

Characterising Many-Body Systems via Quantum Dynamics Simulations



Lucia Vilchez Estevez
Balliol College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity 2025

A mamá y a papá,

Acknowledgements

I would like to begin by expressing my gratitude to my supervisor, Vlatko Vedral, for the freedom he has given me throughout my DPhil. His openness encouraged me to explore different research directions, follow my curiosity, and ultimately develop as an independent researcher.

Throughout my DPhil, I have been fortunate to receive academic support from many people. Above all, I am deeply grateful to my mentor and collaborator, Jinzhao Sun, from whom I have learned not only the practice of research but also much about the field of quantum computing. I also had the privilege of working with Phasecraft, an experience that offered me a unique glimpse into the world of industry while allowing me to pursue exciting scientific questions. I am especially grateful to Filippo Gambetta, Sabrina Wang, and Raul Santos for their guidance, encouragement, and help in navigating the many challenges of the project – particularly the battle against noisy results from hardware runs! Finally, I would like to thank Nicetu Tibau Vidal, whose advice and companionship made the start of my DPhil far less daunting.

Although not all of them were direct collaborators, many people have accompanied me throughout this journey and have been essential to the completion of this thesis. First and foremost, I want to thank everyone in my research group. I count myself very lucky to have spent these years surrounded by such brilliant, kind, and witty people. I am grateful to Abdulla, Aditya, Amy, Cosmin, Giuseppe, Jiaxuan, Lodovico, Maria, Mia, Nicetu, Sam, Simone, and Tom, with whom I shared countless lunch breaks and pub trips that became some of the highlights of my DPhil. A special thank you goes to Lodovico, with whom I began this journey and who has been there unfailingly through every high and low since then.

The people I have shared my life with outside the office in Oxford and across the UK have been just as important as the work itself. It is through them that I find fulfilment and belonging. To my friends Andy, Asang, Ben, Casey, Holly, Jack, Jenny, Joe, and Liv, thank you for all the evenings at the pub, the dinner parties, and the good times. To my housemates Carles, Kim, and Sergi, thanks for making this last year so much fun and for turning home into such a welcoming place to come back to. And to Hasan, thank you for all the amazing adventures and for being one of the main pillars of my time here.

My many visits to Granada throughout the year, and the long stays in Salobreña during the summer, have been a constant source of peace and comfort. No matter where I am, this will always be home, a place to return to. The people there are an essential part of that feeling. I am especially grateful to my friends, Álvaro, Ana, Chema, Gabri, Paloma, and Pedro, who have been by my side since the very beginning. With them, and with the rest of the Spanish crowd, I have shared the best moments of my life.

Finally, I want to thank my family for their endless love and support. Above all, I am deeply grateful to my parents, who built a place where I could be brave. None of this would have been possible without them, and for that, this thesis is theirs.

Abstract

Simulating quantum many-body systems is essential for understanding complex phenomena in condensed matter, chemistry, and materials science. Quantum computing offers a way to overcome the limitations of classical methods by operating directly in exponentially large Hilbert spaces. In particular, classical methods often fail to capture the dynamics of strongly correlated systems, which makes quantum algorithms a promising alternative for exploring these challenging regimes.

This thesis uses quantum dynamics simulations to probe the behaviour of many-body systems. We first introduce a *spectroscopy*-based protocol to extract excitation spectra from the time evolution of observables and analyse its performance on near-term quantum devices. We show that the protocol requires only short evolution times and is resilient to errors, facilitating its implementation on real hardware. Building on this, we simulate fermionic excitation spectra using current quantum devices, introducing resource optimisations – including an approximate free-fermion initial state, a unitary quench without ancillae, and reduced Trotter steps – that significantly lower circuit depth. These improvements enable simulations at system sizes competitive with the state of the art, yielding results that match classical benchmarks.

While most studies focus on unitary dynamics of isolated systems, open quantum dynamics is increasingly relevant for describing non-equilibrium processes across biology, materials science, and quantum control. We extend our scope to systems coupled to an environment and governed by Lindbladian evolution. By studying simulation complexity through physically meaningful quantities – such as correlation length and mixing time – we explore whether open-system dynamics can be as tractable as their unitary counterparts. These findings provide insights into the computational complexity landscape of open-system simulation and highlight regimes where quantum methods offer a clear advantage.

Contents

1	Introduction	1
1.1	Quantum simulation of quantum systems	4
1.2	Quantum dynamics simulations	8
1.3	Classical simulation of quantum systems	12
1.3.1	Matrix Product States and Operators	13
1.3.2	Tensor network algorithms	15
1.4	Thesis structure	18
1.5	Publications	20
2	Probing spectral features of quantum many-body systems	23
2.1	Introduction	24
2.2	Dynamical structure factor	26
2.3	Framework	27
2.4	Quench spectroscopy protocol	31
2.5	Error analysis and resource requirement	36
2.6	Error effect in transition energy estimation	39
2.7	Simulation results	41
2.7.1	Numerical simulations	41
2.7.2	Implementation on a quantum computer	43
2.8	Performance in the presence of noise	46
2.8.1	General stragedy	46
2.8.2	Error mitigation for global depolarising noise	48
2.8.3	Numerical tests	50
2.9	Discussion	51
2.A	Further derivation of the algorithmic error and resource requirement	55
3	Quench spectroscopy of fermonic systems	61
3.1	Introduction	62
3.2	Unitary quench spectroscopy	64
3.3	Classical benchmarking	68
3.3.1	Classical simulation	69
3.3.2	Initial state preparation	70

3.3.3	Trotterization	73
3.4	Quantum circuit implementation	75
3.5	Approximate free-fermion initial state preparation	80
3.5.1	Practical implementation	80
3.6	Results	83
3.6.1	Simulations with different protocols	83
3.6.2	Resource estimation	84
3.6.3	Quantum hardware results	87
3.7	Discussion	91
3.A	(all up)(all down) JW ordering	93
4	Quantum simulation of open system dynamics	97
4.1	Introduction	98
4.2	Methods	100
4.2.1	Formulation of open quantum dynamics	100
4.2.2	Simulating the open-system dynamics	102
4.2.3	Physical measures of quantum algorithmic complexity	105
4.3	Quantum spin chain system	107
4.4	Quantum runtime analysis	109
4.4.1	Time-dependent Lindbladian	113
4.5	Classical runtime analysis	116
4.6	Discussion	119
4.A	Vectorisation of the density matrix	123
4.B	Error bound of the evolution and the algorithmic complexity	124
4.C	Tensor network simulation results for chaotic systems	126
5	Conclusion and outlook	129
5.1	Summary of work	129
5.2	Future research directions	133
5.2.1	Two-dimensional simulations	133
5.2.2	Quantum simulation of radical pairs	136
	Bibliography	138

1

Introduction

The idea of utilising quantum mechanical systems to address computationally challenging problems was first proposed by Richard Feynman in 1981 [1]. In particular, he suggested using a controllable quantum platform, namely a *quantum computer*, to simulate other quantum systems of interest. Quantum systems exhibit a fundamental property: their Hilbert space grows exponentially with system size, making their simulation intractable for classical computers. In contrast, quantum computers are inherently suited to this task, as their computational states also reside in exponentially large Hilbert spaces. This advantage is particularly evident in fields such as chemistry, condensed matter physics, or materials science, where the aim is to simulate quantum many-body phenomena, including electronic or spin systems.

The first idea behind quantum computing was important because it showed that quantum information could be used for real-world applications. However, much remained to be formalised, and there was a need for theoretical evidence of problems that could genuinely benefit from this new paradigm. The next milestone came in 1985, when David Deutsch formalised the concept of a universal quantum computer for the first time [2], a machine capable of simulating any physical system and performing certain calculations faster than a classical Turing machine. Subsequently, the Deutsch-Jozsa algorithm in 1992 [3] provided the first explicit example of a quantum algorithm that is exponentially faster than any deterministic classical

algorithm. Although the problem it solves is of little practical use, the algorithm reveals the possibility of exponential quantum speed-ups.

However, it was Peter Shor’s algorithm in 1994 that demonstrated the first concrete practical application of quantum computing [4], with an efficient quantum method for factoring large integers. Shor’s algorithm has direct implications for cryptographic systems, notably threatening the security of RSA¹ encryption making it vulnerable to quantum attacks. Shortly afterwards, in 1996, Lov Grover introduced a quantum search algorithm [5] that achieved a quadratic speed-up over classical search methods. Grover’s algorithm allows for searching an unstructured database of N items in only $\mathcal{O}(\sqrt{N})$ steps, compared to the classical $\mathcal{O}(N)$ scaling, and applies to a broad class of problems beyond simple database lookup [6–9].

To run and study quantum algorithms, we require access to quantum computers with sufficient scale and quality. By scale we refer primarily to the number of available qubits and the maximum circuit depth that can be executed, while quality encompasses qubit coherence times, gate fidelities, and measurement accuracy. This requires addressing one of the main challenges in quantum computation: noise and decoherence. Due to interactions with the environment, quantum systems are highly susceptible to the loss of coherence, leading to errors in computation. Significant effort must be invested in engineering such computers, including increasing qubit counts and reducing error rates. At the same time, it is imperative to develop theoretical tools to implement quantum algorithms in a fault-tolerant manner and to understand the feasibility of different approaches in the presence of noise. The first error-correcting code to address decoherence was developed by Shor in 1995 [10]. Subsequently, Alexei Kitaev [11] proposed using topological features of matter to protect quantum computation against decoherence. This idea has led to one of the leading approaches for realising a fully fault-tolerant quantum computer.

Note that although algorithms like Shor’s and Grover’s are proven to offer quantum speed-ups over classical algorithms, resource estimates for practical

¹RSA (Rivest, Shamir, and Adleman, 1977) is a public-key cryptographic system used for both encryption and digital signatures. Its security relies on the assumed classical hardness of factoring large integers.

running times reveal that they would require an unreasonable number of physical qubits for fault-tolerant implementation. Demonstrating quantum advantage with these algorithms is therefore very unlikely in the near future. Despite significant advances in the experimental realisation of quantum computers, the requirements for implementing practical quantum algorithms remain far beyond what we can expect to achieve in the coming decades. Current state-of-the-art processors comprise on the order of 10^3 noisy physical qubits [12, 13], with two-qubit gate fidelities typically between 99 and 99.9% and feasible circuit depths limited to a few hundred operations before errors accumulate [14]. In contrast, large-scale fault-tolerant implementations – such as Shor’s algorithm – are estimated to require millions of physical qubits to encode thousands of logical qubits, with logical error rates suppressed to around 10^{-15} over trillions of operations [15]. For this reason, it is imperative to reduce the resource costs of quantum algorithms and to explore new techniques and problems where quantum computers could offer a clear advantage.

There are significant research efforts aimed at running quantum algorithms on current hardware, known as Noisy Intermediate-Scale Quantum (NISQ) devices. These devices are characterised by limited qubit counts, noisy gate operations, and cannot yet do full error correction. The first claim of quantum supremacy² came from Google’s 2019 experiment [17], in which they argued that their quantum processor completed a calculation that would purportedly take classical supercomputers approximately 10,000 years. This claim was later contested in Ref. [18], where the same calculation was performed on a classical supercomputer in just a few days. This situation highlights a common trend in quantum computing: initial claims of quantum advantage often motivate improvements in classical methods which can reduce or even close the gap for certain tasks. In any case, the circuit executed by Google had no significant practical application beyond showcasing the capabilities of their quantum hardware. Although further attempts have been made to demonstrate a quantum speed-up on NISQ devices [19], these efforts have generally focused on

²Quantum supremacy is a term introduced by John Preskill [16] to denote the point at which a quantum device can perform a computation that cannot be completed by a classical processor in a reasonable time.

tasks with little or no physical relevance. A pressing question remains whether there exist problems of practical interest that can be run on NISQ hardware. For instance, simulating quantum properties of microscopic particles, which is highly relevant in fields such as materials science, high-energy physics, or quantum chemistry.

1.1 Quantum simulation of quantum systems

As discussed in the previous section, many research teams are actively seeking quantum advantage, but current achievements in this direction often appear to be tailored to highlight the strengths of specific quantum technologies rather than addressing practical applications directly. This raises an important challenge: applying quantum devices to genuinely useful problems, such as simulating quantum mechanical systems [20, 21].

Exact solutions, such as determining the ground state energy or full dynamical evolution of complex quantum systems, generally remain out of reach for classical approaches due to the enormous quantum information processing required. However, highly effective approximate classical methods have been developed. Density Functional Theory (DFT) [22], for instance, is based on the idea that the properties of a chemical system can be determined from its electron density, rather than its wavefunction. On the other hand, tensor network methods, such as the Density Matrix Renormalization Group (DMRG) [23, 24], capture quantum correlations by compressing the quantum state into a manageable representation. These methods are especially powerful when the system exhibits limited quantum entanglement, for example, when computing low-energy eigenstates or near-equilibrium dynamics of one-dimensional systems with local interactions.

Despite the effectiveness of these classical methods, their approximations can limit accuracy, especially for highly correlated systems. Moreover, when simulating more complex scenarios such as non-equilibrium dynamics and general time evolution, they often suffer from exponentially growing computational costs. Quantum hardware, conversely, naturally encodes quantum correlations and therefore holds promise for simulating such complex scenarios without the same exponential scaling

in memory. Near-term quantum hardware is believed to be capable of simulating simplified models of solid-state materials [25], making some practical quantum simulations within reach today. Quantum chemistry simulations, however, remain more challenging since their computational complexity often scales unfavorably with the number of electron orbitals. Nonetheless, with advancements in quantum hardware, researchers anticipate that these more demanding simulations will also become feasible.

While classical techniques have enabled major progress, key challenges in quantum many-body physics remain unresolved, particularly for systems with strong correlations or far-from-equilibrium dynamics. Quantum simulation offers a promising complementary path, providing new tools to explore phenomena beyond classical reach, guide experimental discovery, and benchmark or refine classical methods as the field advances.

There are two primary approaches to simulating quantum systems with quantum computers: analog and digital quantum simulation. Analog quantum simulation is akin to building a physical wind tunnel to study aerodynamics. In this approach, a controllable quantum system is engineered to directly mimic the dynamics of a target system. Analog simulators can reach relatively large system sizes and have already provided important insights in specific areas [26–28]. However, they are limited in scope, their applicability is typically restricted to problems that closely match the available hardware. The main quantitative limitation is the calibration accuracy of the device, and, crucially, there is no straightforward way to apply error correction. As a result, analog simulators are not broadly programmable or universally applicable.

On the other hand, digital quantum simulation is analogous to modern digital computing. Here, a discrete set of quantum gates is used to manipulate quantum information, allowing for the simulation of a wide variety of quantum systems and algorithms. Digital quantum computers are general-purpose machines [29], with the key theoretical advantage that error correction can be implemented to

counteract gate imperfections. This opens the path to eventually building fault-tolerant quantum computers, capable of reliable large-scale simulations. Currently, we are still in the NISQ era, where the number of qubits is limited and noise remains a significant obstacle. Because of this, early demonstrations of quantum advantage have typically involved problems tailored to the strengths and limitations of specific hardware platforms. Achieving practical quantum advantage for real scientific problems with NISQ devices remains challenging, as noise and error rates limit both the depth and accuracy of implementable algorithms. While some researchers argue that true quantum advantage will only be realised with fault-tolerant devices, there is a substantial effort in the community to explore how near-term quantum computers might still tackle meaningful problems, potentially outperforming classical algorithms, or at least prompting further advances in classical computational methods. Given these considerations, this thesis will focus primarily on digital quantum simulation as a general and programmable framework for quantum computation.

Previously, we have highlighted some of the key applications of quantum computers in calculating properties of quantum systems. Broadly, these applications can be categorised into two main types of problems: static and dynamic. The static problem involves calculating the lowest energy eigenstate (ground state) or specific excited states of a quantum system. The dynamic problem concerns the time evolution of the system.

The static problem can be formally defined as finding the state $|\psi_i\rangle$ such that

$$H |\psi_i\rangle = E_i |\psi_i\rangle, \quad (1.1)$$

where H is the Hamiltonian that governs the system, and E_i are the energy eigenvalues associated with each state, for $i = \{0, 1, \dots, 2^N - 1\}$, where N is the size of the system (i.e., the number of particles). Typically, the static problem is considered to be more manageable by classical computation, particularly for systems with local interactions. In these cases, the relevant space of physical states is much smaller than the full Hilbert space, allowing for compression and efficient

sampling of only the relevant states. As a result, a wide range of sophisticated classical algorithms has been developed across physics, chemistry, and materials science, setting a high and constantly moving threshold for quantum advantage. Nonetheless, classical simulation remains challenging for certain strongly correlated models, such as the Hubbard model [30, 31].

In this thesis, we are interested in exploiting quantum computers to simulate the dynamics of quantum many-body systems. The non-equilibrium dynamics (like the ones in quantum quench experiments that we will explore in Chapter 2 and 3) consist of initialising the quantum states and then letting them evolve with time. The dynamics of an *isolated* system – represented by a pure state $|\psi(t)\rangle$ – can be given by the time-dependent Schrödinger equation

$$\frac{d|\psi(t)\rangle}{dt} = -iH|\psi(t)\rangle, \quad (1.2)$$

where we set $\hbar = 1$. In this case, the time evolution enables the exploration of an exponentially growing (with system size) number of states, which is, in many cases, intractable for classical methods. For a closed system following the Schrödinger equation, the solution at time t is given by the unitary time evolution operator applied to the initial state $|\psi(0)\rangle$:

$$|\psi(t)\rangle = e^{-iHt}|\psi(0)\rangle. \quad (1.3)$$

On another hand, quantum systems do not need to be isolated but can be embedded in an environment, the so-called *open quantum systems*. Instead of describing them with pure states, we now need a density matrix ρ to characterise our mixed states:

$$\rho(t) = \sum_i p_i(t) |\psi_i(t)\rangle \langle \psi_i(t)|, \quad (1.4)$$

which can be interpreted as a mixture of pure states $|\psi_i(t)\rangle$ with probabilities $p_i(t)$. The evolution of a quantum system interacting with a bath (in the Markovian³

³Markovian evolution lacks memory effects, meaning that the system's dynamics depend only on its current state, not on its history. This is typically modelled by assuming weak coupling to the environment and fast environmental relaxation.

case) is described by the Lindblad master equation

$$\frac{d\rho}{dt} = -i[H, \rho] + \mathcal{L}[\rho], \quad (1.5)$$

where $\mathcal{L}[\rho] = \sum_k \gamma_k \left(L_k \rho L_k^\dagger - \frac{1}{2} L_k^\dagger L_k \rho - \frac{1}{2} \rho L_k^\dagger L_k \right)$ is the super-operator describing the coupling with the environment mediated by the operators L_k with coefficients γ_k . These types of systems obviously add a different challenge since they have a non-Hermitian part in the evolution⁴ that we need to handle carefully. To simulate this with a quantum computer, we need to run what is called an imaginary time evolution, where the propagator e^{-Ht} is non-unitary and hence non-norm conserving. The general strategy is to combine the action of the propagator with the normalisation of the state to approximate the open quantum dynamics [32]. We will develop this further in Chapter 4, where we explore the dynamics of correlations in open quantum systems by simulating their Lindblad evolution.

1.2 Quantum dynamics simulations

Since Lloyd's first proposal of a universal quantum algorithm for simulating quantum dynamics [29], there have been significant advances in developing algorithms to use quantum computers for simulating time evolution. These are commonly referred to as Hamiltonian simulation algorithms. For real-time evolution, the task is to efficiently approximate the time evolution operator $U = e^{-iHt}$. In most physical systems, and for the algorithms discussed below, we typically assume that the Hamiltonian can be written as a linear combination of Pauli strings H_j (that is, tensor products of Pauli matrices):

$$H = \sum_{j=1}^L \alpha_j H_j, \quad (1.6)$$

where α_j are real coefficients and L is the number of Pauli strings in the decomposition. The goal is to implement $\tilde{U}$ with an error no greater than ϵ ,

$$|e^{-iHt} - \tilde{U}| \leq \epsilon. \quad (1.7)$$

⁴Strictly speaking, the total Hamiltonian of the system plus environment is Hermitian. The non-Hermitian contribution arises only in the effective description of the reduced system, e.g., in the effective Hamiltonian $H_{\text{eff}} = H_1 - iH_2$, as we will see in Chapter 4.

While the focus here is on real-time evolution, many of the methods discussed can also be adapted to simulate open quantum systems. This typically involves mapping the problem via the substitution $it \rightarrow \tau$, resulting in the imaginary-time Schrödinger equation and a non-unitary propagator $U = e^{-H\tau}$, which requires careful handling of normalisation. We will return to the adaptation of these algorithms for open quantum dynamics at the end of this section.

The first Hamiltonian simulation algorithm [29] uses a product formula approximation, also known as *Trotterization* [33, 34], to approximate the exponential of a sum of operators with a product of exponentials of individual operators. The first-order approximation is given by

$$\tilde{U}_1(t) := \left(\prod_{j=1}^L e^{\frac{-it}{r} \alpha_j H_j} \right)^r \quad (1.8)$$

where r is the number of Trotter steps. The error in the first-order product formula is bounded by

$$\epsilon := \|U(t) - \tilde{U}_1(t)\| = \mathcal{O} \left(\frac{(L\Lambda t)^2}{r} \right), \quad (1.9)$$

where $\Lambda = \max_j \alpha_j$. More generally, the error can be bounded in terms of the commutator structure of the Hamiltonian:

$$\epsilon = \mathcal{O} \left(\sum_{i < j}^L \| [H_i, H_j] \| \frac{(\Lambda t)^2}{r} \right). \quad (1.10)$$

The error can be reduced either by increasing the number of Trotter steps, thus increasing the gate complexity and depth of the circuit, or by using a higher-order Trotter decomposition. In particular, the second-order formula is given by

$$\tilde{U}_2(t) := \left(\prod_{j=1}^L e^{\frac{-it}{2r} \alpha_j H_j} \prod_{k=L}^1 e^{\frac{-it}{2r} \alpha_k H_k} \right)^r, \quad (1.11)$$

with an error upper bound

$$\epsilon := \|U(t) - \tilde{U}_2(t)\| = \mathcal{O} \left(\frac{(L\Lambda t)^3}{r^2} \right). \quad (1.12)$$

More generally, the complexity of the simulation can be improved by using a $(2k)$ th-order product formula, which is defined recursively. Using this higher-order

formula, the error scales as $\mathcal{O}\left(\frac{(L\Lambda t)^{2k+1}}{r^{2k}}\right)$, but the gate count increases exponentially with k , complicating the choice of the optimal Trotter order. The scaling of gate complexity⁵, in terms of the number of steps needed to achieve a given error, is $\mathcal{O}\left(\frac{(t\Lambda)^{1+1/2k}}{\epsilon^{1/2k}}\right)$. One of the main advantages of Trotterization is that, for most physical problems, the dependence of the error on the commutator structure of the Hamiltonian is very favourable, often reducing the scaling with the number of terms L in practice [35]. This improvement has been demonstrated in many practical applications, where the algorithm performs better than what is predicted by the formal worst-case bounds [36–38].

Because of their applicability and simplicity, there have been a lot of efforts to lower the gate count and improve the error scaling of the product formula methods. For instance, randomly permuting the order of the product formula at each step can reduce the error scaling [39]. In this same direction, the quantum stochastic drift (qDRIFT) protocol [40, 41] randomly samples individual hamiltonian terms biased by their weights and can achieve a gate count independent of L .

There exist alternative approaches to product formulas whose performance can be asymptotically more efficient. However, many of these methods come with trade-offs: they often require additional ancillary qubits, and may not take advantage of the commutator structure of the Hamiltonian. As a result, they may not be optimal for simulating quantum systems on near-term devices with limited resources. One of the main post-Trotter approaches is the linear combination of unitaries (LCU), also known as the Taylor expansion method [42–44]. In this method, the propagator is approximated by a truncated Taylor expansion:

$$\tilde{U}(t) := \sum_{k=0}^K \frac{(-itH)^k}{k!} = \sum_{k=0}^K \sum_{j_1 \dots j_k=1}^L \frac{t^k}{k!} \alpha_{j_1} \dots \alpha_{j_k} (-1)^k H_{j_1} \dots H_{j_k}, \quad (1.13)$$

where the Taylor series is expressed as a sum of unitaries. To implement this on a quantum computer, techniques such as amplitude amplification are commonly

⁵The gate complexity (or total runtime) refers to the total number of elementary gate operations required to implement a quantum algorithm, which directly relates to the overall circuit depth. In the case of product formulas, this corresponds to the number of Trotter steps required to approximate the time-evolution operator.

used [44]. The complexity of the LCU method is reduced to $\mathcal{O}\left(\tilde{\Lambda}t \frac{\log(\tilde{\Lambda}t/\epsilon)}{\log \log(\tilde{\Lambda}t/\epsilon)}\right)$, at the cost of introducing ancillary qubits. The number of extra qubits required scales linearly with the truncation order K and with $\log(L)$, which can be a significant drawback for near-term devices.

Moreover, advanced methods such as quantum signal processing (QSP) [45, 46] and qubitization [47] employ quantum walks [48] for Hamiltonian simulation. QSP achieves optimal, additive query complexity $\mathcal{O}\left(\tilde{\Lambda}t + \frac{\log(\epsilon^{-1})}{\log \log(\epsilon^{-1})}\right)$ but requires doubling the number of qubits due to the need to block-encode the Hamiltonian in a larger Hilbert space. Qubitization further improves on this approach by block-encoding the Hamiltonian in a way that maintains the optimal complexity while reducing the ancillary qubit overhead to a constant.

In this thesis, we mainly focus on Trotter-like methods for Hamiltonian simulation, motivated by their unique advantages in practical scenarios. For example, for geometrically local Hamiltonians – those in which each degree of freedom only interacts with nearby sites, as in most lattice models of condensed matter systems – Trotter formulas remain among the most efficient available methods. Their conceptual simplicity, flexibility, and lower constant overhead make them particularly attractive. This simplicity enables straightforward compilation into quantum circuits and makes them well-suited to near-term devices, where coherence times and qubit counts are severely limited. In particular, in Chapter 2 and 3, we implement our algorithms on current hardware, requiring careful resource management to obtain meaningful results. Furthermore, Trotter methods exploit the nested commutator structure of the Hamiltonian, enabling us to make use of commuting or nearly commuting terms, a key feature that is missing in other approaches. While more sophisticated techniques may scale better in some cases and could become preferable as quantum hardware improves, Trotter-based methods offer a reliable and practical option for current experiments and near-term use.

As we have discussed throughout this chapter, many interesting problems in quantum many-body physics extend beyond closed system dynamics to include the evolution of open quantum systems, where interactions with an environment

play a central role. Indeed, this will be the focus of Chapter 4 of this thesis. Moreover, as briefly mentioned earlier, unitary quantum simulation techniques can be adapted to tackle not only real-time (unitary) evolution but also imaginary time evolution [49], which is widely used for ground-state preparation and thermal-state approximation [32]. In particular, methods such as product formulas [50, 51] and Taylor series [52, 53] expansions can be repurposed to approximate open dynamics.

Open dynamics can be explicitly simulated by adding additional degrees of freedom that account for the interaction with the environment. Since the dynamics are described by a master equation for the density matrix (see Equation (1.5)), the use of additional qubits is necessary to obtain a suitable algorithmic representation. This can be achieved via dilation theorems [54, 55], which embed the open dynamics into a larger unitary evolution, or through vectorisation of the density matrix [56, 57], which reformulates the problem in terms of a larger state vector amenable to quantum simulation. The latter approach will be explored in Chapter 4, where we discuss both quantum and classical methods for simulating Lindbladian dynamics.

1.3 Classical simulation of quantum systems

The Hilbert space describing a quantum system grows exponentially with the number of constituent particles, resulting in an intractable computational challenge for classical computers. For instance, representing the full quantum state of a spin system with around 50 particles would already require more memory than is available on any existing supercomputer. As a result, exact classical simulations of generic quantum systems are restricted to very small system sizes. To address this limitation, classical simulation methods employ strategies to compress or sample the relevant regions of the Hilbert space. While this is generally impossible for arbitrary quantum systems, efficient simulation is possible for specific classes of systems that possess particular structure or limited entanglement. The most advanced approaches in this context are based on *tensor network* techniques, which offer a highly efficient representation for a broad class of quantum many-body states.

In this thesis, tensor network algorithms are employed as a central tool for the study, development, and benchmarking of quantum algorithms, including the spectroscopy protocol introduced in Chapter 2 and explicitly implemented for fermionic systems in Chapter 3. These techniques enable the simulation of small systems or short-time evolutions, where classical resources are still sufficient. In addition to providing insight into the conditions under which quantum protocols may excel when realised on quantum hardware, these simulations help to identify the boundaries of classical simulability (see Chapter 4 for a detailed discussion in the context of open quantum systems). In some cases, this can reveal promising avenues for future quantum advantage experiments, especially as quantum hardware improves and larger systems become accessible, a topic explored further in the context of fermionic excitation dynamics in Chapter 3. This chapter provides a brief overview of the classical simulation methods used throughout this thesis, with particular emphasis on one-dimensional tensor network architectures, which underpin the majority of the numerical results presented.

1.3.1 Matrix Product States and Operators

A pure quantum state $|\psi\rangle$ of a system of N sites, each of them with finite physical dimension d , can be written as

$$|\psi\rangle = \sum_{\{i_1 \dots i_N\}=0}^{d-1} c_{i_1 \dots i_N} |i_1\rangle \dots |i_N\rangle, \quad (1.14)$$

i.e. an N -rank tensor, of dimension d^N , where the number of coefficients $c_{i_1 \dots i_N}$ grows exponentially with system size N . The Matrix Product State (MPS) representation can be derived naturally by recursively applying the singular value decomposition (SVD) to the coefficient matrix. At each bipartition of the system, we obtain the Schmidt decomposition, i.e.

$$|\psi\rangle = \sum_{\alpha=1}^D \lambda_{\alpha} |\phi_{\alpha}^{(1 \dots k)}\rangle \otimes |\phi_{\alpha}^{(k \dots N)}\rangle, \quad (1.15)$$

where λ_{α} are the Schmidt coefficients and ϕ_{α} are the orthonormal Schmidt vectors. We apply this SVD successively, starting from the left. The first decomposition gives




Figure 1.1: Tensor network representations. **a** Matrix product state representation with bond dimension χ of a quantum state ψ with physical dimension d . **b** Matrix product operator representation with bond dimension χ of O that also has physical dimension d .

a coefficient tensor shaped as a $d \times d^{N-1}$ matrix. The SVD of such a decomposition can give at most d non-vanishing singular values because the rank is bounded by the smallest dimension of the matrix. The first MPS tensor has dimension $d \times D_1$ with $D_1 \leq d$. Next, we reshaped the tensor again to have dimension $D_1 d \times d^{N-2}$ before performing SVD. Then we will give a maximum number of singular values of $D_2 \leq d^2$. By successively decomposing the state across each site, we arrive at the MPS form where the coefficients can be factorised as a product of tensors:

$$c_{i_1 \dots i_N} = \sum_{\{\alpha\}} A_{\alpha_1}^{(1)i_1} A_{\alpha_1 \alpha_2}^{(2)i_2} \dots A_{\alpha_{N-1}}^{(N)i_N}, \quad (1.16)$$

where each tensor $A^{(k)}$ is the new basis from the SVD decomposition. The bond dimension across each cut is the Schmidt rank of the bipartition across the corresponding bond, with the largest possible value corresponding to the bond closest to the centre, so that $D \leq d^{N/2}$.

As illustrated in Figure 1.1 (a), a quantum state can be represented as a chain of connected tensors in the MPS formalism. The bond dimension D determines the amount of information retained in the MPS, and acts as a measure of the entanglement that can be captured between subsystems. In practical simulations, it is common to introduce a bond dimension cut-off χ , which allows the state to be represented more efficiently by truncating less significant Schmidt coefficients. This truncation leads to an approximate, but much more compact, representation

of the state. The amount of entanglement present in the system is reflected in the required bond dimension: $\chi = 1$ corresponds to a product state (no entanglement), while $\chi = d^{N/2}$ is sufficient to describe any state in the full Hilbert space of N sites with physical dimension d .

The MPS formalism can be extended to mixed states and operators through the Matrix Product Operator (MPO) representation, where each site is associated with a rank-4 tensor. This is visualised analogously to the MPS, with two physical legs (input and output) as shown in Figure 1.1 (b). Expectation values, traces and other physical observables can be efficiently computed by contracting the relevant tensor networks. In diagrammatic notation, this involves connecting the corresponding physical and bond indices of the MPS and MPO tensors, resulting in a network whose full contraction yields the desired quantity.

1.3.2 Tensor network algorithms

Most simulations presented in this thesis involve time-evolving quantum many-body systems. A common prerequisite for such simulations is the preparation of the system in a particular initial state, often the ground state of a model Hamiltonian. For example, as discussed in Chapter 3, we require the ground state of the Fermi-Hubbard model as the starting point for dynamical studies. While exact diagonalisation methods can provide this for small systems, they quickly become intractable for larger system sizes due to the exponential growth of the Hilbert space. To overcome this limitation, we employ the Density Matrix Renormalization Group (DMRG), a variational numerical algorithm designed to efficiently find ground states of one-dimensional systems with local interactions. The DMRG method, originally developed by Steven White in the early 1990s [58], has since become a standard tool for studying strongly correlated quantum systems.

The central idea behind DMRG is to iteratively optimise each tensor in the MPS representation of the wavefunction in order to minimise the energy $\langle \psi | H | \psi \rangle$, subject to a normalisation constraint $\langle \psi | \psi \rangle = 1$. The system's Hamiltonian is represented as an MPO acting on N sites. The optimisation at each site i is carried

out by minimising the cost function $\langle \psi | H | \psi \rangle - \lambda \langle \psi | \psi \rangle$, where λ is a Lagrange multiplier enforcing normalisation. The optimisation condition is given by

$$\frac{\partial}{\partial A^{(i)\dagger}} (\langle \psi | H | \psi \rangle - \lambda \langle \psi | \psi \rangle) = 0, \quad (1.17)$$

where $A^{(i)}$ is the local tensor at site i . Diagrammatically, taking the derivative with respect to $A^{(i)\dagger}$ corresponds to "punching a hole" at the position of the tensor. This leads to a generalised eigenvalue problem for $A^{(i)}$, where the optimal solution is the eigenvector corresponding to the lowest eigenvalue, which can be efficiently obtained using iterative algorithms such as the Lanczos method [59].

The DMRG algorithm proceeds as follows: starting from an initial guess for the MPS (often random), the tensors are updated sequentially site by site, with the rest of the MPS held fixed at each step. The initial MPS has a maximum bond dimension χ , which determines the accuracy and computational cost of the simulation. The procedure is repeated by sweeping back and forth along the chain until convergence is achieved. For a more detailed description of the algorithm and its implementation, the reader is referred to comprehensive reviews on tensor network methods, such as Ref. [60]

The following key step in simulating quantum states with tensor networks is the application of time evolution. A widely used method for this purpose is the Time-Evolving Block Decimation (TEBD) algorithm. This approach relies on the Trotter decomposition, as introduced in Section 1.2. TEBD remains particularly popular for Hamiltonians with nearest-neighbour interactions due to its simplicity and efficiency. The core idea is to approximate the unitary time-evolution operator $U = e^{-iHt}$, where the Hamiltonian can be expressed as a sum of at most two-local terms,

$$H = \sum_i H_i^{(1)} + \sum_i H_{i,i+1}^{(2)}, \quad (1.18)$$

for a one-dimensional chain with N sites. The Hamiltonian can be further partitioned




Figure 1.2: Diagram representation of a single step in the TEBD algorithm. The even (F_i) and odd (G_i) terms of the Hamiltonian are applied consecutively to the matrix product state (MPS) representing the initial state (shown in blue). The entire diagram corresponds to one Trotter step of duration t/r .

into terms acting on even and odd bonds, i.e.,

$$H = H_{\text{even}} + H_{\text{odd}}, \quad (1.19)$$

$$H_{\text{even}} = \sum_{i \text{ even}} F_i = \sum_{i \text{ even}} \left(H_i^{(1)} + H_{i,i+1}^{(2)} \right), \quad (1.20)$$

$$H_{\text{odd}} = \sum_{i \text{ odd}} G_i = \sum_{i \text{ odd}} \left(H_i^{(1)} + H_{i,i+1}^{(2)} \right). \quad (1.21)$$

The individual even and odd terms commute among themselves, that is, $[F_i, F_j] = [G_i, G_j] = 0$, but in general $[H_{\text{even}}, H_{\text{odd}}] \neq 0$. The Trotter decomposition then allows us to approximate the full time-evolution operator using the first-order (Lie-Trotter) product formula as

$$e^{-i(H_{\text{even}}+H_{\text{odd}})t} \approx \left(e^{-iH_{\text{even}}t/r} e^{-iH_{\text{odd}}t/r} \right)^r, \quad (1.22)$$

where r is the number of time steps. Importantly, since the terms in each exponential $e^{-iH_{\text{p}}t/r}$ commute, they can be applied simultaneously, as illustrated in Figure 1.2.

Writing $e^{-iF_it/r}$ or $e^{-iG_it/r}$ in MPO form results in identity tensors with bond dimension $\chi = 1$ everywhere except at sites i and $i + 1$, where the non-trivial two-site gate is applied. After each application of a TEBD layer, a singular value decomposition (SVD) is performed, with the state truncated to a maximum bond dimension χ to control computational resources. All these operations can be carried out efficiently, but there are two primary sources of error. The first is the Trotterization error, as quantified in Equation (1.9). The more significant limitation for real-time evolution is the truncation error that arises from maintaining a fixed

bond dimension. While MPS are known to provide accurate approximations for ground states of local, gapped Hamiltonians (as discussed previously for DMRG), this is generally not the case for time-evolved states, where the entanglement entropy typically grows linearly with time. As a result, simulating long-time dynamics incurs an exponential computational cost. Although several strategies exist to extend simulation times [61], this remains a fundamental challenge. This limitation underscores the potential of quantum computers to simulate quantum dynamics that surpass the capabilities of classical tensor network methods.

In any case, there is an ongoing effort to develop more powerful and general classes of classical simulation methods to compete with the capabilities of quantum hardware. One such approach, which is not restricted to nearest-neighbour interactions, is the Time-Dependent Variational Principle (TDVP) [62, 63]. TDVP offers an alternative to the Trotter decomposition employed in TEBD, with the primary goal of constraining the time evolution generated by the Schrödinger equation to a specified manifold of MPS with a fixed bond dimension. Due to its flexibility and effectiveness, this algorithm is also employed in certain parts of this thesis. However, the technical details of its implementation are omitted here, as they do not significantly contribute to the main discussion; for a comprehensive account, we refer the reader to Ref. [63].

