Randomized Quantum Simulation Of Nuclear Systems On Near-Term Quantum Computers

Kiss, Oriel

UNIVERSITÉ DE GENÈVE FACULTÉ DES SCIENCES

Département de physique nucléaire et corpusculaire Professeur Federico Sanchez Nieto

Randomized Quantum Simulation Of Nuclear Systems On Near-Term Quantum Computers

THÈSE

Présentée à la Faculté des sciences de l'Université de Genève
Pour obtenir le grade de Docteur ès sciences, mention interdisciplinaire

Par

Oriel KISS

de

Genève (GE)

Thèse N°5917

GENÈVE

2025

June 2025

Thèse N°5917

UNIVERSITÉ DE GENÈVE

FACULTÉ DES SCIENCES

DOCTORAT ÈS SCIENCES, MENTION INTERDISCIPLINAIRE

Thèse de Monsieur Oriel Orphée Moïra KISS

intitulée :

«Randomized Quantum Simulation Of Nuclear Systems On Near-Term Quantum Computers»

La Faculté des sciences, sur le préavis de

Monsieur F. SANCHEZ NIETO, professeur ordinaire et directeur de thèse
Département de physique nucléaire et corpusculaire

Madame S. VALLECORSA, docteure et codirectrice de thèse
Conseil européen pour la recherche nucléair (CERN), Genève, Suisse

Monsieur N. BRUNNER, professeur ordinaire
Département de physique appliquée

Monsieur G. MAZZOLA, docteur
Department of Astrophysics, University of Zurich, Zurich, Suisse

autorise l'impression de la présente thèse, sans exprimer d'opinion sur les propositions qui y sont énoncées.

Genève, le 20 juin 2025

Thèse - 5917 -

La Doyenne

N.B. - La thèse doit porter la déclaration précédente et remplir les conditions énumérées dans les "Informations relatives aux thèses de doctorat à l'Université de Genève".

"The arXiv is a wonderful tool, you go there and anyone can write there. So you know, YOU KNOW you get the best quality information."

— Micheal Scott

ABSTRACT

This thesis explores the challenges and opportunities in quantum computing, focusing on mitigating current limitations and advancing quantum simulation techniques. After introducing the fundamentals of quantum simulation, superconducting circuits, and error mitigation we turn to applications of randomized quantum algorithms, emphasizing on dynamics. Key contributions include: improving initial state preparation using the Variational Quantum Eigensolver, extending the capabilities of the qDrift random product formula, and proposing cost-reduction algorithms for quantum simulations. Applications of these methods include constructing stochastic channels for computing the flavor variance of forward scattering neutrinos and observing the Kibble-Zurek mechanism when driving across a quantum phase transition.

The thesis culminates with a framework based on Fourier moments, which facilitates arbitrary Hamiltonian transformations with applications in response functions and ground state energy estimation. We focus on the practical side of the algorithm and discuss how to solve the related challenges using techniques from signal processing, leading to a 10 fold reduction in the number of samples. Combining this framework with DMRG, we demonstrate accurate ground state energy estimation for systems of interacting neutrinos.

While practical quantum advantage remains out of reach, this work lays the foundation for improved algorithms and error mitigation techniques, contributing to the progress needed to unlock the full potential of quantum computing.

RESUMÉ

Cette thèse explore les défis et opportunités de l'informatique quantique, en se concentrant sur l'atténuation des limitations actuelles et l'avancement des techniques de simulation quantique. Après avoir introduit les bases de la simulation quantique et des circuits supraconducteurs, et les protocoles de mitigation d'erreur, nous nous intéressons aux applications des algorithmes quantiques aléatoires, en mettant l'accent sur la dynamique. Les contributions principales incluent : l'amélioration de la préparation d'états initiaux à l'aide de l'algorithme variationnel pour les valeurs propres, l'extension des capacité du protocol aléatoire qDrift, et la proposition d'algorithmes visant à réduire son coût. Ces méthodes ont été appliquées à la construction de canaux stochastiques pour calculer la variance de la diffusion de neutrinos et à l'étude des transitions de phase via le mécanisme de Kibble-Zurek.

La thèse s'achève avec les applications des moments de Fourier, qui permettent des transformations arbitraires de Hamiltoniens. Cela permet en autre de calculer les fonctions de réponse et à estimer des énergies d'état fondamental. Nous mettons l'accent sur les aspects pratiques de l'algorithme et discutons des défis associés, que nous résolvons à l'aide de techniques de traitement de signaux, aboutissant à une réduction d'un facteur de 10 dans le nombre de mesure requis. En combinant ce cadre avec des états initiaux préparés par DMRG, nous calculons l'énergies d'état fondamental de système de neutrinos.

Bien qu'un avantage quantique reste encore hors de portée, cette thèse améliore certain algorithmes ainsi que la mitigation d'erreur, contribuant ainsi aux progrès nécessaires pour exploiter pleinement le potentiel de l'informatique quantique.

PUBLICATIONS

The present thesis is based on following publications.

[1] Oriel Kiss, Michele Grossi, Pavel Lougovski, Federico Sanchez, Sofia Vallecorsa, and Thomas Papenbrock. "Quantum computing of the ^{6}Li nucleus via ordered unitary coupled clusters." In: *Phys. Rev. C* 106 (3 Sept. 2022), p. 034325. DOI: 10.1103/PhysRevC.106.034325. URL: https://link.aps.org/doi/10.1103/PhysRevC.106.034325.

[2] Michele Grossi, Oriel Kiss, Francesco De Luca, Carlo Zollo, Ian Gremese, and Antonio Mandarino. "Finite-size criticality in fully connected spin models on superconducting quantum hardware." In: *Phys. Rev. E* 107 (2 Feb. 2023), p. 024113. DOI: 10.1103/PhysRevE.107.024113. URL: https://link.aps.org/doi/10.1103/PhysRevE.107.024113.

[3] Oriel Kiss, Michele Grossi, and Alessandro Roggero. "Importance sampling for stochastic quantum simulations." In: *Quantum* 7 (Apr. 2023), p. 977. ISSN: 2521-327X. DOI: 10.22331/q-2023-04-13-977. URL: https://doi.org/10.22331/q-2023-04-13-977.

[4] Oriel Kiss, Utkarsh Azad, Borja Requena, Alessandro Roggero, David Wakeham, and Juan Miguel Arrazola. "Early Fault-Tolerant Quantum Algorithms in Practice: Application to Ground-State Energy Estimation." In: *Quantum* 9 (Apr. 2025), p. 1682. ISSN: 2521-327X. DOI: 10.22331/q-2025-04-01-1682. URL: https://doi.org/10.22331/q-2025-04-01-1682.

[5] Oriel Kiss, Michele Grossi, and Alessandro Roggero. "Quantum error mitigation for Fourier moment computation." In: *Phys. Rev. D* 111 (3 Feb. 2025), p. 034504. DOI: 10.1103/PhysRevD.111.034504. URL: https://link.aps.org/doi/10.1103/PhysRevD.111.034504.

[6] Oriel Kiss, Daniil Teplitskiy, Michele Grossi, and Antonio Mandarino. "Statistics of topological defects across a phase transition in a superconducting quantum processor." In: *Quantum Sci. Technol.* (June 2025), p. 035037. DOI: 10.1088/2058-9565/addf75.

ADDITIONAL PUBLICATIONS

Additional published content and contributions which are not covered in the present thesis.

[1] T. Ramazyan, Oriel Kiss, M. Grossi, E. Kajomovitz, and S. Vallecorsa. "Generating muonic force carriers events with classical and quantum neural networks." In: *J. Phys.: Conf. Ser. 2438 012089*. 2023. DOI: 10.1088/1742-6596/2438/1/012089.

[2] Oriel Kiss, Francesco Tacchino, Sofia Vallecorsa, and Ivano Tavernelli. "Quantum neural networks force field generation." In: *Mach. Learn.: Sci. Technol.* 3 (2022), p. 035004. DOI: https://doi.org/10.1088/2632-2153/ac7d3c.

[3] Oriel Kiss, Michele Grossi, Enrique Kajomovitz, and Sofia Vallecorsa. "Conditional Born machine for Monte Carlo event generation." In: *Phys. Rev. A* 106 (2 Aug. 2022), p. 022612. DOI: 10.1103/PhysRevA.106.022612. URL: https://link.aps.org/doi/10.1103/PhysRevA.106.022612.

[4] Manuel S. Rudolph, Sacha Lerch, Supanut Thanasilp, Oriel Kiss, Oxana Shaya, Sofia Vallecorsa, Michele Grossi, and Zoë Holmes. "Trainability barriers and opportunities in quantum generative modeling." In: *npj Quantum Information* 10 (Nov. 2024), p. 116. DOI: 10.1038/s41534-024-00902-0.

[5] Paulin de Schoulepnikoff, Oriel Kiss, Sofia Vallecorsa, Giuseppe Carleo, and Michele Grossi. "Hybrid ground-state quantum algorithms based on neural Schrödinger forging." In: *Phys. Rev. Res.* 6 (2 Apr. 2024), p. 023021. DOI: 10.1103/PhysRevResearch.6.023021. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.6.023021.

[6] Isabel Nha Minh Le, Oriel Kiss, Julian Schuhmacher, Ivano Tavernelli, and Francesco Tacchino. "Symmetry-invariant quantum machine learning force fields." In: *New Journal of Physics* 27.2 (Feb. 2025), p. 023015. DOI: 10.1088/1367-2630/adad0c. URL: https://dx.doi.org/10.1088/1367-2630/adad0c.

[7] Oriel Kiss, Saverio Monaco, Antonio Mandarino, Sofia Vallecorsa, and Michele Grossi. "Quantum phase detection generalization from marginal quantum neural network models." In: *Phys. Rev. B* 107 (8 Feb. 2023), p. L081105. DOI: 10.1103/PhysRevB.107.L081105. URL: https://link.aps.org/doi/10.1103/PhysRevB.107.L081105.

ACKNOWLEDGMENTS

Writing a thesis in the modern age is inherently a collaborative endeavor, and this work is no exception. I am deeply grateful to my advisors, Federico Sanchez and Sofia Vallecorsa, for their unwavering support, and to my adoptive advisor, Alessandro Roggero, for his boundless enthusiasm and inspiring ideas.

Throughout my Ph.D. journey, I had the privilege of being part of several research teams, each contributing to my growth as a researcher in unique ways. I extend my heartfelt thanks to the IBM Quantum Zürich team, particularly Francesco Tacchino and Ivano Tavernelli, for their mentorship and guidance that led to my first publication. Warm thanks also go to Thomas Papenbrock who introduced me to nuclear physics.

I am also immensely grateful to the CERN QTI team - Michele, Massimiliano, Alice, Matteo, Su Yeon, Carla and Vasilis — for the stimulating discussions and countless coffee breaks that enriched my experience.

My appreciation also goes to the Xanadu quantum algorithms team, who welcomed me during a summer and introduced me to a more industrial approach to research. Special thanks to Juan Miguel Arrazola for always challenging me, as well as Utkarsh, Stepan, and Borja for their encouragement, and Frederik and Danial for teaching me table tennis.

I am also grateful to the members of my jury, Guglielmo Mazzola and Nicolas Brunner, for their valuable time and insightful feedback.

Finally, none of this would have been possible without the constant support of my family: my parents, my sister and my partner. Mélanie, to you, with all my love and gratitude—thank you for being by my side through it all.

CONTENTS

LIST OF TABLES

LISTINGS

ACRONYMS

ACDF Approximate Cumulative Distribution Function

ASP Adiabatic state preparation

CC Coupled Cluster

CRG Control Reversal Gate

CDF Cumulative Distribution Function

CS Conpressed Sensing

DD Dynamical Decoupling

DMRG Density Matrix Renormalization Group

ED Exact Diagonalization

EFT Effective Field Theory

ETH Eigenstate Thermalization Hypothesis
EV Echo Verification
FTQC Fault-tolerant Quantum Computing
GSEE Ground-state energy estimation
HF Hartree Fock
JW Jordan-Wigner
KZM Kibble-Zurek Mechanism
LCU Linear Combinations of Unitaries
LL Layerweise Learning
LMG Lipkin-Meshkov-Gick
LT Lin and Tong
MMM Matrix-free Measurement Mitigation
MPS Matrix Product States
NISQ Near-Intermediate Noisy Quantum
NR Noise Renormalization
PF Product Formula
PEV Purified Echo Verification
PQC Parameterized Quantum Circuit
QBE Qubit-based excitation
QEC Quantum Error Correction
QFT Quantum Fourier Transform
QPE Quantum phase estimation
QPT Quantum Phase Transition
QSP Quantum Signal Processing
QSVT Quantum Singular Value Transform
RC Randomized Compiling
SP Symmetry Protection
SPSA Simultaneous Perturbation Stochastic Approximation
SV Symmetry Verification
TFQIM Transverse-Field Quantum Ising Model
TREX Twirled readout error extinction
UCC Unitary Coupled Cluster
VD Virtual Distillation
VQA Variational Quantum Algorithms
VQD Variational Quantum Deflation
VQE Variational Quantum Eigensolver
ZNE Zero-Noise Extrapolation

FREQUENTLY USED SYMBOLS

ϵ	approximation error
n	system size - number of qubits
$\|\psi\rangle$	quantum state
$\|\bar{0}\rangle$	$\|0\rangle^{\otimes n}$
ρ	density matrix
$\mathcal{H}$	Hamiltonian
$\mathcal{O}$	observable
$\mathcal{S}$	symmetry of the Hamiltonian
$\mathcal{O}(\cdot)$	big "O" notation
t	time
$\mathcal{U}(t)$	time evolution operator ($\exp(-i\mathcal{H}t)$)
$\|E_i\rangle$	i-th eigenstate of $\mathcal{H}$ with energy E_i
$\mathbb{1}$	identity matrix
X, Y, Z	Pauli matrices
σ_a, $a \in \{0,1,2,3\}$	Pauli matrices
$\mathcal{N}[\cdot]$	noise channel

1 INTRODUCTION

Digital quantum computing has emerged as a revolutionary framework in science, technology and engineering, using the principles of quantum mechanics to perform computations. The theoretical foundations of quantum computing emerged in the early 1980s, with the famous quote by Feynman [1] "nature is quantum mechanical", and therefore requires a quantum computers to be simulated efficiently. In fact, simulating quantum systems is a difficult task for classical computers. The primary obstacle lies in the exponential increase in time and memory resources necessary to simulate quantum systems, which is in turn rooted in the representation of many-body quantum systems using tensor product of Hilbert spaces. Classical modeling approaches, such as quantum Monte Carlo [2] or tensor networks [3], do have merit when applied on certain type of quantum systems, but typically fail in general, e.g. when strong quantum correlations are present. This lead to the foundation for a novel computing paradigm that would use the characteristics of quantum systems, such as superposition, entanglement and interference. The idea here is to use a controllable quantum system -the quantum computer- to represent other quantum systems of interest.

As the field of quantum computing matures, Near-Intermediate Noisy Quantum (NISQ) devices are now accessible via the cloud through numerous companies, universities, and startups. However, it remains unclear what practical applications can be achieved, either with the current technology, but also in the foreseeable future. In fact, the only known, arguably useful, task with an exponential separation is large-number factoring [4]. While the construction of quantum computers represents a significant scientific milestone in itself, their widespread development can only be justified if they prove to be practically useful. Here, we define "useful" as the ability to solve a problem of scientific or industrial relevance more effectively than any classical computer—whether that be through higher precision, reduced energy consumption, or faster computation. Quantum advantage can be demonstrated in two primary ways: through asymptotic analysis using complexity theory or through heuristic comparisons with the best-known classical algorithms.

Algorithms designed for fault-tolerant quantum architectures, such as Quantum Signal Processing (QSP) [5], are often proven to be asymptotically optimal for a variety of tasks. However, the depth required to execute these algorithms means they are unlikely to be practical in the near-term, as they typically rely on Quantum Error Correction (QEC), which necessitates logical qubits composed of many physical qubits. It remains uncertain when fully fault-tolerant devices will become operational. In contrast, NISQ devices, while significantly limited in capability, are available today. There is optimism that mild improvements in qubit quality, error mitigation techniques, and quantum algorithms could lead to some form of quantum advantage in the near future.

In this thesis, we focus on quantum simulation algorithms tailored for near-term architectures, with particular attention to identifying promising directions to achieving quantum utility. Variational Quantum Algorithms (VQA) [6] have emerged as strong candidates for near-term applications. In VQAs, quantum circuits are optimized to minimize a loss function, which quantifies the success of the algorithm in solving a given problem. This framework has gained considerable popularity in recent years due to its flexibility: variational circuits can be designed with shallow depths, making them suitable for

NISQ devices. They can be employed to prepare quantum wavefunctions or to address tasks such as classification, regression, and generative modeling. While promising, this approach has hidden challenges, particularly the classical-quantum feedback loop required to update the parameters of the quantum circuit. This process has a significant sample complexity, as it involves the costly operations of loading classical data into the quantum conmputer and reading quantum data from it. These overheads are intrinsic to the hybrid quantum-classical paradigm. Hence, it is much more optimal to design algorithms where the interaction between classical and quantum resources is minimal, which is not the case for VQA. Moreover, recent studies have uncovered fundamental issues with quantum machine learning algorithms, particularly the problem of vanishing gradients, also known as barren plateaus [7]. Unlike classical neural networks, quantum variational circuits lack an efficient equivalent of the backpropagation algorithm [8], which is central to the success of modern deep learning models. These limitations significantly reduce the practical appeal of quantum variational circuits, casting doubt on their role in the current research landscape. It remains uncertain whether they can deliver the advantages initially hoped for.

At the outset of my Ph.D. research, the limitations of VQAs were not fully understood, which is why my initial proposal was centered around them. Hence, VQAs hold the advantage of being simple to implement and among the few algorithms compatible with NISQ devices. However, as the challenges of high sample complexity and poor trainability became clearer, I shifted my focus towards the more promising class of quantum simulation algorithms. Being intrinsically quantum in nature, it is more apparent how quantum computers can solve these problem efficiently. Moreover, quantum simulation algorithms do not suffer from the costly classical-quantum feedback loop or trainability issues associated with VQAs.

In quantum simulation, we are interested to in computing specific properties of quantum systems. The workflow involves, prepare a known quantum state, evolving it over time, and estimate physical observables. There are two general approaches to quantum simulation: near-term and long-term methods. Long-term quantum simulation algorithms for closed, time-independent dynamics are considered a solved problem, with QSP providing optimal error scaling that is both additive and linear in time. However, QSP comes with significant overhead that renders it impractical in the near term. In contrast, PF techniques, while offering less impressive asymptotic guarantees, have much lower overhead and are surprisingly efficient in practice. In this dissertation, we will focus on both deterministic and random PF methods, since they are effective and and have significantly less overhead.

While simulating quantum dynamics is invaluable for understanding quantum many-body systems, we will see that it can be use as an ingredient in more sophisticated algorithms. For example, we will demonstrate that the ability to simulate quantum dynamics is a crucial stepping stone towards performing arbitrary Hamiltonian transformations, which hold considerable promise for a variety of applications, including Ground-state energy estimation (GSEE) and response function calculations. We will use an approach similar to QSP but achieved in an incoherent, NISQ-friendly, manner. We argue that performing Hamiltonian transformations in this way has significant potential for near-term quantum computing.

The outline of the thesis is as follows:

1. Chapter 2 introduces the fundamentals of digital quantum computing, along with the underlying physics of superconducting quantum circuits as qubit architectures.

2. Chapter 3 reviews several error suppression and mitigation strategies essential for conducting experiments on NISQ devices. Although this chapter does not contain

original contributions, it serves as a critical review and provides the necessary context for understanding the results presented later in the thesis.

3. Chapter 4 presents the preparation of variational ground states for various nuclear systems, which are utilized throughout the thesis. Additionally, this chapter includes original contribution to improve the optimization of unitary coupled cluster ansätze, as well as results obtained using IBM's quantum hardware.

4. Chapter 5 explores deterministic Trotter-Suzuki PF methods and random qDrift channels, alongside original contributions that generalize qDrift for arbitrary sampling distributions, which can lead to cost reductions. Detailed proofs are provided in Appendix A.

5. Chapter 6 covers two quantum simulation experiments: neutrino thermalization and the observation of the Kibble-Zurek Mechanism (KZM). For the neutrino experiment, original contributions include investigating the scaling behavior of characteristic times for specific observables, as well as introducing a cheap random channel as a proxy for simulating large-scale experiments. In the second experiment, we employ various error mitigation techniques to observe the scaling of cumulants predicted by the KZM across a quantum phase transition.

6. Chapter 7 introduces Fourier moments and discusses their efficient computation on a quantum processor. These moments form the foundation for performing QSP incoherently, a promising application for near-term quantum computers. The chapter then delves into two applications: the computation of response functions using tailored error mitigation techniques and an exploration of the practical challenges associated with GSEE, along with potential solutions.

7. The thesis concludes with a summary of the key findings in Chapter 8, where future perspectives and open research directions are also discussed.

Part I

NEAR-TERM QUANTUM COMPUTERS

In recent years, rapid advancements in quantum hardware, combined with innovative error mitigation techniques, have brought us closer to achieving practical applications on near-term quantum computers. Hardware improvements, such as higher qubit counts, better coherence times, and lower gate error rates, are enabling increasingly complex quantum computations. At the same time, error mitigation strategies are allowing us to deal with noise without full error correction, paving the way for utility-scale applications even in the early fault-tolerant era.

2

DIGITAL QUANTUM COMPUTING

In recent decades, the notion of quantum simulators [9, 10] has emerged as a potent instrument for scientific computing. A quantum simulator is a machine aimed at replicating the dynamic behavior of a quantum system in a closed environment, and takes its motivation in the premise that computing is inherently a physical process. This resulted into two types of quantum simulators: analog and digital. Analog quantum simulators concentrate on physically emulating the behavior of specifics models by implementing them on a controllable quantum platform. This is inherently closer to a physical experiment, as the simulator is built to replicate nature itself. Conversely, digital quantum simulators are versatile, programmable instruments proficient in simulating arbitrary unitary operations, according to the DiVincenzo criterion [11] for universal quantum computing. Digital quantum simulators and inherently more powerful than their analog counter part, as they can run arbitrary quantum algorithms, such as the famous Shor's factoring algorithm [4, 12] with controllable error. However, they are more complex to build and computations performed on them typically possess a larger overhead.

Significant progress has been achieved in both algorithmic innovation and experimental implementation in the recent years. The contemporary realm of digital quantum computers is a rich environment comprised of universities, industries and startups. Many different platforms like trapped ions, superconducting circuits, neutral atom and photonics, among others, are currently explored, as it is still unclear which one will predominate. The field is driven by achieving "quantum advantage", which is the point where a quantum computer can realise a useful task faster or better than any classical computers. The level of maturity required to reach this milestone varies between applications. While factoring typically requires millions of physical qubits [13], only a few hundreds of high-quality noisy qubits could be enough to challenge the state of the arts classical methods [14].

2.1 GATE-BASED QUANTUM COMPUTING

A typical digital quantum simulation experiment is comprised of three main ingredients, as displayed in Figure 2.1. The first one is a description of the physical model, which is often given as an Hamiltonian H, i.e., the energy function of the system, as well as an initial state $|\psi\rangle_0$ and a set of physical observable. The Hamiltonian contains the relevant physical quantities of the systems, such as the frequencies or coupling strength, as well as the locality. It is often expressed using fermionic/bosonic creation and annihilation operation or discretised quantum fields on a lattice. The second ingredient is a formal mapping from the physical to the qubit Hamiltonian. Qubits are two-levels quantum systems, and therefore naturally possess bosonic degree of freedoms. They can be elegantly described using the Pauli algebra defined by the following set of $SU(2)$ matrices

$$\sigma_x \equiv X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_y \equiv Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_z \equiv Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \tag{2.1}$$

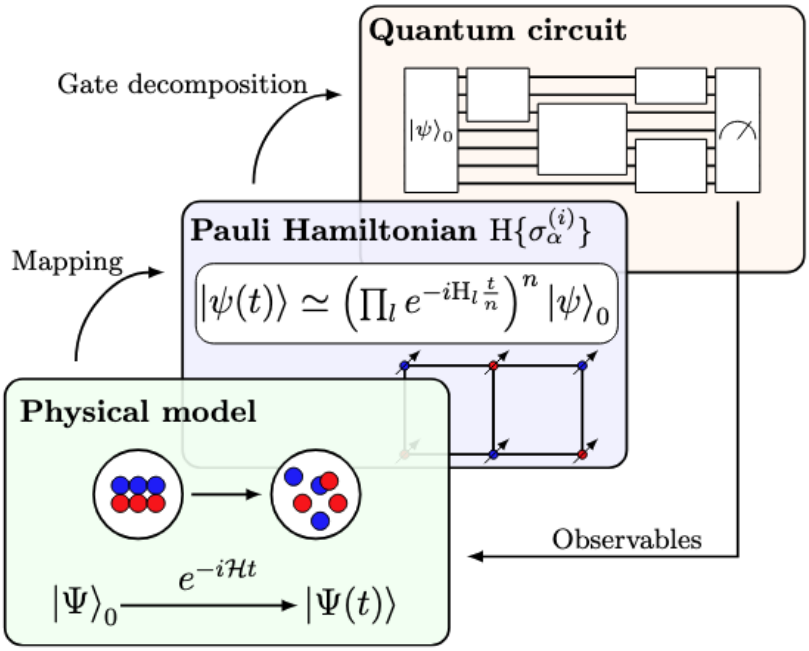

Figure 2.1: **Conceptual illustration of a digital quantum simulator.** A typical quantum simulation task is composed of three ingredients: a physical model to study (green), a mapping to a qubit Hamiltonian (blue) and a gate-based quantum circuit implementing the time evolution (orange). and finally a set of observable which maps the quantum computation back to the physical model. Figure taken from Ref. [15].

also known as Pauli matrices, satisfying the following commutation relations

$$[\sigma_\alpha, \sigma_\beta] = 2i\epsilon_{\alpha\beta\gamma}\sigma_\gamma, \quad \{\sigma_\alpha, \sigma_\beta\} = 2\delta_{\alpha\beta}\mathbb{1}. \tag{2.2}$$

In the above, $\alpha, \beta, \gamma \in \{x, y, z\}$, $\epsilon_{\alpha\beta\gamma}$ denotes the Levi-Civita tensor, and $\delta_{\alpha\beta}$ represents the Kronecker delta. The Pauli algebra is fundamental to the physical characterization of spin-1/2 quantum systems, hence the derived Hamiltonian H often represents a model of interacting spin-1/2 operators. Fermionic systems are more challenging, as they require a non-trivial mapping, such as the Jordan-Wigner (JW) [16] or the Bravyi-Kitaev [17] transformation. The initial state, on the other hand, represents the systems' configuration at the beginning, and is often chosen as a state which is easy to prepare, such as a product state. Finally, the set of observable are quantities of interest, such as the magnetization or entropy, which can help us better understand the underlying physics.

Once the mapping from the original Hamiltonian to a qubit Hamiltonian is performed, it remains to transpile the time-evolution operator $\mathcal{U}(t)$ in a gate-based quantum circuits, e.g. using PF (see Chapter 5) and collect statistics by measuring the state, which can be used to infer the expectation values of interest.

Each hardware architecture possesses a native set of operations that are executed based on the device's physical attributes. Universal quantum computation and simulation can only be achieved if its native set of operations includes a universal set, in the sens that it arbitrary unitary matrices can be built from it [11]. Numerous universal sets of single- and two-qubit gates have been identified, and are all fundamentally equivalent. Processors using various technology platforms may demonstrate unique qubit-qubit interconnectivity and other technical constraints, such as restrictions on gate directionality. These characteristics typically do not impose significant constraints on the platform's computing capacity, since

they may be mitigated using mechanisms such as SWAP operations. Nonetheless, they may provide additional overhead in the overall duration of the simulations, potentially leading to substantial degradation in the outcomes of certain calculations owing to the inherent limitations of NISQ devices concerning noise, cross-talk, and gate errors. Consequently, certain NISQ platforms are more adept at simulating specific physical models (e.g., trapped ions, which possess inherent all-to-all connectivity, can more readily simulate long-range interactions), despite the theoretical capability of any comprehensive digital quantum simulator to replicate arbitrary quantum dynamics.

In the following, we will review the most important quantum gates used for quantum computing and simulation applications [15, 18, 19]. All existing and prospective quantum computing systems enable the manipulation of individual qubits using customized control pulses to execute single qubit gates. The most general single qubit $SU(2)$ operation takes the following form

$$U(\theta,\phi,\lambda) = \begin{pmatrix} \cos(\theta/2) & -e^{i\lambda}\sin(\theta/2) \\ e^{i\phi}\sin(\theta/2) & e^{i(\lambda+\phi)}\cos(\theta/2) \end{pmatrix}, \tag{2.3}$$

and can be obtained, for example, by combining well known single qubit quantum gates such as the Hadamard H[1] and the phase gate $\Phi(\delta)$:

$$H = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \; \Phi(\delta) = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\delta} \end{pmatrix}, \tag{2.4}$$

since we have

$$U(\theta,\phi,\lambda) = e^{-i\theta/2}\Phi\left(\frac{\pi}{2}+\phi\right)H\Phi(\theta)H\Phi\left(-\frac{\pi}{2}+\lambda\right). \tag{2.5}$$

Rotations around the coordinate axes are described by

$$R_\alpha(\theta) = \exp\left(-i\frac{\theta}{2}\sigma_\alpha\right), \quad \alpha = x,y,z. \tag{2.6}$$

These rotations, up to global phase factors, can be realized by selecting specific parameters in $U(\theta,\phi,\lambda)$. For instance, $R_z(\lambda)$ can be expressed as $R_z(\lambda) = e^{-i\lambda/2}\Phi(\lambda) = e^{-i\lambda/2}U(0,0,\lambda)$, while $R_x(\theta)$ is given by $R_x(\theta) = U(\theta,-\pi/2,\pi/2)$, and $R_y(\theta) = U(\theta,0,0)$.

Conversely, any platform capable of performing single-qubit rotations around the coordinate axes can, in principle, implement any arbitrary unitary operation $U(\theta,\phi,\lambda)$ using the Euler decomposition:

$$U(\theta,\phi,\lambda) = R_z(\phi)R_x(\theta)R_z(\lambda) \tag{2.7}$$

Moreover, it is noteworthy that the following relation holds

$$U(\theta,\phi,\lambda) = R_z\left(\phi-\frac{\pi}{2}\right)R_x\left(\frac{\pi}{2}\right)R_z\left(\pi-\theta\right)R_x\left(\frac{\pi}{2}\right)R_z\left(\lambda-\frac{\pi}{2}\right), \tag{2.8}$$

which enables a highly efficient experimental approach where arbitrary single-qubit operations can be implemented using only fixed-phase $R_x(\pi/2)$ gates. This method requires minimal pulse calibrations, while virtual R_z gates—essentially reference frame adjustments in the control software—are error-free [20]. Alternative formal decompositions of generic

1 This notation is ambivalent as 'H' is also used to denote an Hamiltonian. To avoid overloading the notation, we will make sure from context which one is implied.

$SU(2)$ transformations, along with approximate findings using a limited collection of fixed-phase single-qubit operations in lieu of continuous-valued operations, are also documented in Ref. [21].

All universal sets of quantum gates require a minimum of one entangling two-qubit operation. The manner in which this need is fulfilled in actual quantum computers is significantly influenced by the underlying architecture and may differ even across implementations that use the same core technology. Superconducting qubits linked through cross-resonance interactions [22–25] allows universal computation using the native gate set

$$S = \{R_\alpha(\theta), \text{CNOT}\} \tag{2.9}$$

comprising arbitrary single-qubit rotations and the two-qubit CNOT gate

$$\text{CNOT} = \quad = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \tag{2.10}$$

which flip the state of the target qubit, upon the control being in the $|1\rangle$ state. Others two-qubit gates include the XY spin-spin coupling, such as the $R_{xy}(\delta) = e^{-i\delta\sigma_x\otimes\sigma_y}$ and XZ $R_{xz}(\delta) = e^{-i\delta\sigma_x\otimes\sigma_z}$ which are well adapted for superconducting qubits [26, 27]. Both can be represented with two CNOT gates and single qubit rotation as follows. First, we remark that the interaction can be mapped in the z basis via the basis transformations

$$H\sigma_x H = \sigma_z = \Phi\left(\frac{\pi}{2}\right) H\sigma_y H \Phi\left(-\frac{\pi}{2}\right). \tag{2.11}$$

It is therefore sufficient to implement the ZZ interaction, which takes the following form up to a global phase

$$e^{-i\delta\sigma_z\otimes\sigma_z} = R_{zz}(2\delta) = \text{CNOT}\cdot(\mathbb{1}\otimes R_z(2\delta))\cdot\text{CNOT}, \tag{2.12}$$

which can be verified by an explicit conversion in matrix form. Finally, the Heisenberg interaction $XX + YY + ZZ$, which will be useful later on, can be optimally decomposed [28] as

$$e^{-i\alpha\sigma_x\otimes\sigma_x} e^{-i\beta\sigma_y\otimes\sigma_y} e^{-i\gamma\sigma_z\otimes\sigma_z} = \quad = \quad . \tag{2.13}$$

This interaction can be decomposed using only three CNOT gates, which is half of what we would have naively obtained by multiplying the individual interactions using (2.12). This emphasizes the fact that circuit compilation is an important task which can lead to significant improvement.

2.2 INFERRING EXPECTATION VALUES

Reading out information from a quantum state is a crucial step of any quantum algorithm. Hence, quantum states living on a quantum computers are only useful if we can extract classical information, such as expectation values of some meaningful observables. In this way, insights into the underlying physics can be obtained. If the observable is hermitian, expectation values can be easily computed via direct sampling. However, we are sometimes interested in observables $\mathcal{U}$ which are non-hermitian, and only unitary. A typical example is the Loschmidt echo

$$\langle\Psi|e^{-iHt}|\Psi\rangle, \tag{2.14}$$

which is central in the remaining of this thesis. A standard technique in such cases is the Hadamard test [29], which instead encode the expectation value in the amplitudes of a quantum state. We start with the physical state $|\Psi\rangle$ and an ancilla in the state $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. By applying $\mathcal{U}$ on the physical state controlled on the ancilla we obtain

$$\frac{1}{\sqrt{2}}\big(|0\rangle \otimes |\Psi\rangle + |1\rangle \otimes \mathcal{U}|\Psi\rangle\big). \tag{2.15}$$

We proceed by a applying a Hadamard gate on the ancilla, leading to

$$\frac{1}{2}\big(|0\rangle \otimes (\mathbb{1}+\mathcal{U})|\Psi\rangle + |1\rangle \otimes (\mathbb{1}-\mathcal{U})|\Psi\rangle\big). \tag{2.16}$$

The probability of measuring the state $|0\rangle$ is given by

$$P(0) = \frac{1}{4}\langle\Psi|(\mathbb{1}+\mathcal{U}^\dagger)(\mathbb{1}+\mathcal{U})|\Psi\rangle = \frac{1}{2} + \langle\Psi|\mathcal{U}|\Psi\rangle + \langle\Psi|\mathcal{U}^\dagger|\Psi\rangle, \tag{2.17}$$

while

$$P(1) = \frac{1}{4}\langle\Psi|(\mathbb{1}-\mathcal{U}^\dagger)(\mathbb{1}-\mathcal{U})|\Psi\rangle = \frac{1}{2} - \langle\Psi|\mathcal{U}|\Psi\rangle - \langle\Psi|\mathcal{U}^\dagger|\Psi\rangle. \tag{2.18}$$

Taking the difference of the two results in

$$\langle Z\rangle_a = P(0) - P(1) = \frac{1}{2}\langle\Psi|(\mathcal{U}+\mathcal{U}^\dagger)|\Psi\rangle = \Re\langle\Psi|\mathcal{U}|\Psi\rangle. \tag{2.19}$$

The imaginary part can be computed in a similar fashion by staring the ancilla in $\frac{1}{\sqrt{2}}(|0\rangle - i|1\rangle)$. The number of samples requires to estimate the expectation value with accuracy ϵ is of order $\mathcal{O}(1/\epsilon^2)$, which can be improved to $\mathcal{O}(1/\epsilon)$ with amplitude amplification [30].

2.3 QUANTUM PHASE ESTIMATION

The QPE algorithm [31] aims for the estimation of the eigenvalue of a unitary operator U given an eigenstate $U|u\rangle = e^{2\pi i\phi}|u\rangle$. The goal of the phase estimation algorithm is to approximate this unknown phase ϕ with a chosen level of accuracy and probability of success. In the following, we will review this central algorithm, basing ourself on the classic book by Nielsen and Chuang [32].

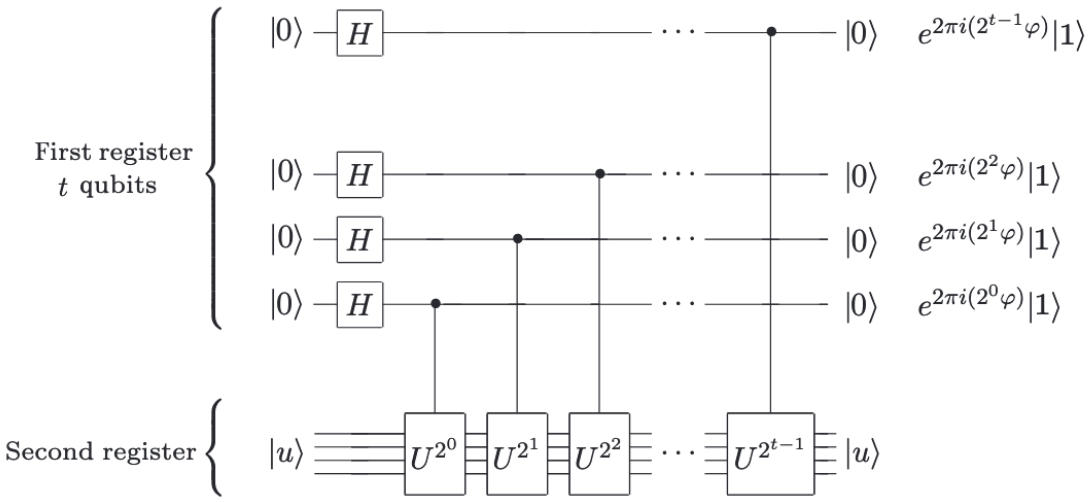

Figure 2.2: **Illustration of the QPE algorithm.** Figure taken from Ref. [32].

2.3.1 *The algorithm*

The QPE procedure assumes access to two essential components:

1. A *state preparation oracle*, capable of preparing the eigenstate $|u\rangle$.
2. A *controlled-unitary operation oracle* capable of performing the controlled-U^{2^j} operation for various non-negative integers j.

These oracles are often represented as black boxes, indicating that QPE is a modular subroutine rather than a stand-alone quantum algorithm. In practical applications, this modularity allows QPE to be combined with other quantum procedures to solve complex problems, where the specific implementation of these black-box operations is adapted for the task at hand.

The QPE algorithm uses two quantum registers. The first register contains t qubits, initialized in the state $|0\rangle$, where t depends on both the desired number of accurate bits for the estimate of ϕ and the probability of success. The second register contains the state $|u\rangle$, the eigenvector of the unitary operator U. We then proceed as follows, which is also illustrated in the circuit in Figure 2.2 from Ref. [32].

1. Prepare the first register in a superposition using Hadamard gates, and apply successive controlled-U^{2^j} operations.
2. Apply the inverse Quantum Fourier Transform (QFT) on the first register.
3. Measure the first register to obtain an estimate $\tilde{\varphi}$ of φ.

We will now dive deeper into the details of each steps. The first step is to apply the QFT

$$|j\rangle \mapsto \frac{1}{\sqrt{n}} \sum_{k=0}^{n-1} e^{2\pi i jk/n} |k\rangle \tag{2.20}$$

to the ancilla register. When the input is the all zero state, this is equivalent to apply the Hadamard transform, creating an uniform superposition of every states. We then proceed

by applying controlled-U^{2^j} operations to the second register. After these operations, the first register's state is given by

$$\frac{1}{2^{t/2}}\left(|0\rangle + e^{2\pi i 2^{t-1}\phi}|1\rangle\right)\left(|0\rangle + e^{2\pi i 2^{t-2}\phi}|1\rangle\right)\cdots\left(|0\rangle + e^{2\pi i 2^{0}\phi}|1\rangle\right) = \frac{1}{2^{t/2}}\sum_{k=0}^{2^t-1} e^{2\pi i\phi k}|k\rangle, \tag{2.21}$$

where we have omited the second register as it remains in the state $|u\rangle$.

We then proceed by applying the inverse QFT to the auxiliary register, enabling the extraction of the binary digits of ϕ. This transformation rearranges the amplitudes in such a way that, if ϕ can be exactly expressed in t bits (i.e., $\phi = \sum_j \phi_j 2^{-j} \equiv 0.\phi_1\phi_2\ldots\phi_t$), the resulting state is given by

$$\frac{1}{2^{t/2}}\left(|0\rangle + e^{2\pi i 0.\phi_t}|1\rangle\right)\left(|0\rangle + e^{2\pi i 0.\phi_{t-1}\phi_t}|1\rangle\right)\cdots\left(|0\rangle + e^{2\pi i 0.\phi_1\phi_2\cdots\phi_t}|1\rangle\right). \tag{2.22}$$

The final step involves measuring the first register in the computational basis, yielding an estimate of ϕ. When t is chosen appropriately, the probability of success for obtaining an accurate estimate of ϕ depends on the number of bits and the desired error tolerance.

An essential feature of the QPE algorithm is the ability of the inverse QFT to transform the superposition state into a measurement-ready state that approximates ϕ with high accuracy. This operation is captured by

$$\frac{1}{2^{t/2}}\sum_{j=0}^{2^t-1} e^{2\pi i\phi j}|j\rangle\otimes|u\rangle \rightarrow |\tilde{\phi}\rangle\otimes|u\rangle, \tag{2.23}$$

where $|\tilde{\phi}\rangle$ denotes a state that serves as a good approximation for ϕ upon measurement. This transformation is the crux of the QPE algorithm.

2.3.2 *Performance and analysis*

The above analysis applies to the ideal case where ϕ can be exactly expressed with a t-bit binary expansion. While it is rarely the case in real situations, iot turns out that QPE can still produce a close approximation to ϕ with high probability.

To better understand why this is the case, let b be the integer in the range 0 to $2^t - 1$ such that $b/2^t = 0.b_1b_2\ldots b_t$ is the best t-bit approximation to ϕ, satisfying $0 \leq \delta = \phi - b/2^t \leq 2^{-t}$. Here, δ denotes the difference between ϕ and $b/2^t$.

Applying the inverse QFT to the state (2.21) results in

$$\frac{1}{2^t}\sum_{k,l=0}^{2^t-1} e^{-2\pi ikl/2^t}e^{2\pi i\phi k}|l\rangle. \tag{2.24}$$

Define α_l as the amplitude of $|(b+l) \bmod 2^t\rangle$,

$$\alpha_l = \frac{1}{2^t}\sum_{k=0}^{2^t-1}\left(e^{2\pi i(\phi-(b+l)/2^t)}\right)^k, \tag{2.25}$$

which is the sum of a geometric series and simplifies to

$$\alpha_l = \frac{1}{2^t}\left(\frac{1-e^{2\pi i(2^t\delta-l)}}{1-e^{2\pi i(\delta-l/2^t)}}\right). \tag{2.26}$$

We suppose now that the final measurement yields the value m. We wish to bound the probability of obtaining m such that $|m-b| > e$, where e is an integer characterizing the error tolerance. The probability of measuring such a m is given by

$$p(|m-b|>e) = \sum_{-2^{t-1}<l\leq -(e+1)} |\alpha_l|^2 + \sum_{e+1\leq l\leq 2^{t-1}} |\alpha_l|^2. \tag{2.27}$$

We know from elementary analysis that for any real θ, it holds that $|1-\exp(i\theta)| \leq 2$. We can therefore bound $|\alpha_l|$ as

$$|\alpha_l| \leq \frac{2}{2^t |1-e^{2\pi i(\delta - l/2^t)}|}. \tag{2.28}$$

Since $|1-\exp(i\theta)| \geq 2|\theta|/\pi$ for $-\pi \leq \theta \leq \pi$, and given that $-2^{t-1} < l \leq 2^{t-1}$, it follows that $-\pi \leq 2\pi(\delta - l/2^t) \leq \pi$, and thus

$$|\alpha_l| \leq \frac{1}{2^{t+1}(\delta - l/2^t)}. \tag{2.29}$$

Combining now (2.27) and (2.29) yields for the error probability

$$p(|m-b|>e) \leq \frac{1}{4}\left[\sum_{l=-2^{t-1}+1}^{-(e+1)} \frac{1}{(l-2^t\delta)^2} + \sum_{l=e+1}^{2^{t-1}} \frac{1}{(l-2^t\delta)^2}\right]. \tag{2.30}$$

Recalling that $0 \leq 2^t\delta \leq 1$, we obtain the follwoing upper bound

$$p(|m-b|>e) \leq \frac{1}{4}\left[\sum_{l=-2^{t-1}+1}^{-(e+1)} \frac{1}{l^2} + \sum_{l=e+1}^{2^{t-1}} \frac{1}{(l-1)^2}\right] \leq \frac{1}{2}\sum_{l=e}^{2^{t-1}-1} \frac{1}{l^2}, \tag{2.31}$$

which simplifies to using an integral approximation

$$p(|m-b|>e) \leq \frac{1}{2}\int_{e-1}^{2^{t-1}-1} \frac{dl}{l^2} = \frac{1}{2(e-1)}. \tag{2.32}$$

If we aim to approximate ϕ to an accuracy of 2^{-n}, it is thus sufficient to choose $e = 2^{t-n}-1$. With $t = n+p$ qubits, we find from (2.32) that the probability of obtaining an approximation with accuracy 2^{-n} is at least $1-1/(2(2^p-2))$. In order to achieve an approximation with n bits of accuracy and success probability $1-\epsilon$, we can choose

$$t = n + \left\lceil \log\left(2 + \frac{1}{2\epsilon}\right)\right\rceil. \tag{2.33}$$

2.3.3 *The initial state is not an eigenstate*

In the above analysis, we assumed that $|u\rangle$ is an eigenstate of the unitary operator U. However, it practice we rarely have access to an exact eigenstate, but only to an approximation. In the following, we consider an initial state $|\psi\rangle$, which can be written as a superposition of U's eigenstates. Specifically, let

$$|\psi\rangle = \sum_u c_u |u\rangle, \tag{2.34}$$

where $|u\rangle$ are eigenstates of U with eigenvalues $e^{2\pi i\phi_u}$, and c_u are the corresponding coefficients of the superposition, satisfying $\sum_u |c_u|^2 = 1$.

When we perform phase estimation with the initial state $|\psi\rangle$, the evolution can be analyzed by decomposing it into the contributions from each eigenstate component $|u\rangle$. In general, we instead have an initial state $|\psi\rangle$ expressed as a superposition of U's eigenstates

$$|\psi\rangle = \sum_u c_u |u\rangle, \tag{2.35}$$

where $|u\rangle$ are eigenstates of U with eigenvalues $e^{2\pi i \phi_u}$, and c_u are the coefficients satisfying $\sum_u |c_u|^2 = 1$. It results that the probability of measuring ϕ_u with $n-$ bit precision is $|c_u|^2(1-\epsilon)$. Therefore, it is essential to have a good approximation of the ground-state, i.e., with $|c_0|$ as close to 1 as possible. Preparing good initial states is an active and important are of research.

2.3.4 *Recent improvements*

Since the original proposal of the QPE algorithm by Kitaev [31], numerous improvement have been performed. Without entering into too much details, we will nevertheless give some pointers.

The first adaptation, based on qubitization [33] is to instead consider the walk operator W from a block-encoding of the Hamiltonian, which yields eigenvalues $e^{\pm i\phi_j}$, with $\phi_j = \arccos(\lambda_j/\lambda)$, λ_j the eigenvalues and $\lambda = \sum_j |\lambda_j|$ the one-norm of the Hamiltonian, which controls the cost of block-encoding techniques. This has several advantages such as

1. Using the block-encoding avoids the error when approximating the time-evolution.
2. The cosine is a smooth and monotone function, which facilities the retrieval of the eigenvalue from the bare phase.

An another improvement is the use of better window function to estimate the phase [34]. While the standard QPE uses the uniform window obtained with the Hadamard transform, one may design optimal filters, e.g. based on the Kaiser window [35]. Using optimal filters yields the advantage of decreasing the tails of the probability distribution of observing the correct phases, which significantly decrease the number of times QPE need to be run.

2.4 OVERVIEW OF SUPERCONDUCTING CIRCUITS

After establishing the theoretical foundations of digital quantum simulations, we will dedicate the remaining of this chapter to a succinct assessment of the principal experimental advancements in superconducting qubits. Although it is challenging to stay abreast of the latest research, we will nonetheless highlight certain essential works that are universally acknowledged as foundational to the field. Numerous alternative architectures for quantum computers, such as trapped ions [36–41], ultracold atoms [42–44], spins in silicon [45–47], quantum dots [48–51], and polarized photons, have been proposed and still investigated today. While each of them have their individual merit, it remains difficult at this stage to assess which one will prevail. We focus on superconducting qubits, mainly because it is the technology I had access to when performing the experiments.

Superconducting circuits are a fast pacing and scalable technology [52–60]. These devices utilize millimeter-scale circuits constructed using lithographic methods derived from traditional integrated circuits. The manufacturing employs a multi-step additive and subtractive patterning technique for thin metallic layers, enabling significant flexibility in circuit design. Another benefit of superconducting circuits is that they use operations

inside the microwave spectrum, where the majority of the existing technology has been created. A disadvantage is that they must be operated at millikelvin temperatures, and therefore necessitate costly dilution refrigerators. In the following, we aim to give a brief overview of the technology, which is treated thoroughly in the following reviews [61–64].

2.4.1 *The transmon*

In superconducitng circuits, qubits are encoded into the low energy spectrum part of collective charge excitation in a non-linear LC resonator. The most successful qubits based on this technology is the transmon [65–71], which consists of two superconducting pads connected by a Josephson junction. From an electrical engineering perspective, it can be modeled as a planar capacitor connected to a nonlinear inductor, and acts as a Cooper pair box in a specific regime that suppresses charge noise.

The Josephson junction operates as a nonlinear inductive component, and is the only nonlinear lossless electrical component at low temperatures. It serves as the most critical element in superconducting circuits since it leads to non-uniform energy spacing, which can be used to control state transitions. It consists of two superconducting pads separated by a thin insulating barrier. The electrical behavior in the barrier is described by the Josephson equations [72]

$$I(t) = I_c \sin\phi(t) \text{ and } \frac{d\phi(t)}{dt} = \frac{2\pi}{\Phi_0} V(t), \tag{2.36}$$

where I_c is the critical current, Φ_0 the superconducting flux quantum, $\phi(t)$ the superconducting phase difference and $V(t)$ the voltage across the junction. We can understand that this acts as a nonlinear (in the phase difference) inductor by rewriting them as

$$\frac{dI}{dt} = \frac{2\pi I_c \cos\phi}{\Phi_0} V \equiv \frac{1}{L_J} V. \tag{2.37}$$

We will start by describing the Cooper pair box [73], which is a superconducting qubit constructed from a superconducting island linked to a pool of superconducting electrons by a Josephson junction. The number of Cooper pairs may be regulated by connecting a DC voltage source using a gate capacitor C_g. If the voltage is high enough, a Cooper pair will tunnel from the reservoir to the island, which will discharge the qubit capacitor. The Hamiltonian of a Cooper pair box connected to the DC voltage can be written as

$$H = 4E_C(\hat{n} - n_g)^2 - E_J \cos\hat{\phi}, \tag{2.38}$$

with $E_J = \Phi_0 I_c/(2\pi)$ the Josesphon energy, n_g being the gate charge, and $E_C = e^2/2C$ the charging energy. The first three energy levels are shown in Figure 2.3 for two different regime: $E_J = E_C$ and $E_J > E_C$. We observe that the energy levels are not equally spaced, on the contrary of a standard quantum harmonic oscillator. Therefore, the first two states can be used as qubits. The sensitivity to charging noise may be mitigated when the Cooper pair box is operated at the sweet point $n_g = 0.5$, when linear effects dissipate. Alternatively, on may augment the ratio E_J/E_C, as seen in the various insets of Figure 2.3, which exponentially diminishes the charge dispersion relation, and hence enhances the qubit dephasing time T_2 [74]. The transmon operates as a Cooper pair box with a ratio of $E_J/E_C \approx 100$. Importantly, the anharmonicity reduction follows a weak power law, so we retain good control on the transmon states transitions.

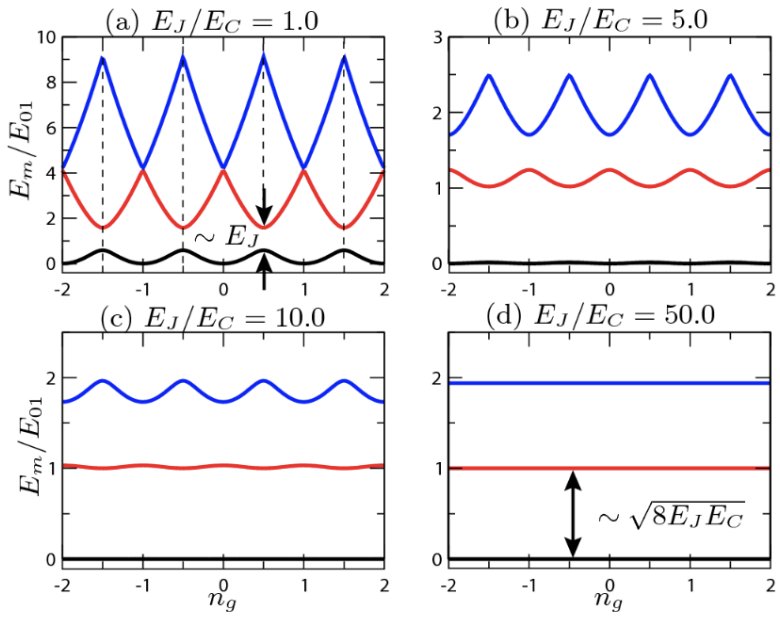

Figure 2.3: **Energy levels of a transmon.** The first three energy levels of a CPB are shown as a function of the gate charge. The different figures refer to different regime for E_J/E_C. We see that the oscillation of the energy levels are reduced when this ratio is increase, which can protect the qubits from charging noise. Figure taken from Ref. [74].

2.4.2 *Driving the transmon*

In this section, we will review the fundamentals required to implements quantum gates at the physical level, as well as performing readout measurements, which can be described in the context of circuit quantum electrodynamics [75].

A qubit can be placed into an arbitrary superposition state by applying an electrical pulse

$$V_d(t) = A\epsilon(t)\sin(\omega_d t + \alpha) \tag{2.39}$$

with carefully controlled amplitude A, envelope $\epsilon(t)$, duration t, and phase α. The driving frequency ω_d is often selected to match the qubit transition frequency, ensuring resonance between the two. The pulse is produced by an external AC power source situated outside the dilution refrigerator. The signal propagates across a coaxial wire and arrives to an on-chip waveguide, which is capacitively connected to the qubit pads. This connection is commonly referred to as the control line or XY line.

To calibrate the amplitude and duration of the electrical pulse required for a specific single-qubit rotation, a time Rabi experiment is typically performed. In this experiment, a qubit is subjected to electrical pulses of varying durations, followed by a measurement. The Rabi frequency can be extracted by fitting the resulting oscillations. In practice, pulses are optimized for rotation with fixed axis and angles, and are then combined to approximate arbitrary rotations.

Two-qubit gates, crucial for producing entanglement, are a basic element of quantum information processing. In superconducting qubit systems, gates are often executed by the application of radio-frequency or flux pulses that facilitate interactions between two qubits. The engagement must be brief and minimal during qubit idling. These methodologies may be generally categorized into two types: flux-tunable gates and all-microwave gates.

Flux-tunable gates depend on the application of direct current or radio-frequency flux pulses via on-chip flux bias lines. The flux pulses modify the frequency of one or both qubits, aligning them in resonance, thereby enabling either a swap interaction or a controlled-phase operation. This method's benefit lies in its very short gate time. This method adds complexity since it necessitates a flux bias line for each qubit and may expose the system to flux noise, which may impair coherence. The iSWAP gate serves as an example of a flux-tunable gate.

All-microwave two-qubit gates are executed without requiring flux bias lines, hence decreasing the number of input lines and streamlining the system design. This technique depends only on the use of microwave pulses to regulate qubit interactions. The absence of a need for flux-tunability in qubits makes them less vulnerable to flux noise, hence extending their dephasing periods. Nonetheless, the gate timings for all-microwave gates are often slower. The cross-resonance gate exemplifies a notable all-microwave two-qubit gate.

2.4.3 *Measuring the transmon*

In circuit quantum electrodynamics, the state of a qubit is measured via its dispersive interaction with a far-detuned microwave resonator. This interaction induces a state-dependent frequency shift in the resonator, which allows one to infer the qubit state indirectly by measuring the response of the resonator to an applied microwave signal. By accurately measuring this phase shift, we can determine whether the qubit is in the ground or excited state. To minimize the impact on qubit coherence (i.e., to reduce dephasing), the microwave signal is typically set to populate the resonator with only a few photons. The signal, after interacting with the resonator, is amplified by cryogenic amplifiers at low temperatures, followed by further amplification at room temperature, and finally processed using heterodyne detection techniques.

2.4.4 *Summary*

In this section, we provided an overview of superconducting circuits for quantum information processing. Superconducting circuits have emerged as a very promising platform for quantum computing, using macroscopic quantum phenomena at cryogenic temperatures to generate highly controlled qubits. Circuits using Josephson junctions have significant benefits: rapid gate timings, scalability, and alignment with contemporary semiconductor manufacturing methods. The most prominent example of superconducting qubits is the transmon, which operates as a Cooper pair box in a certain regime. It enjoys high tolerance to charge noise while maintening good anharmonicity. It his such as promising component of a scalable architecture.

Single-qubit operations are realized by acting with fine-tuned microwave pulses, while two-qubit gates are performed via resonant interactions. Readout is achieved by coupling the qubit to a microwave resonator, and reading the frequency shifts.

The primary problems for superconducting quantum circuits is scaling the device to hundreds or thousands of qubits while preserving high fidelity. Principal study domains are minimizing crosstalk among qubits, enhancing qubit coherence durations, and formulating more effective quantum error correction methodologies. Improvements in cryogenic infrastructure and control electronics are essential for the development of large-scale quantum computers.

The sustained success of superconducting qubit technology will likely hinge on surmounting hardware constraints via hybrid systems, such as the integration of superconducting qubits with other quantum platforms, or by the innovation of innovative materials that minimize loss and decoherence.

3

ERROR MITIGATION

Quantum computers are not closed systems and as such, are subject to noise due to the interaction with the environment. While their classical counterparts also experience noise, e.g. due to interaction with the astrophysical background, this effect is much stronger here. Error correcting codes are employed in classical computers to resolve this problem by incorporating redundancy into the processing of the information. In fact, the error rates can be significantly decreased by employing repetition code, which involves duplicating each bits and correcting possible errors by taking a majority vote after each operation. However, in a quantum computer, the situation becomes more difficult because of the no-cloning theorem [76]. Hence, quantum information cannot be copied, which would for example allow faster than light communication. However quantum information can be protected by being encoded into a larger space. This lead to so-called logical qubits, which are built from many physical qubits and encodes the same information as in one qubit but in an redundant protective way. This protocol is rigorously defined through stabilizer codes, in which logical qubits are encoded in the $+1$ eigenspace of an Abelian subgroup of the Pauli-group. In this way, errors can be detected in real time by measuring each syndromes, and ultimately corrected. Since the logical state is a $+1$ eigenstate, measuring the syndromes does not affect the wavefunction by making it collapses.

QEC codes protect quantum information by encoding logical qubits into entangled states of multiple physical qubits, enabling the detection and correction of errors without disturbing the underlying quantum information. The most widely studied quantum error correction code is the surface code [77–80], which encodes a logical qubit using a 2D grid of physical qubits and can correct both bit-flip and phase-flip errors. Surface codes are particularly suited for superconducting qubit platforms due to their local connectivity requirements and tolerance for low error thresholds. Recent experiments have demonstrated key milestones, including the detection and correction of errors, logical qubit operations, and logical error rates lower than physical qubit error rates, which is critical for achieving fault tolerance. We refer to Ref. [81] for a more thorough introduction on the surface code.

Quantum error correction, however, comes at the cost of substantial overhead in terms of the number of physical qubits required to encode a single logical qubit. The challenge remains in reducing this overhead while maintaining high-fidelity operations across the qubit array. As the field progresses, quantum error correction will evolve from proof-of-concept demonstrations [82–85] into robust, scalable protocols, pushing quantum computers closer to solving problems that are intractable for classical devices. The interplay between hardware development and optimized error correction codes will be key in achieving the long-term vision of fault-tolerant quantum computing [86]. Innovations in hardware—such as improving qubit coherence times, reducing gate errors, and developing noise-resilient architectures—will complement advances in error correction techniques.

Meanwhile, one can employ quantum error mitigation techniques that seek to minimize the impact of noise instead of completely eliminating it. Error correction operates by identifying and rectifying errors in real-time, whereas error mitigation employs a greater quantity of measurements [87] and addresses the impact of noise through data post-processing. Where QEC increases the width and depth of the quantum computation, error

mitigation uses a higher number of measurements, as well as noise-tailored techniques, to diminishes the effect of noise on physical observable. Error mitigation is therefore a purely heuristic approach, as we do not have a precise model of the noise affecting quantum computers. Nevertheless, it is primordial when performing experiments on near-term quantum hardware.

Error mitigation techniques can be categorized into three main groups: error suppression, device-based error mitigation, and state-based error mitigation. Error suppression primarily involves the use of transpilation algorithms to enhance the coherence time of the qubits. They act directly at the circuit level, can be used in synergy with more involved quantum error mitigation strategies and typically do not increase the sample complexity. Device-based error mitigation often involves acquiring attributes from the device and utilizing them to rectify errors during the post-processing phase. The state-based protocols are agnostic to the specific device being used and directly act on the quantum state. We shall now present the various error mitigation approaches that will be utilized throughout this thesis.

3.1 ERROR SUPPRESSION

3.1.1 *Dynamical decoupling*

Dynamical Decoupling (DD), as proposed by Viola and Lloyd [88], aims to enhance the coherence time of quantum devices by utilizing brief, time-dependent control modulation during periods of qubit inactivity. Hence, qubits that are idling tend to decohere more rapidly, a problem that can be alleviated by implementing a series of gates that acts as the identity. As an illustration, we can utilize XX, YY, XY-2, or XY-8 gate sequences [88–90], as benchmarked in Ref. [91], and illustrated in Figure 3.1. Optimizing the pulse sequences, such as the delay between pulses and their relative phases, can be useful for achieving optimal performance with the computation devices, using e.g. genetic algorithms [92]. Based on our empirical observations, XY-8 pulse sequences have demonstrated the highest level of error suppression.

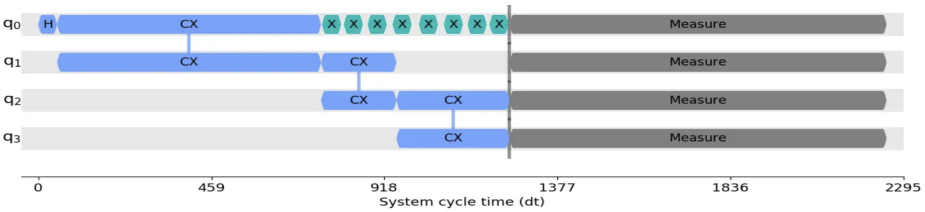

Figure 3.1: Dynamical decoupling pluses (XX-8 sequence) are inserted when the qubits are idling to increase the coherence time. Figure taken from the qiskit textbook.

