Optimization and Quantum Simulation of Fermions

Maarten Stroeks

Optimization and Quantum Simulation of Fermions

Optimization and Quantum Simulation of Fermions

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus, Prof. dr. ir. H. Bijl,
chair of the Board for Doctorates
to be defended publicly on
Thursday, 18 June 2026, 15.00

by

Maarten Eduard Hubertus Marie STROEKS

This dissertation has been approved by the promotors.

Composition of the doctoral committee:

Rector Magnificus,	chairperson
Prof. dr. B.M. Terhal,	Delft University of Technology, *promotor*
Prof. dr. F. Hassler,	RWTH Aachen University, Germany, *promotor*
Independent members:	
Prof. dr. D.C. Gijswijt,	Delft University of Technology
Prof. dr. M. Blaauboer,	Delft University of Technology
Prof. dr. C.W.J. Beenakker,	Leiden University
Prof. dr. S. Gharibian,	Paderborn University, Germany
Prof. dr. M.T. Wimmer,	Delft University of Technology, *reserve member*

Keywords: quantum computing, quantum complexity theory, quantum many-body physics.
Cover by: Maarten Stroeks.

An electronic copy of this dissertation is available at
`https://repository.tudelft.nl/`.

Aan wijlen mijn grootvader, Eduard Stroeks.

Contents

Summary

This thesis studies computational tasks that arise in quantum many-body physics, with an emphasis on fermionic systems. We develop complexity-theoretic classifications and constructive algorithmic results, thereby establishing opportunities and limitations of classical and quantum algorithms for carrying out these tasks.

Chapter 2 investigates the fermionic satisfiability problem, FERMIONIC k-SAT, which asks whether there exists a state in the joint null-space of parity-preserving local fermionic projectors. This serves as a setting for understanding when the satisfiability of some "quantum fermionic constraints" remains classically decidable. We show that the case $k = 2$ is efficiently solvable classically, but that adding global constraints can crucially change the complexity: enforcing particle-number *parity* preserves classical tractability, while fixing the particle number itself makes the problem NP-complete. We also show that at higher locality $k = 9$ the problem becomes QMA_1-hard.

Chapters 3 and 4 address fermionic Hamiltonian optimization through Gaussian approximations, motivated by the broader question of when simple classes of states provably capture a non-vanishing fraction of an interacting ground-state energy. In stark contrast to qubit Hamiltonians – for which product states provably achieve a *constant* fraction of the ground energy – fermionic Gaussian states achieve only a vanishing fraction in general. This raises the question of whether there are certain (physically motivated) structural assumptions that one can impose on fermionic Hamiltonians that avoid this "Gaussian breakdown". Chapter 3 focuses on classically interacting (i.e., density-density interacting) fermionic Hamiltonians, reflecting the diagonal structure of Coulomb terms in electronic-structure models. Here we prove that fermionic Gaussian states achieve a constant approximation ratio of at least $1/3$ for such Hamiltonians, and we develop efficient semidefinite programming methods for Gaussian approximations for several traceless and positive-semidefinite families, including variants that enforce a fixed average particle number. Chapter 4 then shows that sparsity is another decisive structural assumption: for *sparse* fermionic Hamiltonians, Gaussian states also achieve a constant ratio of the ground energy, and the corresponding Gaussian state can be found efficiently.

Chapter 5 develops quantum algorithms for simulating (sparse) free fermions in a compressed representation. While free-fermion physics is classically tractable in an asymptotic sense, practical costs can still become prohibitive when simulating large systems. The key idea behind our quantum algorithms is to block-encode correlation matrices, Green's functions, and related objects into quantum circuits, enabling quantum procedures whose memory costs scale polylogarithmically in the number of modes — an exponential improvement in memory relative to standard classical representations. Depending on geometric structure (lattices versus general graphs), the resulting runtime improvement over the best known classical methods is exponential or strongly polynomial. Because the associated simulation tasks are BQP-hard (establishing a robust

exponential runtime improvement *in general*), the results also imply that compressed free-fermion simulation should be viewed not merely as a convenience for simulating free-fermion physics, but as a route to potential quantum advantage for some families of instances.

Chapters 6 and 7 share the theme of estimating spectra of many-body Hamiltonians from time-domain data, and they connect algorithm design to the practical bottlenecks of extracting spectral information from experiments. Chapter 6 studies how spectral properties can be recovered from signals derived from time evolution, comparing a classical Monte Carlo scheme that estimates imaginary-time decays to quantum procedures that estimate real-time oscillations generated using phase-estimation-type circuits. We give guarantees on resolvability under assumptions relating to spectral gaps, initial states and sign-problem-freeness of the Hamiltonian. Chapter 7 then targets a different problem related to the practical implementation of phase-estimation-type quantum circuits on *near-term* devices. These circuits contain an expensive *controlled*-time-evolution step, and their depth can be significantly reduced by removing this control, making them more suitable to be implemented on such near-term machines. We realize this by employing adapted *phase retrieval* methods, at the cost of having to run such reduced-depth circuits many more times. Numerical studies on the Fermi–Hubbard model illustrate feasibility and noise resilience of our methods.

Samenvatting

Dit proefschrift bestudeert rekentaken die opkomen in de kwantum-veeldeeltjesfysica, met een nadruk op fermionische systemen. We ontwikkelen complexiteitstheoretische indelingen en constructieve algoritmes, en brengen daarmee mogelijkheden en beperkingen in kaart van klassieke en kwantumalgoritmen voor het uitvoeren van deze taken.