1.4 Thesis structure

The development of practical quantum algorithms requires constant progress in adapting methods to the limitations of current hardware. Although large-scale, fault-tolerant quantum computers remain out of reach, ongoing work on near-term algorithms is essential: it both advances our ability to run meaningful computations on existing devices and drives the refinement of classical techniques in parallel. The work in this thesis contributes to this direction by aiming to simulate the dynamics of quantum systems to characterise their behaviour. Specifically, Chapter 2 and 3 focus on extracting spectral information from quantum many-body systems, while Chapter 4 explores correlation phenomena in open quantum systems.

Starting with Chapter 2, we first introduce a spectroscopy-based protocol designed to extract energy differences and excited state energies from observable dynamics. The protocol involves sampling the system dynamics following an initial perturbation of the state, known as a quench, at specific times determined by a probability distribution. From the sampled data, frequency peaks emerge, allowing us to identify energy transitions within the studied system. We then analyse the maximum time complexity of the protocol, highlighting its logarithmic dependence on the inverse of the desired precision. This is particularly advantageous as it implies that short-time evolutions and relatively shallow quantum circuits are sufficient, making it practical for implementation on quantum hardware. Following the theoretical introduction, we present classical simulations to illustrate how spectral accuracy scales with the evolution time, exploring various quantum systems, including spin lattices and molecular models. Additionally, we showcase results from implementing the protocol on actual quantum devices, demonstrating practical feasibility for small-scale systems. Importantly, we also assess the protocol's robustness by evaluating its performance under realistic noise conditions, critical for applications on near-term quantum devices. We find that employing simple error mitigation techniques can significantly enhance accuracy without considerable overhead. We support these findings through simulations designed to quantify the efficiency of these error mitigation strategies.

Chapter 3 builds upon this foundation by applying the spectroscopy-based approach to extract spectral information from dynamics in the Fermi-Hubbard model, focusing specifically on the fermionic excitation spectrum, namely spinons. We demonstrate the protocol experimentally on quantum hardware with a system size of 15 fermions (30 qubits). This simulation scale is competitive with current state-of-the-art quantum hardware dynamics simulations. We introduce an improved variant called unitary quench spectroscopy, which carefully selects a unitary perturbation which is applied to the initial state. This unitary perturbation significantly simplifies implementation in unitary-based computations and enables the exact recovery of the dynamical structure factor (DSF), directly related to the excitation spectrum.

Additionally, we conduct classical simulations to investigate optimal initial state choices and the effects of Trotterization approximation. We also detail a quantum circuit implementation, providing resource estimations achieved by optimal encoding of the fermionic system into qubits. Finally, we present experimental results from hardware runs, successfully recovering the spinon spectrum with remarkable accuracy, even when compared to classical estimates, achieved through minimal error mitigation.

In Chapter 4, we investigate the complexity of simulating open quantum systems (OQS), focusing on how spatial and temporal resources scale with physically meaningful parameters such as correlation length and mixing time. We approach this question through both classical and quantum simulations. On the classical side, we employ tensor network methods to quantify simulation cost via bond dimension and operator entanglement. On the quantum side, we simulate local observables governed by both time-dependent and nonlinear Lindbladian master equations, and use the average correlation length as a proxy for algorithmic complexity. We show that strong dissipation suppresses correlations, making long-time dynamics easier to simulate quantum mechanically, while classical resources grow super-polynomially. These results highlight regimes in which quantum simulations offer a potential advantage over classical methods.

Finally, in Chapter 5 we discuss our results and suggest possible future research directions.

1.5 Publications

Many of the results in this thesis have been published in the works listed below. In each case, the author’s specific contributions are detailed at the beginning of each chapter.

- Chapter 2 is based on work presented in J. Sun, L. Vilchez-Estevez, V. Vedral, A. T. Boothroyd, and M. S. Kim, “Probing spectral features of quantum many-body systems with quantum simulators”, *Nature Communications* **16**, 1403, 2025.

- Chapter 3 is based on work presented in
L. Vilchez-Estevez, R. A. Santos, S. Y. Wang, and F. M. Gambetta, “Extracting the spin excitation spectrum of a fermionic system using a quantum processor” *Phys. Rev. B*, **112**, 045143, 2025.
- Chapter 4 is based on work presented in
L. Vilchez-Estevez, A. Yosifov, and J. Sun, “Direct probing of the simulation complexity of open quantum many-body dynamics”, *arXiv:2508.19959*, 2025.
- During their DPhil, the author worked on the following publications that are not included in this thesis:
N. Tibau Vidal, and L. Vilchez-Estevez, “Creation and annihilation operators for 2D non-abelian anyons”, *arXiv:2304.10462*, 2023.
N. Tibau Vidal, V. Kunte, L. Vilchez-Estevez, M. Lal Bera, M. N. Bera, “Entanglement Asymmetry in non-Abelian Anyonic Systems”, *arXiv:2406.03546*, 2024.

2

Probing spectral features of quantum many-body systems

In this chapter, we present a technique to measure the energy spectrum of quantum many-body systems using quantum simulators. By sampling system dynamics at randomly chosen time intervals, our method acts like a built-in *spectrometer* without requiring ancillary qubits. We show how this approach makes frequency peaks appear naturally, allowing us to estimate energy transitions with only a logarithmic increase in runtime as we demand higher precision. We then examine how our protocol holds up against realistic noise, proving that its total runtime still grows reasonably with accuracy. Numerical tests on spin lattice models confirm the method’s scalability. Finally, we demonstrate the protocol on IBM’s quantum hardware to illustrate its practical potential.

This chapter is based on a publication in collaboration with Jinzhao Sun, Vlatko Vedral, Andrew Boothroyd, and Myungshik S. Kim [64]. The author of this thesis contributed to formulating the central idea and theoretical framework, developed all computational code, performed all numerical simulations, and co-authored the manuscript jointly with J.S. The theoretical aspects and derivation of the total running time in Appendix 2.A were performed by J.S. with input from the author.

2.1 Introduction

Estimating spectral features of quantum many-body systems is a central challenge in condensed matter physics and quantum chemistry. Such spectral properties encode critical information about system dynamics, phase transitions, and excitations, which are essential for understanding complex quantum phenomena. Several experimental and theoretical methods have emerged to probe these properties, including various spectroscopy techniques [65–70], as well as quantum simulation approaches based on engineered quantum devices [71–79] and quantum algorithms [80–83]. However, the complexity of accurately probing many-body systems remains significant, typically requiring extensive resources. For example, experimental neutron scattering spectroscopy needs high-intensity neutron beam facilities, while quantum algorithms like quantum phase estimation demand long coherence times and precise quantum operations [80, 81]. These limitations become particularly pronounced in the current era of noisy intermediate-scale quantum (NISQ) devices [84].

This chapter lays foundational groundwork towards efficiently estimating spectral properties using quantum computers, motivated by the need to overcome the resource-intensive challenges outlined above. Instead of directly measuring eigenenergies or conducting resource-heavy experimental spectroscopy, we leverage the intrinsic relationship between spectral features and observable dynamics. Specifically, spectral information can naturally be extracted by analysing the time evolution of observables [73, 83, 85–88]. For instance, the dynamic structure factor (DSF), central to many spectroscopy experiment, can be obtained by measuring the evolution of certain correlation functions and it reveals clear peaks where energy and momentum conservation conditions are simultaneously satisfied [66, 89].

One intuitive approach involves directly simulating spectroscopy experiments with quantum algorithms, which typically requires evaluating unequal-time correlation functions and thermal state preparation [79, 83, 90, 91]. Another complementary approach focuses on extracting spectral features from post-processing time-dependent signals [73, 85–88, 92–95], which generally reduces ancillary overhead. For example, Zintchenko and Wiebe proposed a randomised unitary-based Bayesian

inference method for spectral gap estimation [93], while Chan *et al.* demonstrated how to efficiently extract spectral properties through Fourier transforms and classical shadows [88, 96].

Nevertheless, these dynamical approaches typically face accuracy limitations when restricted to short evolution times, raising the question of whether spectral features can be probed accurately with fewer quantum resources. In response to this challenge, this chapter introduces a spectroscopic method explicitly designed to efficiently extract high-precision spectral features from short-time dynamics of quantum observables. By carefully processing dynamics with strategically chosen time intervals sampled from a predefined probability distribution, our method significantly reduces quantum resource requirements.

We analyse the maximum time complexity of our protocol, demonstrating a logarithmic dependence on the desired precision, under assumptions analogous to those employed in eigenenergy estimation methods [97–99]. This efficiency allows for precise spectral estimation even with relatively short-time evolutions. A crucial requirement for our approach is the nonvanishing observation of target transitions, which depends on the choice of initial state and observables.

Through classical simulations of spin systems and molecular models, we systematically investigate how spectral accuracy depends on the maximum evolution time. Additionally, we evaluate the algorithm’s performance under realistic noise conditions, highly relevant for near-term quantum hardware. Remarkably, our simulations demonstrate that even simple error mitigation strategies significantly enhance accuracy without substantially increasing the total runtime, which remains polynomial in inverse precision. This approach offers clear advantages compared to established methods such as quantum phase estimation or standard Fourier transforms of time signals [85, 87, 88, 100]. We further quantify the efficacy of error mitigation by incorporating both device and statistical noise into our numerical analyses. Finally, we experimentally validate our method on small spin-system simulations executed on IBM quantum devices, showcasing its practical feasibility.

2.2 Dynamical structure factor

The primary objective of this chapter is to address the extraction of transition energies and excitation spectra of quantum systems. Typically, this is achieved by studying the system's response to external perturbations in spectroscopic experiments. A key feature of this approach is the emergence of frequency excitations, which correspond to elementary transitions between eigenstates. Traditional pump-probe spectroscopy usually imposes specific constraints:

1. The perturbed system ρ must be in equilibrium with its Hamiltonian, satisfying $[\rho, H] = 0$.
2. The perturbation should be weak enough for linear response theory to hold.

Due to the lack of initial state coherence, where the density matrix ρ is diagonal in the eigenbasis of the Hamiltonian ($\langle n|\rho|n'\rangle = \delta_{nn'}e^{-\beta E_n}$), standard spectroscopy methods can only probe equilibrium properties.

In the weak perturbation regime, the system's response functions characterise how energy is absorbed from or emitted into an external probe. These response functions are directly connected to correlation functions known as dynamical structure factors (DSF) [101]. We will examine this connection thoroughly in the subsequent chapter, adapting the theoretical framework specifically to simulate excitation spectra of fermionic systems. At this stage, it is sufficient to recognise that DSFs, obtainable through spectroscopy, predominantly reflect energy resonances between eigenstates $|n'\rangle$ and $|n\rangle$ and can be expressed as follows:

$$S_{AB}(\omega) = \sum_{n,n'} A_{nn'} B_{n'n} \delta(E_{n'} - E_n - \omega) \quad (2.1)$$

with $O_{nn'} = \langle n|\hat{O}|n'\rangle$. In addition, the DSF is the Fourier transform of a two-point unequal time correlation function $\langle \psi|A(t)B|\psi\rangle$ in the equilibrium state ψ ,

$$S_{AB}(\omega) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \langle \psi|A(t)B|\psi\rangle e^{i\omega t} dt. \quad (2.2)$$

A direct method to detect transition energies involves explicitly simulating the spectroscopy procedure (see Refs. [79, 90, 91]). However, this approach tends to

be computationally demanding, as the time complexity for obtaining the spectral function $S(\omega)$ is significant, and ancillary qubits are necessary for measuring unequal-time correlators through a Hadamard-test circuit. Additionally, it requires preparing a thermal equilibrium state ρ , limiting detection strictly to equilibrium properties. This motivates the question of whether it is possible to reduce simulation runtime without compromising accuracy, a key objective addressed throughout this chapter.

2.3 Framework

In this section, we develop a general framework for estimating the transition energies between eigenstates of a quantum many-body system. Transition energies, or energy gaps between eigenstates, are central to understanding the behaviour and excitation spectrum of quantum systems. While conventional spectroscopic methods typically rely on simulating or mimicking physical interactions with external probes, here we take a computational approach that leverages the system's native unitary evolution under its own Hamiltonian.

A key advantage of our approach is that it avoids the need for ancillary qubits or controlled operations, which are often unavailable or resource-intensive on near-term quantum devices. Instead, we extract spectral information by monitoring the time evolution of observables and processing these signals classically. The essential idea is to encode spectral resolution through a frequency-selective channel applied to the system's initial state, followed by a measurement. We define a frequency-resolved operator $\mathcal{G}(\rho, \omega)$, which encodes the contribution of transitions with energy difference ω present in the initial state ρ :

$$\mathcal{G}(\rho, \omega) = \sum_{n, n' \geq 0} \rho^{n'n} |n'\rangle \langle n| p_\tau(E_{n'} - E_n - \omega), \quad (2.3)$$

where $\rho^{n'n} := \langle n' | \rho | n \rangle$ are the matrix elements of the initial state ρ in the energy eigenbasis of the Hamiltonian $H |n\rangle = E_n |n\rangle$, and $p_\tau(\cdot)$ is a frequency windowing function that selects transitions with energy difference $E_{n'} - E_n \approx \omega$. A typical

choice is the Gaussian function

$$p_\tau(\omega) = \exp(-\tau^2\omega^2) \quad (2.4)$$

which localises spectral weight around $\omega = 0$, with resolution controlled by the parameter τ . Operationally, $\mathcal{G}(\rho, \omega)$ can be interpreted as the state ρ filtered in the energy-difference domain, where coherences between eigenstates $|n\rangle$ and $|n'\rangle$ are weighted according to their transition frequency $E_{n'} - E_n$. To extract physically relevant information from this channel, we choose a Hermitian observable $\hat{O}$ and define the following signal:

$$G(\omega) = \text{Tr}[\mathcal{G}(\rho, \omega)\hat{O}]. \quad (2.5)$$

By substituting the explicit form of $\mathcal{G}$, we obtain:

$$G(\omega) := \sum_{n, n' \geq 0} \Gamma_{n', n} p(\tau(E_{n'} - E_n - \omega)), \quad (2.6)$$

where we define the state-and-observable coherence $\Gamma_{n', n} := \rho^{n'n} \langle n | \hat{O} | n' \rangle$, a time-independent quantity that captures the spectral weight associated with the initial state and the observable. Here, we arrange $p_\tau(\omega) = p(\tau\omega)$, such that τ is coupled with ω . The expression above reveals that $G(\omega)$ functions as a spectral detector: for a suitable choice of $p(\cdot)$, each term in the sum contributes a peak centered at the transition energy $\Delta_{n'n} := E_{n'} - E_n$, with intensity determined by $\Gamma_{n'n}$. Therefore, $G(\omega)$ provides access to the spectrum of excitation energies. Moreover, $G(\omega)$ captures features analogous to those of the dynamical structure factor, yet, as we shall show, it can be implemented with significantly reduced quantum resources.

We have now established that realising the quantum operation $\mathcal{G}$ and estimating $G(\omega)$ allows us to probe the energy differences of *arbitrary* quantum systems, contingent on the presence of sufficient state-and-observable coherence. As indicated from the beginning of this thesis, a natural and practical approach to obtaining $G(\omega)$ is via real-time dynamics, which are readily implementable on quantum simulators. To achieve this, we consider the Fourier transform of $G(\omega)$. Specifically,

we need the dual representation of the frequency-windowing function $p(\cdot)$, defined through its Fourier transform:

$$\tilde{g}(t) := \int_{-\infty}^{+\infty} p(\tau\omega) e^{i\tau\omega t} d(\tau\omega), \quad (2.7)$$

along with its inverse transform:

$$p(\tau\omega) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \tilde{g}(t) e^{-i\tau\omega t} dt. \quad (2.8)$$

Since our goal is to construct a spectral detector that can effectively estimate $G(\omega) = \mathbb{E}[\hat{G}(\omega)]$ using an unbiased estimator, it is essential to carefully normalise the dual function $\tilde{g}(t)$, ensuring it behaves as a proper probability distribution. To accomplish this, we introduce a normalised function $g(t) = |\tilde{g}(t)|/c$, with normalisation factor $c := \int_{-\infty}^{\infty} |\tilde{g}(t)| dt$. This normalisation introduces a time-dependent phase factor defined by $e^{i\theta_t} := \tilde{g}(t)/|\tilde{g}(t)|$, allowing us to rewrite the function $g(t)$ explicitly as

$$g(t) = \frac{1}{c} \int_{-\infty}^{+\infty} p(\tau\omega) e^{-i\theta_t} e^{i\tau\omega t} d(\tau\omega), \quad (2.9)$$

with its corresponding dual form given by

$$p(\tau\omega) = \frac{c}{2\pi} \int_{-\infty}^{+\infty} g(t) e^{i\theta_t} e^{-i\tau\omega t} dt. \quad (2.10)$$

Inserting this Fourier representation of $p(\tau\omega)$ into Equation (2.6), the spectral detector takes the form

$$\begin{aligned} G(\omega) &= \frac{c}{2\pi} \sum_{n,n' \geq 0} \Gamma_{n',n} \int_{-\infty}^{\infty} e^{-i\tau(E_{n'} - E_n - \omega)t} g(t) e^{i\theta_t} dt \\ &= \frac{c}{2\pi} \int_{-\infty}^{\infty} \left(\sum_{n,n' \geq 0} \langle n' | e^{iE_{n'}t} \rho e^{-iE_n t} | n \rangle \langle n | \hat{O} | n' \rangle \right) e^{-i\tau\omega t} g(t) e^{i\theta_t} dt \\ &= \frac{c}{2\pi} \int_{-\infty}^{\infty} G(\tau t) g(t) e^{i\theta_t} e^{i\tau\omega t} dt \end{aligned} \quad (2.11)$$

where we define $G(t) := \text{Tr}[\hat{O}\rho(t)]$ in the Schrödinger picture. Thus, Equation (2.11) clearly shows that we can first measure the observable $\hat{O}$ on the time-evolved state

$\rho(\tau t)$, and subsequently reconstruct the spectral signal $G(\omega)$ through classical post-processing of the measured data. Since $g(t)$ is normalised and thus interpretable as a probability distribution, a single-shot unbiased estimator of $\hat{G}(\omega)$ can be defined as:

$$\hat{G}(\omega) = \frac{c}{2\pi} \hat{o}(\tau t_i) e^{i\theta_{t_i}} e^{i\tau\omega t_i}, \quad (2.12)$$

where the time t_i is drawn from the probability distribution $g(t)$, which is τ -independent, and $\hat{o}(\tau t_i)$ is an unbiased estimator of the expectation value $\text{Tr}[\hat{O}\rho(\tau t)]$. It can readily be verified that this estimator is unbiased, satisfying $G(\omega) = \mathbb{E}\hat{G}(\omega)$, where the average is taken over the probability distribution $g(t)$. Note that, in this formulation, we treat $\tau\omega$ collectively as a single variable in Equation (2.11). Consequently, the probability distribution $g(\cdot)$ does not depend explicitly on τ , but the total evaluation time, defined by the maximal required time evolution, is scaled as τt . This extension by a factor τ compared to the original Fourier transform defines the overall time complexity of our protocol. A notable advantage of this approach is that it allows us to systematically evaluate resource requirements for various choices of the windowing function $p(\cdot)$ within a unified and generalisable framework, thereby avoiding case-by-case analyses.

There are two necessary requirements for inferring the transition energy $\Delta_{n',n}$ using the spectral detector in Equation (2.6):

1. a sufficiently large state-and-observable coherence $\Gamma_{n',n}$, and
2. a proper function $p(\omega)$ (or equivalently its dual form $g(t)$) that ensures that $\Delta_{n',n}$ can be distinguished from other transition energies.

The coherence $\Gamma_{n',n}$ is time-independent yet dependent on the state and the chosen observable. In the next section, we discuss the selection of the initial state and observables in order to satisfy the first condition.

The function $p(\cdot)$ plays a role as a spectral detector that can filter $\Delta_{n',n}$ out of the other transition energies. In this work, we focus on the Gaussian function because its rapid decay minimises truncation error, allowing high-precision gap estimation from short-time dynamics. Alternative filter shapes may be used, and their resource

costs can be assessed within the same unified framework. We stress that, while our method reliably locates peaks in the excitation spectrum (hence determining energy differences), the filter itself broadens each peak. Consequently, one cannot infer intrinsic linewidths, such as quasiparticle lifetimes, or distinguish whether observed broadening arises from the finite evolution time or from a continuous spectrum, since both effects manifest as increased peak width.

2.4 Quench spectroscopy protocol

As stated previously, the success of the protocol relies on the coherence $\Gamma_{n',n}$ being nonzero. We now discuss how to choose the initial state and observable to guarantee this. Several works have shown how to probe excitation spectra by engineering controllable quantum systems [71, 86, 87, 102–106]. Here, we summarise the main ideas and extend them to our general framework.

The basis of *quench* spectroscopy protocols is to prepare an initial state $|\psi_0\rangle$ that is a superposition of low-lying energy eigenstates: $|\psi_0\rangle = \sum_j c_j |j\rangle$ where $H|j\rangle = E_j|j\rangle$. Such a superposition can be generated by applying a quench operator $\hat{L}$ to the ground state, $|\psi_0\rangle = \hat{L}|0\rangle$. For a unitary L , the observable coherence between the ground state and an excited state $|n\rangle$ is

$$\langle 0|\hat{O}|n\rangle = \sum_j c_j \langle j|(\hat{L}^{-1})^\dagger \hat{O}|n\rangle, \quad (2.13)$$

so one can ensure $\langle 0|\hat{O}|n\rangle \neq 0$ by a suitable choice of $\hat{L}$. In the common case that $\hat{L}$ is unitary, choosing $\hat{L} = \hat{O}$ immediately gives $\langle 0|\hat{O}|n\rangle = c_n$.

Next, we must ensure that the density-matrix coherences $\rho^{n'n}$ are also non-vanishing. Before doing so, we recall momentum-selection rules: in a translationally invariant system, $[H, \hat{\mathbf{P}}] = 0$, where $\hat{\mathbf{P}}$ is the total momentum operator satisfying $\hat{\mathbf{P}}|n\rangle = \mathbf{p}_n|n\rangle$. Using the translation operator, a local observable at position $\mathbf{x}$ is given by

$$\hat{O}(\mathbf{x}) = e^{-i\hat{\mathbf{P}}\cdot\mathbf{x}} \hat{O} e^{i\hat{\mathbf{P}}\cdot\mathbf{x}} \quad (2.14)$$

with $\hat{O} := \hat{O}(\mathbf{0})$. Assuming the initial state breaks translation invariance, we introduce the momentum-resolved signal by performing a Fourier transform in space this time:

$$\begin{aligned} G_k(\omega) &= \int d\mathbf{x} e^{-i\mathbf{k}\mathbf{x}} G_x(\omega) \\ &= 2\pi \sum_{n,n'=0} \Gamma_{n',n} p(\tau(E_{n'} - E_n - \omega)) \delta(\mathbf{p}_{n'} - \mathbf{p}_n - \mathbf{k}). \end{aligned} \quad (2.15)$$

This result shows that translational symmetry enforces both energy and momentum conservation in any allowed transition: a term in $G_k(\omega)$ survives only if $E_{n'} - E_n = \omega$ and $\mathbf{p}_{n'} - \mathbf{p}_n = \mathbf{k}$. This is the key element in spectroscopy detection, where elementary excitations between eigenstates emerge when the selection rules of energy and momentum are both satisfied.

Although in general a large coherence cannot be guaranteed (as this problem is quantumly hard), there are certain cases where we can manipulate the system to meet this requirement and allow for probing specific types of excitations. Let us first focus on the case of weak quenches, where the low-energy excitations are restricted to a manifold of single-particle (or quasiparticle) states [87, 102]. In this regime, an excited state can be written as a single particle or quasiparticle excitation above the ground state $|0\rangle$, i.e., $|n\rangle = \hat{\gamma}_{\mathbf{q}}^\dagger |0\rangle$, where $\hat{\gamma}_{\mathbf{q}}^\dagger$ is the creation operator for a quasiparticle with momentum $\mathbf{q}$. The quasiparticles may be bosonic or fermionic and satisfy the corresponding (anti)commutation relations.

To probe such excitations, we can consider an observable of the form

$$\hat{O} = \sum_{\mathbf{p}} A_{\mathbf{p}} \hat{\gamma}_{\mathbf{p}} + A_{\mathbf{p}}^* \hat{\gamma}_{\mathbf{p}}^\dagger, \quad (2.16)$$

which couples directly to the quasiparticle modes. The matrix element between the ground and excited state then reads $\langle 0 | \hat{O} | n \rangle = A_{\mathbf{q}}$, assuming $\mathbf{q} \neq \mathbf{0}$, with the ground state defined as the quasiparticle vacuum.

Intuitively, to excite such a mode, we can prepare an initial state containing a small perturbation with momentum $\mathbf{q}$, $|\psi_0\rangle \approx |0\rangle + \beta |n\rangle$ where $\beta \ll 1$ is a small constant [107]. With this choice of initial state and observable, the total

coherence becomes $\Gamma_{n,0} = \beta A_{\mathbf{q}} \neq 0$, ensuring the transition can be detected in the spectral signal.

Such an initial state can be prepared by applying a perturbation of the form $\hat{L} \propto \mathbb{I} + \hat{P}$, where $\hat{P}$ is an operator that couples the ground state to the relevant excited states. This construction guarantees $\langle 0 | \hat{L}^\dagger | 0 \rangle \propto 1$ and $\langle 0 | \hat{L}^\dagger | n \rangle \neq 0$, which implies that the density matrix coherence is

$$\rho^{0n} = \langle 0 | \psi_0 \rangle \langle \psi_0 | n \rangle = \langle 0 | \hat{L}^\dagger | 0 \rangle \langle 0 | \hat{L}^\dagger | n \rangle \neq 0. \quad (2.17)$$

In the particular case of spin systems, the operator $\hat{L}$ will take the form of a spin rotation of an angle θ around the axis $\mathbf{n}$, since it can be written as a linear combination of the identity and a Pauli operator:

$$R_{\mathbf{n}}(\theta) = \cos\left(\frac{\theta}{2}\right) \mathbb{I} - i \sin\left(\frac{\theta}{2}\right) \mathbf{n} \cdot \boldsymbol{\sigma}. \quad (2.18)$$

More generally, let us consider a Hamiltonian $\hat{H}$ that conserves either particle or quasiparticle number. We denote its eigenstates by $|n\rangle$ and the vacuum (ground) state by $|0\rangle$, and assume the system has N modes, either in real space or momentum space. A single-particle excitation created by $\hat{\gamma}_{\mathbf{p}}^\dagger$ acting on the vacuum, i.e. $\hat{\gamma}_{\mathbf{p}}^\dagger |0\rangle$, can be expressed as a linear combination of energy eigenstates:

$$\hat{\gamma}_{\mathbf{p}}^\dagger |0\rangle = \sum_n c_{n\mathbf{p}} |n\rangle, \quad \text{where} \quad c_{n\mathbf{p}} := \langle n | \hat{\gamma}_{\mathbf{p}}^\dagger | 0 \rangle. \quad (2.19)$$

Given that the quasiparticle picture holds, we can prepare an initial state as a weak perturbation of the vacuum:

$$|\psi_0\rangle = \frac{1}{\sqrt{1+\beta^2}} \left(|0\rangle + \beta \hat{\gamma}_{\mathbf{p}}^\dagger |0\rangle \right), \quad (2.20)$$

where $\beta \ll 1$ controls the excitation strength.

The initial state coherence between the vacuum and an eigenstate $|n\rangle$ is given by

$$\rho^{n0} = \langle n | \psi_0 \rangle \langle \psi_0 | 0 \rangle = \frac{\beta \langle n | \hat{\gamma}_{\mathbf{p}}^\dagger | 0 \rangle}{1 + \beta^2} = \frac{\beta c_{n\mathbf{p}}}{1 + \beta^2}. \quad (2.21)$$

The transition amplitude observed by the annihilation operator on the $\mathbf{p}'$ th mode is $\langle 0 | \hat{\gamma}_{\mathbf{p}'} | n \rangle = c_{\mathbf{p}',n}^*$. Therefore the coherence observed by $\hat{O}$ is given by

$$\Gamma_{n,0} = \sum_{\mathbf{p}'} \frac{\beta A_{\mathbf{p}'}}{1 + \beta^2} c_{\mathbf{p}',n}^* c_{n\mathbf{p}}. \quad (2.22)$$

In a similar fashion, one can probe transition energies between two excited states $|n\rangle$ and $|n'\rangle$ that lie within the same particle-number sector. To do so, we prepare an initial state containing two particles by creating two particles as

$$|\psi_0\rangle = \frac{1}{1+\beta^2} (1 + \beta \hat{\gamma}_{\mathbf{p}}^\dagger) (1 + \beta \hat{\gamma}_{\mathbf{p}'}^\dagger) |0\rangle. \quad (2.23)$$

Again, we can decompose the single-particle state $\hat{\gamma}_{\mathbf{p}}^\dagger |0\rangle$ into the single-particle eigenbasis of $|n\rangle$, and we denote the expansion coefficient $\langle n | \hat{\gamma}_{\mathbf{p}}^\dagger | 0 \rangle := c_{n,\mathbf{p}}$ with the normalisation condition $\sum_n |c_{n,\mathbf{p}}|^2 = 1$. We could find that in the case where the number of possible excitations grows polynomially with respect to the system size the coefficient will be nonvanishing.

Similarly, for the two-particle state, we can decompose it into the two-particle eigenbasis of $|n^{(2)}\rangle$, and we denote the expansion coefficient $\langle n^{(2)} | \hat{\gamma}_{\mathbf{p}}^\dagger \hat{\gamma}_{\mathbf{p}'}^\dagger | 0 \rangle := c_{n,\mathbf{p}\mathbf{p}'}^{(2)}$. If we simply wish to observe the transitions in the single-particle manifolds, we can choose the observable $\hat{O} = \sum_{\mathbf{q},\mathbf{q}'} A_{\mathbf{q},\mathbf{q}'} \hat{\gamma}_{\mathbf{q}}^\dagger \hat{\gamma}_{\mathbf{q}'}$ that conserves the particle number. The corresponding transition amplitude in this case can be approximated to $\langle n | \hat{O} | n' \rangle = \sum_{\mathbf{q},\mathbf{q}'} A_{\mathbf{q},\mathbf{q}'} c_{n,\mathbf{q}} c_{\mathbf{q}',n'}^*$, where we have used the fact that the space of eigenstates with different particle numbers is orthogonal, i.e. $\langle n | \hat{\gamma}_{\mathbf{q}}^\dagger | m \rangle = 0$ if $m \neq n$ and $m \neq 0$. This cannot be proved in general, but it holds for most cases where the quasiparticle number is well-defined.

The transitions in the many-particle manifolds are more complicated; we may deliberately choose the observable to infer the transition energy. In our case, the initial state coherence is

$$\rho^{n',n} = \langle n' | \psi_0 \rangle \langle \psi_0 | n \rangle = \frac{\beta^2}{(1+\beta^2)^2} (c_{n',\mathbf{p}'} + c_{n',\mathbf{p}}) (c_{\mathbf{p}',n}^* + c_{\mathbf{p},n}^*).$$

To sum up, the coherence is given by

$$\Gamma_{n',n} = \sum_{\mathbf{q},\mathbf{q}'} \frac{\beta^2 A_{\mathbf{q},\mathbf{q}'} c_{n,\mathbf{q}} c_{\mathbf{q}',n'}^*}{(1+\beta^2)^2} (c_{n',\mathbf{p}'} + c_{n',\mathbf{p}}) (c_{\mathbf{p}',n}^* + c_{\mathbf{p},n}^*). \quad (2.24)$$

Since the Hamiltonian can be engineered, we can specifically engineer quantum devices to detect the energy excitation spectra of various quantum systems. Typical transitions could be observed by preparing initial states consisting of the desired superposition.

The essential message of this section is that if we can uniquely identify the eigenstates of a quantum system, tracking these labels over time allows us to directly determine the energy gap between any two states. Systems supporting quasiparticles inherently satisfy the condition of nonvanishing coherence. For instance, in Fermi liquids, low-energy eigenstates are conveniently described by occupation numbers $n_{\mathbf{p},\sigma} = 0, 1$, and interactions do not cause a single quasiparticle to decay immediately into multiple quasiparticles [108]. Specifically, the overlap $\hat{c}_{k,\sigma}^\dagger$, $|\langle \psi_k | \hat{c}_{k,\sigma}^\dagger | \psi_0 \rangle|^2$ between the quasiparticle state $|\psi_k\rangle$ (with momentum k) and the ground state $|\psi_0\rangle$ remains finite in the thermodynamic limit for momentum values near the Fermi surface. This overlap lies strictly between 0 and 1, reflecting that the quasiparticle state $|\psi_k\rangle$ is a superposition of momentum- k eigenstates. This finite overlap ensures coherent, particle-like excitations. Even strongly interacting states such as antiferromagnets can be interpreted within this quasiparticle framework [108].

In addition to the above example, spectroscopy protocols have demonstrated that excitations can be effectively created in many quantum systems through numerical simulation [86, 87, 109–111] or analogue quantum simulation [71, 85, 102] in cases of the Bose-Hubbard model [85, 86], spin chains [71, 76, 102], and disordered systems [105]. It is worth mentioning how our method differs from those of spectroscopy experiments. In spectroscopy experiments, for instance, the externally injected neutrons act as a weak perturbation, which can probe the intrinsic properties of materials in the equilibrium state. In global or local quench spectroscopy, the eigenstates are assumed to be nearly unchanged after the quench, which is similar to that in spectroscopy experiments, although the state will be driven out of equilibrium. Here, the eigenstates are not assumed to be unchanged after applying the perturbation, and thus our method holds for modifying the initial states and the Hamiltonian, and is applicable for a more general quantum many-body problem, given that the coherence is large.

2.5 Error analysis and resource requirement

In this section, we briefly discuss the computational complexity associated with estimating energy transitions. Our primary goal is to identify the transition energy $\Delta_{n',n}$. For this, the corresponding coherence $\Gamma_{n',n}$ must be nonvanishing, specifically bounded from below by a polynomial in the inverse system size N : $\Gamma_{n',n} \geq \Omega\left(\text{Poly}\left(\frac{1}{N}\right)\right)$. This condition is naturally satisfied in the scenarios considered previously. For instance, in particle-number-conserving cases, when the initial state involves a collection of low-lying excited states, $\Gamma_{n,0}$ remains non-vanishing, whereas $\Gamma_{n,n' \geq 1}$ becomes negligible at higher order. To simplify the analysis, we sort transition energies $\Delta_{n',n}$ in ascending order and denote them as Δ_j . Additionally, we define the spectral gap resolution as $\gamma_j := \min(\Delta_{j+1} - \Delta_j, \Delta_j - \Delta_{j-1})$.

Our objective is then to estimate each transition energy Δ_j within a given accuracy ε , such that $|\hat{\Delta}_j - \Delta_j| \leq \varepsilon$. Intuitively, Δ_j can be identified by locating peaks in the absolute value of the spectral function $G(\omega)$ within the frequency domain ω . To illustrate this clearly, we rewrite $G(\omega)$ as

$$G(\omega) = \Gamma_j p(\tau(\Delta_j - \omega)) + \sum_{i \neq j} \Gamma_i p(\tau(\Delta_i - \omega)). \quad (2.25)$$

It is evident that for sufficiently large τ , the first term dominates, implying that $G(\omega)$ approaches its peak value Γ_j primarily when ω is near Δ_j . Moreover, around the vicinity of Δ_j , specifically $\omega \in [a_L, a_R]$, we observe that $\frac{\partial^2 G(\omega)}{\partial \omega^2} < 0$. This negative curvature can be understood by noting that around $\omega = \Delta_j$, the function $G(\omega)$ is dominated by a Gaussian term. To see this explicitly, consider that $p(x)$ is a Gaussian function: $p(x) = e^{-\frac{x^2}{2\sigma^2}}$. Setting $x = \tau(\Delta_j - \omega)$, we have: $G(\omega) \approx \Gamma_j e^{-\frac{\tau^2(\Delta_j - \omega)^2}{2\sigma^2}}$. The second derivative evaluated at the peak ($\omega = \Delta_j$) is:

$$\left. \frac{\partial^2 G(\omega)}{\partial \omega^2} \right|_{\omega=\Delta_j} = -\frac{\Gamma_j \tau^2}{\sigma^2} < 0, \quad (2.26)$$

confirming the concave-downward curvature at the peak. Thus, the spectral function exhibits a clear and identifiable maximum at the transition frequency. This result is general and applies not only to Gaussian functions but also to any filtering

function $p(x)$ that has a clear peak at $x = 0$ with negative curvature. Consequently, the transition energy can be effectively estimated by identifying the peak of the finite-measurement estimate $\hat{G}(\omega)$, given by:

$$\hat{\Delta}_j = \operatorname{argmax}_{\omega \in [a_L, a_R]} |\hat{G}(\omega)|. \quad (2.27)$$

In the following discussion, we leverage results from spectral filtering methods [88, 93, 97–99] to analyse the computational cost associated with our protocol. The total running time (defined as the product of maximal evolution time and the number of samples) in references [97] and [93], which similarly address the problem of eigenstate preparation and spectral gap estimation, achieves the Heisenberg limit, i.e. an optimal scaling on quantum computers, with complexity $\tilde{\mathcal{O}}(\varepsilon^{-1})$ for eigenenergy estimation. However, these methods are sub-optimal concerning maximal time complexity, a metric of greater practical relevance for near-term quantum devices due to their limited coherence times. The maximal time directly corresponds to circuit depth; thus, optimising this metric is critical, as deeper circuits are more susceptible to errors in noisy quantum hardware.

We will outline only the core idea here and refer the reader to the Appendix 2.A for the complete proof of the complexity scaling for this protocol and extensive technical details of the algorithm. Specifically, for $\tau = \mathcal{O}(\gamma_j^{-1} \log^{\frac{1}{2}}(\varepsilon^{-1}))$, it can be shown that:

1. $|\hat{G}(\omega) - \Gamma_j| < c_1 \tau^2 \varepsilon^2, \forall |\omega - \Delta_j| \leq 0.5\varepsilon;$
2. $|\hat{G}(\omega) - \Gamma_j| > c_2 \tau^2 \varepsilon^2, \forall |\omega - \Delta_j| \in [\varepsilon, 0.1\gamma_j].$

Here, c_1 and c_2 are constants independent of τ and ε , but dependent on Γ_j and the spectral gap resolution γ_j (see Lemma 1 in Appendix 2.A). These inequalities indicate that the difference $d = |\Gamma_j - \hat{G}(\omega)|$ is directly influenced by the estimation error $|\omega - \Delta_j|$. Thus, as the frequency ω approaches the target transition Δ_j , the associated peak in the spectral function becomes distinguishable from other transitions. This argument inherently assumes that Γ_j remains nonvanishing; otherwise, no distinct peak would appear.