3.1.2 *Randomized compiling*

Randomized Compiling (RC) is a process that seeks to convert coherent noise into stochastic noise by averaging over random equivalent circuits [93]. The introduction of stochasticity can significantly decrease the level of errors arising due to the interaction with the environment. Intuitively, this process turns the error accumulation into a random walk, preventing the errors to accumulate coherently and thus proving a quadratic improvement. In practice, two-qubit gates are conjugate with single Pauli operations in such a way as

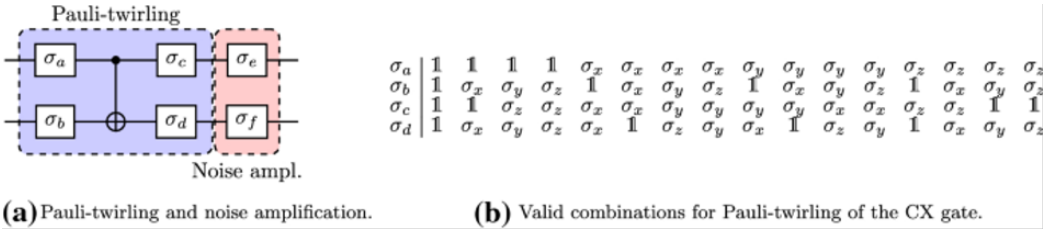

(a) Pauli-twirling and noise amplification. **(b)** Valid combinations for Pauli-twirling of the CX gate.

Figure 3.2: **Illustration of Pauli twirling.** [a] Conjugation of a CNOT gate. [b] Example of possible Pauli twirls. Figure taken from Ref. [95].

not affecting the overall action of the original gate. They exist many possible choices for twirling gates, which have been referenced by [94] in the case of Pauli twirling and is illustrated in Figure 3.2. This process effectively turns the noisy channel into a Pauli channel

$$\mathcal{E}(\rho) = \sum_{P \in \mathcal{P}^{\otimes n}} c_P P \rho P^\dagger, \tag{3.1}$$

where $\mathcal{P} = \{\mathbb{1}, X, Y, Z\}$ is the set of Pauli matrices and c_P the relative error due to P. While running overall all possible choices can be expensive; in practice (3.1) is evaluated via Monte-Carlo sampling. This does not increase the sample complexity, as the shot budget can be distributed among all the twirls .

There are numerous important advantages of deliberately introducing stochasticity in the compilation process [93], such as:

(i) The error channel's off-diagonal components resulting from coherent errors are suppressed, resulting in a reduced error rate.

(ii) The growth of stochastic Pauli error is linear with the circuit depth, while coherent noise accumulates quadratically (in the small error limit). The noise accumulation under RC exhibits the behavior of a random walk, resulting in a quadratic error reduction.

(iii) RC is expected to work in symbiosis with depolarising noise mitigation technique, since it turns the noise model into a Pauli channel, which for the two-qubit case, comprises depolarising channel. This is beneficial as those noise channels tends to be easier to correct.

3.1.3 *Twirled readout error extinction*

The Twirled readout error extinction (TREX) method [96] is a technique designed to mitigate measurement errors by introducing uniform random Pauli bit flips before measurement. By adding randomness, the method helps reduce biases that arise from hardware imperfections during readout. The key idea is that these Pauli bit flips randomize the output channel, making the noise less structured and easier to correct. This employs the same ideas as twirling, but is instead used for readout.

In practice, the random bit flips are tracked, and undone in post-processing. For instance, if a X Pauli gate is applied before the measurement, one may correct it by flipping the measured classical bit. Importantly, these Pauli transformations convert the effect of arbitrary noise into a multiplicative factor that applies to each Pauli observable. This

transformation simplifies the noise model by diagonalizing the measurement channel, making bias correction easier and more efficient.

This method is general enough to be used on a variety of quantum systems, not just single-qubit error models, and requires no detailed model of the readout errors. The only assumption needed is the ability to initialize the circuit in the zero state. The effectiveness of the twirled-readout method depends on the magnitude of the noise factor, with the shot-noise variance of the error-mitigated estimates inversely proportional to this noise strength. By analyzing the underlying noise, one can determine the number of measurement samples required to achieve a desired level of accuracy.

3.1.4 *Pulse-efficient transpilation*

Pulse-efficient transpilation is a technique aimed at optimizing quantum circuits by aligning them more closely with the native gate set of a quantum device [97, 98]. Quantum circuits are often written using a standard gate set, such as CNOT or $\exp(-i\theta Z \otimes Z)$ gates, but the native two-qubit gate on certain IBM's quantum devices is the cross-resonance gate $\exp(-i\theta X \otimes Z)$.

To make circuits more hardware-efficient, transpilation involves rewriting operations, such as $\exp(-i\theta Z \otimes Z)$ gates, in terms of the native cross-resonance gates combined with appropriate qubit rotations. This reduces the need for decomposing gates into more basic native operations, which often requires additional gates and longer gate sequences. By reducing gate count and simplifying operations into shorter-duration pulses, pulse-efficient transpilation enhances the coherence time of quantum devices, thereby mitigating errors caused by decoherence and gate noise.

3.2 DEVICE-BASED ERROR MITIGATION

Device-based error mitigation protocols share the same idea of learning properties of the noise on the device to revert its effect. In the following, we will review the most important techniques, which we will use later to perform experiments.

3.2.1 *Matrix-free measurement mitigation*

An important source of errors is given by readout error, when a qubit is measured in the state $|1\rangle$ while being in the state $|0\rangle$. Such errors introduce a bias which can be mitigated by calibrating the device. The idea is to estimate the level of readout error using simple circuits, and use this information to invert the noisy readout channel. We refer to this process as Matrix-free Measurement Mitigation (MMM). The most precise strategy would be to calibrate all possible measurements and invert the obtained matrix to correct all possible errors. However, this scales exponentially with the number of qubits and is therefore not practical. Instead we follow Ref. Nation et al. [99] and individually invert the error matrices

$$M^k = \begin{pmatrix} P^{(k)}_{0,0} & P^{(k)}_{0,1} \\ P^{(k)}_{1,0} & P^{(k)}_{1,1} \end{pmatrix} \tag{3.2}$$

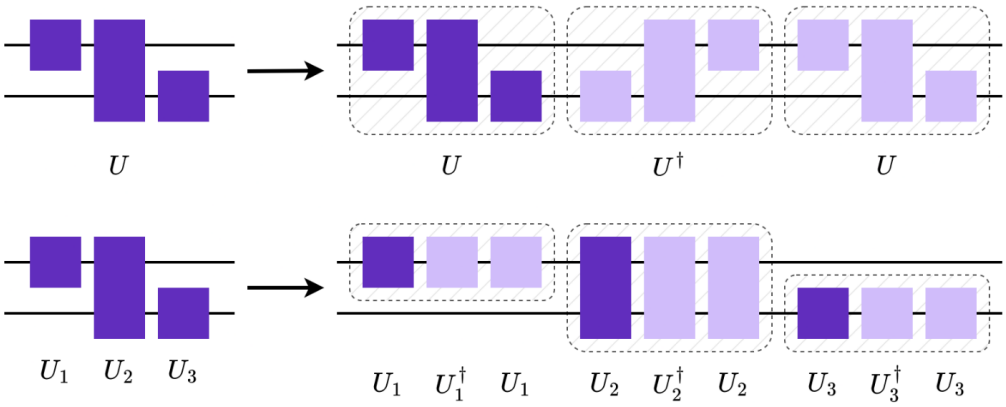

Figure 3.3: **Gate folding.** Local and Global folding for digital noise amplification. Figure from Ref. [100].

and used them to calibrate the samples. Here, $P^{(k)}_{i,j}$ is the probability of the k-th qubit to be in state $j \in \{0,1\}$ while measured in state $i \in \{0,1\}$. The probabilities of measuring the $|0\rangle$ or the $|1\rangle$ state

$$\vec{B}^k = \begin{pmatrix} P_0 \\ P_1 \end{pmatrix} \tag{3.3}$$

obtained by measuring the k-th qubit, can be corrected as

$$\vec{B}^k_{\text{corrected}} = (M^k)^{-1}\vec{B}^k. \tag{3.4}$$

Although this only corrects uncorrelated readout errors, Ref. [99], argues that they remain the most prevalent. As described, this approach only requires to run two additional circuits: the empty circuit and a n-fold Pauli X layer. Additionally, one may select a few more product states and invert the calibration matrices in the spanned subspace. This can be done efficiently if the number of states is low as the calibration matrix remains sparse. We found however that it is often not needed in practice, and thus limit ourself to the simplest case in the remaining of this work.

3.2.2 *Zero-noise extrapolation*

One popular technique is to extrapolate to the noiseless regime by instead increasing the level of noise, which is known at Zero-Noise Extrapolation (ZNE). In the ZNE [101, 102] scheme, the noise is artificially stretched in order to learn a mapping between the level of noise λ and the expectation value of the desired observable y. The noise can then be removed, to some extend, by extrapolating to the zero-noise regime.

One of the easiest way to stretch the noise by a factor of $\lambda = (1 + 2k)$, for $k \in \mathbb{N}$, is by gate folding

$$U \mapsto U(U^\dagger U)^k. \tag{3.5}$$

Since $U^\dagger U = \mathbb{1}$, this only has a non-trivial effect in the presence of noise. In practical situation, we often choose to fold only the two-qubit gates, which are the the most prone to error, by different scaling factors. The gate folding, depicted in Figure 3.3, can be either

implemented at the gate level (lower panel), or at the circuit level (upper panel). Continuous gate folding are also possible by folding a subset of random gates [103], which can be useful when folding the full circuit already exceed the capacity of the quantum hardware.

Subsequently, a fit is conducted to extrapolate the energy value for $\lambda = 0$. Linear extrapolation is perhaps the simplest method and is widely used. It corresponds to the model

$$f(\lambda) = c_0 + c_1\lambda. \tag{3.6}$$

In this case, a simple analytic solution exists, corresponding to the ordinary least squares estimator of the intercept parameter:

$$f(0) = \bar{y} - \frac{S_{\lambda y}}{S_{\lambda\lambda}}\bar{\lambda}, \tag{3.7}$$

where

$$\bar{\lambda} = \frac{1}{m}\sum_j \lambda_j, \;\; \bar{y} = \frac{1}{m}\sum_j y_j, \tag{3.8}$$

$$S_{\lambda y} = \sum_j (\lambda_j - \bar{\lambda})(y_j - \bar{y}), \;\; S_{\lambda\lambda} = \sum_j (\lambda_j - \bar{\lambda})^2. \tag{3.9}$$

Although this option is appropriate for situations when the noise is small, several papers [104–106] have demonstrated that the bias can still be significant in practice. In order to decrease it, one idea would be to employ a polynomial with a higher degree. However, this is likely to be susceptible to overfitting. It has been demonstrated that an exponential function is a more resilient and reliable option in most cases [105, 106].

THE CASE OF DEPOLARISING NOISE We can understand why an exponential fit is a good idea by considering a depolarizing noise channel with noise strength p

$$\mathcal{N}_p[\rho] = (1-p)\rho + p\frac{1}{2^n}\mathbb{1}. \tag{3.10}$$

Folding the gate by a factor λ is equivalent to $p \mapsto p^\lambda$, or

$$\mathcal{N}_{p^\lambda}[\rho] = (1-p^\lambda)\rho + p^\lambda\frac{1}{2^n}\mathbb{1}, \tag{3.11}$$

which can be derived by iteratively applying (3.10). This motivates the parametrization of the noise by a function of the form $f(\lambda) = a + be^{-c\lambda}$, whose parameters can be estimated with a least-square regression. Hence, the effect of depolarizing noise is exponential and therefore it is a natural choice for the ansatz. In geneal, the noise model might be significantly more complex than depolarizing, however it already gives us good insights. Moreover, we can use this protocol in conjunction with Pauli twirling, which has the effect of randomizing the noise model, drifting it closer to a depolarizing one.

3.2.3 *Noise renormalization*

While ZNE attempts to learn a dependence between the observable of interest and the level of noise, we could also learn the effective noise strength and use it to revert the effect of the noise. For example, the noise strength can be estimated by looking at the difference between the empirical and expected values. The NR [107–110] technique effectively estimates the

parameters of a given noise model in order to reverse the impact of the noise. For this purpose, we have to find circuits with known expectation values, such as Clifford circuit, or circuit close to the identity, which are also close to the original one.

In the following, let us assume a depolarising noise model (3.10), which maps the expectation value of a Pauli observable σ to

$$\mathrm{Tr}\left[\sigma\mathcal{N}_p(\rho)\right] = (1-p)\mathrm{Tr}\left[\sigma\rho\right]. \tag{3.12}$$

The noise can thus be corrected by dividing the noisy results by $(1-p)$. The main idea is then to estimate $(1-p)$ by running a circuit with known expectation value. For example, we can make use that the forward time evolution $\mathcal{U}(t)$ followed by a backward time evolution $\mathcal{U}(-t)$ results in the identify circuit. More specifically, we have

$$\langle\psi|\mathcal{O}|\psi\rangle = \langle\psi|\mathcal{U}(t)\mathcal{U}(-t)\mathcal{O}\mathcal{U}(t)\mathcal{U}(-t)|\psi\rangle, \tag{3.13}$$

where $|\psi\rangle$ is an initial state and $\mathcal{O}$ an arbitrary observable chosen such that $\langle\psi|\mathcal{O}|\psi\rangle = 1$, which becomes $(1-p)$ under the depolarising channel.

Therefore, we can correct an even number of Trotter steps by running an additional noise-estimating circuit, with a reference initial state

$$\langle\psi|\mathcal{U}(2t)\mathcal{O}\mathcal{U}(2t)|\psi\rangle_{NR} = \frac{\langle\psi|\mathcal{U}(2t)\mathcal{O}\mathcal{U}(2t)|\psi\rangle}{\langle\psi|\mathcal{U}(t)\mathcal{U}(-t)\mathcal{O}\mathcal{U}(t)\mathcal{U}(-t)|\psi\rangle}. \tag{3.14}$$

A second example is to used a short time evolution $\delta t \ll 1$ such that $\mathcal{U}(\delta t) \approx \mathbb{1}$. The corrected expectation values reads

$$\langle\psi|\mathcal{U}(t)\mathcal{O}\mathcal{U}(t)|\psi\rangle_{NR} = \frac{\langle\psi|\mathcal{U}(t)\mathcal{O}\mathcal{U}(t)|\psi\rangle}{\langle\psi|\mathcal{U}(\delta t)\mathcal{O}\mathcal{U}(\delta t)|\psi\rangle}. \tag{3.15}$$

Although there may be a slight bias, it is expected to fall within the range of statistical uncertainty and has the advantage of being compatible with any number of Trotter steps.

Finally, one might choose a special set of angles which makes the circuit Clifford, and can thus be simulated efficiently.

The attentive reader may question the utility of this method in a practical context, as it relies on the premise that the noise is depolarizing. Indeed, this assumption is crucial, and otherwise, NR is likely to be unsuccessful. However, we can use the same technique as for ZNE, which is to twirl the circuit, thus transforming the noise model into a something close to a depolarizing channel. This is crucial for the practical implementation of this approach.

3.3 STATE-BASED ERROR MITIGATION

All the above techniques had in common that they were trying to learn properties of the hardware noise, and use it to reverse its effect. In the following however, we consider scheme which attempt to mitigate noise at the state-level.

3.3.1 *Echo verification*

Echo Verification (EV) [110, 111] and dual state purification [112] are equivalent techniques that suppress errors by un-preparing the target quantum state and determining whether an error occurred. Using a Hadamard test, this can be accomplished by preparing the target state, deferring the expectation value on the ancilla, un-preparing the state, and projecting

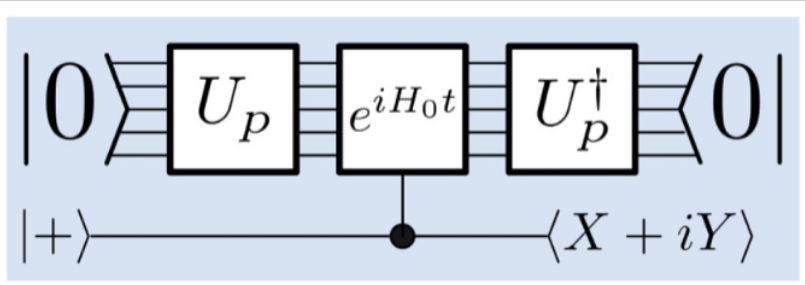

Figure 3.4: **Echo verification procedure**. The state is prepared, the expectation is then read on an ancilla using the Hadamard test. Finally the state is unprepared and post-selected on the zero state. Figure taken from Ref. [111].

it onto the initial state, as displayed in Figure 3.4. This effectively disregards erroneous trials that occur before the post-selection process.

In the subsequent section, we examine expectation values of the form $\langle\bar{0}|B^\dagger\mathcal{U}B|\bar{0}\rangle$, which can be calculated using a Hadamard test (refer to section 2.2). The operator B prepares the target state from the all-zero state $|\bar{0}\rangle = |0\rangle^{\otimes n}$. Following the controlled operation, it is important to remember that the state assumes the following form:

$$|\Phi\rangle = \frac{1}{\sqrt{2}}\left(|\bar{0}\rangle\otimes|0\rangle + B^\dagger\mathcal{U}B|\bar{0}\rangle\otimes|1\rangle\right) \equiv \frac{1}{\sqrt{2}}\left(|\bar{0}\rangle\otimes|0\rangle + |\phi\rangle\otimes|1\rangle\right). \tag{3.16}$$

We remark that $|\phi\rangle \equiv B^\dagger\mathcal{U}B|\bar{0}\rangle$ can be re cast in the basis spanned by $\{|\bar{0}\rangle, |\bar{0}^\perp\rangle\}$

$$|\phi\rangle = \alpha|\bar{0}\rangle + \beta|\bar{0}^\perp\rangle. \tag{3.17}$$

The task is thus reduced to estimating the α coefficient

$$\alpha = \langle\bar{0}|\phi\rangle = \langle\bar{0}|B^\dagger\mathcal{U}B|\bar{0}\rangle\ . \tag{3.18}$$

In the basis composed by the normalized states

$$\{|\bar{0}\rangle\otimes|0\rangle, |\bar{0}^\perp\rangle\otimes|0\rangle, |\bar{0}\rangle\otimes|1\rangle, |\bar{0}^\perp\rangle\otimes|1\rangle\}, \tag{3.19}$$

the density matrix takes the following form

$$|\Phi\rangle\langle\Phi| = \frac{1}{2}\begin{pmatrix} 1 & 0 & \alpha^* & \beta^* \\ 0 & 0 & 0 & 0 \\ \alpha & 0 & |\alpha|^2 & \alpha\beta^* \\ \beta & 0 & \alpha^*\beta & |\beta|^2 \end{pmatrix}, \tag{3.20}$$

which is reduced to the density matrix of the ancilla

$$|\Phi\rangle\langle\Phi|_0 = \mathrm{Tr}\left[|\Phi\rangle\langle\Phi|\left(|\bar{0}\rangle\langle\bar{0}|\otimes\mathbb{1}\right)\right] = \frac{1}{2}\begin{pmatrix} 1 & \alpha^* \\ \alpha & |\alpha|^2 \end{pmatrix} \tag{3.21}$$

when projected onto $|\bar{0}\rangle$. The real and imaginary part of α can then be obtained from Pauli measurements on the ancilla qubit as

$$\langle X_a\rangle_0 = \frac{2\Re\alpha}{1+|\alpha|^2}, \quad \langle Y_a\rangle_0 = \frac{2\Im\alpha}{1+|\alpha|^2}, \quad \langle Z_a\rangle_0 = \frac{1-|\alpha|^2}{1+|\alpha|^2} = \frac{2}{1+|\alpha|^2} - 1\ . \tag{3.22}$$

In the given expression, $\langle \sigma_a \rangle_0$ represents the average value of the Pauli matrix σ on the ancilla, which is selected after the physical system is in the state $|\bar{0}\rangle$. This is calculated by taking the trace of the product of σ_a and ρ_0, and dividing it by the trace of ρ_0. As previously indicated in Ref. [111], it is evident that employing solely post-selection will lead to biased results in the expected values $\langle X_a \rangle_0$ and $\langle Y_a \rangle_0$. Alternatively, the outcome can be unbiased by utilizing the expectation value $\langle Z_a \rangle_0$ as mentioned in Ref. [112].

$$\Re\alpha = \frac{\langle X_a \rangle_0}{1 + \langle Z_a \rangle_0}, \quad \Im\alpha = \frac{\langle Y_a \rangle_0}{1 + \langle Z_a \rangle_0}. \tag{3.23}$$

The main insight is that while projecting onto the state $|\bar{0}\rangle$, the information we are trying to extract, specifically the complex number α, is preserved while simultaneously eliminating some of the errors [111, 112]. Moreover, errors occurring outside of the subspace obtained after post-selection are naturally eliminated.

THE CASE OF DEPOLARISING NOISE This can be seen by considering (again) the impact of depolarizing noise. We recall that the density matrix of the state $|\Phi\rangle$ under a depolarising noise channel with parameter p becomes

$$\mathcal{N}_p[\rho] \equiv \widetilde{\rho} = (1-p)\rho + \frac{p}{2^{n+1}}\mathbb{1}\,. \tag{3.24}$$

We project onto the state $|\bar{0}\rangle$ and obtain

$$\widetilde{\rho}\,(|\bar{0}\rangle\langle\bar{0}|\otimes\mathbb{1}) = (1-p)\rho\,(|\bar{0}\rangle\langle\bar{0}|\otimes\mathbb{1}) + \frac{p}{2^{n+1}}|\bar{0}\rangle\langle\bar{0}|\otimes\mathbb{1}\,. \tag{3.25}$$

The corresponding reduced density matrix of the ancilla qubit reads

$$\widetilde{\rho}_0 = (1-p)\rho_0 + \frac{p}{2^{n+1}}\mathbb{1}\,. \tag{3.26}$$

where now the identity $\mathbb{1}$ is intended as a 2×2 matrix. The noisy expectation values takes now the following form

$$\begin{aligned}
\widetilde{\langle Z_a \rangle}_0 &= \frac{\mathrm{Tr}\,[Z_a\tilde{\rho}_0]}{\mathrm{Tr}\,[\tilde{\rho}_0]} = \frac{1-p+\frac{p}{2^n}}{\frac{1-p}{2}(1+|\alpha|^2)+\frac{p}{2^n}} - 1 \\
\widetilde{\langle X_a \rangle}_0 &= \frac{\mathrm{Tr}\,[X_a\tilde{\rho}_0]}{\mathrm{Tr}\,[\tilde{\rho}_0]} = \frac{(1-p)\Re\alpha}{\frac{1-p}{2}(1+|\alpha|^2)+\frac{p}{2^n}} \\
\widetilde{\langle Y_a \rangle}_0 &= \frac{\mathrm{Tr}\,[Y_a\tilde{\rho}_0]}{\mathrm{Tr}\,[\tilde{\rho}_0]} = \frac{(1-p)\Im\alpha}{\frac{1-p}{2}(1+|\alpha|^2)+\frac{p}{2^n}}.
\end{aligned} \tag{3.27}$$

Putting everything together, we obtain

$$\frac{\widetilde{\langle X_a \rangle}_0}{1 + \widetilde{\langle Z_a \rangle}_0} = \Re\alpha + \mathcal{O}\left(\frac{p}{2^n}\right)\,. \tag{3.28}$$

We observe that the bias has been reduced by a factor $\mathcal{O}(2^n)$, which is exponential in the system size [111].

Although our discussion focused on the resilience of EV against depolarizing noise, it is worth noting that this method is also effective in correcting several other types of noise models. Ref. [112] demonstrates that a correction can be applied to an orthogonal error channel, transforming the probability p into $p^2/(M(1-p^2))$. Here, p represents the likelihood of ending up in a subspace defined by M orthogonal states. Ref [111] presents numerical experiments on depolarizing, amplitude, and phase damping channels, while Ref. [110] consider realistic noise models from qiskit.

OVERHEAD The sole drawback of these techniques is the rise in sample complexity caused by the necessity to estimate the post-selected noisy expectation value with an error of $P_0\epsilon$, where P_0 represents the probability of measuring the $|\bar{0}\rangle$ state (refer to [111] for a more comprehensive analysis). Hence, if the circuit is too deep, errors will occurs with high probability, and few will pass the verification step. Therefore, EV can only extend the reach of NISQ devices up to a point, and is not an alternative for QEC.

3.3.2 *Purified echo verification*

In the absence of noise, the ancilla after post-selection is in a pure state. We can therefore make use of this information to improve our protocol. In fact, we can purify the ancilla by projecting onto the closest pure state, and thus obtain a result closer to the ideal outcome [112, 113]. More precisely, state tomography can be performed on the ancilla qubit, by measuring it in the X, Y, and Z basis, and used it to build the density matrix [114]. The dominant eigenvector can be used as a proxy for the closest pure state, or alternatively, the constraint systems of equations describing the closest pure state with same expectation values can be solved. This variant, already proposed in Ref. [110, 112, 113], shall be referred to as PEV. We show experimentally in Chapter 7, that the improvement procured by purification is comparable to the one from EV compared to the bare results.

3.3.3 *Symmetry verification*

We can design a similar protocol as echo verification using the symmetry of the Hamiltonian $\mathcal{H}$ governing the dynamics. Let us call $\mathcal{U}(t) = \exp(-i\mathcal{H}t)$ the time evolution under $\mathcal{H}$, and let be $\mathcal{S}$ a symmetry of the Hamiltonian, i.e., we have

$$[\mathcal{S}, \mathcal{H}] = 0. \tag{3.29}$$

Since both operator commute, we may block-diagonalise $\mathcal{H}$ within the eigenspace of $\mathcal{S}$. This means in particular that, if the initial state belongs to a specific symmetry sector, it will not leave it. In practice however, noise may shift the state outside the target subspace. We can therefore verify if the measurements belong to the correct symmetry sector, and throw the sample away if this is not the case. We refer to this protocol as Symmetry Verification (SV) [115, 116].

The projector valued measurement $\{\hat{M}_i\}$ gives us a more general framework for verifying symmetries. Given that the state $|\psi\rangle$ is invariant under a symmetry, we can measure $\{\hat{M}_i\}$ on the noisy state ρ and perform post-selection. The resulting state will thus be projected to

$$\rho_s = \frac{\hat{M}_s \rho \hat{M}_s}{\mathrm{Tr}\left[\hat{M}_s \rho\right]}. \tag{3.30}$$

We directly obtain that

$$\mathrm{Tr}\left[\rho_s |\psi\rangle\langle\psi|\right] = \frac{\mathrm{Tr}\left[\rho|\psi\rangle\langle\psi|\right]}{\mathrm{Tr}\left[\hat{M}_s\rho\right]} \geq \mathrm{Tr}\left[\rho|\psi\rangle\langle\psi|\right], \tag{3.31}$$

and we see that the overlap with the target state is amplified. The approach can be readily expanded to many operators $\{\mathcal{S}_1, \mathcal{S}_2, \ldots\}$ provided that $[\mathcal{S}_i, \mathcal{S}_j] = 0$. In the case that the different symmetries do not commute, it becomes inefficient to verify them as it increases the number of trials. It is also possible to verify symmetry in the middle of a computation, as long as we know that the state $|\psi(t)\rangle$ is an eigenstate of $\mathcal{S}$ at any time t.

SYMMETRY PROTECTION In the same manner as in EV, the overhead arises when the probability of passing verification can be prohibitively small. Nevertheless, in this context, we can reduce the likelihood of failing by using Symmetry Protection (SP) [117]. This concept involves using the evolution induced by a symmetry element $\mathcal{S}_i$ at a random time. Given that the state is an eigenstate of $\mathcal{S}$ with eigenvalue s, it only acquires a global phase

$$\mathcal{U}(t_1)\exp\left(-i\mathcal{S}\theta\right)\mathcal{U}(t_2)|\psi\rangle = \exp\left(-is\theta\right)\mathcal{U}(t_1)\mathcal{U}(t_2)|\psi\rangle, \tag{3.32}$$

leaving any expectation value unchanged.

This is useful as the randomness introduced by θ enhances the likelihood of remaining within the subspace. Hence, without SP, the errors caused by noise increase in a linear manner, since they tend to accumulate coherently. However, when SP is applied, the error accumulation follows a random walk pattern, resulting in a quadratic improvement. This is similar to the argument for randomized compiling, with the difference here is that this effectively increase the probability of remaining in the symmetric subspace.

3.3.4 *Virtual distillation*

In section 3.3.2, we have already covered the utilization of purification to enhance the efficacy of echo verification. However, this was only feasible because we were exclusively purifying the individual ancilla qubit. In order to purify the n-qubit quantum state of the physical system, it is necessary to calculate 4^n expectation values, a number that grows exponentially with the size of the system. While more efficient protocols [118] based on classical shadows [119] do exist, we will instead consider VD [120, 121] , which is a process designed to purify a quantum state ρ by doing collective measurements on M copies. To eliminate incoherent errors, one can calculate expectation values using the state $\rho^M/\mathrm{Tr}\left[\rho^M\right]$. The error-free expectation of $\mathcal{O}$ is estimated as

$$\langle\psi|\mathcal{O}|\psi\rangle_{\mathrm{VD}} = \frac{\mathrm{Tr}\left[\rho^M\mathcal{O}\right]}{\mathrm{Tr}\left[\rho^M\right]}. \tag{3.33}$$

We can see, by decomposing ρ in the eigenbasis $\rho = \sum_i p_i|E_i\rangle\langle E_i|$, that the resulting estimator converges exponentially fast in M to the closest pure state. The collective state then reads

$$\frac{\rho^M}{\mathrm{Tr}\left[\rho^M\right]} = \frac{\sum_i p_i^M|E_i\rangle\langle E_i|}{\sum_i p_i^M}, \tag{3.34}$$

and the contribution of the non-dominant eigenvectors are exponentially suppressed. To compute the numerator and denominator of equation (3.33), we can use

$$\mathrm{Tr}\left[\mathcal{O}\rho^M\right] = \mathrm{Tr}\left[\mathcal{O}_i S(M)\rho^{\otimes M}\right]. \tag{3.35}$$

In this context, $\mathcal{O}^{\mathbf{i}}$ represents the observable $\mathcal{O}$ applied to an arbitrary subsystem i, whereas $S(M)$ denotes the cyclic shift operator on a group of M systems.

$$\mathcal{O}^{\mathbf{i}} := \mathbb{1}\otimes\mathbb{1}\cdots\mathcal{O}\cdots\otimes\mathbb{1}, \tag{3.36}$$

$$S(M)|\psi_1\rangle\otimes|\psi_2\rangle\cdots\otimes|\psi_M\rangle = |\psi_2\rangle\otimes|\psi_3\rangle\cdots\otimes|\psi_1\rangle. \tag{3.37}$$

This identity can be verified by expanding the right-hand side, while keeping track of the indices. For the sake of simplicity, we select $i = 1$, resulting in

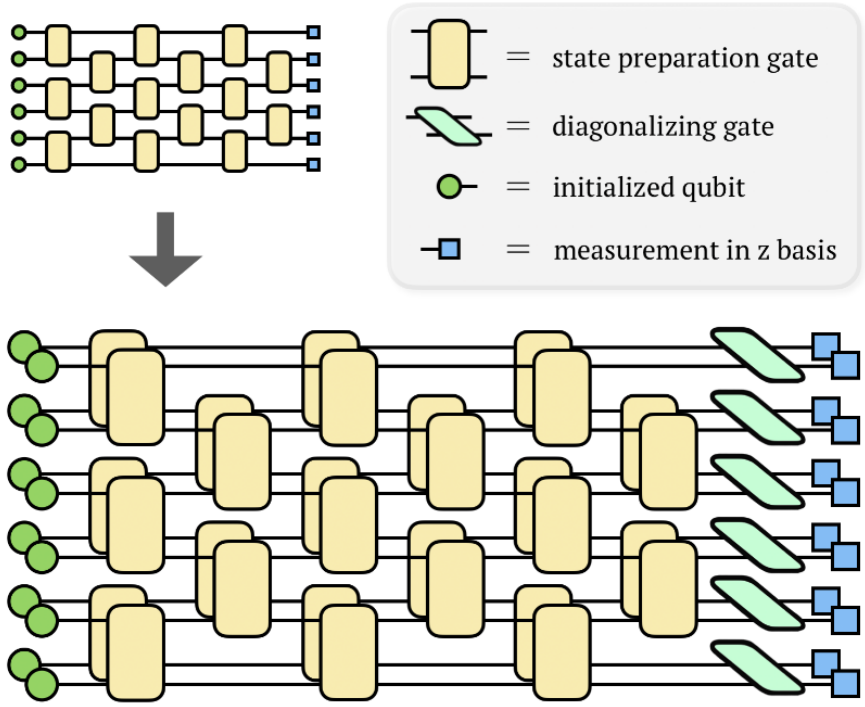

Figure 3.5: **Implementation of** VD via diagonalization measurement. Figure taken from [120].

$$\mathrm{Tr}\left[\mathcal{O}^{\mathbf{1}} S(M)\rho^{\otimes M}\right] = \sum_{\substack{i_1,i_2,\ldots,i_M \\ j_1,j_2,\ldots,j_M \\ k}} \mathcal{O}_{k,j_1}\delta_{j_2,i_1}\cdots\delta_{j_1,i_M}\rho_{i_1,k}\cdots\rho_{i_M,j_M} \tag{3.38}$$

$$= \sum_{\substack{i_1,i_2,\ldots,i_M \\ k}} \rho_{i_1,k}\mathcal{O}_{k,i_M}\rho_{i_M,i_{M-1}}\cdots\rho_{i_2,i_1} \tag{3.39}$$

$$= \mathrm{Tr}\left[\mathcal{O}\rho^M\right]. \tag{3.40}$$

The purpose of this formula is to be able to prepare the states in parallel, and not sequentially, which only leads to an increase in the number of qubits, but not in coherence time.

IMPLEMENTATION An efficient method to implement the VD protocol is by measurement by diagonalization [120]. This involves the simultaneous diagonalization of the observable and shift operator. To simplify the analysis, we will focus on a single qubit observable denoted as $\mathcal{O} = Z_1 \equiv Z$. We start by symmetrizing it by taking the average of M instances of Z denoted as $\mathcal{O}^{(M)} = \frac{1}{M}\sum_{i=1}^{M} Z^{\mathbf{i}}$. To obtain the most generic instance, one can identify a unitary basis transformation denoted as C that satisfies the equation $C\mathcal{O}C^\dagger = Z$. The symmetrized observable $\mathcal{O}^{(M)}$ possesses the advantageous property of commutating with the cyclic shift operator $[\mathcal{O}^{(M)}, S^{(M)}] = 0$. Given that both operators can be expressed as tensor products acting independently on different subsystems, it is possible to simultaneously diagonalize $S^{(M)}$ and $\mathcal{O}^{(M)}$ using an operation that has the same structure. As illustrated in Figure 3.5 for the scenario where $M = 2$, the copies are prepared in parallel, prior to the application of the diagonalizing operator and the subsequent measurement of all components in the computational basis. The numerator

and denominator of (3.33) can be calculated using the same set of measurements, resulting in the corrected expectation value of the original observable $\mathcal{O}$. This procedure can be straightforward applied to larger number of copies and more complex observables, as described more in details in Ref. [120].

3.4 DISCUSSIONS

In this chapter, we have covered many different error suppression and mitigation techniques which can be used to enhance the performance of NISQ devices. We have made a distinction between error suppression, which acts at the transpilation level, device-based, which attempts to learn the noise model to revert it, and state-based error mitigation protocols, which are concerned with using properties of the state to correct errors. This list is by no mean exhaustive. For instance, from recent advances in machine learning [122] and tensor network [123] have emerged new error mitigation protocols. Moreover, the recent quantum utility experiment by Kim, Eddins, Anand, et al. [14] used probabilistic error cancellation [124], which corrects for noise by applying inverse noise operations to cancel errors probabilistically. This is achieved by decomposing noisy operations into ideal (noiseless) operations with some error terms, sampled during the circuit execution. By collecting enough samples and appropriately weighting the measurement outcomes, the expectation value approximates the error-free result, although at the cost of increased variance and measurement complexity. This is however an expensive process which requires a close access to a quantum computer.

While quantum error mitigation has and continue to demonstrate impressive results, it is clear that it is not a replacement for QEC. Hence, its cost is usually exponential in the circuit's depth [87, 125], which is therefore not scalable. However, there is some hope that better qubit quality coupled with such error mitigation techniques may surpass state of the art classical simulation methods.

Part II

RANDOMIZED QUANTUM SIMULATION

The intrinsic randomness of quantum computing invites a natural synergy with randomization techniques, encouraging us to embrace probabilistic designs as integral parts of quantum algorithms. Since quantum computation inherently provides estimators, additional randomness can be seamlessly incorporated, whether in circuit design or parameter selection. In the following, we explore examples where randomization reduces the cost of quantum algorithms.

4

VARIATIONAL STATE PREPARATION

Initial states are important ingredients of any quantum algorithms. In numerous cases, we are interested in preparing ground states, which is a QMA-hard problem [126]. This is because quantum systems tends to be in their ground state, at least at low temperature. Moreover, since this is a QMA-hard problem, many combinatorial optimization problem can be mapped back to solving the ground state of spin glasses Hamiltonian. Many numerical methods have been proposed, with different level of success depending on the systems. For one dimensional spin systems, one successful method is obtained via DMRG [3], which minimize the energy expectation value of MPS trial states. DMRG is highly efficient to generate state with low-entanglement, which comprises ground states usually following an area law. For instance, the entanglement present in ground states of one dimensional system does not depend explicitly on its size. Neural quantum states [127], which are neural-network based, are also strong candidates and are believed to be as effective as DMRG for one- and maybe also two-dimensional systems. In chemistry or nuclear physics, popular methods also include Coupled Cluster (CC) [128] and density functional theory [129]. In this chapter, we explore methods to prepare approximate ground states using variational quantum algorithms [6]. Among them, we find the VQE, proposed by Peruzzo et al. [130], which make use of the Rayleigh-Ritz variational principle.

To be more precise, a parameterized wavefunction $|\psi(\theta)\rangle$, or Parameterized Quantum Circuit (PQC), is prepared on a quantum computer, and its parameters are iteratively updated to minimize the energy expectation value

$$\frac{\langle\psi(\theta)|H|\psi(\theta)\rangle}{\langle\psi(\theta)|\psi(\theta)\rangle}. \tag{4.1}$$

This principle tells us that this is an upper bound to the exact ground state energy E_0, and saturating this bound would gives us E_0. The parameters are typically optimized by gradient descent,

$$\theta \leftarrow \theta - \eta \frac{d}{d\theta}\langle\psi(\theta)|H|\psi(\theta)\rangle, \tag{4.2}$$

either by computing the energy's derivative [131] exactly, or approximating it via finite-difference [132, 133], until convergence is reached. The outline of the algorithm is depicted in Fig. 4.1, where the energy is first estimated on the quantum computer, given to the classical optimizer and the parameters finally updated.

Despite its numerous applications [104, 134–146] ranging from chemistry to nuclear and condensed matter physics, scaling to larger systems remains challenging due to the substantial sample requirements [145, 147–149] and trainability issues [7, 150, 151]. Hence, estimating observable via direct sampling requires a number of measurements of the order $\mathcal{O}(1/\epsilon^2)$, where ϵ is the error tolerance. Since the Hamiltonian typically consists of multiple non-commuting Pauli terms, the number of samples quickly becomes high for systems of practical interest. Moreover, PQC are typically hard to train due to gradients vanishing in all directions; a phenomenon better known as barren plateau. It is known that a poor ansatz design [152] is likely to experience barren plateau, as well as non-local observable [153].

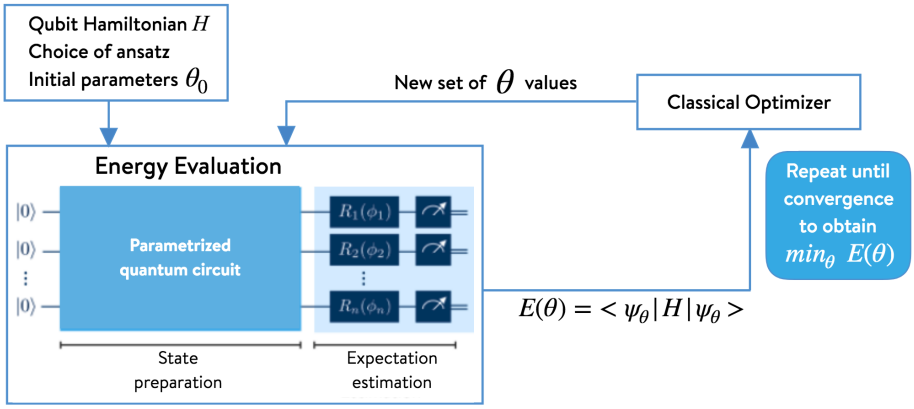

Figure 4.1: **Outline of the VQE.** Figure taken from Ref. [154].

We note that variational methods are not the only way to preparing ground state on a quantum computer. For instance, an other approach is given by Adiabatic state preparation (ASP), which have been extensively researched for the preparation of ground states [155]. They feature numerous innovations such as shortcuts to adiabaticity [156], counterdiabatic driving strategies [157–159], and hybrid methods like counterdiabatic optimal local driving [160, 161]. Nonetheless, the scalability of these methods to extensive system sizes has yet to be verified, as it requires an evolution time scaling inversely to the spectral gap, which is often exponentially small.

Even if variational methods have some caveats, they still have the merit of preparing actual quantum states, and not only estimating its energy. Therefore, they can serve as input for asymptotically optimal quantum algorithms, such as QPE [31] which we touched in Section 2.3. It is expected that variational methods perform better than simpler choices, such as Hartree Fock (HF) or random guesses. Moreover, they are compatible with near term quantum computing since the circuit's depth can be kept arbitrarily shallow. Hence, it remains an important task to design efficient and trainable ansätze, and understand which systems they can solve. In the following, we will prepare the ground and the first excited state of a few different nuclear systems of increasing complexities, which is based on original works presented in Refs. [104, 135].

4.1 LIPKIN-MESHKOV-GLICK MODEL

4.1.1 *Model*

The first system we consider is the LMG model, which was introduced in [162–164] to describe a system of n long-range interacting fermions. The state space constitutes of two n-degenerate shells at fixed energy levels. The fermions can then be exchanged between the two levels, giving rises to a simple, yet non-trivial, description of a nuclei. Using a JW transformation [16], the LMG model can be mapped to a system of interacting spins

$$\mathcal{H} = -\frac{1}{n}\sum_{i<j}^{n} X_i X_j + \gamma Y_i Y_j - B\sum_{i=1}^{n} Z_i, \tag{4.3}$$

which can be directly mapped to a qubit-based quantum computer [138, 165, 166]. The given Hamiltonian describes n spins arranged in a fully linked planar graph, subject to a transverse magnetic field B. The first term in (4.4) denotes a non-uniform interaction in the $x-y$ plane, where each spin is homogeneously interconnected with all other spins. This

serves as a fundamental illustration of long-range interaction. The anisotropy parameter, denoted as $0 \leq \gamma \leq 1$, represents the variation in coupling strengths between the two planar orientations. The system is considered critical and undergoes a second-order phase transition in the thermodynamic limit. This transition occurs between a phase with broken symmetry (disordered) when $B < 1$ and an ordered phase when $B \geq 1$. We introduce a collection of collective-spin operators denoted as

$$S_\sigma = \frac{1}{2}\sum_{i=1}^{n} \sigma_i \quad \text{for } \sigma \in \{X, Y, Z\}. \tag{4.4}$$

.

The LMG Hamiltonian (4.4) is diagonal in the Dicke basis $|j, m\rangle$, which consists of the eigenvectors of $\mathbf{S}^2$ and S_z, denoted as $|j, m\rangle$. We recall that they satisfy the equations $\mathbf{S}^2|j, m\rangle = j(j+1)|j, m\rangle$ and $S_z|j, m\rangle = m|j, m\rangle$. The LMG model interpolates between two easier models: the Lipkin model ($\gamma = 1$) and the fully-connected Ising model ($\gamma = 0$).

4.1.2 *Quantum phase transitions*

Quantum Phase Transition (QPT) in many body systems are defined by an abrupt alteration in the fundamental characteristics of the ground state. Quantum fluctuations are solely responsible for these changes, distinguishing them from thermal phase transitions. These phenomena are commonly attributed to the non-analytic behavior of the ground state properties when the Hamiltonian system undergoes a transition. The phase diagram of this model at zero temperature was obtained in [167] by employing a two-boson Schwinger boson realization of the $SU(1,1)$ Richardson-Gaudin integrable models. Nevertheless, the task of categorizing phase transitions in systems containing a limited number of constituents continues to be a difficult and unresolved issue. For the sake of simplicity, we shall use the terms *critical values* to denote the values of the magnetic field B and the anisotropy γ at which the ground state become degenerate.

4.1.3 *Hardware efficient ansatz*

We use a simple hardware efficient ansatz [144], which has the advantage of being easier to run NISQ devices, without an overhead due to circuit transpilation. For instance, we use an ansatz built upon D repetitions of layers consisting of rotations around the y-axis $R_y(\theta)$, CNOT gates alternating between even and odd sites and a final rotation layer, as depicted in Figure 4.2 for the case of $D = 1$. Despite the simplicity of the ansatz, we are able to obtain accurate energy predictions. Moreover, we do not observe any trainability issue, most probably because we focus on small system sizes.

4.1.4 *Excited states*

The VQE can be rewritten to also target excited states. The approach utilized here is the Variational Quantum Deflation (VQD) method proposed by Higgott, Wang, and Brierley [168]. This iterative method involves first preparing the ground state, and then searching for the state that minimizes the energy, while being orthogonal to the pre-computed ground

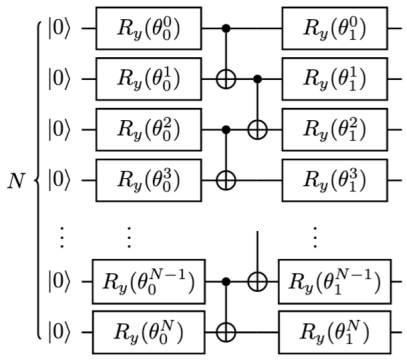

Figure 4.2: **Hardware friendly Ansatz** used for finding the ground state of the LMG model. Figure taken from Ref. [104].

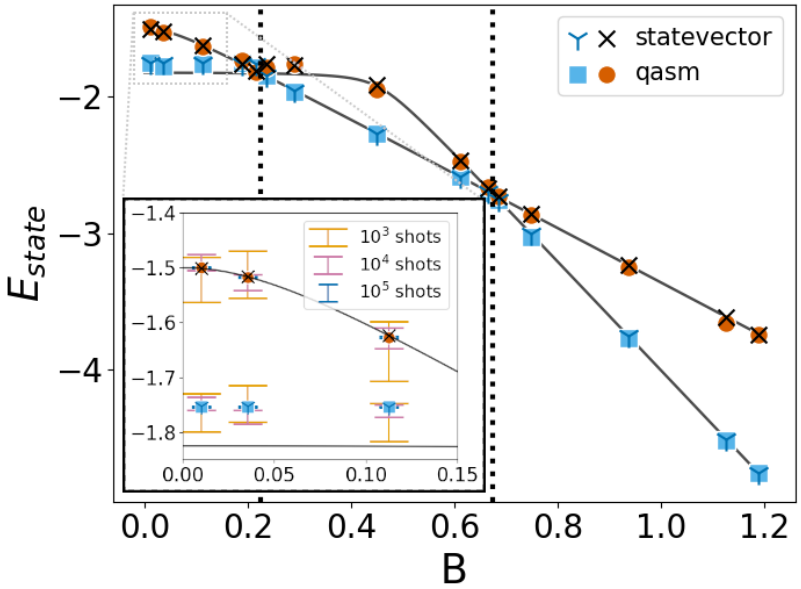

Figure 4.3: **Ground and first excited state energy** of the LMG algorithms obtained with exact and simulated calculation. Source [104].

state. This approach can be generalized for the k-th excited state in an iterative manner. The loss function that is minimized in practice is

$$F(\theta_k) = \langle\psi(\theta_k)|H|\psi(\theta_k)\rangle + \sum_{i=0}^{k-1} \beta_i |\langle\psi(\theta_k)|\psi(\theta_i)\rangle|. \tag{4.5}$$

The wavefunction $|\psi(\theta_i)\rangle$ corresponds to the i-th excited state, and β_i are user-defined coefficients. To ensure that the wavefunction converges to the correct excited state, β_i has to be greater than the energy difference between these two state. In practice, these coefficients can be chosen and adapted by hand.

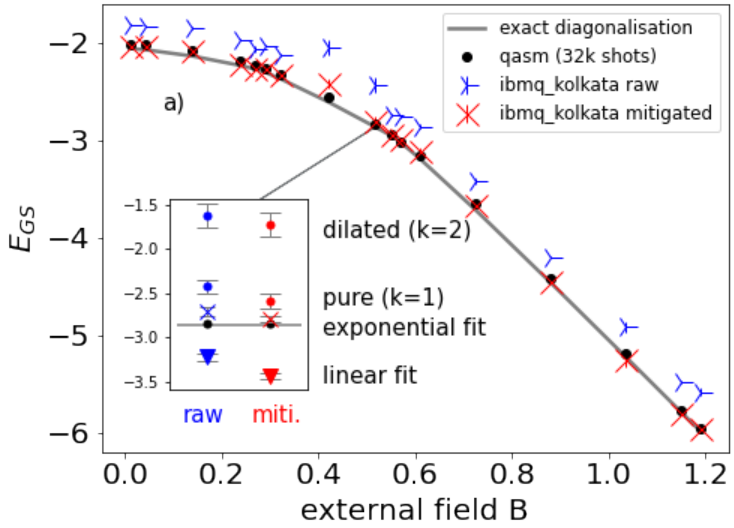

Figure 4.4: **Ground state energy of the LMG model** computed on IBMQ_kolkata. Figure taken from [104].

4.1.5 *Results from classical simulation*

When the number of spins n is small enough, we can diagonalize the Hamiltonian $\mathcal{H}$ and obtain the exact spectrum. We observe multiple crossing points between the energies of the ground and the first excited state. We empirically find that the critical values are found at

$$B_C^k = \frac{n-k}{n}\sqrt{\gamma}, \tag{4.6}$$

for a fixed γ, $k \leq n$ [169], which we aim to verify on a quantum computer. We first start by simulating the quantum algorithm using a statevector approach. We use the previously described hardware efficient ansatz at depth $D = n$ and optimize its parameters with the SLSQP algorithm [170] for a maximum of 2000 iterations. Based on the approach proposed in the adiabatic VQE [171], the converged parameters serve as initial point for the following magnetic field value. Figure 4.3 displays the energy of the ground and first excited states for an LMG Hamiltonian with $n = 4$ spins, at a fixed value of $\gamma = 0.81$, as a function of the magnetic field B. The empirical critical values are represented in dashed lines and correspond to level crossing. The results obtained with the VQE, using finite (qasm) or infinite number of measurement (statevector), are compared to exact diagonalization (black solid lines). We see a good agreements between the exact values and the results given by the VQE.

4.1.6 *Result on IBM quantum hardware*

Ultimately, we conduct an experiment on the quantum device ibm_kolkata with $n = 5$ spins and $\gamma = 0.49$. We first optimize the hardware-efficient ansatz with one layer ($D = 1$) on the noiseless simulator. We further fine tune the ansatz on the quantum hardware using the Simultaneous Perturbation Stochastic Approximation (SPSA) [132] optimizer with maximum of 100 steps. The predicted energy is depicted in Figure 4.4, using 32000 shots, and the error mitigation schemes MMM (see Section 3.2.1), and ZNE (see Section 3.2.2). Solid lines indicate results obtained with exact diagonalization, black points indicate the noiseless

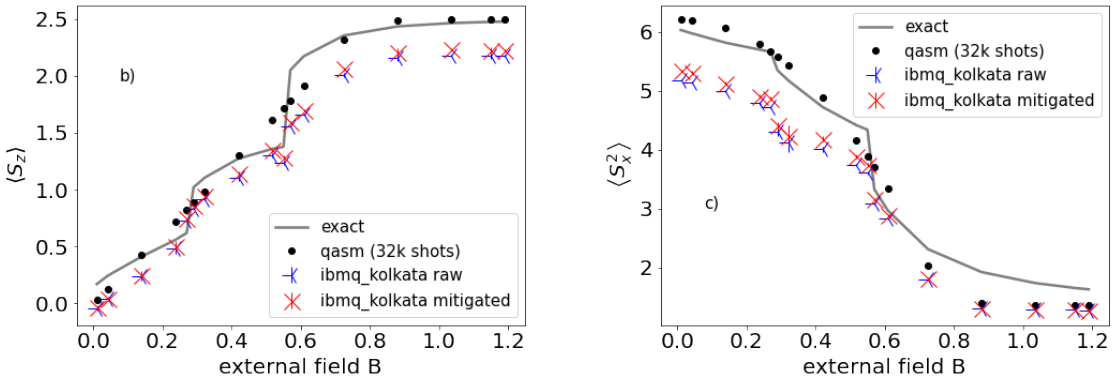

Figure 4.5: Correlators and magnetization using the VQE algorithm on ibmq_kolkata. Figure taken from [104].

qasm simulation with 32000 shots, blue crosses indicate the raw results from the quantum device, and red crosses indicate the extrapolated energies. The error bars represent the 99.5 % confidence interval based on five independent runs. The inset illustrates the impact of various extrapolation functions on a particular point. The result obtained with original circuit is represented by the $k = 1$ point, while the dilated case is represented by $k = 2$. We recanll that in the dilated case, each CNOT is replaced with three CNOT gates. The cross makers display an exponential extrapolation, whereas the triangle markers a linear one. We observe that the linear extrapolation over shoot the target, whereas the exponential fit delivers more accurate results.

Overall, the ground state energy is reproduced with less than 1% error ratio everywhere, suggesting that the quality of current devices is good enough for this particular task.

In addition, two additional observable are measured: the total magnetization S_z and the correlators S_x^2, which are reported in Figure 4.5. They offer additional insights into the phase diagram as they display discontinuities at the critical points.

However, we can see that the additional observable are not as computed as accurately as the energy. There are two parts to the explanation: first, we find that, for large magnetic field, the noiseless simulations also struggle more than for the energy, especially for $\langle S_x^2 \rangle$. This can be expected since operators converge more slowly than the energy as the variational distance decreases [172]. Hence, while the ansatz is able to produce a state with low energy, it seems to not be expressive enough, or trained enough, the to reflect the genuine complexity of the ground state. Hence, increasing the depth yield better solutions but we do not show the results as it would likely be difficult to perform the experiment on real hardware. Secondly, the fine tuning seems to overfit the hardware noise, which indeed minimizes the energy, at the expense of drifting further away from the true ground state.

4.2 LIGHT NUCLEI

In this section, we consider a ^{6}Li nucleus described by the shell model and find the ground and first excited state energy using unitary coupled cluster. It consists of an old problem, which was considered with the first classical computers, which we aim to reproduce here using their quantum counterpart. The model describes light nuclei, from ^{6}He to ^{16}O, built from nucleons in the p shell on top of a frozen ^{4}He core.

4.2.1 *Model*

The model space is comprised of the $0p_{3/2}$ and $0p_{1/2}$ harmonic oscillator orbitals for protons and neutrons, utilizing the Cohen-Kurath interaction [173]. Yet simple, this model has the advantages of being realistic and challenging for current NISQ technologies. We build on Ref. [139] which considered the Deuterium, by expanding it to a larger Hilbert space and a more complex nucleus. More specifically, our work use the same framework as of Ref. [174], which appeared around the same time and computed the energy of ^{6}He and ^{8}Be. The Hamiltonian is written as

$$\mathcal{H} = \sum_i \epsilon_i \hat{a}_i^\dagger \hat{a}_i + \frac{1}{2} \sum_{ijkl} V_{ijlk} \hat{a}_i^\dagger \hat{a}_j^\dagger \hat{a}_k \hat{a}_l \ , \tag{4.7}$$

where $\hat{a}_i^\dagger$ and $\hat{a}_i$ are creation and annihilation operators, respectively. The single-particle energies are denoted by ϵ_i, and V_{ijkl} represent the two-body matrix elements. The quantum state are labeled by $|i\rangle = |r = 0, l = 1, j, j_z, t_z\rangle$. The radial quantum number is denoted by r, the orbital angular momentum by l, the total spin by $j = 1/2, 3/2$ and j_z its projection along the z axis. Finally, the isospin $t_z = \pm 1/2$ differentiate between protons and neutrons. We thus have six proton orbitals and six neutron orbitals in the p shell model space, requiring $n = 12$ qubits. he total spin J, isospin T, and their respective projections J_z and T_z are all preserved by the Cohen-Kurath interaction. We shall make use of these conservation laws in building the wavefunction ansatz.

The shell-model Hamiltonian (4.7) can be converted into a qubit Hamiltonian via the JW [16] transformation

$$\hat{a}_i^\dagger = \frac{1}{2} \left(\prod_{j=0}^{i-1} -Z_j \right) (X_i - iY_i), \text{ and } \hat{a}_i = \frac{1}{2} \left(\prod_{j=0}^{i-1} -Z_j \right) (X_i + iY_i). \tag{4.8}$$

We recall that X_i, Y_i and Z_i are the Pauli matrices acting on the i-th qubit. In this framework, each qubit represents a single-particle state, where $|0\rangle$ means that the orbital is free and $|1\rangle$ occupied. Even if the model is relatively simple, the Hamiltonian (4.7) yet contains 975 non-local Pauli terms. In fact, the number of term scales naively as $\mathcal{O}(n^4)$ because of the two-body interaction. While the conservation of spin and isospin reduce the scaling by an order of magnitude, this number still quickly becomes important. Another challenge of this model is the non-locality of the Pauli terms, since the short-range nuclear interaction becomes non-local when represented in the harmonic-oscillator basis.

4.2.2 *Unitary coupled cluster ansatz*

CC is a numerical method employed to describe many-body systems. It is a post-HF ab initio quantum chemistry method, which is frequently employed in the field of computational chemistry, as well as in nuclear physics [175–177]. In essence, coupled cluster starts from the fundamental HF solution and employs the exponential cluster operator to generate multi-electron wavefunctions that account for nucleon correlation. The cluster operators $\hat{T}$ effectively destroys and recreates particles in different orbitals. It can be decomposed into singles, i.e., one-particle–one-hole (1p-1h), doubles (2p-2h), and higher excitation operators as:

$$\hat{T} = \hat{T}_1 + \hat{T}_2 + \dots \tag{4.9}$$

with

$$\hat{T}_1 = \sum_{i\in\text{virt};\alpha\in\text{occ}} \theta_i^\alpha \hat{a}_i^\dagger \hat{a}_\alpha, \text{ and } \hat{T}_2 = \sum_{i,j\in\text{virt};\alpha,\beta\in\text{occ}} \theta_{ij}^{\alpha\beta} \hat{a}_i^\dagger \hat{a}_j^\dagger \hat{a}_\alpha \hat{a}_\beta. \tag{4.10}$$

Greek indices indicate occupied orbitals of the initial state, whereas Latin indices represent virtual (empty) ones. In essence, the cluster operator builds a superposition of states with the corrects orbitals, by transforming the initial states. To make it compatible with quantum computing, we require the cluster operator to be unitary. Therefore we consider the Hermitian matrix $\hat{T} - \hat{T}^\dagger$, whose complex exponential is unitary, and apply it on the HF state as

$$|\psi(\Theta)\rangle = e^{i\left(\hat{T}(\boldsymbol{\theta}) - \hat{T}^\dagger(\boldsymbol{\theta})\right)} |\psi_0\rangle. \tag{4.11}$$

By making the parameters variational, we can navigate from HF towards the ground state. In order to respect the symmetry, we only examine excitation operators that have a total angular momentum projection of $J_z = 0$. The unitary cluster ansatz is transformed into a qubit operator with adjustable parameters Θ using the JW mapping. It is then transpiled into a universal set of gates by using a PF, which approximate the exponential with arbitrary accuracy, as we will see in Chapter 5. In the following, we simply use a first-order Trotter formula with one step.

4.2.3 *Initial state*

The initial state $|\psi_0\rangle$, which is act upon by the ansatz, is usually chosen as the HF state. The HF state is the product state which minimizes the energy, and can be prepared efficiently in second quantization. Moreover, it is straightforward to start in the right J symmetry sector, by choosing among the product state with the correct property. For example, it is known that the projection of total angular momentum of the ground state fulfills $J_z \in \{-1, 0, 1\}$. We therefore choose an initial state in this symmetry sector. Moreover, this fact can also be used to determine the first excited state in the subspace with total $J_z \in \{-3, -2, 2, 3\}$. This approach provides an advantage compared to other methods, such as the VQD mentioned earlier, because it eliminates the need to calculate the overlap with the ground state.

4.2.4 *Excitation operators ordering*

One degree of freedom which remains in our construction is the ordering of the excitation operators in (4.10). While an exact implementation of the cluster operator does not depend on the ordering, the approximate Trotter formula might since it breaks the symmetry. To get a better understanding of this phenomenon, we train the UCC, using the gradient-free SPSA algorithm [132] for different ordering: random (blue), ascending (green) and descending (red) with respect to the respective coefficients in the Hamiltonian. This means that we associate the weights from (**??**) with the different operators in the UCC ansatz and order them accordingly. The gray area represents solutions achieving a relative error smaller than 1%, i.e. those w

$$|\Delta E/E| \equiv |E_{VQE} - E_{exact}|/|E_{exact}| < 0.01, \tag{4.12}$$

which is usually enough for practical applications. We observe that ordering the operators in descending magnitude order leads to the second optimal solution, hinting that in it is favorable in practice. Intuitively, we can understand this result as prioritizing the most important operators is the most effective choice.

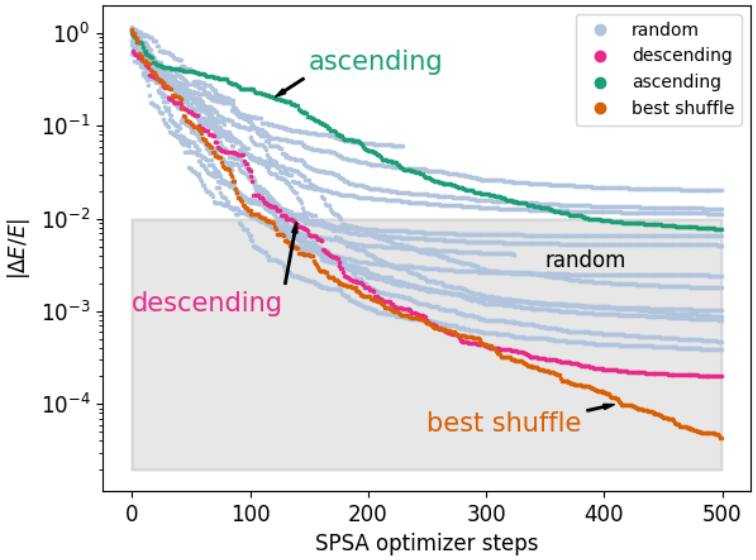

Figure 4.6: Excitation ordering for the energy of the ^{6}Li nuclei. Figure taken from [135].

4.2.5 *Qubit-based excitation*

Due to the JW mapping, the fermionic excitation operators act on a set of $\mathcal{O}(n)$ qubits. Consequently, their implementation on NISQ devices is costly due to the long range interaction, resulting in an increased amount of CNOT and SWAP gates required to match the linear connectivity of IBM's quantum chip. An effective and alternate approach to decrease this gate cost is to consider QBE operators, which apply the excitation operators on a fixed number of qubits by ignoring the Z operators in the JW mapping. Creation and annihilation operators are therefore replaced by

$$\hat{a}_i^\dagger = \frac{1}{2}(X_i - iY_i),\ \hat{a}_i = \frac{1}{2}(X_i + iY_i). \tag{4.13}$$

The main difference is that the resulting operator does not respect fermionic anti-commutation relations, which are ensured via the product of Pauli Z matrices. Despite this fact, QBE-UCC ansätze exhibit a good level of effectiveness for ground state calculations, and are much cheaper to implement because they operate on a fixed number of qubits. In fact, the singles act on two qubits, the doubles on four qubits, and, more generally, the kth excitation operators on 2^k qubits. More precisely, single excitation operators between qubits i and j read

$$U_{ij}(\theta) = \exp\left[i\frac{\theta}{2}(X_iY_j - Y_iX_j)\right] \tag{4.14}$$

and double excitation operator between qubits i, j, k and l are given by

$$\begin{aligned} U_{ijkl}(\theta) = \exp\Big[i\frac{\theta}{8}(&X_iY_jX_kX_l + Y_iX_jX_kX_l + Y_iY_jY_kX_l + Y_iY_jX_kY_l \\ &- X_iX_jY_kX_l - X_iX_jX_kY_l - Y_iX_jY_kY_l - X_iY_jY_kY_l)\Big]. \end{aligned} \tag{4.15}$$

The precise circuit formulation can be found in the original paper [178].

