni and Peter Hänggi. Driven quantum tunneling. *Physics Reports*, 304(5):229–354, 1998.
  - [112] Angelo Russomanno, Alessandro Silva, and Giuseppe E Santoro. Periodic steady regime and interference in a periodically driven quantum system. *Phys. Rev. Lett.*, 109(25):257201, 2012.
  - [113] J. von Neumann and E. P. Wigner. *Über das Verhalten von Eigenwerten bei adiabatischen Prozessen*, pages 294–297. Springer Berlin Heidelberg, Berlin, Heidelberg, 1993.
  - [114] Wei-Kai Shih and Qian Niu. Nonadiabatic particle transport in a one-dimensional electron system. *Phys. Rev. B*, 50(16):11902, 1994.

- 
- [115] Gustavo Rigolin, Gerardo Ortiz, and Victor Hugo Ponce. Beyond the quantum adiabatic approximation: Adiabatic perturbation theory. *Phys. Rev. A*, 78(5):052508, 2008.
- [116] H.P. Breuer and M. Holthaus. Quantum phases and landau-zener transitions in oscillating fields. *Physics Letters A*, 140(9):507 – 512, 1989.
- [117] Clarence Zener. Non-adiabatic crossing of energy levels. *Proc. Roy. Soc. Lond. A*, 137(833):696–702, 1932.
- [118] L.Đ. Landau and E.Ā. Lifshitz. *Quantum Mechanics*. Pergamon, Oxford, 1976.
- [119] Phillip Weinberg, Marin Bukov, Luca D’Alessio, Anatoli Polkovnikov, Szabolcs Vajna, and Michael Kolodrubetz. Adiabatic perturbation theory and geometry of periodically-driven systems. *Physics Reports*, 688:1 – 35, 2017. Adiabatic Perturbation Theory and Geometry of Periodically-Driven Systems.
- [120] HP Breuer and M Holthaus. Adiabatic processes in the ionization of highly excited hydrogen atoms. *Zeitschrift für Physik D Atoms, Molecules and Clusters*, 11(1):1–14, 1989.
- [121] Yi-Fu Zhang, Yun-You Yang, Yan Ju, Li Sheng, Rui Shen, Dong-Ning Sheng, and Ding-Yu Xing. Coupling-matrix approach to the chern number calculation in disordered systems. *Chinese Physics B*, 22(11):117312, 2013.
- [122] M. Puschmann, P. Cain, M. Schreiber, and T. Vojta. Integer quantum Hall transition on a tight-binding lattice. *ArXiv e-prints*, May 2018.
- [123] Tomoki Ozawa and Iacopo Carusotto. Anomalous and quantum hall effects in lossy photonic lattices. *Phys. Rev. Lett.*, 112:133902, Apr 2014.
- [124] Tomoki Ozawa. Steady-state hall response and quantum geometry of driven-dissipative lattices. *Phys. Rev. B*, 97:041108, Jan 2018.
- [125] Matteo M. Wauters, Angelo Russomanno, Roberta Citro, Giuseppe E. Santoro, and Lorenzo Privitera. Localization, topology, and quantized transport in disordered floquet systems. *Phys. Rev. Lett.*, 123:266601, Dec 2019.
- [126] Chyh-Hong Chern, Shigeki Onoda, Shuichi Murakami, and Naoto Nagaosa. Quantum charge pumping and electric polarization in anderson insulators. *Phys. Rev. B*, 76:035334, Jul 2007.
- [127] Jihong Qin and Huaiming Guo. Quantum pumping induced by disorder in one dimension. *Physics Letters A*, 380(29):2317 – 2321, 2016.
- [128] Masaya Nakagawa, Tsuneya Yoshida, Robert Peters, and Norio Kawakami. Breakdown of topological thouless pumping in the strongly interacting regime. *Phys. Rev. B*, 98:115147, Sep 2018.
- [129] T. Haug, L. Amico, L.-C. Kwek, W.J. Munro, and V.M. Bastidas. Duality and topological pumping on entanglement and disorder. *arXiv preprint arXiv:1905.03807*, 2019.
- [130] Mark S Rudner, Netanel H Lindner, Erez Berg, and Michael Levin. Anomalous edge