With these theoretical conditions satisfied, we first estimate $G(\omega)$ using Equation (2.12), and subsequently determine Δ_j by identifying the peak of $\hat{G}(\omega)$ across frequencies ω . Importantly, the estimation of $G(\omega)$ only requires measurements of $\text{Tr}[\hat{O}\rho(\tau t)]$ as given by Equation (2.12), while computing $\hat{G}(\omega)$ itself is performed purely classically and thus does not incur additional quantum resource costs.

We must also address the error introduced by the finite cutoff when evaluating the integral in Equation (2.11). Ideally, the integral extends over $(-\infty, +\infty)$; however, this is not feasible in practice, so we introduce a finite cutoff interval $[-T, +T]$. Within our framework, it is straightforward to analyse the truncation error and derive the requirement on T . One can show that choosing $T = \mathcal{O}(\log^{1/2}(\varepsilon^{-1}))$ suffices to ensure that the cutoff error remains below ε . Consequently, the finite cutoff only contributes a logarithmic factor to the overall circuit complexity. Considering finite τ , a finite cutoff T for the integral, and a finite number of measurements, the algorithmic complexity can be summarised as follows.

Given a nonvanishing Γ_j , to ensure that the estimate $\hat{\Delta}_j$ approximates the true value Δ_j within error ε , we require a maximum evolution time of $\mathcal{O}(\gamma_j^{-1} \log(1/\varepsilon))$ and a total number of measurements $\tilde{\mathcal{O}}(\varepsilon^{-4} \Gamma_j^{-2} \gamma_j^3)^1$. As emphasised earlier, while this protocol does not achieve Heisenberg-limited scaling in the *total* evolution time – as attained by certain phase-estimation-based approaches – it does achieve a maximum evolution time that scales only logarithmically with $1/\varepsilon$, which is the relevant figure of merit for near-term, coherence-limited quantum devices. A rigorous description of these results, along with detailed proofs, is provided in Appendix 2.A.

Finally, we comment on an important practical aspect: the measurement of observable dynamics $G(t)$, which in this protocol can be implemented using Trotterized product formulae without requiring ancillary qubits. It is therefore valuable to discuss the error introduced when implementing this on a quantum computer. We consider using the improved Trotter formula developed in Ref. [112] that uses LCU (as introduced in Chapter 1) to compensate for the Trotter error. Suppose the Hamiltonian can be decomposed into Pauli operators as $H = \sum_{l=1}^L \alpha_l P_l$, with

¹Polylogarithmic dependence on ε is hidden within the $\tilde{\mathcal{O}}$ notation.

P_l denoting Pauli terms and $\tilde{\Lambda} = \sum_l |\alpha_l|$. The gate complexity for simulating time evolution is $O((\tilde{\Lambda}t)^{1+1/(2k+1)} \log(\varepsilon^{-1}))$, where the small overhead in the exponent depends on the order of the algorithm.

For lattice models with nearest-neighbour interactions, the asymptotic scaling is $O((Nt)^{1+2/(2k+1)} \varepsilon^{-1/(2k+1)})$, with N being the system size. In cases where the evolution time is $t = \mathcal{O}(N\gamma^{-1})$, the gate complexity becomes $O(N^{2+2/(2k+1)} \varepsilon^{-1/(2k+1)} \gamma^{-1})$. This scaling is particularly advantageous for lattice models, which are the primary focus in this chapter.

However, in other settings – such as when locality is broken or in fermionic systems (as explored in the next chapter, where fermionic operators must be encoded into qubits) – the choice of encoding or the product formula may significantly affect performance. In such cases, alternative encodings, operator orderings, or different product formulae may be more beneficial. This adaptability highlights a key advantage of our framework: the time evolution algorithm can be flexibly chosen by the user without affecting the general theoretical predictions.

To summarise, this section has analysed the computational complexity of estimating transition energies using our protocol. We showed that the time complexity for transition energy estimation has only a logarithmic dependence on the target simulation accuracy, making it highly favourable for practical implementations. We also discussed the requirements on the number of measurements, finite cutoff effects, and practical considerations for implementing time evolution on quantum hardware using Trotterization.

2.6 Error effect in transition energy estimation

Finally, in this section, we analyse the effects of errors in the estimated energy transitions before moving to simulation results. These errors, whether algorithmic (e.g., from Trotterization) or arising from Hamiltonian imperfections, can be effectively modeled as an additional term δH in the Hamiltonian.

For lattice models, the primary effect of Trotterization is a shift in the energy spectrum, which can be understood within the framework of first-order perturbation

theory. The perturbed Hamiltonian takes the form $H_\nu = H + \delta H$, where $|n\rangle$ are the eigenstates of the original Hamiltonian ($H|n\rangle = E_n|n\rangle$), and $|\nu\rangle$ denote the eigenstates of the perturbed Hamiltonian H_ν . Importantly, this perturbation preserves translation invariance, such that $\hat{\mathbf{P}}|\nu\rangle = \mathbf{p}_\nu|\nu\rangle$. The expectation value of a generic observable is then given by

$$\langle\nu|\hat{O}(\mathbf{x})|\nu'\rangle = e^{i(\mathbf{p}_{\nu'} - \mathbf{p}_\nu)\mathbf{x}} \langle\nu|\hat{O}|\nu'\rangle. \quad (2.28)$$

The momentum selection rule remains valid, enforcing $\mathbf{k} = \mathbf{p}_{\nu'} - \mathbf{p}_\nu$ for excitations between the states $|\nu\rangle$ and $|\nu'\rangle$. When noise or additional perturbations are introduced, the transition energies experience shifts. According to first-order perturbation theory, the perturbed energies and eigenstates are given by

$$E_\nu = E_n + \langle n|\delta H|n\rangle, \quad (2.29)$$

$$|\nu\rangle = |n\rangle + \sum_{m \neq n} A_{mn} |m\rangle \quad (2.30)$$

with $A_{mn} = \Delta_{nm}^{-1} \langle m|\delta H|n\rangle$. From the perturbed energies, it follows that the energy differences become $\Delta_{\nu',\nu} = \Delta_{n',n} + \langle n'|\delta H|n'\rangle - \langle n|\delta H|n\rangle$. We can now revisit the spectral detector $G(\omega)$, considering its expansion up to first order in perturbation theory:

$$G(\omega) = \sum_{\nu',\nu} \Gamma_{\nu',\nu} p(\tau(\Delta_{\nu',\nu} - \omega)), \quad (2.31)$$

Consequently, the presence of error leads to a deviation in the resolved energy differences from their ideal values, as the energy splittings are shifted by $\langle n'|\delta H|n'\rangle - \langle n|\delta H|n\rangle$. Moreover, we can derive the expression for the coherence after the perturbation, which takes the form $\Gamma_{\nu',\nu} = \Gamma_{n',n} + \delta\Gamma_{\nu',\nu}$, where the coherence deviation is given by

$$\delta\Gamma_{\nu',\nu} = \alpha_{n'n} \langle n|\hat{O}|n'\rangle + \beta_{n'n} \rho^{n'n}, \quad (2.32)$$

and the coefficients are defined as

$$\alpha_{n'n} := \sum_m (A_{mn} \rho^{n',m} + A_{mn'}^* \rho^{m,n}), \quad (2.33)$$

$$\beta_{n'n} := \sum_m (A_{mn'} \langle n|\hat{O}|m\rangle + A_{mn}^* \langle m|\hat{O}|n'\rangle) \quad (2.34)$$

which are related to the initial state and the observable. Compared to the unperturbed case, certain eigenstates $|m\rangle$, which do not contribute under the original selection rule, now also contribute to $\Gamma_{n',n}$, thereby altering the spectral weight observed by $G(\omega)$. Furthermore, in the presence of a disordered term that breaks translation invariance, the spectral peaks will become broadened.

2.7 Simulation results

We investigate the performance of our method for probing the energy excitation spectrum of quantum many-body systems by numerical simulation and experimental demonstration on IBMQ devices. We consider a local perturbation applied to the initial product state, which generates a branch of low-lying excitations of the system. Then the system is evolved under unitaries $U = e^{-iHt_i}$ with different times t_i drawn from the Gaussian distribution $g(t)$, followed by measurement.

2.7.1 Numerical simulations

We first demonstrate how transition energies can be estimated. Let us consider a one-dimensional (1D) Heisenberg model with an external field

$$H = J \sum_i \left(\hat{\sigma}_i^x \hat{\sigma}_{i+1}^x + \hat{\sigma}_i^y \hat{\sigma}_{i+1}^y + \hat{\sigma}_i^z \hat{\sigma}_{i+1}^z \right) + h_z \sum_i \hat{\sigma}_i^z, \quad (2.35)$$

with the periodic boundary condition, where $\hat{\sigma}_i^\alpha$ is the Pauli operator along the α axis acting on the site i , $J = -1$ is the ferromagnetic coupling strength, and a field h_z is the external magnetic field. To break the ground state degeneracy and make more excitations emerge, a small magnetic field $h_z = 10^{-2}J$ aligned in the z -direction is applied. The state is first initialised in a product state $|+\rangle^{\otimes N}$, and then a rotation is applied to the central qubit, which takes the form of $\hat{\sigma}_3^x |+\rangle^{\otimes N}$, where the qubit numbering starts from zero in this work. The observable is chosen as $\hat{\sigma}_i^y$ on each site, which reverses the quench operation. The transition energy spectrum of the Heisenberg model is shown in Figure 2.1(a). The peaks become sharper and thus distinguishable with increasing time. We find that the transition energies for the Heisenberg model are estimated very accurately because the dominant transition




Figure 2.1: Transition energy spectra search of the Heisenberg model and the LiH molecule. **a** The 7-site Heisenberg model. The dashed vertical lines show the ideal transition energies whose coherence $\Gamma_{nn'}$ is above a certain threshold, which are calculated by exact diagonalisation (ED). Inset: FWHM dependence on varying τ . The cutoff is chosen as $T_{\text{cut}} = 1$. **b** The spectrum for LiH. The vertical lines show the ideal transition energies $\Delta_{nn'}$ calculated by ED. Inset: Estimation error for the dominant transition energy indicated by the arrow with varying total time.

is distinguishable from the other ones. Figure 2.1 inset shows the full width of the peak at half maximum (FWHM) with increasing time.

We then show results for estimating transition energies for molecular systems, taking LiH as an example. The molecular Hamiltonian of LiH at the bond length $r = 5.0\text{\AA}$ is encoded into six qubits by Jordan-Wigner² mapping. The initial state is prepared as a Hartree-Fock state. In the qubit representation, the Hartree-Fock state corresponds to a computational basis state $|n_1 n_2 \cdots n_M\rangle$, where $n_i \in \{0, 1\}$ encodes the occupation of the i th spin-orbital, representing the single Slater determinant that minimises the mean-field energy. The observable is chosen as a particle-conserving operator $\hat{c}_0^\dagger \hat{c}_1 + \hat{c}_0 \hat{c}_1^\dagger$ with fermionic (creation) annihilation operator

²The Jordan-Wigner transformation maps fermionic creation and annihilation operators to strings of Pauli operators on qubits, thereby preserving the fermionic anticommutation relations.

$\hat{c}$ ($\hat{c}^\dagger$), which considers the transitions between low-lying excited states, in order to make $\langle n|\hat{O}|n'\rangle$ non-vanishing.

The transition states are marked alongside the corresponding energy resonance in Figure 2.1(b). In molecular systems, the excited states of molecular systems are often closely spaced, resulting in interference between different transitions. As a result, the visible peak appears as an addition of different peaks, making it difficult to locate the true transitions from the peak. However, as time increases, the peaks become sharper and more peaks appear, which allows for the distinction of the true transition. The inset of Figure 2.1(b) shows the estimation error with increasing time for the dominant transition. It is worth noting the initial state remains a low-lying excitation above the ground state, and the first few low-lying excited states are degenerate. As molecules become more complex, it is anticipated that more near-degenerate low-lying states will emerge, leading to an increase in peak interference and hence a more intricate spectrum. To distinguish small, adjacent peaks, we need to increase the evolution time.

2.7.2 Implementation on a quantum computer

Next, we show how the excitation spectra of lattice models can be probed with the IBM Kolkata quantum device. We consider the 1D Ising model $H = \sum_i \hat{\sigma}_i^z \hat{\sigma}_{i+1}^z + 2 \sum_i \hat{\sigma}_i^x$ with nearest-neighbour interaction. The ground state is close to the product state when the external field is large, and thus we consider the excitation which is generated by a local perturbation. The initial state is prepared as $|\psi_0\rangle = R_y^i(\frac{\pi}{2})|+\rangle^{\otimes N}$ with the single-qubit rotation operation acting on the central site i . Similarly, we demonstrate the simulation of the 1D Heisenberg model in Equation (2.35) with $h_z = 0$ and $J = 1$.

The time evolution is simulated using the second-order Trotter formula (as introduced in Chapter 1). After some initial tests on the available quantum devices, the Kolkata device performed consistently for this task and thus was selected to execute the quantum circuits. We selected up to 13 qubits with good readout

fidelities and gate fidelities. The real-time dynamics are compiled into single-qubit Pauli rotation gates and CNOT gates.

The excitation spectra of the Ising model and the Heisenberg model are shown in Figure 2.2(a,b), respectively. The energy dispersions for both models are in good agreement with the analytic results for infinitely long spin chains. The experimental simulation errors for the observable dynamics and energy excitations are shown in Figure 2.2(d,e), respectively. The experimental simulation involves over 350 CNOT gates, but the simulation error is maintained at an acceptable level, which indicates the robustness of the method.

The maximum system size is restricted by the large noise and the available size of the hardware. To further show the performance of our method for a relatively large system, we numerically test our method with an increasing system size using tensor network methods. In particular, we initialise the system in the ground state using DMRG and then use TDVP for the time evolution with maximum bond dimension $\chi = 128$. We show the simulation error with different system sizes up to 51 qubits in Figure 2.2(f). The error is decreased by increasing τ and is not increased with increasing system size. In Figure 2.2(c), we show the excitation spectra of the 51-site Heisenberg model for comparison with the experimental result in Figure 2.2(b). Note that the examples studied here are classically simulable by using tensor networks since the central aim is to test the performance of our method for known systems under different conditions. We do not attempt to prove a quantum advantage over classical computing or show its generality for solving arbitrary quantum systems. Instead, we aim to show how the spectroscopic properties of many-body systems can be probed using either analogue or digital quantum simulators as a proof of concept.

We remark that simulating general gapless systems is a widely believed challenge, which poses great challenges to most existing algorithms, such as phase estimation and quantum singular value transformation (QSVT) [46]. The difficulty is that the energies are highly degenerate such that energy is not sufficient to distinguish these states. By introducing other conjugate variables, we could label the states with both the energy and the auxiliary characters (e.g. energy and momentum (E, k))




Figure 2.2: The energy dispersions of 1D lattice models through the demonstration on the IBM quantum devices (a,b,d,e) and tensor-network simulation (c,f). **a** The excitation spectra $G(\omega)$ of an 11-site Ising model. **b** The excitation spectra of a 13-site Heisenberg model. The results are obtained from measurements on IBMQ devices. The red lines in **a** and **b** represent the energy dispersions of infinite long spin chains obtained by analytic calculations. The intensities at $\mathbf{k} = 0$ have been removed. **c** The energy dispersion for the 1D Heisenberg chain of length $L = 51$ with $h_z = 0$ and $J = 1$. The total time evolution is set as $T_{\text{tot}} = \tau = 5$. **d** Experimental simulation error for the measurement results of $\sum_i \langle \hat{\sigma}_i^y(t) \rangle / N$ under real-time evolution conducted on IBMQ device, which are compared to results using noiseless Trotter formulae for both the Ising Hamiltonian (red line) and the Heisenberg Hamiltonian (blue line). **e** Experimental simulation error for the measurement results $G(\omega)$ from IBMQ experiments in the frequency domain averaged over different momenta. The results in **d** and **e** are compared to those using noiseless Trotter formulae. The maximal time length is $T_{\text{tot}} = 5$. **f** Scalability analysis for the gapless Heisenberg chain. The numerical simulation results for different values of τ are compared with the approximately ideal case in which $\tau \rightarrow \infty$. The figure shows the maximum error of the intensity $G(\omega)$ in the momentum space for different time scales τ and system sizes. The error is calculated as $\max_k |G_{k,\tau}(\omega) - G_{k,\tau \rightarrow \infty}(\omega)|$ where the averaging is over ω .

such that the degenerated states can be distinguished. Therefore, this strategy may be useful for analysing gapless systems with certain symmetry conservation, like translation invariance. Although we only test for simple and specific examples in this work, the results might shed light on the resolution of energy spectra of other gapless systems. The spectroscopic method and the examples examined here could stimulate further discussions on estimating the properties of gapless systems.

2.8 Performance in the presence of noise

The main sources of error affecting our computations are algorithmic errors, statistical uncertainty from a finite number of measurements, and imperfections in quantum operations. In the preceding sections, we addressed the first two sources and derived the resource requirements for circuit depth and measurement number necessary to achieve simulation accuracy within precision ϵ . In this section, we analyse errors arising from device noise and discuss possible mitigation strategies.

This consideration is particularly important for near-term quantum hardware, where noise remains a significant challenge. For many quantum algorithms, accounting for error mitigation can make simulations substantially more difficult and require much larger resources. Therefore, it is crucial to study noise-induced errors and their mitigation to ensure the practicality of quantum simulation protocols on current and near-future devices.

2.8.1 General strategy

We first introduce a general quantum error mitigation (QEM) strategy, developed in Ref. [113], based on probabilistic error cancellation, and discuss the associated resource costs in the presence of noise. In our protocol, we primarily require the implementation of e^{-iHt_i} for different time intervals t_i . Assuming Markovian noise arising from unwanted coupling between the device and its environment, the time evolution can be described by the Lindblad master equation for the noisy state $\rho_N(t)$:

$$\frac{d\rho_N(t)}{dt} = -i[H(t), \rho_N(t)] + \lambda \mathcal{L}[\rho_N(t)]. \quad (2.36)$$

Here, $\mathcal{L}[\rho] = \frac{1}{2} \sum_k (2L_k \rho L_k^\dagger - L_k^\dagger L_k \rho - \rho L_k^\dagger L_k)$ is the noise superoperator with error strength λ , characterising the coupling to the environment. Common noise models, such as dephasing and damping, can be represented by local Lindblad terms [21, 114].

The core idea of probabilistic error cancellation [115, 116] is to apply a recovery operation to the noisy state to counteract the effects of noise. Below, we briefly

outline the approach introduced in Ref. [113]. For a small time step δt , the discretised evolution is given by

$$\rho_\alpha(t + \delta t) = \mathcal{E}_\alpha(t)\rho_\alpha(t), \quad (2.37)$$

where $\alpha = I, N$ and $\mathcal{E}_\alpha$ denotes the ideal ($\alpha = I$) or noisy ($\alpha = N$) channel evolving the state from t to $t + \delta t$. A recovery operation $\mathcal{E}_Q$ can approximately map the noisy evolution back to the ideal one up to first order:

$$\mathcal{E}_I(t) = \mathcal{E}_Q\mathcal{E}_N(t) + \mathcal{O}(\delta t^2). \quad (2.38)$$

To realise $\mathcal{E}_Q$, one can probabilistically implement it as a linear combination of physical operators $\mathcal{B}_j$ (basis operators):

$$\mathcal{E}_Q = c \sum_j \alpha_j p_j \mathcal{B}_j, \quad (2.39)$$

where $c = 1 + \mathcal{O}(\delta t)$, $\alpha_j = \pm 1$, and p_j is a normalised probability distribution. Consequently, the ideal state after time T , $\rho_I(T)$, can be approximated as a linear sum of noisy states $\rho_{Q\vec{j}}(T)$:

$$\rho_I(T) = C \sum_{\vec{j}} \alpha_{\vec{j}} p_{\vec{j}} \rho_{Q\vec{j}}(T) + \mathcal{O}(T\delta t), \quad (2.40)$$

where $\vec{j} = (j_1, \dots, j_{n-1})$ and $C = c^n$. For further details on these derivations, see Ref. [113]. The error mitigation scheme incurs a sample overhead of C^2 , which scales as $\mathcal{O}(\exp(\lambda T))$, where λ is the noise strength. Thus, error mitigation is only efficient as long as λT remains bounded. This means that, to effectively implement QEM, the total simulation time must be kept short; if the maximum time is limited, the number of measurements can be correspondingly bounded. Since the maximum time complexity of our method is $\mathcal{O}(\log(1/\epsilon))$, the overhead in the number of measurements introduced by error mitigation is less severe compared to other algorithms. As a result, the total runtime complexity remains $\mathcal{O}(\text{poly}(1/\epsilon))$.

2.8.2 Error mitigation for global depolarising noise

In the following, we illustrate the error mitigation process using a simple noise model: global depolarising noise. In this model, the global depolarising channel acts uniformly on all qubits and, with probability $\lambda\delta t$, replaces the entire quantum state with the maximally mixed state, regardless of its initial configuration. With probability $1 - \lambda\delta t$, the state remains unaffected. This noise channel serves as a simple but illustrative model for the loss of coherence in quantum systems. The evolution under this noise can be described by

$$\mathcal{E}_{\delta t}(\rho) = (1 - \lambda\delta t)\rho + \lambda\delta t\rho_{\text{mix}}, \quad (2.41)$$

where $\rho_{\text{mix}} = \frac{I}{2^N}$ is the maximally mixed state and I is the identity matrix. To model the total noisy evolution, we apply the noise channel after each unitary evolution step, repeating this process until we reach the desired total evolution time $T = m\delta t$ where m is the total number of steps in the evolution. By using the recurrence relation and noting that $\lim_{n \rightarrow \infty} (1 + \frac{x}{n})^n = e^x$, we can derive the following expression for the noisy time-evolved state:

$$\begin{aligned} \rho(t) &= \prod_{\delta t}^T (\mathcal{E}_{\delta t} \circ \mathcal{U}_{\delta t})(\rho_0) = (1 - \lambda\delta t)^m \mathcal{U}_{\delta t}^m(\rho_0) + [1 - (1 - \lambda\delta t)^m] \rho_{\text{mix}} \\ &= \Lambda(T)^{-1} \mathcal{U}_T(\rho) + (1 - \Lambda(T)^{-1}) \rho_{\text{mix}}, \end{aligned} \quad (2.42)$$

where we have defined $\mathcal{U}_t := U(\cdot)U^\dagger$ and $\Lambda(t) = e^{\lambda t}$. To use the exponential approximation, we require $\delta t \rightarrow 0$. The overall action can be described by

$$\mathcal{E} \circ \mathcal{U}_t(\rho) = \Lambda(t)^{-1} U_t \rho U_t^\dagger + (1 - \Lambda(t)^{-1}) \rho_{\text{mix}} \quad (2.43)$$

This can be understood by evolving the state under $\mathcal{U}$ followed by a noise channel $\mathcal{E}(\cdot) = \Lambda(t)^{-1}(\cdot) + (1 - \Lambda(t)^{-1})\rho_{\text{mix}}$. The effect of this process is equivalent to global depolarising noise. Notably, the combined influence of the noise and unitary evolution depends only on the total evolution time t . The expectation value of a Pauli operator is

$$\mathbb{E}\hat{o}_{\text{noisy}}(t) = \text{Tr}(\mathcal{E} \circ \mathcal{U}_t(\rho)) = \Lambda(t)^{-1} \mathbb{E}\hat{o}_{\text{ideal}}(t). \quad (2.44)$$

This result indicates that under depolarising types of noise, the ideal expectation value can be obtained by multiplying the factor $\Lambda(t)$ to the noisy result. In some circumstances, the realistic circuit noise may be converted to a global white noise by Pauli twirling [117, 118] and as studied in Refs. [119, 120], shallow quantum circuits can scramble the noise into global white noise. We numerically verify that the global white noise ansatz could be used to approximate physical noise in certain circumstances, such as local depolarising noise, with simulation results shown in the next section.

To have an estimate of Λ under the global noise model, we can fit the experimental results in a similar vein to randomised benchmarking, which could be robust against state preparation and measurement noise [118]. We run the circuit $(\tilde{\mathcal{U}}_{t/m} \circ \mathcal{U}_{t/m})^m = \mathcal{I}$ where $\tilde{\mathcal{U}}_{t/m} = U^\dagger(t/m)(\cdot)U(t/m)$, which should be an identity channel in the ideal case. In the presence of noise, it becomes

$$(\mathcal{E}_{t/m} \circ \tilde{\mathcal{U}}_{t/m} \circ \mathcal{E}_{t/m} \circ \mathcal{U}_{t/m})^m(\rho) = \Lambda(2t)^{-1}\rho + (1 - \Lambda(2t)^{-1})\rho_{\text{mix}} \quad (2.45)$$

when the noise is gate-independent. We could run it by fitting the result undergoing by the channel $(\mathcal{E}_{t/m} \circ \tilde{\mathcal{U}}_{t/m} \circ \mathcal{E}_{t/m} \circ \mathcal{U}_{t/m})^m$ with different t and m .

Due to error mitigation, the variance of $G(\omega)$ is amplified by a factor of $\max_t \Lambda(t)$. Consequently, to suppress the statistical error to within ϵ , the number of measurements must be increased by this factor. Importantly, in the presence of noise, the number of samples required by this method remains polynomial in the inverse of the desired precision. By contrast, other methods based on the Fourier transform of time signals or quantum phase estimation require an exponentially larger number of samples, scaling as $\mathcal{O}(\exp(1/\epsilon))$. This protocol therefore offers distinct advantages over existing approaches when considering device noise. In the following, we further support this claim with simulation results incorporating both device noise and error mitigation.




Figure 2.3: Simulation data and exponential fit of the circuit evolution with local depolarising noise. Numerical results of the squared circuit $UU^\dagger$ implementation of noisy spectroscopy protocol simulating the 7-site Ising model with $h_x = 2$ and $h_z = 0.1$ for different noise strength p . The depth refers to the repetition numbers. The solid lines represent the best exponential fit for every p . All the regressions have a relative predictive power $R^2 = 0.99$.

2.8.3 Numerical tests

In what follows, we incorporate both device and statistical noise into our simulations. We numerically verify that the global white noise ansatz can approximate certain types of physical noise, such as local depolarising noise, under appropriate conditions. As discussed in the previous section, the total runtime complexity in the presence of device noise remains polynomial, $\mathcal{O}(\text{poly}(\epsilon^{-1}))$, in order to achieve a given precision. For global depolarising noise, the impact on observable estimation can be derived analytically. With knowledge of the error model, it is possible to obtain error-mitigated estimates of observables by fitting the noisy data accordingly.

We test the performance of the simple error mitigation strategy by considering a more practical setup with a depolarising noise model. Specifically, we introduce local depolarising noise after each gate, including both single- and two-qubit gates. We numerically verify the behaviour of this noise model by emulating the noisy quantum circuits. We run the circuit for a Trotter step with time interval δt followed by its inverse: $U(\delta t)U^\dagger(\delta t)$, which is just identity without noise, and repeat it for different times. In Figure 2.3, we show the results of $\sum_i \langle Y_i \rangle$ with an increasing circuit depth. Figure 2.3 indicates that the noisy results can be well approximated by an exponential decay function that aligns with Equation (2.44).

For the noisy simulation and its error mitigation, we consider the 1D Ising model considered in the previous sections with an additional term $h_z = 0.1$. We set the initial state the same as that in the main text, a maximum time evolution of $T = 5$ and the time interval $\delta t = 0.4$. We mitigate the noisy measurement outcomes by using the fitted exponential functions for each qubit. The results under different conditions (noisy, ideal, and error mitigated) are shown in both the time and frequency domains in Figure 2.4. We observe that even this simple error mitigation strategy can improve the results, though error mitigation is still needed for more general types of noise. As expected, the algorithm demonstrates strong resilience to noise, allowing us to recover the spectrum even after noisy evolution. This explains why the error is significantly smaller both in the frequency domain and the momentum space.

2.9 Discussion

In this chapter, we introduce a spectroscopic method for probing transition energies and excitation spectra of quantum many-body systems. The framework presented in this work establishes a connection between spectroscopic techniques and quantum simulation, and thus enables mutual benefits and advancements in both fields. The key element is the realisation of frequency resonance, and one can detect systems with certain symmetries by introducing additional momentum resonance. Our method requires the implementation of short-time dynamics only with the time length logarithmic in precision, and could be experiment-friendly for the current generation of quantum simulators. Advanced measurement algorithms [121, 122] (e.g. classical shadows [123]) can be directly employed to reduce the measurement complexity in measuring observables within this framework.

To reduce the measurement complexity in measuring observables, employing classical shadow methods [123] is straightforward as proposed in Ref. [88]. The total time complexity for measuring N_o low-support observables is $\mathcal{O}(\gamma^3 \varepsilon^{-4} \log(1/\varepsilon) \log(N_o))$, which is only amplified by a logarithmic factor. Advanced measurement algorithms can be directly employed within our framework, such as Pauli grouping methods [121,




Figure 2.4: Noisy and error-mitigated time dynamics of the spectroscopy protocol. **a** Dynamics of the observable expectation value $\langle Y_i \rangle$ at site $i = 1$. The figure shows the results of the noisy (with noise rate $p = 0.005$), the error mitigated (QEM) and the ideal cases for a 7-site Ising model. **b** Root mean square error (RMSE) of the noisy and the error mitigated results in the time domain. **c** Error of the noisy and QEM spectrum in the frequency domain. The green line represents the maximum value of the error in the time domain to compare both results. **d** Error of the spectrum in the k -space.

122], whose efficiency in measuring qubit-wise compatible observables has been verified for various examples.

We numerically test the scalability of the spectroscopic method for lattice models including both gapped and gapless cases. The simulation results obtained for gapless systems with translation invariance imply that our method may be useful for analysing similar systems with certain symmetry conservation, and we leave it to future work. We note that when the ground-state energy is known, finding the transition energy between the excited state $|n\rangle$ and the ground state $|0\rangle$ is equivalent to the problem of finding the eigenenergy of a quantum system. Therefore, resolving the transition energy could be quantumly hard. Our method is efficient when the state and observable dependent coherence is nonvanishing which could be satisfied

in several quantum systems with the conservation of particle numbers. Though this in general does not hold, we tested the coherence for molecules and lattice models.

On the other hand, we remark that although we do not expect that our method can resolve the transition energy for general systems, the spectroscopic protocol can find applications in simulating the energy excitation spectrum; in this task the excitation is a collective behaviour and hence we do not necessarily require resolving a specific energy transition. The simulation result suggests that the spectral properties of several quantum many-body systems can be estimated with high accuracy, even with increasing system size for lattice Hamiltonians considered in this work and when statistical and device noise are present.

We also demonstrated how our method can be applied to probe the excitation spectra of spin Hamiltonians on IBMQ devices. These results indicate that our approach could serve as a quantum computing solution complementary to high-intensity scattering facilities.

Appendices

2.A Further derivation of the algorithmic error and resource requirement

Recall that in the main text we define the spectral detector as

$$G(\omega) := \sum_{n,n' \geq 0} \Gamma_{n',n} p(\tau(E_{n'} - E_n - \omega)), \quad (2.46)$$

which can be simplified as $G(\omega) = 2 \sum_{n < n'} \Re(\Gamma_{n',n}) p(\tau(E_{n'} - E_n - \omega))$. For simplicity, we sort $\Delta_{n',n}$ in ascending order and denote ordered energies as Δ_i and the associated coherence is denoted as $\Gamma_i \leftrightarrow 2\Re(\Gamma_{n',n}) \in \mathbb{R}$. The spectral detector can now be expressed as $G(\omega) = \sum_i \Gamma_i p(\tau(\Delta_i - \omega))$.

Suppose we aim to estimate Δ_j with a nonvanishing Γ_j . Without loss of generality, we take $\Gamma_j > 0$ so that $G(\omega)$ is positive when ω approaches Δ_j (given a large Γ_j). Analogous results hold for negative Γ_j . Intuitively, the spectral information can be extracted from the peak in the intensity of $|G(\omega)|$, as formalised in the following lemma.

Lemma 1. *When $\tau = \frac{1}{0.9\gamma} \sqrt{\ln \left(\frac{20}{\varepsilon^2 \Gamma_j} \right)}$ and $\varepsilon \leq 0.2\gamma$, the two inequalities hold:*

$$\Gamma_j - G(\omega) \leq 0.3\tau^2 \varepsilon^2 \Gamma_j, \forall |\omega - \Delta_j| \leq 0.5\varepsilon \quad (2.47)$$

and

$$\Gamma_j - G(\omega) \geq 0.8\tau^2 \varepsilon^2 \Gamma_j, \forall |\omega - \Delta_j| \in (\varepsilon, 0.1\gamma) \quad (2.48)$$

Proof. When ω is close to Δ_j satisfying $|\omega - \Delta_j| \leq 0.5\varepsilon$, we have

$$\Gamma_j - G(\omega) \leq (1 - p(\tau(\Delta_j - \omega)))\Gamma_j + \max_{i \neq j} p(\tau(\Delta_i - \omega)) \quad (2.49)$$

$$\leq (1 - e^{-\tau^2(0.5\varepsilon)^2})\Gamma_j + 0.05\tau^2 \varepsilon^2 \Gamma_j \leq 0.3\tau^2 \varepsilon^2 \Gamma_j. \quad (2.50)$$

The second inequality uses the following fact that the inequality

$$\max_{i \neq j} p(\tau(\Delta_i - \omega)) \leq e^{-\tau^2(0.9\gamma)^2} \leq 0.05\tau^2\varepsilon^2\Gamma_j, \quad (2.51)$$

holds when $\tau = \frac{1}{0.9\gamma} \sqrt{\ln\left(\frac{20}{\varepsilon^2\Gamma_j}\right)}$ and $\varepsilon \leq 0.2\gamma$. Thus the first inequality by Equation (2.47) holds.

On the other hand, the quantity $\Gamma_j - G(\omega)$ can be lower bounded by

$$\begin{aligned} \Gamma_j - G(\omega) &\geq (1 - p(\tau(\Delta_j - \omega)))\Gamma_j - \max_{i \neq j} p(\tau(\Delta_i - \omega)) \\ &\geq (1 - e^{-\tau^2\varepsilon^2})\Gamma_j - e^{-\tau^2(0.9\gamma)^2} \\ &\geq 0.85\tau^2\varepsilon^2\Gamma_j - 0.05\tau^2\varepsilon^2\Gamma_j = 0.8\tau^2\varepsilon^2\Gamma_j. \end{aligned} \quad (2.52)$$

In the second inequality, we used the fact that $|\Delta_i - \omega| \geq 0.9\gamma$ for $i \neq j$. The third inequality $e^{-\tau^2\varepsilon^2} \leq 1 - 0.85\tau^2\varepsilon^2$ holds when $\tau\varepsilon < 1/2$, which can be achieved when $\varepsilon \leq \tilde{\mathcal{O}}(\gamma(\ln \Gamma_j^{-1})^{-\frac{1}{2}})$. $\square$

The above lemma guarantees that the distance $d = |\Gamma_j - \hat{G}(\omega)|$ is modulated by the estimation error $|\omega - \Delta_j|$. Therefore, Δ_j can be distinguished from the other transitions when ω approaches Δ_j . The next step is to analyse the quantum resources required to have a good estimation of $\hat{G}(\omega)$.

When evaluating the integral in Equation (2.11) in the main text, we set a finite cutoff for the integral range from $[\infty, \infty]$ to $[-T, T]$ in practice. The truncated estimator $\hat{G}(\omega)$ is defined as

$$G^{(T)}(\omega) = \int_{-T}^T G(\tau t) g(t) e^{i\theta t} e^{i\tau\omega t} dt. \quad (2.53)$$

If a Gaussian operation is used, the truncation error has an upper bound

$$|G^{(T)}(\omega) - G(\omega)| \leq 2 \int_T^\infty g(t) dt = \text{erfc}(T/2) \leq \exp(-T^2/4) \quad (2.54)$$

which has been discussed in prior work [97, 99]. Therefore, the truncation error can be bounded by

$$|G^{(T)}(\omega) - G(\omega)| \leq \varepsilon_T, \quad \forall \omega \in \mathbb{R} \quad (2.55)$$

when the cutoff time $T \geq 2\sqrt{\ln(1/\varepsilon_T)}$.

In the following, we consider the error due to a finite number of measurements. A single-shot estimator of $G^{(T)}(\omega)$ takes the form of

$$\hat{G}_i^{(T)}(\omega) = \begin{cases} \hat{o}(\tau t_i) e^{i\theta_{t_i}} e^{i\tau\omega t_i}, & |t_i| \leq T \\ 0, & |t_i| > T \end{cases} \quad (2.56)$$

where the time length t_i is drawn from the probability distribution $g(t)$, which is τ -independent, and $\hat{o}(\tau t_i)$ is an unbiased estimate of $\text{Tr}[\hat{O}(\tau t)\rho]$. We can estimate $G(\omega)$ by

$$\hat{G}(\omega) = \frac{1}{N_s} \sum_{i=1}^{N_s} \hat{G}_i^{(T)}(\omega), \quad (2.57)$$

where N_s is the total number of samples. The transition energy is estimated by

$$\hat{\Delta}_j = \text{argmax}_{\omega \in [\Delta_j - \gamma/2, \Delta_j + \gamma/2]} \hat{G}^{(T)}(\omega). \quad (2.58)$$

Note that it is assumed that $\Gamma_j > 0$. The result is summarised as follows.

Proposition 1. *To guarantee that an estimation $\hat{\Delta}_j$ of the transition energy Δ_j is close to the true value within an error ε , with a failure probability, we require that the maximum time is $\mathcal{O}(\gamma^{-1} \log(1/\varepsilon))$ and the number of measurements $\tilde{\mathcal{O}}(\varepsilon^{-4} \Gamma_j^{-2} \gamma^3 \log(1/\delta))$, where γ is a chosen lower bound of the gap difference $\gamma \leq \min\{\Delta_{j+1} - \Delta_j, \Delta_j - \Delta_{j-1}\}$.*

In this work we use the $\tilde{\mathcal{O}}$ notation where the polylogarithmic dependence is hidden. Proposition 1 indicates that the maximum circuit depth is logarithmic in precision. The result of time complexity also indicates that when we choose to set a small γ of order $\gamma = \mathcal{O}(\varepsilon)$ (that is, a large τ), our result is reduced to algorithmic cooling, in which the total running time reaches $\mathcal{O}(\varepsilon^{-1})$, the so-called Heisenberg limit. However, generally the total running time scales as $\tilde{\mathcal{O}}(\varepsilon^{-4})$, in contrast to the result established in Refs. [93, 97] which is $\tilde{\mathcal{O}}(\varepsilon^{-1})$. The cost of a long total running time can be understood as follows. Recall that the target is to estimate Δ_j with accuracy ε , i.e., $|\hat{\Delta}_j - \Delta_j| \leq \varepsilon$. By the Taylor expansion, we have $|G(\hat{\Delta}_j) - G(\Delta_j)| = \Gamma_j \tau^2 \varepsilon^2 + \mathcal{O}((\tau \varepsilon)^4)$. Due to a finite number of measurements, we can only have an estimate of either $G(\hat{\Delta}_j)$ or $G(\Delta_j)$. According to the Chebyshev

inequality, the sampling number required to achieve accuracy $\tilde{\varepsilon}$ scales as $\mathcal{O}(\tilde{\varepsilon}^{-2})$. To ensure that the energy estimation error is up to ε , we choose to set $\tilde{\varepsilon} = \Gamma_j(\tau\varepsilon)^2$, and consequently, the sampling number scales as $N_s = \tilde{\mathcal{O}}(\varepsilon^{-4}\Gamma_j^{-2}\gamma^4)$.