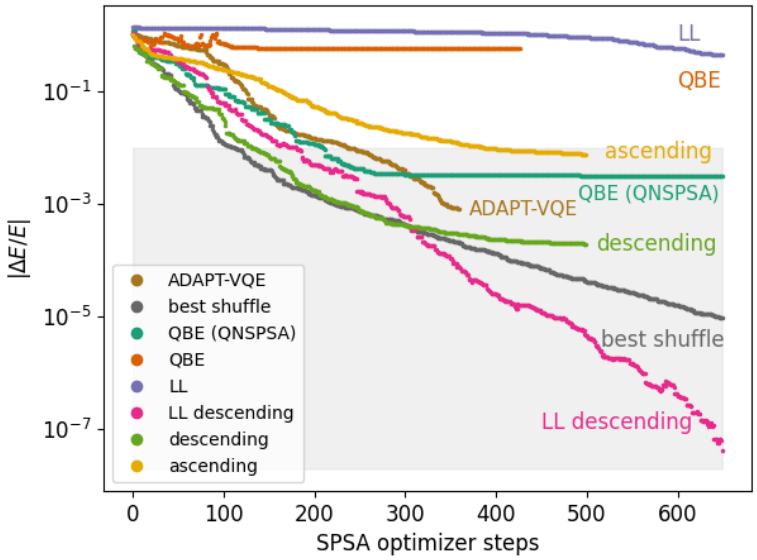

Figure 4.7: **Training strategy** for the energy of the ^{6}Li nuclei. Figure taken from [135].

4.2.6 *Layerweise learning*

Layerweise Learning (LL), as proposed in Ref. [179], is a strategy aimed at addressing trainability issues, and more specifically barren plateaus, at the beginning of the training. Given a pool of operators $\{\tau\}_i$, the objective is to build the ansatz iteratively by choosing from the set and optimize, before adding a new operator until convergence is achieved. We can either choose the next operator by hand, or further enhance this by selecting the subsequent operator that maximizes the gradients' value.

$$\left|\frac{\partial E}{\partial \theta_i}\right|_{\theta_i=0} = |\langle\psi|[H,\hat{\tau}_i]|\psi\rangle|, \tag{4.16}$$

a technique effectively known as ADAPT-VQE [134]. This has the advantage of picking the most promising operators, at the expense of additional gradients computation.

By following these different strategies on a classical simulator, we observe that the most promising method is given by the descending fermionic UCC ansatz optimized in a layerweise fashion, as shown in Figure 4.7. This means in particular, that we first pick the operator with the largest corresponding coefficient in the Hamiltonian, perform $k = 10$ optimization steps, before adding, the second largest operators and so on.

A noteworthy point learned in this experiment is the importance of employing the quantum natural gradients to train the QBE ansatz (shown in orange versus dark green). A quantum natural SPSA optimizer [180] is a variation of SPSA that employs the geometry of the Hilbert space, induced by the quantum geometric tensor [181]

$$G_{ij}(\theta) = \left\langle \frac{\partial\psi_\theta}{\partial\theta_i}, \frac{\partial\psi_\theta}{\partial\theta_j} \right\rangle - \left\langle \frac{\partial\psi_\theta}{\partial\theta_i}, \psi_\theta \right\rangle \left\langle \psi_\theta, \frac{\partial\psi_\theta}{\partial\theta_j} \right\rangle, \tag{4.17}$$

to speed up the convergence. The original version of quantum natural gradient is expensive to implement since it the Hessian of the cost function. Quantum natural SPSA however, approximates it using only six circuit evaluations, which decreases the number of measurements to be independent on the number of parameters.

	No. CNOT	mean	st. deviation	exact	error ratio
raw (gs)	209	-6.27	0.269	-5.529	13.36%
mitigated (gs)	209	-5.319	0.24	-5.529	3.81%
raw (1st es)	41	-2.907	0.87	-3.420	14.97%
mitigated (1st)	41	-3.424	0.08	-3.420	0.12%

Table 4.1: Hardware results of the QBE-UCC ansatz. Source [135].

4.2.7 *Results on IBM quantum hardware*

We assess the performance of the 27 qubit IBM_mumbai superconducting chip using the QBE-UCC ansatz. The ansatz is first trained on the statevector simulator, before being transpiled, to ensure compatibility with the hardware topology, using the SWAP-based BidiRectional (SABRE) heuristic search algorithm [182]. Since it consists of a stochastic algorithms, several iterations are conducted and only the one reducing the overall count of CNOT gates is chosen. Overall, the SABRE method achieves a 50% reduction in CNOT operations compared to the deterministic approach. As far as error mitigation is concerned, we only use MMM to calibrate the readout, as the other ones increase the depth to a point which renders the results useless.

The experiments are repeated ten times for the ground-state and five times for the first excited state, utilizing 8092 shots for each qubit-wise commuting group of Pauli terms. The obtained data [135] is reported in Table 4.1, along with the final CNOT gates count after transpilation. The energy of the ^{6}Li is replicated with an accuracy of 3.81% and 0.12% for the ground and first excited state, respectively, which are within one standard deviation. We observe that the MMM scheme improves accuracy by over 10%.

4.2.8 *Comparison with other works*

While chemistry and condensed matter have been widely explored using the VQE, nuclear physics has been relatively neglected until now. The first demonstration on the deuterium is the one by Dumitrescu et al. [139] and a similar approach to the one described here has been performed in Ref. [174] for the ^{8}Be nuclei. Scaling to larger system size has been extensively performed more recently using the ADAPT-VQE approach [183] up to ^{48}Ca, as well as for the LMG model [141]. Entanglement forging, splitting the contribution from the protons and neutrons, has been explored in Ref. [142], which solve the *p* shell and the *sd* to some extent. However, we note that all of these large scale approaches have only been implemented using statevector simulations, which in particular does not take into account the statistical noise from measurement. It is expected that the measurement overhead to be significant, and it remains of high-interest to estimate how large the shot budget would need to be in order to solve problems of this type.

4.3 PION LESS EFFECTIVE FIELD THEORY

In this section we consider a lattice description of a Triton, i.e. one proton and two neutrons bounded together. This model will be used latter in this thesis, which is why it is introduced here.

4.3.1 *Model*

The model we are examining is a basic toy model for a Triton, which was introduced in [184] and subsequently studied in [110, 185]. This model is inspired by a pionless lattice Effective Field Theory (EFT) [186]. We consider a $2d$ lattice of size $L \times L$ with periodic boundary conditions, with $A = 2$ dynamical nucleons and one static one (infinite mass) on the first site. Despite the model being straightforward and easily simulable, at least for small sizes, it still incorporates a significant amount of the leading order contributions to the interaction, which can provide valuable information about light nuclei and their response function. The Hamiltonian of this model is formally equivalent to a $2d$ Fermi Hubbard model with a hopping term

$$H_{\text{kin}} = -t \sum_{f=\{\uparrow,\downarrow\}} \sum_{\langle i,j \rangle} c_{i,f}^{\dagger} c_{j,f} \ , \tag{4.18}$$

a two-body contact interaction

$$H_{\text{int}} = U \sum_{i=1} n_{i,\uparrow} n_{i,\downarrow} \tag{4.19}$$

and additional one and two-body potentials generated by the static proton at lattice site $i = 1$

$$H_{\text{static}} = U \sum_{f=\{\uparrow,\downarrow\}} n_{1,f} + V n_{1,\uparrow} n_{1,\downarrow} \ . \tag{4.20}$$

Let us remember that the fermionic operator $c_{i,f}$ annihilates a particle of the species f on site i, that $c_{i,f}^{\dagger}$ is the corresponding creation operator, and $n_{i,f} = c_{i,f}^{\dagger} c_{i,f}$ represents the number operator of species f on site i . The kinetic term involves a summation across neighboring sites $\langle i, j \rangle$. The variables U and V denote the magnitudes of the two-body and three-body interactions, respectively. Refer to [184, Table 1] (source [187]) for realistic values of t, U, and V using a physical lattice spacing of $a = 1.4$fm. To obtain a qubit Hamiltonian, we can either go for a second quantisation formalism, with, e.g., the JW transformation as used for the ^{6}Li nucleus. Here, we opt for the more compact representation in first quantization, and use only $2\lceil \log(L^2) \rceil$ qubits. We note that we obtain an exponential reduction in the number of qubits compared to a second quantization approach, at the price of the operators being highly non-local.

To simplify the hopping term, a Gray code ordering of lattice states is used, as described in Refs. [184, 188]. For the specific case of $L = 2$, it consists of the following identification

$$|1\rangle \equiv |00\rangle \ \ |2\rangle \equiv |01\rangle \ \ |3\rangle \equiv |10\rangle \ \ |4\rangle \equiv |11\rangle \ . \tag{4.21}$$

As in Refs. [184, 185], the model is simplified by setting $t = 1$, $U = -7$ and $V = -4U$, leading to most of the terms in the Hamiltonian to cancel each other, and can finally be expressed in the Pauli basis as

$$H = 4.5 \cdot \mathbb{1} - 2 \sum_{i=1}^{4} X_i + 1.75 \left(\sum_{i<j<k} Z_i Z_j Z_k + Z_1 Z_4 + Z_2 Z_3 \right) . \tag{4.22}$$

4.3.2 *Ansatz*

We use the following ansatz, already used in Ref. [184], for the ground state

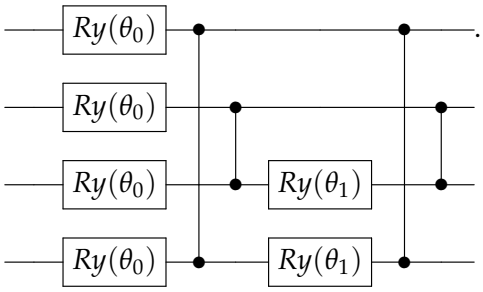

(4.23)

The parameterization of the system involves two angles, θ_0 and θ_1, and only four CZ gates with linear connectivity. The entanglement structure follows the UCCSD (with single and double) format, which would be anticipated for the case $V = 0$, when three-body interactions are absent. Although the ansatz is simple, it achieves an accuracy of less than 10% relative error. The parameters are optimized using the gradient-free optimizer COBYLA (COnstrained Optimization BY Linear Approximations) [133] on a simulator. Extensive experimentation suggests that this ansatz is the most accurate strategy using a shallow quantum circuit, which is probably due to the model's simplicity.

4.3.3 *Excitation operator*

Subsequently, we will examine a scattering experiment in which the probe interacts with the nucleon density by transmitting a momentum $\vec{q}$ and an energy ω to the target. Due to the use of a periodic spatial lattice, the momenta are quantized as

$$\vec{q}_k = \frac{\pi}{La}\vec{x}_k, \tag{4.24}$$

where La denotes the spatial length of the lattice and $\vec{x}_k \in \mathbb{N}^2$ represents the location of the k-th momentum in the reciprocal lattice. We concentrate on probes interacting with the nucleon density $n_{i,f}$, characterized by the operator

$$\hat{O}(\vec{q}_k) = \sum_{f=\uparrow,\downarrow} \rho_f(\vec{q}_k) = \sum_{f=\uparrow,\downarrow} e_f \sum_i e^{i\vec{q}_k\cdot\vec{r}_i} n_{i,f}, \tag{4.25}$$

where e_f denotes the charge of the nucleon f, $\rho_f(\vec{q}_k)$ the nucleon densities in momentum space and $\vec{r}_i$ the position of site i on the spatial lattice.

4.4 DISCUSSIONS

In this chapter, we have discussed how to prepare ground state on a quantum computer using variational methods, and have applied this technique to various nuclear models. While the VQE retains certain advantages, such as being NISQ friendly and could lead to application on near term quantum computers, as demonstrated by the experiments shown here, it also suffers from many problems. The main one is the number of samples required to estimate the energy expectation value to high accuracy, which naively scales as $\mathcal{O}(1/\epsilon^2)$. One hidden issue is that the Hamiltonian is composed of many terms L ($= \mathcal{O}(n^4)$ for the shell model), which increases the scaling roughly to $\mathcal{O}(L/\epsilon^2)$. While amplitude

amplification can reduce the scaling to $\mathcal{O}(1/\epsilon)$, this would also deny the possibility of grouping commutating terms, with the addition of having a large depth overhead. Gonthier et al. [148] estimated the time required to estimate the energy of the CO_2 molecule with chemical accuracy to more than a month of QPU time, which is considerable. Moreover, the optimization of quantum circuits is paved with traps [151] and unfavorable landscapes [7]. While the VQE has generated much excitation in the community due to its ability to run on NISQ devices, there is little evidence that it will remain competitive in the long term. We advocate instead to prepare states obtained with classical methods such as DMRG or CC, load them into the quantum computers and used them as input for more advanced quantum algorithms, as we will explore more in details in Chapter 7.

5

DIGITAL QUANTUM SIMULATION WITH RANDOM PRODUCT FORMULAS

Quantum simulation is the task of using quantum systems to model the behavior of other quantum systems, which can be hard to simulate on classical computers. Numerous challenges in chemistry [189], physics, and materials science [190] —such as determining molecular characteristics, superconductivity, or simulating quantum field theories [191]—are arduous to resolve due to the exponential growth of the Hilbert space. The objective of quantum simulation is to emulate these properties with a controlled quantum system, such as a quantum computer. Quantum systems inherently possess property inherent to quantum physics, making them an appealing venue to simulate the dynamics and characteristics of other quantum systems, hence offering insights would would be otherwise out of reach for conventional computers. Quantum simulations come in two flavors: analog [192] and digital [19]. An analog quantum simulator is a controllable implementation of the system of interest and is tailored to a specific quantum system. It is therefore much closer to an actual experiment and offer no guarantees on the accuracy of the results. Digital simulations requires decomposing the evolution of a quantum system into a sequence of quantum gates executable on a universal quantum computer. They are more versatile and can tackle arbitrary systems, at the cost of being more expensive in term of number of operations.

Performing quantum dynamics efficiently is an important prospect of quantum computing, and is largely motivated by numerous applications, such as understanding quantum systems, or to use as a subroutine in more advanced quantum algorithms. Therefore, it is important to understand how to drive dynamics efficiently. In this chapter, we will solely focus on PF, in a deterministic and random setting. While they may not be asymptotically optimal compared to methods based on Linear Combinations of Unitaries (LCU), they remain remarkably efficient in practice due to their ability to exploit locality in the physical system. Moreover, they often outperform theoretical guarantees [190], making them a good candidate for near-term quantum computers. We will focus on two particular examples in Chapter 6: neutrinos thermalization and QPT dynamics. For now, we will focus on the theory of PF for digital quantum simulations.

In the following, we focus on time-independent problems. The dynamic of a quantum state $|\psi(t)\rangle$ is governed by the time-independent Schrödinger's equation

$$i\hbar\frac{d}{dt}|\psi(t)\rangle = \mathcal{H}|\psi(t)\rangle. \tag{5.1}$$

We work in natural units and set $\hbar = 1$. Assuming the quantum state is known for some time t_0, its time evolution can be computed with the matrix exponential

$$|\psi(t)\rangle = \mathcal{U}(t-t_0)|\psi(t_0)\rangle = \exp\left(-i\mathcal{H}(t-t_0)\right)|\psi(t_0)\rangle, \tag{5.2}$$

where $\mathcal{U}(t)$ is called the time evolution operator. The difficulty resides in the size of the Hamiltonian, which is exponential in the system' size n, and its exponentiation is therefore not straightforward.

In this chapter, we will describe deterministic and random decomposition methods, aiming at approximating the matrix-exponential with tolerance ϵ. We assume that the Hamiltonian can be decomposed in the Pauli basis as

$$\mathcal{H} = \sum_{j=1}^{L} h_j H_j, \text{ where } h_j \in \mathbb{R} \text{ and } H_j = \bigotimes_{l=1}^{n} \sigma_{\alpha_l}, \ \left\|H_j\right\| = 1 \tag{5.3}$$

where $\alpha_i \in \{I, X, Y, Z\}$ being a Pauli letter and σ_{α_i} the corresponding Pauli matrix acting on the i-th qubit.

5.1 TROTTER-SUZUKI DECOMPOSITION

Trotter-Suzuki decompositions are deterministic PF designed to effectively approximate the matrix exponential. When applied to the Hamiltonian described in equation (5.3), the Suzuki formula $\mathcal{S}_n$ provides an estimate of the n-th order in the Taylor series expansion. The first order formula,

$$\mathcal{S}_1(t) = \prod_{j=1}^{L} \exp\left(-itH_j\right) = \mathcal{U}(t) + \mathcal{O}(t^2), \tag{5.4}$$

which consists of the sequential product of the individual terms, is equal to the exact time evolution operator $\mathcal{U}(t)$ up to first order in t. To increase the precision, one may slice the time into r intervals, also known as Trotter steps. Using the triangle inequality, we obtain

$$\left\|\mathcal{U}(t) - \mathcal{S}_1\left(\frac{t}{r}\right)^r\right\| \leq r \left\|\mathcal{U}(t) - \mathcal{S}_1\left(\frac{t}{r}\right)\right\| \leq C_{\text{comm}} \cdot \frac{t^2}{r}, \tag{5.5}$$

where we see that we can reduce the error by increasing the number of steps. The prefactor $C_{\text{comm}} = \sum_{l_1<l_2} \left\|[H_{l_1}, H_{l_2}]\right\|$ is the commutator bound [193]. To increase convergence, we can incorporate high order from the Taylor expansion into the PF. For example, the second order is given by

$$\mathcal{S}_2(t) = \prod_{j=1}^{L} \exp\left(-itH_j/2\right) \prod_{j=L}^{1} \exp\left(-itH_j/2\right) = \mathcal{U}(t) + \mathcal{O}(t^3), \tag{5.6}$$

where we multiply the individual exponential in a forward and backward manner to cancel some errors and compensating for non-commutativity. In general we have the recursive formula for arbitrary even-order PF

$$\mathcal{S}_{2k}(t) = \mathcal{S}_{2k-2}(p_k t)^2 \mathcal{S}_{2k-2}((1-4p_k)t)^2 \mathcal{S}_{2k-2}(p_k t)^2 = \mathcal{U}(t) + \mathcal{O}(t^{2k+1}), \tag{5.7}$$

with $p_k = (4 - 4^{\frac{1}{2k-1}})^{-1}$. We stress out that, in practice, the performance of deterministic PF exceeds the theoretical upper bound [190], which makes it important to compare practical and theoretical performance for a specific systems and initial states. Hence, those are worse case scenarios. In practice, one deal with specific cases, which might be better described by the average case error.

5.2 RANDOM COMPILATION

5.2.1 *The qDrift channel*

Although deterministic formulas exhibit satisfactory performance and are relatively easy to implement, the number of gates required increases linearly with the number of terms L,

which can be substantial. In the example case of chemistry and nuclear physics, the number of terms L increases as $\mathcal{O}(n^4)$. However, it is important to note that the weights in most of these circumstances are not evenly distributed. Therefore, as proposed by Campbell [194], the number of gates can be greatly decreased in some situations by taking into account the relative significance of each term h_i. The main result is given by

Theorem 1 (qDrift bias bound, [194, 195]). *Let $\mathcal{E}(t;N)$ be a qDrift experiment constructed from N samples and probability distribution $p = h/\lambda$. Then, in expectation value, the bias from the exact time-evolution operator is bounded in diamond norm by*

$$\left\|\mathcal{U}(t) - \overline{\mathcal{E}}(t;N)\right\|_\diamond \leq \frac{2\lambda^2 t^2}{N}, \tag{5.8}$$

where $\lambda = \|\mathcal{H}\|_1 = \sum_{l=1}^{L} |h_l|$ is the one-norm of the Hamiltonian.

The results are given in term of the worse case diamond norm.

Definition 1. *The diamond norm, is the worse case distance between two quantum channels, or more precisely, complete positive trace-preserving maps, and reads*

$$\|\mathcal{E}\|_\diamond = \max_\rho \|(\mathcal{E}\otimes\mathcal{I})(\rho)\|_1, \tag{5.9}$$

where the maximum is taken over all possible quantum states and $\|\cdot\|_1$ denotes the one, or trace, norm.

In the following, we will construct the qDrift channel and prove this claim. The qDrift protocol approximates the ideal time evolution operator by constructing a channel $\mathcal{E}(t;N,M)$ by arbitrarily sampling terms from the Hamiltonian based on the magnitude of their coefficients. In practice, a qDrift experiments consists of a product of N unitaries sampled with probability $p(j) = h_j/\lambda$, and is given by

$$V_j(t) = \prod_{k=1}^{N} e^{-i\tau_{j_k} H_{j_k}} \ . \tag{5.10}$$

The probability of drawing this channel is thus

$$p_N(\boldsymbol{j}) = \lambda^{-N} \prod_{k=1}^{N} h_{j_k}. \tag{5.11}$$

In the above, the time steps τ_{j_k} are explicitly chosen to reduce the bias and $\boldsymbol{j} = (j_1, j_2, ..., j_N)$ is a multi-index labeling the experiment.

As summarized in Algorithm 1, the qDrift channel is constructed as the arithmetic average of the M individual experiments

$$\mathcal{E}(t;N,M)[\rho] = \frac{1}{M}\sum_{m}^{M}\left[V_{j^m}\rho V_{j^m}^\dagger\right], \tag{5.12}$$

where the sampling distribution $q(j) = p(j) \coloneqq h_j/\lambda$ is used. We observe that the superscript m in the bold multi-index $\boldsymbol{j}^m$ denotes the multi-index $\boldsymbol{j}$ of the m-th experiment, while j_k^m denotes its corresponding k-th component.

In the asymptotic limit of infinite number of experiments, the channel converges to the average qDrift channel

$$\overline{\mathcal{E}}(t;N)[\rho] = \mathbb{E}_{p_N}\left[\mathcal{E}(t;N,1)[\rho]\right] = \sum_{j_1=1}^{L}\cdots\sum_{j_N=1}^{L} p_N(\boldsymbol{j}) V_{\boldsymbol{j}}(t)\rho V_{\boldsymbol{j}}^\dagger(t)\ , \tag{5.13}$$

Algorithm 1: STOCHASTIC QUANTUM SIMULATION [196]

Input: Hamiltonian $H = \sum_{l=1}^{L} h_l H_l$ with interaction strength $\lambda = \sum_l h_l$, $h_l > 0$ and $\|H_l\| = 1$, total simulation time t, number of samples N, number of experiments M and initial state ρ. Probability density distribution $q(j)$.

```
m ← 1;
E ← 0;
while m ≤ M do
    V ← 𝟙;
    n ← 1;
    while n ≤ N do
        sample j ~ q(j);
        τ_j ← t h_j / (q(j) N);
        V ← e^{-i H_j τ_j} V;
        n ← n + 1;
    end
    E[ρ] ← E[ρ] + V ρ V† ;
    m ← m + 1;
end
ρ ← (1/M) E[ρ];
```

Output: final state $\mathcal{E}(t; N, M)[\rho]$

where we denote with $\mathbb{E}_p[f(x)] = \sum_x p(x) f(x)$ the expectation value of an arbitrary function $f(x)$ sampled with respect to $p(x)$. The special case with $N = 1$ is given by

$$\overline{\mathcal{E}}(t;1)[\rho] = \sum_{j=1}^{L} p(j) e^{-i\tau_j H_j} \rho e^{i\tau_j H_j}, \tag{5.14}$$

which is refereed to as the deterministic qDrift channel.

The strength of each unitary is fixed to a constant $\tau_j := \tau = t\lambda/N$, which is chosen such that the qDrift channel is equal to the exact channel up to first order in the Taylor expansion. We shall now show the anticipated results Theorem 1, i.e., the bounds on the bias. The proof relies heavily on Ref. [195], improving the results from [194]. We show the proof in the main text since it is instructive for the generalisation later. We start by proving an intermediate results about error accumulation.

Lemma 1 (Error accumulation, [195])**.** *Let $\mathbb{E}[V]$ and $U^{1/N}$ be matrices with bounded operator norm, i.e., $\mathbb{E}[V] \leq 1$ and $\|U^{1/N}\| \leq 1$. Then*

$$\|(\mathbb{E}[V])^N - (U^{1/N})^N\| \leq N \|\mathbb{E}[V] - U^{1/N}\|. \tag{5.15}$$

Proof. The triangle inequality and sub-multiplicativity imply

$$\begin{aligned} \|A_1 A_2 - B_1 B_2\| &= \|(A_1 - B_1) A_2 + B_1 (A_2 - B_2)\| && (5.16)\\ &\leq \|A_2\| \|A_1 - B_1\| + \|B_1\| \|A_2 - B_2\|, && (5.17) \end{aligned}$$

for arbitrary matrices A_1, A_2, B_1, B_2 with compatible dimensions. Using the operator norm bounds, we apply this relation recursively to obtain

$$\begin{aligned} \|\mathbb{E}[V]^N - (U^{1/N})^N\| &= \|(\mathbb{E}[V])(\mathbb{E}[V])^{N-1} - U^{1/N}(U^{1/N})^{N-1}\| && (5.18)\\ &\leq \|\mathbb{E}[V] - U^{1/N}\| + \|(\mathbb{E}[V])^{N-1} - (U^{1/N})^{N-1}\| && (5.19)\\ &\leq \cdots \leq N \|\mathbb{E}[V] - U^{1/N}\|, && (5.20) \end{aligned}$$

which is the desired result. □

We can now show Theorem 1.

Proof. To simplify the notation, we define the random matrix $V = \exp(-iX)$, with $X = t\lambda H/N$ and H sampled from the Hamiltonian. Moreover, we divide the time-evolution operator into slices as $U^{1/N} = \exp(-it\mathcal{H}/N)$. It then follows from Lemma 1 that

$$\begin{aligned}
\|(\mathbb{E}[V]) - U\| &= \left\|(\mathbb{E}[V]) - (U^{1/N})^N\right\| && (5.21)\\
&\leq N\left\|(\mathbb{E}[V]) - U^{1/N}\right\| && (5.22)\\
&= N\|\mathbb{E}\left[\exp(-iX) - \mathbb{1} + iX\right] + \left[\mathbb{1} - i\mathbb{E}[X] - \exp(-i\mathbb{E}[X])\right]\| && (5.23)\\
&\leq N\mathbb{E}\left\|\exp(-iX) - \mathbb{1} + iX\right\| + N\left\|\exp(-i\mathbb{E}[X]) - \mathbb{1} + i\mathbb{E}[X]\right\| && (5.24)\\
&\leq \frac{N}{2}\mathbb{E}\left\|X\right\|^2 + \frac{N}{2}\left\|\mathbb{E}[X]\right\|^2 && (5.25)\\
&\leq N\mathbb{E}\left\|X\right\|^2 = \frac{(t\lambda)^2}{N}, && (5.26)
\end{aligned}$$

where we used Jensen's inequality and a generalisation of the bound $|e^{ix} - ix - 1| \leq x^2/2$ from $x \in \mathbb{R}$ to Hermitian matrices. □

5.2.2 *Important sampled qDrift*

IMPORTANCE SAMPLING Importance sampling is a valuable method for calculating expectation values in cases when sampling from the distribution $p(x)$ is challenging, or as a means of decreasing variance. The technique is commonly employed in Monte Carlo integration for variance reduction and involves sampling from a different distribution, denoted as $q(x)$, which is often easier to work with. The samples are then re-weighted accordingly as

$$\mathbb{E}_p[f(x)] = \sum_x p(x)f(x) = \sum_x q(x)\frac{p(x)}{q(x)}f(x) \equiv \mathbb{E}_q[\omega(x)f(x)]. \tag{5.27}$$

Here, $\omega(x)$ represents the re-weighting factor, defined as the ratio $p(x)/q(x)$. It is clear that this provides us with an unbiased estimator of the expected value of interest, whereas the variance is dependent on the choice of $q(x)$. By making a suitable decision, the variance can be substantially diminished, resulting in more cost-effective computations. We recommend referring to [197] for a comprehensive review of the subject. In the following, we present results from Ref. [196] which apply importance sampled to qDrift.

CONSTRUCTION First, we define the qDrift channel using importance sampling. In order to gain more insight of the paradigm shift, we employ the Liouvillian representation

$$\mathcal{L}(\rho) = i(H\rho - \rho H) = i[H, \rho]. \tag{5.28}$$

In this framework, the time evolution channel takes the following form

$$\mathcal{E}(t)[\rho] = e^{iHt}\rho e^{-iHt} \equiv e^{t\mathcal{L}}(\rho) = \sum_{n=0}^{\infty}\frac{t^n\mathcal{L}^n(\rho)}{n!}. \tag{5.29}$$

We first express the qDrift channel, sampled from an arbitrary distribution $q(j)$ as [196]

$$\bar{\mathcal{E}}_q(t;1)[\rho] = \sum_j q(j) e^{-i\tau_j H_j} \rho e^{i\tau_j H_j} \equiv \sum_j q(j) e^{\tau_j \mathcal{L}_j}(\rho) \tag{5.30}$$

$$= \left(1 + \sum_j q(j)\tau_j \mathcal{L}_j + \sum_{n=2}^{\infty} \sum_j q(j)\tau_j^n \mathcal{L}_j^n\right)(\rho)\,. \tag{5.31}$$

We then aim to match the ideal channel to linear order by an appropriate choosing of τ_j

$$\sum_j q(j)\tau_j \mathcal{L}_j \stackrel{!}{=} \frac{t}{N}\sum_j h_j \mathcal{L}_j = \frac{t\lambda}{N}\sum_j \frac{h_j}{\lambda}\mathcal{L}_j = \frac{t\lambda}{N}\sum_j p_j \mathcal{L}_j\,. \tag{5.32}$$

We concentrate on the portion that is linear in time, as the bias at the second order in t/N cannot be matched with any choice of $q(j)$. In order to facilitate notation, we include the constant factor within the generators as

$$\widetilde{\mathcal{L}_j} = \frac{t\lambda}{N}\mathcal{L}_j\,. \tag{5.33}$$

The expected value of a qDrift sample is thus in the correct format and may be expressed as

$$\mathbb{E}_p\left[\widetilde{\mathcal{L}_j}\right] = \sum_j p_j \widetilde{\mathcal{L}_j} = \sum_j q_j \frac{p_j}{q_j}\widetilde{\mathcal{L}_j} = \sum_j q_j \omega_j \widetilde{\mathcal{L}_j} = \mathbb{E}_q\left[\omega_j \widetilde{\mathcal{L}_j}\right]. \tag{5.34}$$

Therefore, the channel can be written as

$$\bar{\mathcal{E}}_q(t;1)[\rho] = \sum_j q_j \exp\left(\omega_j \frac{t\lambda}{N}\mathcal{L}_j(\rho)\right) = \sum_j q_j \exp\left(\frac{t}{N}\frac{h_j}{q_j}\mathcal{L}_j(\rho)\right), \tag{5.35}$$

where we used the explicit expression for p_j, and find

$$\tau_j = \frac{t h_j}{N q_j}. \tag{5.36}$$

In contrast to the conventional qDrift, we observe that the strength of each unitary τ_j is now depending on j. The procedure is briefly summarized in Algorithm 1, which is a generalization of the standard qDrift protocol to accommodate arbitrary sampling distributions.

5.2.3 *Bias, variance and fluctuation bounds*

The bias, concentration, and fluctuation bounds for the importance sampled qDrift channel are computed in this section as a function of the sampling distribution $q(j)$, simulation time t, number of samples N, and number of experiments M. We observe that the standard bounds are recovered when $q(j) = p(j)$, indicating that our framework is a natural extension of the standard qDrift implementation. In order to facilitate notation, the minimum values of N, M, and r, which are integer numbers by definition, are expressed as functions of ϵ, t, and λ. However, these functions may not result in integer values. Therefore, in practical applications, they should be rounded up to the nearest integer. We will now establish an upper bound on the bias error by using the proof strategy from [195, Proposition 3.2]. The comprehensive proofs of the results presented are provided in Appendix A.

Theorem 2 (Bias error bound [196]). *Let $\mathcal{U}(t)$ be the unitary channel of a first-order Trotter product formula, $\overline{\mathcal{E}}_q(t;N)$ an average qDrift channel with importance sampling and $\omega(j) = p(j)/q(j)$ the re-weighting factor. The diamond norm distance between these two channels for $N = 1$ is then upper bounded by*

$$\left\|\mathcal{U}(t) - \overline{\mathcal{E}}_q(t;1)\right\|_\diamond \leq t^2\lambda^2\left(1 + \mathbb{E}_p\left[\omega(j)\right]\right), \tag{5.37}$$

leading to the following result

$$\left\|\mathcal{U}(t) - \overline{\mathcal{E}}_q(t;N)\right\|_\diamond \leq \frac{t^2\lambda^2}{N}\left(1 + \mathbb{E}_p\left[\omega(j)\right]\right) \tag{5.38}$$

for an arbitrary N.

It is now apparent that important sampling can result in an augmentation of the number of qDrift samples N while maintaining a fixed level of precision ϵ. In fact

$$N_q = \left\lceil \frac{t^2\lambda^2}{\epsilon}(1 + \mathbb{E}_p[\omega(j)]) \right\rceil \geq \left\lceil \frac{t^2\lambda^2}{\epsilon}(1+1) \right\rceil = N_p, \tag{5.39}$$

implying $N_q \geq N_p$. Hence, the conventional qDrift channel necessitates fewer samples. Nevertheless, in the following section, we will demonstrate that the overall simulation cost can still be diminished by employing importance sampling, without compromising accuracy, as we can prioritize the sampling of less expensive gates.

Having established the error constraint on the bias in a generalized manner, our next task is to comprehend the concentration of a finite importance sampled qDrift channel around its expectation value. Therefore, this will yield an approximation of the values of M and N, with a specified level of precision ϵ and simulation duration t.

Theorem 3 (Concentration bound [196]). *Let $\mathcal{E}_q(t;N,M)$ be a finite importance sampled qDrift channel on n qubits and NM instances of the V_j unitaries that compose the channel. Their concentration can then be upper bounded $\forall \epsilon \in [0, 4t\lambda]$ as follows*

$$\Pr\left[\left\|\frac{1}{M}\sum_m^M \prod_{k=N}^1 V_{j_k^m} - \mathbb{E}_q\left[V_j\right]^N\right\| \geq \epsilon/2\right] \leq 2^{n+1}\exp -\frac{NM\epsilon^2}{11t^2\lambda^2(1+\max_k \omega(k))^2}. \tag{5.40}$$

In order to guarantee an approximation error $\epsilon/2$ with probability at least $1-\delta$, we can choose

$$NM = 11\frac{t^2\lambda^2}{\epsilon^2}\left(1 + \max_k \omega(k)\right)^2 (n+1)\log\left(\frac{2}{\delta}\right). \tag{5.41}$$

This theorem tells us two essential facts. First, we learn that the resource budget may be divided between M and N, and that, for a constant ϵ accuracy level, the channel quickly converges to the deterministic qDrift with rate NM. Consequently, parallel execution of qDrift diminish the variance exponentially fast. The trade-off between the depth of the circuit and the number of measurements has certain advantages, such as the reduction of hardware errors on NISQ devices. Moreover, it suggests that the qDrift findings obtained in previous studies may be applied to any distribution $q(j)$, provided that the distributions are approximately identical, i.e., if $\omega(j) \approx 1$. This enables the generation of alternative distributions that, for example, lead to less costly circuits or are easier to obtain samples from. For instance, these results may be easily implemented in the continuous qDrift protocol [198], where the complex continuous distribution is replaced with a simpler one. Our findings aid in the identification and understanding of the changes in both bias and

variance resulting from importance sampling. In the subsequent section, we will examine the procedure of picking $q(j)$ to ensure a certain reduction in the simulation cost.

Ultimately, we calculate the limit of the expected fluctuations around the exact evolution $\mathcal{U}_H$, by employing the approach of [195, Proposition 3.4].

Corollary 1 (Fluctuation bound [196]). *Let H be a n-qubit Hamiltonian, $q(j)$ an arbitrary distribution, t the simulation time, N a fixed number of qDrift samples, and M a fixed number of qDrift experiments. Set $\mathcal{U}_H[\rho] = U_H \rho U_H^\dagger$ (with $U_H = e^{-iHt}$) and take the importance sampled qDrift channel $\mathcal{E}_q(t; N, M)$. We have*

$$\begin{aligned}\mathbb{E}\left[\| \mathcal{E}_q(t;N,M) - \mathcal{U}_H\|_\diamond\right] \leq 2\frac{t^2\lambda^2}{N}\left(1+\mathbb{E}_p[\omega]\right) \\ + \alpha\frac{nt\lambda}{NM}\left(1+\max_k \omega(k)\right) + \alpha\sqrt{\frac{n}{NM}}t\lambda\left(1+\max_k\omega(k)\right),\end{aligned} \tag{5.42}$$

with α being a numerical constant depending on H.

We now have a better comprehension of how to optimally choose N and M for a specific distribution $q(j)$ at desired accuracy ϵ. In particular, if we select N to regulate the bias toward ϵ/κ, where $\kappa > 1$, we will have

$$N = 2\kappa\frac{t^2\lambda^2}{\epsilon}(1+\mathbb{E}_p[\omega])\,. \tag{5.43}$$

We then need to choose M so that

$$\alpha t\lambda\left(1+\max_k\omega(k)\right)\left(\frac{n}{NM}+\sqrt{\frac{n}{NM}}\right) \leq \frac{\kappa-1}{\kappa}\epsilon\,. \tag{5.44}$$

Using the choice for N above, we have

$$t\lambda\left(1+\max_k\omega(k)\right) = \sqrt{N\frac{\epsilon}{2\kappa}}\frac{1+\max_k\omega(k)}{\sqrt{1+\mathbb{E}_p[\omega]}}\,, \tag{5.45}$$

resulting in the condition

$$\alpha\sqrt{\frac{n}{2M}}\frac{1+\max_k\omega(k)}{\sqrt{1+\mathbb{E}_p[\omega]}}\left(1+\sqrt{\frac{n}{NM}}\right) \leq (\kappa-1)\sqrt{\frac{\epsilon}{\kappa}}\,. \tag{5.46}$$

Since $N \geq 1$, for $M \geq n$ we obtain

$$M = \frac{n}{\epsilon}\frac{2\alpha^2\kappa}{(\kappa-1)^2}\frac{\left(1+\max_k\omega(k)\right)^2}{1+\mathbb{E}_p[\omega]}\,. \tag{5.47}$$

The parameter κ can be selected to minimize N while ensuring that M does not diverge. In general, the asymptotic scaling is as follows for a given κ and a fixed choice of distribution $q(j)$

$$N = \mathcal{O}\left(\frac{t^2\lambda^2}{\epsilon}\right), \quad M = \mathcal{O}\left(\frac{n}{\epsilon}\right)\,. \tag{5.48}$$

It becomes clear that the concentration bound in Theorem 3 allows for a more efficient scaling of the number of qDrift trials than the $\mathcal{O}(1/\epsilon^2)$ scaling that would be expected from statistical noise alone.

5.2.4 *Cost reduction*

In the preceding section, we gained insight regarding the behavior of qDrift when it is constructed from an arbitrary distribution $q(j)$. In particular, the original distribution provides the optimal distribution in terms of the number of samples. $p(j)_j/\lambda$. Nevertheless, a computational advantage can still be achieved if the various elements in the Hamiltonian have a distinct implementation cost C_j on hardware.

For instance, as proposed in Ref. [196], we can choose

$$q_c(j) = \frac{h_j}{C_j \lambda_c}, \qquad \lambda_c = \sum_l \frac{h_l}{C_l} \,. \tag{5.49}$$

This sampling distribution shall be represented by the subscript c. It is worth noting that the conventional qDrift protocol is recovered when the cost is constant, i.e. $C_j = C$ and that the objective here is to decrease the overall computing cost, rather than the variance. The choice of the cost C_j is determined by the limitations imposed by the hardware, and is especially beneficial if the cost distribution has a high variance. Using Jensen's inequality, we can prove that the predicted cost of one qDrift sample is consistently cheaper when sampling from $q_c(j)$ compared to $p(j)$. The full derivation are reported in Appendix A.2.

Corollary 2. *The expected cost of an importance sampled qDrift channel with $N = 1$ samples and $q(j) = q_c(j)$ is consistently lower than that of the standard qDrift protocol. Specifically, we have that $\mathbb{E}_{q_c}[C] \leq \mathbb{E}_p[C]$.*

Even if, on average, one qDrift sample is less expensive when sampling from this alternative distribution, this may not be sufficient to assert a reduction in the total simulation cost. This is due to the fact that the standard qDrift channel requires fewer samples at a fixed accuracy than an importance sampled one, as demonstrated by Theorem 2. Nevertheless, the cost reduction is generally valid, as demonstrated by the theorem that follows.

Theorem 4 (Cost reduction - pure qDrift [196])**.** *Let $N_p(\epsilon, t)$ and $N_{q_c}(\epsilon, t)$ denote the number of qDrift samples for the two distributions $p(j) = h_j/\lambda$. and $q_c(j) = h_j/(\lambda_c C_j)$ for a given target precision ϵ and propagation time t. Consequently, the expected cost of the important sampled qDrift channel is consistently lower than that of the standard channel.*

$$N_{q_c}(\epsilon, t)\mathbb{E}_{q_c}[C] \leq N_p(\epsilon, t)\mathbb{E}_p[C]. \tag{5.50}$$

The total number of experiments is increased to

$$M_{q_c}(\epsilon) = M_p(\epsilon)\frac{(1 + \mathbb{E}_p[1/C]\max_j C_j)^2}{1 + \mathbb{E}_p[1/C]\mathbb{E}_p[C]} \,, \tag{5.51}$$

which is specifically independent on the total evolution time t.

Consequently, it becomes apparent that the importance sampling procedure outlined in this work ensures a cost savings for any simulation that employs pure qDrift, albeit at the expense of an increase in the number of independent experiments.

5.3 COMPOSITE CHANNELS

A natural extension of the methods covered in this section is in the context of composite channels [199–201], which is a combination of a deterministic and random approach. Hence some methods may be more suitable for some parts of the Hamiltonian, and composite channels can efficiently take use of this. In accordance with [199], our approach will begin by presenting composite channels. We will then integrate these channels with our importance sampling method and calculate the corresponding limits on the bias, variance, and fluctuation.

A composite channel is a combination of multiple channels, each of which is employed to simulate a distinct component of the Hamiltonian. For the sake of simplicity, we will restrict ourself to the bipartite case, in which H can be expressed as $H = A + B$. The first part will be simulated using a deterministic Trotter formula, while the second one will be simulated using qDrift. In order to gain a more comprehensive understanding of this paradigm, we first conduct an outer first-order Trotter decomposition

$$e^{itH}\rho e^{-itH} := e^{itA}e^{itB}\rho e^{-itB}e^{-itA} + E_{A,B}(t), \tag{5.52}$$

where $E_{A,B}$ is the error term. We define $\widetilde{\mathcal{U}}_A(t)$ as an approximation of the time-evolution under A unitary channel $\mathcal{U}_A(t)[\rho] = e^{itA}\rho e^{-itA}$. In this case, we are using a first-order product formula. Additionally, we have $\mathcal{E}_q^B(t; N, M)$, which is the importance sampled qDrift channel for the B term. Let's define the composite channel as $\Lambda_q(t; N, M) = \widetilde{\mathcal{U}}_A(t) \circ \mathcal{E}_q^B(t; N, M)$, and its corresponding average channel as $\overline{\Lambda}_q(t; N)$. It can be demonstrated that the diamond norm distance between this composite channel and the ideal channel $\mathcal{U}_H$ can be upper bounded in the following manner (cf. [199]):

$$\begin{aligned}
\epsilon &:= \left\|\mathcal{U}_H(t) - \overline{\Lambda}_q(t; N_B)\right\|_\diamond && (5.53)\\
&\leq \left\|\mathcal{U}_A(t) - \widetilde{\mathcal{U}}_A(t)\right\|_\diamond + \left\|\mathcal{U}_B(t) - \overline{\mathcal{E}}_q^B(t; N_B)\right\|_\diamond + \|E_{A,B}(t)\|_\diamond && (5.54)\\
&\leq t^2\left(\sum_{i<j} a_i a_j \|[A_i, A_j]\| + \frac{1}{2}\sum_{ij} a_i b_j \|[A_i, B_j]\| + \frac{\lambda_B^2(1 + \mathbb{E}_p[\omega(j)])}{N_B}\right). && (5.55)
\end{aligned}$$

By dividing the overall duration into r intervals and use the customary union bound, we obtain

$$\begin{aligned}
&\left\|\mathcal{U}_H(t) - \overline{\Lambda}_q\left(\frac{t}{r}; N_B\right)^{\circ r}\right\|_\diamond \leq && (5.56)\\
&\frac{t^2}{r}\left(\sum_{i<j} a_i a_j \|[A_i, A_j]\| + \frac{1}{2}\sum_{ij} a_i b_j \|[A_i, B_j]\| + \frac{\lambda_B^2(1 + \mathbb{E}_p[\omega(j)])}{N_B}\right). && (5.57)
\end{aligned}$$

Except from the use of a general importance sample qDrift, and the improved bound Eq. (5.8) obtained in [195], this result was already obtain in [199]. We differ from their approach in that we consider the possibility that evolution under various terms in the

expansion could have a differing gate cost. The total cost of the composite channel can be bound as follows by denoting the cost of the term j in either A or B as C_j^A and C_j^B

$$C(\epsilon,t) \leq r\left(\sum_{l\in I^A} C_l^A + N_B \max_{j\in B}\left(C_j^B\right)\right) \tag{5.58}$$

$$\leq \left(\sum_{l\in I^A} C_l^A + N_B \max_{j\in B}\left(C_j^B\right)\right) \tag{5.59}$$

$$\times \frac{t^2}{\epsilon}\left[\sum_{i<j} a_i a_j \|[A_i,A_j]\| + \frac{1}{2}\sum_{ij} a_i b_j \|[A_i,B_j]\| + \frac{\lambda_B^2(1+\mathbb{E}_p[\omega(j)])}{N_B}\right]. \tag{5.60}$$

$$\tag{5.61}$$

The expected cost can finally be computed as

$$\mathbb{E}_q[C(\epsilon,t)] \leq \left(C_{tot}^A + N_B \mathbb{E}_q\left[C^B\right]\right)\frac{t^2}{\epsilon}\left(\Gamma_{\text{comm}}^{A,B} + \frac{\lambda_B^2\left(1+\mathbb{E}_p[\omega(j)]\right)}{N_B}\right), \tag{5.62}$$

where the total cost for A is denoted by $C_{tot}^A = \sum_{l\in A} C_l^A$, and the contribution containing the commutators among the A terms and between the A and B partitions is denoted by $\Gamma_{\text{comm}}^{A,B}$.

As mentioned before in [199], the amount of qDrift samples N_B each segment is now a fully adjustable parameter. By employing the identical approach used previously, which involves explicitly determining the smallest cost, one may get an ideal number of samples as

$$N_B = \lambda_B \sqrt{\frac{1+\mathbb{E}_p[\omega(j)]}{\mathbb{E}_q\left[C^B\right]}\frac{C_{tot}^A}{\Gamma_{\text{comm}}^{A,B}}}. \tag{5.63}$$

The expected cost therefore becomes

$$\mathbb{E}_q[C(\epsilon,t)] \leq \frac{t^2}{\epsilon}\left(\sqrt{\Gamma_{\text{comm}}^{A,B} C_{tot}^A} + \lambda_B \sqrt{\mathbb{E}_q\left[C^B\right]\left(1+\mathbb{E}_p[\omega(j)]\right)}\right)^2. \tag{5.64}$$

Unfortunately, the concentration bound in Theorem 3 cannot be directly applied to the result, as the r components of the channel $\Lambda_q(t;N,M)$ necessitate M^r experiments to be implemented. To avoid this issue, we define an additional channel

$$\Omega_q(t;N,M,r)[\rho] = \frac{1}{M}\sum_{m=1}^{M} \Lambda_q^m\left(\frac{t}{r};N,1\right)^{\circ r}[\rho], \tag{5.65}$$

obtained by combining M channels, each of which is divided into r segments. The superscript m in Λ_q^m is employed to denote that the channel employs a distinct sample of Nr indices for each experiment indexed by m. It is important to note that in the limit of infinite experiments, $\overline{\Omega}_q(t;N,r) = \overline{\Lambda}_q\left(\frac{t}{r};N_B\right)^{\circ r}$. Nevertheless, we can continue to employ Eq. (5.56) to regulate the bias. The following theorem describes the fluctuations around the average.

Theorem 5 (Concentration bound for composite channels [196]). *Let $\Omega_q(t;N,M,r)$ be a composite channel simulating the Hamiltonian $H = A + B$, where $\widetilde{U}_A \approx e^{-iAt/r}$ is used for the term A and total time t/r, and a finite importance sampled qDrift channel $\mathcal{E}_q(t;N,1)$ for*

the evolution under B and r steps using unitaries $W_j = e^{-iB_j\tau_j/r}$ *with* $\tau_j = (tb_j)/(Nq_j)$. *Its concentration is then upper bounded* $\forall \epsilon \in [0, 4t\lambda_B]$ *by*

$$\begin{aligned}&\Pr\left[\left\|\frac{1}{M}\left[\prod_{s=r}^{1}\left(\widetilde{U}_A\prod_{k=N}^{1}W_{j_{k,s}^m}\right)\right]-\left(\widetilde{U}_A\mathbb{E}\left[W\right]^N\right)^r\right\|\geq\frac{\epsilon}{2}\right]\\ \leq& 2^{n+1}\exp-\frac{NMr\epsilon^2}{11t^2\lambda_B^2(1+\max_k\omega(k))^2}\,.\end{aligned}\tag{5.66}$$

In order to guarantee an approximation error $\epsilon/2$ *with probability at least* $1-\delta$, *it is then sufficient to take*

$$NMr = 11\frac{t^2\lambda_B^2}{\epsilon^2}\left(1+\max_k\omega(k)\right)^2(n+1)\log\left(\frac{2}{\delta}\right)\,.\tag{5.67}$$

Using the new version of the concentration bound Theorem 5, we can also provide an estimate for the expected error of the composite channel.

Corollary 3 (Fluctuation bound for composite channels [196]). *Given* $\mathcal{H} = A + B$ *a bipartite Hamiltonian, an arbitrary distribution* $q(j)$, *a simulation time* t, *a fixed number of qDrift samples* N *and experiments* M, *consider the first-order Trotter approximation of the channel,* $\widetilde{\mathcal{U}}_A(t)$, *as* $\mathcal{U}_H(t)[\rho] = U_H(t)\rho U_H^\dagger(t)$, *where* $U_H(t) = e^{-iHt}$. $\mathcal{U}_A(t)$, $\mathcal{E}_q^B(t;N,M)$ *the importance sampled qDrift channel for the B term and* $\Omega_q(t;N,M,r)$ *the importance sampled composite channel. Then, we have*

$$\begin{aligned}\mathbb{E}\left[\|\ \Omega_q(t;N,M,r)-\mathcal{U}_H\|_\diamond\right]\leq 2\frac{t^2}{r}&\left(\Gamma_{comm}^{A,B}+\frac{\lambda_B^2}{N}\left(1+\mathbb{E}_p[\omega]\right)\right)\\ &+\alpha\frac{nt\lambda_B}{NMr}\left(1+\max_k\omega(k)\right)\\ &+\alpha\sqrt{\frac{n}{NMr}}t\lambda_B\left(1+\max_k\omega(k)\right)\,,\end{aligned}\tag{5.68}$$

where the parameter

$$\Gamma_{comm}^{A,B} = \sum_{i<j}a_ia_j\|[A_i,A_j]\| + \frac{1}{2}\sum_{ij}a_ib_j\|[A_i,B_j]\|\,,\tag{5.69}$$

contains the dependence on commutators.

To guarantee that the predicted error is smaller than ϵ, we require a number of steps of

$$r = 2\kappa\frac{t^2}{\epsilon}\left(\Gamma_{\text{comm}}^{A,B}+\frac{\lambda_B^2}{N}\left(1+\mathbb{E}_p[\omega]\right)\right)\,,\tag{5.70}$$

for some $\kappa > 1$ together with $N = N_B$ from Eq. (5.63) and a number of experiments given by

$$M = \mu_q\frac{n}{\epsilon}\frac{2\alpha^2\kappa}{(\kappa-1)^2}\frac{\lambda_B}{\sqrt{C_{tot}^A}}\,.\tag{5.71}$$

We note that the dependence on $q(j)$ has been absorbed in

$$\mu_q = \frac{\left(1+\max_k\omega(k)\right)^2\sqrt{\mathbb{E}_q\left[C^B\right]/\left(1+\mathbb{E}_p[\omega]\right)}}{\sqrt{\Gamma_{\text{comm}}^{A,B}C_{tot}^A}+\lambda_B\sqrt{\left(1+\mathbb{E}_p[\omega]\right)\mathbb{E}_q\left[C^B\right]}}.\tag{5.72}$$

generator	cost	generator	cost	generator	cost	generator	cost
X_k	0.1	Z_2Z_3	2	Z_1Z_2	6	$Z_1Z_2Z_3$	8
Z_k	0.1	Z_2Z_4	2	$Z_1Z_2Z_4$	4	Z_3Z_4	6
Z_1Z_4	2	$Z_1Z_3Z_4$	6	$Z_2Z_3Z_4$	4	$Z_1Z_2Z_3Z_4$	8
Z_1Z_3	10						

Table 5.1: Implementation cost for the different generators appearing in the two considered Hamiltonians.

A comparable derivation to the one conducted in this context could be applied to more extensive scenarios, as already discussed in [199], when utilizing a higher-order Trotter formula for the outer decomposition in (5.52) or in the simulation of the evolution under A, as well as protocols where the latter is implemented using qubitization. Additional enhancements achieved via the implementation of randomization techniques [195, 202, 203] might also be incorporated.

5.4 NUMERICAL ANALYSIS

To illustrate the efficiency of the cost reduction, we simulate the nuclear EFT Hamiltonian introduced in Sec. 4.3, and compare the diamond norm and the corresponding cost for the standard and important sampled composite qDrift channel. We recall that the Hamiltonian, when expressed in the Pauli basis, takes the following form

$$\begin{aligned} H = 8T + \frac{3u}{4} + \frac{v}{16} - 2T\sum_{i=1}^{4} X_i + \frac{v}{16}\left(\sum_{i<j<k} Z_iZ_jZ_k + Z_1Z_4 + Z_2Z_3\right) \\ + \left(\frac{u}{4} + \frac{v}{16}\right)\cdot\left(\sum_i Z_i + Z_1Z_2 + Z_1Z_3 + Z_2Z_4 + Z_3Z_4 + Z_1Z_2Z_3Z_4\right). \end{aligned} \tag{5.73}$$

We assign a fixed cost to each term in the Hamiltonian, equal to the number of CNOT after tranpilation with linear connectivity. For the single-qubit rotation, we instead choose a cost of 0.1, motivated by the relative fidelity between single and two qubits gates on current quantum hardware. The cost of each terms is summarised in Table 5.1.

To better control the error bounds, we choose uniform coefficients and vary the relative strength between the partitions

$$H^{(j)} = a\sum_{i\in I^A} A_i^{(j)} + b\sum_{i\in I^B} B_i^{(j)}, \tag{5.74}$$

where the subscript j denotes different splitting strategies and I^X a set of indices for X. We consider two different partitions. For the first one, strategy motivated by perturbation theory. In this approach, A represents the simplified model H in the $u = -4v$ configuration, where the majority of the coefficients are zero, and B denotes a minor perturbation that provides a more realistic representation of the system. The first is defined through the following separate contributions

$$A^{(0)} = \sum_{k=1}^{4} X_k + Z_1Z_4 + Z_2Z_3 + \sum_{1\le i<j<k\le 4} Z_iZ_jZ_k \tag{5.75}$$

	$\sum_{i\in I^A} C_i$	$\sum_{i\in I^B} C_i$	$\mathbb{E}_{q_c}[\omega^B]$	$\mathbb{E}_p[C^B]$	$\mathbb{E}_{q_c}[C^B]$	$N_p\mathbb{E}_p[C^B]\cdot\frac{\epsilon}{t^2\lambda^2}$	$N_{q_c}\mathbb{E}_{q_c}[C^B]\cdot\frac{\epsilon}{t^2\lambda^2}$
model 0	28.4	30.4	15.43	3.38	0.22	6.76	3.6
model 1	10.5	48.3	15.02	4.83	0.32	9.66	5.15

Table 5.2: Expectation value of the cost of simulating the two terms in the partition A and B for the two different models $j = 0$, $j = 1$, using either first order Trotterization or (importance sampled) qDrift.

$$B^{(0)} = \sum_{k=1}^{4} Z_k + Z_1Z_2 + Z_1Z_3 + Z_2Z_4 + Z_3Z_4 + Z_1Z_2Z_3Z_4 \,. \tag{5.76}$$

The second partition is chosen as to reduce the expected cost of each circuit, as a way to show case the utility of importance sampling. Part A comprises terms that necessitate minimal resources for simulation on quantum devices, characterized by a low cost C in terms of the number of two-qubit CNOT gates, whereas part B consists of the most challenging terms while still incorporating some straightforward single-qubit rotations to reduce the expected cost. We have

$$A^{(1)} = \sum_{k=1}^{4} X_k + Z_1 + Z_1Z_4 + Z_2Z_3 + Z_2Z_4 + Z_1Z_2Z_4 \tag{5.77}$$

$$B^{(1)} = \sum_{k=2}^{4} Z_k + Z_1Z_2 + Z_1Z_3 + Z_3Z_4 + Z_1Z_2Z_3 + Z_2Z_3Z_4 + Z_1Z_3Z_4 + Z_1Z_2Z_3Z_4. \tag{5.78}$$

The expected costs, on the two partitions, for Trotter and (importance sampled) composite channels, are presented in Table 5.2. We calculate the deterministic cost, the expected cost per step, and the total cost at a certain precision ϵ and time t. Importance sampling can reduce the overall simulation cost by a factor of two relative to plain qDrift, while the implementation of the qDrift channel facilitates a cost reduction by an order of magnitude compared to Trotterization.

We drive quantum simulations of the models $j = 0$ and $j = 1$ on a noiseless simulator for two distinct simulation time, $t = 0.05a$ and $t = 0.1a$, and various values of the strength coefficient $b \in [0.005, 0.5]$, employing a composite channel $\Omega_q(t;\, N = 1,\, M,\, r = 1)$. Part A is evolved using a single iteration of the first-order Trotter formula, whilst part B is evolved employing a qDrift channel with $N = 1$ step. We examine two distinct sampling distributions: the usual one $q = p$ and $q = q_c$, with the cost defined as previously stated. The diamond norm is shown in Figure 5.1 as a function of b/a, with the upper line representing calculations for model $j = 0$ and the lower line for model $j = 1$. The diamond distance to the optimal simulation is represented by a dashed black line for the complete Trotterized channel, a red line for the deterministic qDrift $\overline{\mathcal{E}}(t;\, N = 1)$ channel, a dashed blue line for the standard ($q = p$) composite channel with $M = 1$, a dashed yellow line for the standard composite channel with $M = 10$, and violet dots for the importance sampled ($q = q_c$) composite channel with $M = 1$, alongside green dots for the same channel with $M = 10$.

The errors of the various composite channels are similar and align with the pure Trotter channel, except for $b \approx 0.05$, when Trotterization exhibits a marginal advantage. This

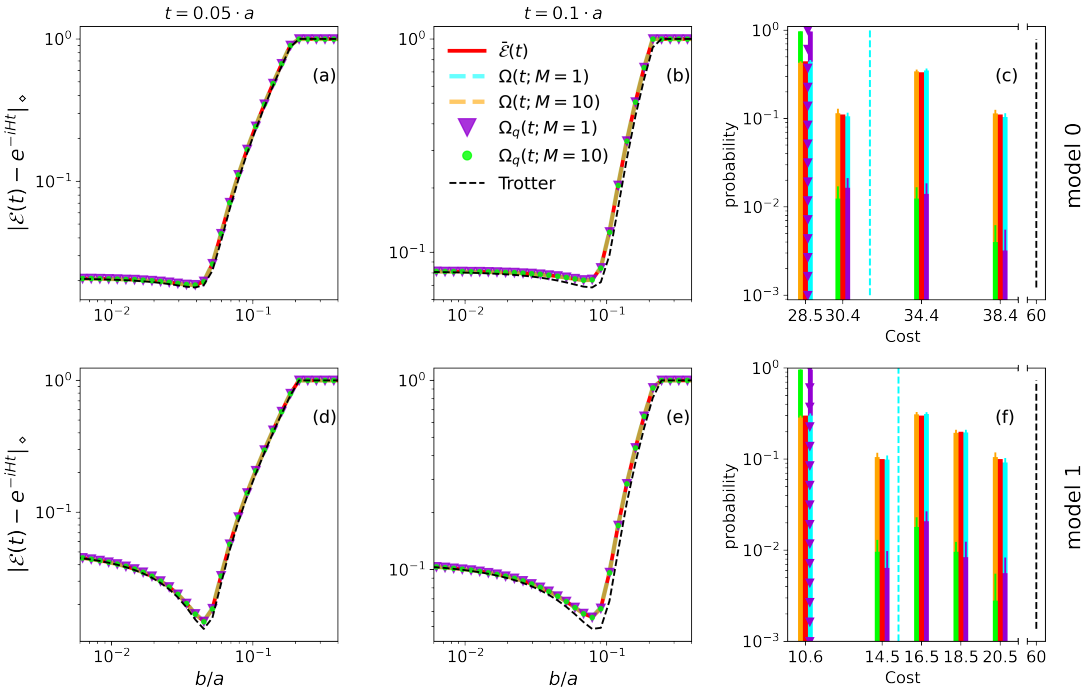

Figure 5.1: **Cost comparison.** [Left] Diamond norm against the coefficient ratio b/a and histograms [right] of the cost of the quantum circuits in terms of CNOT gates for the different composite channels: deterministic (red), standard (dashed lines) and importance sampled (dots). Figure taken from [196].

indicates that the various channels possess comparable accuracy, allowing us to examine the costs associated with their implementation, as seen in panels (c) and (f) in the last column. We provide the histograms of the cost across $R = 50$ qDrift channels derived from the sampled circuits utilizing a composite qDrift channel (importance sampled) with $M = 1$ or $M = 10$ experiments, with error bars representing a 95% confidence interval calculated from 50 distinct cost histograms. To enhance readability, contiguous bars of varying colors, positioned with a little offset along the cost axis, share identical cost labels beneath them. Furthermore, we see that only a specific number of different bars are occupied, corresponding to each potential cost inside the chosen partition. Furthermore, the dashed and dotted lines represent the expected value of the standard and importance-sampled qDrift channel, respectively, which remains consistent for both $M = 1$ and $M = 10$. This highlights that the cost of the importance-sampled channel approximates the predicted cost, but the cost of the conventional channel exhibits greater heterogeneity. The fundamental figure of merit is the allocation of costs. Consequently, the significance of the sampled qDrift mostly consists of low-cost circuits, since the likelihood of selecting an inexpensive circuit is elevated, but the conventional qDrift samples terms irrespective of cost, resulting in a more uniform cost distribution. Significantly, we achieve a decrease in the necessary resource allocation by a factor of 1.8 (2) for model $j = 0$ utilizing the (importance sampled) composite channel, and 3.8 (5) for model $j = 1$, without compromising accuracy, in comparison to the complete Trotterization. The cost for simulating the B component is diminished by one (or two) orders of magnitude when using the pure (or importance sampled) qDrift channel, in contrast to direct first-order Trotterization, as indicated in Table 5.2.

5.5 DISCUSSIONS

This study extends previous findings on concentration boundaries of random product formulas, to arbitrary sampling distribution and average of qDrift channels. We study the protocol's bias and concentration, building upon prior works. Specifically, we demonstrated that a qDrift channel converges exponentially rapidly in NM to its expectation value, where N represents the number of qDrift samples and M the number of experiments, thereby facilitating an effective distribution between circuit's depth (governed by N) and number of samples (governed by M). These results shows that it is efficient to perform simultaneous simulations of a qDrift channel. We can therefore maintain each circuit sufficiently shallow to reduce noise and execution time. Furthermore, by integrating the specific implementation cost for evolution associated with each Hamiltonian term within an appropriate sampling distribution $q_c(j)$, we demonstrate that importance sampling ensures a definitive reduction in total cost, such as the quantity of CNOT gates, thereby facilitating a more efficient implementation of qDrift in the near term. Under plausible assumptions, comparable cost reductions may also be achieved while aiming to decrease the number of T gates, which is advantageous for error-corrected devices.