- states and the bulk-edge correspondence for periodically driven two-dimensional systems. *Phys. Rev. X*, 3(3):031005, 2013.
- [131] Andrew C. Potter, Takahiro Morimoto, and Ashvin Vishwanath. Classification of interacting topological Floquet phases in one dimension. *Physical Review X*, 6(4):041001, 2016.
- [132] Dominic V. Else and Chetan Nayak. Classification of topological phases in periodically driven interacting systems. *Physical Review B*, 93(20):201103, 2016.
- [133] Oles Shtanko and Ramis Movassagh. Stability of periodically driven topological phases against disorder. *Phys. Rev. Lett.*, 121:126803, Sep 2018.
- [134] Netanel H Lindner, Erez Berg, and Mark S Rudner. Universal chiral quasisteady states in periodically driven many-body systems. *Phys. Rev. X*, 7(1):011018, 2017.
- [135] T. Gulden, E. Berg, M. S. Rudner, and N. H. Lindner. Exponentially long lifetime of universal quasi-steady states in topological floquet pumps. *arXiv preprint arXiv:1901.08385*, 2019.
- [136] Dmitry A. Abanin, Ehud Altman, Immanuel Bloch, and Maksym Serbyn. Many-body localization, thermalization, and entanglement. *arXiv*, page 1804.11065, 2018.
- [137] Fenner Harper, Rahul Roy, Mark S Rudner, and SL Sondhi. Topology and broken symmetry in floquet systems. *arXiv preprint arXiv:1905.01317*, 2019.
- [138] E. Abrahams, P. W. Anderson, D. C. Licciardello, and T. V. Ramakrishnan. Scaling theory of localization: Absence of quantum diffusion in two dimensions. *Phys. Rev. Lett.*, 42:673–676, Mar 1979.
- [139] J. T. Edwards and D. J. Thouless. Numerical studies of localization in disordered systems. *J. Phys. C*, 5:807, 1972.
- [140] Emil Prodan, Taylor L. Hughes, and B. Andrei Bernevig. Entanglement spectrum of a disordered topological chern insulator. *Phys. Rev. Lett.*, 105:115501, Sep 2010.
- [141] Emil Prodan. Disordered topological insulators: a non-commutative geometry perspective. *Journal of Physics A: Mathematical and Theoretical*, 44(11):113001, feb 2011.
- [142] Kun Yang and R. N. Bhatt. Floating of extended states and localization transition in a weak magnetic field. *Phys. Rev. Lett.*, 76:1316–1319, Feb 1996.
- [143] Ivar Martin, Gil Refael, and Bertrand Halperin. Topological frequency conversion in strongly driven quantum systems. *Phys. Rev. X*, 7:041008, Oct 2017.
- [144] Vedika Khemani, Rahul Nandkishore, and S. L. Sondhi. Nonlocal adiabatic response of a localized system to local manipulations. *Nature Physics*, 11:560, 2015.
- [145] Matteo Ippoliti and R. N. Bhatt. Dimensional crossover of the integer quantum hall plateau transition and disordered topological pumping. *arXiv:1905.13171*, May 2019.
- [146] Janos K. Asboth, Laszlo Oroszlany, and Andras Palyi. *A Short Course on Topological*

- Insulators*, volume 919 of *Lecture Notes in Physics*. Springer International Publishing, 2016.
- [147] Matteo M. Wauters, Glen Bigan Mbeng, and Giuseppe E. Santoro. Polynomial scaling of qaoa for ground-state preparation of the fully-connected p-spin ferromagnet. 2020.
- [148] Antonio Acín, Immanuel Bloch, Harry Buhrman, Tommaso Calarco, Christopher Eichler, Jens Eisert, Daniel Esteve, Nicolas Gisin, Steffen J Glaser, Fedor Jelezko, Stefan Kuhr, Maciej Lewenstein, Max F Riedel, Piet O Schmidt, Rob Thew, Andreas Wallraff, Ian Walmsley, and Frank K Wilhelm. The quantum technologies roadmap: a european community view. *New Journal of Physics*, 20(8):080201, aug 2018.
- [149] John Preskill. Quantum computing and the entanglement frontier. 2012.
- [150] Edward Farhi and Aram W. Harrow. Quantum Supremacy through the Quantum Approximate Optimization Algorithm. *arXiv e-prints*, Oct 2016.
- [151] 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 Val-lalonga, 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, 2019.
- [152] Bettina Heim, Troels F. Rønnow, Sergei V. Isakov, and Matthias Troyer. Quantum versus classical annealing of ising spin glasses. *Science*, 348(6231):215–217, 2015.
- [153] E. Crosson and A. W. Harrow. Simulated quantum annealing can be exponentially faster than classical simulated annealing. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 714–723, Oct 2016.
- [154] Vasil S. Denchev, Sergio Boixo, Sergei V. Isakov, Nan Ding, Ryan Babbush, Vadim Smelyanskiy, John Martinis, and Hartmut Neven. What is the computational value of finite-range tunneling? *Phys. Rev. X*, 6:031015, Aug 2016.
- [155] Tameem Albash and Daniel A. Lidar. Demonstration of a scaling advantage for a quantum annealer over simulated annealing. *Phys. Rev. X*, 8:031016, Jul 2018.
- [156] Yuki Susa, Yu Yamashiro, Masayuki Yamamoto, and Hidetoshi Nishimori. Exponential