The above result is guaranteed by the following lemma, which ensures the efficiency of energy estimation.

Lemma 2. *The estimation $\hat{\Delta}_j$ is related to the true transition energy Δ_j by $|\hat{\Delta}_j - \Delta_j| \leq \varepsilon$. with a failure probability of δ , when $\tau \geq \frac{1}{0.9\gamma} \sqrt{\ln\left(\frac{10}{\varepsilon^2\Gamma_j}\right)}$, the cutoff $T \geq 2\sqrt{2\ln(\sqrt{10}\tau^{-1}\varepsilon^{-1})}$ and the number of measurements $N_s \geq 200(\varepsilon^4\Gamma_j^2\tau^4)^{-1} \ln(4/\delta)$.*

Proof. Suppose we already have a rough estimation of Δ_j which lies in the range of $[\Delta_j^L, \Delta_j^R]$ with $\Delta_j^L = \Delta_j - \gamma/2$ and $\Delta_j^R = \Delta_j + \gamma/2$. Then we refine the grid and let

$$\omega_k = \Delta_j^L - 0.5\gamma + \frac{\gamma(k-1)}{M} \quad (2.59)$$

for $k = 1, 2, \dots, M$, and $M = 2\lceil\gamma/\varepsilon\rceil + 1$. It is worth noting that computing $\hat{G}^{(T)}(\omega)$ with different values of ω_k requires purely classical computing, and M is irrelevant to the measurement numbers N_s . One can use a much smaller grid size without demanding quantum measurements.

The following shows that the transition energy can be estimated within a certain precision given the resources listed in Lemma 2. With the grid in Equation (2.59), the estimation is determined by

$$\hat{\Delta}_j = \operatorname{argmax}_k \hat{G}^{(T)}(\omega_k). \quad (2.60)$$

The question is whether the estimation has a bounded error satisfying $|\hat{\Delta}_j - \Delta_j| \leq \varepsilon$.

The error due to a finite number of measurements can be bounded using the Hoeffding inequality. The estimator $\hat{G}^{(T)}(\omega)$ is related to ideal $G^{(T)}(\omega)$ by

$$|G^{(T)}(\omega) - \hat{G}^{(T)}(\omega)| \leq \Gamma_j \varepsilon_n, \quad \forall \omega \in \mathbb{R} \quad (2.61)$$

with a failure probability $\delta/2$ when $N_s = 2(\varepsilon_n\Gamma_j)^{-2} \ln(\frac{4}{\delta})$. Here we used the fact that $G_i^{(T)}(\omega) \leq 1$. Since the grid resolution is $\varepsilon/2$, there exists k_m , such that

$$|\omega_{k_m} - \Delta_j| \leq 0.5\varepsilon. \quad (2.62)$$

Combining Equation (2.53) and Equation (2.61), we have

$$\begin{aligned} \Gamma_j - \hat{G}^{(T)}(\omega_{k_m}) &\leq (\Gamma_j - G(\omega_{k_m})) + |G(\omega_{k_m}) - G^{(T)}(\omega_{k_m})| + |G^{(T)}(\omega_{k_m}) - \hat{G}^{(T)}(\omega_{k_m})| \\ &\leq 0.3\tau^2\varepsilon^2\Gamma_j + \varepsilon_T\Gamma_j + \Gamma_j\varepsilon_n \leq 0.5\tau^2\varepsilon^2\Gamma_j \end{aligned} \quad (2.63)$$

with a failure probability $\delta/2$. In the last inequality, we have set $\varepsilon_n = 0.1\tau^2\varepsilon^2$ and $\varepsilon_T \leq 0.1\tau^2\varepsilon^2$. The latter condition can be satisfied by setting $T \geq 2\sqrt{2\ln(\sqrt{10}\tau^{-1}\varepsilon^{-1})}$.

This indicates that

$$\hat{G}^{(T)}(\hat{\Delta}_j) \geq \hat{G}^{(T)}(\omega_{k_m}) > (1 - 0.5\tau^2\varepsilon^2)\Gamma_j. \quad (2.64)$$

On the other hand, if the estimator determined by Equation (2.60) fails to give an accurate estimation up to ε , that is, $|\hat{\Delta}_j - \Delta_j| > \varepsilon$, then using Lemma 1 we have $\Gamma_j - G(\hat{\Delta}_j) \geq 0.8\tau^2\varepsilon^2\Gamma_j$. We have

$$\begin{aligned} \Gamma_j - \hat{G}^{(T)}(\hat{\Delta}_j) &\leq (\Gamma_j - G(\hat{\Delta}_j)) + |G(\hat{\Delta}_j) - G^{(T)}(\hat{\Delta}_j)| + (|G^{(T)}(\hat{\Delta}_j) - \hat{G}^{(T)}(\hat{\Delta}_j)|) \\ &\geq 0.8\tau^2\varepsilon^2\Gamma_j - \varepsilon_T\Gamma_j - \varepsilon_n\Gamma_j \geq 0.6\tau^2\varepsilon^2\Gamma_j \end{aligned} \quad (2.65)$$

with a failure probability $\delta/2$. The truncated estimator thus can be bounded

$$\hat{G}^{(T)}(\hat{\Delta}_j) \leq \Gamma_j(1 - 0.6\tau^2\varepsilon^2), \quad (2.66)$$

which violates Equation (2.64). Consequently, the assumption does not hold, which proves that $\Pr(|\hat{\Delta}_j - \Delta_j| > \varepsilon) < \delta$ with a failure probability at most δ . $\square$

Proposition 1 naturally flows from Lemma 2. According to Lemma 2, it is easy to verify that the maximum time is $t = \tau \times T = \mathcal{O}(\gamma^{-1}\log(\varepsilon^{-1}))$ and $N_s = \tilde{\mathcal{O}}(\varepsilon^{-4}\gamma^4\Gamma_j^{-2}\ln(1/\delta))$. The total running time scales as $\tau \times T \times N_s = \tilde{\mathcal{O}}(\varepsilon^{-4}\gamma^3\log(\varepsilon^{-1})\log(\delta^{-1}))$.

3

Quench spectroscopy of fermionic systems

In this chapter, we continue exploring spectroscopy-based techniques to study quantum many-body systems, now focusing specifically on low-energy excitations in fermionic models. We examine the spin excitation spectrum of the one-dimensional Fermi-Hubbard model using a digital quantum simulator. Our tailored protocol begins with a specific quantum quench and efficiently reconstructs the retarded spin Green's function from simple time-evolved observables. By designing the method to be resilient to perturbations in the initial state, we greatly simplify initial state preparation, enabling us to devote most quantum resources to the accurate simulation of time dynamics via Trotterization. The intrinsic robustness of our protocol to noise ensures accurate spectral reconstructions even for large system sizes. We demonstrate the experimental feasibility of our approach using up to 30 qubits on an IBM Heron r2 quantum device.

This chapter is based on a publication in collaboration with Raul A. Santos, Sabrina Yue Wang, and Filippo Maria Gambetta [124]. The author developed and optimised the algorithm with input from all co-authors, performed the classical benchmarking and resource analysis, and carried out the numerical simulations and hardware experiments. The DGA code was provided by S.Y.W., but executed by the author, who also developed all other code. The manuscript was drafted by the author and revised jointly with F.M.G.

3.1 Introduction

A central theme of this thesis is the inherent potential of quantum computers to tackle problems that are otherwise intractable with classical methods. While there has been significant progress in demonstrating quantum advantage, these achievements often involve carefully tailored problems with limited physical or scientific relevance [17, 125, 126]. A different goal, however, is to use quantum computers to address resource-intensive problems that have genuine scientific or societal impact. A historically significant example is the simulation of quantum mechanical systems. These problems are classically hard due to the exponential growth of the Hilbert space with system size, the complexity of entanglement, and, for fermions, the infamous sign problem [1, 127]. Quantum devices are believed to be naturally suited to these tasks, as originally envisioned by Feynman, who argued that quantum systems can efficiently simulate other quantum systems. This capability is particularly relevant for simulating the dynamics of electrons in materials and large molecules, which are central to many open questions in physics and chemistry. As such, the simulation of fermionic systems is of great value across a range of disciplines.

The characterisation of fermionic systems, especially through their dynamical properties, is essential for understanding quantum materials. As discussed in Chapter 2, extracting information from real-time dynamics enables the study of spectroscopic properties of many-body systems. Of particular interest is the dynamical structure factor, which encodes the system's response to external probes and describes how it absorbs or emits energy. The dynamical structure factor is primarily determined by the excitation spectrum of modes that couple to the applied perturbation. Experimentally, these quantities can be accessed through measurements of the doubly-differential cross-section in inelastic scattering experiments, such as inelastic neutron scattering (INS) [128], resonant inelastic X-ray spectroscopy (RIXS) [129], and electron energy-loss spectroscopy (EELS) [130]. The direct correspondence between theoretical predictions and

experimental probes makes dynamical structure factors a powerful tool for testing microscopic models of materials.

In this chapter, we adapt the quench spectroscopy protocol introduced previously to investigate properties of fermionic systems. We show that, by carefully choosing the quench operator, this approach gives access to the exact retarded spin correlation function, and thus the dynamical spin structure factor, provided the initial state does not break spin reflection symmetry. We apply this method to study the spin excitation spectrum of the one-dimensional Fermi-Hubbard model on an IBM digital quantum computer.

Despite their apparent simplicity, Fermi-Hubbard models are iconic prototypes of strongly correlated electronic systems and have been identified as ideal candidates for simulation on near-term quantum computers [131–143]. These models are widely believed to hold the key to understanding high-temperature superconductivity in cuprates [30, 144–150]. In this context, a detailed understanding of elementary excitations is crucial, and the computation of dynamical (density and spin) structure factors has proven to be an essential tool for their characterisation [151–155].

When implementing the quench spectroscopy protocol on a near-term quantum computer, it is crucial to account for limited computational resources and the presence of noise. Resource constraints, such as the number of available qubits, the gate count, and circuit depth, must therefore be carefully optimised. The most resource-intensive components of quench spectroscopy experiments are the initial state preparation and the simulation of time evolution. In particular, preparing non-trivial initial states, such as the ground state of the system, typically demands significant resources and poses a major experimental challenge. In this work, we exploit the remarkable robustness of the protocol to perturbations in the initial state, which allows us to minimise the resources required for state preparation. As a result, we are able to allocate the majority of computational resources to the Hamiltonian simulation, which is implemented using a Trotter decomposition. Additionally, we explore different fermion-to-qubit encodings to further reduce the depth of individual Trotter circuit layers. Together, these optimisations enable us

to perform time-dynamics simulations of Fermi-Hubbard models of up to $L = 15$ sites (corresponding to $n = 30$ qubits) on an IBM quantum computer, successfully recovering all the features of the spinon spectrum.

3.2 Unitary quench spectroscopy

In the previous chapter, we discussed how standard spectroscopic protocols probe the excitation spectrum of quantum systems by measuring their linear response to weak, external perturbations. These methods, rooted in equilibrium physics, ultimately provide access to the dynamical structure factor, a key quantity that encodes both the energy and spectral weight of transitions between eigenstates. However, implementing these traditional protocols on quantum computers presents several challenges, which we address (or at least significantly improve upon) with the protocol introduced in Chapter 2.

Here, we build upon the approaches discussed previously, as well as on existing work in local quench spectroscopy [87], and introduce a slight modification: we initiate the protocol with a tailored *unitary* perturbation. This adaptation makes it possible to efficiently simulate fermionic systems, which is the primary focus of this chapter, on quantum computers. By examining the dynamical structure factor and its connection to the *retarded Green's function* (i.e., unequal-time correlators), we show that our protocol can exactly resolve the dynamical structure factor for such systems. In contrast, earlier protocols based on the quench spectral function (the Fourier transform of time-dependent expectation values) can generally only approximate the dynamical structure factor (see discussion in Ref. [86]). While such approximations are sufficient for resolving energy differences, as demonstrated in the previous chapter, our approach places greater emphasis on the dynamical structure factor itself, an object of fundamental importance for characterising quantum many-body systems.

In a linear response theory setting, an external perturbing field (the probe) couples linearly to a system's observable B and generates a change in another

observable A . The resulting system response is captured by the retarded Green's function (also referred to as the retarded linear response function) [156, 157]

$$G_{AB}(t) := -i \langle \Psi | [A(t), B] | \Psi \rangle \quad \text{with } t \geq 0, \quad (3.1)$$

where $|\Psi\rangle$ usually denotes the ground state of the system, $[A, B] := AB - BA$ represents the commutator, and $A(t)$ is the operator in the Heisenberg picture under the Hamiltonian H , namely $A(t) := e^{-iHt} A e^{iHt}$. For simplicity, we have set $\hbar = 1$ and omitted any explicit spatial dependence.

We can now focus on the important case where $B = A^\dagger$, the *fluctuation-dissipation theorem* connects the imaginary part of the Fourier transform of $G_{AA^\dagger}(t)$ with the system's dynamical structure factor $\mathcal{S}_{AA^\dagger}(\omega)$. The latter can be obtained from inelastic scattering experiments [158] and is defined as the Fourier transform of the correlation function $\mathcal{S}_{AA^\dagger}(t) := \langle \Psi | A(t) A^\dagger | \Psi \rangle$. When working around the equilibrium state (at zero-temperature) the fluctuation-dissipation theorem reads [156]

$$\text{Im} G_{AA^\dagger}(\omega) = -2\pi \text{sign}(\omega) \mathcal{S}_{AA^\dagger}(\omega). \quad (3.2)$$

In particular, one can find that the spectral information about low-lying excitations of the system that couple to the operator A can be obtained by looking at $\text{Im} G_{AA^\dagger}(\omega)$ for $\omega > 0$ [156, 157]:

$$\text{Im} G_{AA^\dagger}(\omega) = -\pi \sum_n |A_{0n}|^2 [\delta(\omega - E_n) - \delta(\omega + E_n)], \quad (3.3)$$

where $A_{0n} = \langle \Psi | A | E_n \rangle$, with $|E_n\rangle$ eigenstate of H with eigenvalue E_n . Note that the ground-state energy E_0 is set to zero.

We can now expand the operator B in Equation (3.1) as an expansion in terms of a string of Pauli operators (i.e., $B := \sum_j \alpha_j^B P_j$). So, we can focus on $G_a(t) := i \langle \Psi | [P_a, A(t)] | \Psi \rangle$ for a Pauli string P_a . Taking into account that Pauli operators square to the identity, we obtain that

$$G_a(t) = \sum_{\sigma=\pm} \sigma \langle \Psi | e^{i\sigma \frac{\pi}{4} P_a} A(t) e^{-i\sigma \frac{\pi}{4} P_a} | \Psi \rangle. \quad (3.4)$$

Examining this relation, we see that it describes the dynamics of the operator A following a local quench of the state $|\Psi\rangle$ by the unitary operator $e^{i\sigma \frac{\pi}{4} P_a}$. Thus, the

retarded Green's function G_{AB} can be computed by measuring the time evolution of A on the quenched state. This approach generalises the local quench spectroscopy protocol introduced in Refs. [86, 87]. As previously discussed, this protocol enables the recovery of the spectrum of low-lying excitations in many-body quantum systems by tracking the non-equilibrium dynamics of single-particle observables after a quantum quench, making it ideally suited for current quantum simulators [64, 83, 159–161]. The number of measurements required in this protocol depends on the number of terms present in the expansion of the operator A .

We can optimise this further by considering the particular structure of the problem. To this end, we specialise on spin observables in fermionic systems. The spin operators are defined as bilinears in the fermion creation (annihilation) operators $c_{i\sigma}^\dagger (c_{i\sigma})$ at site i and spin σ ,

$$\hat{S}_i^\alpha := \sum_{\sigma\sigma'} c_{i,\sigma}^\dagger S_{\sigma\sigma'}^\alpha c_{i,\sigma'}, \quad (3.5)$$

where S^α is the $\alpha = x, y, z$ Pauli matrix. The retarded spin Green's function is then defined as

$$G_{jk}^x(t) := -i \langle \Psi | [\hat{S}_k^x(t), \hat{S}_j^x] | \Psi \rangle \quad \text{with } t \geq 0. \quad (3.6)$$

Using the fermionic anticommutation relations, it is easy to prove that $(\hat{S}_i^x)^{2k} = (\hat{S}_i^x)^2$ and $(\hat{S}_i^x)^{2k+1} = \hat{S}_i^x$. From these relations it follows that

$$e^{i\theta \hat{S}_j^x} = 1 + (\hat{S}_j^x)^2 (\cos \theta - 1) + i \sin \theta \hat{S}_j^x. \quad (3.7)$$

The action of this unitary operator on *any* state $|\Psi\rangle$ with local spin reflection symmetry¹ is

$$e^{i\theta \hat{S}_j^x} |\Psi\rangle = (\cos \theta + i \sin \theta \hat{S}_j^x) |\Psi\rangle, \quad (3.8)$$

¹This is a very restrictive condition. In our case, we are interested in systems with *global* spin-reflection symmetry instead. The derivation of the approximate dynamical structure factor under this weaker assumption, as well as the precise conditions required to access the exact Fourier transform of the retarded Green's function, are discussed in the revised version of Ref. [124].

where we have used that $(S_j^x)^2 |\Psi\rangle = |\Psi\rangle$. Defining the state $|\Psi_\theta^j\rangle := e^{i\theta\hat{S}_j^x} |\Psi\rangle$ and using Equation (3.7) and Equation (3.8) for $\theta = \pi/4$, we can show that

$$\begin{aligned} \langle \Psi_\pi^j | \hat{S}_k^x(t) | \Psi_\pi^j \rangle &= \frac{1}{2} \langle \Psi | (1 - i\hat{S}_j^x) \hat{S}_k^x(t) (1 + i\hat{S}_j^x) | \Psi \rangle \\ &= \frac{1}{2} \langle \Psi | \hat{S}_k^x(t) + i\hat{S}_k^x(t) \hat{S}_j^x - i\hat{S}_j^x \hat{S}_k^x(t) + \hat{S}_j^x \hat{S}_k^x(t) \hat{S}_j^x | \Psi \rangle. \end{aligned}$$

For any state with total fixed spin, the expectation of an odd number of $\hat{S}_j^x$ operators is zero, as they raise or lower the spin by one. This means that

$$\langle \Psi_\pi^j | \hat{S}_k^x(t) | \Psi_\pi^j \rangle = \frac{i}{2} \langle \Psi | [\hat{S}_k^x(t), \hat{S}_j^x] | \Psi \rangle = \frac{1}{2} G_{jk}^x(t). \quad (3.9)$$

As discussed in the previous chapter, the quench spectroscopy framework makes use of the space-time Fourier transform of the expectation value $\langle \hat{S}_k^x(t) \rangle$ measured on the quenched state $|\Psi_\pi^j\rangle$, the so-called quench spectral function (QSF) in the original local quench spectroscopy literature [87]. Generally, the QSF only approximates the true dynamical structure factor, and extracting meaningful spectral information depends on carefully matching the quench operator and the measured observable. However, the unitary quench protocol described here, allows us to directly access the retarded spin Green's function for any fermionic system whose ground state is symmetric under spin reflection. According to Equation (3.3), the spectral properties of low-lying spin excitations are encoded in the imaginary part of the Fourier transform, which, by the fluctuation-dissipation theorem of Equation (3.2), is directly related to the dynamical spin structure factor. For consistency, throughout what follows, we refer to the protocol presented here as the quench spectroscopy protocol, and to the imaginary part of the space-time Fourier transform of $\langle \hat{S}_k^x(t) \rangle$ in Equation (3.9) as the QSF.

In what follows, we apply this protocol to the 1D Fermi-Hubbard model on L sites with open boundary conditions and lattice spacing equal to 1, with Hamiltonian

$$\begin{aligned} H_{\text{FH}} &:= -J \sum_{i=0}^{L-2} \sum_{\sigma} (c_{i,\sigma}^\dagger c_{i+1,\sigma} + \text{h.c.}) + U \sum_{i=0}^{L-1} n_{i,\uparrow} n_{i,\downarrow} \\ &:= H_0 + H_{\text{int}}, \end{aligned} \quad (3.10)$$

where J is the hopping strength, U is the on-site interaction parameter (or Coulomb potential), and $n_{i,\sigma} = c_{i,\sigma}^\dagger c_{i,\sigma}$ is the density operator for the spin $\sigma \in \{\uparrow, \downarrow\}$ on site i . In this work, we will fix $J = 1$ and $U/J = 3$ unless stated otherwise. All energies (times) will be measured in units of J ($1/J$). Low-energy excitations of the Fermi-Hubbard model can be generated through multiple mechanisms. When the total number of electrons per spin sector is conserved, electron-hole pairs can be formed. Charged excitations arise by either adding or removing particles with spin-up or spin-down from the ground state. Here, we focus on spin excitations, which are produced by flipping the spin of an electron and can be described in terms of a pair of spinons [162]. The ground state of the 1D Fermi-Hubbard model has definite total spin and is symmetric under spin reflection, meaning that we can access their dynamical spin structure factor using Equation (3.9). Moreover, as we will show in Section 3.3, for the values of the system parameters we consider here one can recover the spectroscopic features of the low-energy excitations of the Fermi-Hubbard model by applying the quench to the ground state of a free-fermion model or even to a shallow circuit approximation of the latter (see Section 3.5). Both states have definite total spin and are symmetric under spin reflection, which means that Equation (3.9) still holds.

3.3 Classical benchmarking

The protocol outlined above consists of four main stages. First, initialise the system to the ground state (or to a quantum state that closely approximates it). Second, apply a local quench to introduce a low-energy excitation, whose spectrum is the focus of our study. Third, evolve in time the perturbed state over a specified duration t . Finally, measure the expectation value of the spin operator $\hat{S}_j^x$ to analyse the resulting dynamics.

As discussed in the previous section, the quench operator is given by

$$\hat{Q}_j := e^{i\frac{\pi}{4}\hat{S}_j^x}, \quad (3.11)$$

To minimise the effect of the open boundary conditions, we apply the quench operator to the central site $j = \lfloor L/2 \rfloor$. After the quench, the state is evolved under the time-evolution operator associated with Equation (3.10). Once the time-evolved state $|\Psi_{\frac{\pi}{4}}^j(t)\rangle = e^{-iHt} |\Psi_{\frac{\pi}{4}}^j\rangle$ is constructed, we measure the expectation value of the $\hat{S}_i^x$ operator for each site $i \in \{0, \dots, L-1\}$.

3.3.1 Classical simulation

An example of the protocol for a system with $L = 51$ sites and long-time evolution is shown in Figure 3.1. The simulation has been performed using the implementation of the time-dependent variational principle (TDVP) algorithm [62, 63]. Here, we fixed $U = 3$ and filling fraction $n_e := N_e/L = 2/3$, with $N_e = \sum_{i,\sigma} \langle n_{i,\sigma} \rangle_0$ the total number of electrons in the initial state $|\Psi\rangle$. Here, $\langle \dots \rangle_0$ denotes the expectation value on the Fermi-Hubbard model ground state, $|\Psi\rangle$. Once $\langle \hat{\mathbf{S}}^x(t) \rangle = (\langle \hat{S}_0^x(t) \rangle, \dots, \langle \hat{S}_{L-1}^x(t) \rangle)$ (with $\langle \dots \rangle$ the expectation value on $|\Psi_{\frac{\pi}{4}}^j\rangle$) is obtained for all the time steps t_m in the time evolution (with $m = 1, \dots, M$ and $t_{M-1} = T$), the spin excitation spectrum can be resolved by taking a space-time Fourier transform of $\langle \hat{\mathbf{S}}^x(t) \rangle$ and examining the data in the momentum-frequency domain. This gives access to the QSF, as discussed previously in this thesis. In this chapter, all Fourier transforms are computed using the zero-padding technique [163] with respect to time, whereby a number $N'_{\text{samples}} - N_{\text{samples}}$ of zeros is appended to the original time signal. In particular, N'_{samples} is chosen in such a way that $2N_{\text{samples}}$ frequencies are present in the frequency interval $\omega \in [0, 6J]$.

The presence of noise in current near-term quantum computers limits the maximum evolution time that can be reached on a quantum computer and makes long-time evolutions impractical due to the accumulation of errors. Therefore, we aim to determine the shortest evolution time that still allows us to recover the main properties of the spin excitation spectrum. For example, one peculiar feature of the two-spinon excitation spectrum is the dome-shaped structure with zero-frequency cusps at $k = 0$ and $k = \pm 2k_F = \pm \pi n_e$, which can be used in experiments to derive information about the doping of the system. Here, $k_F := \pi n_e/2$ is the Fermi




Figure 3.1: TDVP simulation of the quench protocol for the 1D Fermi-Hubbard model with $L = 51$ sites. **a, c** Dynamics of the observable $\langle \hat{S}_i^x(t) \rangle$ for $U = 3$ and filling $n_e = 2/3$ for total running times $T = 10$ and $T = 3$, respectively. **b, d** Normalised absolute value of the corresponding quench spectral functions in the momentum-frequency domain after applying a two-dimensional Fourier transform to the space-time domain data. The maximum bond dimension for the TDVP simulation is $\chi = 128$.

momentum. From Figure 3.1, we observe that although the resolution decreases and the spectral lines broaden, a shorter evolution time of $T = 3$ is sufficient to recover the spectrum. Shorter times result in greater uncertainty in features such as the precise location of the zero-frequency cusps. However, we argue that this level of accuracy is often sufficient and can significantly reduce computational resource requirements. Furthermore, for smaller systems, the resolution is inherently limited by the system size, meaning that increasing the total evolution time does not necessarily lead to an improved spectrum compared with the thermodynamic limit.

3.3.2 Initial state preparation

To compute the QSF using the protocol introduced in Section 3.2, the quench operator $\hat{Q}_j$ should be applied to the ground state $|\Psi\rangle$ of the 1D Fermi-Hubbard

model. However, preparing the ground state of interacting many-body systems on a quantum simulator is a challenging problem on its own [164, 165] and, in general, costly in terms of circuit implementation [136, 166–168]. On the other hand, as discussed for the local quench spectroscopy protocol in Ref. [86], the spectrum of a given low-energy excitation can be reconstructed by computing the QSF on a different state $|\Psi'\rangle$ such that $\hat{Q}_j |\Psi'\rangle \approx \sum_{\mathbf{q}} \alpha_{\mathbf{q}} \gamma_{\mathbf{q}}^\dagger |\Psi\rangle$, with $\gamma_{\mathbf{q}}^\dagger$ the creation operator in momentum space associated with the excitation under investigation. This holds if the state $\hat{Q}_j |\Psi'\rangle$ contains at most a single excitation. Note that in our case, the low-energy excitations are pairs of spinons. Hence, the above requirement becomes $\hat{Q}_j |\Psi'\rangle \approx \sum_{\mathbf{k}} \alpha_{\mathbf{k}} \sum_{\mathbf{q}} (c_{\mathbf{q},\uparrow}^\dagger c_{\mathbf{k}+\mathbf{q},\downarrow} + c_{\mathbf{q},\downarrow}^\dagger c_{\mathbf{k}+\mathbf{q},\uparrow}) |\Psi\rangle$, where we used the momentum space form of the spin operator $\hat{S}_j^x$.

Here, we examine the performance of the protocol when the quench operator is applied to the ground state of the free-fermion Hamiltonian, $|\psi_{\text{FF}}\rangle$. This alternative is well-suited for quantum computing, as it can be efficiently implemented using Givens rotations [133] and can serve as a good approximation of the Fermi-Hubbard Hamiltonian for small values of U . Remarkably, we find that this approximation remains effective even in the presence of strong interactions.

As we can see in Figure 3.2 (a-b), the energy of the Fermi-Hubbard ground state and the expectation value of H_{FH} on the free-fermion ground state diverge fairly quickly as a function of the on-site interaction strength U , suggesting that the overlap between the two states decreases with increasing U . As can be seen by comparing the blue and orange curves, the expectation value of H_{FH} on the free-fermion ground state is larger than the energy of the state $\hat{Q}_{\lfloor L/2 \rfloor} |\Psi\rangle$ for $U \gtrsim 2$. However, from panel (a) we can see that, for a system with $L = 15$, the state $\hat{Q}_{\lfloor L/2 \rfloor} |\psi_{\text{FF}}\rangle$ (orange curve) is within the energy range of the subspace of H_{FH} containing at most a single excitation (shaded region) up to $U \lesssim 4$. This means that, for larger values of U , $\hat{Q}_{\lfloor L/2 \rfloor} |\psi_{\text{FF}}\rangle$ may contain more elementary excitations, resulting in a QSF with reduced accuracy.

On the other hand, the structural similarity index measure (SSIM) [169, 170] between the spectra obtained through the quench spectroscopy protocol starting




Figure 3.2: Simulation of the quench spectroscopy protocol for a 1D Fermi-Hubbard model initialized in the free-fermion ground state. **a** Energy per site for different initial states in a system with $L = 15$ sites, setting the ground state energy of the Fermi-Hubbard Hamiltonian H_{FH} to zero. The green curve represents the energy of the state obtained by applying the quench to the Fermi-Hubbard ground state, $\hat{Q}_{\lfloor L/2 \rfloor} |\Psi\rangle$, the blue curve shows the expectation value of the Fermi-Hubbard Hamiltonian H_{FH} on the free-fermion ground state $|\psi_{\text{FF}}\rangle$, and the orange curve corresponds to the energy of the state obtained by applying the quench to the free-fermion ground state, $\hat{Q}_{\lfloor L/2 \rfloor} |\psi_{\text{FF}}\rangle$, all as functions of the on-site interaction U . Also shown is the energy range (shaded region) of the subspace spanned by $|\Psi\rangle$ and $\hat{Q}_j |\Psi\rangle$ with $j = 0, \dots, L - 1$, i.e., the low-energy subspace of H_{FH} containing at most a single spin excitation. **b** Same as for (a), but for a system with $L = 51$ sites. **c** Structural similarity index measure (SSIM) between the space-time signals obtained starting from the Fermi-Hubbard and free-fermion ground state with $L = 51$ (blue curve) and their Fourier transforms (yellow curve) for a total evolution time of $T = 10$. Values of the SSIM close to 1 indicate that two images are similar, with $\text{SSIM} = 1$ for identical signals.

from $|\Psi\rangle$ and $|\psi_{\text{FF}}\rangle$ remains close to 1 over a considerable range of U even for a large system size (see Figure 3.2 (c)). We recall that $-1 \leq \text{SSIM} \leq 1$, with $\text{SSIM} = 1$ for two identical images. This suggests that additional excitations does not significantly affect the expectation value of $\hat{S}_i^x(t)$, confirming the robustness of the quench spectroscopy protocol against perturbations to the system's ground state [86, 87] and justifying the use of the free-fermion ground state in our simulations and

subsequent implementation. For a more in-depth comparison between the QSFs obtained initialising the system in the Fermi-Hubbard model and free-fermion ground state see Figure 3.7 in Section 3.6.

We note here that initialising the protocol in the free-fermion ground state instead of the Fermi-Hubbard model's one can be interpreted in two different ways: Either as the exact quenched non-equilibrium spin Green's function, $\langle \Psi_{\text{FF}} | [\hat{S}_k^x(t), \hat{S}_j^x] | \Psi_{\text{FF}} \rangle$, or as an approximation to the equilibrium retarded spin Green's function discussed in Section 3.2. In principle, both contain important information about the nature of the excitations of a many-body quantum system. For instance, in the second interpretation, while the approximation is not formally controlled, it can still reveal the breakdown of the quasiparticle picture, as it is always possible to compute the phase space available to the quasiparticle excitations and compare it with the signal extracted from a quantum computer. Furthermore, at the end of Section 3.5.1 we will discuss how this approximation can be turned into a controlled one at the price of a higher quantum computational cost.

3.3.3 Trotterization

To successfully implement the protocol on a quantum device and obtain the spin excitation spectrum, we need to time-evolve the quenched quantum state. We do this using a product (Trotter) formula, introduced in Chapter 1, which approximates the time-evolution operator of a Hamiltonian $H = \sum H_i$ as a product of its summands $e^{-iH_i t}$. The evolution operator is thus decomposed into a product of N_{Trotter} small time evolution operators, $\mathcal{U}(t) = \exp(-iHt) = \mathcal{U}(\delta t) \dots \mathcal{U}(\delta t)$, with $\delta t = t/N_{\text{Trotter}}$. In the case of the Fermi-Hubbard model and using a first-order Trotter formula, we can further discretise every $\mathcal{U}(\delta t)$ as a product of terms that can be implemented simultaneously on a quantum device, i.e., a product of the interacting part of the Hamiltonian H_{int} , the hopping term on even sites H_0^{even} , and the




Figure 3.3: Trotterized time evolution of the 1D Fermi-Hubbard model with $L = 15$, $U/J = 3$ and $n_e = 2/3$, using the ground state as the initial state. The latter has been obtained using the density-matrix renormalization group (DMRG) algorithm. The exact simulation has been performed using the TDVP algorithm, while the Trotterized simulations have been obtained with time-evolving block-decimation (TEBD) and expressing the Hamiltonian as a product of the interacting part and the hopping term on even and odd sites. DMRG and the time evolution algorithms were executed with maximum bond dimension $\chi = 64$. Top: (a) exact and Trotterized time dynamics with (b) three and (c) five fixed Trotter steps using the first-order Trotter formula. Bottom: (e), (f), (g) absolute value of the QSF obtained by Fourier transforming and normalizing the time-dependent signals in panels (a), (b), (c), respectively. In the last column, (d) SSIM and (h) root mean square error (RMSE) between the QSFs obtained from exact and trotterized time-evolution simulations are shown for increasing numbers of Trotter steps. Both first- and second- order Trotter are simulated.

hopping term on odd sites H_0^{odd} :

$$\begin{aligned} \mathcal{U}(t) = & (e^{-i\delta t H_0^{\text{even}}} e^{-i\delta t H_0^{\text{odd}}} e^{-i\delta t H_{\text{int}}})^{N_{\text{Trotter}}} \\ & + O\left(\frac{t^2}{N_{\text{Trotter}}}\right). \end{aligned} \quad (3.12)$$

Since our main purpose is to find an optimal implementation of the quench spectroscopy protocol in NISQ devices, we will avoid using higher-order Trotter formulas due to higher gate counts and deeper circuits [171]. Hence, in the following, we will only consider first- and second- order Trotter formulas. In view of implementation, it is crucial to determine the number of Trotter steps required to accurately resolve the spectrum. In some cases, a large number of steps may be necessary to reduce errors below a certain threshold, or the required number of steps might scale with the system size. Testing this is essential to ensure that

the protocol remains viable for NISQ devices. In Figure 3.3, we show that the Trotterization error in the spectrum reaches an acceptable threshold with just four Trotter steps for a reasonably large system size. Notably, the first-order Trotter method performs as well as (or even better than) the second-order method, allowing us to justify its use to minimise the depth of the Hamiltonian evolution ansatz.

The number of Trotter steps naturally depends on the total evolution time. For smaller systems, where long evolution times are unnecessary, the number of Trotter steps can be reduced while maintaining the same level of accuracy. In the following, the maximum number of Trotter steps that we use is $N_{\text{Trotter}} = 6$ (see Table 2).

In passing, we note that, alternatively, we could fix the Trotter step size, δt , instead of the number of Trotter steps, N_{Trotter} . Although fixing δt may be beneficial to improve the quality of the space-time signal, maintaining a consistent level of noise across all time values is more effective when performing the Fourier transform. This consistency ensures that, after normalisation, the Fourier transform removes any constant damping factors in a given observable (see Section 3.6.3). For this reason, we opted for the fixed Trotter steps approach.

3.4 Quantum circuit implementation

The encoding of fermions to qubits has an impact on the resources needed for the simulation. In this work, we use the Jordan-Wigner (JW) encoding, which maps every fermionic mode labelled by a site and spin index to one qubit. In this encoding, the creation and annihilation operators become [136, 138]:

$$\begin{aligned} c_i &\rightarrow Z_0 \dots Z_i \frac{X_i + iY_i}{2}, \\ c_i^\dagger &\rightarrow Z_0 \dots Z_i \frac{X_i - iY_i}{2}, \end{aligned} \tag{3.13}$$

where the index $i = 0, \dots, 2L - 1$ now includes both the site and the spin degree of freedom. The hopping term between qubits i and j (with $i < j$) and on-site repulsion thus become

$$\begin{aligned} c_i^\dagger c_j + c_j^\dagger c_i &\rightarrow \frac{1}{2}(X_i X_j + Y_i Y_j) Z_{i+1} \dots Z_{j-1}, \\ n_i n_j = c_i^\dagger c_i c_j^\dagger c_j &\rightarrow \frac{1}{4}(I - Z_i)(I - Z_j). \end{aligned} \tag{3.14}$$




Figure 3.4: Complete circuit for a fermionic system of size $L = 3$. Note the fermion-to-qubit ordering is changed after the initialisation with the Givens rotations. We first encode all the spin up modes in the first L qubits and the spin down modes in the last L qubits. We then move the qubits around with FSWAP gates (represented as swap gates in red) to encode every pair of modes $(i, \uparrow)$ and $(i, \downarrow)$ next to each other. After this reordering, the quench operator $\hat{Q}_j$ is applied to the two central qubits, corresponding to the spin up and down modes of fermion at $i = \lfloor L/2 \rfloor$ site. The time-evolution operator is implemented using a first-order Trotter formula, consisting of even and odd hopping, and on-site interaction terms. Note that due to the (up, down)(up, down) JW ordering, implementing each hopping layers requires 3 layers of FSWAP gates per Trotter step. In the figure, only one Trotter step is shown. Finally, the unitary B_{XY} diagonalizing the spin operator $\hat{S}_i^x$ is applied to each pair of qubits, which are then measured in the computational basis.