Furthermore, we broaden our findings to encompass composite channels [199], whereby the Hamiltonian is divided into A and B, which are individually simulated using a deterministic approach (such as a Trotter-Suzuki Product Formula) and a qDrift channel, respectively. We demonstrate that the same significance sampling distribution $q_c(j)$ may be utilized in both instances to diminish the quantum resources necessary for implementation. In a standard scenario where evolution under varying conditions in the total Hamiltonian entails distinct implementation costs, the explicit incorporation of this information in our framework allows for the potential enhancement of savings attainable through composite channels by optimizing partitioning strategies [199]. Generally, it may be advantageous to manage Hamiltonian terms, which are costly yet possess modest norms, in a stochastic manner with qDrift.

The theoretical findings are demonstrated with numerical simulations of a basic triton model on a (2×2) lattice in first quantization. A substantial cost decrease $(5\times)$ may be achieved with the utilization of composite channels and importance sampling. The methodology is resilient over varying strength levels of the partition's two components. The quadratic scaling of the qDrift component with simulation time renders the protocol especially advantageous for applications necessitating relatively brief evolution durations, such as protocols for measuring observables through QSP [204].

The important sampling technique introduced here is advantageous due to its simplicity and guaranteed cost reduction. Nevertheless, further research on more direct numerical optimizations of the sampling distribution and the partitioning strategy are required for the construction of optimal composite channels.

6 QUANTUM SIMULATIONS IN NUCLEAR PHYSICS

In the previous chapter, we discussed how to simulate the dynamics governed by a Hamiltonian using deterministic and random PF. While we focused on theoretical aspects, we shall now exemplify the developed techniques on two different applications in nuclear physics: neutrinos thermalization and quenching across a QPT.

6.1 NEUTRINO FLAVOR OSCILLATION

In astronomical contexts characterized by high temperature and density, such as core collapse supernovae and binary neutron star mergers, the emission of neutrinos in significant fluxes is crucial for dynamic and chemical development. Electron flavor neutrinos and antineutrinos interact weakly with nearby nucleons, affecting the proton-to-neutron ratio. A thorough understanding of neutrino flavor composition is crucial for in-depth studies of the development of these systems [205–208]. In dense environments, neutrinos may undergo coherent forward scattering, a phenomenon acutely responsive to their quantum mechanical flavor states. Flavor-dependent relative phases may occur during scattering on charged leptons, resulting in flavor exchange with other neutrinos via neutral current coherent forward scattering. In situations with a sufficiently elevated neutrino number density, this flavor exchange effect is anticipated to predominate the flavor development of the neutrino gas, resulting in innovative coherent phenomena such as flavor spectrum splits and swaps [206, 209, 210].

In the following, we follow Ref. [211] and examine the flavor evolution of a neutrino gas under the assumption that the potential from coherent $\nu - \nu$ forward scattering is the predominant influence. This situation is often designated in the literature as the *fast* flavor oscillation limit [212–217]. The redistribution of flavor among neutrinos is expected to occur at length scales of around one meter. Our investigation focuses on this particular regime, seeking to thoroughly examine the related dynamics. We investigate this regime using two methods: converged deterministic PF, and a qDrift-inspired stochastic channel.

The coherent forward scattering potential for neutrinos exhibits an all-to-all Heisenberg-like interaction, which is expected to be non-integrable in general. The Hamiltonian that dictates the coherent evolution of neutrino flavor content comprises three components [218]: the vacuum potential, arising from the linear combinations of neutrino mass states and flavor states; the matter potential [219, 220], produced by coherent forward scattering via charged current interactions with local charged leptons; and the $\nu - \nu$ potential, which results from neutral current coherent forward scattering among neutrinos [205, 221]. We assume that the vacuum oscillation potential is negligible compared to the other two potentials. We also assume that the matter profile is homogeneous in the relevant environmental areas, enabling us to choose a corotating frame to eliminate its overall influence on each neutrino. Subsequent to these adjustments, only the $\nu - \nu$ coherent forward scattering Hamiltonian persists. We postulate axial symmetry for the velocity coupling term, ensuring that the velocity components perpendicular to the momentum symmetry axis (designated as z) average to zero.

This leads to the following Hamiltonian for n neutrinos

$$\hat{H}_{\nu\nu} = \frac{\mu}{2n}\sum_{i<j}(1-\vec{v}_i\cdot\vec{v}_j)\hat{\vec{\sigma}}_i\cdot\hat{\vec{\sigma}}_j \equiv H_0 + \frac{\mu}{n}\sum_{i<j}h_{i,j}P_{i,j}. \tag{6.1}$$

Here, $\hat{\vec{\sigma}}$ is the usual vector of Pauli operators, $P_{i,j}$ a swap gate, H_0 an irrelevant constant and μ is the scale of the $\nu-\nu$ interaction. Since we are only considering two neutrino families, we can directly map the physical Hamiltonian to a qubit Hamiltonian by identifying the $|0\rangle$ state with the first family and $|1\rangle$ with the second. If a third family would be present, we could either use qutrit or two qubits to encode them. Given a typical core-collapse supernova flux at a radius of 50 km, a simple estimate gives $\mu \approx 1\,\mathrm{cm}^{-1}$. As it is the only dimensionful parameter in the problem, we hereafter measure all distances and times in units of μ and set $\mu = 1$. We consider a forward peak distribution with

$$\begin{aligned} v_x &\sim \sqrt{1-v_z^2}\cos\theta \\ v_y &\sim \sqrt{1-v_z^2}\sin\theta \\ v_z &\sim 1 - \frac{2}{\sqrt{\pi}\sigma}\Theta(x)e^{-x^2/2\sigma^2}, \end{aligned} \tag{6.2}$$

where variance $\sigma = 1/2$ and $\Theta(x)$ the Heaviside step function. This choice of parameters describe neutrino evolving along one axis, with random fluctuations in the plane parallel to the main direction.

6.1.1 *Thermalization*

Non integrable many-body Hamiltonians are generally expected to conform to the Eigenstate Thermalization Hypothesis (ETH) [222–224]; see Ref. [225] for a comprehensive review. In systems characterized by all-to-all contacts or without a distinct lattice structure, we refer to Refs. [226, 227] as illustrative examples. With the exceptions of integrable systems, the ETH may only apply to the majority of eigenstates. For generic Hamiltonain, specific eigenstates, known as quantum scars, may possess supplementary symmetries beyond those of the Hamiltonian, hence rendering the ETH inapplicable [228, 229]. Nevertheless, these states are anticipated to possess measure zero in the density of states at the thermodynamic limit, unless there is specific symmetry protection. Incorporating extra symmetry into the Hamiltonian, without achieving integrability, leads to the fragmentation of the Hilbert space into several block-diagonal subsectors, each of which is insufficiently large for the ETH [230].

The ETH fundamentally characterizes the structure of "few-body" operators that act on a subextensive portion of local Hilbert spaces, which are combined via tensor product to form the full many-body Hilbert space [211]. For a few-body operator $\hat{O}$, the matrix element between two eigenstates $|E_\alpha\rangle$ and $|E_\beta\rangle$ of the Hamiltonian is conjectured to follow the ansatz,

$$\begin{aligned} \langle E_\alpha|\hat{O}|E_\beta\rangle &= \hat{O}(E_{\mathrm{ave}})\delta_{\alpha\beta} + e^{-S(E_{\mathrm{ave}})/2}f_{\hat{O}}(E_{\mathrm{ave}},\Delta_{\alpha\beta})R_{\alpha\beta} \\ S(E_{\mathrm{ave}}) &= \ln\sum_\alpha E_\alpha\delta\left(\epsilon(E_{\mathrm{ave}}-E_\alpha)\right), \end{aligned} \tag{6.3}$$

where $E_{\mathrm{ave}} = \frac{E_\alpha+E_\beta}{2}$, and $\Delta_{\alpha\beta} = E_\alpha - E_\beta$. The microcanonical average of the operator $\hat{O}$ over states close to E_{ave} is denoted as $\hat{O}(E_{\mathrm{ave}})$, where S represents the corresponding

microcanonical entropy, which counts the number of eigenstates around E_{ave} within a small interval defined by ϵ. The term $R_{\alpha\beta}$ is an independent random variable for each pair α, β, with an expected value of zero and variance of one. In large-dimensional Hilbert spaces, we expect S to be large, even for small windows, so in the thermodynamic limit, where $\epsilon \to 0$, the matrix elements of the operator in eigenstates are predominantly determined by the microcanonical average. The function $f_{\hat{O}}$ relates to the linear response correlation function of the thermodynamic system when perturbed by the operator $\hat{O}$ [230], and for a fixed E_{ave}, it decreases exponentially with respect to the energy difference $\omega_{\alpha\beta}$ in chaotic systems. In general, the ETH ansatz (6.3) for matrix elements is expected to hold in chaotic systems but not in integrable ones.

6.1.2 *Simulation using PF*

We start our investigation with a first order product formula using a SWAP network and small time steps such as the trotterized evolution is a proxy for the exact dynamics. The SWAP network [231, 232] is a protocol to perform non local operations, when only next-neighbor connectivity is, and is optimal for a fully connected graph. It involves swapping every qubits after letting them interact. After $n/2$ layers of even-odd and odd-even interaction, all qubit will have interacted exactly once, and the ordering will be inverted. We can express a full Trotter step in terms of nearest neighbor partial swaps if we add a full swap for every pair, leading to

$$\begin{aligned} U_T &= \prod_{l=1}^{r}\prod_{k=1}^{n/2}\left[\prod_{j\ even} e^{-it\frac{\mu}{Nr}h_{j,j+1}P_{j,j+1}}P_{j,j+1}\right]\left[\prod_{j\ odd} e^{-it\frac{\mu}{Nr}h_{j,j+1}P_{j,j+1}}P_{j,j+1}\right] \\ &= \prod_{l=1}^{r}\prod_{k=1}^{n/2}\left[\prod_{j\ even} e^{-it\frac{\mu}{Nr}(h_{j,j+1}+\frac{\pi}{2})P_{j,j+1}}\right]\left[\prod_{j\ odd} e^{-it\frac{\mu}{Nr}(h_{j,j+1}+\frac{\pi}{2})P_{j,j+1}}\right], \end{aligned} \tag{6.4}$$

with r the number of Trotter steps. The initial state is chosen as the wall state configuration $|\Phi\rangle = |0\rangle^{\otimes n/2} \otimes |1\rangle^{\otimes n/2}$, which correspond to all of the neutrinos to be in the first family and the rest in the second. Note that due to permutation symmetry the ordering does not have an effect on the physics, but it will have one on the algorithmic of the random channel later on. We consider three different observables

1. $Z_j(t) = \langle\Psi|\mathcal{U}(t)Z_j\mathcal{U}(t)^\dagger|\Psi\rangle$ (single-qubit observable)
2. $L(t) = |\langle\Psi|\mathcal{U}(t)|\Psi\rangle|^2$ (Loschmidt echo)
3. $S_j(t) = -\text{Tr}\left[\rho(t)\log\rho_j(t)\right]$ with $\rho(t) = |\Psi(t)\rangle\langle\Psi(t)|$ (single-qubit entropy)
4. $V(t) = \frac{1}{n}\sum_j Z_j(t)^2$ (variance),

which will help us to study thermalization. It is already known that the model from (6.1) thermalizes [211]: the observable converge to the prediction from the canonical microcanonical ensemble. Guided by Ref. [211], we define the characteristic time as the time required to reach the first local optimum (in the case of 1. and 2.), to reach 0.95 (for 3.) and 0.1 (for 4.). We currently care about understanding the scaling of this characteristic time with the system size, which could give us insights into the large scale dynamics. As en example, we display the evolution for these four observable in a single realization of the Hamiltonian in (6.1) and $n = 20$ in Figure 6.1.

Panel a) shows the thermalization of the single-qubit observables, with the inset displaying all possible qubits, while panel b) shows the entropy (violet), Loschmit echo (orange)

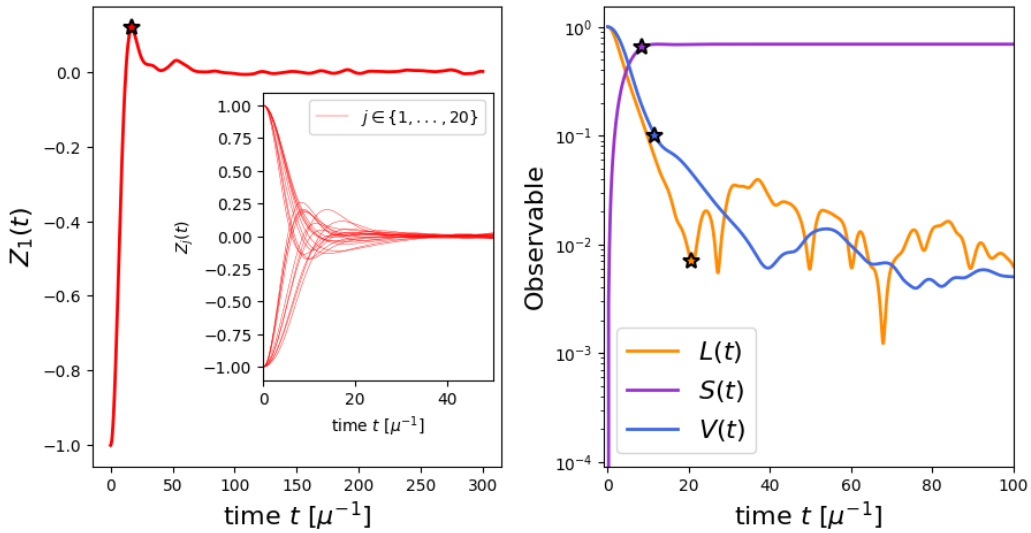

Figure 6.1: **Exact flavor evolution** $(n = 20)$**.** The flavor evolution simulated with a converged PF is shown for [left] the single-qubit observable, and [right] the Loschmidt echo, single-qubit entropy and variance. The stars denote the characteristic time.

and Variance (blue). While the three first observables have already been studies and are known to exhibits sign of thermalization, it is unclear if $V(t)$ behaves in the same way. To better understand this aspect, we show the scaling of the characteristic time for the four observables as a function of the system size in Figure 6.2. The dynamics is performed using a Trotter decomposition with a time step $\delta t = 0.1$, such that the error is at most 0.01, which is negligible at this scale. The errorbars correspond to the interquartile range, meaning that it contains 50% of every possible configuration. While this might appear as a large interval, it is required because of the presence of outliers, especially for the single-qubit observable $Z_j(t)$. We can see on the plot that the variance behaves similarly to the single-qubit entropy up to $n = 28$ neutrinos, and that we can be hopeful that this trend will persist. Hence, since we start in the wall state configuration, the single qubit entropy only depends on $Z_j(t)$ as the variance.

From those small size experiments, it is difficult to extract the scaling of the characteristic time with the system sizes, as even a constant function would be compatible within the errorbars. It is therefore crucial to perform larger scale simulations to better understand the scaling. However, since the size of a Trotter step grows linearly with n, the task becomes increasingly difficult, and brute force Trotterization remains out of reach of current quantum hardware, but also state of the art MPS simulator. For this reason, we will propose, in the next section, a new approximate method which can empirically reproduce the variance, at a cost of the order $\mathcal{O}(t)$.

6.1.3 *Simulation using a random dynamics*

Simulating neutrinos is a difficult task considering the number of interaction. Hence, the depth grows at least linearly with the number of neutrinos [233], which makes it difficult for near term quantum computer. For this reason, we aim at constructing a black box, with reasonable depth, acting as close as possible to the real evolution. The black box consists of brick wall quantum circuit, whose building blocks are Heisenberg evolution

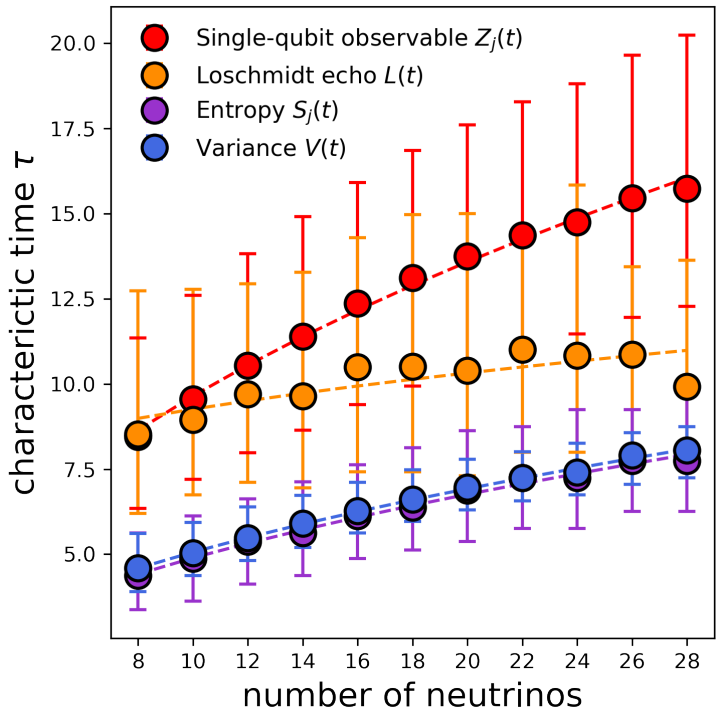

Figure 6.2: Scaling of the characteristic time for the single-qubit (red), Loschmit echo (orange), entropie (violet) and variance (blue) as a function of the system size using a Trotterized evolution with step size $\delta t = 0.2$.

and a swap, acting on an input state. The angles of the evolution are chosen uniformly from the Hamiltonian's coefficient.

In the construction of the trotterized evolution (6.4), the subdivision of these unitary into blocks of size $n/2$ is somewhat artificial if we only look at ensemble properties of one body observables, since we can consider the final unitary as layers of random partial swaps stacked together. Loosing track of each qubit's label, does not have any effect if we consider observables of the form

$$O_{ens} = \frac{1}{n}\sum_{i=1}^{n} F\left[Tr[O_i\rho]\right] \ , \tag{6.5}$$

since these are invariant for label permutations. The full dynamics, at least at the level of averaged one-qubit observables, can be seen as being approximated by products of individual even-odd layers with random partial swaps

$$U_{\text{rdm}}(L) = \prod_{k=1}^{L}\left[\prod_{j\ even} e^{-i\left(\alpha\widetilde{h}_{jk}+\frac{\pi}{2}\right)P_{j,j+1}}\right]\left[\prod_{j\ odd} e^{-i\left(\alpha\widetilde{h}_{jk}+\frac{\pi}{2}\right)P_{j,j+1}}\right] \tag{6.6}$$

where we only require the *time-step* α to be small so that this reproduces the Trotter decomposition, and the coupling constants h_{jk} to have the same distribution of the original $h_{i,j}$. The dependence on n is not important because for a small μt all the prefactor does is to guarantee a step of $\mathcal{O}(1)$ after the $\mathcal{O}(n)$ interactions in a step. For a given choice of step α the time evolution is mapped into a discrete evolution in layers. In order to map this correctly into the desired time evolution for a system of n qubits, we can use a global scale factor that maps $\alpha L \to t$ such as

$$O_{ens}(t) \approx O_{ens}(\tau_n \alpha L) \tag{6.7}$$

and the prefactor τ_n can be fixed looking at the short time dynamics, which can be done either analytically or numerically in an efficient way. In addition, since the exact dynamics is invariant for any permutation of the initial state, while the random dynamics does not need to be, we will also sample equivalent initial state by applying permutations.

To fix the value of α, we analytical compute the first two derivatives at time $t = 0$. Since the initial state is a product state, we can explicitly show that the first derivative is zero. The second derivative can be computed as

$$g \equiv \frac{d^2}{dt^2} V(t)|_{t=0} = \frac{2}{n} \sum_{ij} \mu_{ij}^2 \left(1 - \langle Z_i \rangle \langle Z_j \rangle\right) . \tag{6.8}$$

The value of α is then computed for each configuration of the random channel as $\alpha = \sqrt{g_{\text{ex}}/g_{\text{rdm}}}$. We note that the derivative can be computed analytically from the coefficient only. While it is straightforward in the exact case, it requires some more thoughts in the random one. The idea is to see the random channel as one trotter step of an effective Hamiltonian, extract its coefficient and compute the derivative using (6.8). In summary, we are build our random channel to respect most of what we know about the exact evolution, e.g. permutation symmetry, particle number conserving, short term dynamic and randomize over the rest.

While we have physical understanding of why our construction make sens, we are lacking a theoretical proof and performance guarantees. We can not use the qDRIFT formalism introduced in the previous section as we are not converging towards the target density matrices, but are only able to reproduce global expectation values. We therefore empirically demonstrate in Figure 6.3 for $n = 20$ neutrinos that our randomized approach can replicate global properties of the exact evolution, up to intermediate time. Panel a shows that the exact evolution (black line), is comprised between the first and third inter quartile range of the randomized approach. The grey curves shows the results obtained with a cost equivalent PF, i.e. where the depth is the same as the randomized approach. We can appreciate that at cost equivalent the random approach is more accurate. In the second panel, we show the dependence on the number of layers, hinting that we require $L = \mathcal{O}(t)$. Explicitly, the different colors show the maximal number of layer L, leading to different time step $\delta t = t/L$. Only the choice of $\delta t = 1$ is consistent with the exact dynamic. This yields $L = t_{\text{max}}$. We note that increasing the number of layers does not converge better, and in fact we see that for $L = 100$, the predictions significant deviates at long time. We can expect this behavior as we are effectively converging towards the evolution of the average Hamiltonian, which can be shown using a construction a la qDRIFT. We note that it does not break any complexity theory assumption, since we expect the cost to be at least linear in time.

We are now in a position to simulate larger models, both using quantum devices and MPS simulators.

SIMULATION USING MPS For simulations using MPS, we employ `quimb` to simulate quantum circuits at fixed bond dimension χ, i.e. the number of singular values kept during the contraction. This approach is particularly efficient when the system's entanglement is limited, as the bond dimension χ constrains the maximum entanglement that can be represented. However, in this simulation, entanglement grows linearly with the circuit depth, as interactions occur only between neighboring qubits in each layer. Consequently, this poses a significant challenge for MPS. Nevertheless, for system sizes up to 100 spins, we can brute-force the problem by choosing high bond dimensions $\chi \in \{2^8, 2^9, 2^{10}, 2^{11}\}$ and extrapolating to the limit $\chi \to \infty$ using an exponential fit, $f(\chi) = a + 2^{-b \cdot \chi^2}$, where

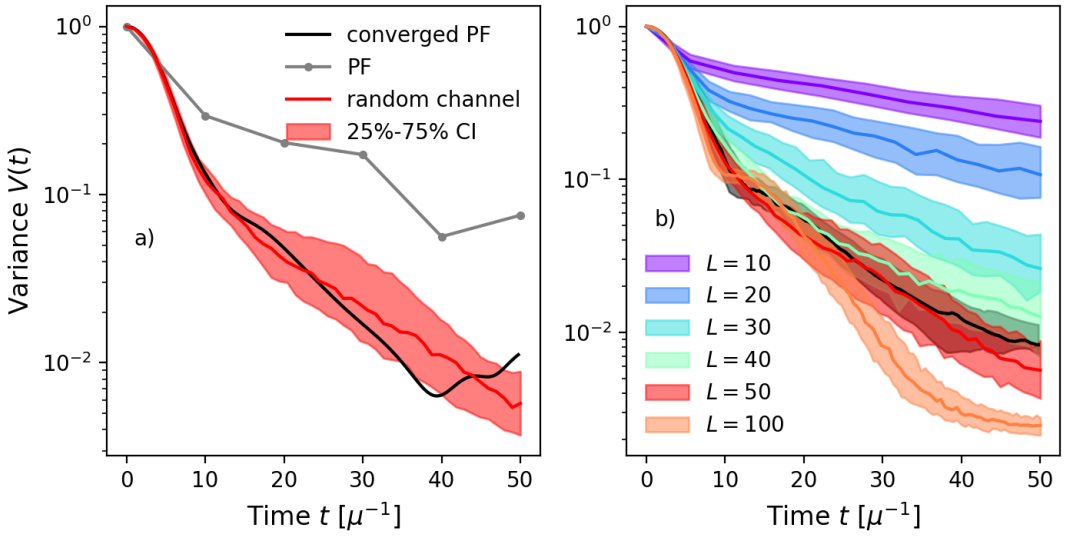

Figure 6.3: **Empirical validation of the random channel.** The variance of the single qubit observable is shown as a function of time. The black line displays the exact dynamic computed with a converged product formula, while the colored one shows the random channel with a band corresponding to the inter-quartile range. a) The left panel motivates that the random channel reproduce the exact dynamic. The gray curves represent a Trotter product formula with the same number of gates as the random channel, emphasizing that randomness is a key ingredient to reduce the number of operators. B) The right panels show the difference in the random channel when changing the number of layers. It should be noted that L in the legend has to be interpreted as the number of steps, and that the channels are built by adding layers over time.

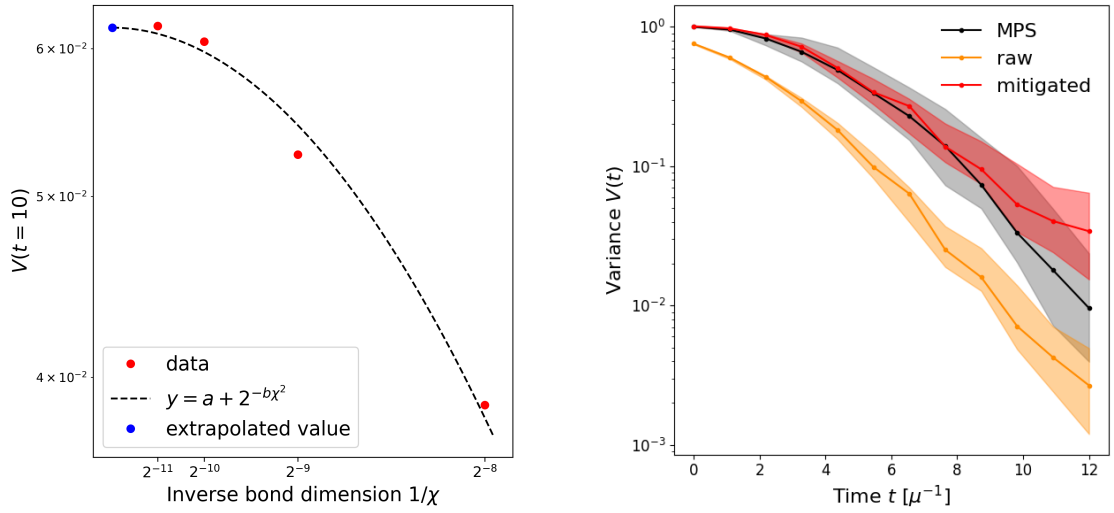

(a) MPS (b) ibm_brisbane

Figure 6.4: **Variance of the flavor of $n = 32$ interacting neutrinos on MPS simulation and IBM quantum hardware.** [Left] The variance (red) at $t = 10$ computed with MPS are shown as a function of the inverse bond dimension. The dotted line denotes an exponential fit and the blue dot is the extrapolated value at infinite bond dimension. [Right] The raw data are shown in orange, and the mitigated one in red, while the reference solution obtained with MPS is shown in black. The dashed area correspond to the inter-quartile range over 20 different Hamiltonian.

the coefficients a and b are determined via least squares regression. As an example, we show in Figure 6.4 [left] the extrapolation procedure for $n = 32$ spins and $t = L = 10$. We understand that this calculation is challenging for MPS as entanglement builds rapidly. Hence, scaling to over one hundred sites becomes increasingly difficult, as the quality of the fit deteriorates and the results become less reliable. We cannot further increase the bond dimension due to computational limitations, which is why we need to perform the same calculation on a quantum computer.

SIMULATION ON A QUANTUM COMPUTER We now perform the same calculation on the 127-qubit chip `ibm_brisbane`. The Heisenberg interaction is compiled using three echoed cross-resonance gates $e^{iX\otimes Z}$, which are equivalent to CNOT gates up to single-qubit rotations. Our goal is to compute the expectation values of a family of random circuits. These circuits are naturally self-twirling, making them well-suited for our NR error mitigation strategy.

To estimate the noise parameter, we evaluate the same circuit on the all-zero initial state, which is invariant under the Hamiltonian evolution. Additionally, we enhance the convergence by including 75 instances of Pauli twirling and TREX. While symmetry verification could also be applied, it has minimal and inconsistent impact in this context. For this reason, we rely solely on the renormalization strategy.

As a proof of concept, we compute the evolution for $n = 32$, shown in Figure 6.4 [right], and compare it with the extrapolated MPS evolution. We observe that both calculations agree up to $t = 10$, covering the region of interest. Importantly, error mitigation is crucial to recover the expected dynamics. We do not currently have access to more computing resources on the ibm machines, which is why wo only show this proof of concept in the

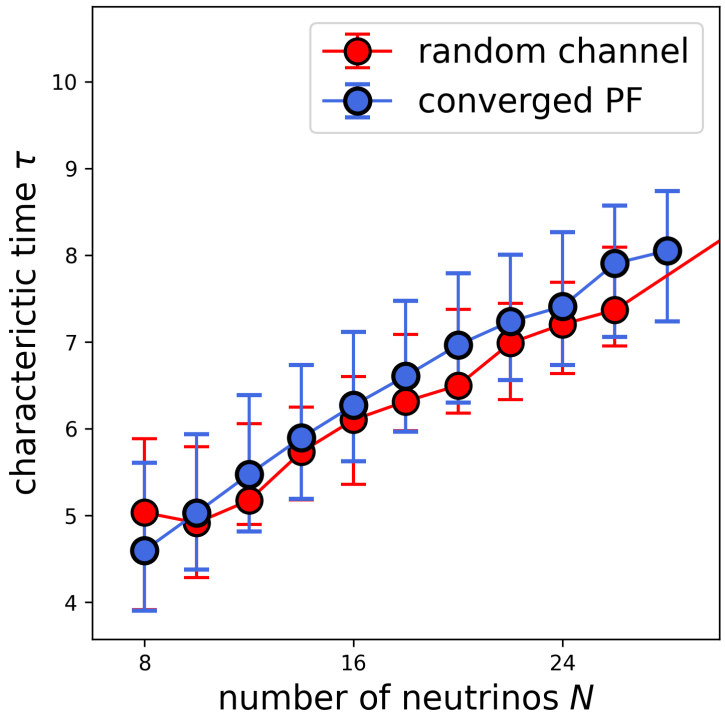

Figure 6.5: **Scaling of the variance with the system size.** The results with the converged PF are displayed in blue, while the one with random channel in red. The dotted lines display a logarithmic (violet) and square root (pink) fit, evaluated on the maximum range of the error bars.

present thesis. However, we are currently making a case to get privileged access and run the same calculation for $n = 64, 96$ and 128 on the last generation of devices.

In Figure 6.5, we compare the scaling of the variance of the observable using the Trotter expansion (blue) and the random channel approach (red), with error bars indicating the interquartile range. System sizes up to 28 neutrinos are computed using statevector simulators. We observe that our random channel yields the same characteristic time as the converged PF within confidence interval. Finally, we compute in Figure 6.6 larger-scale computation using an IBM quantum computer, as well as MPS simulation and semi-classical phase space approach (PSA) [234]. We observe the the QPU is bounded by the PSA which exhibits a square root scaling, and the MPS with has a log scaling.

6.1.4 *Discussions*

In this section, we focus on computing the thermalization time of forward-scattering neutrinos. This problem is challenging due to the high number of interactions, both for MPS simulations and quantum computers. We propose constructing an effective stochastic channel that captures the global behavior of the system, such as the variance across the neutrinos. Empirically, we show that this approach matches predictions obtained using converged PF.

We then simulate large system sizes using both MPS and the IBM quantum computer, observing agreement up to intermediate times. We are able to simulate above one hundred neutrinos using the latest IBM Heron devices.

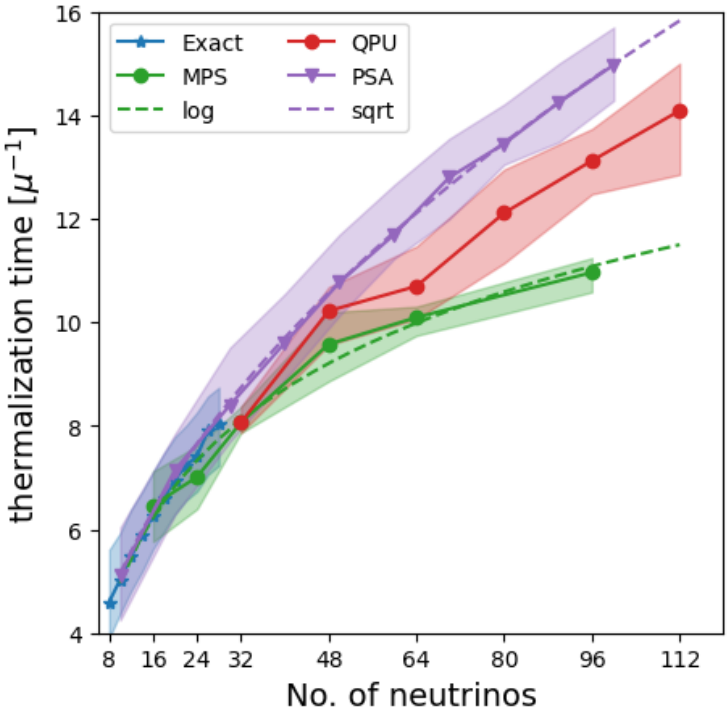

Figure 6.6: **Scaling of the variance with the system size.** The results with the QPU (red) are shown alongside the MPS calculation (green) and the semi classical phase space approach (violet).

We note that our proposal simplifies the computation of global properties, benefiting both quantum computers and MPS simulations. This is feasible because we focus on highly specific systems, and our method is thus not general. However, for this particularly difficult problem, we provide an efficient approach to handle larger system sizes.

6.2 DYNAMICS OF A QUANTUM PHASE TRANSITION

Crossing a quantum phase transition within a short duration leads to critical slowing down, which interrupts adiabatic dynamics and causes the emergence of topological defects. The mean density of these defects correlates with the quench rate, according to a universal power law as predicted by the KZM. In this Section, we aim to study the counting statistics of kink density in the one-dimensional Transverse-Field Quantum Ising Model (TFQIM). We demonstrate on a 20-qubit quantum processing unit, that higher-order cumulants follow a universal power law scaling as a function of the quench time. The following is based on Ref. [235].

6.2.1 *The Kibble-Zurek mechanism*

The KZM is a fundamental theory in non-equilibrium statistical physics, extensively used to explain the dynamics of systems experiencing continuous phase transitions. Kibble's foundational research [236], originally driven by cosmological by structure formation in the early Universe, enabled Zurek [237] to apply these concepts to superfluid helium, thereby facilitating cosmological experiments. The KZM stipulate that a system traversing a

quantum phase transition within a finite quench time τ_Q, cannot maintain equilibrium due to the phenomenon of critical slowing down [238, 239].

At the critical point h_c, systems undergoing a second-order phase transition display either a zero mass gap or an infinite correlation length, resulting in their ground state no longer evolving adiabatically [238, 240]. As the system approaches the critical point, the equilibrium relaxation time becomes infinite, hindering the system's ability to adapt swiftly enough to maintain equilibrium. This produces non-adiabatic dynamics, resulting in the emergence of excitations such as topological defects. These defects, or kinks, are the vestiges of the original symmetries.

The KZM predicts a universal scaling rule for the average density of defects, $\hat{n}$, as a function of the quench time, namely $\hat{n} \propto \tau_Q^{-\alpha}$. The exponent α is dictated by the system's geometry and the two essential exponents z and ν, representing the dynamical and correlation lengths, respectively, so that the logarithm of the mass gap scales as their product. The scaling behavior has been experimentally reproduced in both classical and quantum systems, establishing the KZM as a reliable instrument for examining non-equilibrium dynamics [241, 242].

The KZM has been extensively studied in quantum systems experiencing phase transitions, e.g the one-dimensional TFQIM, where the exact dynamic of a quantum phase transition can be analytically derived through a series of Landau-Zener transitions [243]. The distinctive features of these systems are that, unlike the original cosmological derivation and its superfluid counterpart, the fundamental symmetry is discrete ($\mathbf{Z}_2$). Additionally, a unitary, dissipation-free evolution results in a non-equilibrium state where all cumulants—beyond merely the first (the mean)—of the probability distribution of the topological defects display a KZM-like scaling. Campo [244] have demonstrated that in the thermodynamic limit, the complete counting statistics of kinks in the TFQIM exhibit universal behavior for sufficiently slow quenches [244], while for rapid quenches, there is a failure of KZM-scaling characterized by kink formation dynamics that is independent of τ_Q, resulting in a constant defect density [245]. Furthermore, they quantified and identified the biases induced by symmetry breaking [246] and the deviations that arise near equilibrium [247]. While prior research was performed analytically on one-dimensional systems, Schmitt et al. [248] examined the KZM in two-dimensional systems using numerical methods.

Investigations into the KZM have been performed on multiple platforms, including ultracold gases [249], Rydberg atoms [241], trapped ions [242], as well as both analog [250] and digital [251] quantum devices utilizing superconducting qubits. These investigations often include quantifying the statistical characteristics of topological defects generated during the phase transition, so providing empirical validation for the KZM and its wider ramifications in statistical physics. Ultimately, the KZM serves as a fundamental theory, providing a mathematical framework for comprehending the emergence of topological defects and the dynamics of symmetry breaking across many physical systems.

6.2.2 *The model*

The one-dimensional TFQIM is a paradigmatic system in the study of quantum phase transitions and critical phenomena. It consists of a linear chain of spins, each governed by two competing interactions: a transverse field and nearest-neighbor coupling. We consider the TFQIM of the form

$$H(t) = -J(t) \sum_{i=1}^{n-1} X_i X_{i+1} - h(t) \sum_{i=1}^{n} Z_i \tag{6.9}$$

with $J(t) = J_0 t/\tau_Q$ and $h(t) = (1 - t/\tau_Q)h_0$. The competition among these interactions induces a QPT at a critical field strength $h_c/J_c = 1$, wherein the system transitions from a ferromagnetic phase, marked by ordered spin alignment, to a paramagnetic phase, in which the spins align with the transverse field. The model is analytically solvable, establishing it as a fundamental model for comprehending quantum criticality and non-equilibrium dynamics in low-dimensional systems. We establish $h_0 = J_0 = 1$ to ensure that the quantum phase transition occurs at the midpoint of the quench, namely at $t = \tau_Q/2$. The initial state is chosen as the ground state at $t = 0$, represented by the all spin-up state $|0\rangle^{\otimes N} \equiv |\bar{0}\rangle$. We use the kink operator to quantify the amount of topological defects

$$\hat{n} = \frac{1}{2n} \sum_{i=1}^{n-1} (\mathbb{1} - X_i X_{i+1}) \,. \tag{6.10}$$

The cumulants are the coefficients in the series expansion of the moment-generating function for kink distributions. We focus on the first two, namely:

$$\kappa_1 = \langle \hat{n} \rangle, \tag{6.11}$$
$$\kappa_2 = \langle (\hat{n} - \langle \hat{n} \rangle)^2 \rangle = \langle \hat{n}^2 \rangle - \kappa_1^2 \tag{6.12}$$
$$\tag{6.13}$$

This observable can be estimated within an error ϵ by measuring the wavefunction in the Hadamard basis at least $\mathrm{Var}[\hat{n}^k]/\epsilon^2 \leq 1/\epsilon^2$ times. All correlators can be assessed on this basis and hence be calculated at the same time. The primary problem occurs during long quench durations, since the expectation values are anticipated to be small, necessitating an increased number of statistics to discern them from shot noise. Nonetheless, there exists no inherent issue with calculating expected values in this manner.

6.2.3 *Time-dependent quantum simulation*

The quench is executed by evolving the initial state using a time-dependent Hamiltonian $H(t)$. The temporal evolution is derived from the time-ordered unitary evolution originating from the Hamiltonian as

$$\mathcal{U}(t_f, t_0) = \mathcal{T} \exp\left(-i \int_{t_0}^{t_f} H(t) dt\right), \tag{6.14}$$

where $\mathcal{T}$ is the time ordering operator to account for the non-commutativity at different time $[H(t), H(t')] \neq 0$. Various methods exist for implementing $\mathcal{U}(t_f, t_0)$ on digital quantum computers, with the most direct approach being the quasistatic approximation [252]. Within this framework, the simulation time is divided into r Trotter steps of size $\Delta t = (t_f - t_0)/r$, where $H(t) \approx H(k'\Delta t)$, with $k' = k + 1/2$ and $t \in [k\Delta t, (k+1)\Delta t]$. The time-evolution operator may thus be approximated by

$$\mathcal{U}(t_f, t_0) = \prod_{k=1}^{r} \exp\left[-i\Delta t H(t_0 + k'\Delta t)\right] . \tag{6.15}$$

The matrix exponential $\exp(-i\Delta t H)$ can be decomposed into a set of universal gates using for example PF developed in Chapter 5. In the following, we shall consider a second order

Trotter-Suzuki decomposition. For simplicity, let us assume that the Hamiltonian breaks down into two groups of pairwise commutating term,

$$H(t) = f(t)A + g(t)B, \tag{6.16}$$

such that $\exp(iA)$ and $\exp(iB)$ are exactly implementable. We use a second-order product formula, which is given in the static case by

$$\exp(-iHt) = \exp(-if(t)At/2)\exp(-ig(t)Bt)\exp(-if(t)At/2) + \mathcal{O}(t^3). \tag{6.17}$$

Merging both approximations (6.15) and (6.17), we obtain

$$\begin{aligned}\mathcal{U}(t) &\equiv \mathcal{U}(t_f = t, t_0 = 0) \\ &= \prod_{k=1}^{r} \exp[-if(k'\Delta t)A\Delta t/2] \cdot \exp[-ig(k'\Delta t)B\Delta t]\exp[-if(k'\Delta)tA\Delta t/2].\end{aligned} \tag{6.18}$$

Straightforward extensions of the quasistatic PF method include adaptive step sizing [253], random compilation techniques like qDrift [194, 196] with one-norm scaling [254], and leveraging the interaction picture [255]. While asymptotically superior methods exist—such as those based on the Dyson series [256], Magnus expansions [257, 258], discrete clocks [259], or flow equations [260]—the PF method is notably effective in practice, especially for systems with a high degree of locality [190].

6.2.4 *Quench on a quantum processor*

We begin by performing the quench on the IQM Garnet superconducting quantum computer [261], which consists of twenty transmon qubits and achieves a two-qubit gate fidelity of 99.5% [262]. Due to the chip's connectivity, displayed in Fig. 6.9, we work with a system of $N = 19$ spins. We implement the second-order PF method, varying the number of Trotter steps r up to 20 across different quench times $\tau_Q \in [0.1, 20]$. The first two cumulants are presented in Fig. 6.7, with statevector simulations shown as blue squares, raw data as purple dots, and error-mitigated results as red dots. The continuous line represents a near-exact evolution, simulated on a statevector simulator using 2000 Trotter steps. While this is not an exact result, we choose the size of the Trotter steps such that the error is at most of the order 10^{-6}. Higher-order moments are not displayed, as they are several orders of magnitude smaller and would require significantly more measurements to be accurately estimated.

We perform a second 50 spins and 100 spins experiment on the ibm_kingston and ibm_aachen Heron r2 devices. They are heavy-hexagonal lattice chips consisting of 156 qubits, with two-qubit gates fidelity between $[99.83\%, 99.94\%]$. The results for the first cumulant as shown in Fig. 6.8, and are compared to equivalent MPS simulations. The qubits are selected such as to minimize the worst two-qubit gate fidelity. The second order cumulant can not be resolved because of the limited amounts of samples collected, which should grows as $\mathcal{O}(N^4)$. For the larger experiments, we noted that it was more effective to choose the renormalization factor as the maximum over every observable

$$(1-p) = \max_{0\leq i<N-1} \langle \mp | \tilde{\mathcal{U}}^\dagger(t) X_i X_{i+1} \tilde{\mathcal{U}}(t) | \mp \rangle, \tag{6.19}$$

which appears as a solid line in the corresponding inset.

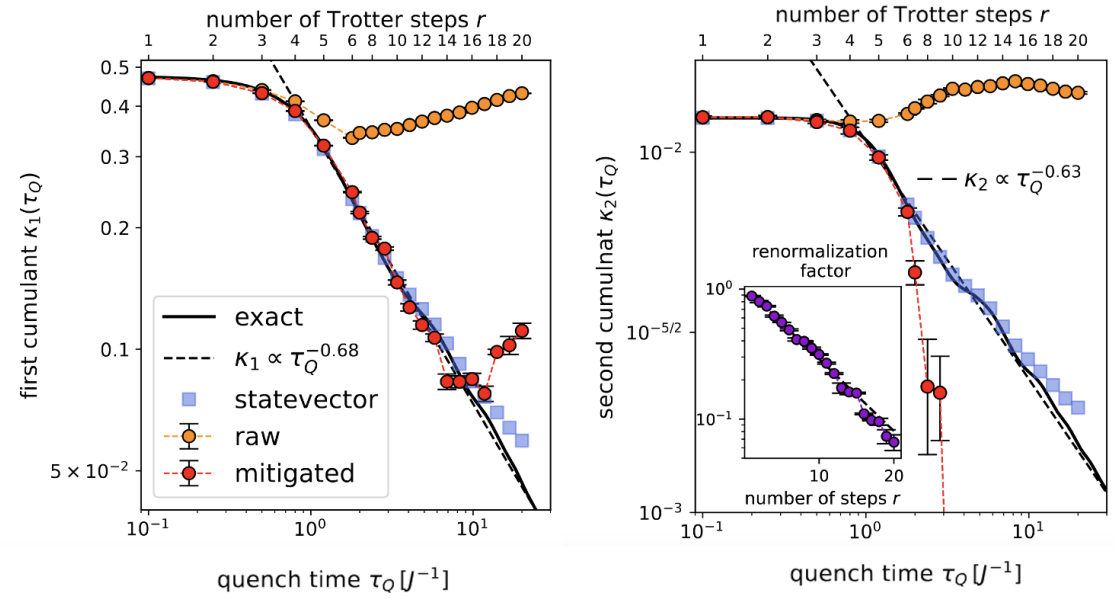

Figure 6.7: **Quantum simulation on IQM Garnet ($n = 19$).** The first two cumulants of the kink density are shown as a function of the quench duration and the number of Trotter steps. The blue squares represent statevector simulations, the orange points indicate raw data, and the red ones shows results processed with the complete error mitigation pipeline. The complexity of the circuits increases by one extra Trotter each time. The solid black line illustrates the exact evolution, while the decay rate is shown by a dotted line. The inset illustrates the renormalization factor (purple) as a function of the circuit is depth. Figure taken from Ref. [235].

Error mitigation is implemented through noise renormalization techniques [107–110, 263], see Section 3.2.3, enhanced by Pauli twirls [94] and readout [96] twirling, as well as MMM [99]. We normalize each expectation value utilizing the identity

$$\langle \mp | \mathcal{U}^\dagger(t) \hat{n} \mathcal{U}(t) | \mp \rangle = 1, \tag{6.20}$$

employing the alternative scheduling function $h(t) = \pi/\Delta t$ for the field, thereby rendering the associated evolution a symmetry of the Hamiltonian. The renormalization factor exhibits an exponential decay with the depth, as seen in the inset of the right plot in Figure 6.7. Consequently, this approach cannot be used for circuits with too high depth; nevertheless, in this instance, it allows us to expand it by a factor of four. Each circuit is executed 2000 times, and the error bars represent a 99% confidence interval calculated by Bayesian data augmentation. The accuracy limit due to finite statistics is $10^{-5/2}$, as seen on the y-axis.

6.2.5 *Discussions*

Until a few years ago, simulations of critical dynamical behavior were only feasible in specialized laboratories. The emergence of on-demand quantum digital processing units has enabled a wider array of researchers to investigate the universality of scaling principles and their deviations in complex systems.

We perform a quantum simulation of the dynamics of a QPT using a digital superconducting quantum computer, with a circuit depth of up to one hundred consecutive CNOT

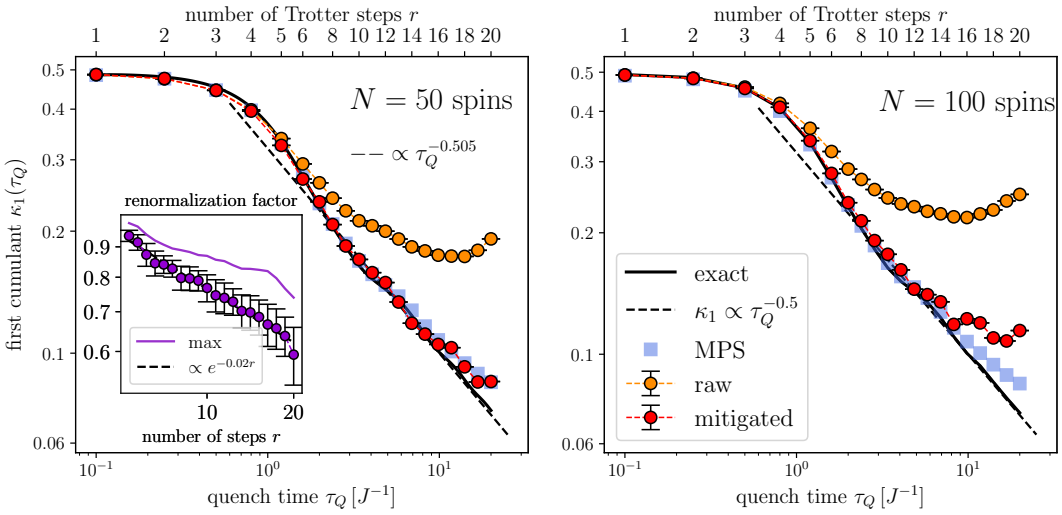

Figure 6.8: **Quantum simulation on IBM Kingston ($N = 50$) and IBM Aachen ($N = 100$).** The first cumulant of the kink density is shown as a function of the quench time, as well as the number of Trotter steps. The blue squares denote MPS simulation of the exact same circuits, the orange points the raw data, and the red the results obtained after the full pipeline of error mitigation. The complexity of the circuits is increased as one additional Trotter step per data point. The black continuous line shows a proxy for the exact evolution, which is obtained with a maximal of 100 Trotter steps, while the decay rate appears as a dotted line. The inset in the $N = 50$ plot refers to the renormalization factor computed according to the expression given in (6.19).

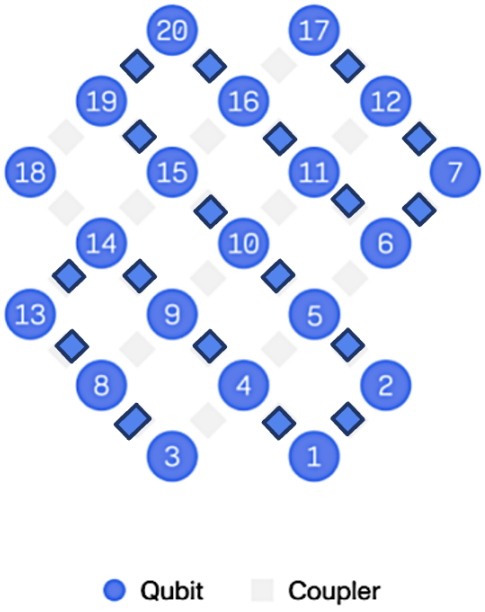

Figure 6.9: **Topology of IQM Garnet.** The active connections used in the experiment are displayed with the shaded blue squares, and the 18th qubit is inactive since we can not find a full 20 qubit line.

gates, while implementing various error mitigation strategies. Through error mitigation, we increase the number of Trotter steps by a factor of four, at least for the first cumulant. Calculating higher cumulants requires more measurements, since they require greater precision due to their lower values, which we did not have the budget for. Despite the exponential increase in error mitigation costs [87, 125], rendering it unsustainable in the long run, it is an essential instrument for leveraging quantum computers in the present time. We argue that error mitigation serves as a bridge to the age of early fault-tolerance [264].

7

APPLICATIONS OF FOURIER MOMENTS

We are now in a position to bring all the ingredients discussed in the present thesis together, and discuss one promising application of near-term quantum computers: incoherent Hamiltonian transformation through Fourier moments. The Hamiltonian moments in Fourier space, i.e. expectation values of the unitary evolution operator governed by a Hamiltonian at various times, offer a powerful framework for apprehending quantum systems. They provide insights into e.g. energy distribution, higher-order dynamics, response functions, and correlations. However, in the context of quantum algorithms, the main reason why Fourier moments are interesting, is that they allow to perform polynomial transformation of a given Hamiltonian, using relatively low circuit's depth circuits. This process is performed incoherently by first analytically expanding the function of interest into its Fourier serie, and computing the phases, or Fourier moments, on the quantum processor. This is a more convenient framework for near-term quantum computers, as only the time evolution is performed there, while the rest of the work is done on a classical device. While this does not have the optimal asymptotic scaling of QSP [204], which is instead coherently realized on the quantum computer using LCU, it is much more favorable near-term.

In this chapter, we will see how to compute Fourier moments on a quantum processor [110], and present two applications: response function [265] in Section **??** and ground-state energy estimation [266, 267] in Section 7.4.1. We will start by defining the core component of those two applications, namely Fourier moments, and how to effectively compute them on a quantum processor.

7.1 REVIEW OF QUANTUM SIGNAL PROCESSING

QSP provides an efficient method for performing arbitrary Hamiltonian transformation in a coherent way. Specifically, QSP outputs a degree-d polynomial $f \in \mathbb{C}[x]$ using a sequence of unitary operators [33, 268, 269]

$$U_{\mathrm{QSP}} := \prod_{k=1}^{d} S(\phi_k) W(x) = \begin{pmatrix} f(x) & * \\ * & * \end{pmatrix}, \tag{7.1}$$

where

$$S(\phi) := \begin{pmatrix} e^{i\phi} & 0 \\ 0 & e^{-i\phi} \end{pmatrix}, \quad W(x) := \begin{pmatrix} x & \sqrt{1-x^2} \\ \sqrt{1-x^2} & -x \end{pmatrix}. \tag{7.2}$$

Here, we follow the convention in Corollary 8 of [269] , where $W(x)$ acts as a reflection operator. The $*$ denote arbitrary entries making the full matrix unitary. This is referred to as block-encoding, where a non-unitary operation is encoded into a unitary acting on a larger space. For any polynomial $f(x)$ that meets certain criteria [269], there exists a set of QSP angles $\{\phi_k\}$. These criteria are

1. f must have parity $(d \mod 2)$.

2. $|f(x)| \leq 1$ for all $x \in [-1,1]$.
3. $|f(x)| \geq 1$ for all $x \in (-\infty,-1] \cup [1,\infty)$.
4. $f(ix)f^*(ix) \geq 1$ for all $x \in \mathbb{R}$ if d is even.

The function $f(x)$ is then encoded in $\langle 0|U_{\text{QSP}}|0\rangle$. Since polynomials are dense in the space of smooth functions, it is easy to see that any (smooth) Hamiltonian transformation can be efficiently implemented in this way. While the depth of this circuit is only proportional to the degree d, it is still difficult to implement QSP on noisy devices. Some early experiments include the computation of bipartite entanglement entropies in Ref. [270]. It has been shown in Ref. [5] that the main quantum algorithms known up to this day (QPE, Quantum eigenvalue transformation, factoring, ...) fall under the umbrella of Quantum Singular Value Transform (QSVT), which is an extension of QSP for arbitrary (even non-square) matrices.

As a direct application, QSP provides an efficient approach to Hamiltonian simulation. The main idea is to approximate e^{-ixt} over an interval $I \subseteq [-1,1]$ by a polynomial of fixed degree. For a given time $t > 0$ and accuracy ϵ_{poly}, we seek a polynomial f such that

$$\max_{x \in I} |f(x) - e^{-ixt}| \leq \epsilon_{\text{poly}}.$$

One way to find f is through the Jacobi-Anger expansion [268] of the exponential function:

$$\begin{aligned} e^{-ixt} &= \cos(xt) - i\sin(xt) \\ \cos(xt) &= J_0(t) + 2\sum_{k=1}^{\infty} J_{2k}(t)T_{2k}(x) \\ \sin(xt) &= 2\sum_{k=1}^{\infty} J_{2k+1}(t)T_{2k+1}(x), \end{aligned} \tag{7.3}$$

where $J_i(t)$ denotes the i-th order Bessel function and $T_i(x)$ is the Chebyshev polynomial of order i. By tolerating an error ϵ_{poly}, the polynomial can be truncated to degree

$$d = \Theta\left(\frac{t + \log(1/\epsilon_{\text{poly}})}{\log(e + \log(1/\epsilon_{\text{poly}})/t)}\right),$$

which scales nearly linearly with t and logarithmically with $1/\epsilon_{\text{poly}}$. The goal is to apply this polynomial transformation to the spectrum of the Hamiltonian H. This is achieved by block encoding H, i.e. embedding it within a unitary operator $W(H)$ acting on a larger Hilbert space, e.g. using LCU [271].

While QSP techniques are symptomatically optimal, and the method of choice in the Fault-tolerant Quantum Computing (FTQC) regime, the required depth remains out of reach in general for near-term quantum computer. It is therefore important to study more practical, though less optimal, methods. One of them is based on Fourier moments, as we will now discuss.

7.2 FOURIER MOMENTS

Given a the time evolution operator $\mathcal{U}(t) = \exp(-iHt)$, an initial state $|\Psi_0\rangle$, and a time step τ, the j-th Fourier moment is defined as

$$g_j(\tau) = \langle \Psi_0 | \mathcal{U}(j\tau) | \Psi_0 \rangle, \tag{7.4}$$

whose real (imaginary) part can be computed using an Hadamard test, introduced in Section 2.2, with the following circuit [29]

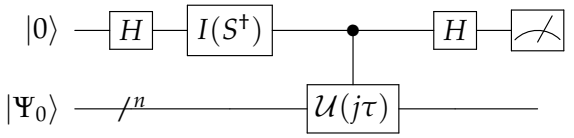

. (7.5)

One see that it all boils down to the implementation of control time-evolution, which can still be significantly more difficult than the vanilla version, applied on an initial state. Usually, initial states are chosen close to the ground state, which can be prepared on near-term quantum computer using for example variational methods, discussed in Chapter 4. For the dynamics, we will consider deterministic PF, see Chapter 5, since they are easier to analyse. We note however that the cost can be significantly reduced, when specific systems are considered, using random product formula as discussed in Section 5.2.2. While the overhead of controlling the time evolution might be restrictive in general, it can be reduced when considering the specifics of the physical system. In general we have for arbitrary unitaries A and V

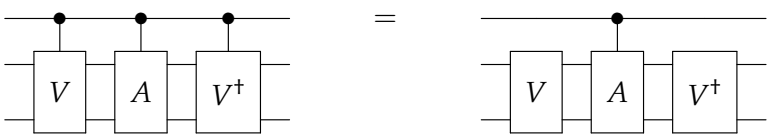

. (7.6)

It turns out that many interactions can be written in this form. For instance, when considering physical Hamiltonian, it is sufficient to only consider terms composed of tensor product of Pauli Z matrices, as we may perform arbitrary basis transformation using single-qubit gate. In the case of two-qubits, the controlled interaction can be written as

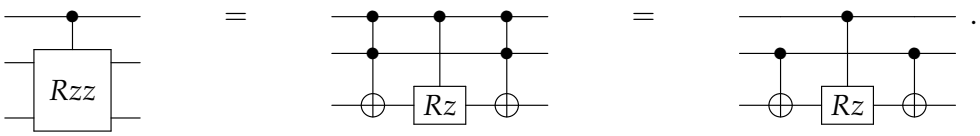

. (7.7)

This identity can be easily verified by considering the ancilla being in the state $|0\rangle$ or $|1\rangle$. Hence, only the rotation needs to be controlled, leading to an overhead of 2 CNOT gates per term, to implement the control rotation, assuming arbitrary connectivity. The situation is analog for multi-qubit interaction beyond two qubits.

While this implementation only add two CNOT gates per term in the optimal scenarios, the situation can be more severe if the quantum processor has limited connectivity. Hence, we are required to swap qubits around, which increases the cost of controlling each rotation on the ancilla. Therefore, it is important to find alternatives to directly controlling the time-evolution, which can be used in any algorithms based on controlled time evolution.

7.2.1 *Control-free implementation*

Although the CNOT overhead from control operations is only linear in the number of terms, it is important, especially when using actual devices, to minimize this cost as much as possible. Hence, this overhead appears in many quantum algorithms and will remain in the FTQC era. Although ancilla-free approaches have been proposed in this context, see Refs. [111, 272, 273], we choose to minimize overhead by using Control Reversal Gate (CRG), as discussed previously in [274, 275]. A reversal gate R is a combination of Pauli matrices that anti-commutes with the Hamiltonian, expressed as $\{H, R\} = 0$. The control reversal gate comprises its controlled version on an ancilla in the state $|0\rangle$ and effectively toggles the flow of time in either the forward or backward direction. For instance, if the ancilla is in the zero state, we apply

$$R \exp(-iHt) R^\dagger = R \sum_n \frac{(-itH)^n}{n!} R^\dagger = \sum_n \frac{(itH)^n}{n!} R R^\dagger = \exp(itH), \tag{7.8}$$

which is a backward time evolution, and simply a forward time evolution $\mathcal{U}(t)$ otherwise. The CRG framework has the additional advantage of being faster by a factor of two, since the phase kickback contribution occurs twice. Consequently, just half the number of Trotter steps is necessary to achieve a given precision. To compute the real component of $\mathcal{U}(j\tau)$, we can use the following circuit

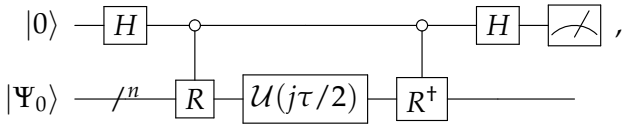

, (7.9)

while the imaginary part can be computed by adding a S gate on the ancilla wire.

It has to be noted, that the backward time evolution has to be the inverse of the forward time evolution, i.e., $\mathcal{U}(-t)\mathcal{U}(t) = \mathbb{1}$, and hence, $\mathcal{U}(-t) = \mathcal{U}(t)^{-1}$. This is inherently fulfilled for product formulae of even order. Importantly, the overhead of the CRG does not scale with the number of Trotter steps.

The attentive reader may have notice that this is only true if we are able to find CRG for the whole Hamiltonian, which is not always possible. In that case, we can partition H into M groups such that $H = \sum_m^M H_m$, and find distinct CRG R_m for each group. We can now simply insert the CRGs between the evolution of each term. Unfortunately, this approach now scales linearly with the number of Trotter steps. However, even in the most unfavorable situation, we can demonstrate that it is sufficient to set $M = 2n$, where n represents the number of qubits.

The grouping can be achieved as follows: for every term in the Hamiltonian, we first group together all terms that have an X or Y operation at location 0, and in a separate group, we gather all terms with a Z operation at location 0. We then choose $R = Z_0$ for the first group and $R = X_0$ for the second group. This procedure is repeated for all n locations, obtaining $M = 2n$ groups that can be reversed using single-qubit Pauli operators.

Using this strategy, for a general n-qubit Hamiltonian, a single second-order Trotter step controlled on an ancilla qubit can be implemented with an additional cost of $8n$ CNOT gates over the base cost of implementing the first Trotter step without a control. In comparison, a naive implementation of a controlled first-order Trotter step for a Hamiltonian H composed of L Pauli strings, where each term is exponentiated individually, requires $4L$ CNOT gates.

When $L > 2n$, this CRG approach provides an advantage over a direct implementation. Specifically, the overhead scales at most linearly with the system size, whereas a direct

implementation for a k-local Hamiltonian requires $\mathcal{O}(n^k)$ additional entangling gates. To the best of our knowledge, this is a novel result that may be of interest in applications beyond Fourier moment estimation.