- speedup of quantum annealing by inhomogeneous driving of the transverse field. *J. Phys. Soc. Jpn.*, 87(2):023002, 2018.
- [157] G. Passarelli, V. Cataudella, and P. Lucignano. Improving quantum annealing of the ferromagnetic  $p$ -spin model through pausing. *Phys. Rev. B*, 100:024302, Jul 2019.
- [158] Gianluca Passarelli, Ka-Wa Yip, Daniel A. Lidar, Hidetoshi Nishimori, and Procolo Lucignano. Reverse quantum annealing of the  $p$ -spin model with relaxation. *Phys. Rev. A*, 101:022331, Feb 2020.
- [159] G. Passarelli, V. Cataudella, R. Fazio, and P. Lucignano. Counterdiabatic driving in the quantum annealing of the  $p$ -spin model: A variational approach. *Phys. Rev. Research*, 2:013283, Mar 2020.
- [160] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O'brien. A variational eigenvalue solver on a photonic quantum processor. *Nature communications*, 5:4213, 2014.
- [161] C Kokail, C Maier, R van Bijnen, T Brydges, MK Joshi, P Jurcevic, CA Muschik, P Silvi, R Blatt, CF Roos, et al. Self-verifying variational quantum simulation of lattice models. *Nature*, 569(7756):355, 2019.
- [162] Zhihui Wang, Stuart Hadfield, Zhang Jiang, and Eleanor G. Rieffel. Quantum approximate optimization algorithm for maxcut: A fermionic view. *Phys. Rev. A*, 97:022304, Feb 2018.
- [163] Zhang Jiang, Eleanor G. Rieffel, and Zhihui Wang. Near-optimal quantum circuit for grover's unstructured search using a transverse field. *Phys. Rev. A*, 95:062317, Jun 2017.
- [164] Tommaso Caneva, Rosario Fazio, and Giuseppe E. Santoro. Adiabatic quantum dynamics of a random ising chain across its quantum critical point. *Phys. Rev. B*, 76:144427, Oct 2007.
- [165] Glen Bigan Mbeng, Luca Arceci, and Giuseppe E. Santoro. Optimal working point in digitized quantum annealing. *Phys. Rev. B*, 100:224201, Dec 2019.
- [166] Wen Wei Ho, Cheryne Jonay, and Timothy H. Hsieh. Ultrafast State Preparation via the Quantum Approximate Optimization Algorithm with Long Range Interactions. *arXiv e-prints*, page arXiv:1810.04817, Oct 2018.
- [167] H.J. Lipkin, N. Meshkov, and A.J. Glick. Validity of many-body approximation methods for a solvable model. *Nuclear Physics*, 62:188, Feb 1965.
- [168] Michael Streif and Martin Leib. Comparison of QAOA with Quantum and Simulated Annealing. *arXiv e-prints*, page arXiv:1901.01903, Jan 2019.
- [169] Jorge Nocedal and Stephen Wright. *Numerical optimization*. Springer Science & Business Media, 2006.
- [170] Matteo M. Wauters, Emanuele Panizon, Glen B. Mbeng, and Giuseppe E. Santoro.