A common choice to minimise the weight of the Pauli operators corresponding to the hopping term of a fermionic Hamiltonian is to assign the first L values to the modes with spin up and the last L to the modes with spin down. In what follows, we will refer to this configuration as (all up)(all down). As we will show below, this may not be the optimal choice.

As stated above, our simulation proceeds by the following steps. We start in the ground state of the free-fermion model H_0 (i.e., Equation (3.10) when $U = 0$) for a given electronic occupation N_e which can be prepared efficiently [133, 172]. As we discuss below, such a state can be prepared with a sequence of Givens rotations acting on pairs of adjacent modes of the same spin sector [133]. A Givens rotation is a general rotation operation in the two-dimensional complex subspace. It is composed by a θ -rotation to change the amplitude, and a ϕ -phase operator. However, since we are dealing with real Slater determinants due to the time-reversal symmetry of the Hamiltonian, we can set $\phi = 0$. The rotation matrix for two


adjacent modes labelled 1 and 2 is then given by

$$G(\theta) := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \frac{\theta}{2} & -\sin \frac{\theta}{2} & 0 \\ 0 & \sin \frac{\theta}{2} & \cos \frac{\theta}{2} & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} = e^{i\frac{\theta}{2}(X_1Y_2 - Y_1X_2)}.$$

The operation above can be implemented on two qubits by the rotation

$$G(\theta) = e^{-i\frac{\pi}{4}Z_2}N(\theta/2, 0, \theta/2)e^{i\frac{\pi}{4}Z_2}, \quad (3.15)$$


with $N(\alpha, \beta, \gamma) := e^{i(\alpha X_1X_2 + \beta Y_1Y_2 + \gamma Z_1Z_2)}$. For $\beta = 0$, the two-qubit gate $N(\alpha, \beta, \gamma)$ can be implemented using two CNOTs as [173, 174]



Because the layers of Givens rotations are implemented to each spin sector independently, for this step it is convenient to map the spin up sector to the first L qubits of the JW encoding and the spin down sector to the remaining qubits, i.e., choosing the (all up)(all down) ordering.

Next, we need to find an implementation of the quench operator defined in Equation (3.11) and applied to the central site. The weight of this operator in the JW encoding can be very sensitive to the specific ordering we choose for the fermionic modes. Since we have chosen to start with an (all up)(all down) convention (which is convenient to minimise the depth of the Givens rotations circuit), the two qubits representing the spin up and down modes of the fermion at $j = \lfloor L/2 \rfloor$ are not next to each other. In this case, the quench operator would have weight $k = L + 1$ and would require a decomposition into $2\lceil \log_2(k) \rceil - 1$ two-qubit gates to be implemented [175]. Moreover, while the weight of the encoded on-site term of H_{FH} is not affected by the JW ordering (see Equation (3.14)), it is important to keep in mind that in order to implement such interactions in a quantum computer with local connectivity, such as the IBM device we employ in this work, the qubits involved should be physically close to each other. In general, this is difficult to achieve using the (all up)(all down) JW ordering. To address both the issues, we

can rearrange the qubits to the configuration (up, down)(up, down), where the spin up and spin down modes of each site i are represented by consecutive qubits in the JW encoding. We also establish that this is the most resource-efficient ordering for our time-dynamics simulation (see Appendix 3.A for other possible orderings). To switch from the (all up)(all down) to the (up, down)(up, down) ordering, we introduce a fermionic SWAP (FSWAP) network layer after the Givens rotations which rearranges the initial JW ordering such that the modes $(i, \uparrow)$ and $(i, \downarrow)$ are next to each other [176]. Using IBM's native gate set, each FSWAP can be decomposed as



After applying the FSWAP network, the quench operator is local in the (up, down)(up, down) JW encoding and can be implemented in terms of qubit gates as

$$\hat{Q}_j = e^{\frac{i\theta}{2}(X_j X_{j+1} + Y_j Y_{j+1})}, \quad (3.16)$$

for $\theta = \pi/4$ and $j = \lfloor L/2 \rfloor$. This unitary operator can be implemented on a quantum computer in the same way as the time-evolution operator of the hopping term of the Hamiltonian, $H_{12}(\theta) := e^{\frac{-i\theta}{2}(X_1 X_2 + Y_1 Y_2)} = N(-\theta/2, 0, -\theta/2)$ that, as discussed above, can be done using 2 CNOT gates.

To implement the product formula of Equation (3.12), we note that the even and odd hopping terms of the Hamiltonian H_0 in Equation (3.10), H_0^{even} and H_0^{odd} (defined by constraining the summation over sites in Equation (3.10) to even and odd values of i , respectively), act on adjacent fermionic modes with the same spin, which, in our (up, down)(up, down) JW ordering, are represented by qubits that are one qubit apart. Consequently, we need an additional layer of FSWAPS both before and after each $H(\theta)$. On another hand, the time-evolution operator associated with the on-site interaction part of the Hamiltonian, H_{int} , which act between two qubits representing the two spin modes of the same fermion site, can

be implemented with the onsite gate $O(\phi)$,

$$\begin{array}{c} \text{---} \bullet \text{---} \\ | \\ \boxed{O(\phi)} \\ | \\ \text{---} \end{array} = \begin{array}{c} \text{---} \bullet \text{---} \\ | \\ \boxed{e^{-i\phi}} \\ | \\ \text{---} \end{array},$$

which in terms of CNOT gates becomes

$$\begin{array}{c} \text{---} \boxed{R_z(\frac{\phi}{2})} \text{---} \bullet \text{---} \bullet \text{---} \\ | \quad | \\ \oplus \quad \oplus \\ \text{---} \boxed{R_z(-\frac{\phi}{2})} \text{---} \boxed{R_z(\frac{\phi}{2})} \text{---} \end{array}.$$

Following the Trotter decomposition in Equation (3.12), we repeat the application of the even and odd hopping terms, and the onsite term, N_{Trotter} times.

After the Trotterized time evolution, we need to make the necessary measurements on the time-evolved state to infer the quantities of interest. The observable $\langle \hat{S}_i^x \rangle = \frac{1}{2} \langle X_{i\uparrow} X_{i\downarrow} + Y_{i\uparrow} Y_{i\downarrow} \rangle$ needs to be measured for each pair of fermionic modes $(i, \uparrow)$ and $(i, \downarrow)$. Here, $X_{i\sigma}$ and $Y_{i\sigma}$ denote Pauli operators acting on the qubit encoding the fermionic mode (i, σ) . Recalling that these modes are next to each other in the (up, down)(up, down) JW ordering of the qubits, we can apply the following unitary before the measurement [136]

$$\boxed{B_{XY}} = \begin{array}{c} \text{---} \bullet \text{---} \boxed{H} \text{---} \bullet \text{---} \\ | \quad | \\ \oplus \quad \oplus \\ \text{---} \bullet \text{---} \end{array},$$

which can be implemented with 3 CNOT gates and diagonalizes $\hat{S}_i^x$ as $|01\rangle \langle 01| - |10\rangle \langle 10|$, allowing us to compute the expectation value of $\hat{S}_i^x$ from the measurement results.

Importantly, measuring the two-point correlators $\langle X_i X_{i+1} + Y_i Y_{i+1} \rangle$ using the above change of basis allows us to exploit the particle number conservation of the Fermi-Hubbard Hamiltonian for post-selection as a simple error mitigation technique [138]. In particular, since in experiments on real quantum hardware we run many shots of the same circuit, we can retain only those bit strings containing the correct number of fermions, providing an effective way to filter out noisy shots.

In Figure 3.4, we show a simple circuit with all the components discussed in this section.

3.5 Approximate free-fermion initial state preparation

In the previous sections, we initialised the quench spectroscopy protocol to the free-fermion model ground state, as it provides a hardware-compatible alternative to the Fermi-Hubbard ground state for the circuit ansatz. However, the circuit to prepare the free-fermion ground state exactly with Givens rotations requires $O(L^2)$ gates at depth $O(L)$ [133]. In practical implementations, this depletes a significant amount of quantum resources. We propose here the Dense Givens Approximate state preparation scheme which is a simple, alternative approach that reduces quantum resource requirements while maintaining the quench spectroscopy algorithm's performance. Here, the free-fermion ground state is approximated with n_{layers} layers of Givens rotations. Each layer is maximally dense in Givens gates whereby every qubit is connected to its two nearest physical neighbours. Note that each layer has constant depth with respect to the system size. A schematic of such a single layer is shown in Figure 3.5 for a single spin sector with $L = 6$ at half-filling. The Givens angles are obtained by using the Variational Quantum Eigensolver (VQE) algorithm [139, 166, 177, 178] either by minimising the energy with respect to the free-fermion Hamiltonian or by maximising the overlap fidelity with the exact free-fermion ground state. For a desired occupation N_e , the fermions are initially maximally spread across both sectors, and this choice intuitively comes from the fact that the true ground state is homogeneous in particle density.

3.5.1 Practical implementation

The advantage of using the DGA scheme is that we can calibrate the number of Givens rotation layers included in the approximation. Regardless of system size, each layer adds a fixed two-qubit depth of eight. Naturally, adding more layers improves fidelity, but our simulations show that for the system sizes considered in this work, no more than three layers are needed to achieve a SSIM ≥ 0.9 when comparing the resulting spectrum to the exact one. This is shown in Figure 3.6, where we compare the dynamics of $\langle \hat{S}^x(t) \rangle$ and the QSF obtained starting the




Figure 3.5: Approximate free-fermion initial state preparation for a Fermi-Hubbard model with $L = 6$ in a single spin sector at half filling ($N_e = 3$) with a single layer of the DGA state preparation scheme. The gates $G(\theta_i)$ are Givens rotations whose angles are optimised via VQE.

quench protocol from the free-fermion ground state (panels (a) and (e)), with the ones obtained using a DGA approximate state with $n_{\text{layers}} = 2$ (panels (b) and (f)) and $n_{\text{layers}} = 3$ (panels (c) and (d)). This allows us to significantly reduce the depth of the circuit without compromising the final outcome of the protocol, making the DGA state an efficient choice.

To evaluate the quality of the initial state approximation, we also examine the fidelity $F_{\text{DGA}} = |\langle \psi_{\text{DGA}} | \psi_{\text{FF}} \rangle|^2$ between the approximate DGA state, $|\psi_{\text{DGA}}\rangle$, and the free-fermion ground state $|\psi_{\text{FF}}\rangle$, across different system sizes. As can be seen in Figure 3.6 (d), the fidelity for a given number of layers n_{layers} decreases with increasing system size.

We observe a clear correlation between the fidelity of the initial state and the SSIM of the resulting spectrum obtained at the end of the quench protocol: the better the initial state approximation, the closer the resulting spectrum is to the exact one.

As a result of the observations above, the number of layers required to maintain fidelity at an acceptable level increases with system size. However, this increase in circuit depth is significantly smaller compared to the growth that would result from using the full Givens rotation layers required for the exact free-fermion ground state preparation. For these reasons, we conclude that the DGA approximation of the




Figure 3.6: Trotterised time evolution with $N_{\text{Trotter}} = 5$ Trotter steps of the Fermi-Hubbard model with $L = 15$, $U = 3$ and $n_e = 2/3$, using the DGA approximation of the free-fermion ground state as the initial state. Top: time dynamics starting with the (a) free-fermion ground state and DGA free-fermion state with (b) two and (c) three layers. Bottom: QSFs corresponding to panels (a-c) obtained after a space-time Fourier transform. In the last column, the fidelity of the approximate DGA free-fermion state with respect to the exact free-fermion ground state is shown for different number of DGA layers, n_{layers} . The fidelity is shown as a function of (d) the system size for a fermion filling fraction close to $n_e = 2/3$ and (f) the electronic occupation N_e for a fixed system size $L = 15$. Data in panels (a-c) and (e-g) have been obtained by an exact simulation of the quantum circuits with 4096 shots. Errors in the SSIM are computed using the bootstrap technique with 10^4 instances [179].

free-fermion ground state serves as an effective initial state for our protocol, as it enables good (and, in some cases, better) results while requiring shallower circuits.

In Section 3.3 we have discussed how it is possible to retrieve a good approximation of the spin excitation spectrum of the 1D Fermi-Hubbard model using the unitary quench spectroscopy protocol by replacing the ground state of the Fermi-Hubbard model with the one of the corresponding free-fermion model. In this section, we have introduced the DGA scheme to prepare a high-fidelity approximation of the latter, allowing us to allocate the majority of the quantum resources to the time-evolution part of the protocol. It is important to note that, in principle, the state preparation part of the quantum circuit can be replaced by a VQE ansatz to directly prepare an approximation of the ground state of the Fermi-Hubbard model. This would make the initial state preparation a controlled approximation: as the

number of VQE layers increases, the variational state is expected to converge to the real ground state [136, 177, 178]. However, this approach would push the required resources beyond the capabilities of current hardware. For example, implementing the VQE algorithm to prepare the ground state of the Fermi-Hubbard model within the framework of the Hamiltonian Variational Ansatz (HVA) requires initialising the variational ansatz to the ground state of the corresponding free-fermion model [132, 136, 138]. Therefore, in the next generation of NISQ devices, one can consider adding a few layers of the HVA after the DGA part of the circuit. We expect that that this will improve the efficacy of our algorithm in the strongly interacting regime.

3.6 Results

3.6.1 Simulations with different protocols

In this section, we begin by presenting a comprehensive analysis of the two-spinon spectrum for the 1D Fermi-Hubbard model obtained with different approaches and on-site interaction strengths. The results are shown in Figure 3.7. Here, the first row shows the QSF obtained starting the quench spectroscopy protocol from the ground state of the Fermi-Hubbard model, $|\Psi\rangle$, as given by Equation (3.9). As discussed in Section 3.3, one possibility to reduce the resources to implement the protocol on a real quantum device is to apply the quench to the ground state of the free-fermion model, $|\psi_{\text{FF}}\rangle$, in order to save computational resources for the time-dynamics simulation. As discussed there, we expect this protocol to give a good approximation of the system's excitation spectrum if only a single spin excitation (i.e., a single pair of spinons) is present in the quenched state, which occurs for small values of U (see Figure 3.2 (a)). By comparing the first and second rows in Figure 3.7, we can see that the QSF computed using the Fermi-Hubbard ground state and the one obtained by applying the quench to the free-fermion ground state are in excellent agreement for $U < 4$. Remarkably, the agreement between the two QSFs remains good even for larger value of U , where all the main features of the dynamical spin structure factor are still visible in the approximate QSF (such as the softening of the spin excitations for large U and the cusp with $\omega \rightarrow 0$ for $k \rightarrow \pm 2k_F$).

In particular, we notice that the main discrepancies arise for $k \rightarrow \pm 2k_F$. This is expected since for these values of k spin excitations become gapless and many pairs of spinons with $k \approx \pm 2k_F$ can potentially be excited with a vanishing amount of energy, leading to a breakdown of the quench spectroscopy protocol.

As discussed above, we can use the DGA state as an alternative to the free-fermion ground state to further reduce circuit depth and number of two-qubit gates, making it more suitable for near-term hardware implementations. As long as the fidelity of the DGA state, $|\psi_{\text{DGA}}\rangle$, with respect to the free-fermion ground state is sufficiently high, the QSF computed using the DGA state captures all the key spectral features of the spin excitation spectrum with an accuracy comparable with the QSF obtained using the free-fermion ground state for $U < 6$. Note that in panels (k)-(t) of Figure 3.7 we also incorporate the effects of Trotterization, which are the major responsible for the loss of quality of the approximate QSF for larger values of U .

Finally, the last row of Figure 3.7 shows the QSF computed using the DGA initial state and the parameters that we used in our experiments on IBM hardware, described in Section 3.6.3 below. Here, in order to reduce circuit depth and two-qubit gate counts, we run simulations up to $T = 3$ using six Trotter steps. These parameters are sufficient to recover a satisfactory spin excitation spectrum for $U = 3$ (see Figure 3.8). On the other hand, to access larger values of U more Trotter steps and a longer time evolution are needed. These observations emphasise the balance required between initial DGA state fidelity, circuit depth and two-qubit gate counts, and time-evolution accuracy to optimise the performance of the unitary quench protocol.

3.6.2 Resource estimation

Following the gate decomposition and circuit ansatz described in Section 3.4, we can determine the gate count and circuit depth. These metrics depend on the chosen qubit ordering. For example, operations like the quench must be applied to adjacent qubits to avoid handling the JW string. Therefore, if the selected ordering places modes $(\lfloor L/2 \rfloor, \uparrow)$ and $(\lfloor L/2 \rfloor, \downarrow)$ on distant qubits, as in the (all




Figure 3.7: QSFs for various protocols and values of U . Panels (a)-(e) display the QSF obtained by applying the unitary quench spectroscopy protocol to the ground state of the 1D Fermi-Hubbard model, (f)-(j) show results for a quench from the free-fermion model ground state, and (k)-(o) correspond to a quench from the DGA state with $n_{\text{layers}} = 3$ using a Trotterized dynamics with $N_{\text{Trotter}} = 10$ Trotter steps. All results in (a)-(t) are obtained by taking the absolute value of the normalized Fourier transform of $\langle \hat{S}^x(t) \rangle$ after a total evolution time $T = 5$ on a system with $L = 15$. The final row (p)-(t) shows the QSF obtained using the DGA approximate state, using $T = 3$, $n_{\text{layers}} = 3$ DGA layers, and $N_{\text{Trotter}} = 6$ Trotter steps, matching the settings used for hardware results discussed in Section 3.6.3. In all panels, the red dots denote the maximum of the QSF for each value of k . Data in panels (a)-(j) have been obtained using the TDVP algorithm with maximum bond dimension $\chi = 128$. Data in panels (k)-(t) are obtained by simulation of the quantum circuit with 4096 shots.

up)(all down) encoding we discussed in Section 3.4, we need to introduce a layer of FSWAPs to bring them together, which increases both gate count and circuit depth. Specifically, each FSWAP layer adds a factor of $O(L)$ to the depth.

As discussed, due to the constraints and characteristics of the IBM hardware, we prioritise the (up, down)(up, down) ordering shown in Figure 3.4 to perform the quench, the time evolution, and the measurement part of the unitary quench spectroscopy protocol. However, we start from the (all up)(all down) ordering in which the initialisation part of the circuit becomes straightforward, as the Givens rotations are applied in independent layers to the spin-up and spin-down sectors. Specifically, for the initialisation to the free-fermion ground state, $|\psi_{\text{FF}}\rangle$, we have

$N_G = (L - N_e/2)N_e$ gates [133]. This number is maximised at $L^2/2$ in the half-filled case, where $N_e = L$. We use this maximum to establish an upper bound on the gate count for any initial state. This gate block consists of $L - 1$ layers of Givens gates, and since each Givens rotation can be implemented with a depth of 2 two-qubit gates, this section of the circuit has a total depth of $2(L - 1)$. As discussed above, we can instead use a n_{layers} DGA approximation of the free-fermion ground state that brings the depth of this part of the circuit down to $4n_{\text{layers}}$. We connect every pair of fermions in every Givens layer (as in Figure 3.5) so the total number of two-qubit gates in this case is $2(L - 1)n_{\text{layers}}$.

We then require an FSWAP layer consisting of $(L^2 - L)/2$ gates and a two-qubit gate depth of $2(L - 1)$ to rearrange the qubits into the (up, down)(up, down) ordering. The subsequent quench operator adds only a two-qubit gate depth factor of two. Likewise, each Hamiltonian evolution gate has a depth of two. Each Trotter step consists of three layers (even and odd hopping, and on-site interaction terms) that can be implemented simultaneously. Taking into account the FSWAPs applied before and after the each hopping term, the total two-qubit gate depth for the time evolution part amounts to $12N_{\text{Trotter}}$.

Finally, the unitary operator B_{XY} , which diagonalises the spin operator $\hat{S}_i^x$, is applied to each pair of qubits. As B_{XY} can be implemented with three CNOT gates, this final layer contributes a constant two-qubit gate depth of 3 and a total of $3L$ two-qubit gates.

The advantage of the (up, down)(up, down) encoding is that it eliminates the need to rearrange the qubits before measurement, as all modes associated with the same fermionic sites are already paired adjacently. Additionally, in this scheme, the on-site interaction gates are applied locally, meaning no additional CNOT layers are required when compiling to a real device with local connectivity. Consequently, the total circuit depth sums to $2L + 12N_{\text{Trotter}} + 4n_{\text{layers}} + 3$, and the number of two-qubit gates is $L^2 + L(10N_{\text{Trotter}} + 2n_{\text{layers}} + 2) + (2 - 8N_{\text{Trotter}} - 2n_{\text{layers}})$. Alternatively, the gate count and depth of the (all up)(all down) encoding are presented in Appendix 3.A, where they are found to be larger and therefore less resource-efficient.

3.6.3 Quantum hardware results

In this section, we present the spin excitation spectrum of the 1D Fermi-Hubbard model obtained using the unitary quench spectroscopy algorithm on IBM quantum hardware. The circuit in Figure 3.4 is implemented starting from the approximate free-fermion state obtained via the DGA scheme described in Section 3.5, with $n_{\text{layers}} = 2$ for $L \in \{9, 11, 13\}$, and $n_{\text{layers}} = 3$ for $L = 15$. The number of layers for the DGA scheme is chosen to ensure that the fidelity between the DGA and the free-fermion ground state is close to a threshold of 0.9 (see Table 2 for details). Interestingly, for certain system sizes, we find that even when the fidelity falls below this threshold, the spectrum is still recoverable, yielding excellent results. Although some spectral features may appear weaker when using a shallower DGA ansatz with lower fidelity, they can often be enhanced by increasing the number of Trotter steps as discussed in Section 3.6.1. After initialisation, we apply the FSWAP layer, the quench operator, the Trotterized time evolution, and finally we measure in the $XX + YY$ basis, enabling particle number post-processing (see Section 3.4). In order to mitigate noise in the quantum simulations, we employ techniques such as an XX dynamical decoupling sequence [180, 181], and Pauli twirling on two-qubit gates and measurements [117] as implemented in the Qiskit SDK [182]. In particular, we expect the synergy between the quench spectroscopy protocol and Pauli twirling to make the algorithm very robust to noise. Pauli twirling converts coherent noise into Pauli noise [117, 183] which effectively acts on single-site Pauli operators P_i as a depolarising channel, $\langle P_i(t) \rangle_{\text{noisy}} = [1 - p_i(t)] \langle P_i(t) \rangle$, with $p_i(t)$ a time-dependent error rate. If we assume that $p_i(t) \approx p$, $\forall i, t$, then the noisy space-time signal retrieved from a quench experiment will be $\langle \hat{\mathbf{S}}^x(t) \rangle_{\text{noisy}} \approx (1 - p) \langle \hat{\mathbf{S}}^x(t) \rangle$, since $\langle \hat{\mathbf{S}}^x(t) \rangle$ consists of single-site Pauli operators. Since the QSF is normalised so that the maximum of its absolute value is one, the constant damping factor $(1 - p)$ does not affect its features, regardless of the damping of the space-time signal. To ensure that $p_i(t)$ is as constant as possible for the different values of t , in this work we use the same number of Trotter steps $\forall t \in [0, T]$. On the




Figure 3.8: Simulation and hardware results from quench spectroscopy simulations of the 1D Fermi-Hubbard model for system sizes $L = 9, 11, 13, 15$. The circuit details and the parameters used in these experiments are reported in Table 1 and Table 2, respectively. **a-d** Spectra from noiseless simulations of the quantum circuits. **e-h** Spectra from experiments. All the circuits are executed on the IBMQ device `ibm_fez` using a total of 10^6 shots and 8 Pauli twirling randomizations per circuit. **i-l** Histograms showing the distribution of counts across different fermion numbers at the end of the computation for the same system sizes, with the initial fermion count N_e highlighted in orange. In panels (**a-h**), the theoretical position of the cusp of the spin excitation spectrum at $2k_F = \pm\pi n_e$ is marked by a dashed line, while the red dots denote the maximum of the QSF for each value of k . In panels (**e-h**) errors on the SSIM have been computed using the bootstrap technique over 10^4 instances [179].

other hand, thanks to Pauli twirling, we expect $p_i(t)$ to be homogeneous on all qubits to a good degree of approximation.

Figure 3.8 shows the results obtained from implementing the unitary quench spectroscopy circuit on the IBM quantum processor `ibm_fez`. The recovered spectra are in excellent agreement with the theoretical predictions for all the system sizes we considered, ranging from $L = 9$ (with 18 qubits) to $L = 15$ (with 30 qubits). As the system size L increases, errors in the computation due to hardware noise become more apparent. As observed in the particle number histograms in panels (i-l), computational errors can cause the inclusion of spectral signals from unintended




Figure 3.9: Hardware results without post-selection. Quench spectroscopy simulations of the 1D Fermi-Hubbard model are presented for system sizes $L = 9, 11, 13, 15$ (from left to right). The circuit details and parameters are the same as in Figure 3.8 and reported in Tables 1 and 2.

particle number sectors, potentially interfering with the reconstruction of the spin excitation spectrum. The distribution of counts in the histograms is compatible with noise driving the system to a maximally mixed state, where the bit strings distribute as a Gaussian around half filling $N_e = L$. On the other hand, readout errors on short circuits, where the system state does not become fully mixed, could explain the skewness of the distribution. In this work we exploit the particle number conservation of the Fermi-Hubbard model to post-select only those bit strings containing the correct number of fermions (see the orange bars in the histograms in Figure 3.8 (i-l)). This results in an overall improvement of the quality of the reconstructed spectra. However, we also note that the protocol has some intrinsic resilience to bit-flip errors, which are the main responsible for the change in fermionic particle numbers. As an example, in Figure 3.9 we show the hardware results of panels (e-h) in Figure 3.8 with-out particle number post-selection. Intuitively, a bit-flip error channel acting on the time-evolved density matrix induces a multiplicative error in the QSF, which leaves its structure unaffected once normalised – analogous to the behaviour observed earlier for Pauli twirling and the depolarising channel. This fact makes this approach suitable even for trapped ions and neutral-atom architectures, where collecting a large number of samples is often impractical due to their inherently slower operation speed [140].

Due to the nature of the noise and the small depth of the circuit, we expect errors in the signal $\langle \hat{S}^x(t) \rangle$ to be localised in space and time. In this case, they would mainly affect the short-wavelength/high-energy region of the spectrum, as can

L	N	Depth	Gates
9	18	89	543
11	22	105	797
13	26	109	977
15	30	117	1193

Table 1: Number of qubits, two-qubit depth and total number of two-qubit gates for the experiments shown in Figure 3.8.

L	T	N_e	n_{layers}	F_{DGA}	N_{Trotter}
9	3	6	2	0.93	5
11	3	6	2	0.90	6
13	3.5	6	2	0.88	6
15	3	10	3	0.95	6

Table 2: System size L , maximum evolution time T , fermion occupation number N_e , number of Trotter steps N_{Trotter} , number of DGA layers n_{layers} , and fidelity between the approximate DGA state and the free-fermion ground state for the experiments shown in Figure 3.8.

be seen, for instance, from the vertical stripes at large momenta $k \approx \pm\pi$ on the sides of the $L = 13$ and $L = 15$ spectra in panels (g) and (h) of Figure 3.8, respectively. Despite this, we can successfully recover the main features of the spin excitation spectrum, such as correctly identifying the cusp for $k = \pm 2k_F$ and locating the maximum of the (absolute value of the) QSF as a function of k . In addition to the long-wavelength/high-energy noise, we also note some discrepancies in the position of the maximum of the QSF in the recovered spectra for $k \approx \pm 2k_F$. In particular, see the case with $L = 9$ in panels (a) and (e) in Figure 3.8. These can be explained by noting that even in the simulated results in panels (a-d) the QSF for $k \approx \pm 2k_F$ is spread over a large window of energy values. Therefore, we expect the position of the maximum of the QSF to be sensitive to small fluctuations in the space-time signal. Moreover, as discussed at the beginning of Section 3.6, in this region the spin excitations are almost gapless and deviations from the two-spinon continuum are more likely due to the presence of a potentially larger number of spin excitations.

Note that for the simulation with $L = 13$ we choose a smaller occupation number N_e because due to the limited momentum resolution even the spectrum obtained

from a noiseless simulation does not match the one obtained from the corresponding TDVP simulation. We expect this issue to become negligible for large system sizes. On the other hand, we take advantage of this change in occupation number to show how the protocol is capable of correctly capturing the softening of the spin excitations at lower densities even when it is executed on quantum hardware.

3.7 Discussion

In this chapter, we introduced the unitary quench spectroscopy protocol which allowed us to obtain the spectral properties of the spin excitations of a 1D Fermi-Hubbard model from a digital quantum processor, without resorting to auxiliary qubits or extra measurements. We exploited the intrinsic robustness of the protocol to perturbations to the initial state to significantly reduce the resources required to prepare the initial state. This has been achieved by introducing the DGA scheme to prepare a high-fidelity approximation of a free-fermion ground state using a shallow circuit with depth independent of the system's size. Crucially, this allowed us to save the majority of the resources for the time-dynamics simulations.

Since spectral properties of quasiparticle excitations are encoded into the peaks of the corresponding retarded Green's function, dynamical structure factor, or quench spectral function, they can be extracted from normalized signals. This makes it possible to effectively mitigate the effects of global depolarizing noise. In this context, we demonstrated that the unitary quench spectroscopy protocol is extremely effective when combined with simple noise-tailoring techniques, such as Pauli twirling, which turns an unknown noise channel into a Pauli channel. All of this, combined with a careful choice of the Jordan-Wigner ordering of the fermionic modes, allowed us to extract useful information about the spin excitation spectra of large instances of the 1D Fermi-Hubbard model [135, 138, 140, 142]. Moreover, the two-qubit gate depth and count of the circuits we considered are not far from state-of-the-art experiments implementing sophisticated error mitigation techniques [19, 184]. We have therefore established that even in the presence of noise, useful spectroscopic information of fermionic systems can be extracted from noisy quantum devices.

Appendices

3.A (all up)(all down) JW ordering

In the main text, we employ the (up, down)(up, down) JW ordering to encode fermionic modes into qubits. However, alternative encodings may be more efficient depending on the hardware. For example, the (all up, all down) encoding can be utilised not only during initialization (as discussed in Section 3.4 and shown in Figure 3.4) but throughout the entire computation. In this scheme, all fermionic modes with spin up are encoded in the first L qubits, while the spin down modes are encoded in the remaining qubits.




Figure 3.A.1: Complete ansatz circuit for a fermionic system of size $L = 3$ in the (all up)(all down) JW ordering, where we use we encode all the spin up modes in the first L qubits and the spin down modes in the last L .

Here, the initialisation part of the circuit is straightforward because the Givens rotations are applied in layers independently to the spin up and spin down sector. This part of the circuit can be implemented with $L - 1$ layers of Givens gates. Since every Givens rotation can be implemented with depth 2, this part of the circuit has two-qubit gate depth $2(L - 1)$.

After that, we need a FSWAP layer to bring the modes $(\lfloor L/2 \rfloor, \uparrow)$ and $(\lfloor L/2 \rfloor, \downarrow)$ next to each other (both in the JW encoding and in the hardware) to implement

the quench operator. Then, the FSWAP layer is applied again to go back to the original ordering. Every FSWAP layer contains $L - 1$ FSWAPs. If we assume that the FSWAP gate and $\hat{Q}_{\lfloor L/2 \rfloor} = H(\pi/2)$ can be implemented with two-qubit gate depth 2, then the part of the circuit used to implement the quench will have two-qubit gate depth $4(L - 1) + 2 = 4L - 2$.

Similarly, assuming all-to-all connectivity, every term of the Trotterized time evolution has two-qubit gate depth 2. Each Trotter step consists of three layers (odd hoppings, even hoppings, and on-site interactions) that can be implemented simultaneously. Therefore, the circuit for a time evolution implemented using a first-order Trotter formula with N_{Trotter} Trotter steps has overall two-qubit gate depth $6N_{\text{Trotter}}$.

Finally, we need an extra FSWAP layer (this time consisting of $(L^2 - L)/2$ FSWAP gates, but still with two-qubit gate depth $2(L - 1)$) to reorder the qubits so that all pairs of qubits corresponding to fermionic modes on a given site i and opposite spin are next to each other. We then measure all of them in the $(XX + YY)$ -basis by applying the unitary B_{XY} to each pair of qubits. This is implemented with constant two-qubit gate depth 3 and total number of two-qubit gates $3L$. The total two-qubit gate depth of the ansatz circuit will then be $8L + 6N_{\text{Trotter}} - 3$. The total circuit is shown in Figure 3.A.1.

Alternatively, we can leave out the last FSWAP layer at the price of increasing the number of circuits we need to run. That is, for every observable $\hat{S}_i^x$ with $i \in \{0, \dots, L - 1\}$, we can consider a circuit where we measure the high-weight Pauli strings $\langle X_i Z_{i+1} \dots Z_{i+L-1} X_{i+L} \rangle$ and $\langle Y_i Z_{i+1} \dots Z_{i+L-1} Y_{i+L} \rangle$, and add them together. Note that this strings arise as a consequence of JW encoding when two fermionic modes are not adjacent to each other. In this case, the depth will be $6L + 6N_{\text{Trotter}} - 4$ instead. However, the sampling overhead for this approach grows with the system size and the expectation values $\langle \hat{S}_i^x \rangle$ are more sensitive to noise. Moreover, since in this case one has to measure the two parts of $\langle \hat{S}_i^x \rangle$ independently, post-selection on fermionic occupation number is not possible. For all these reasons, we have not considered this approach in this work.

Note that, using the DGA scheme to prepare an approximate free-fermion ground state, the overall two-qubit gate depths become $4n_{\text{layers}} + 6L + 6N_{\text{Trotter}} - 1$ and $4n_{\text{layers}} + 4L + 6N_{\text{Trotter}} - 2$ with and without the last FSWAP layer, respectively.

Comparing these resource requirements with the one discussed in Section 3.4, we can see that the (all up)(all down) encoding is more efficient in the presence of small system size and large number of Trotter steps. For the values of L and N_{Trotter} considered in this work (see Table 2), we have found that the (up, down)(up, down) JW ordering is always the most convenient. Moreover, one has to keep in mind that in order to implement the time-evolution operator associated with H_{int} on a device with local connectivity, as the one we consider in this work, the interacting qubits must be physically close to each other. In the (all up)(all down) encoding, this requires additional SWAP gates. On the other hand, the implementation in the (up, down)(up, down) encoding is straightforward.

4

Quantum simulation of open system dynamics

In this chapter, we investigate the computational complexity of simulating open quantum systems, with a focus on two physically-relevant parameters, correlation length and mixing time, and explore whether it can be comparable (or even lower) to that of simulating their closed counterparts. Our results suggest that dissipation reduces complexity, making open quantum systems easier to simulate at long times. Moreover, we examine the classical complexity for tensor network simulations of the open system and compare it with the required quantum resources within the Lindbladian formalism. We observe numerically that in classical tensor network simulations, classical complexity does not decrease with stronger dissipation, revealing a separation between quantum and classical resource scaling.

This chapter is based on a publication in collaboration with Alexander Yosifov, and Jinzhao Sun [185]. The author of this thesis contributed to formulating the central idea and framework, developed all computational code, performed all numerical simulations, and wrote the manuscript with contributions from A.Y. and J.S. The author was also involved in the theoretical developments in Appendix 4.B, while the derivation of the sampling complexity was carried out by J.S., with input from the author.

4.1 Introduction

In the previous part of this thesis, we focused on simulating the dynamics of isolated quantum systems. Yet in reality, no system is truly closed: all quantum systems inevitably interact with their surrounding environment [186, 187]. In this final chapter, we therefore broaden our scope to study the resulting non-unitary evolution, known as open quantum dynamics. Such dynamics have attracted growing interest for their central role in modelling non-equilibrium processes across a wide range of contexts, from biology [188, 189] and materials science [190], to thermodynamics [191, 192] and quantum control [193, 194]. However, simulating open quantum systems remains a significant challenge, as the coupling to an environment induces non-unitary evolution governed by dissipative mechanisms such as relaxation and thermalisation [195–198].

The most widely adopted methods for simulating open quantum systems (OQS) rely on dilation techniques, such as Stinespring dilation [54, 199] and purification methods [200]. These techniques generally recast the non-unitary dynamics of an OQS into a unitary evolution within an enlarged Hilbert space that explicitly includes the environment. By subsequently tracing out the environmental degrees of freedom, one recovers the original non-unitary dynamics of the reduced system. Such approaches have significantly advanced algorithmic design for OQS, enabling simulations of noisy quantum channels [201] and facilitating ancilla-based implementations on contemporary noisy quantum hardware [202]. However, dilation-based methods can mask the intrinsic complexity of the reduced system’s evolution, making it challenging to clearly differentiate inherent complexity from artifacts introduced by the dilation process itself.

Recent research has also explored how dissipation impacts the complexity of OQS dynamics. While dissipation typically leads to coherence loss and the suppression of entanglement [203–206], it has been observed to increase complexity in certain scenarios. Prosen and Pizorn [207], for example, identified that OQS interacting with Markovian environments far from equilibrium can exhibit quantum phase transitions which lead to extensive correlations in the non-equilibrium steady state, driving a

linear increase in the operator-space entanglement entropy. Moreover, Aloisio et al. [208] found that sampling from classical stochastic quantum processes involving multiple time steps can match the computational complexity of sampling complex quantum many-body states. Bhattacharjee et al. [209] further demonstrated that dissipation's effect on Krylov complexity in the SYK model involves a nuanced interplay: an initial phase characterised by exponential complexity growth ultimately transitions to a dissipative suppression phase, indicating significant intermediate computational costs.

Thus, an essential open question is whether we can directly access and characterise the intrinsic complexity of OQS dynamics and precisely determine how dissipation contributes to complexity, using experimentally accessible and physically meaningful parameters without relying on dilation-based techniques.

In this chapter, we address the issue of quantifying and understanding the complexity of simulating OQS, focusing on physically meaningful parameters – namely, the correlation length and mixing time – across both short- and long-time regimes [210–213], as suggested in the context of quantum imaginary-time evolution (QITE) algorithms [32]. We employ two complementary approaches: (1) classical simulations using tensor networks (TNs), and (2) quantum simulations based on time-dependent [214] and nonlinear Lindbladian master equations [215]. In the classical case, complexity is quantified by the bond dimension and operator entanglement entropy, which reflect the growth of correlations and the computational resources required to efficiently represent system states. In the quantum context, we simulate local spin observable evolution in dissipative spin chains, interpreting complexity through the mixing time, defined as the duration for the OQS to reach a steady state under specific simulation algorithms. This measure provides an upper bound for the algorithmic complexity rather than directly characterising the inherent dynamics of the system. We are particularly interested in how temporal computational resources scale with respect to dissipation strength, temporal duration, and the initial conditions of the system.