This is, however, the worst-case situation, and much better implementation can be found on a case-by-case basis. For example, one-dimensional Heisenberg models only need the two following groups by [274]

$$R_Z = \bigotimes_{i=0}^{\lfloor n/2 \rfloor} Z_{2i} \qquad R_X = \bigotimes_{i=0}^{\lfloor n/2 \rfloor} X_{2i} \,. \tag{7.10}$$

It turns out, that for our specific EFT model from Section 4.3, we can find an efficient CRG sequence, as we will see later.

7.2.2 *Interpolation via compressive sensing*

Since the computation of Fourier moments is an expensive task, it is beneficial to develop ways to reduce the computation cost. CS is a technique aiming at reconstructing a signal of size $d+1$, from very few samples. To be more precise, let us assume we are given a signal $f(t) = \sum_k x_k \psi_k(t) \in \mathbb{R}^d$, which is S-sparse in some basis. Then, we can reconstruct the signal with very high probability using only $\mathcal{O}(S \log d/S)$ data points. For context, we assume the signal we wish to recover is a time serie of Fourier moments

$$f(t) = \mathrm{Tr}\left[\rho e^{-iH\tau j}\right] = \sum_{k\in\mathcal{S}} |c_k|^2 e^{-iE_k\tau k} = \sum_{k\in\mathcal{S}} |c_k|^2 (\cos E_k\tau k + i \sin E_k \tau k), \tag{7.11}$$

where $\rho = \sum c_j |j\rangle$ is a given initial state written in the eigenbasis of H, with some support $\mathcal{S}$. If the signal is sparse in the eigenbasis, e.g. ρ lives in a small subspace and we have $|\mathcal{S}| = \mathcal{O}(\mathrm{poly}(n))$. CS is able to perform two tasks: under sampling, and shot noise mitigation. It would be also useful to be able to extrapolate the signal in time. However, our experience shows that this is a much difficult task which seems to be out of the reach of CS.

7.2.2.1 *Under sampling*

In the under sampling scenario, we aim to reconstruct the signal f, by only having access to $m \in \Omega$ samples. By assuming that f is S-sparse in the $\{\psi_k\}$ basis, we simply consider $f^* = \Psi x^*$, where x^* is the solution of the convex optimization problem

$$\min_{\tilde{x}\in\mathbb{R}^d} \|\tilde{x}\|_1 \text{ subject to } f_k = (\Psi\tilde{x})_k \,\forall k \in \Omega. \tag{7.12}$$

Im fact, we are optimising for the sparest solution compatible with the data. For the basis transformation, we can simply choose the discrete cosine transform $\Psi = \mathrm{DCT}(\mathbb{1})$ for the real part, and the discrete sine transform for the imaginary one

$$\begin{aligned} \mathrm{DCT}(x)_k &= \sum_{j=0}^{d} x_j \cos \frac{\pi k(j+1/2)}{(d+1)} \\ \mathrm{DST}(x)_k &= \sum_{j=0}^{d} x_j \sin \frac{\pi (k+1)(j+1/2)}{(d+1)} \end{aligned} \tag{7.13}$$

From [276], we have following guarantees:

Theorem 6. *Let U be an orthogonal matrix (the conversion matrix between the measurement basis and the basis where the signal is sparse), with $|U_{k,j}| \leq \mu(U)$ and some samples $m = |\Omega|$. If the number of samples is chosen such that*

$$m \geq C_0 \mu(U)^2 S \log d/\delta \tag{7.14}$$

$$m \geq C_0' \log^2 d/\delta, \tag{7.15}$$

for some constant C_0, C_0', then every signal of sparsity S can be recovered with probability $(1-\delta)$.

Since the optimization can be recast as a linear program, and that the optimization is convex, efficient methods do exist. Moreover, it has been shown that this technique is optimal, i.e., that the signal can not be recovered from fewer samples.

We note that this technique can be improved with re-weighting the sparsity [277]. Here the idea is to introduce a weight vector ω, and to solve instead

$$\min_{\tilde{x}\in\mathbb{R}^d} \|\omega \cdot \tilde{x}\|_1 \text{ subject to } f_k = (\Psi\tilde{x})_k \, \forall k \in \Omega. \tag{7.16}$$

The weights are initially chosen to be all ones, and are iteratively updated to $\omega_i = 1/(||x_i^{(j)}|| + \epsilon)$, where $x^{(j)}$ is the solution obtained at the j-th optimization procedure, and ϵ a small constant required for numerical stability. This will effectively enforce the sparsity, as small coefficients are given more weights and are thus pushed towards zero.

7.2.2.2 *Robustness*

CS has also the powerful ability to deal with noise, which includes the unavoidable statistical shot noise. Hence, we only have access to the unbiased estimator of the true signal, which can be described as $\hat{f}(t) = f(t) + z(t)$, where $z(t)$ is a normal distributed random variable around 0 and variance depending on the initial state. It has been shown in Ref. [278], that the obtained solution is robust to noise, but also to approximate sparsity. To better understand the claim, we first need to introduce the restricted isometry property (RIP)

Definition 2 (RIP). *For each integer $S = 1, 2, \ldots$, define the isometry constant δ_S of a matrix A as the smallest number such that*

$$(1-\delta_S)\|x\|_2^2 \leq \|Ax\|_2^2 \leq (1+\delta_S)\|x\|_2^2, \tag{7.17}$$

holds for all S-sparse vector x.

The intuition is that any S column vector of A are approximately orthogonal.
We then solve the relaxed following minimization problem

$$\min_{\tilde{x}\in\mathbb{R}^d} \|\omega \cdot \tilde{x}\|_1 \text{ subject to } \|f_k - (\Psi\tilde{x})_k\|_2 \leq \epsilon \, \forall k \in \Omega. \tag{7.18}$$

where ϵ quantifies the amount of noise. We finally have the following result

Theorem 7. *Assuming that Ψ obeys the RIP with constant $\delta_{2S} \leq \sqrt{2}-1$, then the solution obeys*

$$\|x^* - x\|_2 \leq C_0 \|x - x_S\|_1/\sqrt{S} + C_1\epsilon, \tag{7.19}$$

where x is the exact solution and x_S the S-sparsified solution.

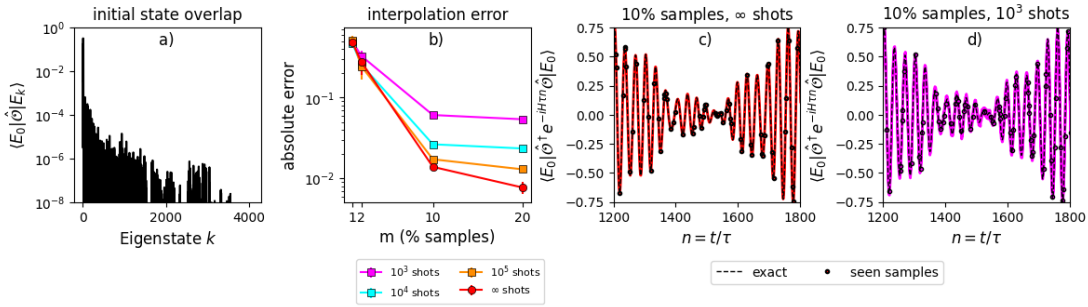

Figure 7.1: **Compressive Sensing.** Performance of CS for a 12-qubit Hamiltonian.

7.2.2.3 *Numerical results*

For illustration, we examine three nucleons on a 8×8 grid, with their interactions characterized by the EFT outlined in Section 4.3. The initial states consist of the exact ground state, which is excited via the momentum transfer operator $\hat{O}(\vec{q})$, as defined in (4.25). The results of the interpolation are shown in Figure 7.1.The left panel displays the overlap of the initial states on the eigenbasis, revealing that it is sparse up to a cutoff of 10^{-5}. Panel b) illustrates the interpolation error as a function of the sample size, with color indicating the varying number of measurements for each unique sample. Ultimately, the two rightmost panels show the interpolation using 10% of the data.

We observe, that for an initial state of physical significance, we can decrease the number of samples by an order of magnitude and still retrieve the full signal, even in the presence of noise.

Interpolation is a useful tool which can reduce the actual cost of quantum algorithms based on Fourier moments, with a low classical overhead. However, it is quite sensitive to shot noise, and therefore will only work when all the moments are estimated with high precision. Thus, it is mainly useful when the Fourier coefficients are fairly of the same order of magnitude. Hence, if they decay quickly, it is much favorable to evaluate the sum as a whole with importance sampling, and therefore not having precise estimate of the individual moments. It turns out that Gaussian functions, which are used for response functions, fall into this category, while the Heaviside function, used for GSEE, does not benefit much from this approach.

7.3 RESPONSE FUNCTIONS

Response functions, which characterize the linear response of a many-body system after an excitation, are often difficult to calculate from first principles, making them an attractive application for quantum computers. We guide the reader to Ref. [**review_standard_model_roggero**] for a review of the subject. For example, they can help us better understand scattering mechanisms by examining the internal structure of the target [279], since they include the same information as an inclusive reaction cross-section. They are often defined via the spectral density operator $\delta(\omega - (H - E_0))$, where E_0 denotes the target ground state energy and H represents the Hamiltonian describing the target. The interaction between the target and the external probe is characterized by an excitation operator $\hat{O}(\vec{q})$, which represents the momentum transfer $\vec{q}$ of the scattering process. At high momenta, the response functions are expected to be mostly influenced by the momentum distribution or spectral functions of the target. Nonetheless, for lower momentum transfer, two-body

interactions are crucial for assessing neutrino characteristics in neutrino-nuclei studies, such as MiniBooNE, MicroBooNE, T2K, and DUNE [280]. The following is largely based on Ref. [110].

Given an Hamiltonian H, an initial state $|\Psi_0\rangle$ describing the state of the target and the momentum transfer excitation operator $\hat{O}(\vec{q})$, we define the frequency dependent *response function* as

$$S(\omega,\vec{q}) = \langle\Psi_0|\hat{O}(\vec{q})^\dagger\delta(\omega-(H-E_0))\hat{O}(\vec{q})|\Psi_0\rangle = \sum_f \left|\langle\Psi_0|\hat{O}(\vec{q})|f\rangle\right|^2\delta(\omega-(E_f-E_0)), \tag{7.20}$$

where $\{|f\rangle\}_f$ are the eigenstates of H with corresponding energies $\{E_f\}_f$. Calculating this quantity is often quite difficult in practice, since it requires the whole spectrum. An effective approach to address this issue is to use an integral transform of the type

$$\Phi(\nu,\vec{q}) = \int d\omega K(\nu,\omega)S(\omega,\vec{q}) = \langle\Psi_0|\hat{O}(\vec{q})^\dagger K(\nu,(H-E_0))\hat{O}(\vec{q})|\Psi_0\rangle. \tag{7.21}$$

This enables a direct computation, using simpler ground state methods and appropriate kernels. In Quantum Monte Carlo computations, the Laplace kernel is often used because to its association with imaginary-time correlation functions [280–283], while the Lorentz integral transform [284–288] is more suited for CC or exact diagonalization methods.

To effectively reconstruct the energy dependence of the response function with minimal errors, one typically necessitates integral kernels that serve as smooth approximations of the energy delta function from Eq. (7.20), possessing a finite width (refer to [289] for a more comprehensive characterization), such as Gaussian or Lorentzian functions. Utilizing a translationally invariant integral kernel, namely where $K(\nu,\omega) = K(|\nu-\omega|)$, and representing it via a collection of orthogonal polynomials ϕ_j, we get

$$\Phi(\nu) = \sum_j^\infty c_j(\nu)\langle\Psi_0|\hat{O}(\vec{q})^\dagger\phi_j(H)\hat{O}(\vec{q})|\Psi_0\rangle\ , \tag{7.22}$$

where the coefficients $c_j(\nu)$ specify the particular kernel function employed. For this approach to be efficient, we shall truncate the serie expansion while keeping a controllable small error. A Gaussian kernel, represented using either Chebyshev or Fourier orthogonal polynomials, usually fulfills this criterion [265, 289, 290]. The integral transform can now be expressed using the Fourier basis by truncating the series expansion to orders D as

$$\Phi_N(\nu) = \sum_{j=-D}^{D} c_j(\nu)m(j\tau),\ \text{with}\ m(j\tau) = \langle\Psi_0|\,\hat{O}(\vec{q})^\dagger e^{-iHj\tau}\hat{O}(\vec{q})\,|\Psi_0\rangle\,. \tag{7.23}$$

We recall that the Fourier moments $m(t)$ are the only values intended to be computed on a quantum computer. The selection of the time-step τ and the truncation order D influences the lowest energy resolution Δ attainable with the integral transform at a certain precision. Specifically, it can be demonstrated that, utilizing the Gaussian kernel, the total evolution time $T = D\tau$ must scale as $T = \mathcal{O}(1/\Delta\sqrt{\log(1/\epsilon)})$ in order to maintain the approximation error below ϵ. Here, Δ refers to the energy resolution allowed by the kernel (for further details, refer to [265, 289]).

7.3.1 *Moment of a nuclear EFT*

In this section, we present the computations of the first few moments on an IBM quantum processor, using echo verification (Sec. 3.3.2) and noise renormalization (Sec. 3.2.3) error mitigation techniques. We conduct an experiment with a single double Trotter step with stretched time and another with multiple double steps at fixed time. By *double Trotter step*, we refer to two Trotter steps with half the time each, $\mathcal{U}(j\tau) = \mathcal{U}(j\tau/2)\mathcal{U}(j\tau/2)$, as our error mitigation scheme limits us to performing an even number of steps.

7.3.1.1 *Explicit circuits*

We first start by explicitly [110] constructing the circuit to perform the Hadamard test using CRG. We discuss the optimization decomposition as well as the implementation of the control reversal gates. We recall that the Hamiltonian of interest, see Eq. (5.73), reads

$$H = \alpha \sum_k X_k + \beta \left(Z_1 Z_4 + Z_2 Z_3 + \sum_{i<j<k} Z_i Z_j Z_k \right) . \tag{7.24}$$

We want to execute the Hadamard test with CRG, thus we are seeking a gate that anti-commutes with the Hamiltonian. For the single X rotations, we can either use a Y or a Z, while the diagonal component requires a X or a Y. With linear connectivity, we can do the following

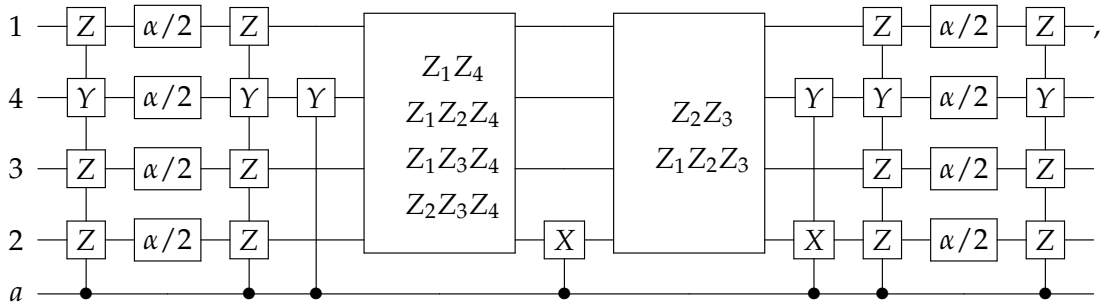

, (7.25)

where the box with $\alpha/2$ is the X rotation of angle $\alpha/2$ for the one body and the other multi-qubit gates implement the rotations proportional to β. We can simplify this, by using the fact that Z commutes with the big boxes, to

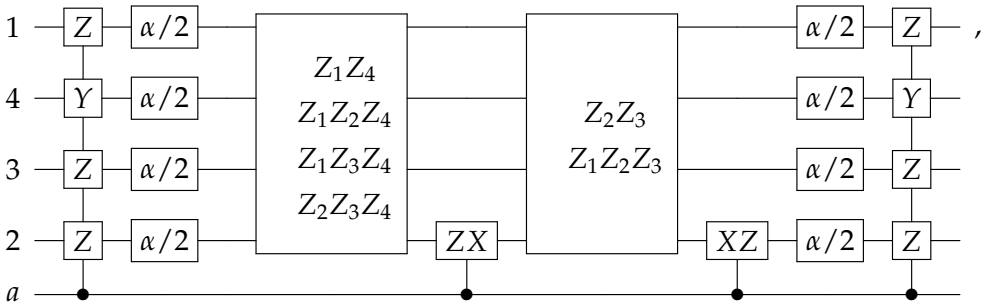

, (7.26)

Now, since $XZ = -ZX$ we have

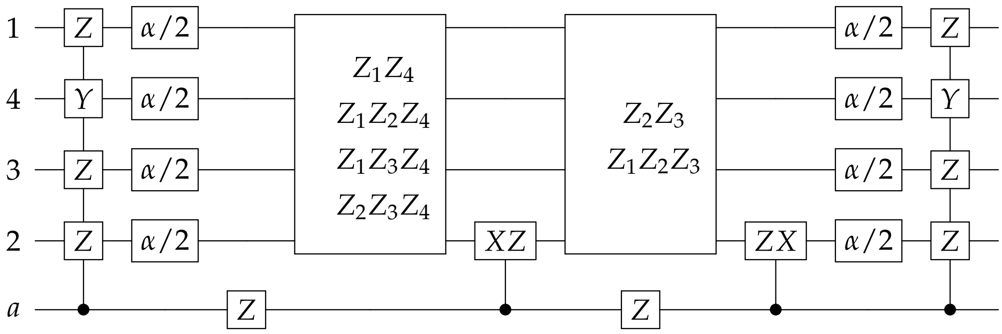

, (7.27)

which simplifies to

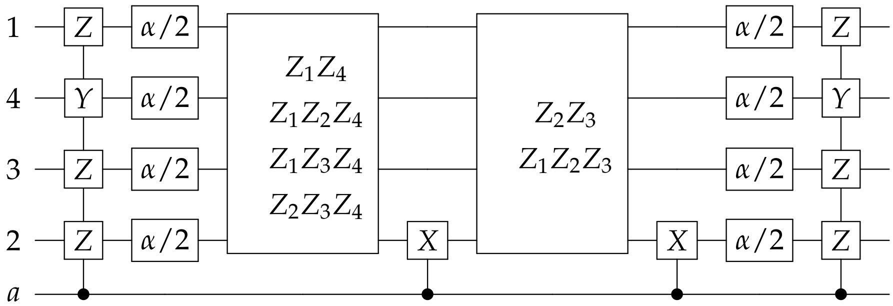

. (7.28)

The controlled Pauli operations at the edges need to be executed just twice, regardless of the number of steps. This may be accomplished using linear connection with seven CNOT gates

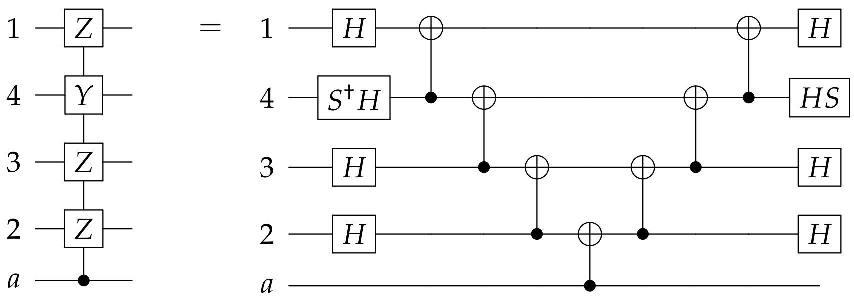

. (7.29)

We note that the three Hadamard gates on the right may be repositioned beyond the X rotations, so transforming them into more elementary Z rotations. For the interaction components, we begin with three terms involving $Z_1 Z_4$ and record their parity in qubit 2 for the two three-body terms, resulting in

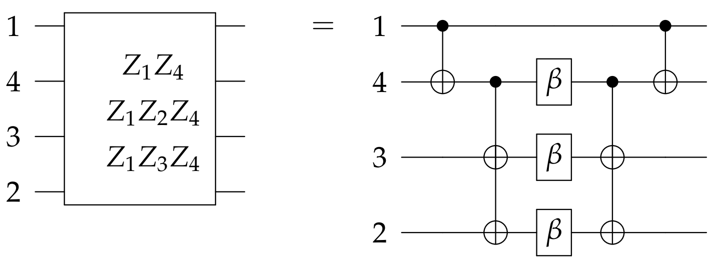

. (7.30)

so six CNOT gates with all-to-all connectivity (four can be removed using non-local R_{ZX} gates). With linear connectivity this takes only two more CNOT gates

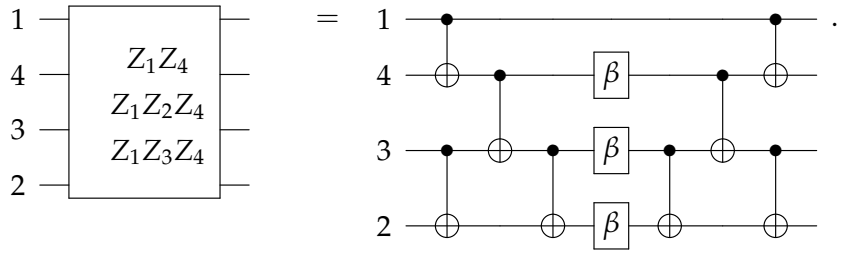

(7.31)

The $Z_2Z_3Z_4$ term can be done directly in linear connectivity with four CNOT gates (or two plus one R_{ZX}) as

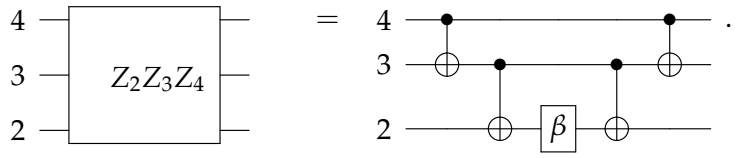

(7.32)

Ultimately, we do the final operation as described in Eq. (7.31), but inverted and performing just two rotations. It is important to observe that no explicit SWAP gates have been used. The reader may find this is a painful process and he would be correct. Hence, finding the best circuit transpilation is a difficult problem, and requires significant progress for quantum computers to be practical in general. Right now, automated tools are not efficient enough when the circuits lacks structure and are unable to efficiently use indirect swapping operation. We find that it is currently better to optimize the circuits by hand to achieve the best results. As a comparison, we surpass the Qiskit transpiler [291] by an order of magnitude.

7.3.2 *Noise model*

Prior to using the actual quantum computer, we evaluate PEV and NR on two noise models with varying strength. Figure 7.2 shows the absolute error relative to the intensity of the noise model for a depolarizing channel in panel a) and a scaled realistic fake backend from Qiskit [291] in panel b). In the first scenario, the circuit comprises a single double Trotter step including 56 CNOT gates, with $O_i = O_j = Z_1$. Conversely, in the second case we perform 12 double steps consisting of 452 CNOT gates in total. The black line shows an upper bound on the expected shot noise barrier. The shaded area correspond to a 95% confidence interval obtained over 30 repetitions.

The purified PEV approach demonstrates superior performance, crossing the shot noise barrier, at a value of p which is one hundred times bigger than without the error mitigation technique. Panel b) illustrates a little improvement for the NR approach attributable to Pauli twirling.

7.3.3 *Single step on hardware*

We now report data obtained on the superconducting quantum device ibmq_kolkata [292], which comprises 27 fixed-frequency transmon qubits, exhibiting fundamental transition frequencies of about 5 GHz and anharmonicities of -340 MHz. The median qubit lifespan T_1 is 109.86 μs, while the median coherence time T_2 is 58.95 μs. The qubits used for the

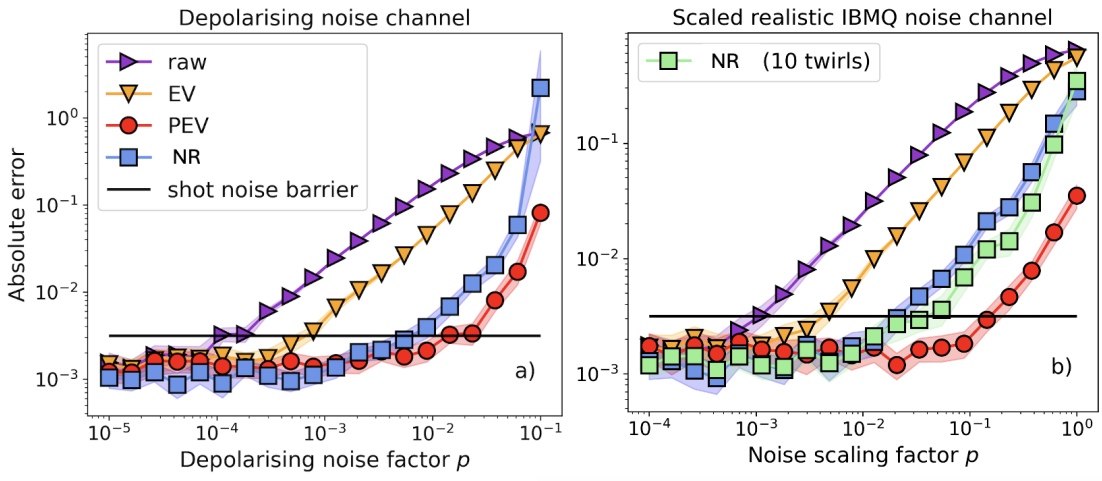

Figure 7.2: **Noise model:** Absolute error against the strength p of the noise model for different error mitigation strategy. Figure taken from Ref. [110]

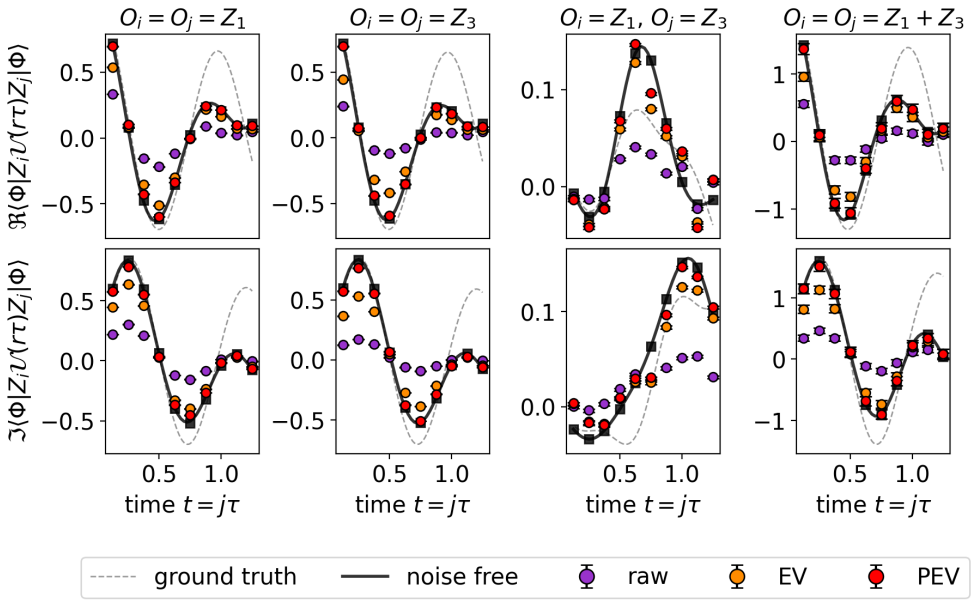

Figure 7.3: **Echo verification - single step:** The first two columns exhibit the diagonal components, the third column presents the off-diagonal component, and the last column contains the reconstructed moments. The real and imaginary components are shown in the first and second rows, respectively. The gray curve represents the ground truth from exact diagonalisation, the black dots denote the statevector simulation, while the colorful dots are derived from ibmq_kolkata using EV (orange) and PEV (red). The error bars represent one standard deviation calculated by Bayesian data augmentation. Figure taken from Ref. [110].

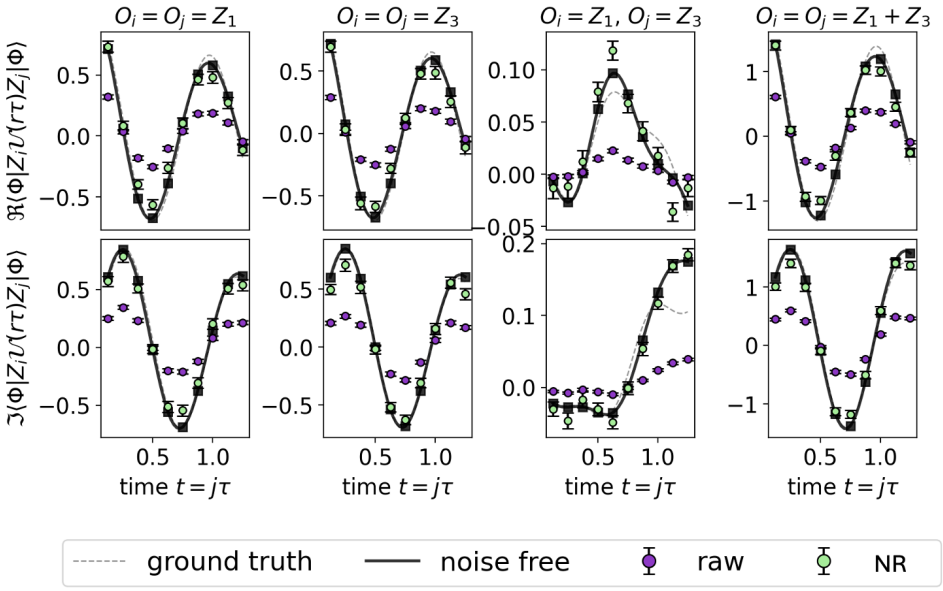

Figure 7.4: **Noise renormalization - single step:** The first two columns show the diagonal component, the third column illustrates the off-diagonal component, and the last column depicts the whole moment. The real and imaginary components are shown in the first and second rows, respectively. The gray curve represents the ground truth from exact diagonalisation, the black dots denote the statevector simulation, while the green dots are derived from ibmq_kolkata using NR. The error bars represent one standard deviation calculated by Bayesian inference. Figure taken from Ref. [110].

experiments are manually picked, exhibiting a CNOT error ranging from 6×10^{-3} to 9×10^{-3}, a readout error between 7×10^{-3} and 4×10^{-2}, and a sx error from 3×10^{-4} to 4×10^{-4}.

In the first experiment, we calculate the first ten moments using a double second-order Trotter step with $\tau = 0.125$. Error suppression and calibration methodologies, including XY-8 dynamical decoupling [88–92], pulse-efficient transpilation [97], Pauli twirling [93, 94] (16 twirls), control mitigation [111], and readout error mitigation [293], are consistently incorporated unless explicitly stated otherwise. Each circuit is executed using 10^5 shots. Information on the error mitigation and suppression strategies are available in Chapter 3. Although twirling is not essential when using PEV, it is used by default since it is customary practice.

Figure 7.3 illustrates the results obtained by PEV. The first two columns exhibit the diagonal components, the third column presents the off-diagonal component, and the last column displays the complete reconstructed moments. The real and imaginary components are shown in the first and second rows, respectively. The error bars represent one standard deviation calculated by Bayesian inference [106]. The results obtained with PEV align with the statevector simulation, while the raw data is severely damped. Secondly, the Trotter approximation becomes erroneous after $t = 0.5$, signifying that further Trotter steps are necessary to computer higher moments.

Figure 7.4 illustrates the same calculations using the NR approach. In this case NR seems to be equally successful, despite the off-diagonal terms exhibiting more variance. This may be explained by the error amplification phenomenon caused by estimating $(1 - p)$ on the diagonal elements, with slightly different circuits. Although the impact of Pauli whirling is not substantial, it seems to diminish the variation. Ultimately, the Trotter approximation

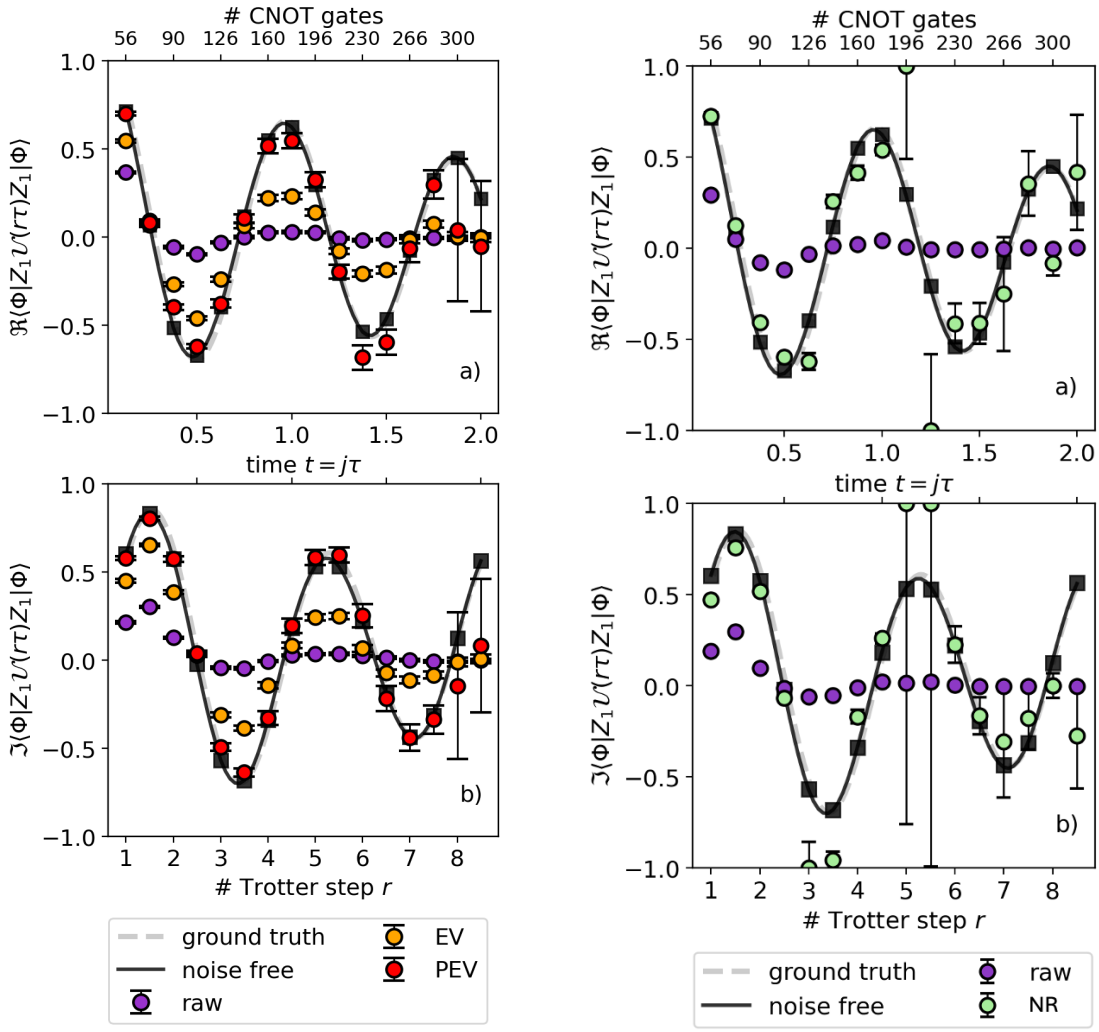

(a) PEV (b) NR

Figure 7.5: **Hardware results - multiple steps:** The upper (lower) panel shows the real (imaginary) part of the first diagonal contribution to the moment, with $O_l = O_k = O_1 = Z_1$ using PEV (left) and NR (right). The grey curve shows the ground truth from exact diagonalisation, while the black dotes are from statevector simulation. Figure taken from Ref. [110].

maintains its accuracy for extended time because of the factor of two fast forwarding from CRG.

7.3.3.1 *Multiple Trotter steps*

In our second experiment, we increase the number of Trotter steps to up to the capabilities' limit of the quantum computer. For simplicity, we only calculate the real and imaginary components of the diagonal contribution, with $O_l = O_k = O_1 = Z_1$, as shown on the left and right sides of Figure 7.5 for the PEV and NR mitigation strategies, respectively. We execute a maximum of eight double Trotter steps, culminating in a total of 300 CNOT gates.

Utilizing PEV yields correct results for up to seven steps, equivalent to 266 CNOT gates, after which significant deviations occur. The high variance after the breaking point may be explained by examining the purity of the ancilla qubit, shown in the left panel of Figure 7.6.

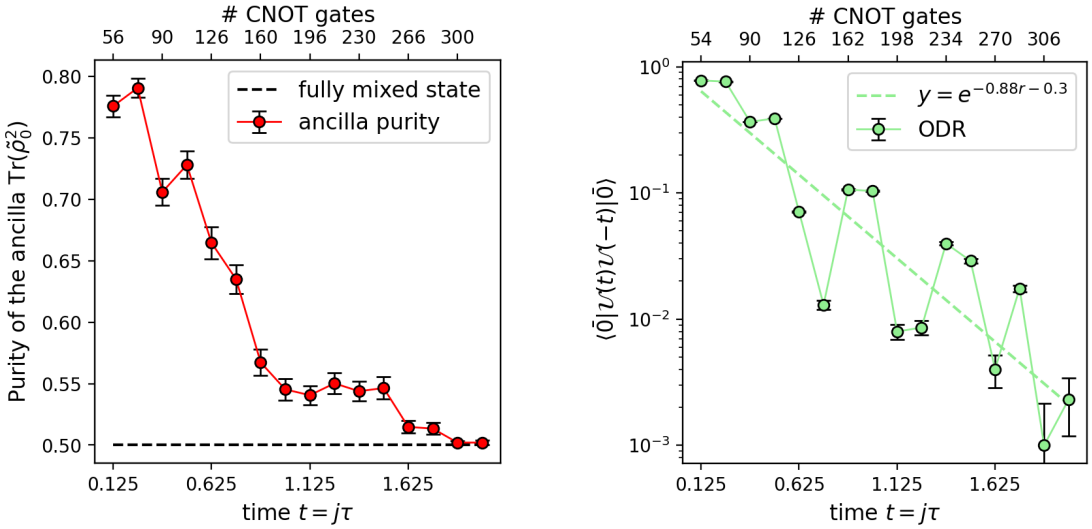

(a) Purity of the ancilla

(b) Renormalization factor

Figure 7.6: **Hardware results - Analysis:** The left panel shows the purity of than ancilla qubit as a function of the number of Trotter steps, while the right panel displays the renormalization factor. Figure taken from Ref. [110].

We observe that the purity is decaying to 0.5, which is the sign that we reached a completely mixed state. When purifying, the state is randomly projected onto the two eigenstates. Importantly, the purity serves as a effective tool to diagnostic the reliability of the strategy.. Hence, one may verify whether the purity exceeds 0.5 to provide a degree of assurance about the protocol's success.

The NR approach is less reliable, failing after just two double steps, and demonstrates more variation, particularly in the absence of Pauli twirling. This is explained by the fact that, when the noise level is high, the division by $(1 - p)$ becomes unstable. To investigate this effect, we plot the renormalization factor as a function of time (and the number of Trotter steps r) in the right panel of Figure 7.6. We observe that it decays exponentially with depth, making the sampling more costly as the number of Trotter steps increases.

7.3.4 *Discussions*

To summaries, our study highlights the importance of Fourier moments in understanding physical properties, especially their use in calculating response functions. While quantum computers inherently enable the computation of Fourier moments, performing controlled time evolution remains challenging, which results in an increase in CNOT gate overhead. In this section, we studied means to overcome these challenges by using CRG. We incorporate the PEV and NR protocols with this framework for the calculation of moments. We conclude that PEV is an effective strategy for this task, and is reliable until the purity of the ancilla reaches 0.5. One drawback of PEV is that it only a single observable might be inferred from the measurement, which is not the case for NR. This is not an issue when computing Fourier moments, but somewhat limits its application to more popular problems such as variational state preparation, where commutating terms in the Hamiltonian can be evaluated simultaneously.

7.4 GROUND STATE ENERGY ESTIMATION

In Section 4, we previously examined the use of the variational eigensolver for the task of GSEE. Notwithstanding its extensive applications [104, 134, 136, 139–143, 146, 294], scaling to larger systems poses significant challenges due to considerable sample requirements [145, 147] and issues related to trainability [150, 151, 153]. Therefore, the cost of running the VQE becomes quite extensive when we seek to estimate precisely ground-state energy of large systems. However, VQE might still be useful to prepare initial state for more sophisticated algorithms.

One important example of such algorithms is QPE [31], which samples eigenvalues from the distribution generated by an initial state. It comes with theoretical guarantees, and a significant overhead in circuit's depth.

Alternatives to QPE have been developed to improve applications on early fault-tolerant quantum computers. One example, is to perform a particular Hamiltonian transformation that which exhibits its spectrum. For instance, the method presented by Lin and Tong [266], drawing inspiration from Ref. [295], constructs the CDF of the spectral measure. The discontinuities are then associated with the eigenvalues of the underlying Hamiltonian. This algorithm, referred to in the following as the Lin and Tong (LT) algorithm, computes the expectation value of the Heaviside function $\Theta(\cdot)$ of the Hamiltonian for a given initial state $\langle\Psi_0|\Theta(H)|\Psi_0\rangle$, by expressing the Heaviside function as a Fourier series and calculating the Fourier moments on a quantum computer.

These methods have been further extended and improved by using the error function instead of the Heaviside [296], introducing Gaussian kernels [**ding_gaussian**, 297, 298], employing the quantum eigenvalue transform of unitaries (QETU) [274], via rejection sampling [299] or by exploring implementation on real quantum hardware [300, 301] or against noise models [302]. Even if these methods aim at addressing the problem of limited resources, they all assume access to a lower bound η on the overlap between the initial state and true ground state. It remains unclear how this algorithm would perform in practice without this information. In the following, we instead advocate on *improving* the energy estimation associated with an initial state obtained from classical methods such as DMRG [3], rather than solving for the exact ground state [303], which is known to be a QMA-complete problem [126, 304–306]. The following is based on Ref. [267].

7.4.1 *Construction*

We will now explain the LT algorithm with greater details, aiming to draw the state of the art protocol of this algorithm. We begin by defining the problem at hand and outlining the methodologies presented in Refs. [266, 296, 297]. We study a Hamiltonian H characterized by the following eigen decomposition

$$\tau H = \sum_k \lambda_k \Pi_k, \tag{7.33}$$

where τ is chosen to normalize H as $\tau||H|| < \pi/2$, and $\Pi_k = |E_k\rangle\langle E_k|$ are the projectors onto eigenspace associated with scaled eigenvalues $\lambda_k = \tau E_k$, ordered as $\lambda_0 \leq \lambda_1 \leq \cdots \leq \lambda_k$. Furthermore, we presume access to an initial state $\rho = |\Psi_0\rangle\langle\Psi_0|$ and express its spectral measure for H in accordance with Eq. (7.33) as

$$p(x) = \sum_k p_k \tilde{\delta}(x - \lambda_k) = \sum_k \mathrm{Tr}[\rho \Pi_k] \tilde{\delta}(x - \lambda_k), \tag{7.34}$$

where $\tilde{\delta}(\cdot)$ is the Dirac delta function and $0 < \eta \leq p_0 \equiv |\langle E_0|\Psi\rangle|^2$ a lower limit on the ground-state overlap. We aim to solve the following problem.

Problem 1. *Given a precision $\delta > 0$ and lower bound on the overlap parameter $\eta > 0$, we seek to decide if*

$$Tr[\rho \Pi_{\leq x-\delta}] < \eta \quad or \quad Tr[\rho \Pi_{\leq x+\delta}] > 0. \tag{7.35}$$

More simply put, we want to determine if the ground state energy satisfies $\lambda_0 \in [x - \delta, x + \delta]$. Addressing this issue allows for the estimation of the ground state energy by binary search throughout an energy grid [266]. The LT algorithm approximates the CDF of the spectral measure as

$$C(x) = \sum_{i:\lambda_i \leq x} p_i, \tag{7.36}$$

with a Fourier series whose moments can be obtained on a quantum computer. The approximation error ϵ is related to the precision as $\epsilon = \delta/\tau$. We can subsequently extract the eigenvalues of the Hamiltonian, since they manifest as discontinuities, or *jumps*, in the CDF.

The calculation of the CDF entails convolving the Heaviside step function $\Theta(x)$ with the spectral density $p(x) = \sum_k p_k \tilde{\delta}(x - \tau\lambda_k)$, resulting in $C(x) = p(x) * \Theta(x)$. We calculate its approximation as a Fourier series by assessing Fourier moments on the quantum computer and then summing them, weighted by the Fourier coefficients, on a classical device. The Fourier series approximation of the error function $(\cdot)$, with precision fixed by the parameter β, is derived as [296]

$$F(x;\beta) = \sum_{|k| \leq D} F_k(\beta) e^{ikx}, \tag{7.37}$$

where D is the number of Fourier moments, or runtime, of the algorithm, which dictates the approximation's accuracy, and $F_k(\beta)$ [296] the Fourier coefficients

$$\begin{aligned} F_0(\beta) &= 1/2, \\ F_{2j+1}(\beta) &= -i\sqrt{\frac{\beta}{2\pi}} e^{-\beta} \frac{I_j(\beta) + I_{j+1}(\beta)}{2j+1}, \text{ and} \\ F_{2d+1}(\beta) &= -i\sqrt{\frac{\beta}{2\pi}} e^{-\beta} \frac{I_d(\beta)}{2d+1}. \end{aligned} \tag{7.38}$$

Here d is related to D as $2d + 1 = D$, and $I_n(\beta)$ is the n-th modified Bessel function of the first kind. The smoothed Heaviside function can be seen as a scaled and shifted version of the error function, while the parameter β controls convergence between the scaled error function towards the Heaviside function. In order to guarantee an approximation error ϵ, one can take [296]

$$\beta = \max\left[1, \frac{1}{4\sin^2\delta} W_0\left(\frac{2}{\pi\epsilon^2}\right)\right], \tag{7.39}$$

where $W_0(\cdot)$ is the principal branch of the Lambert-W function, together with a number of terms scaling as

$$D = \mathcal{O}\left(\delta^{-1}\log\left(\epsilon^{-1}\right)\right). \tag{7.40}$$

We direct the reader to [296, Appendix 1] for more information and a precise formulation of this mathematical construction. In practice, β interpolates between the scaled error function and the Heaviside. An approximate periodic CDF can be constructed using these coefficients as

$$\tilde{C}(x) = \int_{-\pi/2}^{\pi/2} p(y)F(x-y)dy = \sum_{|k|\leq D} F_k e^{ikx}\langle\Psi_0|e^{-i\tau k\mathcal{H}}|\Psi_0\rangle = \sum_{|k|\leq D} F_k e^{ikx} m_k(\tau). \tag{7.41}$$

Noting that $m_{-k} = m_k^*$, $F_{2k} = 0$, $F_k = -i|F_k|$ for $k > 0$ and $F_k = i|F_k|$ for $k < 0$, the ACDF can be written as

$$\tilde{C}(x) = \frac{1}{2} + 2\sum_{k=1}^{d} |F_j|\left(\Re[g_j]\sin jx + \Im[g_j]\cos jx\right), \tag{7.42}$$

where $j \equiv 2k+1$ and m_j are the Fourier moments.

Finally, to bring the convergence from $\mathcal{O}(\epsilon^{-2})$ down to $\mathcal{O}(\epsilon^{-1})$, and therefore reach the Heisenberg scaling limit, the sum can be computed using importance sampling [266]. More precisely, we sample a set of M integers $\{k_1, \ldots, k_M\}$, with each $k \in \{1, 2, \ldots, \lceil (D-1)/2 \rceil\}$. The probability of observing k is equal to the normalized Fourier coefficients $|F_k|/\mathcal{F}$ (with the normalization factor $\mathcal{F} = \sum_{k=1}^{d} |F_{2k+1}|$). The corresponding Fourier moments g_k are then measured once and averaged together to build the ACDF)

$$G(x) = \frac{1}{2} + \frac{2\mathcal{F}}{M}\sum_{i=1}^{M}\left[\Re[g_{k_i}(\tau)]\sin(k_i x) + \Im[g_{k_i}(\tau)]\cos(k_i x)\right]. \tag{7.43}$$

Here, M corresponds to the total sample complexity (i.e. number of shots) required by the algorithm. Following the same steps as in [266], one can easily show that the variance of G is bounded by $2\mathcal{F}^2/M$. Once an estimate $\tilde{E}$ is found, we can refine it by computing the derivative of the ACDF and take the maximum on an interval $[\tilde{E}-\delta, \tilde{E}+\delta]$ of the CDF's derivative [300]

$$G'(x) = \frac{2\mathcal{F}}{M}\sum_{i=1}^{M} k_i\left[\Re[g_{k_i}(\tau)]\cos k_i x - \Im[g_{k_i}(\tau)]\sin k_i x\right]. \tag{7.44}$$

We note that this is similar in spirit to identifying the maxima of a Gaussian kernel placed around the energy guess $\tilde{E}$ [297].

An overview of the whole algorithm is shown in Figure 7.7, showing the input parameters, reconstruction of ACDF and the detection of inflection point, which we cover in greater detail in the next section. The Fourier moments, see Eq. (7.4), are calculated on the quantum computer using time indices j sampled according to the Fourier coefficients of the approximation error function. The ACDF is constructed by summing the moments weighted with their respective coefficients, as seen in Eq. (7.42). The ground state energy $\tilde{E}_0$ is ultimately determined by first identifying the inflection point using *ruptures*, validated by the statistical ANOVA method, and subsequently refining this estimate using the point exhibiting the largest gradient inside its δ-neighborhood (Sec. 7.4.2.3).

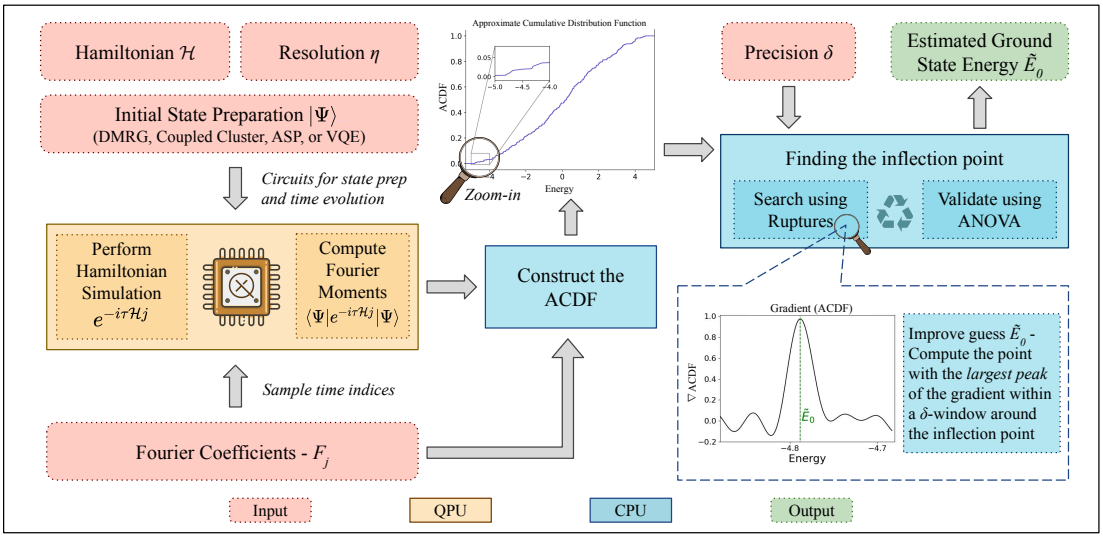

Figure 7.7: **Workflow of the algorithm.** Figure taken from Ref. [267].

7.4.2 *The LT algorithm in practice*

The LT algorithm serves as a prospective ground-state energy estimate method for early FTQC devices. Despite the established capability of this technique to ascertain the ground-state energy using $M = \mathcal{O}(1/\eta^2)$ samples and a maximal depth of $D = \mathcal{O}(\delta^{-1} \log \delta^{-1} \eta^{-1})$ [266], several obstacles must be addressed for its practical applicability. This section delineates the practical issues and proposes feasible solutions.

Many practical issues stem from the presumption that we possess a starting state exhibiting a minimum overlap of η with the true ground state, where η is sufficiently large. Nonetheless, this is often unrealistic for two reasons:

1. preparing starting states with substantial overlaps is challenging.
2. there is an absence of methods to evaluate the constraint η on the overlap parameter p_0.

7.4.2.1 *Initial State Preparation*

The quality of the initial state, particularly regarding its overlap with the actual ground state, substantially limits the effectiveness of QPE and its variations. Consequently, the technique used to create the initial states is crucial. The two main methodologies entail either developing a known quantum state using quantum techniques or directly implementing a wave function derived from conventional methods.

ASP techniques have been extensively researched for the production of ground states [155]. They include innovations such as shortcuts to adiabaticity [156], counterdiabatic driving strategies [157–159], and hybrid methods like counterdiabatic optimal local driving [160, 161]. Furthermore, VQE [130] and dissipative [307–311] methodologies retain significant potential. Nonetheless, the scalability of these methods to extensive system sizes has yet to be verified. For example, ASP requires an evolution time that is inversely related to the spectral gap, which might be exponentially small. In contrast, VQE suffers from trainability issues notably the barren plateau phenomena and high sample complexity. Dissipative approaches need connection to a bath or transformation to the Hamiltonian via QSP which

may be costly to execute. We delved more into the VQE in Chapter 4, as it is the most NISQ friendly initial state preparation algorithms.

The second category includes approaches that directly load state vectors, represented classically using methods such as DMRG [3, 312] and CC [128], onto the quantum computer. These algorithms provide a significant benefit by having cutoff parameters, such as the bond size χ for DMRG or the quantity of interacting orbitals for CC. Despite their significant scaling with system size, it is always possible to determine a cutoff such that the algorithms can run in reasonable time. This often produces a state of far better quality than a random guess or a mean-field state such as HF. This highly facilitates the extraction of low-lying eigenstate components by methods like LT or QPE. A crucial step in using conventional methods is the transfer of the quantum state to the quantum computer. Indeed, if the state vector is arbitrary, the cost is often exponential in the number of qubits. Better scaling can be attained if the state has certain distinctive structure. For instance, Refs [303, 313] devised a technique to load multi-Slater determinants and CC states, respectively. Conversely, states acquired using DMRG have an overhead which is at most polynomial in the bond dimension [314–317]. Additionally, techniques have been developed to augment initial state overlap, such as optimizing the orbital set [318] or eliminating undesirable contributions [319], with heuristics detailed in Ref. [320].

We choose to tackle this difficulty by sparsifying (or truncating) the classical wavefunction and use the sparse quantum state preparation proposed by Gleinig and Hoefler [321]. The sparsification process involves keeping only the S biggest components, and then renormalizing. The S-sparsified state, represented as $|\Psi_S\rangle$, may be efficiently loaded into the quantum computer with just $\mathcal{O}(SN)$ two-qubit gates and $\mathcal{O}(S \log S + N)$ single-qubit gates, hence avoiding exponential scaling. Although DMRG states may be loaded with a depth of $\mathcal{O}(n\chi^2)$ [314–317] in the worst cases, the sparsification technique may be beneficial if the prefactor is high. Furthermore, sparsifying the state allows us to investigate the regime in which DMRG fails to converge and acts as a proxy to examine the performance of LT algorithms in this context. While loading MPS states is efficient, situations might still exist where the sparsification is favorable. Particularly if the state only contains a few major contribution. However, there is no fundamental problems in using the full MPS state.

7.4.2.2 *Hamiltonian Simulation*

It is standard practice to measure resource needs in relation to calls to the time evolution oracle $\mathcal{U}(n\tau)$. To maintain Heisenberg scaling, it is essential to employ asymptotically optimal techniques, e.g. by QSP[33], typically scaling as $\mathcal{O}(\|\mathcal{H}\|_1 \tau D + \log \epsilon^{-1})$ queries to the LCU decomposition. Nevertheless, they entail considerable disadvantages, including the need for several supplementary qubits, high connectivity, and often a large prefactor. Consequently, we choose to use PF, see Section 5, based on p-order Trotter-Suzuki decomposition [322]. These PF exhibit an error that scales as $\epsilon \leq C(\tau D)^{p+1}/r^p$, where r denotes the number of Trotter steps and C is a prefactor contingent upon the commutators [193], which can be specifically computed for each system, such as spins [190], the Hubbard model [323], or neutrinos [324]. This decision is driven by two main factors: firstly, PFs do not need ancilla overhead or expensive control operations, and they often exceed theoretical guarantees, even when constrained to the low-energy regime [325]. To limit the error to below ϵ, we may choose

$$r = \left\lceil C^{1/p} (\tau D)^{(1+1/p)} \epsilon^{-1/p} \right\rceil . \tag{7.45}$$

Algorithm 2: FIND SMALLEST BREAKPOINT

1.25
Input: signal $\{y_i\}_{i=1}^{n}$, failure probability $\alpha \in [0,1]$
`FindBP`(y, α) *[r]Finds breakpoint (Eq. 7.47) $b \leftarrow n$;
significant $\leftarrow$ True ;
while significant **do**
  $\tilde{b} =$ `FindBP`$(y[:b], \alpha)$ *[r]Optimization significant = `ANOVA` (y, b, α) *[r]Validation
  if significant **then**
    $b \leftarrow \tilde{b}$ *[r]Upon successful validation
  end
end
Output: b

In doing so, we forfeit Heisenberg scaling, but keep many important advantages suiting the limitations of early FTQC devices. For d-local Hamiltonians, PFs may use vanishing commutators to get superior scaling compared to qubitization. Furthermore, due to the decay of Fourier coefficients at a rate of $1/x$, we are mainly conducting short-time simulations, which is where PFs shine the most. This perception is corroborated by our numerical tests in Section 7.4.2.5, where accurate energy estimate is achieved even for circuits limited to relatively shallow depths. In the present thesis, we only use PFs as they are surprisingly efficient in practice.

While we do not explicitly explore these techniques here, numerous enhancements for product formulas have been proposed, including but not limited to randomization [202, 326, 327], multi-product formulas [203, 328], qDRIFT compilation [194–196, 329], and composite formulas [199, 330]. As these strategies exchange circuit depth for supplementary measurements, they are particularly appropriate for early FTQC, necessitating further exploration in more targeted studies.

7.4.2.3 *Detection of discontinuities*

The second primary obstacle we encounter is the identification of discontinuities in the ACDF, which is directly linked to the difficulty of calculating a precise lower bound η on the overlap with the actual ground state. The issue is twofold:

1. It remains challenging to estimate η, which is necessary for inverting the CDF as outlined in [266]. We note that recent advances are proposing efficient ways to do so when using DMRG [34], by comparing initial states of low and high bond dimension.

2. Since the number of jumps tend to increase exponentially with the system size, the CDF becomes quasi-continuous, making the jump detection procedure difficult.

This challenges are closely associated with the orthogonality catastrophe [331–333], which asserts that the overlap between a random beginning state and any eigenstate diminishes exponentially rapidly in bigger systems. The continuous characteristic of the CDF is influenced by the magnitude of the support of the initial state in the eigenbasis. Consequently, bound states, or more broadly, states that are sparse in the eigenbasis, will result in a well behaved step function. Nevertheless, since the preparation of initial states of this caliber continues to be an unresolved issue, it is crucial to examine states with exponential support.

Algorithm 3: ANOVA

1.25

Input: signal $\{y_i\}_{i=1}^n$, breaking point $b \in [1, n]$, failure probability $\alpha \in [0, 1]$

$\bar{y} = \frac{1}{n}\sum_{i=1}^{n} y_i$ *[r]Mean of signal components $\bar{y}_0 = \frac{1}{b}\sum_{i=1}^{b} y_i$ *[r]Mean of first segment $\bar{y}_1 = \frac{1}{n-b}\sum_{i\geq b} y_i$ *[r]Mean of second segment $SS_w = \sum_{i=1}^{n}(y_i - \bar{y})^2$ *[r]Sum of squares within segments $SS_b = \sum_{i<b}(y_i - \bar{y}_0)^2 + \sum_{i\geq b}(y_i - \bar{y}_1)^2$ *[r]Sum of squares between segments $f = (n-2)SS_b/SS_w$ *[r]F-ratio significant $\leftarrow$ False ;

if $F(f, 1, n-2) > (1-\alpha)$ **then**

  if $\bar{y}_1 > \bar{y}_0$ **then**

    significant $\leftarrow$ True *[r]Successful F-test & monotone function

  end

end

Output: significant $\in$ {True, False}

We adopt a different perspective, seeking to understand the practical implications of employing LT with an initial state of unknown quality. Rather than striving for the exact determination of the ground state energy, we focus on enhancing the energy estimate of the initial state.

To this end, we aim to locate the inflection point of the CDF, which we define as the first energy value which has a non-zero density with high probability. We argue that, in general, the inflection point provides an improvement over the energy expectation value of the initial state [303]. In fact, the expectation value correspond to the mean energy of the state, while the inflection point correspond to the low-energy part in the density of states. Moreover, since we make sure that the density is non-zero using statistical tests, we ensure that the inflection upper bound the true ground-state energy with high enough probability. We address this new task by identifying a statistically significant increase in the ACDF comprised of small contributions from neighboring eigenstates. Intuitively, we are turning the question around and ask ourself what is best estimate $E_{c(\eta)}$ we can do within the available resources. Given a maximal depth D, the minimal increase which can be detected is directly lower bounded by the approximation error of the Heaviside function $\epsilon(D)$. We can then define the projector on the low-energy regime

$$\begin{aligned} \mathcal{P}(\eta) &= \sum_{i=0}^{c(\eta)} |E_i\rangle\langle E_i| \\ c(\eta) &= \min n \in \mathbb{N} : \sum_{i=0}^{n} |\langle\Psi|E_i\rangle\langle E_i|\Psi\rangle| \geq \eta > 2\epsilon(D) \end{aligned} \tag{7.46}$$

where $c(\eta)$ is a truncation such that the overlap of the initial states in the low-energy manifold is higher than η. By turning the problem around, we are asking what is the best we can do given limited resources, which is much more close to a practical scenario. We emphasize that the accumulation is technical not required in the algorithm, but rather serve as a pedagogical tool illustrating its aim.

We propose a procedure based on a kernel change-point detection method [334, 335]; this is a signal processing technique to detect changes in the mean of a given signal, implemented with the software package *ruptures* [336]. Despite its simplicity, the technique can be used to perform complex time series analysis, yielding great results across multiple scientific domains [337–340]. In this scenario, we execute an iterative procedure (Algorithm 2) that comprises two primary steps:

1. Dividing the signal into two segments, which entails identifying the breakpoint that optimally separates the time series.

2. Evaluating the statistical significance of the split using the ANOVA-based scheme.

If the step is validated, we repeat on the left part of the signal (up to the recently detected split) and find a new breakpoint, until one is finally rejected. In other words, we aim to find the first significant step of the ACDF.

To find the breakpoint, the signal $y \in \mathbb{R}^n$ is first mapped onto a reproducing Hilbert space with $\phi : \mathbb{R} \to \mathbb{H}$, implicitly defined as $\phi(y) = K(y, \cdot)$. The kernel function K, typically chosen to be Gaussian, induces the metric on the Hilbert space. We then solve the following minimization problem, by embedding the signal in the Hilbert space, to find the smallest breakpoint:

$$\min_{b \in \{1,\dots,n\}} \sum_{t=1}^{b} \|\phi(y_t) - \bar{y}_{1:b}\| + \sum_{t=b+1}^{n} \|\phi(y_t) - \bar{y}_{b+1:n}\|\,, \tag{7.47}$$

where $\bar{y}_{a:b} = \sum_{x=a}^{b} y_x/(b-a)$ is the mean of the signal between a and b. In the context of this work, the signal is $y_t = C(t) = p(t) * \Theta(t)$, where $\pi/2 < t < \pi/2$ belongs to the energy grid used to compute the first estimate. We guide the reader to Ref. [334] for an in-depth description of this technique.

The ANOVA algorithm (Algorithm 3) is used to decide if the breakpoint is statically significant by computing the mean of the two segments, as well as their standard deviations and their ratio f. The breakpoint is accepted if the p value of an F-test [341] with input $(f, 1, n-2)$ is greater than $1-\alpha$, with α being a small tolerated failure probability. We recall that the F-test is a statistical test for comparing the variances or standard deviations from two populations using the F-distribution. Since we require the CDF to be monotone increasing, the breakpoint b is only accepted if $\bar{y}_{1:b} > \bar{y}_{b+1:n}$, in addition to passing the F-test.