- Reinforcement-learning-assisted quantum optimization. *Phys. Rev. Research*, 2:033446, Sep 2020.
- [171] Volodymyr Mnih, Adrià Puigdomènech Badia, Mehdi Mirza, Alex Graves, Timothy P. Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. 2016.
- [172] Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. Safe, multi-agent, reinforcement learning for autonomous driving. 2016.
- [173] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Praseem Goyal, Lawrence D. Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, Xin Zhang, Jake Zhao, and Karol Zieba. End to end learning for self-driving cars. 2016.
- [174] Shixiang Gu, Ethan Holly, Timothy Lillicrap, and Sergey Levine. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. 2016.
- [175] OpenAI, Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, Jonas Schneider, Nikolas Tezak, Jerry Tworek, Peter Welinder, Lilian Weng, Qiming Yuan, Wojciech Zaremba, and Lei Zhang. Solving rubik’s cube with a robot hand, 2019.
- [176] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359, 2017.
- [177] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dhharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018.
- [178] Andrew Lucas. Ising formulations of many np problems. *Frontiers in Physics*, 2:5, 2014.
- [179] V. Bapst, L. Foini, F. Krzakala, G. Semerjian, and F. Zamponi. The quantum adiabatic algorithm applied to random optimization problems: The quantum spin glass perspective. *Phys. Rep.*, 523:127–206, 2013.
- [180] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1):4812, 2018.
- [181] Leo Zhou, Sheng-Tao Wang, Soonwon Choi, Hannes Pichler, and Mikhail D. Lukin. Quantum Approximate Optimization Algorithm: Performance, Mechanism, and Implementation on Near-Term Devices. *arXiv e-prints*, page arXiv:1812.01041, Dec 2018.
- [182] Domenico D’Alessandro. *Introduction to quantum control and dynamics*. Chapman and Hall/CRC, 2007.

- [183] Jens Kober, J Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.
- [184] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [185] 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.
- [186] Thomas Fösel, Petru Tighineanu, Talitha Weiss, and Florian Marquardt. Reinforcement learning with neural networks for quantum feedback. *Phys. Rev. X*, 8:031084, Sep 2018.
- [187] Moritz August and JosÁ© Miguel HernÁ©ndez-Lobato. Taking gradients through experiments: Lstms and memory proximal policy optimization for black-box quantum control. 2018.
- [188] Sami Khairy, Ruslan Shaydulin, Lukasz Cincio, Yuri Alexeev, and Prasanna Balaprakash. Learning to optimize variational quantum circuits to solve combinatorial problems. 2019.
- [189] Artur Garcia-Saez and Jordi Riu. Quantum observables for continuous control of the quantum approximate optimization algorithm via reinforcement learning. 2019.
- [190] Jiahao Yao, Marin Bukov, and Lin Lin. Policy gradient based quantum approximate optimization algorithm. 2020.
- [191] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347, 2017.
- [192] Glen Bigan Mbeng, Rosario Fazio, and Giuseppe E. Santoro. Optimal quantum control with digitized quantum annealing. 2019.
- [193] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Leo Zhou. The quantum approximate optimization algorithm and the sherrington-kirkpatrick model at infinite size. 2019.
- [194] Joshua Achiam. Spinning Up in Deep Reinforcement Learning. 2018.
- [195] A. P. Young and H. Rieger. Numerical study of the random transverse-field Ising spin chain. *Phys. Rev. B*, 53(13):8486–8498, April 1996.
- [196] V. Vitale, G. De Filippis, A. de Candia, A. Tagliacozzo, V. Cataudella, and P. Lucignano. Assessing the quantumness of the annealing dynamics via leggett garg’s inequalities: a weak measurement approach. *Scientific Reports*, 9(1):13624, 201.
- [197] J. H. Shirley. Solution of Schrödinger equation with a hamiltonian periodic in time. *Phys. Rev.*, 138:B979, 1965.
- [198] N. G. Dickson et al. Thermally assisted quantum annealing of a 16-qubit problem. *Nat.*



*Commun.*, 4:1903, 2013.

- [199] Kostyantyn Kechedzhi and Vadim N. Smelyanskiy. Open-system quantum annealing in mean-field models with exponential degeneracy. *Phys. Rev. X*, 6:021028, May 2016.
- [200] Luca Arceci, Simone Barbarino, Rosario Fazio, and Giuseppe E. Santoro. Dissipative landau-zener problem and thermally assisted quantum annealing. *Phys. Rev. B*, 96:054301, Aug 2017.
- [201] Johann Carl Friedrich Gauss. *Disquisitiones arithmeticae*, volume 157. Yale University Press, 1966.
- [202] Wikipedia contributors. Multiplicative group of integers modulo  $n$  — Wikipedia, the free encyclopedia, 2019.