We show both analytically and numerically that increasing dissipation reduces the correlation length over time, indicating that OQS dynamics may be easier to simulate at long times than the real-time evolution of Hamiltonian systems. In stark contrast to the quantum case, we observe numerically that the classical simulation complexity is not affected by the growth of dissipation. This suggests the existence of regions, where quantum methods are advantageous over classical techniques. Finally, we consider an extension to ground-state preparation under time-dependent Lindbladians, where the jump operators are applied randomly at different sites over time in contrast to the dissipative methods in Refs. [216, 217]. These results can help identify physically meaningful regimes where quantum algorithms outperform classical methods in simulating OQS.

4.2 Methods

4.2.1 Formulation of open quantum dynamics

Many quantum processes can be captured by quantum Markov dynamics, where the evolution is recast into a quantum dynamical semigroup that generalises the Schrödinger equation. We consider a quantum system S , described by Hamiltonian H and density operator ρ , which are Hermitian in the operator space $\mathcal{A}(\mathcal{H})$ with Hilbert space $\mathcal{H}$; and suppose the evolution of S is monitored by a detector with efficiency $\eta_k \in [0, 1]$, see Ref. [215]. The detector term is a new feature compared to the Lindblad equations discussed earlier in this thesis. It introduces non-linear behaviour, providing a more complete description of the open dynamics. The linear form is recovered by setting $\eta_k = 0$. Formally, the dynamics of S is governed by a Lindblad master equation $\partial_t \rho = \mathcal{L}[\rho]$, where the Lindbladian is given by

$$\mathcal{L}[\rho] = -i[H, \rho] + \gamma \mathcal{N}[\rho], \quad (4.1)$$

which consists of a unitary and dissipative components, denoted by $-i[H, \rho]$ and $\gamma \mathcal{N}[\rho]$, respectively, and $\gamma \in \mathbb{R}^+$ is the interaction strength which dictates

the rate of dissipation. The detector efficiency implies a nonlinear term in the Lindbladian. Explicitly

$$\mathcal{N}[\rho] = \sum_k \alpha_k \left(-\frac{1}{2} \{L_k^\dagger L_k, \rho\} + (1 - \eta_k) L_k \rho L_k^\dagger + \eta_k \langle L_k L_k^\dagger \rangle_\rho \rho \right), \quad (4.2)$$

where $\|\mathcal{N}\| = 1$. The set of jump operators $\{L_k\} \subset \mathcal{A}(\mathcal{H})$ describe the dissipation in S , i.e., the decay of the off-diagonal terms of the positive semidefinite matrix ρ , while $\{\alpha_k\} \in \mathbb{R}$ are the corresponding decay coefficients; in the Markovian regime $\alpha_k > 0$. Here, each term in Equation (4.1) is time-dependent and the coefficients are omitted for notational simplicity. Unlike the linear case, this Lindbladian can capture non-Hermitian dynamics and, from a computational perspective, its formulation does not allow for straightforward vectorisation.

The nonlinear process can be described in several ways. One way is to formalise it within a quantum instrument framework as $\rho(t + dt) = \mathcal{E}(\rho(t))$, where $\mathcal{E} = \sum_k \mathcal{E}_k$ denotes a family of dynamical maps [215]. Here, we consider an extended Hilbert space $S \otimes R$, where R stores the classical record of the detector with orthogonal labels $\{|0\rangle, |k, \text{jp}\rangle, |k, \text{disc}\rangle\}_k$, denoting the three possible outcomes at step k : no jump, a detected jump, or a jump takes place but remains undetected, respectively. Each of the three outcomes is associated with a corresponding operator acting on $S \otimes R$. Specifically

$$\begin{aligned} K_0 &= |0\rangle\langle 0|_R \otimes \left(\mathbb{I} - iH \delta t - \frac{1}{2} \sum_k \gamma_k L_k^\dagger L_k \delta t \right), \\ K_{k, \text{jp}} &= |k, \text{jp}\rangle\langle 0|_R \otimes \sqrt{\gamma_k(1 - \eta_k) \delta t} L_k, \\ K_{k, \text{disc}} &= |k, \text{disc}\rangle\langle 0|_R \otimes \sqrt{\gamma_k \eta_k \delta t} L_k. \end{aligned} \quad (4.3)$$

It is easy to check that the reduced map on the system space $\Lambda(\rho_S) = \text{Tr}_R(\sum_j K_j(\rho_S \otimes |0\rangle\langle 0|) K_j^\dagger)$ is CPTP up to first order of δt . Given S is in state ρ at time t , the unnormalised post-selected state (labelled by $r \in \{0\} \cup \{(k, \text{jp})\}$) at $t + \delta t$ is

$$\rho^{\text{un}}(t + \delta t) = \rho + \delta t \left(-i[H, \rho] + \gamma \sum_k (1 - \eta_k) \mathcal{D}_k[\rho] \right) + \mathcal{O}(\delta t^2),$$

where $\mathcal{D}_k[\rho] = L_k \rho L_k^\dagger - \frac{1}{2}\{L_k^\dagger L_k, \rho\}$ is the dissipator. The corresponding success probability is then $p_{\text{keep}} = 1 - \sum_k \gamma_k \eta_k \delta t \langle L_k^\dagger L_k \rangle + \mathcal{O}(\delta t^2)$. Normalising Equation (4.4) yields the evolution given by Equation (4.2), where the last term is exactly the nonlinear drift that appears solely due to the renormalisation after discarding observed-click runs.

Another way to describe Equation (4.2) makes use of the fact that it admits a natural interpretation in terms of quantum jumps [218–220]. We again use a detector with efficiency η_k to monitor the occurrence of quantum jumps, where we suppose that a fraction of the data (given by η_k) is discarded as post-processing. The role of η_k (and the post-selection) here is to interpolate the evolution between the normal dissipative and non-Hermitian dynamics. Considering the normalisation, a jump happens with probability $\frac{(1-\eta_k)\delta p_k}{1-\eta_k\delta p_k} \approx (1-\eta_k)\delta p_k$, whereas a “no jump” event occurs with probability $\frac{1-\delta p_k}{1-\eta_k\delta p_k} \approx 1 - (1-\eta_k)\delta p_k$. From here, it is easy to see that this can be naturally described by a quantum jump process. A system S in pure state $|\psi(t)\rangle$ evolves for each small timestep δt according to the stochastic dynamics: with probability $(1-\eta_k)\delta p_k(t)$ the state jumps to $L_k |\psi(t)\rangle / \|L_k |\psi(t)\rangle\|$, while with probability $1 - (1-\eta_k)\delta p_k(t)$ it evolves to $e^{-\frac{1}{2}L_k^\dagger L_k \delta t} |\psi(t)\rangle / \|e^{-\frac{1}{2}L_k^\dagger L_k \delta t} |\psi(t)\rangle\|$, where $\delta p_k(t) = \delta t \text{Tr}(L_k^\dagger L_k \rho)$. It was shown in Ref. [221] that this leads to Equation (4.2).

4.2.2 Simulating the open-system dynamics

Here, we present a stochastic Monte Carlo sampling method to simulate the nonlinear, time-dependent stochastic Schrödinger equation in Equation (4.2), see, e.g., Refs. [222–224]. We describe the algorithm below.

The jump time is computed classically based on the total jump probability (related to α_k and η_k) using the sampling method in Ref. [225], where the key task is to compute the probability that there will be no jump in a given time period. The total jump probability in the short duration from t to $t + \delta\tau$ is $P_{\text{jp}}(t) = \gamma \sum_k (1 - \eta_k) \delta p_k(t)$. Given the state undergoes a quantum jump at time t , the probability that there will be no jump until $t + \delta\tau$ is $Q(t) = \lim_{\delta\tau \rightarrow 0} \prod_i^{\tau/\delta\tau} (1 - P_{\text{jp}}(t + i\delta\tau)) = e^{-\int_t^{t+\tau} P_{\text{jp}}(\tau') d\tau'}$. From this, we can compute the sequence of jump

times t_i , which in turn determines the ensemble of trajectories. In case there is no jump in a given time period, the state continues to evolve under the effective operator

$$H_{\text{eff}} = H + \frac{i\gamma}{2} \sum_k \left(L_k^\dagger L_k - \langle L_k^\dagger L_k \rangle \right), \quad (4.4)$$

where H is the Hamiltonian of S , and the subtraction of $\langle L_k^\dagger L_k \rangle$ is due to the normalisation after applying $e^{-\frac{1}{2}L_k^\dagger L_k \delta t}$. This is the only part of the algorithm that must be implemented on a quantum device. For each trajectory, we start from an initial state $|\psi(0)\rangle$ which we evolve under Equation (4.4) by using the QITE algorithm [32] until the jump time t_i . This process is repeated until the final simulation time T is reached, generating a single trajectory. Observable quantities are then estimated by averaging over many such trajectories. The complexity of the simulation is mainly determined by the cost of simulating the non-Hermitian operator H_{eff} on a digital quantum computer.

Here, one can see that the vectorised Lindbladian in Ref. [56] takes a similar form to H_{eff} , which is divided into Hermitian and non-Hermitian parts. Quantum jumps interrupt the dynamics only through interleaving with the jump operator L_k , which has a singular value decomposition $L_k = U_k D_k V_k$ with unitaries U_k and V_k , and a diagonal operator $D_k = \sum_j a_{k,j} |j\rangle \langle j|$. As shown in Ref. [225], the diagonal operator can be realised by a nonunitary exponentiated operator that can be approximated efficiently using the QITE algorithm as well. We remark that the average number of jump events is $\mathcal{O}(\gamma T \sum_k \alpha_k)$, which is linear in T . For more information about this, we again refer to Ref. [225]. The realisation of the jump integrates naturally with the QITE framework. Therefore, in what follows, we primarily consider the cost associated with simulating H_{eff} and consider a continuous evolution up to time T .

To enable a fair comparison between open quantum system and unitary dynamics later in Section 3.6, we introduce the operator $\tilde{H}_{\text{eff}} = H_{\text{eff}} / \|H_{\text{eff}}\|$, i.e., a normalised version of Equation (4.4). Explicitly, we divide Equation (4.4) into the Hermitian and non-Hermitian part, $\tilde{H}_{\text{eff}} = \lambda_1 H_1 - i\lambda_2 H_2$, where $\lambda_1^2 + \lambda_2^2 = 1$, and $\|H_1\| = \|H_2\| = 1$. By normalising the effective operator, the final time can be regarded as stretched and we denote this rescaled final time by T . Overall, the procedure involves

decomposing H_{eff} into H_1 and H_2 , allowing the representation of the time-evolution operator as a product of exponentials $|\psi(t)\rangle = \exp[-i(\lambda_1 H_1 t - i\lambda_2 H_2 t)] |\psi(0)\rangle$. Using Trotterization, the nonunitary dynamics is approximated by an alternating sequence of unitary operations, generated by H_1 , and QITE-based steps that incorporate the effects of H_2 as $|\psi(t)\rangle = [\exp(-i\lambda_1 H_1 \delta t) \exp(-\lambda_2 H_2 \delta t)]^{T/\delta t} |\psi(0)\rangle + \mathcal{O}(T^2/\delta t)$. Here a first-order Trotter decomposition is applied as an example, while K th-order Trotter decomposition is used in deriving the analytical bound in Appendix 4.B. The first part H_1 can be implemented on a quantum computer by applying standard simulation techniques for Trotterization [50, 226]. In contrast, the nonunitary component H_2 cannot be directly realised on quantum hardware and is instead simulated using QITE proposed in Ref. [32]. QITE-based simulations of dissipation typically enlarge the effective support of the evolution operator, suggesting a higher complexity compared to purely unitary dynamics. Still, this increase is not universal. In the case of strong dissipation – where the dissipative term acts effectively like a projector – the complexity is determined mainly by the projector’s support and becomes independent of the total evolution time. Consequently, for a fixed t , increasing γ can, in fact, reduce the required support, and thereby lower the QITE simulation complexity. Moreover, even for weaker dissipation, its effects accumulate over time and can eventually dominate the dynamics, again leading to reduced complexity. This supports our conjecture that there exists a timescale beyond which the complexity of simulating open quantum systems is below that of Hamiltonian systems. Yet, at very short times, the situation is more subtle: the use of a unitary approximation to represent the nonunitary evolution can temporarily increase the complexity.

Before continuing, we comment on another strategy for simulating the Lindbladian master equation by density matrix vectorisation, where the density matrix is mapped to a pure state in an extended Hilbert space [227]. The dynamics is then simulated by applying the Lindbladian superoperator to this *vectorised* state [228]. As with the generalised evolution operator, the Lindbladian can also be decomposed into normalised Hermitian and non-Hermitian parts. Our stochastic approach is

different from the methods for open quantum systems proposed in Refs. [56, 229], where the vectorised form of the system's density matrix ρ is encoded in a larger state vector. A key consequence is that they can only preserve $\text{Tr}(\rho^2)$, rather than $\text{Tr}(\rho)$. As a result, the observable expectation value is calculated by $\langle O \rangle / \text{Tr}(\rho)$, with $\text{Tr}(\rho)$ being measured at each timestep. Note that the complexity analysis of the full evolution of the vectorised density matrix remains the same for the pure-state evolution, as it adopts a similar form with unitary and nonunitary components.

In what follows, we assume the quantum device implements only the evolution generated by H_{eff} ; however, the complexity analysis remains valid for the full evolution of the vectorised density matrix. For further details about this approach, see Appendix 4.A.

4.2.3 Physical measures of quantum algorithmic complexity

We now wish to analyse the quantum algorithmic complexity in terms of physically interpretable quantities. It is well known that in the short-time regime, open quantum systems exhibit growing spatial correlations due to information propagation, which influences the computational resources required to simulate them [230, 231].¹ For a quantum system S in state $|\Psi_s\rangle$, usually the following definition is adopted (see Refs. [233, 234]): we say that $|\Psi_s\rangle$ exhibits ξ -exponential decay of correlations if there exists a correlation length $\xi > 0$ such that for all pairs of disjoint regions (A, B) with separation $d(A, B) \geq l_0$, where $l_0 \in \mathbb{R}^+$ is a fixed distance threshold, and for all operators $\hat{A} \in \mathcal{A}(\mathcal{H}_A)$, $\hat{B} \in \mathcal{A}(\mathcal{H}_B)$, the bound

$$\text{Cor}(\hat{A}, \hat{B}) \leq \|\hat{A}\| \|\hat{B}\| e^{-\frac{d}{\xi}}$$

holds. Here, ξ corresponds to the slowest decaying correlation in S across all region pairs and all times [235], and the correlation function reads

$$\text{Cor}(\hat{A}, \hat{B}) := \langle \Psi_s | \hat{A} \otimes \hat{B} | \Psi_s \rangle - \langle \Psi_s | \hat{A} \otimes \mathbb{I} | \Psi_s \rangle \langle \Psi_s | \mathbb{I} \otimes \hat{B} | \Psi_s \rangle,$$

¹In closed many-body systems, this growth is constrained by the Lieb–Robinson bounds [232].

where for a pair (A, B) at distance d , one can always define a correlation length $C_{A,B} := -\frac{d(A,B)}{\log\left(\frac{|\text{Cor}(A,B)|}{\|A\|\|B\|}\right)} \leq \xi$ that satisfies the definition of correlation length for a given choice of regions. This bound applies for a fixed state $|\Psi_s\rangle$, i.e., fixed $\mathcal{H}$.

In this work, we wish to establish a tighter bound that can track correlation decay during the time evolution of S . Particularly, we consider a pair-specific, distance- and time-dependent correlation length as a replacement of ξ . More formally, we take $C_{A,B}(t)$ to be the infimum over all such $\text{Cor}(\hat{A}, \hat{B})$

$$C_{A,B}(t) := \inf \left\{ C > 0 : \left| \text{Cor}(\hat{A}, \hat{B}) \right| \leq \|\hat{A}\| \|\hat{B}\| e^{-\frac{d}{C}} \right\}, \quad (4.5)$$

where C controls how fast the correlations between A and B decay with d . From all pair-specific correlation lengths, we can thereby recover ξ as the worst-case across all pairs of regions $\xi(t) := \sup_{A,B} C_{A,B}(t)$.

To define the average of Equation (4.5), let $\mathcal{P}_d$ denote the set of all pairs of regions (A, B) at distance d , then

$$\overline{C(d, t)} := \frac{1}{|\mathcal{P}_d|} \sum_{(A,B) \in \mathcal{P}_d} C_{A,B}(t). \quad (4.6)$$

From the concavity of the function $f(x) = e^{-d/x}$ when $x \geq d/2$, and using Jensen's inequality, we find that $\overline{\text{Cor}(d, t)} \leq \overline{e^{-d/C_{A_i, B_i}}} |_{A_i - B_i = d} \leq e^{-d/\bar{C}}$ when $\min_i C_i \geq 1/2$. The average correlation length is introduced to capture the general case when $C_{A,B}(t)$ varies across different pairs of regions. The correlation length $\xi(t)$ captures the worst-case behaviour, whereas the average behaviour is captured by $C_{A,B}(t) \leq \xi(t)$, which gives a tighter bound for that pair as it is defined by the smallest value of C that satisfies the inequality in Equation (4.5). Physically, the decay of correlations quantifies how rapidly local information is lost with distance under the system's dynamics. Fast correlation decay implies an effectively local description of the state, whereas slow decay signals the build-up of long-range correlations, which is known to significantly increase the difficulty of classical and quantum simulation. In the context of our protocol, correlation lengths directly control the scaling of the required resources, making them a natural quantity for characterising the complexity of simulating open-system dynamics.

These quantities will later be used to relate the error bounds and algorithmic complexity. As we shall see later, the algorithmic complexity is mainly determined by the average correlation length $\bar{C}$. In Appendix 4.B, we detail its scaling, including Trotter errors and the dependence of the total simulation error on $\bar{C}$. In brief, the sampling complexity of simulating the overall evolution for a total time T can be expressed as the sum of the real- and imaginary-time contributions. For the latter, the complexity scales polynomially with T and the inverse precision ϵ^{-1} , and exponentially with $\bar{C}$. This formulation sharpens the dependence derived previously in Ref. [32], since the averaged correlation $\bar{C}$ can be much smaller than $\max_i C_i$ for $i \in \{1, N\}$. The circuit depth, on the other hand, is simply determined by the K th-order Trotterization of the evolution and scales as $T^{1+1/K}$ and $\epsilon^{-1/K}$, as expected. In what follows, we explore in more detail how these correlations behave with increasing time and dissipation strength, in order to gain a clearer picture of the hardness of simulating such systems.

On the other hand, in the asymptotic long-time limit, S approaches a unique steady state ρ_∞ . Where the rate of convergence is controlled by the spectral gap Δ of the Liouvillian $\mathcal{L}$, and the mixing time satisfies: $t_{\text{mixing}} \sim \Delta^{-1}$. After t_{mixing} , the quantum simulation no longer requires evolving time-dependent correlations. The required ancilla size saturates to C_{mixing} , and the number of steps becomes independent of the initial state. Thus, the quantum complexity is saturated up to a certain timescale.

4.3 Quantum spin chain system

We now turn to numerical simulations to explore how the complexity of open-system dynamics depends on key physical parameters: dissipation strength γ , evolution time t , and system size N . As a concrete setting, we consider a quantum spin chain system S with N sites, described by a one-dimensional Ising Hamiltonian with transverse and longitudinal fields

$$H_I = J \sum_i X_i X_{i+1} + h_x \sum_i X_i + h_z \sum_i Z_i, \quad (4.7)$$




Figure 4.1: Illustration of the system under study. The boundary-driven quantum spin chain S is coupled with strength γ to two baths at its edges, $i = 0$ and $j = N - 1$, and the set of jump operators $\{L_k\}$ acts locally at the boundary sites. Spatial separation of $d = 1$ means two sites are nearest neighbours with coupling J .

where J is the nearest-neighbour coupling strength, and h_x and h_z denote the magnetic fields in the X and Z directions, respectively. We set $J = 1$ and choose $h_x = h_z = -2$, as in this regime the expectation values of the local operators typically exhibit oscillatory behaviour, making the model suitable for testing our conjecture about the Lindbladian evolution.

Here, S is in contact with an external environment, as described by the time-independent version of Equation (4.1), see Figure 4.1. We couple S to two baths (modelled as local Markovian reservoirs) at its two edges, $i = 0$ and $j = N - 1$, and make a simple but common choice:

$$\begin{aligned}
 L_1 &= \sigma_0^+, & \alpha_1 &= \gamma(1 + \mu), \\
 L_2 &= L_1^\dagger = \sigma_0^-, & \alpha_2 &= \gamma(1 - \mu), \\
 L_3 &= \sigma_{N-1}^+, & \alpha_3 &= \gamma(1 - \mu), \\
 L_4 &= L_3^\dagger = \sigma_{N-1}^-, & \alpha_4 &= \gamma(1 + \mu),
 \end{aligned} \tag{4.8}$$

where γ is the system-bath coupling strength which determines the dissipation rate, μ is the driving parameter, and $\sigma^\pm = \sigma^x \pm i\sigma^y$. In Appendix 4.C we present additional results with $h_x = 0.5$ and $h_z = -1.05$, corresponding to a chaotic regime. The correlation length there exhibits the same trend, becoming increasingly damped with growing dissipation strength and total evolution time, but more noticeably than in the case studied in the main text.

We fix the driving part of the Lindbladian and are primarily interested in how γ impacts the dynamics of S . The jump operators $\{L_k\}$ are assumed to act locally at the boundaries of the spin chain, leaving the coherent Hamiltonian-driven

evolution in the bulk unaffected. Such boundary-driven setups are conceptually appealing as they are analogous to classical nonequilibrium lattice models. While a rigorous microscopic derivation of these local operators from weak system-bath coupling typically yields nonlocal terms, this approach remains physically insightful. It captures key features of non-equilibrium steady states, quantum transport, and dissipation-induced dynamics. This makes the framework well-motivated for exploring questions in quantum thermodynamics and boundary-driven open quantum systems [236, 237]. In the simplest case of single-site boundary driving, one typically introduces two Lindblad operators that flip a spin up or down with different probabilities, thereby inducing a net magnetisation at the boundary site.

To simulate the open quantum system dynamics, the vectorised density matrix is encoded as a matrix product state (MPS) with physical dimension 4 for each of the N sites, corresponding to an extended Hilbert space of dimension $(2^N)^2 = 4^N$ (see Appendix 4.A for details). The Liouvillian superoperator $\hat{\mathcal{L}}$ is represented as a matrix product operator (MPO) with the same enlarged dimension. Time evolution is carried out using the time-evolving block decimation (TEBD) algorithm, where the exponential $e^{\hat{\mathcal{L}}\delta t}$ is approximated via a second-order Trotter–Suzuki decomposition into sequentially applied local exponentials. The timestep δt is dynamically adjusted to ensure consistent numerical accuracy throughout the simulation. The classical dependence is related to the bond dimension χ , while to ensure the simulation accuracy (i.e., the error converges with increasing χ), we also need to keep track of χ .

4.4 Quantum runtime analysis

In this section, we study the influence of dissipation on two physically relevant parameters in both the short- and long-time regimes, which will inform us about the quantum algorithmic complexity of simulating the system of interest S . In particular, we will focus on the *mixing time*, which characterises how quickly S reaches its steady state, and the *correlation length*, which quantifies the spatial extent of quantum correlations within S . Both provide physical insight into the dissipative dynamics and serve as key indicators for assessing the simulability of S .




Figure 4.2: Time evolution of spin observables and trace distance for a dissipative boundary-driven spin chain system S of size $N = 25$ under increasing dissipation rates γ . **a-c** Expectation values of the local spin operators $\langle \sigma_r^x(t) \rangle$, $\langle \sigma_r^y(t) \rangle$, and $\langle \sigma_r^z(t) \rangle$, respectively, under Lindbladian dynamics for different values of γ . **d** Trace distance between consecutive timesteps, $D = \|\rho(t) - \rho(t - \delta t)\|$, indicating convergence toward a steady state. In all panels, $\gamma \in [0, 3]$, with step size of 0.05, where increasing values are shown in progressively lighter shades of blue. Simulations were performed using MPO methods with $\chi = 250$.

We numerically analyse the dependence of simulation complexity on the dissipation strength γ , evolution time t , and system size N .

Before analysing global quantities, it is useful to first examine how dissipation manifests locally in S . Local *magnetisation* observables offer a direct window into the interplay between coherent dynamics and environmental noise: they reveal how quickly local excitations lose their quantum character and relax towards equilibrium. By studying their time evolution for different values of γ , we can build physical intuition for the mechanisms that ultimately determine the simulability of S . Here, for each site r and spin component $a \in \{X, Y, Z\}$, the expectation

value is given by $\langle \sigma_r^a(t) \rangle = \text{Tr}[\sigma_r^a \rho(t)]$, where $\rho(t)$ is the density matrix evolved under Equation (4.1) with Equation (4.7) as the Hamiltonian. Figure 4.2(a)-(c) show the time evolution of the expectation values of these local Pauli operators for various γ . As expected, increasing γ results in stronger damping of the oscillations, reflecting how stronger environmental coupling suppresses coherent dynamics and drives S towards a steady state at long times. In the closed case, oscillations persist over long times, reflecting coherent spin exchange in the Heisenberg model. As γ increases, these oscillations decay more rapidly and their amplitude is strongly suppressed, signalling the progressive loss of phase coherence and the dominance of relaxation processes. Physically, this means that information about the initial state is erased faster, S approaches a featureless steady state sooner, and the dynamical complexity is reduced. In the broader context of simulability, such damping hints at shorter correlation lengths and mixing times, which can make long-time dynamics more accessible to classical methods and reduce the potential advantage of quantum algorithms in this regime.

The damping of local magnetisation oscillations observed in Figure 4.2(a)-(c) already suggests that dissipation accelerates the relaxation of S towards equilibrium. To quantify this effect in a more global and model-independent way, we turn to the *mixing time* – the timescale over which S effectively reaches a steady state. This provides a direct indicator of how long coherent or correlated dynamics persist, and therefore how challenging the long-time regime may be to simulate. We quantify this by computing the trace distance $D(t) = \|\rho(t) - \rho(t - \delta t)\|$ between density matrices at successive timesteps for various values of γ . Defining a small threshold ε , S is considered to have effectively reached its steady state when $D < \varepsilon$. The results in Figure 4.2(d) show that stronger dissipation leads to faster convergence, i.e., shorter mixing time, consistent with the relaxation of local observables in Figure 4.2(a)-(c).

We then proceed to the relation between the *correlation length* and simulation complexity. Spin correlations serve as key indicators of the required computational resources as the presence of long-range correlations across distant sites generally means




Figure 4.3: Evolution of spin- Z correlations in a dissipative boundary-driven spin chain S with $N = 25$ sites as a function of γ , t , and d . **a** Two-point correlation $C_d^z(t = T)$ at the final time $T = 10$, plotted for spatial separations $d = 1, \dots, N - 1$ (curves range from deep blue at $d = 1$ to deep red at $d = 24$); labels indicate the shortest and longest distances. Inset showing the distance-averaged peak correlation $\frac{1}{N-1} \sum_{d=1}^{N-1} \max_t |C_d^z(t)|$ versus γ . **b** Maximum correlation across distances as a function of t . The data ranges for $\gamma \in [0, 0.3]$ with step size of 0.05.

increased complexity [207]. Crucially, the correlation length bounds the number of measurements. The spin correlation function encodes these spatial dependencies

$$c_{ij} = \text{Tr}(\sigma_i^z \sigma_j^z \rho) - \text{Tr}(\sigma_i^z \rho) \text{Tr}(\sigma_j^z \rho),$$

where i and j denote different sites. Figure 4.3(a) shows the correlation length $C_d^z(t)$ at distances $d \in 1, \dots, N - 1$ for a fixed long evolution time. The results demonstrate a systematic decrease in $C_d^z(t)$ as γ increases. This behaviour indicates suppression of long-range correlations, and consequently, a reduction in the complexity of simulating strongly dissipative open quantum systems. Besides, to provide a broader perspective, Figure 4.3 shows the *average* spatial correlations, which more clearly highlight the reduction of correlation length – and hence simulation

complexity – with increasing dissipation at fixed time. Here we present results at relatively long times, where this behaviour is most pronounced, although a similar (albeit less drastic) trend is already visible at shorter times. Since, according to Section 4.2.3, the running time for simulating open quantum system dynamics on a quantum computer is bounded by the maximum correlation length across all distances at each timestep, we also compute this quantity throughout the evolution of S for each γ , see Figure 4.3(b). We note that correlations are not necessarily suppressed by long-time evolution alone: due to the oscillatory behaviour of the Hamiltonian, they can persist or even revive at later times. However, dissipation introduces a damping effect that becomes increasingly dominant with time, leading to a systematic reduction of correlations. This trend is consistent with the average correlation results in Figure 4.3(a), where dissipation clearly shortens the correlation length. Together with the analytical scaling in Section 4.2.3, which links time complexity to the correlation length, these results indicate that dissipation effectively reduces the simulation complexity – even if correlations themselves do not always decay monotonically with total time evolution.

4.4.1 Time-dependent Lindbladian

In this section, we also investigate time-dependent Lindbladian dynamics, where either the dissipation strength, the Pauli operators, or the sites on which the jump operators act vary during evolution. Overall, we observe qualitatively similar behaviour to the results discussed here, with no significant deviations. This further supports the generality of our conclusions and provides motivation for applications such as dissipative quantum state preparation, where time-dependent Lindbladians are often required to engineer the desired steady state [216].

In the first protocol, the system evolves under a time-dependent Lindblad equation, where the dissipation strength $\gamma(t)$ varies smoothly during the evolution according to

$$\gamma(t) = \gamma_{\max} \left| \sin \left(\frac{\pi t}{t_{\text{total}}} \right) \right|. \quad (4.9)$$




Figure 4.4: Dynamics under time-dependent and random-site Lindbladian protocols for system S of size $N = 11$ and maximum bond dimension $\chi = 150$. Top row: Time evolution of the local magnetization $\langle \sigma_{N/2}^z(t) \rangle$ at the center of the chain for (a) $\gamma_{\max} = 0.1$ and (b) $\gamma_{\max} = 0.5$. Bottom row: Connected correlation $C_{0,N/2}^z(t)$ between the ends and the center, for (c) $\gamma_{\max} = 0.1$ and (d) $\gamma_{\max} = 0.5$. Results are shown for three cases: time-independent dissipation ($\gamma = \gamma_{\max}$), time-dependent dissipation $\gamma(t)$, and the random Lindblad protocol (random operators and sites).

Here, $\gamma_{\max}$ sets the maximum dissipation rate reached during the dynamics. The Lindblad operators are fixed as $\sigma^\pm$ acting at the two ends of the chain, as in the previous section.

In the second protocol, we fix γ to a constant value but randomise both the operator type (choosing randomly from X , Y , or Z Pauli operators) and the sites at which the Lindblad operators act. For each simulation run, two sites are selected at random, and each is assigned a randomly chosen Pauli operator as the Lindblad operator. Figure 4.4 displays the time evolution of the local magnetisation $\langle \sigma_i^z(t) \rangle$ at the center of the chain (top row) and the connected correlation function $C_{0,N/2}^z(t)$ (bottom row), for both protocols and for two values $\gamma_{\max} = 0.1$ and $\gamma_{\max} = 0.5$. For the correlations, we focus on the distance d between site $i = 0$ and the centre of the chain, but note that the observed behaviour is representative of other values of d as well.

These results highlight how both the temporal profile and spatial structure of the dissipation channel can influence relaxation and steady-state properties of local

observables in open quantum systems. Notably, we find that the time-dependence of γ in the first protocol does not lead to significant differences in either local magnetisation or correlations when compared to the corresponding time-independent case with $\gamma = \gamma_{\max}$. This is expected, as the overall effect of γ can be understood as accumulating over time, but never exceeding $\gamma_{\max}$, so that both the relaxation and steady-state behaviour are largely determined by $\gamma_{\max}$. The correlations show a similar insensitivity to the time dependence of the dissipation, being primarily controlled by γ itself. This behaviour persists in different choices of d .

In contrast, the second protocol, where Lindblad operators act at random sites with randomly chosen Pauli types, produces more noticeable and consistent differences in the observable dynamics. Here, the randomness in the spatial and operator structure of the dissipation can substantially alter the system's relaxation pathways, leading to a richer variety of dynamical behaviours in the observables, depending on the particular realisation of the dissipative channel. Overall, we observe qualitatively similar behaviour to the results discussed here, with no significant deviations. At the same time, we find that randomised application of jump operators can accelerate relaxation processes, for instance by enhancing the decay of correlations, thereby modifying the system's dynamical pathways. This further supports the generality of our conclusions and highlights the potential of such protocols for applications such as dissipative quantum state preparation, where time-dependent Lindbladians are often required to engineer the desired steady state [216].

Finally, we wish to remark that while the discussion in this section has focused on linear Lindbladian dynamics, nonlinear master equations could be even more challenging to simulate classically, as their explicit matrix representation is already difficult to construct. Importantly, the quantum algorithm we introduce in this chapter can efficiently simulate both linear and nonlinear evolutions, making our conclusions directly relevant to the broader setting.

4.5 Classical runtime analysis

We now analyse the classical resource requirements – bond dimension and operator entanglement entropy – and provide simulation results to illustrate how simulation errors depend on time t , system size N , and bond dimension χ . Generally, we expect simulation accuracy to improve for larger χ at the cost of increased computational runtime. The objective here is twofold: (1) to determine whether the growth trend differs between quantum and classical algorithms, and (2) to assess whether classical algorithms exhibit rapid growth in complexity with increasing N and t .

Figure 4.5 illustrates the classical simulability of open quantum systems through that accuracy-complexity trade-off. For smaller systems, we quantify simulation errors by comparing results at finite χ against the ground-truth solution computed using a sufficiently large bond dimension, chosen here as $\chi = 250$ to ensure convergence of the expectation values of the observable (see Appendix 4.C). The accuracy of TEBD simulations at given χ with respect to the ground truth is measured via spin-spin correlation functions, with deviations quantified as the correlation error

$$\varepsilon = \left(\frac{\sum_{ij} (c_{ij} - \tilde{c}_{ij})^2}{\sum_{ij} \tilde{c}_{ij}^2} \right)^{1/2}, \quad (4.10)$$

where $\tilde{c}_{ij}$ are the ground-truth correlations. In Figure 4.5(a), we show how the correlation error ε evolves over t for $N = 25$ and $\chi \in [50, 225]$. As expected, larger χ consistently yield better accuracy, especially at longer times. Since S eventually reaches a steady state, ε saturates to a constant value, which is why we think the error remains bounded over long simulation times. In a similar manner, in Figure 4.5(b) we fix $t = 3$ and examine ε with respect to N for $N \in [10, 25]$. The results confirm the linear growth observed earlier. The main contribution to the complexity of the algorithm, however, arises from its scaling with χ , as it determines the fidelity between our MPS state and the evolved quantum state. In Figure 4.5(c), ε is plotted on a logarithmic scale as a function of χ for $N \in [10, 25]$ and fixed $t = 3$. Evidently, for a given χ , smaller systems exhibit smaller ε , making them easier to simulate. For practical purposes, in Figure 4.5(d) we set an error threshold




Figure 4.5: TNs simulations for estimating the complexity of classical simulations. In all panels $\gamma = 0.05$. We compute the correlation error, indicative of the classical simulability of S : (a) as a function of t for $N = 25$ and $\chi \in [50, 225]$, (b) as a function of χ for $N \in [10, 25]$ and $t = 3$; inset: growth of χ over t , and (c) as a function of $N \in [10, 25]$ for $\chi \in [50, 225]$ and $t = 3$. (d) Minimum $\chi_{\min}$ to reach $\epsilon_{\chi_{\min}} \leq 0.05$ as a function of $N \in [10, 25]$ for $t \in \{3, 4, 5\}$.

$\epsilon_{\chi_{\min}} \leq 0.05$ and identify the minimal bond dimension $\chi_{\min}$ necessary to reach that threshold for every N . Naturally, $\chi_{\min}$ depends on t , as errors consistently increase with longer simulation times, illustrated clearly in Figure 4.5(d). Overall, as expected, we see that larger χ yield more precise results, progressively approaching exact solutions. This confirms that dynamics can be approximated with arbitrary accuracy given sufficiently large χ . In our case, we find that $\chi = 200$ is sufficient to reach the error threshold for the system sizes considered, so working with samples




Figure 4.6: Operator entanglement entropy plot for $N = 25$, $\chi = 250$, and $\gamma \in [0, 0.3]$ with step size of 0.05. Each data point is obtained by calculating the operator space entanglement entropy of the bipartition between the two halves of the MPS.

up to $\chi = 250$ is well justified. This is consistent with the results in Appendix 4.C, where we observe clear convergence of the observables with increasing bond dimension, further supporting that $\chi = 250$ provides a reliable estimate.

The classical simulability of TNs is known to be fundamentally constrained by the entanglement entropy of the quantum state. To assess this limitation, we analyse the growth of entanglement entropy as an indicator of computational complexity. In Figure 4.6, we study the operator space entanglement entropy of the time-evolved state. We find that although the entanglement entropy initially increases, it then exhibits qualitatively different behaviour depending on γ . For γ above a certain threshold, the entanglement entropy decreases over time, consistent with our expectations and the results presented in Ref. [238]. However, when considering simulation complexity, a more relevant metric is the bond dimension χ required to represent the state. Surprisingly, we find that χ remains largely constant across different values of γ . This result is somewhat counter-intuitive compared to the behaviour of the entanglement entropy. The key difference lies in the fact that the entanglement entropy characterises the final state, whereas simulating the full dynamics requires tracking the entire evolution, which can be more demanding. As a result, χ does not decrease with t even for different values of γ .

To give a brief definition of the operator entanglement, we define our density matrix in MPO form as a Schmidt decomposition between the left L and right R block

$$\hat{\rho} = \sum_i \lambda_i \hat{o}_i^{[L]} \hat{o}_i^{[R]}, \quad (4.11)$$

where $\text{Tr}(\hat{o}_i^{[L/R]} \hat{o}_j^{[L/R]}) = \delta_{ij}$ and λ_i denotes the Schmidt coefficients. Then, define the operator entanglement as $S_{\text{OP}} = -\sum_i \lambda_i^2 \log(\lambda_i^2)$. S_{OP} is not strictly connected with quantum entanglement between blocks but it is an important quantity nonetheless that gives insights into the efficiency of the MPO representation. Recall that, for the purpose of our simulation, we have encoded the density matrix MPO into a MPS with double the quantum dimension. In this case, the operator entropy of ρ is equivalent to calculating the entanglement entropy of the evolved MPS, carefully handling the extra physical dimension.

A noteworthy observations is that the complexity of TNs does not decrease with t and γ in certain regimes, in stark contrast to the complexity associated with quantum algorithms. This suggests there may exist a range of dissipation strength, where the separation between quantum and classical computational complexity is maximised.