Furthermore, we avoid overshooting by stopping the procedure if the left part of the breakpoint is, on average, smaller than a threshold. To this end, we stop the point search if $\bar{y}_{b-l:b} > k\sigma + \tilde{\epsilon}$, for $k \in \mathbb{N}$, l a small window (for example, $l = 20$), and σ and $\tilde{\epsilon}$ are the empirical standard deviation and error, respectively, of the signal computed on the energy region $-\pi \leq x < -\pi/2$, where we know that no eigenvalues are present. Using these empirical estimates has the advantage of being closer to the real deviation than the theoretical upper bounds. The value of k sets the confidence of the predictions and is usually chosen as $k \in \{1, 2, 3\}$. In order to make the scheme more robust, we repeat the above procedure multiple times using different data for the ACDF and compute the median of means. Once the inflection point $\tilde{b}$ is known, the energy guess is chosen as the maxima of the gradient of the approximate CDF around a small window δ around the inflection point, $[\tilde{b} - \delta/2,\ \tilde{b} + \delta/2]$, as depicted in Figure 7.7.

In summary, the signal is processed using ruptures to find a breakpoint, which is then validated according to a statistical test subject to a monotone function condition. If accepted, we perform the same analysis on the early part of the signal up to the breakpoint and discard the rest. We repeat these steps until a breakpoint is discarded and return the last proposed breakpoint. While this method does not have any theoretical guarantees, it can be employed without knowing η a priory, and is therefore applicable in practice. Furthermore, it is statistically robust and efficient in practice.

Since the eigenvalue is situated at the middle of the jump, it is important to look at the gradient of the CDF. The first significant maximum of the gradient exactly corresponds to

an eigenvalue, while the secondary peaks are due to the approximation with the Fourier series. The whole procedure described in this section can be summed up as the following steps [267]

1. Prepare the best possible ground state $|\Psi_0\rangle$ using the available methods in your toolbox, e.g., DMRG [3], CC [128], etc.
2. Load the state on the quantum computer, e.g., by sparsification [321], or by encoding them as sums of Slater determinants or MPS [303].
3. Sample M time indices from the coefficients of the Fourier series [Eq. (7.38)].
4. Use a product formula to compute one sample of the corresponding Fourier moment with a suitable number of steps [Eq. (7.45)].
5. Build the ACDF using M samples [Eq. (7.43)].
6. Identify the inflection point x of the ACDF using *ruptures* [Eq. (7.47)].
7. The energy estimate is the maxima of the ACDF's derivative in an δ-window around x.

7.4.2.4 *Resources estimates*

Given the difficulties associated with the practical application of the LT method, particularly in ensuring tight lower limits η and the possibly little overlap between the initial state and the true ground state, we endeavor to determine the optimal outcomes achievable with constrained resources. We address the following issue.

Problem 2. *Considering a maximum depth D and a maximal shot budget M, what is the optimal upper bound on the ground-state energy that the algorithm can yield?*

This section presents quantitative resource estimates for the LT algorithm, categorized into two parts: the highest Fourier moment D, which is linked with the maximal evolution time of the Hamiltonian simulation through $T_{\max} = \tau D$, and influences the bias in approximating the CDF; and the number of samples M, linked to the variance of the output energy. We see that D is closely associated with the maximum depth of the quantum circuits, which is why we also refer to it as maximal runtime. The first result is an estimation of the maximal runtime [267], adapted from [296].

Theorem 8. *For a given accuracy $\epsilon \in (0, 1)$, the maximal runtime D required to approximate the error function $(\cdot)$ with ϵ-error is given by*

$$D = 2 \cdot \left\lceil \sqrt{f\left(\beta,\, \min\left[1,\, 4e^{-w(\epsilon)/2}\right]\right) \cdot w(\epsilon)} \right\rceil + 1, \tag{7.48}$$

where β is chosen according to Eq. (7.39) with

$$w(\epsilon) = W_0\left(\frac{18}{\pi\epsilon^2}\right) \text{ and } f(\beta, \epsilon) = -\frac{\ln\epsilon + \beta}{W_0\left(-\left[1 + \beta^{-1}\ln\epsilon\right] e^{-1}\right)}, \tag{7.49}$$

with $W_0(\cdot)$ being the principal branch of the Lambert-W function.

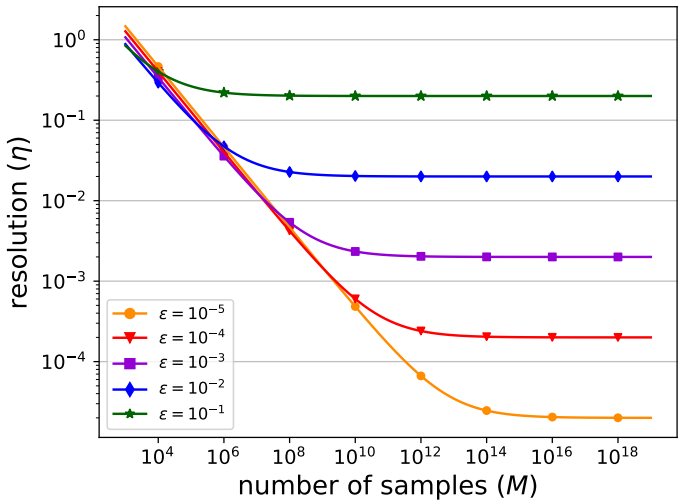

Figure 7.8: **Estimating the resolution.** Estimation of the resolution η of a discernible jump in the context of the fully connected Heisenberg model with 26 spins, as mentioned in Sec. 7.4.2.5, using a specified number of samples M with a confidence level of $1 - \vartheta = 95\%$ for various precision values ϵ. Figure taken from Ref. [267].

Proof. The proof [267] of this theorem is based on Ref. [296, *Theorem 3*, Appendix A], which provides a Fourier approximation to the Heaviside function with maximal error $\max |\Theta(x) - F(x;\beta)| \leq (\epsilon_1 + \epsilon_2 + \epsilon_3)/2 = \epsilon$ [296, Eq. (A12)]. Here, $\epsilon_{1,2}$ represent the errors arising from finite truncation d and the scaling parameter β, respectively, in approximating the error function, while ϵ_3 accounts for the error in approximating the Heaviside function with the error function.

For the choice $\epsilon_1 = \epsilon_2 = \epsilon_3 = 2\epsilon/3$ in this approximation, and using the values of β and d such that of $-(\epsilon_1 + \epsilon_2)/2 \leq F_d(x,\beta) \leq 1 + (\epsilon_1 + \epsilon_2)/2$ [296, Eq. (A13)], we get $d \geq \sqrt{t \cdot w(\epsilon)}$ [296, Eq. (A6)] for an integer t where

$$t = \begin{cases} f(\beta, 4e^{-w(\epsilon)/2}) & \text{for } w(\epsilon) \geq 2\ln 4 \\ \beta & \text{for } w(\epsilon) < 2\ln 4 \end{cases}, \tag{7.50}$$

with the function $f(\cdot,\cdot)$ defined above. Since $f(\beta,1) = \beta$ and $\epsilon \in (0,1)$, one can rewrite Eq. (7.50) as $t = f(\beta, \min\left[1, 4e^{-w(\epsilon)/2}\right])$ and obtain the minimum maximal runtime $D = 2d + 1$ as $2\left\lceil \sqrt{t \cdot w(\epsilon)} \right\rceil + 1$. □

This maximum runtime can now be used to estimate the number of samples necessary to resolve an accumulation of size η. It is essential to note that η is no longer a lower bound on the overlap but rather a user-defined parameter that quantifies the desired resolution.

Theorem 9. *For a given maximal runtime D and accuracy $\epsilon \in (0,1)$, the number of samples M required to guarantee the correct result with probability $1 - \vartheta$ is*

$$M = \left\lceil 2 \cdot \left[\frac{2.07\pi^{-1}(\log 4D + \gamma) + 1}{\eta - 2\epsilon} \right]^2 \right\rceil. \tag{7.51}$$

where $\eta > 2\epsilon$ is the accumulation on the low-energy part of the spectrum and γ is the Euler-Mascheroni constant. The last assumption is required since a jump smaller than the error threshold cannot be resolved.

Proof. We make use of $\mathbb{E}[G] = \tilde{C}(x)$ from Eq. (7.42) to solve Problem 1 (Eq. (7.35)) as:

$$\begin{aligned} G < \zeta &\Rightarrow \mathrm{Tr}[\rho \Pi_{\leq x-\delta}] < \eta \\ G \geq \zeta &\Rightarrow \mathrm{Tr}[\rho \Pi_{\leq x+\delta}] > 0. \end{aligned} \tag{7.52}$$

We use $\zeta = \eta/2$ to attune for errors in estimating $G(x)$ from sampling (Eq. (7.43)) and decide whether $\tilde{C}(x) = 0$ or $\tilde{C}(x) \geq \eta$. Consequently, we can define the above as the probability $\mathbb{P}[G < \eta/2] < \kappa$ conditioned on $\tilde{C}(x) < \eta - \epsilon$ which implies $C(x-\delta) < \eta$, and use

the one-sided Chebyshev inequality to find

$$\begin{aligned} \mathbb{P}[G < \eta/2] < \mathbb{P}[G \leq \eta/2] &= \mathbb{P}[G \leq \tilde{C} - \eta/2 + \epsilon] \\ &\leq \frac{\mathrm{Var}[G]}{\mathrm{Var}[G] + (\eta/2 - \epsilon)^2} < \frac{\mathrm{Var}[G]}{(\eta/2-\epsilon)^2}. \end{aligned} \tag{7.53}$$

Using the upper bound on the variance of G, we finally obtain

$$\mathbb{P}[G < \eta/2] < \frac{8\mathcal{F}^2}{M(\eta - 2\epsilon)^2}. \tag{7.54}$$

To ensure that this probability is bounded by κ, it is sufficient to choose

$$M \geq \left(\frac{2\sqrt{2}\mathcal{F}}{\eta - 2\epsilon}\right)^2 \left(\frac{1}{\kappa}\right). \tag{7.55}$$

We can improve the dependence on κ by noting that the individual summands in the average defining G are bounded by $2\mathcal{F}/M$, which allows us to apply the Hoeffding inequality as follows.

$$\mathbb{P}[\bar{G} < \eta/2] < \exp\left(-\frac{2\cdot(\eta/2-\epsilon)^2}{M\cdot(2\mathcal{F}/M)^2}\right) \Rightarrow M \geq \left(\frac{2\sqrt{2}\mathcal{F}}{\eta-2\epsilon}\right)^2 \log\left(\frac{1}{\kappa}\right). \tag{7.56}$$

We can now use the upper bound for the norm of the Fourier coefficients $\mathcal{F}$ computed as [296]:

$$\mathcal{F} = \sum_{|j|\leq d} |\hat{F}_j(\beta)| \leq \frac{2.07}{2\pi}(H_D + 2\ln 2) + \frac{1}{2} \approx \frac{2.07}{2\pi}(\log(4D) + \gamma) + \frac{1}{2}, \tag{7.57}$$

where H_k is the k^{th} Harmonic number, which can be estimated as shown above using the Euler-Mascheroni constant $\gamma = 0.57721567$ [342].

Finally, we obtain the desired result by noting that the algorithm must be executed $\mathcal{O}(\log \delta^{-1})$ times to solve Problem 1 [266]. If the algorithm fails with probability at most κ, then to estimate the ground state of the system with a success probability of at least $1 - \vartheta$, we must have $\kappa^{-1} \leq \vartheta^{-1} \log \delta^{-1}$, where $\delta = \tau\epsilon$, as previously defined.

□

It is important to note that these estimates are meant to be used in estimating the error tolerance and resolution which can be achieved using limited resources D and M. For instance, we use the Theorem 9 in Fig. 7.8 to show the inflections of size η that can be

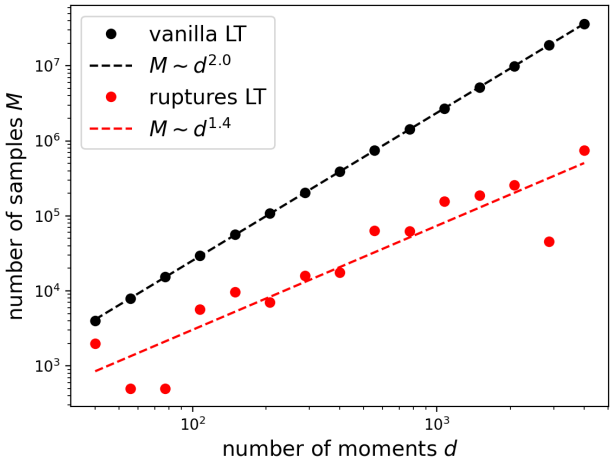

Figure 7.9: **Resources comparison.** Minimal number of samples required to detect a step of size $\eta = 4\epsilon(d)$ for an artificial signal. The black dots shows the vanilla LT procedure using the resources estimates of Theorem 9, while the red dots are the empirical results using ruptures, ANOVA and the safeguard. The failure probability is chosen as 0.05.

resolved for different error tolerances ϵ with a given sample budget M. We observe that the resolution which can be achieved plateaus after a certain sample budget, meaning that for further refined resolution, we need to increase the simulation time (i.e., decrease ϵ). In particular, this means that we can not trade depth with samples indefinitely, and that to reach higher resolution, it is required to increase the depth. As stated in Theorem 8, the maximum depth D is primarily determined by the error ϵ when approximating the step function. We leverage these resource estimates to determine the maximum number of samples required for step detection. Given a maximal number of moments d, which can be determined from the maximum depth supported by the device, the minimal detectable step size is given by $\eta = 4\epsilon(d)$, accounting for the approximation error. The number of samples required to achieve this level of precision is determined by Theorem 9. However, using the proposed signal processing technique, *ruptures*, we demonstrate that significantly fewer samples are required.

To validate this, we construct an artificial signal with two eigenvalues and a ground state overlap $\eta = 4\epsilon(d)$, and calculate the minimal number of samples required to find the first eigenvalue with precision $\delta = 0.05$. To enhance robustness, we employ a median-of-means estimator, incorporating its additional overhead into the total sample count. The results, presented in Fig. 7.9, illustrate that our automatic step detection procedure achieves a quadratic improvement over the vanilla approach. We will now delve deeper into this observation by conducting numerical simulations on a challenging system.

7.4.2.5 *Numerical simulations*

We consider a fully connected Heisenberg model with random couplings over N spins

$$\mathcal{H} = \frac{1}{N} \sum_{i<j} \sum_{a \in \{x,y,z\}} J_a^{ij} \cdot \sigma_a^i \sigma_a^j, \text{ where } J_a^{ij} \sim \mathcal{N}(0,1). \tag{7.58}$$

Here, σ_a^i is the corresponding Pauli matrix applied to the i-th qubit.

This model is universal and as such, can approximate any two-local Hamiltonian. Furthermore, tensor network approaches are anticipated to run into difficulties when

contracting indices due to the extensive amount of connections, making it an effective testbed for quantum algorithms. In this system, we simulate the dynamics implementing a second-order Trotter-Suzuki formula [322] with a time step of $\Delta t = \tau/r$, where $r = 8$ denotes the number of Trotter steps, and employ a SWAP networks-based [231, 232] circuit construction to achieve the most efficient circuit compression possible.

Employing a product formula for estimating time evolution has the benefit of maintaining a shallow circuit, thereby rendering it more appropriate for early FTQC systems. The depth of the quantum circuits used in the subsequent analyses is $2NrD$. We did not use any of the prospective enhancements to Hamiltonian simulation using product formulae outlined in Sec. 7.4.2.2, which are likely to reduce the maximum depth at the cost of requiring extra samples.

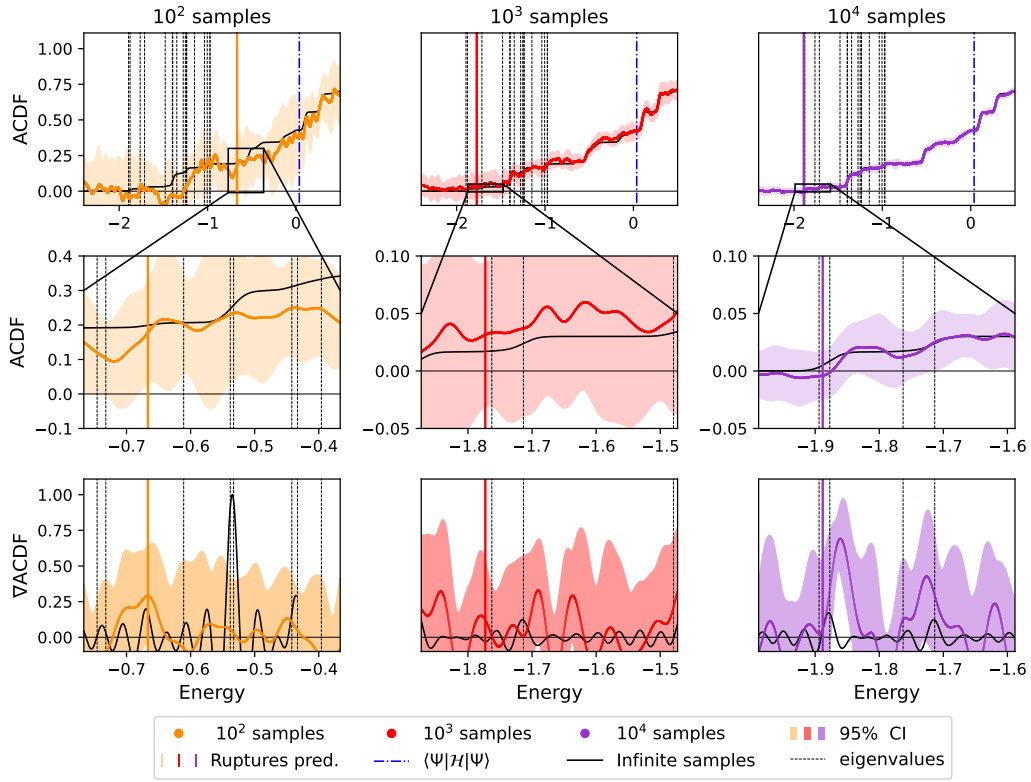

Figure 7.10: **Six spins, random initial state.** ACDF (top row), zoomed ACDF (middle row), and its gradient (∇ACDF, bottom row) calculated with varying sample sizes (per column). The first 20 eigenvalues derived with Exact Diagonalization (ED) are shown with dashed vertical lines, while the thick blue vertical line indicates the energy of the initial state. The colored line denotes the energy obtained using the *ruptures* technique. The black line is the ACDF (and its gradient) calculated using infinite statistics, while the shaded region indicates a 95% confidence interval derived from 10 independent runs. Figure taken from Ref. [267].

SMALL SYSTEM SIZES We start with a small system composed of six spins ($N = 6$). We choose a random initial state with $p_0 = 0.0014$ and $p_1 = 0.015$. The reader should keep in mind that a MPS state obtained with DMRG is likely to perform better, even for low bond dimension. We set $\epsilon = 0.055$ and hence require $D = 350$ Fourier moments for the construction of the approximate CDF.

The ACDF and its gradient are shown in Figure 7.4.2.5 for various number of sample M. The solid black line, shows the infinite statistics limit (the regime of infinite samples),

where al the moments are exact and appear in the summation. The ground-state energy is detected using the ruptures algorithms, explained in Section 7.4.2.3, and is later refined by maximizing the gradient in its neighborhood. We note that the inability to identify the ground state arises from the approximation error of the step function ϵ exceeding the overlap. Nonetheless, we determine the energy of the first excited state using only 10^4 shots and a random initial state.

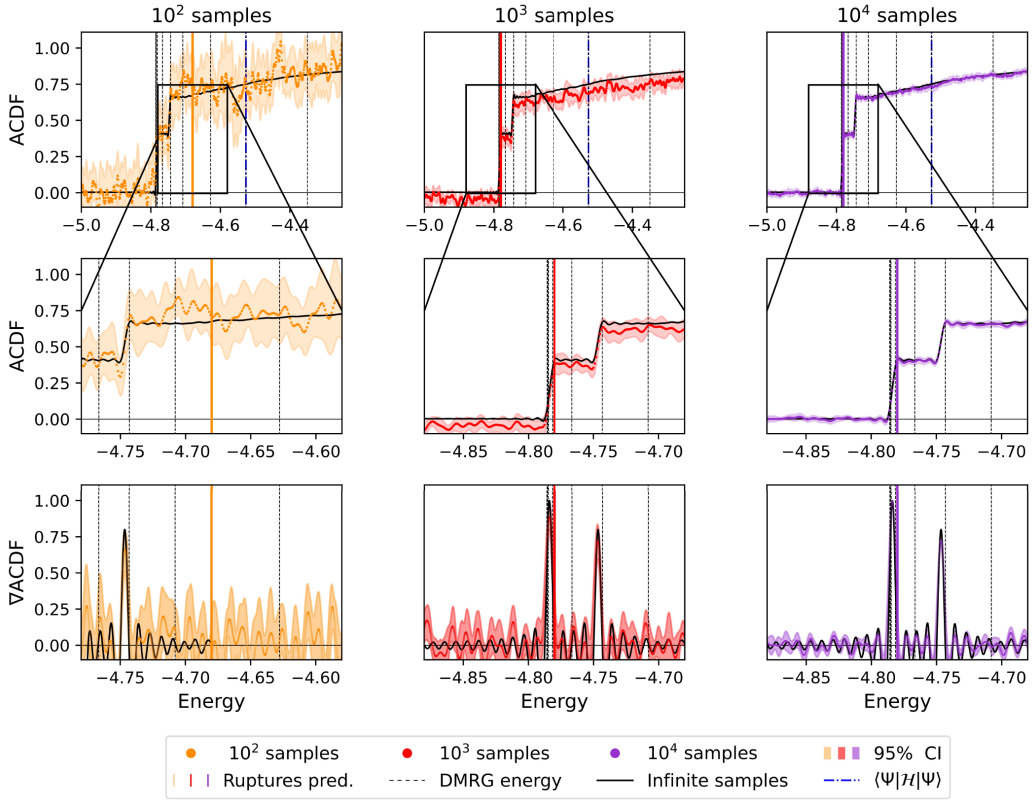

Figure 7.11: **26 spins, DMRG ($\chi = 10$).** ACDF (top row), magnified ACDF (middle row), and its gradient (∇ACDF, bottom row) calculated with varying sample sizes (per column). The energies derived from untruncated DMRG states with varying bond dimensions $5 \leq \chi \leq 2000$ are shown by dashed vertical lines, while the thick blue vertical line represents the energy of the initial state. The colored line denotes the energy obtained using the *ruptures* technique. The black line is the ACDF (and its gradient) calculated using infinite statistics, while the shaded area denotes a 95% confidence interval derived from 10 independent runs. Figure taken from Ref. [267].

LARGE SYSTEM SIZES We then shift to a larger model with $N = 26$ spins. Due to the prohibitive cost of ED in this regime, we use DMRG with a bond dimension of $\chi = 2000$ to define the target energy. We choose $\epsilon = 0.019$, resulting in $D = 6600$. We conduct two experiments: the first with the initial state chosen as $|\Psi_0\rangle = \text{DMRG}(\chi = 10)$, and the second from its sparsification with $S = 13$. The motivation for using DMRG with a low bond dimension is to get a good quality initial state, which can be calculated even for large system sizes. When scaling up, we are likely to increase the bond dimension but we will never be able to reproduce the exact ground state of challenging systems because of computing restrictions.

The outcomes using the DMRG state are shown in Figure 7.11, which presents the approximate CDF calculated with 10^2, 10^3, and 10^4 samples. The data illustrates the primary

contributions from low-energy eigenstates and a continuous contribution from higher-energy states. Our approach enables the recovery of the ground-state energy comparable to DMRG with bond dimension $\chi = 2000$. It is crucial to note that the CDF clearly displays two discontinuities in the lower segment of the spectrum. This results from the high quality of the initial state, including two low-energy bound states and a tail of high-energy residue. Consequently, the left segment of the CDF displays a stepwise configuration, while the right segment is continuous. Nonetheless, as seen in Figure 7.4.2.5, the cumulative distribution function derived from the sparsified DMRG state seems continuous, with discontinuities only becoming apparent upon magnification by a factor of one hundred. Crucially, in both instances, we enhance the DMRG estimate (blue dotted vertical line) by 5% and 32%, respectively.

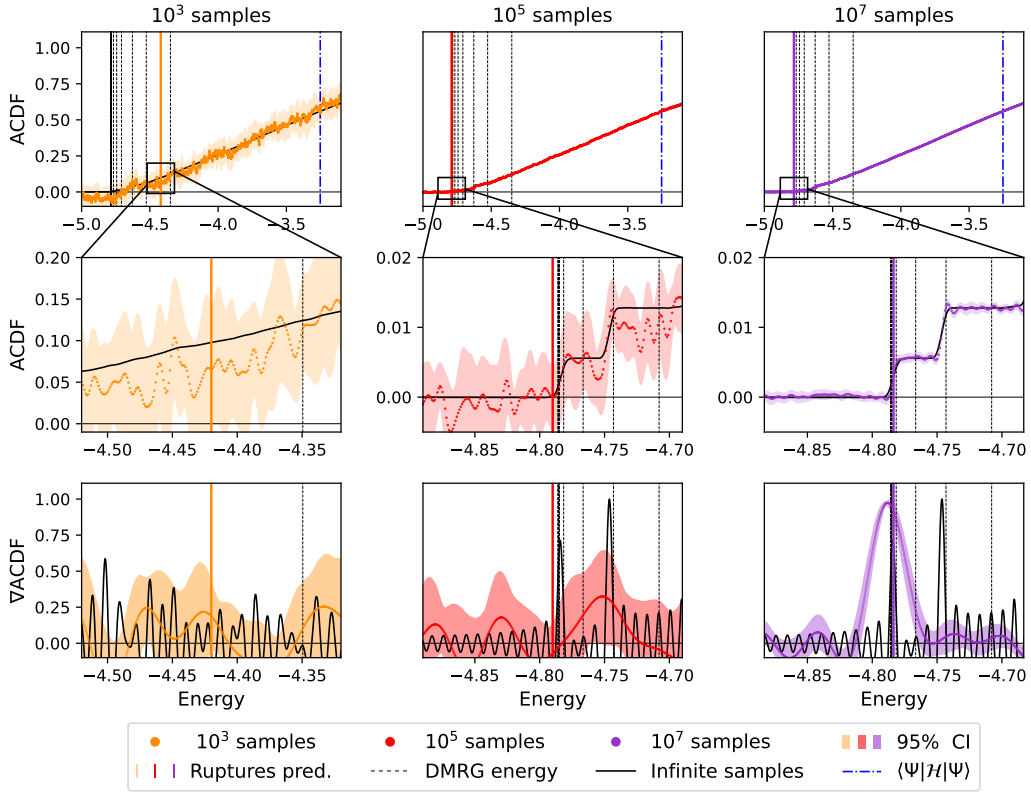

Figure 7.12: **26 spins, sparsified DMRG ($\chi = 10$).** ACDF (top row), zoomed ACDF (middle row), and its gradient (∇ACDF, bottom row) calculated with varying sample sizes (per column). The energies derived from untruncated DMRG states with varying bond dimensions $5 \leq \chi \leq 2000$ are shown by dashed vertical lines, while the thick blue vertical line represents the energy of the sparsified intial state with $S = 13$. The colored line denotes the energy obtained using the *ruptures* technique. The black line is the ACDF (and its gradient) calculated using infinite statistics, while the shaded area denotes a 95% confidence interval derived from ten independent runs. Figure taken from Ref. [267].

The primary difference in the aforementioned situation is that the approximate CDF does not exhibit explicit steps. This renders necessary to use our ruptures technique. We observe that a larger number of samples is required to estimate the ground state energy, of 10^5 and 10^7 samples to get the energy of DMRG with $\chi = 1000$ and $\chi = 2000$, respectively.

We examine the quality of the sparsified initial state $|\Psi_S\rangle$ by presenting the L^2 distance and the overlap with the original DMRG initial state as a function of the truncation parameter k in Figure 7.13. The observed gain in quality diminishes after $S = 13$, which justifies the

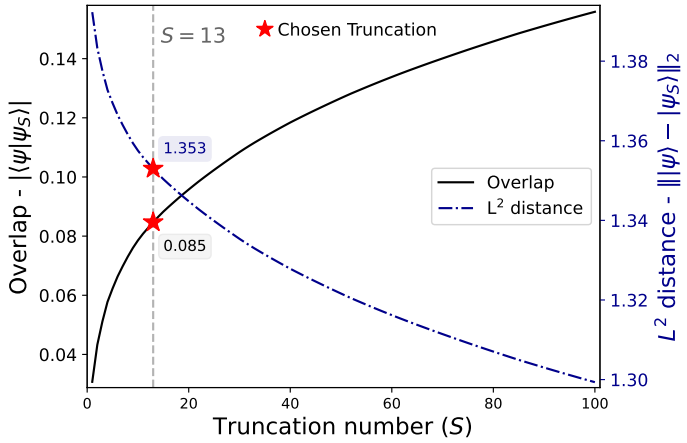

Figure 7.13: **Quality of the sparsification.** The left axis (black) shows the overlap of $|\Psi_S\rangle$ with the DMRG state, while the right axis (blue) depicts the L^2 distance. The red stars indicate the selected truncation number for the experiments. Figure taken from Ref. [267].

selection of this truncation. Moreover, we see that the squared overlap p_0 between the (sparsified) initial state and the converged DMRG state is $(2 \times 10^{-5})\ 7 \times 10^{-6}$, necessitating $M = 10^{13}$ (using Theorem 9) samples to accurately identify the precise ground state with 95% confidence. This is the main motivation behind ruptures, to target inflection points instead of the exact first step. Hence, it requires far less resources than theoretically expected for detecting discontinuities, while delivering comparable performance. The crucial aspect here is the paradigm shift from targeting the exact ground state energy, which is arduous and expensive, to identifying a significant accumulation, which is far cheaper and sufficient for most purposes.

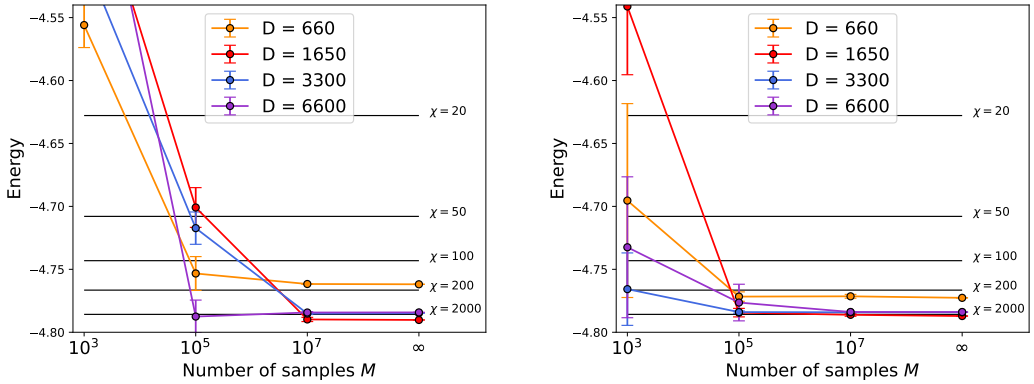

(a) 26 qubits, DMRG ($\chi = 10$)

(b) 26 qubits, sparsified DMRG ($\chi = 10$)

Figure 7.14: **Energy predictions** using the technique described in Section 7.4.2.3, employing the maxima of the gradient of the ACDF during a $\delta = 0.01$ ball. The horizontal lines represent the values predicted by DMRG with varying bond dimensions χ. The various colors indicate the maximum time evolution used to calculate the CDF, while the x-axis represents the number of samples. The error bars indicate one standard deviation calculated from 10 experiments. The infinity sign denotes the domain of infinite samples. Figure taken from Ref. [267]

Ultimately, we provide the energy predictions as a function of the sample quantity and maximum depth in Figure 7.14. We note that using just one-fifth of the depth used in

the previous figures, we can get equivalent energy estimates utilizing only 10^3 and 10^5 samples, respectively, for the DMRG and sparse initial states. This suggests that identifying the inflection point needs might require less precision than previously expected.

7.4.3 *Summary*

Based on our numerical simulations conducted on the non-trivial random fully-connected Heisenberg model, LT-type algorithms present themselves as a promising algorithm for GSEE on FTQC devices. Our main contributions are to instead focus on identifying the inflection point of the spectral CDF, which is an easier target than the exact ground state energy, while being efficient, statically robust and accurate. Furthermore, we advocate to use initial states obtained with one of the best classical methods available: DMRG. This paradigm shift allows to relax the high precision requirements on the initial states, which influences the scaling of depth and sample size. Hence, in this setting, the LT algorithm seems to require far less resources as previously anticipated, Although QPE is expected to surpass the LT algorithm in the FTQC era, our results indicate that LT algorithms can enhance classical solutions with only constrained quantum resources ($10^4 - 10^5$ samples), significantly fewer than theoretical predictions. In conclusion, the LT algorithm is a promising and effective quantum algorithm, connecting the NISQ and FTQC eras.

7.5 DISCUSSION

In this chapter, we have explored incoherent ways to perform Hamiltonian transformation using Fourier moments. Instead of performing the calculation coherently, as with QSP, we instead advocate to break the target function into its Fourier decomposition, and evaluate the phases, or Fourier moments, on the quantum computer. While this is asymptotically less optimal than QSP, it is better suited for near-term quantum devices. We explained how to compute the Fourier moments in a control-free way, and show case our approach on a real quantum computer. We then focused on ground state energy estimation using the Heaviside function, where eigenvalues can be computed by detecting steps in the spectral CDF.

Fourier moments are a versatile tool, and can be applied to a variety of cases. Therefore, it has a high potential for applications. Moreover, its compatibility with near-term devices makes it interesting in the close future, until full QEC is achieved. It is arguably one of the best usage of quantum computers before fault-tolerance.

8

CONCLUSIONS

Quantum information introduces a revolutionary computing paradigm with the potential to address certain problems far more efficiently than classical methods. However, significant challenges remain: first, decoherence rapidly degrades the quantum information stored in qubits, preventing the execution of deep circuits; and second, apart from factoring, there are no end-to-end useful applications that demonstrate a clear advantage over classical approaches. While the first issue can be solved using quantum error correction, and mitigated in the mean time using a variety of techniques. The second issue is however more severe, as it dictates wherever quantum computers are worth building.

While resolving these challenges is beyond the scope of a single thesis, this work aims to better understand the current limitations. After introducing the framework of quantum simulation and superconducting circuits, we focus on error mitigation. We classify error mitigation protocols into three categories: error suppression, device-based techniques, and state-based methods. Although error mitigation is not scalable and cannot replace quantum error correction in the long term, we argue that it is a crucial tool for conducting experiments on quantum devices today and will likely remain so for the years to come.

In the second part of this thesis, we explore applications involving randomized quantum algorithms. Quantum simulation involves three main components: an initial state, a time evolution operator, and an observable. We begin by addressing initial state preparation using the Variational Quantum Eigensolver, a heuristic algorithm that approximates ground states by minimizing the energy expectation value of a parameterized wavefunction. While VQE lacks performance guarantees, we successfully applied it to problems in nuclear shell theory and the Lipkin model, proposing a Hamiltonian-informed training approach for the unitary coupled cluster ansatz. Although VQE may not yield accurate solutions for more complex systems, it can serve as a valuable starting point for more advanced algorithms.

Next, we examine the compilation of time evolution operators. Trotter-Suzuki decompositions are well-established and highly effective for systems with substantial structure, but they can be computationally expensive for nuclear physics problems. To address this, we investigate the random product formula, known as qDrift, whose circuit depth depends not on the number of Hamiltonian terms but on its one-norm. We extend bias and concentration bounds for qDrift when sampling from arbitrary distributions and averaging over multiple circuit realizations. Furthermore, we propose an algorithm that reduces simulation costs compared to the standard approach.

We demonstrate two quantum simulation applications. In the first, we construct a stochastic channel, using the theoretical knowledge of the exact evolution, to compute the flavor variance in forward-scattering neutrinos. While we do not explicitly approximate the exact channel, we empirically show that permutation-invariant observables can be reproduced with a circuit depth independent of system size. In the second experiment, we perform a quench across a phase transition and observe the Kibble-Zurek mechanism.

Finally, we utilize the developed tools to propose an end-to-end application based on Fourier moments—expectation values of time evolution operators. These moments are versatile quantities enabling arbitrary Hamiltonian transformations. This framework has broad applications, including response functions and ground state energy estimation. For the latter, we address practical challenges, such as detecting small overlaps with the true

ground state, and introduce a signal processing-based heuristic that achieves a quadratic reduction in sample complexity. Combining this algorithm with initial states prepared using DMRG, we demonstrate the estimation of ground state energies for challenging systems from low-bond-dimension states.

While achieving a meaningful quantum advantage remains a distant goal, this work presents a promising framework based on Fourier moments. However, the feasibility of implementing the required circuit depths on near-term quantum devices remains uncertain, necessitating further advancements in quantum algorithms, quantum error correction and hardware to bridge this gap.

A critical focus moving forward is the identification of specific physical systems of interest, such as lattice gauge theories or chemistry, and the fine-tuning of algorithms tailored to these cases. For instance, one must consider the optimal implementation of time evolution operators suited to particular initial states and observables, alongside a precise evaluation of the computational costs. While QSP is theoretically optimal for quantum simulation, the practical pre-factors remain uncertain, particularly when compared to simpler PFs. Addressing these uncertainties requires adopting an engineering perspective, focusing on optimizing every aspect of the workflow. Incremental improvements across various components could collectively yield substantial performance gains. Although such an approach may lack the allure of achieving exponential advantage, it represents a vital and methodical step toward enhancing the practical utility of quantum algorithms. Examples include designing state-dependent qDrift protocols or leveraging composite channels more effectively.

Realizing quantum advantage will likely require a co-design approach that integrates algorithms with the hardware's intrinsic strengths and limitations. This includes exploiting hardware topology, developing system-specific error mitigation strategies, and optimizing compilation to maximize performance on the target device. Furthermore, it is important to understand how to optimally compile quantum algorithms within quantum error correction frameworks, which is rather different than for NISQ devices. While significant progress is required to address practical problems, ongoing advancements in quantum error correction and algorithm development provide grounds for optimism that important problems may be solved in the near future with the help of quantum computers.

BIBLIOGRAPHY

[1] R.P. Feynman. "Simulating physics with computers." In: *Int J Theor Phys* 21 (1982), pp. 467–488. DOI: https://doi.org/10.1007/BF02650179.

[2] J. E. Gubernatis, N. Kawashima, and P. Werner. *Quantum Monte Carlo Methods*. Cambridge, UK: Cambridge University Press, 2016.

[3] Steven R. White. "Density matrix formulation for quantum renormalization groups." In: *Phys. Rev. Lett.* 69 (19 Nov. 1992), pp. 2863–2866. DOI: 10.1103/PhysRevLett.69.2863. URL: https://link.aps.org/doi/10.1103/PhysRevLett.69.2863.

[4] P.W. Shor. "Algorithms for quantum computation: discrete logarithms and factoring." In: *Proceedings 35th Annual Symposium on Foundations of Computer Science*. 1994, pp. 124–134. DOI: 10.1109/SFCS.1994.365700.

[5] John M. Martyn, Zane M. Rossi, Andrew K. Tan, and Isaac L. Chuang. "Grand Unification of Quantum Algorithms." In: *PRX Quantum* 2 (4 Dec. 2021), p. 040203. DOI: 10.1103/PRXQuantum.2.040203. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.040203.

[6] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. "The theory of variational hybrid quantum-classical algorithms." In: *New Journal of Physics* 18.2 (2016), p. 023023. DOI: 10.1088/1367-2630/18/2/023023. URL: https://doi.org/10.1088/1367-2630/18/2/023023.

[7] Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. "Barren plateaus in quantum neural network training landscapes." In: *Nature communications* 9.1 (2018), p. 4812.

[8] Amira Abbas, Robbie King, Hsin-Yuan Huang, William J. Huggins, Ramis Movassagh, Dar Gilboa, and Jarrod McClean. "On quantum backpropagation, information reuse, and cheating measurement collapse." In: *Advances in Neural Information Processing Systems*. Ed. by A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine. Vol. 36. Curran Associates, Inc., 2023, pp. 44792–44819. URL: https://proceedings.neurips.cc/paper_files/paper/2023/file/8c3caae2f725c8e2a55ecd600563d172-Paper-Conference.pdf.

[9] I. M. Georgescu, S. Ashhab, and Franco Nori. "Quantum simulation." In: *Rev. Mod. Phys.* 86 (1 Mar. 2014), pp. 153–185. DOI: 10.1103/RevModPhys.86.153. URL: https://link.aps.org/doi/10.1103/RevModPhys.86.153.

[10] Iulia Buluta and Franco Nori. "Quantum Simulators." In: *Science* 326 (2009), pp. 108–111. DOI: 10.1126/science.1177838.

[11] David P DiVincenzo. "The physical implementation of quantum computation." In: *Fortschritte der Physik: Progress of Physics* 48.9-11 (2000), pp. 771–783.

[12] Peter W. Shor. "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer." In: *SIAM Journal on Computing* 26.5 (1997), pp. 1484–1509. DOI: 10.1137/S0097539795293172. eprint: https://doi.org/10.1137/S0097539795293172. URL: https://doi.org/10.1137/S0097539795293172.

[13] Craig Gidney and Martin Ekerå. "How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits." In: *Quantum* 5 (Apr. 2021), p. 433. ISSN: 2521-327X. DOI: 10.22331/q-2021-04-15-433. URL: https://doi.org/10.22331/q-2021-04-15-433.

[14] Y. Kim, A. Eddins, S. Anand, et al. "Evidence for the utility of quantum computing before fault tolerance." In: *Nature* 618 (2023), pp. 500–505. DOI: 10.1038/s41586-023-06096-3.

[15] Francescso Tacchino. "Digital quantum simulations and machine learning on near-term quantum processors." In: *Ph.D. thesis, Univesity of Pavia* (2019). URL: https://iris.unipv.it/bitstream/11571/1317093/2/Tacchino_PhDThesis.pdf.

[16] P. Jordan and E. Wigner. "Über das Paulische Äquivalenzverbot." In: *Z. Physik* 47 (1928), pp. 631–651. DOI: https://doi.org/10.1007/BF01331938.

[17] Jacob T. Seeley, Martin J. Richard, and Peter J. Love. "The Bravyi-Kitaev transformation for quantum computation of electronic structure." In: *J. Chem. Phys.* 137.22 (Dec. 2012), p. 224109. DOI: 10.1063/1.4768229.

[18] Francesco Tacchino, Alessandro Chiesa, Stefano Carretta, and Dario Gerace. "Quantum Computers as Universal Quantum Simulators: State-of-the-Art and Perspectives." In: *Advanced Quantum Technologies* 3.3 (2020), p. 1900052. DOI: https://doi.org/10.1002/qute.201900052.

[19] A. Miessen, P. J. Ollitrault, F. Tacchino, and Ivano Tavernelli. "Quantum algorithms for quantum dynamics." In: *Nature Computational Science* 3 (2023), pp. 25–37. DOI: 10.1038/s43588-022-00374-2. URL: https://doi.org/10.1038/s43588-022-00374-2.

[20] David C. McKay, Christopher J. Wood, Sarah Sheldon, Jerry M. Chow, and Jay M. Gambetta. "Efficient Z gates for quantum computing." In: *Phys. Rev. A* 96 (2 Aug. 2017), p. 022330. DOI: 10.1103/PhysRevA.96.022330. URL: https://link.aps.org/doi/10.1103/PhysRevA.96.022330.

[21] Jerry M. Chow et al. "Universal Quantum Gate Set Approaching Fault-Tolerant Thresholds with Superconducting Qubits." In: *Phys. Rev. Lett.* 109 (6 Aug. 2012), p. 060501. DOI: 10.1103/PhysRevLett.109.060501. URL: https://link.aps.org/doi/10.1103/PhysRevLett.109.060501.

[22] G. S. Paraoanu. "Microwave-induced coupling of superconducting qubits." In: *Phys. Rev. B* 74 (14 Oct. 2006), p. 140504. DOI: 10.1103/PhysRevB.74.140504. URL: https://link.aps.org/doi/10.1103/PhysRevB.74.140504.

[23] Chad Rigetti and Michel Devoret. "Fully microwave-tunable universal gates in superconducting qubits with linear couplings and fixed transition frequencies." In: *Phys. Rev. B* 81 (13 Apr. 2010), p. 134507. DOI: 10.1103/PhysRevB.81.134507. URL: https://link.aps.org/doi/10.1103/PhysRevB.81.134507.

[24] Jerry M. Chow et al. "Simple All-Microwave Entangling Gate for Fixed-Frequency Superconducting Qubits." In: *Phys. Rev. Lett.* 107 (8 Aug. 2011), p. 080502. DOI: 10.1103/PhysRevLett.107.080502. URL: https://link.aps.org/doi/10.1103/PhysRevLett.107.080502.

[25] Sarah Sheldon, Easwar Magesan, Jerry M. Chow, and Jay M. Gambetta. "Procedure for systematically tuning up cross-talk in the cross-resonance gate." In: *Phys. Rev. A* 93 (6 June 2016), p. 060302. DOI: 10.1103/PhysRevA.93.060302. URL: https://link.aps.org/doi/10.1103/PhysRevA.93.060302.

[26] Y. Salathé et al. "Digital Quantum Simulation of Spin Models with Circuit Quantum Electrodynamics." In: *Phys. Rev. X* 5 (2 June 2015), p. 021027. DOI: 10.1103/PhysRevX.5.021027. URL: https://link.aps.org/doi/10.1103/PhysRevX.5.021027.

[27] David C. McKay, Stefan Filipp, Antonio Mezzacapo, Easwar Magesan, Jerry M. Chow, and Jay M. Gambetta. "Universal Gate for Fixed-Frequency Qubits via a Tunable Bus." In: *Phys. Rev. Appl.* 6 (6 Dec. 2016), p. 064007. DOI: 10.1103/PhysRevApplied.6.064007. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.6.064007.

[28] Farrokh Vatan and Colin Williams. "Optimal quantum circuits for general two-qubit gates." In: *Phys. Rev. A* 69 (3 Mar. 2004), p. 032315. DOI: 10.1103/PhysRevA.69.032315. URL: https://link.aps.org/doi/10.1103/PhysRevA.69.032315.

[29] R. Somma, G. Ortiz, J. E. Gubernatis, E. Knill, and R. Laflamme. "Simulating physical phenomena by quantum networks." In: *Phys. Rev. A* 65 (4 Apr. 2002), p. 042323. DOI: 10.1103/PhysRevA.65.042323. URL: https://link.aps.org/doi/10.1103/PhysRevA.65.042323.

[30] Gilles Brassard, Peter Hoyer, Michele Mosca, and Alain Tapp. "Quantum Amplitude Amplification and Estimation." In: *Quantum Computation and Quantum Information* 305 (2002), pp. 53–74. DOI: https://doi.org/10.1090/conm/305/05215.

[31] A Yu Kitaev. "Quantum measurements and the Abelian stabilizer problem." In: *arXiv e-prints* (1995). arXiv: quant-ph/9511026 [quant-ph].

[32] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.

[33] Guang Hao Low and Isaac L. Chuang. "Hamiltonian Simulation by Qubitization." In: *Quantum* 3 (July 2019), p. 163. ISSN: 2521-327X. DOI: 10.22331/q-2019-07-12-163. URL: http://dx.doi.org/10.22331/q-2019-07-12-163.

[34] Dominic W. Berry et al. *Rapid initial state preparation for the quantum simulation of strongly correlated molecules*. 2024. arXiv: 2409.11748 [quant-ph]. URL: https://arxiv.org/abs/2409.11748.

[35] J. Kaiser and R. Schafer. "On the use of the Io-sinh window for spectrum analysis." In: *IEEE Transactions on Acoustics, Speech, and Signal Processing* 28.1 (1980), pp. 105–107. DOI: 10.1109/TASSP.1980.1163349.

[36] Colin D. Bruzewicz, John Chiaverini, Robert McConnell, and Jeremy M. Sage. "Trapped-ion quantum computing: Progress and challenges." In: *Applied Physics Reviews* 6.2 (2019), p. 021314. DOI: 10.1063/1.5088164. URL: https://doi.org/10.1063/1.5088164.

[37] K. Wright et al. "Benchmarking an 11-qubit quantum computer." In: *Nature Communications* 10 (2019), p. 5464. DOI: 10.1038/s41467-019-13534-2. URL: https://doi.org/10.1038/s41467-019-13534-2.

[38] J. I. Cirac and P. Zoller. "Quantum computations with cold trapped ions." In: *Phys. Rev. Lett.* 74 (1995), pp. 4091–4094. DOI: 10.1103/PhysRevLett.74.4091. URL: https://doi.org/10.1103/PhysRevLett.74.4091.

[39] R. Blatt and C. F. Roos. "Quantum simulations with trapped ions." In: *Nat. Phys.* 8 (2012), pp. 277–284. DOI: 10.1038/nphys2252. URL: https://doi.org/10.1038/nphys2252.

[40] R. Blatt and D. Wineland. "Entangled states of trapped atomic ions." In: *Nature* 453 (2008), pp. 1008–1015. DOI: 10.1038/nature07125. URL: https://doi.org/10.1038/nature07125.

[41] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. "Quantum dynamics of single trapped ions." In: *Rev. Mod. Phys.* 75 (2003), pp. 281–324. DOI: 10.1103/RevModPhys.75.281. URL: https://doi.org/10.1103/RevModPhys.75.281.

[42] C. Gross and I. Bloch. "Quantum simulations with ultracold atoms in optical lattices." In: *Science* 357 (2017), pp. 995–1001. DOI: 10.1126/science.aaf6846. URL: https://doi.org/10.1126/science.aaf6846.

[43] I. Bloch, J. Dalibard, and W. Zwerger. "Many-body physics with ultracold gases." In: *Rev. Mod. Phys.* 80 (2008), pp. 885–964. DOI: 10.1103/RevModPhys.80.885. URL: https://doi.org/10.1103/RevModPhys.80.885.

[44] M. Lewenstein, A. Sanpera, V. Ahufinger, B. Damski, A. Sen, and U. Sen. "Ultracold atomic gases in optical lattices: mimicking condensed matter physics and beyond." In: *Adv. Phys.* 56 (2007), pp. 243–379. DOI: 10.1080/00018730701223200. URL: https://doi.org/10.1080/00018730701223200.

[45] R. Vrijen, E. Yablonovitch, K. Wang, H. W. Jiang, A. Balandin, V. Roychowdhury, T. Mor, and D. DiVincenzo. "Electron-spin-resonance transistors for quantum computing in silicon-germanium heterostructures." In: *Phys. Rev. A* 62 (2000), p. 012306. DOI: 10.1103/PhysRevA.62.012306. URL: https://doi.org/10.1103/PhysRevA.62.012306.

[46] A. Morello, J. J. Pla, F. A. Zwanenburg, K. W. Chan, K. Y. Tan, H. Huebl, M. Möttönen, C. D. Nugroho, C. Yang, J. A. van Donkelaar, et al. "Single-shot readout of an electron spin in silicon." In: *Nature* 467 (2010), pp. 687–691. DOI: 10.1038/nature09406. URL: https://doi.org/10.1038/nature09406.

[47] B. E. Kane. "A silicon-based nuclear spin quantum computer." In: *Nature* 393 (1998), pp. 133–137. DOI: 10.1038/30263. URL: https://doi.org/10.1038/30263.

[48] E. Knill, R. Laflamme, and G. J. Milburn. "A scheme for efficient quantum computation with linear optics." In: *Nature* 409 (2001), pp. 46–52. DOI: 10.1038/35051009. URL: https://doi.org/10.1038/35051009.

[49] J. D. Franson, M. Donegan, M. Fitch, B. Jacobs, and T. Pittman. "High-fidelity quantum logic operations using linear optical elements." In: *Physical Review Letters* 89.13 (2002), p. 137901. DOI: 10.1103/PhysRevLett.89.137901. URL: https://doi.org/10.1103/PhysRevLett.89.137901.

[50] Lars S. Madsen, Fabian Laudenbach, Mohammad F. Askarani, and et al. "Quantum computational advantage with a programmable photonic processor." In: *Nature* 606 (2022), pp. 75–81. DOI: 10.1038/s41586-022-04725-x. URL: https://doi.org/10.1038/s41586-022-04725-x.

[51] Juan Miguel Arrazola, Ville Bergholm, Kamil Brádler, and et al. "Quantum circuits with many photons on a programmable nanophotonic chip." In: *Nature* 591 (2021), pp. 54–60. DOI: 10.1038/s41586-021-03202-1. URL: https://doi.org/10.1038/s41586-021-03202-1.

[52] M. Mariantoni, H. Wang, T. Yamamoto, M. J. Schwarz, I. Chiorescu, Y. Nakamura, J. S. Tsai, S. Lloyd, and J. Martinis. "Implementing the Quantum von Neumann Architecture with Superconducting Circuits." In: *Science* 334 (2011), pp. 61–65. DOI: 10.1126/science.1208371. URL: https://doi.org/10.1126/science.1208371.

[53] J. Koch, T. M. Yu, J. M. Gambetta, A. A. Houck, K. C. Fong, S. Choi, J. M. Martinis, M. H. Devoret, and J. M. Clarke. "Charge-insensitive qubit design derived from the Cooper pair box." In: *Physical Review A* 76 (2007), p. 042319. DOI: 10.1103/PhysRevA.76.042319. URL: https://doi.org/10.1103/PhysRevA.76.042319.

[54] J. Gambetta, W. A. Braff, A. Wallraff, K. M. S. Jones, J. M. Gambetta, J. M. Martinis, M. H. Devoret, and J. M. Clarke. "Protocols for optimal readout of qubits using a continuous quantum nondemolition measurement." In: *Physical Review A* 76 (2007), p. 012325. DOI: 10.1103/PhysRevA.76.012325. URL: https://doi.org/10.1103/PhysRevA.76.012325.

[55] J. Gambetta, A. Blais, M. Boissonneault, I. Martin, and R. J. Schoelkopf. "Quantum trajectory approach to circuit QED: Quantum jumps and the Zeno effect." In: *Physical Review A* 77 (2008), p. 012112. DOI: 10.1103/PhysRevA.77.012112. URL: https://doi.org/10.1103/PhysRevA.77.012112.

[56] J. Majer et al. "Coupling superconducting qubits via a cavity bus." In: *Nature* 449 (2007), pp. 443–447. DOI: 10.1038/nature06184. URL: https://doi.org/10.1038/nature06184.

[57] C. Rigetti, J. M. Gambetta, S. Poletto, J. M. Chow, J. M. Clarke, and J. M. Martinis. "Superconducting qubit in a waveguide cavity with a coherence time approaching 0.1 ms." In: *Physical Review B* 86 (2012), 100506(R). DOI: 10.1103/PhysRevB.86.100506. URL: https://doi.org/10.1103/PhysRevB.86.100506.

[58] A. Nersisyan, S. Poletto, N. Alidoust, C. Rigetti, and M. Mariantoni. "Manufacturing low dissipation superconducting quantum processors." In: *arXiv:1901.08042* (2019). URL: https://arxiv.org/abs/1901.08042.

[59] J. Kelly et al. "State preservation by repetitive error detection in a superconducting quantum circuit." In: *Nature* 519 (2015), pp. 66–69. DOI: 10.1038/nature14270. URL: https://doi.org/10.1038/nature14270.

[60] R. Barends et al. "Digitized adiabatic quantum computing with a superconducting circuit." In: *Nature* 534 (2016), pp. 222–225. DOI: 10.1038/nature17656. URL: https://doi.org/10.1038/nature17656.

[61] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver. "A Quantum Engineer's Guide to Superconducting Qubits." In: *arXiv preprint arXiv:1904.06560* (2019). URL: https://arxiv.org/abs/1904.06560.

[62] L. Lamata, A. Parra-Rodriguez, M. Sanz, and E. Solano. "Digital-analog quantum simulations with superconducting circuits." In: *Advances in Physics: X* 3 (2018), p. 1457981. DOI: 10.1080/23746149.2018.1457981. URL: https://doi.org/10.1080/23746149.2018.1457981.

[63] G. Wendin. "Quantum information processing with superconducting circuits: a review." In: *Reports on Progress in Physics* 80 (2017), p. 106001. DOI: 10.1088/1361-6633/aa7e1a. URL: https://doi.org/10.1088/1361-6633/aa7e1a.

[64] Riccardo Manenti and Mario Motta. *Quantum Information Science*. Oxford University Press, Aug. 2023. ISBN: 978-0-19-878748-8.

[65] A. Córcoles, E. Magesan, S. J. Srinivasan, A. W. Cross, M. Steffen, J. M. Gambetta, and J. M. Chow. "Demonstration of a quantum error detection code using a square lattice of four superconducting qubits." In: *Nature Communications* 6 (2015), p. 6979. DOI: 10.1038/ncomms7979. URL: https://doi.org/10.1038/ncomms7979.

[66] C. Rigetti, J. M. Gambetta, S. Poletto, B. L. Plourde, J. M. Chow, A. D. Córcoles, J. A. Smolin, S. T. Merkel, J. R. Rozen, G. A. Keefe, et al. "Superconducting qubit in a waveguide cavity with a coherence time approaching 0.1 ms." In: *Physical Review B* 86.10 (2012), p. 100506. DOI: 10.1103/PhysRevB.86.100506.

[67] C. Wang, Y. Y. Gao, I. M. Pop, U. Vool, C. Axline, T. Brecht, R. W. Heeres, L. Frunzio, M. H. Devoret, G. Catelani, et al. "Measurement and control of quasiparticle dynamics in a superconducting qubit." In: *Nature Communications* 5 (2014), pp. 1–7. DOI: 10.1038/ncomms6836.

[68] M. Hutchings, J. B. Hertzberg, Y. Liu, N. T. Bronn, G. A. Keefe, M. Brink, J. M. Chow, and B. Plourde. "Tunable superconducting qubits with flux-independent coherence." In: *Phys. Rev. Applied* 8 (2017), p. 044003. DOI: 10.1103/PhysRevApplied.8.044003.

[69] M. Takita, A. W. Cross, A. Córcoles, J. M. Chow, and J. M. Gambetta. "Experimental demonstration of fault-tolerant state preparation with superconducting qubits." In: *Phys. Rev. Lett.* 119 (2017), p. 180501. DOI: 10.1103/PhysRevLett.119.180501.

[70] D. Rosenberg, D. Kim, R. Das, D. Yost, S. Gustavsson, D. Hover, P. Krantz, A. Melville, L. Racz, G. O. Samach, et al. "3D integrated superconducting qubits." In: *npj Quantum Information* 3 (2017), pp. 1–5. DOI: 10.1038/s41534-017-0044-0.

[71] M. Ware, B. R. Johnson, J. M. Gambetta, T. A. Ohki, J. M. Chow, and B. L. T. Plourde. "Cross-resonance interactions between superconducting qubits with variable detuning." In: (2019). arXiv: arXiv:1905.11480 [quant-ph]. URL: https://arxiv.org/abs/1905.11480.

[72] R. P. Feynman, R. B. Leighton, and M. L. Sands. *The Feynman Lectures on Physics, Volume 3*. Addison-Wesley, 1963.

[73] M. Büttiker. "Zero-current persistent potential drop across small-capacitance Josephson junctions." In: *Physical Review B* 36 (1987), p. 3548. DOI: 10.1103/PhysRevB.36.3548.

[74] Jens Koch, Terri M. Yu, Jay Gambetta, A. A. Houck, D. I. Schuster, J. Majer, Alexandre Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. "Charge-insensitive qubit design derived from the Cooper pair box." In: *Phys. Rev. A* 76 (4 Oct. 2007), p. 042319. DOI: 10.1103/PhysRevA.76.042319. URL: https://link.aps.org/doi/10.1103/PhysRevA.76.042319.

[75] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R. S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf. "Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics." In: *Nature* 431 (2004), pp. 162–167. DOI: 10.1038/nature02851. URL: https://doi.org/10.1038/nature02851.

[76] J.L. Park. "The concept of transition in quantum mechanics." In: *Foundations of Physics* 1 (1970), pp. 23–33. DOI: 10.1007/BF00708652. URL: https://doi.org/10.1007/BF00708652.

[77] S. B. Bravyi and A. Y. Kitaev. "Quantum codes on a lattice with boundary." In: (1998). arXiv: quant-ph/9811052 [quant-ph]. URL: https://arxiv.org/abs/quant-ph/9811052.

[78] E. Dennis, A. Y. Kitaev, A. Landahl, and J. Preskill. "Topological quantum memory." In: *Journal of Mathematical Physics* 43 (2002), pp. 4452–4505. DOI: 10.1063/1.1499754. arXiv: quant-ph/0110143 [quant-ph]. URL: http://arxiv.org/abs/quant-ph/0110143.

[79] A. Y. Kitaev. "Quantum error correction with imperfect gates." In: *Quantum Communication, Computing, and Measurement*. Ed. by Osamu Hirota, Alexander S. Holevo, and Carlton M. Caves. Plenum Publishing Corporation, 1997. Chap. Quantum error correction with imperfect gates, pp. 181–188. DOI: 10.1007/978-1-4615-5923-8_22.

[80] "Suppressing quantum errors by scaling a surface code logical qubit." In: *Nature* 614.7949 (2023), pp. 676–681.

[81] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. "Surface codes: Towards practical large-scale quantum computation." In: *Phys. Rev. A* 86 (3 Sept. 2012), p. 032324. DOI: 10.1103/PhysRevA.86.032324. URL: https://link.aps.org/doi/10.1103/PhysRevA.86.032324.

[82] Dolev Bluvstein, Simon J Evered, Alexandra A Geim, Sophie H Li, Hengyun Zhou, Tom Manovitz, Sepehr Ebadi, Madelyn Cain, Marcin Kalinowski, Dominik Hangleiter, et al. "Logical quantum processor based on reconfigurable atom arrays." In: *Nature* 626.7997 (2024), pp. 58–65.

[83] Dolev Bluvstein et al. "Logical quantum processor based on reconfigurable atom arrays." In: *Nature* 626 (2024), pp. 58–65. DOI: 10.1038/s41586-023-06927-3. URL: https://doi.org/10.1038/s41586-023-06927-3.

[84] S. J. Evered et al. "High-fidelity parallel entangling gates on a neutral-atom quantum computer." In: *Nature* 622 (2023), pp. 268–272. DOI: 10.1038/s41586-023-06481-y. URL: https://doi.org/10.1038/s41586-023-06481-y.

[85] Ciaran Ryan-Anderson, Justin G Bohnet, Kenny Lee, Daniel Gresh, Aaron Hankin, JP Gaebler, David Francois, Alexander Chernoguzov, Dominic Lucchetti, Natalie C Brown, et al. "Realization of real-time fault-tolerant quantum error correction." In: *Physical Review X* 11.4 (2021), p. 041058.

[86] Earl Campbell. "A series of fast-paced advances in Quantum Error Correction." In: *Nature Reviews Physics* 6 (2024), pp. 160–161. DOI: 10.1038/s42254-024-00706-3. URL: https://doi.org/10.1038/s42254-024-00706-3.

[87] Ryuji Takagi, Suguru Endo, Shintaro Minagawa, and Mile Gu. "Fundamental limits of quantum error mitigation." In: *npj Quantum Information* 8 (2022), p. 114. DOI: https://doi.org/10.1038/s41534-022-00618-z.

[88] Lorenza Viola and Seth Lloyd. "Dynamical suppression of decoherence in two-state quantum systems." In: *Phys. Rev. A* 58 (4 Oct. 1998), pp. 2733–2744. DOI: 10.1103/PhysRevA.58.2733. URL: https://link.aps.org/doi/10.1103/PhysRevA.58.2733.

[89] K. Khodjasteh and D. A. Lidar. "Fault-Tolerant Quantum Dynamical Decoupling." In: *Phys. Rev. Lett.* 95 (18 Oct. 2005), p. 180501. DOI: 10.1103/PhysRevLett.95.180501. URL: https://link.aps.org/doi/10.1103/PhysRevLett.95.180501.

[90] Götz S. Uhrig. "Keeping a Quantum Bit Alive by Optimized π-Pulse Sequences." In: *Phys. Rev. Lett.* 98 (10 Mar. 2007), p. 100504. DOI: 10.1103/PhysRevLett.98.100504. URL: https://link.aps.org/doi/10.1103/PhysRevLett.98.100504.