4.6 Discussion

In this chapter, we assess the algorithmic complexity of simulating open quantum systems under the general conjecture that there exists a threshold timescale beyond which open quantum systems within the Lindbladian formalism are computationally easier to simulate than Hamiltonian systems. Instead of relying on dilation methods [54, 199, 200] that can obscure the intrinsic complexity of the evolution of the reduced system, we do so directly by focusing on two physically relevant parameters: correlation length and mixing time. To make our analysis more comprehensive, we employ two complementary approaches: classical TNs and quantum jump representation via QITE. Methodologically, we present a quantum algorithm that treats both linear and nonlinear Lindbladians in the same way, and allows for the non-Hermitian dynamics to be simulated deterministically without the need for extra ancilla; thereby we can perform stochastic simulations with no additional

Trotterization cost by only sampling jump times. This stochastic simulation method goes beyond existing methods that either need to introduce qubits each timestep or can simulate nonlinear evolution only probabilistically. Moreover, we derived a tighter bound for the correlation decay of a time-evolving system in terms of its average value.

To test the conjecture, we numerically studied a boundary-driven dissipative quantum spin chain system, analysing both local observables and global quantities. We observe that increasing the dissipation strength reduces the correlation length of local operators at long times and shortens the mixing time (the time it takes a system to reach its steady state), indicating faster relaxation and suppressed long-range correlations. These trends support the view that, under appropriate conditions, long-time open quantum system dynamics is less complex to simulate than real-time unitary dynamics. We further examined the relation between classical and quantum resources. We show that in our classical simulations the runtime is dominated by the bond dimension and displays super-polynomial growth with time (at some fixed accuracy), whereas the quantum runtime scales linearly with time. Interestingly, and in contrast with the quantum case, we observe that increasing the dissipation strength did not affect the bond dimension needed for accurate TN simulations, even though operator-space correlations tend to decrease, highlighting a distinction between asymptotic correlation decay and the transient entanglement that governs truncation error.

This chapter provides insight into how dissipation shapes both classical and quantum algorithmic complexity via physically meaningful measures. These results can pave the way for designing more efficient algorithms for simulating open quantum systems, and identifying regimes where quantum algorithms can outperform classical ones. Some interesting directions for future investigation emerge. In [Section 4.4.1](#) we highlighted the role of randomness in the steady-state properties of local observables. A key question here is how to exploit that to algorithmically design randomised dissipation protocols that systematically reduce mixing times. Achieving this can

reduce circuit depth and gate count, enabling scalable preparation of complex many-body states on near-term hardware with applications across quantum chemistry, materials science, and condensed matter physics. On a related note, it is also worthwhile to determine whether the observed linear scaling of quantum algorithms persists in realistic settings like those with nonlocal jump operators or strongly correlated finite-temperature environments. We envision that physically relevant measures, such as correlation length and mixing time, will see broader adoption as reliable guides for estimating algorithmic complexity and informing the design of efficient algorithms for open-system dynamics.

Appendices

4.A Vectorisation of the density matrix

To simulate the quantum dynamics governed by a linear Lindbladian, we encode the density matrix as a pure state through vectorisation. This transformation allows us to represent mixed-state evolution in a form suitable for quantum circuits. The same technique is also applied for classical TN simulations of open quantum systems, as described in Section 3.6. The mapping sends ρ to a vector in an extended Hilbert space

$$|\rho\rangle\rangle = \sum_{i,j} \rho_{ij} |i\rangle \otimes |j\rangle.$$

Under this transformation, the Lindblad master equation takes a Schrödinger-like form

$$\frac{d}{dt} |\rho\rangle\rangle = \hat{\mathcal{L}} |\rho\rangle\rangle,$$

where the Liouvillian superoperator $\hat{\mathcal{L}}$ acts linearly as

$$\begin{aligned} \hat{\mathcal{L}} = & -i(\mathbb{I} \otimes H - H^* \otimes \mathbb{I}) \\ & + \sum_k \left(L_k \otimes L_k - \frac{1}{2} \mathbb{I} \otimes L_k^\dagger L_k - \frac{1}{2} L_k^\dagger L_k \otimes \mathbb{I} \right). \end{aligned} \quad (4.12)$$

To compare open-system ($\gamma > 0$) with unitary ($\gamma = 0$) dynamics on equal footing, we consider a normalised Liouvillian

$$\hat{\mathcal{L}} = -i(\lambda_1 H_1 - i\lambda_2 H_2), \quad (4.13)$$

where $\|H_1\| = \|H_2\| = 1$ and $\lambda_1^2 + \lambda_2^2 = 1$, with H_1 and H_2 is the matrix representation of the superoperator. In analogy with the operator H_{eff} in Equation (4.4), the time evolution of the vectorised density matrix can be written as

$$|\rho(t)\rangle\rangle = \exp[-i(\lambda_1 H_1 t - i\lambda_2 H_2 t)] |\rho(0)\rangle\rangle. \quad (4.14)$$

The derivation then proceeds by applying a first-order Trotter decomposition, as an example

$$|\rho(t)\rangle\rangle = \left[e^{-i\lambda_1 H_1 \delta t} e^{-\lambda_2 H_2 \delta t} \right]^{T/\delta t} |\rho(0)\rangle\rangle + \mathcal{O}\left(\frac{T^2}{\delta t}\right),$$

which adopts the same form as in the quantum jump situation. In Appendix 4.B, K th-order Trotter decomposition is applied. The dynamics of the vectorised state is then approximated as alternating applications of unitary operations and operations derived from the QITE approach.

4.B Error bound of the evolution and the algorithmic complexity

Below we analyse the simulation complexity of open quantum system dynamics. We first divide the total time evolution T into N_2 segments of duration $\delta t = T/N_2$. For the non-Hermitian component H_2 , non-commuting terms introduce a Trotter error in QITE. To control Trotter error, we further divide each δt into N_1 sub-steps. Here, we consider a K th-order Trotterization at all levels (with even K). The number of sub-steps N_1 is set by the condition

$$\frac{N_2(\lambda_2 \delta t)^{K+1}}{N_1^K} = \frac{\varepsilon}{3} = \mathcal{O}(\varepsilon), \quad (4.15)$$

where we decompose the total error ϵ into three contributions: the splitting of the evolution into Hermitian and non-Hermitian parts, the Trotter error in the real-time evolution, and the unitary decomposition of the imaginary-time evolution operator. Each contribution is controlled to be at most ϵ , which gives

$$N_1 = \mathcal{O}\left(\frac{(\lambda_2 T)^{\frac{K+1}{K}}}{N_2 \varepsilon^{\frac{1}{K}}}\right). \quad (4.16)$$

The simulation error is governed by the average correlation length $\overline{C}$. For the effective non-Hermitian Hamiltonian H_{eff} , we write

$$H_2 = \sum_{i=1}^{M_k} h_{2,i}. \quad (4.17)$$

Following the argument in Ref. [32], let Ψ_p denote the pure state at the p th step

$$|\Psi_p\rangle = \frac{\prod_{k=1}^{M_k} e^{-\delta t h_{2,k}} |\Psi_{p-1}\rangle}{\|\prod_{k=1}^{M_k} e^{-\delta t h_{2,k}} |\Psi_{p-1}\rangle\|}, \quad (4.18)$$

with density matrix $\rho_p = |\Psi_p\rangle\langle\Psi_p|$. Consider a region R_v consisting of all sites within distance v of the support of h_k . The distance between successive reduced states satisfies

$$\begin{aligned} \|\text{Tr}_{\setminus R_v} \rho_p - \text{Tr}_{\setminus R_v} \rho_{p-1}\|_1 &\leq \sum_{i=1}^{M_k} \left\| \prod_{k=i+1}^{M_k} e^{-\delta t h_{2,k}} \rho_{p-1} - \prod_{k=i}^{M_k} e^{-\delta t h_{2,k}} \rho_{p-1} \right\|_1 \\ &\leq \sum_{i=1}^{M_k} e^{-\frac{v}{\bar{c}_i}} \leq M_k e^{-\frac{v}{\bar{c}}}. \end{aligned} \quad (4.19)$$

Using the domain size derived in Ref. [32], but in this case using Jensen's inequality as introduced in the main text, we set $v = \mathcal{O}\left(\bar{C} \ln(M_k N_2 \varepsilon^{-1})\right)$. We can bound the error from approximating the Lindbladian by a sequence of unitaries. The relevant parameter here is the average correlation length $\bar{C}$, rather than the maximal one. The sample complexity per timestep for the non-Hermitian evolution is

$$S_q^{\text{NH}}(\delta t) = N_1 M_K 3^{\max v^{\text{dim}}} = \mathcal{O}(N_1 N_2 M_K^2 e^{\bar{C}^{\text{dim}}} \varepsilon^{-1}), \quad (4.20)$$

where we make use of state tomography. Any quantum state can be decomposed in the Pauli basis, which consists of the three non-trivial measurement operators (X , Y , and Z). For a system of N_q qubits, this yields 3^{N_q} measurement settings that are informationally complete. Consequently, the sampling complexity scales exponentially with the domain size as $3^{\max v^{\text{dim}}}$.

At a higher level, N_2 is fixed by the Trotter error between H_1 and H_2

$$\epsilon_{\text{Trotter}} = \mathcal{O}\left((\lambda_1 \lambda_2)^{K+1} \frac{T^{K+1}}{N_2^K}\right) = \frac{\varepsilon}{3}, \quad (4.21)$$

so that

$$N_1 = \mathcal{O}(\lambda_1^{-(1+\frac{1}{K})}), \quad N_2 = \mathcal{O}\left(\frac{(\lambda_1 \lambda_2 T)^{1+\frac{1}{K}}}{\varepsilon^{\frac{1}{K}}}\right).$$

Thus the sampling cost due to the non-Hermitian part becomes

$$S_q^{\text{NH}}(\delta t) = \mathcal{O}\left(M_K^2 (\lambda_2 T)^{1+\frac{1}{K}} e^{\bar{C}^{\text{dim}}} \varepsilon^{-(1+\frac{1}{K})}\right). \quad (4.22)$$

In many physically relevant cases, environmental interactions are described by local Lindbladians with commuting terms in H_2 . Then, no N_1 subdivision is required and the complexity reduces to

$$S_q^{\text{NH}}(T) = \mathcal{O}\left(M_K^2(\lambda_1\lambda_2 T)^{2+\frac{2}{K}} e^{\bar{C}^{\text{dim}}} \varepsilon^{-(1+\frac{2}{K})}\right). \quad (4.23)$$

For the real-time part, the number of subsegments per δt is $(\lambda_1 T)^{1+\frac{1}{K}} N_2^{-1} \varepsilon^{-\frac{1}{K}}$, leading to a gate complexity of $(\lambda_1 T)^{1+\frac{1}{K}} \varepsilon^{-\frac{1}{K}}$. Including both real- and imaginary-time contributions, the total sampling and gate complexity is

$$S_q^{\text{tot}} = \mathcal{O}\left((\lambda_1 T \varepsilon^{-1})^{1+o(1)} + (\lambda_2 T \varepsilon^{-1})^{2+o(1)} M_K^2 e^{\bar{C}^{\text{dim}}}\right). \quad (4.24)$$

We note that the above expression characterises the total sample complexity accumulated over all steps. It also reflects the computational complexity associated with the stepwise growth of the circuit (with each block of duration δt).

Finally, note that $\bar{C} \leq \xi$, where ξ is the universal correlation length of the system. Physically, $\bar{C}$ measures correlations at a fixed distance d with all possible operators, whereas ξ captures correlations at all distances. Our formulation sharpens the dependence on correlation length, since $\bar{C}$ can be much smaller than $\max_i C_i$.

4.C Tensor network simulation results for chaotic systems

Although the main text focusses on the Heisenberg model with oscillatory behaviour in its Pauli observables, in Figure 4.C.1 we show the case for $h_x = 0.5$, $h_z = -1.05$, corresponding to a chaotic regime. Here, the coherent oscillations present at weak dissipation are rapidly suppressed as γ increases, leading to faster relaxation towards steady state. This behaviour illustrates how dissipation tends to erase the signatures of chaos in local observables, effectively driving the system into featureless dynamics at long times. Evidently, the correlation length exhibits the same trend, becoming increasingly damped with γ – more noticeably than in the case studied in the main text – and with total evolution time, as shown in Figure 4.C.2. This confirms that dissipation not only suppresses temporal oscillations but also washes out spatial




Figure 4.C.1: Time evolution under Lindbladian dynamics of spin observables for S of size $N = 25$ under increasing dissipation rates $\gamma \in [0, 0.3]$ with step size of 0.05. a–c Expectation values of the local spin operators $\langle \sigma_r^x(t) \rangle$, $\langle \sigma_r^y(t) \rangle$, and $\langle \sigma_r^z(t) \rangle$, respectively.

correlations across the chain, reinforcing the picture that strong coupling to the environment reduces the effective complexity of the dynamics.

On a different note, in this work, we set the bond dimension to $\chi = 250$ as the ground truth. In Figure 4.C.3 we show that this choice is well justified, with the observable dynamics clearly converging as χ increases. The convergence is particularly evident at intermediate times, where entanglement growth is more pronounced, yet the difference between $\chi = 200$ and $\chi = 250$ remains negligible. This validates our numerical accuracy and ensures that all qualitative conclusions drawn in the main text are robust with respect to truncation errors.




Figure 4.C.2: Spin- Z correlations in a dissipative spin chain of with $N = 25$ sites as a function of γ . **a** Two-point correlation $C_d^z(t = T)$ at the final time $T = 10$, plotted for separations $d = 1, \dots, N - 1$ (curves range from deep red at $d = 1$ to deep blue at $d = N - 1$); labels indicate the shortest and longest distances. **b** Maximum correlation across distances as a function of t . Similarly, $\gamma \in [0, 0.3]$ with step size of 0.05.




Figure 4.C.3: Time evolution of spin observables for a system of size $N = 25$ for increasing bond dimension χ . **a-c** Expectation values of local spin operators $\langle \sigma_r^x(t) \rangle$, $\langle \sigma_r^y(t) \rangle$, and $\langle \sigma_r^z(t) \rangle$, respectively, under Lindbladian dynamics for $\chi \in [50, 250]$, with increasing values shown in progressively darker shades of blue.

5

Conclusion and outlook

In this thesis, we have explored how quantum dynamics simulations can be harnessed to probe the physical properties of complex many-body systems, with a particular focus on regimes that are challenging for classical computation. This is a key direction in quantum computing, as it represents one of the few areas where quantum devices may offer a genuine advantage in solving practical and scientifically relevant problems. However, current quantum computers are still small and noisy, making such tasks highly non-trivial, and the search for practical applications in the NISQ era remains elusive. The work presented in this thesis helps pave the way forward by introducing efficient algorithms tailored for quantum hardware, studying their implementation on real quantum devices, and, crucially, analysing their complexity to validate their successful performance.

5.1 Summary of work

Since the overarching goal of this thesis is to simulate quantum many-body systems in order to characterise their behaviour, we have examined both the dynamics of isolated systems through unitary evolution and of open systems described by Lindbladian dynamics. This dual focus is a distinctive feature of the work, as relatively few studies in the literature address the characterisation of open-system

dynamics in the context of quantum computing. Our results in Chapter 4 contribute to laying the foundations for simulating open quantum dynamics on quantum devices, indicating that in some cases these simulations can be less challenging than their closed-system counterparts, while also offering insights that can guide future algorithmic and experimental developments in this area. By contrast, Chapter 2 and 3 focus on unitary dynamics, exploiting dynamical properties to extract static quantities, most notably the excitation spectrum of many-body systems, through efficient spectroscopy-based protocols.

The first contribution of this thesis, presented in Chapter 2, is the development of a protocol for probing transition energies and excitation spectra of quantum many-body systems, establishing a connection between spectroscopic techniques and quantum simulation that enables advances in both fields. The method detects frequency resonances from short-time dynamics, requiring evolution times only logarithmic in the target precision, making it well suited to current quantum hardware. The main advances and contributions of this chapter are:

- **Applicability.** The protocol extends beyond existing spectroscopy or analogue quantum simulation schemes, as it can probe spectral features of systems that do not conserve particle number.
- **Complexity advantage.** We show that the circuit complexity is near optimal for spectral property estimation, and that in the presence of noise the total runtime (circuit depth $\times$ samples) scales polynomially with the inverse precision, $\mathcal{O}(\text{Poly}(\epsilon^{-1}))$, in contrast to the exponential sample complexity $\mathcal{O}(\exp(\epsilon^{-1}))$ of previous methods. We support this theoretical claim with classical simulations of noisy evolutions, demonstrating that the spectral information can be recovered using only simple error-mitigation techniques.
- **Versatility and implementation.** We demonstrate the protocol on a range of systems, including large-scale numerical simulations to study scaling with system size, and real-device experiments on IBM quantum hardware to probe the excitation spectra of spin Hamiltonians.

In Chapter 3, we generalised the quench spectroscopy protocol to study spin excitation properties of fermionic systems, showing that with a carefully chosen quench operator it is possible to access the exact retarded spin correlation function and dynamical spin structure factor when the initial state preserves spin reflection symmetry. We apply this approach to the one-dimensional Fermi–Hubbard model on an IBM quantum computer, leveraging the protocol’s robustness to initial state perturbations to minimise state preparation costs and allocate most resources to Hamiltonian simulation via a first-order Trotter formula. While higher-order Trotter methods can in principle reduce discretisation error, they incur substantially higher gate counts and circuit depths on NISQ devices [171]. Since our main aim is an efficient implementation of the quench spectroscopy protocol, and because we find that first-order Trotterization often performs as well as or better than second-order formulas in practice, we restrict to the first order. With this choice and by exploring different fermion-to-qubit encodings to reduce circuit depth, we simulate systems up to $L = 15$ sites and recover the full two-spinon continuum spectrum in close agreement with theoretical predictions. This implementation would not have been possible without the following contributions and optimisations developed in this chapter:

- Approximate free-fermion initial state obtained via a variational quantum eigensolver, where the angles of a fixed number of Givens rotation layers are optimised with respect to fidelity, yielding an efficient approximation to the free-fermion ground state.
- Choice of a unitary quench operator, removing the need for block encoding or ancillary qubits and reducing circuit complexity.
- Detailed classical simulations to analyse Trotter errors and total time evolution, enabling extrapolation to larger systems and guiding optimal parameter choices prior to hardware runs.
- Selection of the optimal Jordan–Wigner ordering when mapping fermions to qubits, chosen from among many possibilities to minimise circuit depth.

- Minimal but effective noise suppression techniques, including dynamical decoupling and Pauli twirling. The latter is particularly successful here, as it approximates coherent noise as Pauli noise, which transforms into a constant factor in the Fourier domain and vanishes upon normalisation.

Finally, in Chapter 4, we shift our focus from isolated systems to open quantum systems, investigating how the complexity of simulating open quantum dynamics compares both to their closed-system counterparts and to classical simulation methods. In particular, we explore the algorithmic complexity of simulating open quantum many-body dynamics without resorting to dilation techniques, which often obscure the intrinsic structure of the problem by embedding it into a larger unitary evolution. Instead, we work directly in the Lindbladian framework, where the relevant physical quantities, such as correlation length and mixing time, emerge naturally as the key parameters controlling complexity. We propose a quantum algorithm that enables the stochastic simulation of open-system dynamics and provide numerical results that reinforce the theoretical predictions, specifically demonstrating the decay of correlations and the acceleration of mixing time with increasing dissipation strength. The main strengths of the work in this chapter are:

- Unified framework. We present a general quantum algorithm applicable to both linear and non-linear Lindbladians, including time-dependent cases. The approach requires only the quantum simulation of an effective non-Hermitian Hamiltonian, achieved by applying QITE followed by standard unitary dynamics.
- Computational complexity of open dynamics. We show that the damping of correlations due to dissipation reduces the quantum simulation complexity at long times, in contrast to the closed-system case.
- Tensor-network simulations of spin-bath models. We perform classical simulations of spin systems coupled to baths with both fixed and time-dependent couplings. We find that time dependence has little impact on complexity,

except when the bath coupling site is randomly varied at each step, in which case relaxation is accelerated through faster correlation decay. These results support the conclusion that randomised dissipation can provide a more efficient route to state preparation.

- Bridging theory and numerics. Large-scale tensor-network simulations serve two purposes: (i) they support and extend the theoretical predictions, providing estimates of experimentally relevant timescales; and (ii) they expose the limitations of classical methods, where bond-dimension growth and operator entanglement entropy set clear boundaries. These limitations delineate the regimes where quantum simulations can offer a genuine advantage.

Together, the methods and results presented in this thesis move us a step closer to realising quantum simulation algorithms capable of providing a new and meaningful understanding of complex quantum many-body phenomena.

5.2 Future research directions

In addition to the results presented in this thesis, several natural extensions and open challenges arise. In what follows, we outline two main directions that represent natural continuations of this work and that we expect to be of particular relevance for the quantum simulation of many-body systems. These questions have been present throughout the development of this thesis, and we hope that the results obtained here will provide a foundation for pursuing them further.

5.2.1 Two-dimensional simulations

From the implementation of the quench algorithm on real quantum hardware in Chapter 3, we were impressed by the quality of the results, even though several approximations were required to reduce circuit depth. We managed to push the simulation to 30 qubits (equivalent to 15 fermions), which represents one of the largest demonstrations of fermionic time dynamics on a quantum device to date [140, 141]. Of course, this was carried out in one dimension, a case that

can still be efficiently simulated classically. A natural and pressing question is whether such protocols can be extended to two dimensions while preserving the error resilience observed in 1D.

There is currently great interest in exploring quantum advantage through real-time dynamics of higher-dimensional quantum systems [137, 239, 240], where tensor-network methods often become inefficient. However, careful attention must be paid to implementation details. Hardware connectivity is a key factor: for example, on IBM devices moving from 1D to 2D would require a large number of SWAP operations, significantly increasing circuit depth and gate count. Other platforms, such as Quantinuum’s H2 trapped-ion quantum processor, which provide all-to-all connectivity, could therefore be promising candidates for carrying out such experiments.

An interesting future direction is to exhaustively explore tensor-network methods for simulating the dynamics of 2D systems, in particular fermionic models. This would allow us to identify the regions where classical methods fail to achieve a given accuracy, or where the maximum accessible evolution time decreases as a function of the bond dimension. Once such regimes are mapped out, we can assess whether they are within reach of current quantum hardware. Addressing this question would involve several considerations, including the choice of encoding, circuit compilation strategies, and the lattice topology supported by the device.

As an example, in Figure 5.1 we show a simulation of the quench protocol applied to a two-dimensional spin model. Specifically, we simulate the time dynamics of a $J_1 - J_2$ Hamiltonian which has been widely used in describing magnetism:

$$H = J_1 \sum_{\langle i,j \rangle} \sigma_i^z \sigma_j^z + J_2 \sum_{\langle\langle i,j \rangle\rangle} \sigma_i^z \sigma_j^z + h \sum_i \sigma_x, \quad (5.1)$$

where σ_i^α ($\alpha = x, y, z$) are the Pauli matrices acting on site i . The coupling constant J_1 corresponds to nearest-neighbour interactions, summed over all adjacent pairs $\langle i, j \rangle$ on the 2D lattice, while J_2 accounts for next-nearest-neighbour interactions, summed over pairs $\langle\langle i, j \rangle\rangle$. In our simulations, we represent this model using a MPS. From the tensor-network perspective, this is a rather naive approach, since




Figure 5.1: Simulation of the $J_1 - J_2$ model Hamiltonian applied on a square lattice with $L_x = L_y = 11$, for coupling values $J_1 = 1$, $J_2 = 0.5$, and external field $h = 0.01$. **a Evolution of the expectation value of the spin operator $\langle \sigma_i^y(t) \rangle$. **b** The normalised modulus of the spectral function.**

the 2D lattice is encoded into a one-dimensional tensor train. Even with an optimal site ordering, many local interactions in the Hamiltonian become highly non-local in the MPS, requiring larger bond dimensions to achieve accurate approximations. In some instances, a more natural choice for this task would be to employ Projected Entangled Pair States (PEPS) [241]. PEPS generalise MPS to higher dimensions, accommodating arbitrary geometries, while preserving the locality of interactions in two dimensions. This makes them particularly well suited for representing strongly correlated 2D systems, though at a significantly higher computational cost.

In the simulation, we can clearly observe the application of the quench at the centre and the subsequent spread of the perturbation over time. After performing the Fourier transform (now in two spatial dimensions), we obtain the spectral representation. However, the additional spatial dimension comes with a limitation: the momentum-space resolution is directly tied to the system size. In particular, the number of accessible points in $k_{x,y}$ is restricted by the number of qubits along the $L_{x,y}$ direction of the lattice, meaning that even for relatively large qubit counts, the resolution in momentum space may remain coarse.

An exciting direction for further exploration would be to build upon this idea and extend such simulations to fermionic systems, where understanding the interplay

between dimensionality, tensor-network efficiency, and quantum resources could shed light on regimes beyond the reach of classical methods.

5.2.2 Quantum simulation of radical pairs

Another fascinating extension of the work on open quantum dynamics in Chapter 4 is the simulation of radical pairs using quantum computers. Radical pairs, consisting of two radicals each with an unpaired electron, play a central role in both biology and technology. In biology, they are implicated in magnetoreception – the hypothesised quantum mechanism behind the navigational abilities of migratory birds, mediated by radical pairs in cryptochrome molecules [242, 243]. In technology, they underpin phenomena such as dynamic nuclear polarisation in NMR [244], and strongly influence the efficiency of OLEDs and photovoltaic devices [245], where their spin dynamics determine performance. The dynamics of radical pairs are governed by the interplay between their coherent spin evolution and spin-selective recombination, which can be described by the Haberkorn master equation [246],

$$\frac{d}{dt}\hat{\rho}(t) = -\frac{i}{\hbar}[\hat{H}, \hat{\rho}(t)] - \{\hat{K}, \hat{\rho}(t)\}, \quad (5.2)$$

where $\hat{H}$ is the effective spin Hamiltonian and $\hat{K}$ encodes recombination into singlet and triplet channels. The Hamiltonian itself includes exchange and dipolar coupling between the electron spins, Zeeman interactions with external magnetic fields, and hyperfine couplings to surrounding nuclear spins. Importantly, these couplings are modulated by molecular motion, leading to stochastic fluctuations and spin relaxation [247].

Classically, such dynamics are typically simulated using the Stochastic Liouville Equation or Lindblad-type formalisms. While these approaches have yielded insight, they are severely limited by the exponential scaling of the nuclear spin Hilbert space, as well as by approximations that obscure essential dynamical features. Stochastic Schrödinger equation (SSE) techniques offer improvements, but still face challenges in capturing the full non-linear, environment-driven dynamics.

As we discussed in Chapter 4, quantum computers provide a route to efficiently simulate open quantum dynamics. A prospective line of work would be to investigate whether the open-system dynamics of radical pairs can be directly simulated on quantum hardware [248, 249]. Algorithms such as quantum signal processing (QSP) or QITE could be adapted to simulate Lindbladian or stochastic dynamics beyond classical reach. A particularly exciting possibility is to explore connections between non-linear Lindblad dynamics and stochastic spin models of radical pairs, and to design quantum algorithms capable of faithfully capturing these processes. Given the severe limitations of classical approaches, even small-scale quantum devices could already provide valuable insights into these biologically and technologically relevant systems.

To conclude, the field of quantum computing is advancing at an extraordinary pace. While large-scale, fault-tolerant devices remain a long-term goal, current NISQ processors already offer a valuable platform for testing and refining quantum algorithms. Continued focus on algorithmic development is therefore essential: not only does it allow us to extract the most from existing hardware, but it also prepares the ground for applications of genuine scientific and technological relevance. In particular, areas such as materials science and chemistry demand the ability to model complex quantum behaviour, including open-system dynamics. By addressing some of these challenges, the contributions of this thesis highlight both the opportunities and the challenges that lie ahead, and point towards ways in which quantum computers may ultimately deliver significant insights.

Bibliography

- [1] R. P. Feynman, “Simulating physics with computers”, *International Journal of Theoretical Physics* **21**, 467–488 (1982).
- [2] D. Deutsch, “Quantum theory as a universal physical theory”, *International Journal of Theoretical Physics* **24**, 1–41 (1985).
- [3] D. Deutsch, “Rapid solution of problems by quantum computation”, *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences* **439**, 553–558 (1992).
- [4] P. W. Shor, “Algorithms for quantum computation: Discrete logarithms and factoring”, *Proceedings - Annual IEEE Symposium on Foundations of Computer Science, FOCS*, 124–134 (1994).
- [5] L. K. Grover, “A fast quantum mechanical algorithm for database search”, *Proceedings of the Annual ACM Symposium on Theory of Computing* **Part F129452**, 212–219 (1996).
- [6] G. Brassard, P. Hoyer, and A. Tapp, “Quantum Algorithm for the Collision Problem”, *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **1380**, 163–169 (1997).
- [7] W. P. Baritompa, D. W. Bulger, and G. R. Wood, “Grover’s Quantum Algorithm Applied to Global Optimization”, <https://doi.org/10.1137/040605072> **15**, 1170–1184 (2006).
- [8] G. Sun, S. Su, and M. Xu, “Quantum algorithm for polynomial root finding problem”, *Proceedings - 2014 10th International Conference on Computational Intelligence and Security, CIS 2014*, 469–473 (2015).
- [9] M. Fürer, “Solving NP-Complete Problems with Quantum Search”, *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **4957 LNCS**, 784–792 (2008).
- [10] P. W. Shor, “Scheme for reducing decoherence in quantum computer memory”, *Physical Review A* **52**, R2493 (1995).
- [11] A. Y. Kitaev, “Quantum computations: Algorithms and error correction”, *Russian Mathematical Surveys* **52**, 1191–1249 (1997).
- [12] I. Shaik and J. Van de Pol, “Optimal Layout Synthesis for Deep Quantum Circuits on NISQ Processors with 100+ Qubits”, *Leibniz International Proceedings in Informatics, LIPIcs* **305** (2024).
- [13] D. Castelvechi, “IBM releases first-ever 1,000-qubit quantum chip”, *Nature* **624**, 238 (2023).

- [14] G. González-García, R. Trivedi, and J. I. Cirac, “Error Propagation in NISQ Devices for Solving Classical Optimization Problems”, *PRX Quantum* **3**, 040326 (2022).
- [15] C. Gidney and M. Ekerå, “How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits”, *Quantum* **5**, 1–31 (2021).
- [16] J. Preskill, “Quantum computing and the entanglement frontier”, (2012).
- [17] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, et al., “Quantum supremacy using a programmable superconducting processor”, *Nature* **574**, 505–510 (2019).
- [18] E. Pednault, J. A. Gunnels, G. Nannicini, L. Horesh, and R. Wisnieff, “Leveraging Secondary Storage to Simulate Deep 54-qubit Sycamore Circuits”, (2019).
- [19] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, et al., “Evidence for the utility of quantum computing before fault tolerance”, *Nature* **618**, 500–505 (2023).
- [20] J. I. Cirac and P. Zoller, “Goals and opportunities in quantum simulation”, *Nature Physics* 2012 8:4 **8**, 264–266 (2012).
- [21] I. M. Georgescu, S. Ashhab, and F. Nori, “Quantum simulation”, *Reviews of Modern Physics* **86**, 153 (2014).
- [22] R. O. Jones, “Density functional theory: its origins, rise to prominence, and future”, *Reviews of modern physics* **87**, 897 (2015).
- [23] G. K. L. Chan and S. Sharma, “The density matrix renormalization group in quantum chemistry”, *Annual Review of Physical Chemistry* **62**, 465–481 (2011).
- [24] A. Baiardi, “Electron Dynamics with the Time-Dependent Density Matrix Renormalization Group”, *Journal of Chemical Theory and Computation* **17**, 3320–3334 (2021).
- [25] J. Quintanilla and C. Hooley, “The strong-correlations puzzle”, *Physics World* **22**, 32 (2009).
- [26] J. Y. Choi, S. Hild, J. Zeiher, P. Schauß, A. Rubio-Abadal, T. Yefsah, V. Khemani, D. A. Huse, I. Bloch, and C. Gross, “Exploring the many-body localization transition in two dimensions”, *Science* **352**, 1547–1552 (2016).
- [27] C. S. Chiu, G. Ji, A. Bohrdt, M. Xu, M. Knap, E. Demler, F. Grusdt, M. Greiner, and D. Greif, “String patterns in the doped Hubbard model”, *Science* **365**, 251–256 (2019).
- [28] J. Koepsell, J. Vijayan, P. Sompet, F. Grusdt, T. A. Hilker, E. Demler, G. Salomon, I. Bloch, and C. Gross, “Imaging magnetic polarons in the doped Fermi–Hubbard model”, *Nature* 2019 572:7769 **572**, 358–362 (2019).
- [29] S. Lloyd, “Universal quantum simulators”, *Science* **273**, 1073–1078 (1996).
- [30] J. P. F. LeBlanc, A. E. Antipov, F. Becca, I. W. Bulik, G. K.-L. Chan, C.-M. Chung, Y. Deng, M. Ferrero, T. M. Henderson, C. A. Jiménez-Hoyos, et al. (Simons Collaboration on the Many-Electron Problem), “Solutions of the two-dimensional Hubbard model: benchmarks and results from a wide range of numerical algorithms”, *Phys. Rev. X* **5**, 041041 (2015).

- [31] B. X. Zheng, C. M. Chung, P. Corboz, G. Ehlers, M. P. Qin, R. M. Noack, H. Shi, S. R. White, S. Zhang, and G. K. L. Chan, “Stripe order in the underdoped region of the two-dimensional Hubbard model”, *Science* **358**, 1155–1160 (2017).
- [32] M. Motta, C. Sun, A. T. Tan, M. J. O’Rourke, E. Ye, A. J. Minnich, F. G. Brandao, and G. K.-L. Chan, “Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution”, *Nat. Phys.* **16**, 205–210 (2020).
- [33] H. F. Trotter, “On the Product of Semi-Groups of Operators”, *Proceedings of the American Mathematical Society* **10**, 545 (1959).
- [34] M. Suzuki, “Generalized Trotter’s formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems”, *Communications in Mathematical Physics* **51**, 183–190 (1976).
- [35] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, “Theory of trotter error with commutator scaling”, *Physical Review X* **11**, 011020 (2021).
- [36] R. Babbush, J. McClean, D. Wecker, A. Aspuru-Guzik, and N. Wiebe, “Chemical basis of Trotter-Suzuki errors in quantum chemistry simulation”, *Physical Review A - Atomic, Molecular, and Optical Physics* **91**, 022311 (2015).
- [37] S. Raeisi, N. Wiebe, and B. C. Sanders, “Quantum-circuit design for efficient simulations of many-body quantum dynamics”, *New Journal of Physics* **14**, 103017 (2012).
- [38] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, “Elucidating reaction mechanisms on quantum computers”, *Proceedings of the National Academy of Sciences of the United States of America* **114**, 7555–7560 (2017).
- [39] A. M. Childs, A. Ostrander, and Y. Su, “Faster quantum simulation by randomization”, *Quantum* **3**, 182 (2019).
- [40] E. Campbell, “Random Compiler for Fast Hamiltonian Simulation”, *Physical Review Letters* **123**, 070503 (2019).
- [41] C. F. Chen, H. Y. Huang, R. Kueng, and J. A. Tropp, “Concentration for Random Product Formulas”, *PRX Quantum* **2**, 040305 (2021).
- [42] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, “Simulating Hamiltonian Dynamics with a Truncated Taylor Series”, *Physical Review Letters* **114**, 090502 (2015).
- [43] A. M. Childs and N. Wiebe, “Hamiltonian Simulation Using Linear Combinations of Unitary Operations”, *Quantum Information and Computation* **12** (2012).
- [44] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, “Exponential improvement in precision for simulating sparse Hamiltonians”, *Proceedings of the Annual ACM Symposium on Theory of Computing*, 283–292 (2014).
- [45] G. H. Low and I. L. Chuang, “Optimal Hamiltonian Simulation by Quantum Signal Processing”, *Physical Review Letters* **118**, 010501 (2017).
- [46] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, “Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics”, in *Proceedings of the 51st annual acm sigact symposium on theory of computing* (2019), pp. 193–204.

- [47] G. H. Low and I. L. Chuang, “Hamiltonian simulation by qubitization”, *Quantum* **3**, 163 (2019).
- [48] A. M. Childs, “Universal computation by quantum walk”, *Physical Review Letters* **102**, 180501 (2009).
- [49] M. Kliesch, T. Barthel, C. Gogolin, M. Kastoryano, and J. Eisert, “Dissipative quantum Church-Turing theorem”, *Physical Review Letters* **107**, 120501 (2011).
- [50] J. Han, W. Cai, L. Hu, X. Mu, Y. Ma, Y. Xu, W. Wang, H. Wang, Y. Song, C.-L. Zou, et al., “Experimental simulation of open quantum system dynamics via trotterization”, *Phys. Rev. Lett.* **127**, 020504 (2021).
- [51] H. Wang, S. Ashhab, and F. Nori, “Quantum algorithm for simulating the dynamics of an open quantum system”, *Physical Review A - Atomic, Molecular, and Optical Physics* **83**, 062317 (2011).
- [52] R. Cleve and C. Wang, “Efficient Quantum Algorithms for Simulating Lindblad Evolution”, *Leibniz International Proceedings in Informatics, LIPIcs* **80** (2016).
- [53] A. W. Schlimgen, K. Head-Marsden, L. M. Sager, P. Narang, and D. A. Mazziotti, “Quantum Simulation of Open Quantum Systems Using a Unitary Decomposition of Operators”, *Physical Review Letters* **127**, 270503 (2021).
- [54] Z. Hu, R. Xia, and S. Kais, “A quantum algorithm for evolving open quantum dynamics on quantum computing devices”, *Sci. Rep.* **10**, 3301 (2020).
- [55] K. Head-Marsden, S. Krastanov, D. A. Mazziotti, and P. Narang, “Capturing non-Markovian dynamics on near-term quantum computers”, *Physical Review Research* **3**, 013182 (2021).
- [56] H. Kamakari, S.-N. Sun, M. Motta, and A. J. Minnich, “Digital quantum simulation of open quantum systems using quantum imaginary-time evolution”, *PRX Quantum* **3**, 010320 (2022).
- [57] N. Ramusat and V. Savona, “A quantum algorithm for the direct estimation of the steady state of open quantum systems”, *Quantum* **5**, 399 (2021).
- [58] S. R. White, “Density matrix formulation for quantum renormalization groups”, *Physical Review Letters* **69**, 2863 (1992).
- [59] C. Lanczos, “An iteration method for the solution of the eigenvalue problem of linear differential and integral operators”, *Journal of Research of the National Bureau of Standards* **45**, 255–282 (1959).
- [60] G. Catarina and B. Murta, “Density-matrix renormalization group: a pedagogical introduction”, *The European Physical Journal B* 2023 96:8 **96**, 1–30 (2023).
- [61] S. Paeckel, T. Köhler, A. Swoboda, S. R. Manmana, U. Schollwöck, and C. Hubig, “Time-evolution methods for matrix-product states”, *Annals of Physics* **411**, 167998 (2019).
- [62] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pi žorn, H. Verschelde, and F. Verstraete, “Time-dependent variational principle for quantum lattices”, *Phys. Rev. Lett.* **107**, 070601 (2011).
- [63] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, “Unifying time evolution and optimization with matrix product states”, *Phys. Rev. B* **94**, 165116 (2016).