[91] Mustafa Ahmed Ali Ahmed, Gonzalo A. Álvarez, and Dieter Suter. "Robustness of dynamical decoupling sequences." In: *Phys. Rev. A* 87 (4 Apr. 2013), p. 042309. DOI: 10.1103/PhysRevA.87.042309. URL: https://link.aps.org/doi/10.1103/PhysRevA.87.042309.

[92] Gregory Quiroz and Daniel A. Lidar. "Optimized dynamical decoupling via genetic algorithms." In: *Phys. Rev. A* 88 (5 Nov. 2013), p. 052306. DOI: 10.1103/PhysRevA.88.052306. URL: https://link.aps.org/doi/10.1103/PhysRevA.88.052306.

[93] Akel Hashim et al. "Randomized Compiling for Scalable Quantum Computing on a Noisy Superconducting Quantum Processor." In: *Phys. Rev. X* 11 (4 Nov. 2021), p. 041039. DOI: 10.1103/PhysRevX.11.041039. URL: https://link.aps.org/doi/10.1103/PhysRevX.11.041039.

[94] Z. Cai and S. Benjamin. "Constructing Smaller Pauli Twirling Sets for Arbitrary Error Channels." In: *Sci Rep* 9 (2019), p. 1128. DOI: https://doi.org/10.1038/s41598-019-46722-7.

[95] F. G. Fuchs, V. Falch, and C. Johnsen. "Quantum Poker—a game for quantum computers suitable for benchmarking error mitigation techniques on NISQ devices." In: *Eur. Phys. J. Plus* 135.353 (2020). DOI: 10.1140/epjp/s13360-020-00360-5.

[96] Ewout van den Berg, Zlatko K. Minev, and Kristan Temme. "Model-free readout-error mitigation for quantum expectation values." In: *Phys. Rev. A* 105 (3 Mar. 2022), p. 032620. DOI: 10.1103/PhysRevA.105.032620. URL: https://link.aps.org/doi/10.1103/PhysRevA.105.032620.

[97] Nathan Earnest, Caroline Tornow, and Daniel J. Egger. "Pulse-efficient circuit transpilation for quantum applications on cross-resonance-based hardware." In: *Phys. Rev. Res.* 3 (4 Oct. 2021), p. 043088. DOI: 10.1103/PhysRevResearch.3.043088. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.3.043088.

[98] André Melo, Nathan Earnest-Noble, and Francesco Tacchino. "Pulse-efficient quantum machine learning." In: *Quantum* 7 (Oct. 2023), p. 1130. ISSN: 2521-327X. DOI: 10.22331/q-2023-10-09-1130. URL: https://doi.org/10.22331/q-2023-10-09-1130.

[99] Paul D. Nation, Hwajung Kang, Neereja Sundaresan, and Jay M. Gambetta. "Scalable Mitigation of Measurement Errors on Quantum Computers." In: *PRX Quantum* 2 (4 Nov. 2021), p. 040326. DOI: 10.1103/PRXQuantum.2.040326. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.040326.

[100] Ritajit Majumdar, Pedro Rivero, Friedrike Metz, Areeq Hasan, and Derek S Wang. "Best practices for quantum error mitigation with digital zero-noise extrapolation." In: *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*. Vol. 1. IEEE. 2023, pp. 881–887.

[101] Andre He, Benjamin Nachman, Wibe A. de Jong, and Christian W. Bauer. "Zero-noise extrapolation for quantum-gate error mitigation with identity insertions." In: *Phys. Rev. A* 102 (1 July 2020), p. 012426. DOI: 10.1103/PhysRevA.102.012426. URL: https://link.aps.org/doi/10.1103/PhysRevA.102.012426.

[102] Abhinav Kandala, Kristan Temme, Antonio D. Córcoles, Antonio Mezzacapo, Jerry M. Chow, and Jay M. Gambetta. "Error mitigation extends the computational reach of a noisy quantum processor." In: *Nature* 567 (2019), pp. 491–495. DOI: https://doi.org/10.1038/s41586-019-1040-7.

[103] Vincent R. Pascuzzi, Andre He, Christian W. Bauer, Wibe A. de Jong, and Benjamin Nachman. "Computationally efficient zero-noise extrapolation for quantum-gate-error mitigation." In: *Phys. Rev. A* 105 (4 Apr. 2022), p. 042406. DOI: 10.1103/PhysRevA.105.042406. URL: https://link.aps.org/doi/10.1103/PhysRevA.105.042406.

[104] Michele Grossi, Oriel Kiss, Francesco De Luca, Carlo Zollo, Ian Gremese, and Antonio Mandarino. "Finite-size criticality in fully connected spin models on superconducting quantum hardware." In: *Physical Review E* 107.2 (Feb. 2023). Publisher: American Physical Society, p. 024113. DOI: 10.1103/PhysRevE.107.024113. URL: https://link.aps.org/doi/10.1103/PhysRevE.107.024113 (visited on 04/07/2023).

[105] Tudor Giurgica-Tiron, Yousef Hindy, Ryan LaRose, Andrea Mari, and William J Zeng. "Digital zero noise extrapolation for quantum error mitigation." In: *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE. 2020, pp. 306–316.

[106] Benjamin Hall, Alessandro Roggero, Alessandro Baroni, and Joseph Carlson. "Simulation of collective neutrino oscillations on a quantum computer." In: *Phys. Rev. D* 104 (6 Sept. 2021), p. 063009. DOI: 10.1103/PhysRevD.104.063009. URL: https://link.aps.org/doi/10.1103/PhysRevD.104.063009.

[107] Sarmed A Rahman, Randy Lewis, Emanuele Mendicelli, and Sarah Powell. "Self-mitigating Trotter circuits for SU(2) lattice gauge theory on a quantum computer." In: *Phys. Rev. D* 106 (7 Oct. 2022), p. 074502. DOI: 10.1103/PhysRevD.106.074502. URL: https://link.aps.org/doi/10.1103/PhysRevD.106.074502.

[108] Miroslav Urbanek, Benjamin Nachman, Vincent R. Pascuzzi, Andre He, Christian W. Bauer, and Wibe A. de Jong. "Mitigating Depolarizing Noise on Quantum Computers with Noise-Estimation Circuits." In: *Phys. Rev. Lett.* 127 (27 Dec. 2021), p. 270502. DOI: 10.1103/PhysRevLett.127.270502. URL: https://link.aps.org/doi/10.1103/PhysRevLett.127.270502.

[109] Roland C. Farrell, Marc Illa, Anthony N. Ciavarella, and Martin J. Savage. "Scalable Circuits for Preparing Ground States on Digital Quantum Computers: The Schwinger Model Vacuum on 100 Qubits." In: *PRX Quantum* 5 (2 Apr. 2024), p. 020315. DOI: 10.1103/PRXQuantum.5.020315. URL: https://link.aps.org/doi/10.1103/PRXQuantum.5.020315.

[110] Oriel Kiss, Michele Grossi, and Alessandro Roggero. "Quantum error mitigation for Fourier moment computation." In: *Phys. Rev. D* 111 (3 Feb. 2025), p. 034504. DOI: 10.1103/PhysRevD.111.034504. URL: https://link.aps.org/doi/10.1103/PhysRevD.111.034504.

[111] Thomas E. O'Brien, Stefano Polla, Nicholas C. Rubin, William J. Huggins, Sam McArdle, Sergio Boixo, Jarrod R. McClean, and Ryan Babbush. "Error Mitigation via Verified Phase Estimation." In: *PRX Quantum* 2 (2 May 2021), p. 020317. DOI: 10.1103/PRXQuantum.2.020317. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.020317.

[112] Mingxia Huo and Ying Li. "Dual-state purification for practical quantum error mitigation." In: *Phys. Rev. A* 105 (2 Feb. 2022), p. 022427. DOI: 10.1103/PhysRevA.105.022427. URL: https://link.aps.org/doi/10.1103/PhysRevA.105.022427.

[113] "Purification-based quantum error mitigation of pair-correlated electron simulations." In: *Nat. Phys.* 19 (2023), pp. 1787–1792. DOI: doi.org/10.1038/s41567-023-02240-y.

[114] John A. Smolin, Jay M. Gambetta, and Graeme Smith. "Efficient Method for Computing the Maximum-Likelihood Quantum State from Measurements with Additive Gaussian Noise." In: *Phys. Rev. Lett.* 108 (7 Feb. 2012), p. 070502. DOI: 10.1103/PhysRevLett.108.070502. URL: https://link.aps.org/doi/10.1103/PhysRevLett.108.070502.

[115] X. Bonet-Monroig, R. Sagastizabal, M. Singh, and T. E. O'Brien. "Low-cost error mitigation by symmetry verification." In: *Phys. Rev. A* 98 (6 Dec. 2018), p. 062339. DOI: 10.1103/PhysRevA.98.062339. URL: https://link.aps.org/doi/10.1103/PhysRevA.98.062339.

[116] R. Sagastizabal et al. "Experimental error mitigation via symmetry verification in a variational quantum eigensolver." In: *Phys. Rev. A* 100 (1 July 2019), p. 010302. DOI: 10.1103/PhysRevA.100.010302. URL: https://link.aps.org/doi/10.1103/PhysRevA.100.010302.

[117] Minh C. Tran, Yuan Su, Daniel Carney, and Jacob M. Taylor. "Faster Digital Quantum Simulation by Symmetry Protection." In: *PRX Quantum* 2 (1 Feb. 2021), p. 010323. DOI: 10.1103/PRXQuantum.2.010323. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.010323.

[118] Sumner N. Hearth, Michael O. Flynn, Anushya Chandran, and Chris R. Laumann. "Efficient Local Classical Shadow Tomography with Number Conservation." In: *Phys. Rev. Lett.* 133 (6 Aug. 2024), p. 060802. DOI: 10.1103/PhysRevLett.133.060802. URL: https://link.aps.org/doi/10.1103/PhysRevLett.133.060802.

[119] H. Y. Huang, R. Kueng, and J. Preskill. "Predicting many properties of a quantum system from very few measurements." In: *Nat. Phys.* 16 (2020), pp. 1050–1057. DOI: 10.1038/s41567-020-0932-7.

[120] William J. Huggins, Sam McArdle, Thomas E. O'Brien, Joonho Lee, Nicholas C. Rubin, Sergio Boixo, K. Birgitta Whaley, Ryan Babbush, and Jarrod R. McClean. "Virtual Distillation for Quantum Error Mitigation." In: *Phys. Rev. X* 11 (4 Nov. 2021), p. 041036. DOI: 10.1103/PhysRevX.11.041036. URL: https://link.aps.org/doi/10.1103/PhysRevX.11.041036.

[121] Bálint Koczor. "Exponential Error Suppression for Near-Term Quantum Devices." In: *Phys. Rev. X* 11 (3 Sept. 2021), p. 031057. DOI: 10.1103/PhysRevX.11.031057. URL: https://link.aps.org/doi/10.1103/PhysRevX.11.031057.

[122] Haoran Liao, Derek S Wang, Iskandar Sitdikov, Ciro Salcedo, Alireza Seif, and Zlatko K Minev. "Machine learning for practical quantum error mitigation." In: *arXiv preprint arXiv:2309.17368* (2023).

[123] Sergei Filippov, Matea Leahy, Matteo AC Rossi, and Guillermo García-Pérez. "Scalable tensor-network error mitigation for near-term quantum computing." In: *arXiv preprint arXiv:2307.11740* (2023).

[124] Youngseok Kim, Christopher J. Wood, Theodore J. Yoder, Seth T. Merkel, Jay M. Gambetta, Kristan Temme, and Abhinav Kandala. "Scalable error mitigation for noisy quantum circuits produces competitive expectation values." In: *Nat. Phys.* (2023). DOI: https://doi.org/10.1038/s41567-022-01914-3.

[125] Y. Quek, D. Stilck França, S. Khatri, et al. "Exponentially tighter bounds on limitations of quantum error mitigation." In: *Nature Physics* (2024). DOI: 10.1038/s41567-024-02536-7. URL: https://doi.org/10.1038/s41567-024-02536-7.

[126] Daniel Nagaj. "Local Hamiltonians in quantum computation." In: *arXiv preprint* (2008). arXiv: 0808.2117 [quant-ph].

[127] Giuseppe Carleo and Matthias Troyer. "Solving the quantum many-body problem with artificial neural networks." In: *Science* 355.6325 (2017), pp. 602–606.

[128] Rodney J. Bartlett and Monika Musiał. "Coupled-cluster theory in quantum chemistry." In: *Rev. Mod. Phys.* 79 (1 Feb. 2007), pp. 291–352. DOI: 10.1103/RevModPhys.79.291. URL: https://link.aps.org/doi/10.1103/RevModPhys.79.291.

[129] W. Kohn, A. D. Becke, and R. G. Parr. "Density Functional Theory of Electronic Structure." In: *The Journal of Physical Chemistry* 100.31 (Aug. 1996), pp. 12974–12980. DOI: 10.1021/jp960669l.

[130] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O'brien. "A variational eigenvalue solver on a photonic quantum processor." In: *Nature communications* 5.1 (2014), p. 4213. DOI: 10.1038/ncomms5213. URL: http://dx.doi.org/10.1038/ncomms5213.

[131] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. "Evaluating analytic gradients on quantum hardware." In: *Phys. Rev. A* 99 (3 Mar. 2019), p. 032331. DOI: 10.1103/PhysRevA.99.032331. URL: https://link.aps.org/doi/10.1103/PhysRevA.99.032331.

[132] J.C. Spall. "Implementation of the simultaneous perturbation algorithm for stochastic optimization." In: *IEEE Transactions on Aerospace and Electronic Systems* 34.3 (1998), pp. 817–823. DOI: 10.1109/7.705889.

[133] M.J.D. Powell. *Advances in Optimization and Numerical Analysis*. Springer, 1994.

[134] Harper R. Grimsley, Sophia E. Economou, Edwin Barnes, and Nicholas J. Mayhall. "An adaptive variational algorithm for exact molecular simulations on a quantum computer." In: *Nat Commun* 10 (2019), p. 3007. DOI: https://doi.org/10.1038/s41467-019-10988-2.

[135] Oriel Kiss, Michele Grossi, Pavel Lougovski, Federico Sanchez, Sofia Vallecorsa, and Thomas Papenbrock. "Quantum computing of the ^{6}Li nucleus via ordered unitary coupled clusters." In: *Physical Review C* 106.3 (Sept. 2022). Publisher: American Physical Society, p. 034325. DOI: 10.1103/PhysRevC.106.034325. URL: https://link.aps.org/doi/10.1103/PhysRevC.106.034325 (visited on 04/07/2023).

[136] Alexey Uvarov, Jacob D. Biamonte, and Dmitry Yudin. "Variational quantum eigensolver for frustrated quantum systems." In: *Phys. Rev. B* 102 (7 Aug. 2020), p. 075104. DOI: 10.1103/PhysRevB.102.075104. URL: https://link.aps.org/doi/10.1103/PhysRevB.102.075104.

[137] Panagiotis Kl. Barkoutsos et al. "Quantum algorithms for electronic structure calculations: Particle-hole Hamiltonian and optimized wave-function expansions." In: *Phys. Rev. A* 98 (2 Aug. 2018), p. 022322. DOI: 10.1103/PhysRevA.98.022322. URL: https://link.aps.org/doi/10.1103/PhysRevA.98.022322.

[138] A. M. Romero, J. Engel, Ho Lun Tang, and Sophia E. Economou. "Solving nuclear structure problems with the adaptive variational quantum algorithm." In: *Phys. Rev. C* 105 (6 June 2022), p. 064317. DOI: 10.1103/PhysRevC.105.064317. URL: https://link.aps.org/doi/10.1103/PhysRevC.105.064317.

[139] E. F. Dumitrescu, A. J. McCaskey, G. Hagen, G. R. Jansen, T. D. Morris, T. Papenbrock, R. C. Pooser, D. J. Dean, and P. Lougovski. "Cloud Quantum Computing of an Atomic Nucleus." In: *Phys. Rev. Lett.* 120 (21 May 2018), p. 210501. DOI: 10.1103/PhysRevLett.120.210501. URL: https://link.aps.org/doi/10.1103/PhysRevLett.120.210501.

[140] Saverio Monaco, Oriel Kiss, Antonio Mandarino, Sofia Vallecorsa, and Michele Grossi. "Quantum phase detection generalization from marginal quantum neural network models." In: *Phys. Rev. B* 107 (8 Feb. 2023), p. L081105. DOI: 10.1103/PhysRevB.107.L081105. URL: https://link.aps.org/doi/10.1103/PhysRevB.107.L081105.

[141] Axel Pérez-Obiol, AM Romero, J Menéndez, A Rios, A García-Sáez, and B Juliá-Díaz. "Nuclear shell-model simulation in digital quantum computers." In: *Scientific Reports* 13.1 (2023), p. 12291. DOI: 10.1038/s41598-023-39263-7. URL: http://dx.doi.org/10.1038/s41598-023-39263-7.

[142] Paulin de Schoulepnikoff, Oriel Kiss, Sofia Vallecorsa, Giuseppe Carleo, and Michele Grossi. "Hybrid ground-state quantum algorithms based on neural Schrödinger forging." In: *Phys. Rev. Res.* 6 (2 Apr. 2024), p. 023021. DOI: 10.1103/PhysRevResearch.6.023021. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.6.023021.

[143] Shweta Sahoo, Utkarsh Azad, and Harjinder Singh. "Quantum phase recognition using quantum tensor networks." In: *The European Physical Journal Plus* 137.12 (Dec. 2022). ISSN: 2190-5444. DOI: 10.1140/epjp/s13360-022-03587-6. URL: http://dx.doi.org/10.1140/epjp/s13360-022-03587-6.

[144] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M. Chow, and Jay M. Gambetta. "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets." In: *Nature* 549 (2017), pp. 242–246. DOI: https://doi.org/10.1038/nature23879.

[145] Alexis Ralli, Peter J. Love, Andrew Tranter, and Peter V. Coveney. "Implementation of measurement reduction for the variational quantum eigensolver." In: *Phys. Rev. Res.* 3 (3 Aug. 2021), p. 033195. DOI: 10.1103/PhysRevResearch.3.033195. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.3.033195.

[146] Utkarsh Azad and Harjinder Singh. "Quantum chemistry calculations using energy derivatives on quantum computers." In: *Chemical Physics* 558 (June 2022), p. 111506. ISSN: 0301-0104. DOI: 10.1016/j.chemphys.2022.111506. URL: http://dx.doi.org/10.1016/j.chemphys.2022.111506.

[147] Dave Wecker, Matthew B. Hastings, and Matthias Troyer. "Progress towards practical quantum variational algorithms." In: *Phys. Rev. A* 92 (4 Oct. 2015), p. 042303. DOI: 10.1103/PhysRevA.92.042303. URL: https://link.aps.org/doi/10.1103/PhysRevA.92.042303.

[148] Jérôme F. Gonthier, Maxwell D. Radin, Corneliu Buda, Eric J. Doskocil, Clena M. Abuan, and Jhonathan Romero. "Measurements as a roadblock to near-term practical quantum advantage in chemistry: Resource analysis." In: *Phys. Rev. Res.* 4 (3 Aug. 2022), p. 033154. DOI: 10.1103/PhysRevResearch.4.033154. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.4.033154.

[149] Zhenyu Cai. "Resource Estimation for Quantum Variational Simulations of the Hubbard Model." In: *Phys. Rev. Appl.* 14 (1 July 2020), p. 014059. DOI: 10.1103/PhysRevApplied.14.014059. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.14.014059.

[150] Michael Ragone, Bojko N. Bakalov, Frédéric Sauvage, Alexander F. Kemper, Carlos Ortiz Marrero, Martin Larocca, and M. Cerezo. *A Unified Theory of Barren Plateaus for Deep Parametrized Quantum Circuits*. 2023. arXiv: 2309.09342 [quant-ph].

[151] Eric R Anschuetz and Bobak T Kiani. "Quantum variational algorithms are swamped with traps." In: *Nature Communications* 13.1 (2022), p. 7760. DOI: 10.1038/s41467-022-35364-5. URL: http://dx.doi.org/10.1038/s41467-022-35364-5.

[152] Zoë Holmes, Kunal Sharma, M. Cerezo, and Patrick J. Coles. "Connecting Ansatz Expressibility to Gradient Magnitudes and Barren Plateaus." In: *PRX Quantum* 3 (1 Jan. 2022), p. 010313. DOI: 10.1103/PRXQuantum.3.010313. URL: https://link.aps.org/doi/10.1103/PRXQuantum.3.010313.

[153] Marco Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J Coles. "Cost function dependent barren plateaus in shallow parametrized quantum circuits." In: *Nature communications* 12.1 (2021), p. 1791. DOI: 10.1038/s41467-021-21728-w. URL: http://dx.doi.org/10.1038/s41467-021-21728-w.

[154] Microsoft. *Azure Quantum Development Kit*. https://github.com/microsoft/qsharp. If you use this software in your work, please cite it using the following metadata. 2024.

[155] Tameem Albash and Daniel A. Lidar. "Adiabatic quantum computation." In: *Reviews of Modern Physics* 90.1 (Jan. 2018). ISSN: 1539-0756. DOI: 10.1103/revmodphys.90.015002. URL: http://dx.doi.org/10.1103/RevModPhys.90.015002.

[156] Erik Torrontegui, Sara Ibáñez, Sofia Martínez-Garaot, Michele Modugno, Adolfo del Campo, David Guéry-Odelin, Andreas Ruschhaupt, Xi Chen, and Juan Gonzalo Muga. "Shortcuts to Adiabaticity." In: *Advances in Atomic, Molecular, and Optical Physics*. Elsevier, 2013, pp. 117–169. DOI: 10.1016/b978-0-12-408090-4.00002-5. URL: http://dx.doi.org/10.1016/B978-0-12-408090-4.00002-5.

[157] M V Berry. "Transitionless quantum driving." In: *Journal of Physics A: Mathematical and Theoretical* 42.36 (Aug. 2009), p. 365303. DOI: 10.1088/1751-8113/42/36/365303. URL: https://dx.doi.org/10.1088/1751-8113/42/36/365303.

[158] Mustafa Demirplak and Stuart A. Rice. "Adiabatic Population Transfer with Control Fields." In: *The Journal of Physical Chemistry A* 107.46 (2003), pp. 9937–9945. DOI: 10.1021/jp030708a. URL: https://doi.org/10.1021/jp030708a.

[159] Mustafa Demirplak and Stuart A. Rice. "Assisted Adiabatic Passage Revisited." In: *The Journal of Physical Chemistry B* 109.14 (2005), pp. 6838–6844. DOI: 10.1021/jp040647w.

[160] Ieva Čepaitė, Anatoli Polkovnikov, Andrew J. Daley, and Callum W. Duncan. "Counterdiabatic Optimized Local Driving." In: *PRX Quantum* 4 (1 Jan. 2023), p. 010312. DOI: 10.1103/PRXQuantum.4.010312. URL: https://link.aps.org/doi/10.1103/PRXQuantum.4.010312.

[161] Francesco Pio Barone, Oriel Kiss, Michele Grossi, Sofia Vallecorsa, and Antonio Mandarino. "Counterdiabatic optimized driving in quantum phase sensitive models." In: *New Journal of Physics* 26.033031 (2024). URL: http://iopscience.iop.org/article/10.1088/1367-2630/ad313e.

[162] Harry J Lipkin, N Meshkov, and AJ Glick. "Validity of many-body approximation methods for a solvable model:(I). Exact solutions and perturbation theory." In: *Nuclear Physics* 62.2 (1965), pp. 188–198.

[163] N Meshkov, AJ Glick, and HJ Lipkin. "Validity of many-body approximation methods for a solvable model:(II). Linearization procedures." In: *Nuclear Physics* 62.2 (1965), pp. 199–210.

[164] AJ Glick, HJ Lipkin, and N Meshkov. "Validity of many-body approximation methods for a solvable model:(III). Diagram summations." In: *Nuclear Physics* 62.2 (1965), pp. 211–224.

[165] Kenneth Robbins and Peter J. Love. "Benchmarking near-term quantum devices with the variational quantum eigensolver and the Lipkin-Meshkov-Glick model." In: *Phys. Rev. A* 104 (2 Aug. 2021), p. 022412. DOI: 10.1103/PhysRevA.104.022412. URL: https://link.aps.org/doi/10.1103/PhysRevA.104.022412.

[166] Asahi Chikaoka and Haozhao Liang. "Quantum computing for the Lipkin model with unitary coupled cluster and structure learning ansatz." In: *Chinese Physics C* 46.2 (Feb. 2022), p. 024106. DOI: 10.1088/1674-1137/ac380a. URL: https://doi.org/10.1088/1674-1137/ac380a.

[167] S. Lerma H. and J. Dukelsky. "The Lipkin–Meshkov–Glick model as a particular limit of the SU(1,1) Richardson–Gaudin integrable models." In: *Nuclear Physics B* 870.2 (2013), pp. 421–443. ISSN: 0550-3213. DOI: https://doi.org/10.1016/j.nuclphysb.2013.01.019. URL: https://www.sciencedirect.com/science/article/pii/S0550321313000564.

[168] Oscar Higgott, Daochen Wang, and Stephen Brierley. "Variational Quantum Computation of Excited States." In: *Quantum* 3 (2019), p. 156. DOI: https://doi.org/10.22331/q-2019-07-01-156.

[169] Gang Chen and JQ Liang. "Unconventional quantum phase transition in the finite-size Lipkin–Meshkov–Glick model." In: *New Journal of Physics* 8.12 (2006), p. 297.

[170] D. Kraft. *A Software Package for Sequential Quadratic Programming*. Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt Köln: Forschungsbericht. Wiss. Berichtswesen d. DFVLR, 1988. URL: https://books.google.co.uk/books?id=4rKaGwAACAAJ.

[171] Stuart M. Harwood, Dimitar Trenev, Spencer T. Stober, Panagiotis Barkoutsos, Tanvi P. Gujarati, Sarah Mostame, and Donny Greenberg. "Improving the Variational Quantum Eigensolver Using Variational Adiabatic Quantum Computing." In: *ACM Transactions on Quantum Computing* 3.1 (Jan. 2022). ISSN: 2643-6809. DOI: 10.1145/3479197. URL: https://doi.org/10.1145/3479197.

[172] Federico Becca and Sandro Sorella. *Quantum Monte Carlo Approaches for Correlated Systems*. Cambridge University Press, 2017.

[173] S. Cohen and D. Kurath. "Effective interactions for the 1p shell." In: *Nuclear Physics* 73.1 (1965), pp. 1–24. ISSN: 0029-5582. DOI: https://doi.org/10.1016/0029-5582(65)90148-3. URL: https://www.sciencedirect.com/science/article/pii/0029558265901483.

[174] I. Stetcu, A. Baroni, and J. Carlson. "Variational approaches to constructing the many-body nuclear ground state for quantum computing." In: *Phys. Rev. C* 105 (6 June 2022), p. 064308. DOI: 10.1103/PhysRevC.105.064308. URL: https://link.aps.org/doi/10.1103/PhysRevC.105.064308.

[175] Francesco A. Evangelista, Garnet Kin-Lic Chan, and Gustavo E. Scuseria. "Exact parameterization of fermionic wave functions via unitary coupled cluster theory." In: *The Journal of Chemical Physics* 151.24 (2019), p. 244112. DOI: 10.1063/1.5133059. URL: https://doi.org/10.1063/1.5133059.

[176] Jonathan Romero, Ryan Babbush, Jarrod R McClean, Cornelius Hempel, Peter J Love, and Alán Aspuru-Guzik. "Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz." In: *Quantum Sci. Technol.* 4.1 (2019), p. 014008. DOI: 10.1088/2058-9565/aad3e4. URL: https://doi.org/10.1088/2058-9565/aad3e4.

[177] Qing Guo and Ping-Xing Chen. "Optimization of VQE-UCC Algorithm Based on Spin State Symmetry." In: *Frontiers in Physics* 9 (2021). DOI: 10.3389/fphy.2021.735321. URL: https://www.frontiersin.org/article/10.3389/fphy.2021.735321.

[178] Yordan S. Yordanov, David R. M. Arvidsson-Shukur, and Crispin H. W. Barnes. "Efficient quantum circuits for quantum computational chemistry." In: *Phys. Rev. A* 102 (6 Dec. 2020), p. 062612. DOI: 10.1103/PhysRevA.102.062612. URL: https://link.aps.org/doi/10.1103/PhysRevA.102.062612.

[179] Andrea Skolik, Jarrod R. McClean, Masoud Mohseni, Patrick van der Smagt, and Martin Leib. "Layerwise learning for quantum neural networks." In: *Quantum machine intelligence* 3 (2021), p. 5. DOI: https://doi.org/10.1007.

[180] Julien Gacon, Christa Zoufal, Giuseppe Carleo, and Stefan Woerner. "Simultaneous Perturbation Stochastic Approximation of the Quantum Fisher Information." In: *Quantum* 5 (2021), p. 567. DOI: https://quantum-journal.org/papers/q-2021-10-20-567/.

[181] James Stokes, Josh Izaac, Nathan Killoran, and Giuseppe Carleo. "Quantum Natural Gradient." In: *Quantum* 4 (May 2020), p. 269. ISSN: 2521-327X. DOI: 10.22331/q-2020-05-25-269. URL: https://doi.org/10.22331/q-2020-05-25-269.

[182] Gushu Li, Yufei Ding, and Yuan Xie. "Tackling the Qubit Mapping Problem for NISQ-Era Quantum Devices." In: *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. ASPLOS '19. Providence, RI, USA: Association for Computing Machinery, 2019, pp. 1001–1014. ISBN: 9781450362405. DOI: 10.1145/3297858.3304023. URL: https://doi.org/10.1145/3297858.3304023.

[183] A. Pérez-Obiol, A. M. Romero, J. Menéndez, A. Rios, A. García-Sáez, and B. Juliá-Díaz. "Nuclear shell-model simulation in digital quantum computers." In: *Scientific Reports* 13.1 (2023), p. 12291. DOI: 10.1038/s41598-023-39263-7. URL: https://doi.org/10.1038/s41598-023-39263-7.

[184] Alessandro Roggero, Andy C. Y. Li, Joseph Carlson, Rajan Gupta, and Gabriel N. Perdue. "Quantum computing for neutrino-nucleus scattering." In: *Phys. Rev. D* 101 (7 Apr. 2020), p. 074038. DOI: 10.1103/PhysRevD.101.074038. URL: https://link.aps.org/doi/10.1103/PhysRevD.101.074038.

[185] A. Baroni, J. Carlson, R. Gupta, Andy C. Y. Li, G. N. Perdue, and A. Roggero. "Nuclear two point correlation functions on a quantum computer." In: *Phys. Rev. D* 105 (7 Apr. 2022), p. 074503. DOI: 10.1103/PhysRevD.105.074503. URL: https://link.aps.org/doi/10.1103/PhysRevD.105.074503.

[186] Paulo F. Bedaque and Ubirajara van Kolck. "EFFECTIVE FIELD THEORY FOR FEW-NUCLEON SYSTEMS." In: *Annual Review of Nuclear and Particle Science* 52.1 (2002), pp. 339–396. DOI: 10.1146/annurev.nucl.52.050102.090637. eprint: https://doi.org/10.1146/annurev.nucl.52.050102.090637. URL: https://doi.org/10.1146/annurev.nucl.52.050102.090637.

[187] A Rokash, E Epelbaum, H Krebs, D Lee, and U-G Meißner. "Finite volume effects in low-energy neutron–deuteron scattering." In: *Journal of Physics G: Nuclear and Particle Physics* 41.1 (Dec. 2013), p. 015105. DOI: 10.1088/0954-3899/41/1/015105. URL: https://dx.doi.org/10.1088/0954-3899/41/1/015105.

[188] Olivia Di Matteo, Anna McCoy, Peter Gysbers, Takayuki Miyagi, R. M. Woloshyn, and Petr Navrátil. "Improving Hamiltonian encodings with the Gray code." In: *Phys. Rev. A* 103 (4 Apr. 2021), p. 042405. DOI: 10.1103/PhysRevA.103.042405. URL: https://link.aps.org/doi/10.1103/PhysRevA.103.042405.

[189] Hongbin Liu, Guang Hao Low, Damian S Steiger, Thomas Häner, Markus Reiher, and Matthias Troyer. "Prospects of quantum computing for molecular sciences." In: *Materials Theory* 6.1 (2022), p. 11. DOI: 10.1186/s41313-021-00039-z. URL: http://dx.doi.org/10.1186/s41313-021-00039-z.

[190] Andrew M. Childs, Dmitri Maslov, Yunseong Nam, Neil J. Ross, and Yuan Su. "Toward the first quantum simulation with quantum speedup." In: *Proceedings of the National Academy of Sciences* 115.38 (2018), pp. 9456–9461. DOI: 10.1073/pnas.1801723115. URL: https://www.pnas.org/doi/abs/10.1073/pnas.1801723115.

[191] Stephen P. Jordan, Keith S. M. Lee, and John Preskill. "Quantum Algorithms for Quantum Field Theories." In: *Science* 336.6085 (2012), pp. 1130–1133. DOI: 10.1126/science.1217069. eprint: https://www.science.org/doi/pdf/10.1126/science.1217069. URL: https://www.science.org/doi/abs/10.1126/science.1217069.

[192] J. Argüello-Luengo, A. González-Tudela, T. Shi, P. Zoller, and J. I. Cirac. "Analogue quantum chemistry simulation." In: *Nature* 574 (2019), pp. 215–218. DOI: 10.1038/s41586-019-1614-4.

[193] Andrew M. Childs, Yuan Su, Minh C. Tran, Nathan Wiebe, and Shuchen Zhu. "Theory of Trotter Error with Commutator Scaling." In: *Physical Review X* 11.1 (Feb. 2021), p. 011020. DOI: 10.1103/PhysRevX.11.011020. URL: https://link.aps.org/doi/10.1103/PhysRevX.11.011020 (visited on 04/06/2023).

[194] Earl Campbell. "Random Compiler for Fast Hamiltonian Simulation." In: *Phys. Rev. Lett.* 123 (7 Aug. 2019), p. 070503. DOI: 10.1103/PhysRevLett.123.070503. URL: https://link.aps.org/doi/10.1103/PhysRevLett.123.070503.

[195] Chi-Fang Chen, Hsin-Yuan Huang, Richard Kueng, and Joel A. Tropp. "Concentration for Random Product Formulas." In: *PRX Quantum* 2 (4 Oct. 2021), p. 040305. DOI: 10.1103/PRXQuantum.2.040305. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.040305.

[196] Oriel Kiss, Michele Grossi, and Alessandro Roggero. "Importance sampling for stochastic quantum simulations." In: *Quantum* 7 (Apr. 2023), p. 977. ISSN: 2521-327X. DOI: 10.22331/q-2023-04-13-977. URL: https://doi.org/10.22331/q-2023-04-13-977.

[197] Surya T. Tokdar and Robert E. Kass. "Importance sampling: a review." In: *WIREs Computational Statistics* 2.1 (2010), pp. 54–60. DOI: https://doi.org/10.1002/wics.56. eprint: https://wires.onlinelibrary.wiley.com/doi/pdf/10.1002/wics.56. URL: https://wires.onlinelibrary.wiley.com/doi/abs/10.1002/wics.56.

[198] Dominic W. Berry, Andrew M. Childs, Yuan Su, Xin Wang, and Nathan Wiebe. "Time-dependent Hamiltonian simulation with L^1-norm scaling." In: *Quantum* 4 (Apr. 2020), p. 254. ISSN: 2521-327X. DOI: 10.22331/q-2020-04-20-254. URL: https://doi.org/10.22331/q-2020-04-20-254.

[199] Matthew Hagan and Nathan Wiebe. "Composite Quantum Simulations." In: *Quantum* 7 (Nov. 2023), p. 1181. ISSN: 2521-327X. DOI: 10.22331/q-2023-11-14-1181. URL: https://doi.org/10.22331/q-2023-11-14-1181.

[200] Shi Jin and Xiantao Li. "A Partially Random Trotter Algorithm for Quantum Hamiltonian Simulations." In: *ArXiv e-prints* 2109.07987 (2021). DOI: 10.48550/ARXIV.2109.07987. URL: https://arxiv.org/abs/2109.07987.

[201] Abhishek Rajput, Alessandro Roggero, and Nathan Wiebe. "Hybridized Methods for Quantum Simulation in the Interaction Picture." In: *Quantum* 6 (Aug. 2022), p. 780. ISSN: 2521-327X. DOI: 10.22331/q-2022-08-17-780. URL: https://doi.org/10.22331/q-2022-08-17-780.

[202] Andrew M. Childs, Aaron Ostrander, and Yuan Su. "Faster quantum simulation by randomization." In: *Quantum* 3 (2019). ISSN: 2521-327X. DOI: 10.22331/q-2019-09-02-182. URL: https://doi.org/10.22331/q-2019-09-02-182.

[203] Paul K. Faehrmann, Mark Steudtner, Richard Kueng, Mária Kieferová, and Jens Eisert. "Randomizing multi-product formulas for Hamiltonian simulation." In: *Quantum* 6 (Sept. 2022), p. 806. ISSN: 2521-327X. DOI: 10.22331/q-2022-09-19-806. URL: https://doi.org/10.22331/q-2022-09-19-806.

[204] Guang Hao Low and Isaac L. Chuang. "Optimal Hamiltonian Simulation by Quantum Signal Processing." In: *Phys. Rev. Lett.* 118 (1 Jan. 2017), p. 010501. DOI: 10.1103/PhysRevLett.118.010501. URL: https://link.aps.org/doi/10.1103/PhysRevLett.118.010501.

[205] G. Sigl and G. Raffelt. "General kinetic description of relativistic mixed neutrinos." In: *Nuclear Physics B* 406.1 (1993), pp. 423–451. ISSN: 0550-3213. DOI: https://doi.org/10.1016/0550-3213(93)90175-O. URL: https://www.sciencedirect.com/science/article/pii/055032139390175O.

[206] Sergio Pastor and Georg Raffelt. "Flavor Oscillations in the Supernova Hot Bubble Region: Nonlinear Effects of Neutrino Background." In: *Phys. Rev. Lett.* 89 (19 Oct. 2002), p. 191101. DOI: 10.1103/PhysRevLett.89.191101. URL: https://link.aps.org/doi/10.1103/PhysRevLett.89.191101.

[207] Sergio Pastor, Georg Raffelt, and Dmitry V. Semikoz. "Physics of synchronized neutrino oscillations caused by self-interactions." In: *Phys. Rev. D* 65 (5 Feb. 2002), p. 053011. DOI: 10.1103/PhysRevD.65.053011. URL: https://link.aps.org/doi/10.1103/PhysRevD.65.053011.

[208] Nicole F. Bell, Andrew A. Rawlinson, and R. F. Sawyer. "Speedup through entanglement: Many body effects in neutrino processes." In: *Phys. Lett. B* 573 (2003), pp. 86–93. DOI: 10.1016/j.physletb.2003.08.035. arXiv: hep-ph/0304082.

[209] James Pantaleone. "Neutrino oscillations at high densities." In: *Physics Letters B* 287.1-3 (1992), pp. 128–132.

[210] Georg G. Raffelt and Alexei Yu. Smirnov. "Adiabaticity and spectral splits in collective neutrino transformations." In: *Phys. Rev. D* 76 (2007), p. 125008. DOI: 10.1103/PhysRevD.76.125008. arXiv: 0709.4641 [hep-ph].

[211] Joshua D. Martin, Duff Neill, A. Roggero, Huaiyu Duan, and J. Carlson. "Equilibration of quantum many-body fast neutrino flavor oscillations." In: *Phys. Rev. D* 108 (12 Dec. 2023), p. 123010. DOI: 10.1103/PhysRevD.108.123010. URL: https://link.aps.org/doi/10.1103/PhysRevD.108.123010.

[212] R.F. Sawyer. "'Classical' instabilities and 'quantum' speed-up in the evolution of neutrino clouds." In: (Aug. 2004). arXiv: hep-ph/0408265.

[213] Sovan Chakraborty, Rasmus Sloth Hansen, Ignacio Izaguirre, and Georg G Raffelt. "Self-induced neutrino flavor conversion without flavor mixing." In: *Journal of Cosmology and Astroparticle Physics* 2016.03 (2016), p. 042.

[214] Basudeb Dasgupta, Alessandro Mirizzi, and Manibrata Sen. "Fast neutrino flavor conversions near the supernova core with realistic flavor-dependent angular distributions." In: *Journal of Cosmology and Astroparticle Physics* 2017.02 (2017), p. 019.

[215] Joshua D. Martin, J. Carlson, and Huaiyu Duan. "Spectral swaps in a two-dimensional neutrino ring model." In: *Phys. Rev.* D101.2 (2020), p. 023007. DOI: 10.1103/PhysRevD.101.023007. arXiv: 1911.09772 [hep-ph].

[216] Alessandro Roggero, Ermal Rrapaj, and Zewei Xiong. "Entanglement and correlations in fast collective neutrino flavor oscillations." In: *Phys. Rev. D* 106 (4 Aug. 2022), p. 043022. DOI: 10.1103/PhysRevD.106.043022. URL: https://link.aps.org/doi/10.1103/PhysRevD.106.043022.

[217] Zewei Xiong and Yong-Zhong Qian. "Stationary solutions for fast flavor oscillations of a homogeneous dense neutrino gas." In: *Phys. Lett. B* 820 (2021), p. 136550. DOI: 10.1016/j.physletb.2021.136550. arXiv: 2104.05618 [astro-ph.HE].

[218] Maria Cristina Volpe. "Neutrinos from dense: flavor mechanisms, theoretical approaches, observations, new directions." In: (Jan. 2023). arXiv: 2301.11814 [hep-ph].

[219] L. Wolfenstein. "Neutrino oscillations in matter." In: *Phys. Rev. D* 17 (9 May 1978), pp. 2369–2374. DOI: 10.1103/PhysRevD.17.2369. URL: https://link.aps.org/doi/10.1103/PhysRevD.17.2369.

[220] SP Mikheev and A Yu Smirnov. "Resonance amplification of oscillations in matter and spectroscopy of solar neutrinos." In: *Yadernaya Fizika* 42.6 (1985), pp. 1441–1448.

[221] A.B. Balantekin and Y. Pehlivan. "Neutrino-Neutrino Interactions and Flavor Mixing in Dense Matter." In: *J. Phys. G* 34 (2007), pp. 47–66. DOI: 10.1088/0954-3899/34/1/004. arXiv: astro-ph/0607527.

[222] J. M. Deutsch. "Quantum statistical mechanics in a closed system." In: *Phys. Rev. A* 43 (4 Feb. 1991), pp. 2046–2049. DOI: 10.1103/PhysRevA.43.2046. URL: https://link.aps.org/doi/10.1103/PhysRevA.43.2046.

[223] Mark Srednicki. "Chaos and quantum thermalization." In: *Phys. Rev. E* 50 (2 Aug. 1994), pp. 888–901. DOI: 10.1103/PhysRevE.50.888. URL: https://link.aps.org/doi/10.1103/PhysRevE.50.888.

[224] Mark Srednicki. "Thermal fluctuations in quantized chaotic systems." In: *Journal of Physics A: Mathematical and General* 29.4 (1996), pp. L75–L79. DOI: 10.1088/0305-4470/29/4/003.

[225] Anatoli Polkovnikov Luca D'Alessio Yariv Kafri and Marcos Rigol. "From quantum chaos and eigenstate thermalization to statistical mechanics and thermodynamics." In: *Advances in Physics* 65.3 (2016), pp. 239–362. DOI: 10.1080/00018732.2016.1198134.

[226] Lydzba Patrycja, Yicheng Zhang, Marcos Rigol, and Lev Vidmar. "Single-particle eigenstate thermalization in quantum-chaotic quadratic Hamiltonians." In: *Phys. Rev. B* 104 (21 Dec. 2021), p. 214203. DOI: 10.1103/PhysRevB.104.214203. URL: https://link.aps.org/doi/10.1103/PhysRevB.104.214203.

[227] C. L. Baldwin, C. R. Laumann, A. Pal, and A. Scardicchio. "Clustering of Nonergodic Eigenstates in Quantum Spin Glasses." In: *Phys. Rev. Lett.* 118 (12 Mar. 2017), p. 127201. DOI: 10.1103/PhysRevLett.118.127201. URL: https://link.aps.org/doi/10.1103/PhysRevLett.118.127201.

[228] Naoto Shiraishi and Takashi Mori. "Systematic Construction of Counterexamples to the Eigenstate Thermalization Hypothesis." In: *Phys. Rev. Lett.* 119 (3 July 2017), p. 030601. DOI: 10.1103/PhysRevLett.119.030601. URL: https://link.aps.org/doi/10.1103/PhysRevLett.119.030601.

[229] K. Pakrouski, P. N. Pallegar, F. K. Popov, and I. R. Klebanov. "Many-Body Scars as a Group Invariant Sector of Hilbert Space." In: *Phys. Rev. Lett.* 125 (23 Dec. 2020), p. 230602. DOI: 10.1103/PhysRevLett.125.230602. URL: https://link.aps.org/doi/10.1103/PhysRevLett.125.230602.

[230] Sanjay Moudgalya, B. Andrei Bernevig, and Nicolas Regnault. "Quantum many-body scars and Hilbert space fragmentation: a review of exact results." In: *Reports on Progress in Physics* 85.8 (2022), p. 086501. DOI: 10.1088/1361-6633/ac73a0.

[231] Ian D. Kivlichan, Jarrod McClean, Nathan Wiebe, Craig Gidney, Alán Aspuru-Guzik, Garnet Kin-Lic Chan, and Ryan Babbush. "Quantum Simulation of Electronic Structure with Linear Depth and Connectivity." In: *Phys. Rev. Lett.* 120 (11 Mar. 2018), p. 110501. DOI: 10.1103/PhysRevLett.120.110501. URL: https://link.aps.org/doi/10.1103/PhysRevLett.120.110501.

[232] Benjamin Hall, Alessandro Roggero, Alessandro Baroni, and Joseph Carlson. "Simulation of collective neutrino oscillations on a quantum computer." In: *Phys. Rev. D* 104 (6 Sept. 2021), p. 063009. DOI: 10.1103/PhysRevD.104.063009. URL: https://link.aps.org/doi/10.1103/PhysRevD.104.063009.

[233] Valentina Amitrano, Alessandro Roggero, Piero Luchi, Francesco Turro, Luca Vespucci, and Francesco Pederiva. "Trapped-ion quantum simulation of collective neutrino oscillations." In: *Phys. Rev. D* 107 (2 Jan. 2023), p. 023007. DOI: 10.1103/PhysRevD.107.023007. URL: https://link.aps.org/doi/10.1103/PhysRevD.107.023007.

[234] Denis Lacroix, Angel Bauge, Bulent Yilmaz, Mariane Mangin-Brinet, Alessandro Roggero, and A. Baha Balantekin. "Phase-space methods for neutrino oscillations: Extension to multibeams." In: *Phys. Rev. D* 110 (10 Nov. 2024), p. 103027. DOI: 10.1103/PhysRevD.110.103027. URL: https://link.aps.org/doi/10.1103/PhysRevD.110.103027.

[235] Oriel Kiss, Daniil Teplitskiy, Michele Grossi, and Antonio Mandarino. "Statistics of topological defects across a phase transition in a superconducting quantum processor." In: *Quantum Sci. Technol.* (June 2025), p. 035037. DOI: 10.1088/2058-9565/addf75.

[236] T.W.B. Kibble. "Some implications of a cosmological phase transition." In: *Physics Reports* 67.1 (1980), pp. 183–199. ISSN: 0370-1573. DOI: https://doi.org/10.1016/0370-1573(80)90091-5. URL: https://www.sciencedirect.com/science/article/pii/0370157380900915.

[237] Wojciech H. Zurek. "Cosmological experiments in superfluid helium?" In: *Nature* 317 (1985), pp. 505–508. DOI: 10.1038/317505a0. URL: https://doi.org/10.1038/317505a0.

[238] Jacek Dziarmaga. "Dynamics of a quantum phase transition and relaxation to a steady state." In: *Advances in Physics* 59.6 (2010), pp. 1063–1189. DOI: https://doi.org/10.1080/00018732.2010.514702.

[239] Adolfo Del Campo and Wojciech H Zurek. "Universality of phase transition dynamics: Topological defects from symmetry breaking." In: *International Journal of Modern Physics A* 29.08 (2014), p. 1430018. DOI: https://doi.org/10.1142/S0217751X1430018X.

[240] Anatoli Polkovnikov, Krishnendu Sengupta, Alessandro Silva, and Mukund Vengalattore. "Colloquium: Nonequilibrium dynamics of closed interacting quantum systems." In: *Rev. Mod. Phys.* 83 (3 Aug. 2011), pp. 863–883. DOI: 10.1103/RevModPhys.83.863. URL: https://link.aps.org/doi/10.1103/RevModPhys.83.863.

[241] Alexander Keesling et al. "Quantum Kibble-Zurek mechanism and critical dynamics on a programmable Rydberg simulator." In: *Nature* 568 (2019), pp. 207–211. DOI: 10.1038/s41586-019-1070-1. URL: https://doi.org/10.1038/s41586-019-1070-1.

[242] Jin-Ming Cui, Fernando Javier Gómez-Ruiz, Yun-Feng Huang, Chuan-Feng Li, Guang-Can Guo, and Adolfo del Campo. "Experimentally testing quantum critical dynamics beyond the Kibble-Zurek mechanism." In: *Communications Physics* 3 (2020), p. 44. DOI: 10.1038/s42005-020-0306-6. URL: https://doi.org/10.1038/s42005-020-0306-6.

[243] Jacek Dziarmaga. "Dynamics of a Quantum Phase Transition: Exact Solution of the Quantum Ising Model." In: *Phys. Rev. Lett.* 95 (24 Dec. 2005), p. 245701. DOI: 10.1103/PhysRevLett.95.245701. URL: https://link.aps.org/doi/10.1103/PhysRevLett.95.245701.

[244] Adolfo del Campo. "Universal Statistics of Topological Defects Formed in a Quantum Phase Transition." In: *Phys. Rev. Lett.* 121 (20 Nov. 2018), p. 200601. DOI: 10.1103/PhysRevLett.121.200601. URL: https://link.aps.org/doi/10.1103/PhysRevLett.121.200601.

[245] Hua-Bi Zeng, Chuan-Yin Xia, and Adolfo del Campo. "Universal Breakdown of Kibble-Zurek Scaling in Fast Quenches across a Phase Transition." In: *Phys. Rev. Lett.* 130 (6 Feb. 2023), p. 060402. DOI: 10.1103/PhysRevLett.130.060402. URL: https://link.aps.org/doi/10.1103/PhysRevLett.130.060402.

[246] Marek M. Rams, Jacek Dziarmaga, and Wojciech H. Zurek. "Symmetry Breaking Bias and the Dynamics of a Quantum Phase Transition." In: *Phys. Rev. Lett.* 123 (13 Sept. 2019), p. 130603. DOI: 10.1103/PhysRevLett.123.130603. URL: https://link.aps.org/doi/10.1103/PhysRevLett.123.130603.

[247] Federico Balducci, Mathieu Beau, Jing Yang, Andrea Gambassi, and Adolfo del Campo. "Large Deviations beyond the Kibble-Zurek Mechanism." In: *Phys. Rev. Lett.* 131 (23 Dec. 2023), p. 230401. DOI: 10.1103/PhysRevLett.131.230401. URL: https://link.aps.org/doi/10.1103/PhysRevLett.131.230401.

[248] Markus Schmitt, Marek M. Rams, Jacek Dziarmaga, Markus Heyl, and Wojciech H. Zurek. "Quantum phase transition dynamics in the two-dimensional transverse-field Ising model." In: *Science Advances* 8.37 (2022), eabl6850. DOI: 10.1126/sciadv.abl6850.

[249] I-Kang Liu, Jacek Dziarmaga, Shih-Chuan Gou, Franco Dalfovo, and Nick P. Proukakis. "Kibble-Zurek dynamics in a trapped ultracold Bose gas." In: *Phys. Rev. Res.* 2 (3 Aug. 2020), p. 033183. DOI: 10.1103/PhysRevResearch.2.033183. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.2.033183.

[250] Trond I Andersen, Nikita Astrakhantsev, Amir Karamlou, Julia Berndtsson, Johannes Motruk, Aaron Szasz, Jonathan A Gross, Tom Westerhout, Yaxing Zhang, Ebrahim Forati, et al. *Thermalization and Criticality on an Analog-Digital Quantum Simulator*. 2024. arXiv: 2405.17385 [quant-ph].

[251] Alexander Miessen, Daniel J Egger, Ivano Tavernelli, and Guglielmo Mazzola. *Benchmarking digital quantum simulations and optimization above hundreds of qubits using quantum critical dynamics*. 2024. arXiv: 2404.08053 [quant-ph].

[252] Max Born and Vladimir Fock. "Beweis des Adiabatensatzes." In: *Zeitschrift für Physik* 51 (1928), pp. 165–180. DOI: 10.1007/BF01343193.

[253] Hongzheng Zhao, Marin Bukov, Markus Heyl, and Roderich Moessner. "Adaptive Trotterization for Time-Dependent Hamiltonian Quantum Dynamics Using Piecewise Conservation Laws." In: *Phys. Rev. Lett.* 133 (1 July 2024), p. 010603. DOI: 10.1103/PhysRevLett.133.010603. URL: https://link.aps.org/doi/10.1103/PhysRevLett.133.010603.

[254] Dominic W. Berry, Andrew M. Childs, Yuan Su, Xin Wang, and Nathan Wiebe. "Time-dependent Hamiltonian simulation with L^1-norm scaling." In: *Quantum* 4 (Apr. 2020), p. 254. ISSN: 2521-327X. DOI: 10.22331/q-2020-04-20-254. URL: https://doi.org/10.22331/q-2020-04-20-254.

[255] Abhishek Rajput, Alessandro Roggero, and Nathan Wiebe. "Hybridized Methods for Quantum Simulation in the Interaction Picture." In: *Quantum* 6 (Aug. 2022), p. 780. ISSN: 2521-327X. DOI: 10.22331/q-2022-08-17-780. URL: https://doi.org/10.22331/q-2022-08-17-780.

[256] Mária Kieferová, Artur Scherer, and Dominic W Berry. "Simulating the dynamics of time-dependent Hamiltonians with a truncated Dyson series." In: *Physical Review A* 99.4 (Apr. 2019), p. 042314. DOI: 10.1103/PhysRevA.99.042314.

[257] Kunal Sharma and Minh C Tran. *Hamiltonian Simulation in the Interaction Picture Using the Magnus Expansion*. 2024. arXiv: 2404.02966 [quant-ph].

[258] Pablo Antonio Moreno Casares, Modjtaba Shokrian Zini, and Juan Miguel Arrazola. *Quantum simulation of time-dependent Hamiltonians via commutator-free quasi-Magnus operators*. 2024. arXiv: 2403.13889 [quant-ph].

[259] Jacob Watkins, Nathan Wiebe, Alessandro Roggero, and Dean Lee. *Time Dependent Hamiltonian Simulation Using Discrete Clock Constructions*. 2022. arXiv: 2203.11353 [quant-ph].

[260] Boyuan Shi and Florian Mintert. "Quantum simulations of time-dependent Hamiltonians beyond the quasistatic approximation." In: *Phys. Rev. Res.* 6 (2 Apr. 2024), p. 023097. DOI: 10.1103/PhysRevResearch.6.023097. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.6.023097.

[261] Leonid Abdurakhimov, Janos Adam, Hasnain Ahmad, Olli Ahonen, Manuel Algaba, Guillermo Alonso, Ville Bergholm, Rohit Beriwal, Matthias Beuerle, Clinton Bockstiegel, et al. *Technology and Performance Benchmarks of IQM's 20-Qubit Quantum Computer*. 2024. arXiv: 2408.12433 [quant-ph].

[262] Fabian Marxer et al. "Long-Distance Transmon Coupler with cz-Gate Fidelity above 99.8%." In: *PRX Quantum* 4 (1 Feb. 2023), p. 010314. DOI: 10.1103/PRXQuantum.4.010314. URL: https://link.aps.org/doi/10.1103/PRXQuantum.4.010314.

[263] A. H. Z. Kavaki and R. Lewis. "From square plaquettes to triamond lattices for SU(2) gauge theory." In: *Communications Physics* 7 (2024), p. 208. DOI: 10.1038/s42005-024-01697-4. URL: https://doi.org/10.1038/s42005-024-01697-4.

[264] Amara Katabarwa, Katerina Gratsea, Athena Caesura, and Peter D Johnson. "Early fault-tolerant quantum computing." In: *arXiv preprint arXiv:2311.14814* (2023).

[265] Jeremy Hartse and Alessandro Roggero. "Faster spectral density calculation using energy moments." In: *The European Physical Journal A* 59.3 (2023), p. 41.

[266] Lin Lin and Yu Tong. "Heisenberg-Limited Ground-State Energy Estimation for Early Fault-Tolerant Quantum Computers." In: *PRX Quantum* 3.1 (Feb. 2022). DOI: 10.1103/prxquantum.3.010318. URL: https://doi.org/10.1103/prxquantum.3.010318.

[267] Oriel Kiss, Utkarsh Azad, Borja Requena, Alessandro Roggero, David Wakeham, and Juan Miguel Arrazola. "Early Fault-Tolerant Quantum Algorithms in Practice: Application to Ground-State Energy Estimation." In: *Quantum* 9 (Apr. 2025), p. 1682. ISSN: 2521-327X. DOI: 10.22331/q-2025-04-01-1682. URL: https://doi.org/10.22331/q-2025-04-01-1682.

[268] Guang Hao Low and Isaac L Chuang. "Optimal Hamiltonian simulation by quantum signal processing." In: *Physical review letters* 118.1 (2017), p. 010501.

[269] András Gilyén, Yuan Su, Guang Hao Low, and Nathan 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.

[270] Y. Kikuchi, C. Mc Keever, L. Coopmans, Michael Lubasch, and Marcello Benedetti. "Realization of quantum signal processing on a noisy quantum computer." In: *npj Quantum Information* 9 (2023), p. 93. DOI: 10.1038/s41534-023-00762-0.

[271] Andrew M. Childs and Nathan Wiebe. "Hamiltonian simulation using linear combinations of unitary operations." In: *Quantum Info. Comput.* 12.11–12 (Nov. 2012), pp. 901–924. ISSN: 1533-7146. DOI: 10.5555/2481569.2481570.

[272] Sirui Lu, Mari Carmen Bañuls, and J. Ignacio Cirac. "Algorithms for Quantum Simulation at Finite Energies." In: *PRX Quantum* 2 (2 May 2021), p. 020321. DOI: 10.1103/PRXQuantum.2.020321. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.020321.

[273] Yanwu Gu, Yunheng Ma, Nicolò Forcellini, and Dong E. Liu. "Noise-Resilient Phase Estimation with Randomized Compiling." In: *Phys. Rev. Lett.* 130 (25 June 2023), p. 250601. DOI: 10.1103/PhysRevLett.130.250601. URL: https://link.aps.org/doi/10.1103/PhysRevLett.130.250601.

[274] Yulong Dong, Lin Lin, and Yu Tong. "Ground-State Preparation and Energy Estimation on Early Fault-Tolerant Quantum Computers via Quantum Eigenvalue Transformation of Unitary Matrices." In: *PRX Quantum* 3 (4 Oct. 2022), p. 040305. DOI: 10.1103/PRXQuantum.3.040305. URL: https://link.aps.org/doi/10.1103/PRXQuantum.3.040305.

[275] Max Bee-Lindgren, Zhengrong Qian, Matthew DeCross, Natalie C. Brown, Christopher N. Gilbreth, Jacob Watkins, Xilin Zhang, and Dean Lee. *Controlled Gate Networks Applied to Eigenvalue Estimation*. 2022. arXiv: 2208.13557 [quant-ph].

[276] Emmanuel Candès and Justin Romberg. "Sparsity and incoherence in compressive sampling." In: *Inverse Problems* 23.3 (Apr. 2007), p. 969. ISSN: 0266-5611. DOI: 10.1088/0266-5611/23/3/008. URL: https://dx.doi.org/10.1088/0266-5611/23/3/008 (visited on 06/01/2023).

[277] Emmanuel J. Candes, Michael B. Wakin, and Stephen P. Boyd. *Enhancing Sparsity by Reweighted L1 Minimization*. arXiv:0711.1612 [math, stat]. Nov. 2007. URL: http://arxiv.org/abs/0711.1612 (visited on 06/01/2023).

[278] Emmanuel J. Candès, Justin K. Romberg, and Terence Tao. "Stable signal recovery from incomplete and inaccurate measurements." In: *Communications on Pure and Applied Mathematics* 59.8 (2006), pp. 1207–1223. ISSN: 1097-0312. DOI: 10.1002/cpa.20124. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.20124 (visited on 06/01/2023).

[279] Alessandro Roggero and Joseph Carlson. "Dynamic linear response quantum algorithm." In: *Phys. Rev. C* 100 (3 Sept. 2019), p. 034610. DOI: 10.1103/PhysRevC.100.034610. URL: https://link.aps.org/doi/10.1103/PhysRevC.100.034610.

[280] A. Lovato, J. Carlson, S. Gandolfi, N. Rocco, and R. Schiavilla. "Ab Initio Study of (ν_ℓ, ℓ^-) and $(\overline{\nu}_\ell, \ell^+)$ Inclusive Scattering in ^{12}C: Confronting the MiniBooNE and T2K CCQE Data." In: *Phys. Rev. X* 10 (3 Sept. 2020), p. 031068. DOI: 10.1103/PhysRevX.10.031068. URL: https://link.aps.org/doi/10.1103/PhysRevX.10.031068.

[281] J. Carlson and R. Schiavilla. "Euclidean proton response in light nuclei." In: *Phys. Rev. Lett.* 68 (25 June 1992), pp. 3682–3685. DOI: 10.1103/PhysRevLett.68.3682. URL: https://link.aps.org/doi/10.1103/PhysRevLett.68.3682.

[282] J. Carlson, S. Gandolfi, F. Pederiva, Steven C. Pieper, R. Schiavilla, K. E. Schmidt, and R. B. Wiringa. "Quantum Monte Carlo methods for nuclear physics." In: *Rev. Mod. Phys.* 87 (3 Sept. 2015), pp. 1067–1118. DOI: 10.1103/RevModPhys.87.1067. URL: https://link.aps.org/doi/10.1103/RevModPhys.87.1067.

[283] A. Lovato, S. Gandolfi, J. Carlson, Steven C. Pieper, and R. Schiavilla. "Electromagnetic Response of ^{12}C: A First-Principles Calculation." In: *Phys. Rev. Lett.* 117 (8 Aug. 2016), p. 082501. DOI: 10.1103/PhysRevLett.117.082501. URL: https://link.aps.org/doi/10.1103/PhysRevLett.117.082501.

[284] Victor D. Efros, Winfred Leidemann, and Giuseppina Orlandini. "Response functions from integral transforms with a Lorentz kernel." In: *Physics Letters B* 338.2 (1994), pp. 130–133. ISSN: 0370-2693. DOI: https://doi.org/10.1016/0370-2693(94)91355-2. URL: https://www.sciencedirect.com/science/article/pii/0370269394913552.

[285] V D Efros, W Leidemann, G Orlandini, and N Barnea. "The Lorentz integral transform (LIT) method and its applications to perturbation-induced reactions." In: *Journal of Physics G: Nuclear and Particle Physics* 34.12 (Oct. 2007), R459. DOI: 10.1088/0954-3899/34/12/R02. URL: https://dx.doi.org/10.1088/0954-3899/34/12/R02.

[286] S. Bacca, N. Barnea, G. Hagen, G. Orlandini, and T. Papenbrock. "First Principles Description of the Giant Dipole Resonance in ^{16}**O**." In: *Phys. Rev. Lett.* 111 (12 Sept. 2013), p. 122502. DOI: 10.1103/PhysRevLett.111.122502. URL: https://link.aps.org/doi/10.1103/PhysRevLett.111.122502.

[287] S. Bacca, N. Barnea, G. Hagen, M. Miorelli, G. Orlandini, and T. Papenbrock. "Giant and pigmy dipole resonances in ^{4}He, ^{16,22}O, and ^{40}Ca from chiral nucleon-nucleon interactions." In: *Phys. Rev. C* 90 (6 Dec. 2014), p. 064619. DOI: 10.1103/PhysRevC.90.064619. URL: https://link.aps.org/doi/10.1103/PhysRevC.90.064619.

[288] J. E. Sobczyk, B. Acharya, S. Bacca, and G. Hagen. "Ab Initio Computation of the Longitudinal Response Function in ^{40}Ca." In: *Phys. Rev. Lett.* 127 (7 Aug. 2021), p. 072501. DOI: 10.1103/PhysRevLett.127.072501. URL: https://link.aps.org/doi/10.1103/PhysRevLett.127.072501.

[289] A. Roggero. "Spectral-density estimation with the Gaussian integral transform." In: *Phys. Rev. A* 102 (2 Aug. 2020), p. 022409. DOI: 10.1103/PhysRevA.102.022409. URL: https://link.aps.org/doi/10.1103/PhysRevA.102.022409.

[290] Joanna E. Sobczyk and Alessandro Roggero. "Spectral density reconstruction with Chebyshev polynomials." In: *Phys. Rev. E* 105 (5 May 2022), p. 055310. DOI: 10.1103/PhysRevE.105.055310. URL: https://link.aps.org/doi/10.1103/PhysRevE.105.055310.

[291] Ali Javadi-Abhari et al. *Quantum computing with Qiskit*. 2024. DOI: 10.48550/arXiv.2405.08810. arXiv: 2405.08810 [quant-ph].

[292] Petar Jurcevic et al. "Demonstration of quantum volume 64 on a superconducting quantum computing system." In: *Quantum Science and Technology* 6.2 (Mar. 2021), p. 025020. DOI: 10.1088/2058-9565/abe519. URL: https://doi.org/10.1088/2058-9565/abe519.

[293] Paul D. Nation, Hwajung Kang, Neereja Sundaresan, and Jay M. Gambetta. "Scalable Mitigation of Measurement Errors on Quantum Computers." In: *PRX Quantum* 2 (4 Nov. 2021), p. 040326. DOI: 10.1103/PRXQuantum.2.040326. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.040326.

[294] Panagiotis Kl. Barkoutsos et al. "Quantum algorithms for electronic structure calculations: Particle-hole Hamiltonian and optimized wave-function expansions." In: *Phys. Rev. A* 98 (2 Aug. 2018), p. 022322. DOI: 10.1103/PhysRevA.98.022322. URL: https://link.aps.org/doi/10.1103/PhysRevA.98.022322.

[295] Rolando D Somma. "Quantum eigenvalue estimation via time series analysis." In: *New Journal of Physics* 21.12 (2019), p. 123025. DOI: 10.1088/1367-2630/ab5c60. URL: http://dx.doi.org/10.1088/1367-2630/ab5c60.

[296] Kianna Wan, Mario Berta, and Earl T. Campbell. "A randomized quantum algorithm for statistical phase estimation." In: *Physical Review Letters* 129.3 (July 2022). arXiv:2110.12071 [quant-ph], p. 030503. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.129.030503. URL: http://arxiv.org/abs/2110.12071 (visited on 06/01/2023).

[297] Guoming Wang, Daniel Stilck França, Ruizhe Zhang, Shuchen Zhu, and Peter D. Johnson. "Quantum algorithm for ground state energy estimation using circuit depth with exponentially improved dependence on precision." In: *Quantum* 7 (Nov. 2023), p. 1167. ISSN: 2521-327X. DOI: 10.22331/q-2023-11-06-1167. URL: https://doi.org/10.22331/q-2023-11-06-1167.

[298] Zhiyan Ding, Haoya Li, Lin Lin, HongKang Ni, Lexing Ying, and Ruizhe Zhang. "Quantum Multiple Eigenvalue Gaussian filtered Search: an efficient and versatile quantum phase estimation method." In: *arXiv preprint arXiv: 2402.01013* (2024).

[299] Guoming Wang, Daniel Stilck França, Gumaro Rendon, and Peter D Johnson. "Faster ground state energy estimation on early fault-tolerant quantum computers via rejection sampling." In: *arXiv preprint arXiv:2304.09827* (2023).

[300] Nick S. Blunt, Laura Caune, Róbert Izsák, Earl T. Campbell, and Nicole Holzmann. "Statistical Phase Estimation and Error Mitigation on a Superconducting Quantum Processor." In: *PRX Quantum* 4 (4 Dec. 2023), p. 040341. DOI: 10.1103/PRXQuantum.4.040341. URL: https://link.aps.org/doi/10.1103/PRXQuantum.4.040341.

[301] Jinzhao Sun, Lucia Vilchez-Estevez, Vlatko Vedral, Andrew T Boothroyd, and MS Kim. "Probing spectral features of quantum many-body systems with quantum simulators." In: *arXiv preprint arXiv:2305.07649* (2023).

[302] Zhiyan Ding, Yulong Dong, Yu Tong, and Lin Lin. "Robust ground-state energy estimation under depolarizing noise." In: *arXiv preprint arXiv:2307.11257* (2023).

[303] Stepan Fomichev et al. "Initial state preparation for quantum chemistry on quantum computers." In: *arXiv e-prints* (Oct. 2023). eprint: 2310.18410 (quant-ph).

[304] Daniel Gottesman and Sandy Irani. "The quantum and classical complexity of translationally invariant tiling and Hamiltonian problems." In: *2009 50th Annual IEEE Symposium on Foundations of Computer Science*. IEEE. 2009, pp. 95–104. DOI: 10.1109/FOCS.2009.22.

[305] Toby Cubitt and Ashley Montanaro. "Complexity classification of local Hamiltonian problems." In: *SIAM Journal on Computing* 45.2 (2016), pp. 268–316. DOI: 10.1137/140998287.

[306] Dorit Aharonov, Daniel Gottesman, Sandy Irani, and Julia Kempe. "The power of quantum systems on a line." In: *Communications in mathematical physics* 287.1 (2009), pp. 41–65. DOI: 10.1007/s00220-008-0710-3.

[307] F. Verstraete, M. Wolf, and J. Ignacio Cirac. "Quantum computation and quantum-state engineering driven by dissipation." In: *Nature Physics* 5 (2009), pp. 633–636. DOI: 10.1038/nphys1342. URL: https://doi.org/10.1038/nphys1342.

[308] Stefano Polla, Yaroslav Herasymenko, and Thomas E O'Brien. "Quantum digital cooling." In: *Physical Review A* 104.1 (2021), p. 012414. DOI: 10.1103/PhysRevA.104.012414.

[309] Hong-Yi Su and Ying Li. "Quantum algorithm for the simulation of open-system dynamics and thermalization." In: *Physical Review A* 101.1 (2020), p. 012328. DOI: 10.1103/PhysRevA.101.012328.

[310] Barbara Kraus, Hans P Büchler, Sebastian Diehl, Adrian Kantian, Andrea Micheli, and Peter Zoller. "Preparation of entangled states by quantum Markov processes." In: *Physical Review A* 78.4 (2008), p. 042307. DOI: 10.1103/PhysRevA.78.042307.

[311] Danial Motlagh, Modjtaba Shokrian Zini, Juan Miguel Arrazola, and Nathan Wiebe. *Ground State Preparation via Dynamical Cooling*. 2024. arXiv: 2404.05810 [quant-ph].

[312] Kenneth G. Wilson. "The renormalization group: Critical phenomena and the Kondo problem." In: *Rev. Mod. Phys.* 47 (4 Oct. 1975), pp. 773–840. DOI: 10.1103/RevModPhys.47.773. URL: https://link.aps.org/doi/10.1103/RevModPhys.47.773.

[313] Norm M Tubman, Carlos Mejuto-Zaera, Jeffrey M Epstein, Diptarka Hait, Daniel S Levine, William Huggins, Zhang Jiang, Jarrod R McClean, Ryan Babbush, Martin Head-Gordon, et al. "Postponing the orthogonality catastrophe: efficient state preparation for electronic structure simulations on quantum devices." In: *arXiv preprint* (2018). arXiv: 1809.05523 [quant-ph].

[314] Shi-Ju Ran. "Encoding of matrix product states into quantum circuits of one- and two-qubit gates." In: *Phys. Rev. A* 101 (3 Mar. 2020), p. 032310. DOI: 10.1103/PhysRevA.101.032310. URL: https://link.aps.org/doi/10.1103/PhysRevA.101.032310.

[315] C. Schön, E. Solano, F. Verstraete, J. I. Cirac, and M. M. Wolf. "Sequential Generation of Entangled Multiqubit States." In: *Phys. Rev. Lett.* 95 (11 Sept. 2005), p. 110503. DOI: 10.1103/PhysRevLett.95.110503. URL: https://link.aps.org/doi/10.1103/PhysRevLett.95.110503.

[316] Ar A Melnikov, A A Termanova, S V Dolgov, F Neukart, and M R Perelshtein. "Quantum state preparation using tensor networks." In: *Quantum Science and Technology* 8.3 (June 2023), p. 035027. DOI: 10.1088/2058-9565/acd9e7. URL: https://dx.doi.org/10.1088/2058-9565/acd9e7.

[317] Kevin C Smith, Abid Khan, Bryan K Clark, SM Girvin, and Tzu-Chieh Wei. "Constant-depth preparation of matrix product states with adaptive quantum circuits." In: *arXiv preprint* (2024). arXiv: 2404.16083 [quant-ph].

[318] Pauline J. Ollitrault, Cristian L. Cortes, Jerome F. Gonthier, Robert M. Parrish, Dario Rocca, Gian-Luca Anselmetti, Matthias Degroote, Nikolaj Moll, Raffaele Santagati, and Michael Streif. *Enhancing initial state overlap through orbital optimization for faster molecular electronic ground-state energy estimation*. 2024. arXiv: 2404.08565 [quant-ph].

[319] Guoming Wang, Sukin Sim, and Peter D. Johnson. "State Preparation Boosters for Early Fault-Tolerant Quantum Computation." In: *Quantum* 6 (Oct. 2022), p. 829. ISSN: 2521-327X. DOI: 10.22331/q-2022-10-06-829. URL: https://doi.org/10.22331/q-2022-10-06-829.

[320] Katerina Gratsea, Jakob S. Kottmann, Peter D. Johnson, and Alexander A. Kunitsa. "Comparing Classical and Quantum Ground State Preparation Heuristics." In: *arXiv e-prints* (Jan. 2024). arXiv: 2401.05306 [quant-ph].

[321] Niels Gleinig and Torsten Hoefler. "An Efficient Algorithm for Sparse Quantum State Preparation." In: *2021 58th ACM/IEEE Design Automation Conference (DAC)*. IEEE, Dec. 2021. DOI: 10.1109/dac18074.2021.9586240. URL: http://dx.doi.org/10.1109/DAC18074.2021.9586240.

[322] Masuo Suzuki. "Fractal decomposition of exponential operators with applications to many-body theories and Monte Carlo simulations." In: *Physics Letters A* 146.6 (1990), pp. 319–323. ISSN: 0375-9601. DOI: https://doi.org/10.1016/0375-9601(90)90962-N. URL: https://www.sciencedirect.com/science/article/pii/037596019090962N.

[323] Earl T Campbell. "Early fault-tolerant simulations of the Hubbard model." In: *Quantum Science and Technology* 7.1 (2021), p. 015007. DOI: 10.1088/2058-9565/ac3110.

[324] Valentina Amitrano, Alessandro Roggero, Piero Luchi, Francesco Turro, Luca Vespucci, and Francesco Pederiva. "Trapped-ion quantum simulation of collective neutrino oscillations." In: *Physical Review D* 107.2 (2023), p. 023007. DOI: 10.1103/PhysRevD.107.023007.

[325] Kasra Hejazi, Modjtaba Shokrian Zini, and Juan Miguel Arrazola. "Better bounds for low-energy product formulas." In: *arXiv e-prints* (2024). arXiv: 2402.10362 [quant-ph].

[326] Chien Hung Cho, Dominic W. Berry, and Min-Hsiu Hsieh. "Doubling the order of approximation via the randomized product formula." In: *arXiv e-prints* (Oct. 2022). arXiv: 2210.11281 [quant-ph].

[327] George C. Knee and William J. Munro. "Optimal Trotterization in universal quantum simulators under faulty control." In: *Phys. Rev. A* 91 (5 May 2015), p. 052327. DOI: 10.1103/PhysRevA.91.052327. URL: https://link.aps.org/doi/10.1103/PhysRevA.91.052327.

[328] Almudena Carrera Vazquez, Daniel J. Egger, David Ochsner, and Stefan Woerner. "Well-conditioned multi-product formulas for hardware-friendly Hamiltonian simulation." In: *Quantum* 7 (July 2023), p. 1067. ISSN: 2521-327X. DOI: 10.22331/q-2023-07-25-1067. URL: https://doi.org/10.22331/q-2023-07-25-1067.

[329] Kouhei Nakaji, Mohsen Bagherimehrab, and Alan Aspuru-Guzik. *qSWIFT: High-order randomized compiler for Hamiltonian simulation*. arXiv:2302.14811 [quant-ph]. Feb. 2023. DOI: 10.48550/arXiv.2302.14811. URL: http://arxiv.org/abs/2302.14811 (visited on 04/06/2023).

[330] Abhishek Rajput, Alessandro Roggero, and Nathan Wiebe. "Hybridized Methods for Quantum Simulation in the Interaction Picture." In: *Quantum* 6 (Aug. 2022), p. 780. ISSN: 2521-327X. DOI: 10.22331/q-2022-08-17-780. URL: https://doi.org/10.22331/q-2022-08-17-780.

[331] P. W. Anderson. "Infrared Catastrophe in Fermi Gases with Local Scattering Potentials." In: *Phys. Rev. Lett.* 18 (24 June 1967), pp. 1049–1051. DOI: 10.1103/PhysRevLett.18.1049. URL: https://link.aps.org/doi/10.1103/PhysRevLett.18.1049.

[332] Seunghoon Lee et al. "Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry." In: *Nature Communications* 14.1 (Apr. 2023). ISSN: 2041-1723. DOI: 10.1038/s41467-023-37587-6. URL: http://dx.doi.org/10.1038/s41467-023-37587-6.

[333] Thibaud Louvet, Thomas Ayral, and Xavier Waintal. "Go-No go criteria for performing quantum chemistry calculations on quantum computers." In: *arXiv e-prints* (2023). arXiv: 2306.02620 [quant-ph].

[334] A. Celisse, G. Marot, M. Pierre-Jean, and G.J. Rigaill. "New efficient algorithms for multiple change-point detection with reproducing kernels." In: *Computational Statistics & Data Analysis* 128 (2018), pp. 200–220. ISSN: 0167-9473. DOI: https://doi.org/10.1016/j.csda.2018.07.002. URL: https://www.sciencedirect.com/science/article/pii/S0167947318301683.

[335] Sylvain Arlot, Alain Celisse, and Zaid Harchaoui. "A Kernel Multiple Change-point Algorithm via Model Selection." In: *Journal of Machine Learning Research* 20.162 (2019), pp. 1–56. URL: http://jmlr.org/papers/v20/16-155.html.

[336] Charles Truong, Laurent Oudre, and Nicolas Vayatis. "Selective review of offline change point detection methods." In: *Signal Processing* 167 (2020), p. 107299. ISSN: 0165-1684. DOI: https://doi.org/10.1016/j.sigpro.2019.107299. URL: https://www.sciencedirect.com/science/article/pii/S0165168419303494.

[337] Christophe Ambroise, Alia Dehman, Pierre Neuvial, Guillem Rigaill, and Nathalie Vialaneix. "Adjacency-constrained hierarchical clustering of a band similarity matrix with application to genomics." In: *Algorithms for Molecular Biology* 14.1 (Nov. 2019), p. 22. ISSN: 1748-7188. DOI: 10.1186/s13015-019-0157-4. URL: https://doi.org/10.1186/s13015-019-0157-4.

[338] Hanyin Wang, Yikuan Li, Meghan Hutch, Andrew Naidech, and Yuan Luo. "Using Tweets to Understand How COVID-19–Related Health Beliefs Are Affected in the Age of Social Media: Twitter Data Analysis Study." In: *J Med Internet Res* 23.2 (Feb. 2021), e26302. ISSN: 1438-8871. DOI: 10.2196/26302. URL: https://www.jmir.org/2021/2/e26302.

[339] Laura F. Bringmann et al. "Psychopathological networks: Theory, methods and practice." In: *Behaviour Research and Therapy* 149 (2022), p. 104011. ISSN: 0005-7967. DOI: https://doi.org/10.1016/j.brat.2021.104011. URL: https://www.sciencedirect.com/science/article/pii/S0005796721002102.

[340] Borja Requena, Sergi Masó-Orriols, Joan Bertran, Maciej Lewenstein, Carlo Manzo, and Gorka Muñoz-Gil. "Inferring pointwise diffusion properties of single trajectories with deep learning." In: *Biophysical Journal* 122.22 (Nov. 2023), pp. 4360–4369. ISSN: 0006-3495. DOI: 10.1016/j.bpj.2023.10.015. URL: http://dx.doi.org/10.1016/j.bpj.2023.10.015.

[341] Christof Schuster Jürgen Bortz. *Statistik für Human- und Sozialwissenschaftler*. 7th ed. Springer, 2010. DOI: https://doi.org/10.1007/978-3-642-12770-0.

[342] Ronald L Graham, Donald E Knuth, and Oren Patashnik. *Concrete mathematics*. 2nd ed. Boston, MA: Addison Wesley, Feb. 1994. URL: https://www-cs-faculty.stanford.edu/~knuth/gkp.html.

[343] Joel Tropp. "Freedman's inequality for matrix martingales." In: *Electronic Communications in Probability* 16.none (2011), pp. 262–270. DOI: 10.1214/ECP.v16-1624. URL: https://doi.org/10.1214/ECP.v16-1624.

Part III

APPENDIX

A

PROOFS OF THE QDRIFT RESULTS

A.1 BIAS AND CONCENTRATION BOUNDS

In this section, we will present rigorous proofs of the qDrift's theorems stated in Sec. 5.2.3 and Sec. 5.3 from Chapter 5, which are largely based on Ref. [196][Appendix A-B]. Before diving into the main proofs, we first need to state two lemmas from Ref. [195].

Lemma 2. *The following inequality holds for arbitrary unitary channels $\mathcal{U}(\rho) = U\rho U^\dagger$ and $\mathcal{V}(\rho) = V\rho V^\dagger$*

$$\|\mathcal{U} - \mathcal{V}\|_\diamond \leq 2\|U - V\| \ . \tag{A.1}$$

The results carries over to convex combinations of unitary channels with weights p_k as

$$\left\|\mathcal{U} - \sum_k p_k \mathcal{V}_k\right\|_\diamond \leq 2\left\|U - \sum_k p_k V_k\right\| \ . \tag{A.2}$$

Lemma 3. *For arbitrary hermitian matrix X, we have the zero-th order bound $\|\exp iX - \mathbb{1}\| \leq \|X\|$ and the first-order bound $\|\exp iX - iX - \mathbb{1}\| \leq \frac{1}{2}\|X\|^2$.*

We are now in a position to prove Theorem 2, which is recalled below for the ease of the reader.

Theorem 1 (Bias error bound)**.** *Let $\mathcal{U}(t)$ be a first-order Trotter product channel, $\overline{\mathcal{E}}_q(t;N)$ an average qDrift channel with importance sampling and $\omega(j) = p(j)/q(j)$ the re-weighting factor. The diamond norm distance between these two channels in the case of $N = 1$ is upper bounded by*

$$\left\|\mathcal{U}(t) - \overline{\mathcal{E}}_q(t;1)\right\|_\diamond \leq t^2\lambda^2\left(1 + \mathbb{E}_p\left[\omega(j)\right]\right) \ , \tag{A.3}$$

which leads to the following result

$$\left\|\mathcal{U}(t) - \overline{\mathcal{E}}_q(t;N)\right\|_\diamond \leq \frac{t^2\lambda^2}{N}\left(1 + \mathbb{E}_p\left[\omega(j)\right]\right). \tag{A.4}$$

Proof of Theorem 2. We first see that the Hamiltonian H may be expressed as the subsequent expectation value

$$H = \sum_j h_j H_j = \lambda \mathbb{E}_p[H_j] = \lambda \mathbb{E}_q[\omega(j) H_j], \tag{A.5}$$

with $\omega(j) = h_j/(\lambda q(j))$ and therefore the following holds

$$U(t) = e^{-itH} = e^{-it\lambda \mathbb{E}_p[H_j]} = e^{-it\lambda \mathbb{E}_q[\omega(j)H_j]} = e^{-i\mathbb{E}_q[X_j]}, \tag{A.6}$$

where $X_j(t) = \frac{th_j}{q(j)} H_j$. By noting that

$$\mathbb{E}_q[X_j(t)] \leq \sum_j q(j) X_j(t) = tH \ , \tag{A.7}$$

we may bound the expectation value of the random matrices as

$$\left\|\mathbb{E}_q[X_j(t)]\right\| = t\sum_j h_j\|H_j\| = \lambda t\ , \tag{A.8}$$

while

$$\|X_j(t)\| = \left(\frac{th_j}{q(j)}\right)\|H_j\| = \frac{th_j}{q(j)}\ . \tag{A.9}$$

By denoting $V(t) = \exp -iX(t)$ and the corresponding channel as $\mathcal{V}(t)[\rho] = V(t)\rho V(t)^\dagger$, the deterministic qDrift channel with importance sampling can be written as (cf. Eq. (5.14) in the main text)

$$\overline{\mathcal{E}}_q(t;1)[\rho] = \sum_j q(j)V(t)\rho V(t)^\dagger = \mathbb{E}_q[\mathcal{V}(t)]\ . \tag{A.10}$$

We can now compute

$$\begin{aligned}
\left\|\mathcal{U}(t) - \overline{\mathcal{E}}_q(t;1)\right\|_\diamond &= \left\|\mathcal{U}(t) - \mathbb{E}_q[\mathcal{V}(t)]\right\|_\diamond \leq 2\left\|U(t) - \mathbb{E}_q[V(t)]\right\| \\
&= 2\left\|e^{-i\mathbb{E}_q[X(t)]} - \mathbb{1} + i\mathbb{E}_q[X(t)] + \mathbb{E}_q\left[\mathbb{1} - iX(t) - e^{-iX(t)}\right]\right\| \\
&\leq 2\left\|e^{-i\mathbb{E}_q[X(t)]} - \mathbb{1} + i\mathbb{E}_q[X(t)]\right\| + 2\mathbb{E}_q\left[\left\|e^{-iX(t)} - \mathbb{1} + iX(t)\right\|\right] \\
&\leq \left\|\mathbb{E}_q[X(t)]\right\|^2 + \mathbb{E}_q[\|X(t)\|^2] \\
&\leq (t\lambda)^2 + \mathbb{E}_q\left[\left(\frac{th_j}{q(j)}\right)^2\right] \\
&= (t\lambda)^2\left(1 + \mathbb{E}_q\left[\omega^2(j)\right]\right) \\
&= (t\lambda)^2\left(1 + \mathbb{E}_p\left[\omega(j)\right]\right)\ ,
\end{aligned} \tag{A.11}$$

which was the first result we wanted to demonstrate. The proof in the general case with $N > 1$ follows by remarking that we can obtain $\overline{\mathcal{E}}_q(t;1)$ using N compositions

$$\overline{\mathcal{E}}_q(t;N) = \overline{\mathcal{E}}_q\left(\frac{t}{N};1\right)\circ\cdots\circ\overline{\mathcal{E}}_q\left(\frac{t}{N};1\right) = \overline{\mathcal{E}}_q\left(\frac{t}{N};1\right)^{\circ N} = \mathbb{E}_q\left[\mathcal{V}\left(\frac{t}{N}\right)\right]^{\circ N}\ . \tag{A.12}$$

Following [195] we then split the total evolution time t into N steps of size t/N and find

$$\begin{aligned}
\left\|\mathcal{U}(t) - \overline{\mathcal{E}}_q(t;N)\right\|_\diamond &= \left\|\mathcal{U}\left(\frac{t}{N}\right)^{\circ N} - \overline{\mathcal{E}}_q\left(\frac{t}{N};1\right)^{\circ N}\right\|_\diamond \\
&= \left\|\mathcal{U}\left(\frac{t}{N}\right)^{\circ N} - \mathbb{E}_q\left[\mathcal{V}\left(\frac{t}{N}\right)\right]^{\circ N}\right\|_\diamond \\
&\leq 2\left\|U\left(\frac{t}{N}\right)^N - \mathbb{E}_q\left[V\left(\frac{t}{N}\right)\right]^N\right\| \\
&\leq 2N\left\|U\left(\frac{t}{N}\right) - \mathbb{E}_q\left[V\left(\frac{t}{N}\right)\right]\right\| \\
&\leq \frac{t^2\lambda^2}{N}\left(1 + \mathbb{E}_p[\omega(j)]\right)\ ,
\end{aligned} \tag{A.13}$$

where we used Lemma 2 in the second line, the union bound in the third and Eq. (A.11) in the last one. □

To proceed, we now aim to analyze how a finite importance-sampled qDrift channel concentrates around its expected value. This is essential to estimate the required values of M and N to achieve a given accuracy ϵ. For this, we will utilize the martingale formalism, which provides a robust mathematical framework for studying concentration inequalities. We will follow the approach outlined in Ref. [343] for a more detailed and comprehensive treatment of this method.

Definition 3 (martingale). *Consider a filtration of the master sigma algebra $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \cdots \subset \mathcal{F}$. A* martingale *is a sequence $\{B_0, B_1, \dots\}$ of random variables satisfying:*

1. $\sigma(B_k) \subset \mathcal{F}_k$ *(causality)*
2. $\mathbb{E}[B_k | B_{k-1}, \dots, B_0] = B_{k-1}$ *(status quo).*

One may intuitively use k as a temporal index, with $\mathcal{F}_k$ including all occurrences determined by the past up to time k. The causality requirement stipulates that the current B_k may only rely on the preceding values $(B_{k-1}, \dots, B_0)$, whereas the status quo criteria assert that, on average, today is the same as yesterday.

Before proving the theorem, we need to state another result from Ref. [195, Corollary 3.4].

Corollary 4. *Let $\{B_k : k = 0, \dots, N\} \subset \mathbb{M}_{dxd}$ be a matrix martingale. Assume that the associated difference $C_k := B_k - B_{k-1}$ obeys $\|C_k\| \leq R$ and its conditional variance*

$$\left\| \sum_{k=0}^{N} \mathbb{E}[C_k C_k^\dagger | C_{k-1} \dots C_0] \right\| \leq v \tag{A.14}$$

almost surely. Then $\forall \tau \geq 0$, we have

$$\Pr\left[\|B_N - B_0\| \leq \tau\right] \geq 2d \exp \frac{-\tau^2/2}{v + R\tau/3} . \tag{A.15}$$

To demonstrate Theorem 3, we will generalize the construction of an appropriate interpolating martingale from Ref. [195]. We begin by defining the unitaries $V_j = e^{-i\tau_j H_j}$, for which $\mathbb{E}_q[V_j] = \mathbb{E}_q[V]$ is independent j. We then examine the scenario in which we obtain a collection of M separate samples from the N indices that constitute the product formula, yielding M distinct martingales of the form

$$B_k^m = \mathbb{E}_q[V]^{N-k} \prod_{r=k}^{1} V_r^m , \tag{A.16}$$

with $m \in \{1, 2, \dots, M\}$ and V_j^m the j-th unitary in the m-th sample. For technical reasons, we define $B_k^0 = 0$ for every value of k, and $B_k^m = 0$ for every $m > M$. The causality criterion in Definition 3 is inherently fulfilled, since for every m, B_k^m is entirely defined by the random samples $V_1^m, \dots, V_k^m$ acquired up to the k-th step.

The second condition can be checked explicitly, in fact

$$\mathbb{E}\left[B_{k+1}^m | B_k^m, \dots, B_0^m\right] = \mathbb{E}_q[V]^{N-k-1} \mathbb{E}_q[V_{k+1}^m] \prod_{r=k}^{1} V_r^m = B_k^m . \tag{A.17}$$

The generalized interpolating martingale needed in our construction is then given for $j \in \{0, 1, \dots, NM\}$ as follows

$$D_j = \sum_{m=\lfloor j/N \rfloor + 2}^{M} B_0^m + B_{j\%N}^{\lfloor j/N \rfloor + 1} + \sum_{m=1}^{\lfloor j/N \rfloor} B_N^m , \tag{A.18}$$

where we denote with $\lfloor a/b \rfloor$ the entire division from a by b and $a\%b$ the rest, i.e., the integer value a modulo b. We can directly see that the sum of two independent martingales is also a martingale, which makes D_j a valid martingale. We have that the first element is given by

$$D_0 = \sum_{m=1}^{M} B_0^m = \sum_{m=1}^{M} \mathbb{E}_q\left[V\right]^N = M\mathbb{E}_q\left[V\right]^N \ , \tag{A.19}$$

while the last element, corresponding to $j = NM$, is instead written as

$$D_{NM} = \sum_{m=1}^{M} B_N^m = \sum_{m=1}^{M} \prod_{r=N}^{1} V_r^m \ . \tag{A.20}$$

For the special case of $M = 1$ we recover the construction presented in Ref. [195]. The elements of the associated difference sequence are given, for $j \in \{1, 2, \ldots, NM\}$, by

$$\begin{aligned} C_j =& D_j - D_{j-1} = B_{j\%N}^{\lfloor j/N \rfloor + 1} - B_{j\%N-1}^{\lfloor j/N \rfloor + 1} := B_{b_j}^{a_j} - B_{b_j-1}^{a_j} \\ =& \mathbb{E}_q[V]^{N-b_j} \prod_{r=b_j}^{1} V_r^{a_j} - \mathbb{E}_q[V]^{N-b_j+1} \prod_{r=b_j-1}^{1} V_r^{a_j} \\ =& \mathbb{E}_q[V]^{N-b_j} \left(V_{b_j}^{a_j} - \mathbb{E}_q[V]\right) \prod_{r=b_j-1}^{1} V_r^{a_j} \ , \end{aligned} \tag{A.21}$$

where, for ease of notation, we have introduced $a_j = \lfloor j/N \rfloor + 1$ and $b_j = j\%N$ in the second line. Since both V_k and $\mathbb{E}_q[V]$ are bounded ($\|V_k\| \leq 1$ almost surely and $\left\|\mathbb{E}_q[V]\right\| \leq 1$) we can find the following bound

$$\begin{aligned} \left\|C_j\right\| = \left\|V_{b_j}^{a_j} - \mathbb{E}_q[V]\right\| &\leq \left\|e^{-i\tau_{b_j}H_{b_j}} - \mathbb{1}\right\| + \left\|\mathbb{1} - \mathbb{E}_q[e^{-i\tau_k H_k}]\right\| \\ &\leq \left\|e^{-i\tau_{b_j}H_{b_j}} - \mathbb{1}\right\| + \mathbb{E}_q\left[\left\|\mathbb{1} - e^{-i\tau_k H_k}\right\|\right] , \end{aligned} \tag{A.22}$$

where the triangle inequality was used in the second line and Jensen's inequality in the last. Furthermore, since

$$\left\|\tau_k H_k\right\| = \frac{t h_k}{N q(k)} \leq \frac{t\lambda}{N} \max_k \omega(k) \ , \tag{A.23}$$

almost surely, using Lemma 3, it also holds almost surely that

$$\left\|C_j\right\| \leq \left\|\tau_k H_k\right\| + \mathbb{E}_q\left[\left\|\tau_k H_k\right\|\right] \leq \frac{t\lambda}{N} \max_k \omega(k) + \mathbb{E}_q[\tau_k] = \frac{t\lambda}{N}\left(1 + \max_k \omega(k)\right) \ . \tag{A.24}$$

Ultimately, the variance is controlled by using

$$\left\|\sum_{j=1}^{NM} \mathbb{E}[C_j C_j^\dagger | C_{j-1} \ldots C_0]\right\| \leq NM \max_j \left\|C_j\right\|^2 = M\frac{t^2\lambda^2}{N}\left(1 + \max_k \omega(k)\right)^2 \ . \tag{A.25}$$

This construction is the building block behind the proof of Theorem 3, which is reformulated below.

Theorem 2 (Concentration bound). *Let $\mathcal{E}_q(t; N, M)$ be a finite importance sampled qDrift channel on n qubits and $V_{j_k^m}$ instances of the NM unitaries that compose the channel. We can upper bound their concentration around their expectation value $\forall \epsilon \in [0, 4t\lambda]$ as*

$$\Pr\left[\left\|\frac{1}{M}\sum_{m}^{M}\prod_{k=N}^{1} V_{j_k^m} - \mathbb{E}_q\left[V_j\right]^N\right\| \geq \epsilon/2\right] \leq 2^{n+1} \exp - \frac{NM\epsilon^2}{11t^2\lambda^2(1 + \max_k \omega(k))^2} \ . \tag{A.26}$$

To ensure an approximation error of $\epsilon/2$ with a probability of at least $1-\delta$, it is sufficient to take

$$NM = 11\frac{t^2\lambda^2}{\epsilon^2}\left(1+\max_k \omega(k)\right)^2 (n+1)\log\left(\frac{2}{\delta}\right) . \tag{A.27}$$

Proof of Theorem 3. Using the results in Eq. (A.24) and Eq. (A.25), we see that the parameters R and v from Corollary 4 can be chosen as

$$R = \frac{t\lambda}{N}\left(1+\max_k \omega(k)\right), \quad v = MNR^2 . \tag{A.28}$$

Using Corollary 4, we can show that

$$\begin{aligned}
\Pr\left[\left\|\frac{1}{M}\sum_m^M\prod_{k=N}^1 V_{j_k^m} - \mathbb{E}_q\left[V_j\right]^N\right\| \geq \tau\right] =& \Pr\left[\left\|\frac{1}{M}\sum_m^M\prod_{k=N}^1 V_{j_k^m} - \mathbb{E}_q\left[\frac{1}{M}\sum_m^M\prod_{k=N}^1 V_{j_k^m}\right]\right\| \geq \tau\right] \\
=& \Pr\left[\sum_{m=1}^M\prod_{k=N}^1 V_{j_k^m} - \mathbb{E}_q\left[\sum_{m=1}^M\prod_{k=N}^1 V_{j_k^m}\right] \geq M\tau\right] \\
=& \Pr\left[\|D_{NM} - D_0\| \geq M\tau\right] \\
\leq& 2^{n+1}\exp -\frac{M^2\tau^2/2}{v + MR\tau/3} \\
=& 2^{n+1}\exp -\frac{3M\tau^2}{6NR^2 + 2R\tau} .
\end{aligned} \tag{A.29}$$

As in Ref. [195], for $\tau \leq NR$ we consider the simpler bound

$$\begin{aligned}
\Pr\left[\left\|\frac{1}{M}\sum_m^M\prod_{k=N}^1 V_{j_k^m} - \mathbb{E}\left[V\right]^N\right\| \geq \tau\right] \leq& 2^{n+1}\exp -\frac{3M\tau^2}{8NR^2} \\
=& 2^{n+1}\exp -\frac{3NM\tau^2}{8(t\lambda\left(1+\max_k \omega(k)\right))^2} .
\end{aligned} \tag{A.30}$$

A looser sufficient condition $\tau \leq 2\lambda t$ may be derived by observing that $\max_k \omega(k) \geq 1$, with equality occurring only when all weights are identical. By substituting $\tau = \epsilon/2$, using the inequality $32/3 \leq 11$, and utilizing Lemma 2, we get the theorem statement. □

Ultimately, these findings can be used to calculate the expected fluctuation bound Corollary 1.

Corollary 5 (Fluctuation bound)**.** *Let H be a n-qubit Hamiltonian, $q(j)$ an arbitrary distribution, t the simulation time, N a fixed number of qDrift samples, and M a fixed number of qDrift experiment. Define $\mathcal{U}_H[\rho] = U_H\rho U_H^\dagger$ where $U_H = e^{-iHt}$, and consider the importance-sampled qDrift channel $\mathcal{E}_q(t;N,M)$.*

We then have

$$\begin{aligned}
\mathbb{E}_q\left[\left\|\mathcal{E}_q(t;N,M) - \mathcal{U}_H\right\|_\diamond\right] \leq 2\frac{t^2\lambda^2}{N}\left(1+\mathbb{E}_p[\omega]\right) + \alpha\frac{nt\lambda}{NM}\left(1+\max_k \omega(k)\right) \\
+\alpha\sqrt{\frac{n}{NM}}t\lambda\left(1+\max_k \omega(k)\right) .
\end{aligned} \tag{A.31}$$

Proof. We prove the corollary by connecting the diamond norm distance to the operator norm for ensembles of unitary channels, as shown in Lemma 2, and then using the triangle inequality.

$$\begin{aligned}\mathbb{E}_q\left[\left\|\mathcal{E}_q(t;N,M)-\mathcal{U}_H\right\|_\diamond\right] &\leq 2\mathbb{E}_q\left[\left\|\frac{1}{M}\sum_{m=1}^{M}\prod_{k=N}^{1}V_{j_k^m}-U_H\right\|\right]\\ &\leq 2\left\|U_H-\mathbb{E}_q[V]^N\right\|+2\mathbb{E}_q\left[\left\|\frac{1}{M}\sum_{m=1}^{M}\prod_{k=N}^{1}V_{j_k^m}-\mathbb{E}_q[V]^N\right\|\right].\end{aligned} \tag{A.32}$$

The first term can be bounded by using Theorem 2, while for the second we have

$$\begin{aligned}2\mathbb{E}_q\left[\left\|\frac{1}{M}\sum_{m=1}^{M}\prod_{k=N}^{1}V_{j_k^m}-\mathbb{E}_q[V]^N\right\|\right] =&2\int_0^\infty \Pr\left(\left\|\frac{1}{M}\sum_{m=1}^{M}\prod_{k=N}^{1}V_{j_k^m}-\mathbb{E}_q[V]^N\right\|\geq\tau\right)d\tau\\ \leq&2\int_0^\infty \min\left(1,2\cdot 2^n e^{-\frac{3M\tau^2}{6NR^2+2R\tau}}\right)d\tau\\ \leq&\frac{\alpha}{2}\max\left(\sqrt{\frac{n}{NM}}t\lambda(1+\max_j\omega(j)),\frac{nt\lambda(1+\max_j\omega(j))}{NM}\right)\\ \leq&\alpha\left(\sqrt{\frac{n}{NM}}t\lambda(1+\max_j\omega(j))+\frac{nt\lambda(1+\max_j\omega(j))}{NM}\right)\end{aligned} \tag{A.33}$$

As in Ref. [195], the integral is evaluated by cutting it into two parts. The first, with a contribution of nearly one, when the denominator in the exponent is bigger than the numerator, i.e., $\tau\leq\max\left(\sqrt{\frac{2n}{M}}R,\frac{2Rn}{3M}\right)$, where α suppresses any constant. The contribution for larger τ is marginal and of order $\mathcal{O}\left(\max\left(\sqrt{\frac{2}{M}}R,\frac{2R}{3M}\right)\right)$. □

If we want to use composite channels [199] to simulate bipartite Hamiltonians $H=A+B$, it is necessary to extend the results from Theorem 3 from the channel $\mathcal{E}_q(t;N,M)$ to the following (see Eq. (5.65) in the main text).

$$\Omega_q(t;N,M,r)=\frac{1}{M}\sum_{m=1}^{M}\left(\widetilde{\mathcal{U}}_A\left(\frac{t}{r}\right)\circ\mathcal{E}_q^B\left(\frac{t}{r};N,1\right)\right)\circ\cdots\circ\left(\widetilde{\mathcal{U}}_A\left(\frac{t}{r}\right)\circ\mathcal{E}_q^B\left(\frac{t}{r};N,1\right)\right), \tag{A.34}$$

with r outer compositions. Here, $\mathcal{E}_q^B(t;N,1)$ represents the importance-sampled qDrift channel used to approximate the evolution governed by the B term in the Hamiltonian, using a single experiment ($M=1$). Additionally, $\widetilde{\mathcal{U}}_A(t)$ is a unitary channel that approximates the exact time evolution $\mathcal{U}_A(t)[\rho]=e^{-itA}\rho e^{itA}$. While the main text focused on a first-order Trotter approximation, the following theorem applies to more general settings, including cases where $\widetilde{\mathcal{U}}_A$ is an arbitrary unitary matrix.

Theorem 3 (Concentration bound for composite channels)**.** *Let $\Omega_q(t;N,M,r)$ be a composite channel acting on n qubits for the Hamiltonian $H=A+B$, where the evolution under A is approximated by an arbitrary unitary $\widetilde{U}_A\approx e^{-iAt/r}$ for a time step t/r. To approximate the evolution under B, a finite importance-sampled qDrift channel $\mathcal{E}_q(t;N,1)$ is used, and the channel*

evolves over r steps with unitaries $W_j = e^{-iB_j\tau_j/r}$, where $\tau_j = (tb_j)/(Nq_j)$. Its concentration around its expectation value can be upper bounded $\forall \epsilon \in [0, 4t\lambda_B]$ as

$$\begin{aligned}&\Pr\left[\left\|\frac{1}{M}\left[\prod_{s=r}^{1}\left(\widetilde{U}_A\prod_{k=N}^{1}W_{j_{ks}^m}\right)\right]-\left(\widetilde{U}_A\mathbb{E}\left[W\right]^N\right)^r\right\|\geq\epsilon/2\right]\\&\leq 2^{n+1}\exp-\frac{NMr\epsilon^2}{11t^2\lambda_B^2(1+\max_k\omega(k))^2}\,.\end{aligned}\tag{A.35}$$

In order to guarantee an approximation error $\epsilon/2$ with probability at least $1-\delta$, it is then sufficient to take

$$NMr = 11\frac{t^2\lambda_B^2}{\epsilon^2}\left(1+\max_k\omega(k)\right)^2(n+1)\log\left(\frac{2}{\delta}\right)\,.\tag{A.36}$$

Proof. The result can be shown by extending the martingales introduced in the proof of the concentration bound Theorem 3. We begin by noting that $\mathbb{E}_q[W_j] = \mathbb{E}_q[W]$, which holds independently of j. For each of the M experiments, we consider r sets of N indices corresponding to the unitaries in the product, and denote the j-th unitary in the s-th block for the m-th experiment as W_{sj}^m. Using this notation, we generalize the martingale from Eq. (A.16) as follows:

$$\mathcal{B}_k^m = \left(\prod_{s=r}^{s_k+1}\widetilde{U}_A\mathbb{E}_q[W]^N\right)\left(\widetilde{U}_A\mathbb{E}_q[W]^{N-j_k}\prod_{j=j_k}^{1}W_{s_kj}^m\right)\left(\prod_{s=s_k-1}^{1}\widetilde{U}_A\prod_{j=N}^{1}W_{sj}^m\right)\,,\tag{A.37}$$

where $k = \{0, 1, \ldots, Nr\}$. The indices are defined as $s_k = \lfloor (k-1)/N \rfloor + 1$, and $j_0 = 0$. For $k > 0$, $j_k = (k-1)\%N + 1$. The definition of j_0 differs to account for the edge case where $k = 0$. We also enforce that $\mathcal{B}_k^0 = 0$ and $\mathcal{B}_k^m = 0$ for $m > M$, similar to how it was done for the B_k^m martingale. The causality condition in Definition 3 is satisfied automatically since, for any m, $\mathcal{B}_k^m$ is entirely determined by the random samples $W_{11}^m, \ldots, W_{s_kj_k}^m$ gathered up to the k-th step. The second condition can be checked explicitly, for $\lfloor (k)/N \rfloor = \lfloor (k-1)/N \rfloor$, implying that $s_{k+1} = s_k$ and $j_{k+1} = j_k + 1$, we have

$$\begin{aligned}&\mathbb{E}\left[\mathcal{B}_{k+1}^m|\mathcal{B}_k^m,\ldots,\mathcal{B}_0^m\right]\\&=\left(\prod_{s=r}^{s_k+1}\widetilde{U}_A\mathbb{E}_q[W]^N\right)\left(\widetilde{U}_A\mathbb{E}_q[W]^{N-j_{k+1}}\mathbb{E}_q[W_{s_kj_{k+1}}]\prod_{j=j_{k+1}-1}^{1}W_{s_kj}^m\right)\left(\prod_{s=s_k-1}^{1}\widetilde{U}_A\prod_{j=N}^{1}W_{sj}^m\right)\\&=\left(\prod_{s=r}^{s_k+1}\widetilde{U}_A\mathbb{E}_q[W]^N\right)\left(\widetilde{U}_A\mathbb{E}_q[W]^{N-j_k}\prod_{j=j_k}^{1}W_{s_kj}^m\right)\left(\prod_{s=s_k-1}^{1}\widetilde{U}_A\prod_{j=N}^{1}W_{sj}^m\right)\\&=\mathcal{B}_k^m\end{aligned}\tag{A.38}$$

while, for $\lfloor k/N \rfloor = \lfloor (k-1)/N \rfloor + 1$, i.e., $k = nN$ for some integer n and therefore $s_{k+1} = s_k + 1$ together with $j_k = N$ and $j_{k+1} = 1$, we have the following

$$\begin{aligned}
&\mathbb{E}\left[\mathcal{B}_{k+1}^m | \mathcal{B}_k^m, \ldots, \mathcal{B}_0^m\right] \\
&= \left(\prod_{s=r}^{s_{k+1}+1} \widetilde{U}_A \mathbb{E}_q[W]^N\right) \left(\widetilde{U}_A \mathbb{E}_q[W]^{N-j_{k+1}} \mathbb{E}_q[W_{s_k j_{k+1}}] \prod_{j=j_{k+1}-1}^{1} W_{s_k j}^m\right) \left(\prod_{s=s_{k+1}-1}^{1} \widetilde{U}_A \prod_{j=N}^{1} W_{sj}^m\right) \\
&= \left(\prod_{s=r}^{s_{k+1}+1} \widetilde{U}_A \mathbb{E}_q[W]^N\right) \left(\widetilde{U}_A \mathbb{E}[W]^N\right) \left(\widetilde{U}_A \prod_{j=N}^{1} W_{(s_{k+1}-1)j}^m\right) \left(\prod_{s=s_{k+1}-2}^{1} \widetilde{U}_A \prod_{j=N}^{1} W_{sj}^m\right) \\
&= \left(\prod_{s=r}^{s_k+1} \widetilde{U}_A \mathbb{E}_q[W]^N\right) \left(\widetilde{U}_A \mathbb{E}_q[W]^{N-j_k} \prod_{j=j_k}^{1} W_{s_k j}^m\right) \left(\prod_{s=s_k-1}^{1} \widetilde{U}_A \prod_{j=N}^{1} W_{sj}^m\right) \\
&= \mathcal{B}_k^m \,,
\end{aligned} \tag{A.39}$$

by noting that $N - j_k = 0$.

Finally, the new interpolating martingale for the different experiments takes the following form

$$\mathcal{D}_j = \sum_{m=\lfloor j/N \rfloor + 2}^{M} \mathcal{B}_0^m + \mathcal{B}_{j\%N}^{\lfloor j/N \rfloor + 1} + \sum_{m=1}^{\lfloor j/N \rfloor} \mathcal{B}_N^m \,, \tag{A.40}$$

with the special case for $j = 0$ and $j = NMr$ given, similarly to before, by

$$\mathcal{D}_0 = \sum_{m=1}^{M} \mathcal{B}_0^m = M \left(\prod_{s=r}^{1} \widetilde{U}_A \mathbb{E}_q[W]^N\right) \qquad \mathcal{D}_{NMr} = \sum_{m=1}^{M} \mathcal{B}_{Nr}^m = \sum_{m=1}^{M} \left(\prod_{s=r}^{1} \widetilde{U}_A \prod_{j=N}^{1} W_{sj}^m\right) . \tag{A.41}$$

Since between consecutive indices, the martingale $\mathcal{D}_j$ changes by only a single unitary, the difference sequence $\mathcal{C}_j = \mathcal{D}_j - \mathcal{D}_{j-1}$ have similar properties as C_j in Eq. (A.21) above. In particular

$$\|\mathcal{C}_j\| \leq \left\| \frac{\tau_{k_j}}{r} B_{k_j} \right\| + \mathbb{E}_q \left[\left\| \frac{\tau_k}{r} B_k \right\| \right] \leq \frac{t\lambda_B}{Nr} \left(1 + \max_k \omega(k)\right) \,, \tag{A.42}$$

using the same strategy employed to arrive to Eq. (A.24). For the variance, we instead obtain

$$\left\| \sum_{j=1}^{NMr} \mathbb{E}[\mathcal{C}_j \mathcal{C}_j^\dagger | \mathcal{C}_{j-1} \ldots \mathcal{C}_0] \right\| \leq NMr \max_j \|\mathcal{C}_j\|^2 = M \frac{t^2 \lambda_B^2}{Nr} \left(1 + \max_k \omega(k)\right)^2 . \tag{A.43}$$

The result follows then by taking the parameters R and v from Corollary 4 as

$$R = \frac{t\lambda_B}{Nr} \left(1 + \max_k \omega(k)\right) , \quad v = MNrR^2 \,, \tag{A.44}$$

and using Corollary 4 as was done to show Theorem 3. □

We are now ready to establish an upper bound for the expected error of a composite channel.

Corollary 6 (Fluctuation bound for composite channels). *Let $H = A + B$ be an n-qubit Hamiltonian, $q(j)$ an arbitrary distribution, t the simulation time, N a fixed number of qDrift samples, and M a fixed number of qDrift experiments. Define $\mathcal{U}_H(t)[\rho] = U_H(t)\rho U_H^\dagger(t)$, where $U_H(t) = e^{-iHt}$, and let $\widetilde{\mathcal{U}}_A(t)$ be a first-order Trotter approximation of the channel $\mathcal{U}_A(t)$. Furthermore, let $\mathcal{E}_q^B(t; N, M)$ represent the importance-sampled qDrift channel for the B term, and $\Omega_q(t; N, M, r)$ denote the importance-sampled composite channel. Then, we have*

$$\begin{aligned}\mathbb{E}\left[\| \ \Omega_q(t;N,M,r) - \mathcal{U}_H\|_\diamond\right] \leq & 2\frac{t^2}{r}\left(\Gamma_{comm}^{A,B} + \frac{\lambda_B^2}{N}\left(1+\mathbb{E}_p[\omega]\right)\right) \\ & + \alpha\frac{nt\lambda_B}{NMr}\left(1+\max_k \omega(k)\right) + \alpha\sqrt{\frac{n}{NMr}}t\lambda_B\left(1+\max_k\omega(k)\right),\end{aligned} \tag{A.45}$$

where the parameter

$$\Gamma_{comm}^{A,B} = \sum_{i<j} a_i a_j \|[A_i, A_j]\| + \frac{1}{2}\sum_{ij} a_i b_j \|[A_i, B_j]\|, \tag{A.46}$$

contains the dependence on commutators.

Proof. We prove this corollary in the same manner as Corollary 1, by linking the diamond norm distance to the operator norm for ensembles of unitary channels (see Lemma 2), and then applying the triangle inequality.

$$\begin{aligned}&\mathbb{E}\left[\| \ \Omega_q(t;N,M,r) - \mathcal{U}_H\|_\diamond\right] \leq 2\mathbb{E}\left[\left\|\frac{1}{M}\sum_{m=1}^{M}\prod_{s=r}^{1}\prod_{k=N}^{1}\tilde{U}_A W_{j_{ks}^m} - U_H\right\|\right] \\ &\leq 2\left\|U_H - \left(\widetilde{U}_A\mathbb{E}_q[W]^N\right)^r\right\| + 2\mathbb{E}_q\left[\left\|\frac{1}{M}\sum_{m=1}^{M}\prod_{s=r}^{1}\prod_{k=N}^{1}\tilde{U}_A W_{j_{ks}^m} - \left(\widetilde{U}_A\mathbb{E}_q[W]^N\right)^r\right\|\right].\end{aligned} \tag{A.47}$$

The first term can be bounded by using Theorem 2 and the error bounds for composite channel from Eq. (5.56) to

$$2\left\|U_H - \mathbb{E}_q[\left(\widetilde{U}_A\mathbb{E}_q[W]^N\right)^r\right\| \leq 2\frac{t^2}{r}\left(\Gamma_{\text{comm}}^{A,B} + \frac{\lambda_B^2}{N}\left(1+\mathbb{E}_p[\omega]\right)\right), \tag{A.48}$$

while for the second we have

$$\begin{aligned}&\mathbb{E}_q\left\|\frac{1}{M}\sum_{m=1}^{M}\prod_{s=r}^{1}\prod_{k=N}^{1}\tilde{U}_A W_{j_{ks}^m} - \left(\widetilde{U}_A\mathbb{E}_q[W]^N\right)^r\right\| \\ &= 2\int_0^\infty \Pr\left(\left\|\frac{1}{M}\sum_{m=1}^{M}\prod_{s=r}^{1}\prod_{k=N}^{1}\tilde{U}_A W_{j_{ks}^m} - \left(\widetilde{U}_A\mathbb{E}_q[W]^N\right)^r\right\| \geq \tau\right)d\tau \\ &\leq 2\int_0^\infty \min\left(1, 2^{n+1}e^{-\frac{3M\tau^2}{6NrR^2+2R\tau}}\right)d\tau \\ &\leq \frac{\alpha}{2}\max\left(\sqrt{\frac{n}{NMr}}t\lambda_B(1+\max_k\omega(k)), \frac{nt\lambda_B(1+\max_k\omega(k))}{NMr}\right) \\ &\leq \alpha\left(\sqrt{\frac{n}{NMr}}t\lambda_B(1+\max_k\omega(k)) + \frac{nt\lambda_B(1+\max_k\omega(k))}{NMr}\right),\end{aligned} \tag{A.49}$$

With $R = \frac{t\lambda_B}{Nr}(1 + \max_k \omega(k))$ from above, the integral is evaluated, as in Ref. [195], by splitting it into two parts. The first part contributes nearly one when the denominator in the exponent exceeds the numerator, i.e., $\tau \leq \max\left(\sqrt{\frac{2n}{M}}R, \frac{2Rn}{3M}\right)$, where α suppresses any constant. The contribution from larger τ is marginal, of order $\mathcal{O}\left(\max\left(\sqrt{\frac{2}{M}}R, \frac{2R}{3M}\right)\right)$. □

A.2 COST REDUCTION

In this section, we will recall and present the results for the specific distribution $q_c(j) = \frac{h_j}{C_j}$, where C_j represents the implementation cost of the corresponding term.

Corollary 7. *The expected cost of an importance-sampled qDrift channel with $N = 1$ sample and $q(j) = q_c(j)$ is always lower than that of the standard qDrift.*

$$\mathbb{E}_{q_c}[C] \leq \mathbb{E}_p[C]. \tag{A.50}$$

Proof of Corollary 2.

$$\mathbb{E}_{q_c}[C] = \sum_{j=1}^{L} q(j)C_j = \frac{\sum_{j=1}^{L} \frac{h_j}{C_j}C_j}{\sum_{j=1}^{L} \frac{h_j}{C_j}} = \frac{\sum_{j=1}^{L} h_j}{\sum_{j=1}^{L} \frac{h_j}{C_j}} = \frac{1}{\mathbb{E}_p[\frac{1}{C}]} \leq \mathbb{E}_p[C], \tag{A.51}$$

where we used in the last step Jensen's inequality with $\varphi(C) = 1/C$, which is convex for real positive numbers $C > 0$. □

This result indicates that, on average, unitaries sampled from $q(j) = q_c(j)$ are less costly to implement compared to those sampled from $q(j) = p(j)$. It remains to be demonstrated whether the total implementation cost at a fixed accuracy is also reduced with this choice of distribution.

Theorem 4 (Cost reduction - pure qDrift)**.** *Let N_p and N_{q_c} denote the number of qDrift samples for the two distributions $p(j) = \frac{h_j}{\lambda}$ and $q_c(j) = \frac{h_j}{\lambda_c C_j}$, respectively, for a given target precision ϵ. The expected cost of the importance-sampled qDrift channel is thus always less than that of the standard one:*

$$N_{q_c}\mathbb{E}_{q_c}[C] \leq N_p\mathbb{E}_p[C]. \tag{A.52}$$

However, the number of experiments increases according to

$$M_{q_c}(\epsilon, t) = M_p(\epsilon, t)\frac{(1 + \mathbb{E}_p[1/C]\max_j C_j)^2}{1 + \mathbb{E}_p[1/C]\mathbb{E}_p[C]}. \tag{A.53}$$

More specifically, this increase is independent of the total evolution time t.

Proof of Theorem 4. The factor between N_p and N_{q_c}, see Eq. (5.39), can be expanded as follows

$$\begin{aligned}\mathbb{E}_p\left[\omega(j)\right] &= \sum_j \frac{h_j}{\lambda}\omega(j) = \sum_j \lambda_c \frac{h_j C_j}{\lambda^2} \\ &= \mathbb{E}_p[C]\frac{\lambda_c}{\lambda} = \mathbb{E}_p[C]\mathbb{E}_p[1/C].\end{aligned} \tag{A.54}$$

Utilizing this result, along with Corollary 2 and Theorem 2, we determine that a sufficient condition for N_{q_c} to ensure an error of ϵ over a total time t is at least

$$N_{q_c} = \frac{t^2\lambda^2}{\epsilon}\left(1 + \mathbb{E}_p[\omega(j)]\right). \tag{A.55}$$

This leads to an average total cost of

$$C_{q_c} = \frac{t^2\lambda^2}{\epsilon}\left(1 + \mathbb{E}_p[\omega(j)]\right)\mathbb{E}_{q_c}[C] = \frac{t^2\lambda^2}{\epsilon}\frac{1 + \mathbb{E}_p[C]\mathbb{E}_p[1/C]}{\mathbb{E}_p[1/C]}. \tag{A.56}$$

In contrast, for the standard qDrift approach, we have

$$C_p = 2\frac{t^2\lambda^2}{\epsilon}\mathbb{E}_p[C]. \tag{A.57}$$

In order to guarantee a cost reduction we then need

$$\frac{1 + \mathbb{E}_p[C]\mathbb{E}_p[1/C]}{\mathbb{E}_p[1/C]} \leq 2\mathbb{E}_p[C]\,, \tag{A.58}$$

or equivalently

$$\mathbb{E}_p[C]\mathbb{E}_p[1/C] \geq 1\,. \tag{A.59}$$

This condition is always satisfied, as can be directly demonstrated using Jensen's inequality. The conclusion regarding the increase in the number of experiments directly stems from the definition of the distribution $q_c(j)$ and the sufficient condition given by Eq. (5.47), which is derived from Corollary 1.

□

We will finally demonstrate that the cost reduction remains valid when considering composite channels.

Theorem 5 (Cost reduction - composite channel)**.** *Let $C_p(\epsilon, t)$ and $C_{q_c}(\epsilon, t)$ represent the expected costs for implementing the composite channels $\Omega_p(t; N, M, r)$ and $\Omega_{q_c}(t; N, M, r)$, using the distributions $p(j) = h_j/\lambda$ and $q_c(j) = h_j/(\lambda_c C_j)$, respectively, for a specified target precision ϵ and propagation time t. The following relation holds:*

$$\mathbb{E}_{q_c}[C(\epsilon, t)] \leq \mathbb{E}_p[C(\epsilon, t)]. \tag{A.60}$$

However, the number of experiments increases, such that $M_{q_c}(\epsilon) \geq M_p(\epsilon)$, while maintaining the same scaling with respect to the error ϵ and system size n, and remaining independent of the total evolution time t.

Proof. From Eq. (5.64), we know that the cost of the composite channel can be expressed as

$$C(\epsilon, t) = \frac{t^2}{\epsilon}\left(\sqrt{\Gamma^{A,B}_{\text{comm}} C^A_{tot}} + \lambda_B\sqrt{\mathbb{E}_{q_c}\left[C^B\right]\left(1 + \mathbb{E}_p[\omega(j)]\right)}\right)^2. \tag{A.61}$$

Since the first term is independent of the choice of the sampling distribution, we only need to focus on the second term, which also appears in Theorem 4 and can be bounded as follows:

$$\mathbb{E}_{q_c}\left[C^B\right]\left(1 + \mathbb{E}_p[\omega(j)]\right) \leq 2\mathbb{E}_p\left[C^B\right]. \tag{A.62}$$

For the number of experiments instead, using Eq. (5.71) and Eq. (5.72), for the $p(j)$ distribution we have

$$M_p(\epsilon) = \frac{n}{\epsilon}\frac{2\alpha^2\kappa}{(\kappa-1)^2}\frac{\lambda_B}{\sqrt{C^A_{tot}}}\frac{2^{3/2}\sqrt{\mathbb{E}_p\left[C^B\right]}}{\sqrt{\Gamma^{A,B}_{\text{comm}}C^A_{tot}}+\lambda_B\sqrt{2\mathbb{E}_p\left[C^B\right]}}, \tag{A.63}$$

while for the importance sampled distribution $q_c(j)$ we find

$$\begin{aligned}
M_{q_c}(\epsilon) &= \frac{n}{\epsilon}\frac{2\alpha^2\kappa}{(\kappa-1)^2}\frac{\lambda_B}{\sqrt{C^A_{tot}}}\frac{\left(1+\mathbb{E}_p[1/C^B]\max_j C^B_j\right)^2\sqrt{\frac{1}{\left(1+\mathbb{E}_p[1/C^B]\mathbb{E}_p[C^B]\right)\mathbb{E}_p[1/C^B]}}}{\sqrt{\Gamma^{A,B}_{\text{comm}}C^A_{tot}}+\lambda_B\sqrt{\left(1+\mathbb{E}_p[1/C^B]\mathbb{E}_p[C^B]\right)/\mathbb{E}_p[1/C^B]}} \\
&\geq \frac{n}{\epsilon}\frac{2\alpha^2\kappa}{(\kappa-1)^2}\frac{\lambda_B}{\sqrt{C^A_{tot}}}\frac{\left(1+\mathbb{E}_p[1/C^B]\max_j C^B_j\right)^2\sqrt{\frac{1}{\left(1+\mathbb{E}_p[1/C^B]\mathbb{E}_p[C^B]\right)\mathbb{E}_p[1/C^B]}}}{\sqrt{\Gamma^{A,B}_{\text{comm}}C^A_{tot}}+\lambda_B\sqrt{2\mathbb{E}_p[C^B]}} \\
&\geq \frac{n}{\epsilon}\frac{2\alpha^2\kappa}{(\kappa-1)^2}\frac{\lambda_B}{\sqrt{C^A_{tot}}}\frac{2\sqrt{1+\mathbb{E}_p[1/C^B]\max_j C^B_j}\sqrt{\mathbb{E}_p[C^B]}}{\sqrt{\Gamma^{A,B}_{\text{comm}}C^A_{tot}}+\lambda_B\sqrt{2\mathbb{E}_p[C^B]}} \\
&\geq \frac{n}{\epsilon}\frac{2\alpha^2\kappa}{(\kappa-1)^2}\frac{\lambda_B}{\sqrt{C^A_{tot}}}\frac{2^{3/2}\sqrt{\mathbb{E}_p\left[C^B\right]}}{\sqrt{\Gamma^{A,B}_{\text{comm}}C^A_{tot}}+\lambda_B\sqrt{2\mathbb{E}_p\left[C^B\right]}} = M_p(\epsilon),
\end{aligned} \tag{A.64}$$

where the second line is obtained by remarking that, due to Jensen's inequality,

$$\frac{1+\mathbb{E}_p[C^B]\mathbb{E}_p[1/C^B]}{\mathbb{E}_p[1/C^B]} \leq 2\mathbb{E}_p[C^B]\,. \tag{A.65}$$

In order to get to the third line we used the inequality

$$\frac{\left(1+\mathbb{E}_p[1/C^B]\max_j C^B_j\right)^3}{1+\mathbb{E}_p[1/C^B]\mathbb{E}_p[C^B]} \geq 4\mathbb{E}_p[1/C^B]\mathbb{E}_p[C^B]\,, \tag{A.66}$$

and the last inequality follows from $1+\mathbb{E}_p[1/C^B]\max_j C^B_j \geq 2$. □

COLOPHON

This document was typeset using the typographical look-and-feel `classicthesis` developed by André Miede and Ivo Pletikosić. The style was inspired by Robert Bringhurst's seminal book on typography *"The Elements of Typographic Style"*. `classicthesis` is available for both LaTeX and LyX:

https://bitbucket.org/amiede/classicthesis/