- [64] J. Sun, L. Vilchez-Estevez, V. Vedral, A. T. Boothroyd, and M. S. Kim, “Probing spectral features of quantum many-body systems with quantum simulators”, *Nature Communications* **16**, 1403 (2025).
- [65] S. Mukamel, *Principles of nonlinear optical spectroscopy*, 6 (Oxford University Press, USA, 1995).
- [66] S. W. Lovesey, *Theory of neutron scattering from condensed matter* (1984).
- [67] A. T. Boothroyd, *Principles of neutron scattering from condensed matter* (Oxford University Press, 2020).
- [68] J. Als-Nielsen and D. McMorrow, *Elements of modern x-ray physics* (John Wiley & Sons, 2011).
- [69] M. Knap, A. Kantian, T. Giamarchi, I. Bloch, M. D. Lukin, and E. Demler, “Probing real-space and time-resolved correlation functions with many-body ramsey interferometry”, *Physical review letters* **111**, 147205 (2013).
- [70] Y. Wan and N. Armitage, “Resolving continua of fractional excitations by spinon echo in thz 2d coherent spectroscopy”, *Physical Review Letters* **122**, 257401 (2019).
- [71] C. Senko, J. Smith, P. Richerme, A. Lee, W. Campbell, and C. Monroe, “Coherent imaging spectroscopy of a quantum many-body spin system”, *Science* **345**, 430–433 (2014).
- [72] P. Jurcevic, B. P. Lanyon, P. Hauke, C. Hempel, P. Zoller, R. Blatt, and C. F. Roos, “Quasiparticle engineering and entanglement propagation in a quantum many-body system”, *Nature* **511**, 202–205 (2014).
- [73] A. Valdés-Curiel, D. Trypogeorgos, E. Marshall, and I. Spielman, “Fourier transform spectroscopy of a spin-orbit coupled bose gas”, *New journal of physics* **19**, 033025 (2017).
- [74] M. Dalmonte, B. Vermersch, and P. Zoller, “Quantum simulation and spectroscopy of entanglement hamiltonians”, *Nature Physics* **14**, 827–831 (2018).
- [75] S. R. Manmana, S. Wessel, R. M. Noack, and A. Muramatsu, “Time evolution of correlations in strongly interacting fermions after a quantum quench”, *Physical Review B* **79**, 155104 (2009).
- [76] B. Yoshimura, W. Campbell, and J. Freericks, “Diabatic-ramping spectroscopy of many-body excited states”, *Physical Review A* **90**, 062334 (2014).
- [77] P. Richerme, Z.-X. Gong, A. Lee, C. Senko, J. Smith, M. Foss-Feig, S. Michalakis, A. V. Gorshkov, and C. Monroe, “Non-local propagation of correlations in quantum systems with long-range interactions”, *Nature* **511**, 198–201 (2014).
- [78] R. Menu and T. Roscilde, “Quench dynamics of quantum spin models with flat bands of excitations”, *Physical Review B* **98**, 205145 (2018).
- [79] C.-K. Lee, J. W. Zhong Lau, L. Shi, and L. C. Kwek, “Simulating energy transfer in molecular systems with digital quantum computers”, *Journal of Chemical Theory and Computation* **18**, 1347–1358 (2022).
- [80] A. Miessen, P. J. Ollitrault, F. Tacchino, and I. Tavernelli, “Quantum algorithms for quantum dynamics”, *Nature Computational Science* **3**, 25–37 (2023).

- [81] Y. Ge, J. Tura, and J. I. Cirac, “Faster ground state preparation and high-precision ground energy estimation with fewer qubits”, *Journal of Mathematical Physics* **60**, 022202 (2019).
- [82] G. H. Low and I. L. Chuang, “Hamiltonian simulation by uniform spectral amplification”, *arXiv preprint arXiv:1707.05391* (2017).
- [83] C.-K. Lee, C.-Y. Hsieh, S. Zhang, and L. Shi, “Simulation of condensed-phase spectroscopy with near-term digital quantum computers”, *Journal of Chemical Theory and Computation* **17**, 7178–7186 (2021).
- [84] K. Bharti, A. Cervera-Liarta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, et al., “Noisy intermediate-scale quantum algorithms”, *Reviews of Modern Physics* **94**, 015004 (2022).
- [85] J. Tangpanitanon, A. Megrant, Z. Chen, B. Chiaro, A. Dunsworth, et al., “Spectroscopic signatures of localization with interacting photons in superconducting qubits”, *Science* **358**, 1175–1179 (2017).
- [86] L. Villa, J. Despres, and L. Sanchez-Palencia, “Unraveling the excitation spectrum of many-body systems from quantum quenches”, *Phys. Rev. A* **100**, 063632 (2019).
- [87] L. Villa, J. Despres, S. J. Thomson, and L. Sanchez-Palencia, “Local quench spectroscopy of many-body quantum systems”, *Phys. Rev. A* **102**, 033337 (2020).
- [88] H. H. S. Chan, R. Meister, M. L. Goh, and B. Koczor, “Algorithmic shadow spectroscopy”, *arXiv preprint arXiv:2212.11036* (2022).
- [89] M. Kowalewski, B. P. Fingerhut, K. E. Dorfman, K. Bennett, and S. Mukamel, “Simulating coherent multidimensional spectroscopy of nonadiabatic molecular processes: from the infrared to the x-ray regime”, *Chemical reviews* **117**, 12165–12226 (2017).
- [90] A. Chiesa, F. Tacchino, M. Grossi, P. Santini, I. Tavernelli, D. Gerace, and S. Carretta, “Quantum hardware simulating four-dimensional inelastic neutron scattering”, *Nature Physics* **15**, 455–459 (2019).
- [91] J. P. Stenger, G. Ben-Shach, D. Pekker, and N. T. Bronn, “Simulating spectroscopy experiments with a superconducting quantum computer”, *Physical Review Research* **4**, 043106 (2022).
- [92] Y. Matsuzaki, H. Hakoshima, K. Sugisaki, Y. Seki, and S. Kawabata, “Direct estimation of the energy gap between the ground state and excited state with quantum annealing”, *Japanese Journal of Applied Physics* **60**, SBBI02 (2021).
- [93] I. Zintchenko and N. Wiebe, “Randomized gap and amplitude estimation”, *Physical Review A* **93**, 062306 (2016).
- [94] K. Sugisaki, K. Toyota, K. Sato, D. Shiomi, and T. Takui, “Quantum algorithm for the direct calculations of vertical ionization energies”, *The Journal of Physical Chemistry Letters* **12**, 2880–2885 (2021).
- [95] K. Sugisaki, C. Sakai, K. Toyota, K. Sato, D. Shiomi, and T. Takui, “Bayesian phase difference estimation: a general quantum algorithm for the direct calculation of energy gaps”, *Physical Chemistry Chemical Physics* **23**, 20152–20162 (2021).

- [96] C. Huang, F. Zhang, M. Newman, J. Cai, X. Gao, Z. Tian, J. Wu, H. Xu, H. Yu, B. Yuan, et al., “Classical simulation of quantum supremacy circuits”, arXiv preprint arXiv:2005.06787 (2020).
- [97] P. Zeng, J. Sun, and X. Yuan, “Universal quantum algorithmic cooling on a quantum computer”, arXiv preprint arXiv:2109.15304 (2021).
- [98] L. Lin and Y. Tong, *Heisenberg-limited ground state energy estimation for early fault-tolerant quantum computers*, 2021.
- [99] G. Wang, D. Stilck-França, R. Zhang, S. Zhu, and P. D. Johnson, “Quantum algorithm for ground state energy estimation using circuit depth with exponentially improved dependence on precision”, arXiv preprint arXiv:2209.06811 (2022).
- [100] Y. Gu, Y. Ma, N. Forcellini, and D. E. Liu, “Noise-resilient phase estimation with randomized compiling”, *Physical Review Letters* **130**, 250601 (2023).
- [101] R. Kubo, “The fluctuation-dissipation theorem”, *Rep. Prog. Phys.* **29**, 255 (1966).
- [102] P. Jurcevic, P. Hauke, C. Maier, C. Hempel, B. Lanyon, R. Blatt, and C. F. Roos, “Spectroscopy of interacting quasiparticles in trapped ions”, *Physical Review Letters* **115**, 100501 (2015).
- [103] P. Bleicker and G. S. Uhrig, “Strong quenches in the one-dimensional fermi-hubbard model”, *Physical Review A* **98**, 033602 (2018).
- [104] J. Vijayan, P. Sompet, G. Salomon, J. Koepsell, S. Hirthe, A. Bohrdt, F. Grusdt, I. Bloch, and C. Gross, “Time-resolved observation of spin-charge deconfinement in fermionic hubbard chains”, *Science* **367**, 186–189 (2020).
- [105] L. Villa, S. Thomson, and L. Sanchez-Palencia, “Finding the phase diagram of strongly correlated disordered bosons using quantum quenches”, *Physical Review A* **104**, 023323 (2021).
- [106] L. Villa, “Out-of-equilibrium dynamics and quench spectroscopy of ultracold many-body quantum systems”, PhD thesis (Institut polytechnique de Paris, 2021).
- [107] P. Jurcevic, H. Shen, P. Hauke, C. Maier, T. Brydges, C. Hempel, B. Lanyon, M. Heyl, R. Blatt, and C. Roos, “Direct observation of dynamical quantum phase transitions in an interacting many-body system”, *Physical review letters* **119**, 080501 (2017).
- [108] X.-G. Wen, *Quantum field theory of many-body systems: from the origin of sound to an origin of light and electrons* (Oxford university press, 2004).
- [109] J. Bonča, S. A. Trugman, and M. Berciu, “Spectral function of the holstein polaron at finite temperature”, *Physical Review B* **100**, 094307 (2019).
- [110] D. Jansen, J. Bonča, and F. Heidrich-Meisner, “Finite-temperature density-matrix renormalization group method for electron-phonon systems: thermodynamics and holstein-polaron spectral functions”, *Physical Review B* **102**, 165155 (2020).
- [111] F. Schrodi, P. Silvi, F. Tschirsich, R. Fazio, and S. Montangero, “Density of states of many-body quantum systems from tensor networks”, *Physical Review B* **96**, 094303 (2017).

- [112] P. Zeng, J. Sun, L. Jiang, and Q. Zhao, “Simple and high-precision hamiltonian simulation by compensating trotter error with linear combination of unitary operations”, arXiv preprint arXiv:2212.04566 (2022).
- [113] J. Sun, X. Yuan, T. Tsunoda, V. Vedral, S. C. Benjamin, and S. Endo, “Mitigating realistic noise in practical noisy intermediate-scale quantum devices”, *Physical Review Applied* **15**, 034026 (2021).
- [114] A. A. Houck, H. E. Türeci, and J. Koch, “On-chip quantum simulation with superconducting circuits”, *Nature Physics* **8**, 292 (2012).
- [115] S. Endo, S. C. Benjamin, and Y. Li, “Practical quantum error mitigation for near-future applications”, *Physical Review X* **8**, 031027 (2018).
- [116] K. Temme, S. Bravyi, and J. M. Gambetta, “Error mitigation for short-depth quantum circuits”, *Physical review letters* **119**, 180509 (2017).
- [117] J. J. Wallman and J. Emerson, “Noise tailoring for scalable quantum computation via randomized compiling”, *Phys. Rev. A* **94**, 052325 (2016).
- [118] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, “Randomized benchmarking of quantum gates”, *Physical Review A—Atomic, Molecular, and Optical Physics* **77**, 012307 (2008).
- [119] A. M. Dalzell, N. Hunter-Jones, and F. G. Brandão, “Random quantum circuits transform local noise into global white noise”, *Communications in Mathematical Physics* **405**, 78 (2024).
- [120] J. Foldager and B. Koczor, “Can shallow quantum circuits scramble local noise into global white noise?”, *Journal of Physics A: Mathematical and Theoretical* **57**, 015306 (2023).
- [121] H.-Y. Huang, R. Kueng, and J. Preskill, “Efficient estimation of pauli observables by derandomization”, *Physical review letters* **127**, 030503 (2021).
- [122] B. Wu, J. Sun, Q. Huang, and X. Yuan, “Overlapped grouping measurement: a unified framework for measuring quantum states”, *Quantum* **7**, 896 (2023).
- [123] H.-Y. Huang, R. Kueng, and J. Preskill, “Predicting many properties of a quantum system from very few measurements”, *Nature Physics* **16**, 1050–1057 (2020).
- [124] L. Vilchez-Estevez, R. A. Santos, S. Y. Wang, and F. M. Gambetta, “Extracting the spin excitation spectrum of a fermionic system using a quantum processor”, *Phys. Rev. B* **112**, 045143 (2025).
- [125] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, et al., “Quantum computational advantage using photons”, *Science* **370**, 1460–1463 (2020).
- [126] A. Morvan, B. Villalonga, X. Mi, S. Mandrà, A. Bengtsson, P. V. Klimov, Z. Chen, S. Hong, C. Erickson, I. K. Drozdov, et al., “Phase transitions in random circuit sampling”, *Nature* **634**, 328–333 (2024).
- [127] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*, 10th Anniversary (Cambridge University Press, Cambridge, 2010).

- [128] I. A. Zaliznyak and J. M. Tranquada, “Neutron scattering and its application to strongly correlated systems”, in *Strongly correlated systems: experimental techniques*, edited by A. Avella and F. Mancini (Springer Berlin Heidelberg, Berlin, Heidelberg, 2015), pp. 205–235.
- [129] L. J. P. Ament, M. van Veenendaal, T. P. Devereaux, J. P. Hill, and J. van den Brink, “Resonant inelastic x-ray scattering studies of elementary excitations”, *Rev. Mod. Phys.* **83**, 705–767 (2011).
- [130] R. F. Egerton, “Electron energy-loss spectroscopy in the TEM”, *Rep. Prog. Phys.* **72**, 016502 (2008).
- [131] D. Wecker, M. B. Hastings, N. Wiebe, B. K. Clark, C. Nayak, and M. Troyer, “Solving strongly correlated electron models on a quantum computer”, *Phys. Rev. A* **92**, 062318 (2015).
- [132] D. Wecker, M. B. Hastings, and M. Troyer, “Progress towards practical quantum variational algorithms”, *Phys. Rev. A* **92**, 042303 (2015).
- [133] Z. Jiang, K. J. Sung, K. Kechedzhi, V. N. Smelyanskiy, and S. Boixo, “Quantum algorithms to simulate many-body physics of correlated fermions”, *Physical Review Applied* **9**, 044036 (2018).
- [134] A. Montanaro and S. Stanisic, *Compressed variational quantum eigensolver for the Fermi-Hubbard model*, 2020.
- [135] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, A. Bengtsson, S. Boixo, M. Broughton, B. B. Buckley, et al., *Observation of separated dynamics of charge and spin in the Fermi-Hubbard model*, 2020.
- [136] C. Cade, L. Mineh, A. Montanaro, and S. Stanisic, “Strategies for solving the Fermi-Hubbard model on near-term quantum computers”, *Phys. Rev. B* **102**, 235122 (2020).
- [137] L. Clinton, J. Bausch, and T. Cubitt, “Hamiltonian simulation algorithms for near-term quantum hardware”, *Nat. Commun.* **12**, 4989 (2021).
- [138] S. Stanisic, J. L. Bosse, F. M. Gambetta, R. A. Santos, W. Mruczkiewicz, T. E. O’Brien, E. Ostby, and A. Montanaro, “Observing ground-state properties of the Fermi-Hubbard model using a scalable algorithm on a quantum computer”, *Nat. Commun.* **13**, 5743 (2022).
- [139] A. M. Dalzell, S. McArdle, M. Berta, P. Bienias, C.-F. Chen, A. Gilyén, C. T. Hann, M. J. Kastoryano, E. T. Khabiboulline, A. Kubica, et al., *Quantum algorithms: a survey of applications and end-to-end complexities*, 2023.
- [140] K. Hémary, K. Ghanem, E. Crane, S. L. Campbell, J. M. Dreiling, C. Figgatt, C. Foltz, J. P. Gaebler, J. Johansen, M. Mills, et al., “Measuring the Loschmidt amplitude for finite-energy properties of the Fermi-Hubbard model on an ion-trap quantum computer”, *PRX Quantum* **5**, 030323 (2024).
- [141] R. Nigmatullin, K. Hemery, K. Ghanem, S. Moses, D. Gresh, P. Siegfried, M. Mills, T. Gatterman, N. Hewitt, E. Granet, et al., *Experimental demonstration of break-even for the compact fermionic encoding*, 2024.
- [142] T. A. Bespalova, K. Delić, G. Pupillo, F. Tacchino, and I. Tavernelli, *Simulating the fermi-hubbard model with long-range hopping on a quantum computer*, 2024.

- [143] R. Au-Yeung, B. Camino, O. Rathore, and V. Kendon, “Quantum algorithms for scientific computing”, Rep. Prog. Phys. **87**, 116001 (2024).
- [144] M. Imada, A. Fujimori, and Y. Tokura, “Metal-insulator transitions”, Rev. Mod. Phys. **70**, 1039–1263 (1998).
- [145] P. A. Lee, N. Nagaosa, and X.-G. Wen, “Doping a Mott insulator: physics of high-temperature superconductivity”, Rev. Mod. Phys. **78**, 17–85 (2006).
- [146] B. Keimer, S. A. Kivelson, M. R. Norman, S. Uchida, and J. Zaanen, “From quantum matter to high-temperature superconductivity in copper oxides”, Nature **518**, 179–186 (2015).
- [147] E. Fradkin, S. A. Kivelson, and J. M. Tranquada, “Colloquium: theory of intertwined orders in high temperature superconductors”, Rev. Mod. Phys. **87**, 457–482 (2015).
- [148] H.-C. Jiang and T. P. Devereaux, “Superconductivity in the doped Hubbard model and its interplay with next-nearest hopping t' ”, Science **365**, 1424–1428 (2019).
- [149] D. P. Arovas, E. Berg, S. A. Kivelson, and S. Raghu, “The hubbard model”, Annu. Rev. Condens. Matter Phys. **13**, 239–274 (2022).
- [150] M. Qin, T. Schäfer, S. Andergassen, P. Corboz, and E. Gull, “The Hubbard model: a computational perspective”, Annu. Rev. Condens. Matter Phys. **13**, 275–302 (2022).
- [151] I. A. Zaliznyak, H. Woo, T. G. Perring, C. L. Broholm, C. D. Frost, and H. Takagi, “Spinons in the strongly correlated copper oxide chains in SrCuO_2 ”, Phys. Rev. Lett. **93**, 087202 (2004).
- [152] M. Eschrig, “The effect of collective spin-1 excitations on electronic spectra in high-Tc superconductors”, Adv. Phys. **55**, 47–183 (2006).
- [153] A. C. Walters, T. G. Perring, J.-S. Caux, A. T. Savici, G. D. Gu, C.-C. Lee, W. Ku, and I. A. Zaliznyak, “Effect of covalent bonding on magnetism and the missing neutron intensity in copper oxide compounds”, Nat. Phys. **5**, 867–872 (2009).
- [154] C. J. Jia, E. A. Nowadnick, K. Wohlfeld, Y. F. Kung, C. -. Chen, S. Johnston, T. Tohyama, B. Moritz, and T. P. Devereaux, “Persistent spin excitations in doped antiferromagnets revealed by resonant inelastic light scattering”, Nat. Commun. **5**, 3314 (2014).
- [155] S. Li, A. Nocera, U. Kumar, and S. Johnston, “Particle-hole asymmetry in the dynamical spin and charge responses of corner-shared 1D cuprates”, Commun. Phys. **4**, 217 (2021).
- [156] G. Giuliani and G. Vignale, *Quantum theory of the electron liquid* (Cambridge University Press, 2005).
- [157] R. M. Martin, L. Reining, and D. M. Ceperley, *Interacting electrons: theory and computational approaches* (Cambridge University Press, 2016).
- [158] K. Sturm, “Dynamic structure factor: an introduction”, Z. Naturforsch. A **48**, 233–242 (1993).

- [159] M. L. Baez, M. Goihl, J. Haferkamp, J. Bermejo-Vega, M. Gluza, and J. Eisert, “Dynamical structure factors of dynamical quantum simulators”, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 26123–26134 (2020).
- [160] P. Roushan, C. Neill, J. Tangpanitanon, V. M. Bastidas, A. Megrant, R. Barends, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, et al., “Spectroscopic signatures of localization with interacting photons in superconducting qubits”, *Science* **358**, 1175–1179 (2017).
- [161] C. Chen, G. Emperauger, G. Bornet, F. Caleca, B. Gély, M. Bintz, S. Chatterjee, V. Liu, D. Barredo, N. Y. Yao, et al., *Spectroscopy of elementary excitations from quench dynamics in a dipolar XY Rydberg simulator*, 2024.
- [162] J. Sólyom, *Fundamentals of the physics of solids*, Vol. 3 (Springer Berlin Heidelberg, 2010).
- [163] R. Lyons, *Understanding digital signal processing* (Pearson Education International, 2011).
- [164] N. Schuch and F. Verstraete, “Computational complexity of interacting electrons and fundamental limitations of density functional theory”, *Nat. Phys.* **5**, 732–735 (2009).
- [165] S. Gharibian, Y. Huang, Z. Landau, and S. W. Shin, “Quantum Hamiltonian complexity”, *Found. Trends Theor. Comput. Sci.* **10**, 159–282 (2015).
- [166] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, “A variational eigenvalue solver on a photonic quantum processor”, *Nat. Commun.* **5**, 4213 (2014).
- [167] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, “The theory of variational hybrid quantum-classical algorithms”, *New J. Phys.* **18**, 023023 (2016).
- [168] Z. Cai, “Resource estimation for quantum variational simulations of the Hubbard model”, *Phys. Rev. Appl.* **14**, 014059 (2020).
- [169] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, “Image quality assessment: from error visibility to structural similarity”, *IEEE Trans. Image Process.* **13**, 600–612 (2004).
- [170] A. K. Venkataramanan, C. Wu, A. C. Bovik, I. Katsavounidis, and Z. Shahid, “A hitchhiker’s guide to structural similarity”, *IEEE Access* **9**, 28872–28896 (2021).
- [171] J. Ostmeier, “Optimised Trotter decompositions for classical and quantum computing”, *J. Phys. A* **56**, 285303 (2023).
- [172] F. Verstraete, J. I. Cirac, and J. I. Latorre, “Quantum circuits for strongly correlated quantum systems”, *Phys. Rev. A* **79**, 032316 (2009).
- [173] B. Kraus and J. I. Cirac, “Optimal creation of entanglement using a two-qubit gate”, *Phys. Rev. A* **63**, 062309 (2001).
- [174] A. Smith, M. S. Kim, F. Pollmann, and J. Knolle, “Simulating quantum many-body dynamics on a current digital quantum computer”, *npj Quantum Inf.* **5**, 106 (2019).
- [175] L. Clinton, T. Cubitt, B. Flynn, F. M. Gambetta, J. Klassen, A. Montanaro, S. Piddock, R. A. Santos, and E. Sheridan, “Towards near-term quantum simulation of materials”, *Nat. Commun.* **15**, 211 (2024).

- [176] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. K.-L. Chan, and R. Babbush, “Quantum simulation of electronic structure with linear depth and connectivity”, *Phys. Rev. Lett.* **120**, 110501 (2018).
- [177] K. Bharti, A. Cervera-Liarta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, et al., “Noisy intermediate-scale quantum algorithms”, *Rev. Mod. Phys.* **94**, 015004 (2022).
- [178] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, et al., “The variational quantum eigensolver: a review of methods and best practices”, *Phys. Rep.* **986**, 1–128 (2022).
- [179] M. Newman and G. Barkema, *Monte carlo methods in statistical physics* (Clarendon Press, 1999).
- [180] L. Viola and S. Lloyd, “Dynamical suppression of decoherence in two-state quantum systems”, *Phys. Rev. A* **58**, 2733 (1998).
- [181] B. Pokharel, N. Anand, B. Fortman, and D. A. Lidar, “Demonstration of Fidelity Improvement Using Dynamical Decoupling with Superconducting Qubits”, *Phys. Rev. Lett.* **121**, 220502 (2018).
- [182] A. Javadi-Abhari, M. Treinish, K. Krsulich, C. J. Wood, J. Lishman, J. Gacon, S. Martiel, P. D. Nation, L. S. Bishop, A. W. Cross, et al., *Quantum computing with Qiskit*, 2024.
- [183] A. Hashim, R. K. Naik, A. Morvan, J.-L. Ville, B. Mitchell, J. M. Kreikebaum, M. Davis, E. Smith, C. Iancu, K. P. O’Brien, et al., “Randomized compiling for scalable quantum computing on a noisy superconducting quantum processor”, *Phys. Rev. X* **11**, 041039 (2021).
- [184] J. Robledo-Moreno, M. Motta, H. Haas, A. Javadi-Abhari, P. Jurcevic, W. Kirby, S. Martiel, K. Sharma, S. Sharma, T. Shirakawa, et al., *Chemistry beyond exact solutions on a quantum-centric supercomputer*, 2024.
- [185] L. Vilchez-Estevez, A. Yosifov, and J. Sun, “Direct probing of the simulation complexity of open quantum many-body dynamics”, arXiv:2508.19959 (2025).
- [186] H.-P. Breuer and F. Petruccione, *The theory of open quantum systems* (Oxford University Press, 2002).
- [187] A. Rivas and S. F. Huelga, *Open quantum systems*, Vol. 10 (Springer, 2012).
- [188] M. Mohseni, Y. Omar, G. S. Engel, and M. B. Plenio, *Quantum effects in biology* (Cambridge University Press, 2014).
- [189] M. Mohseni, A. Shabani, S. Lloyd, Y. Omar, and H. Rabitz, “Geometrical effects on energy transfer in disordered open quantum systems”, *J. Chem. Phys.* **138**, 204309 (2013).
- [190] B. Bauer, S. Bravyi, M. Motta, and G. K.-L. Chan, “Quantum algorithms for quantum chemistry and quantum materials science”, *Chem. Rev.* **120**, 12685–12717 (2020).
- [191] N. Yunger Halpern, P. Faist, J. Oppenheim, and A. Winter, “Microcanonical and resource-theoretic derivations of the thermal state of a quantum system with noncommuting charges”, *Nat. Commun.* **7**, 12051 (2016).

- [192] Á. Rivas, “Strong coupling thermodynamics of open quantum systems”, *Phys. Rev. Lett.* **124**, 160601 (2020).
- [193] H. M. Wiseman and G. J. Milburn, *Quantum measurement and control* (Cambridge University Press, 2009).
- [194] S. Kallush, R. Dann, and R. Kosloff, “Controlling the uncontrollable: quantum control of open-system dynamics”, *Sci. Adv.* **8**, eadd0828 (2022).
- [195] M. Cattaneo, M. A. Rossi, G. García-Pérez, R. Zambrini, and S. Maniscalco, “Quantum simulation of dissipative collective effects on noisy quantum computers”, *PRX Quantum* **4**, 010324 (2023).
- [196] G. Akemann, M. Kieburg, A. Mielke, and T. Prosen, “Universal signature from integrability to chaos in dissipative open quantum systems”, *Phys. Rev. Lett.* **123**, 254101 (2019).
- [197] I. Reichental, A. Klemmner, Y. Kafri, and D. Podolsky, “Thermalization in open quantum systems”, *Phys. Rev. B* **97**, 134301 (2018).
- [198] L. Sá, P. Ribeiro, and T. Prosen, “Integrable nonunitary open quantum circuits”, *Phys. Rev. B* **103**, 115132 (2021).
- [199] F. Ticozzi and L. Viola, “Quantum and classical resources for unitary design of open-system evolutions”, *Quantum Sci. Technol.* **2**, 034001 (2017).
- [200] A. W. Schlingens, K. Head-Marsden, L. M. Sager-Smith, P. Narang, and D. A. Mazziotti, “Quantum simulation of open quantum systems using density-matrix purification”, *arXiv:2207.07112* (2022).
- [201] D. Kretschmann, D. Schlingemann, and R. F. Werner, “The information-disturbance tradeoff and the continuity of stinespring’s representation”, *IEEE Trans. Inf. Theory* **54**, 1708–1717 (2008).
- [202] A. W. Schlingens, K. Head-Marsden, L. M. Sager-Smith, P. Narang, and D. A. Mazziotti, “Quantum state preparation and nonunitary evolution with diagonal operators”, *Phys. Rev. A* **106**, 022414 (2022).
- [203] Z. Wang, R. Qi, Y. Lu, Z. Wu, and J. Jie, “Absence of correlations in dissipative interacting qubits: a no-go theorem”, *Phys. Rev. B* **110**, 155129 (2024).
- [204] V. Alba and F. Carollo, “Spreading of correlations in markovian open quantum systems”, *Phys. Rev. B* **103**, L020302 (2021).
- [205] Y. Zhao, Y. Tian, J. Ye, Y. Wu, Z. Zhao, Z. Chi, T. Tian, H. Yao, J. Hu, Y. Chen, et al., “Universal dissipative dynamics in strongly correlated quantum gases”, *Nat. Phys.* **21**, 530–535 (2025).
- [206] G. Barontini, R. Labouvie, F. Stubenrauch, A. Vogler, V. Guarrera, and H. Ott, “Controlling the dynamics of an open many-body quantum system with localized dissipation”, *Phys. Rev. Lett.* **110**, 035302 (2013).
- [207] T. Prosen and I. Pižorn, “Quantum phase transition in a far-from-equilibrium steady state of an xy spin chain”, *Phys. Rev. Lett.* **101**, 105701 (2008).
- [208] I. A. Aloisio, G. A. White, C. D. Hill, and K. Modi, “Sampling complexity of open quantum systems”, *PRX Quantum* **4**, 020310 (2023).
- [209] B. Bhattacharjee, X. Cao, P. Nandy, and T. Pathak, “Operator growth in open quantum systems: lessons from the dissipative syk”, *JHEP* **2023**, 54 (2023).

- [210] D. Fang, J. Lu, and Y. Tong, “Mixing time of open quantum systems via hypocoercivity”, *Phys. Rev. Lett.* **134**, 140405 (2025).
- [211] S. Alipour, A. Rezakhani, A. Babu, K. Mølmer, M. Möttönen, and T. Ala-Nissila, “Correlation-picture approach to open-quantum-system dynamics”, *Phys. Rev. X* **10**, 041024 (2020).
- [212] A. Babu, S. Alipour, A. Rezakhani, and T. Ala-Nissila, “Unfolding system-environment correlation in open quantum systems: revisiting master equations and the born approximation”, *Phys. Rev. Res.* **6**, 013243 (2024).
- [213] D. Gribben, A. Strathearn, G. E. Fux, P. Kirton, and B. W. Lovett, “Using the environment to understand non-markovian open quantum systems”, *Quantum* **6**, 847 (2022).
- [214] H.-P. Breuer and J. Piilo, “Stochastic jump processes for non-markovian quantum dynamics”, *EPL* **85**, 50004 (2009).
- [215] Y.-G. Liu, H. Fan, and S. Chen, “Digital quantum simulation of the nonlinear lindblad master equation based on quantum trajectory averaging”, *arXiv:2504.00121* (2025).
- [216] L. Lin, “Dissipative preparation of many-body quantum states: towards practical quantum advantage”, *arXiv:2505.21308* (2025).
- [217] Z. Ding, C.-F. Chen, and L. Lin, “Single-ancilla ground state preparation via lindbladians”, *Phys. Rev. Res.* **6**, 033147 (2024).
- [218] C. Gneiting, A. V. Rozhkov, and F. Nori, “Jump-time unraveling of markovian open quantum systems”, *Phys. Rev. A* **104**, 062212 (2021).
- [219] X. Niu, J. Li, S. Wu, and X. Yi, “Effect of quantum jumps on non-hermitian systems”, *Phys. Rev. A* **108**, 032214 (2023).
- [220] G. Peretto, F. Carollo, and I. Lesanovsky, “Thermodynamics of quantum-jump trajectories of open quantum systems subject to stochastic resetting”, *SciPost Phys.* **13**, 079 (2022).
- [221] Y.-G. Liu and S. Chen, “Lindbladian dynamics with loss of quantum jumps”, *Phys. Rev. B* **111**, 024303 (2025).
- [222] R. Dum, A. Parkins, P. Zoller, and C. Gardiner, “Monte carlo simulation of master equations in quantum optics for vacuum, thermal, and squeezed reservoirs”, *Phys. Rev. A* **46**, 4382 (1992).
- [223] J. Dalibard, Y. Castin, and K. Mølmer, “Wave-function approach to dissipative processes in quantum optics”, *Phys. Rev. Lett.* **68**, 580 (1992).
- [224] K. Mølmer, Y. Castin, and J. Dalibard, “Monte carlo wave-function method in quantum optics”, *J. Opt. Soc. Am. B* **10**, 524–538 (1993).
- [225] S. Endo, J. Sun, Y. Li, S. C. Benjamin, and X. Yuan, “Variational quantum simulation of general processes”, *Physical Review Letters* **125**, 010501 (2020).
- [226] M. Pocrnic, D. Segal, and N. Wiebe, “Quantum simulation of lindbladian dynamics via repeated interactions”, *J. Phys. A: Math. Theor.* **58**, 305302 (2025).
- [227] D. An, A. M. Childs, and L. Lin, “Quantum algorithm for linear non-unitary dynamics with near-optimal dependence on all parameters”, *arXiv:2312.03916* (2023).

- [228] A. Gilchrist, D. R. Terno, and C. J. Wood, “Vectorization of quantum operations and its use”, arXiv:0911.2539 (2009).
- [229] H. H. S. Chan, D. M. Ramo, and N. Fitzpatrick, “Simulating non-unitary dynamics using quantum signal processing with unitary block encoding”, arXiv:2303.06161 (2023).
- [230] D. Poulin, “Lieb-robinson bound and locality for general markovian quantum dynamics”, *Phys. Rev. Lett.* **104**, 190401 (2010).
- [231] N. Schuch, M. M. Wolf, K. G. H. Vollbrecht, and J. I. Cirac, “On entropy growth and the hardness of simulating time evolution”, *New J. Phys.* **10**, 033032 (2008).
- [232] E. H. Lieb and D. W. Robinson, “The finite group velocity of quantum spin systems”, *Commun. Math. Phys.* **28**, 251–257 (1972).
- [233] F. G. Brandao and M. Horodecki, “Exponential decay of correlations implies area law”, *Commun. Math. Phys.* **333**, 761–798 (2015).
- [234] F. G. Brandao and M. J. Kastoryano, “Finite correlation length implies efficient preparation of quantum thermal states”, *Commun. Math. Phys.* **365**, 1–16 (2019).
- [235] M. B. Hastings and T. Koma, “Spectral gap and exponential decay of correlations”, *Commun. Math. Phys.* **265**, 781–804 (2006).
- [236] F. Benatti, R. Floreanini, and L. Memarzadeh, “Exact steady state of the open xx-spin chain: entanglement and transport properties”, *PRX Quantum* **2**, 030344 (2021).
- [237] G. T. Landi, D. Poletti, and G. Schaller, “Nonequilibrium boundary-driven quantum systems: models, methods, and properties”, *Rev. Mod. Phys.* **94**, 045006 (2022).
- [238] D. Wellnitz, G. Preisser, V. Alba, J. Dubail, and J. Schachenmayer, “Rise and fall, and slow rise again, of operator entanglement under dephasing”, *Phys. Rev. Lett.* **129**, 170401 (2022).
- [239] J. Haferkamp, D. Hangleiter, A. Bouland, B. Fefferman, J. Eisert, and J. Bermejo-Vega, “Closing Gaps of a Quantum Advantage with Short-Time Hamiltonian Dynamics”, *Physical Review Letters* **125**, 250501 (2020).
- [240] Y. Wu, Y. Zhang, and X. Yuan, “An Efficient Classical Algorithm for Simulating Short Time 2D Quantum Dynamics”, (2024).
- [241] F. Verstraete and J. I. Cirac, “Renormalization algorithms for quantum-many body systems in two and higher dimensions”, arXiv preprint cond-mat/0407066 (2004).
- [242] C. T. Rodgers and P. J. Hore, “Chemical magnetoreception in birds: The radical pair mechanism”, *Proceedings of the National Academy of Sciences of the United States of America* **106**, 353–360 (2009).
- [243] C. Kerpel, S. Richert, J. G. Storey, S. Pillai, P. A. Liddell, D. Gust, S. R. Mackenzie, P. J. Hore, and C. R. Timmel, “Chemical compass behaviour at microtesla magnetic fields strengthens the radical pair hypothesis of avian magnetoreception”, *Nature Communications* 2019 10:1 **10**, 1–7 (2019).

- [244] A. Javed, R. Jabbour, S. V. Sadasivan, S. Alsaghir, A. Alhussni, M. Jhamnani, and A. Equbal, “Dynamic nuclear polarization: State of the art and future possibilities with light activation”, *Chemical Physics Reviews* **6** (2025).
- [245] T. Scharff, W. Ratzke, J. Zipfel, P. Klemm, S. Bange, and J. M. Lupton, “Complete polarization of electronic spins in OLEDs”, *Nature Communications* 2021 12:1 **12**, 1–8 (2021).
- [246] R. Haberkorn, “Density matrix description of spin-selective radical pair reactions”, *Molecular Physics* **32**, 1491–1493 (1976).
- [247] B. Adams, I. Sinayskiy, and F. Petruccione, “An open quantum system approach to the radical pair mechanism”, *Scientific Reports* 2018 8:1 **8**, 1–10 (2018).
- [248] P. H. Alvarez, F. T. Chowdhury, L. D. Smith, T. J. Brokowski, C. D. Aiello, D. R. Kattnig, and M. C. de Oliveira, “Simulating spin biology using a digital quantum computer: Prospects on a near-term quantum hardware emulator”, *APL Quantum* **1** (2024).
- [249] Y. Zhang, Z. Hu, Y. Wang, and S. Kais, “Quantum Simulation of the Radical Pair Dynamics of the Avian Compass”, *Journal of Physical Chemistry Letters* **14**, 832–837 (2023).