Hoofdstuk 2 bestudeert het fermionische satisfiabiliteitsprobleem, FERMIONIC k-SAT, wat vraagt of er een toestand bestaat in de gezamenlijke nulruimte van pariteitsbehoudende lokale fermionische projectoren. Dit dient als context om te begrijpen wanneer de satisfiabiliteit van bepaalde "kwantum-fermionische restricties" klassiek beslisbaar blijft. We laten zien dat het geval $k = 2$ klassiek efficiënt oplosbaar is, maar dat het toevoegen van globale restricties de complexiteit cruciaal kan veranderen: het afdwingen van de *pariteit* van het aantal deeltjes behoudt klassieke tractabiliteit, terwijl het vastzetten van het aantal deeltjes zelf het probleem NP-compleet maakt. Ook tonen we aan dat het probleem bij hogere localiteit $k = 9$ QMA_1-hard wordt.

Hoofdstukken 3 en 4 bekijken de optimalisatie van fermionische Hamiltonianen via Gaussische benaderingen, gemotiveerd door de bredere vraag wanneer eenvoudige klassen van toestanden aantoonbaar een constant deel van de energie van een interagerende grondtoestand vangen. In contrast met qubit-Hamiltonianen — waarvoor producttoestanden aantoonbaar een *constante* fractie van de grondenergie behalen — behalen fermionische Gaussische toestanden in het algemeen (asymptotisch) nul energie. Dit roept de vraag op of er bepaalde (fysisch gemotiveerde) structurele aannames zijn die men aan fermionische Hamiltonianen kan opleggen om deze "Gaussische instorting" te vermijden. Hoofdstuk 3 richt zich op klassiek interagerende (d.w.z. dichtheid-dichtheid interagerende) fermionische Hamiltonianen, wat de diagonale structuur weerspiegelt van Coulomb-termen in modellen uit elektronische-structuurtheorie. Hier bewijzen we dat fermionische Gaussische toestanden voor zulke Hamiltonianen een constante benaderingsratio van ten minste $1/3$ behalen, en ontwikkelen we efficiënte semidefiniteprogrammingmethoden voor Gaussische benaderingen voor verschillende trace-loze en positief-semidefinitie families, inclusief varianten waarbij het gemiddelde aantal deeltjes vastgezet wordt. Hoofdstuk 4 laat vervolgens zien dat ijlheid een andere doorslaggevende structurele aanname is: voor *ijle* fermionische Hamiltonianen behalen Gaussische toestanden eveneens een constante fractie van de grondenergie, en kan de bijbehorende Gaussische toestand efficiënt worden gevonden.

In Hoofdstuk 5 ontwikkelen we kwantumalgoritmen voor het simuleren van (ijle) vrije fermionen in een gecomprimeerde representatie. Hoewel vrije-fermionfysica asymptotisch klassiek hanteerbaar is, kunnen de praktische kosten bij de simulatie van grote systemen alsnog te kostbaar worden. Het kernidee achter onze kwantumalgoritmen is het block-encoden van correlatiematrices, Green's-functies en verwante objecten in kwantumcircuits, waardoor kwantumprocedures gerealiseerd kunnen worden waarvan

de geheugenkosten polylogaritmisch schalen in het aantal modes — een exponentiële verbetering in geheugen ten opzichte van standaard klassieke representaties. Afhankelijk van de geometrische structuur (roosters versus algemene grafen) is de resulterende runtime-verbetering ten opzichte van de best bekende klassieke methoden exponentieel of sterk polynomiaal. Omdat de bijbehorende simulatietaken BQP-hard zijn (wat *in het algemeen* een robuuste exponentiële runtime-verbetering betekent), impliceren de resultaten bovendien dat gecomprimeerde vrije-fermionsimulatie niet enkel als een handigheid voor het simuleren van vrije-fermionfysica moet worden gezien, maar ook als een route naar potentiële kwantumvoorsprong voor sommige gevallen.

Hoofdstukken 6 en 7 bekijken beiden het probleem van het afschatten van spectra van veeldeeltjes-Hamiltonianen uit tijdsdomeindata, en verbinden algoritmeontwikkeling met de praktische bottlenecks behorende bij het onttrekken van spectrale informatie uit experimenten. Hoofdstuk 6 bestudeert hoe spectrale eigenschappen kunnen worden teruggewonnen uit signalen die voortkomen uit tijdsevolutie, en vergelijkt een klassiek Monte Carlo-schema dat imaginaire-tijdvervallen afschat met kwantumprocedures die reële-tijdoscillaties afschatten die worden gegenereerd aan de hand van circuits van het phase-estimation-type. Bovendien geven we garanties voor de kwaliteit van deze afschattingen onder aannames die betrekking hebben op spectrale gaps, begintoestanden en de af/aanwezigheid van een sign-probleem voor de Hamiltoniaan. Hoofdstuk 7 richt zich vervolgens op een ander probleem dat samenhangt met de praktische implementatie van phase-estimation-achtige kwantumcircuits op *near-term* apparaten. Deze circuits bevatten een kostbare stap van *gecontroleerde* tijdsevolutie, en hun diepte kan aanzienlijk worden verminderd door deze controle te verwijderen, waardoor ze geschikter worden om op zulke near-term machines te worden geïmplementeerd. We realiseren dit door aangepaste *phase retrieval*-methoden toe te passen, ten koste van het feit dat dergelijke circuits met gereduceerde diepte veel vaker moeten worden uitgevoerd. Numerieke studies van het Fermi–Hubbard-model illustreren de haalbaarheid en ruisbestendigheid van onze methoden.

1

Introduction

1.1. Preface

Quantum computers are anticipated to carry out particular computational tasks much more efficiently than their classical counterparts. Error-corrected quantum machinery can store and process large amounts of quantum information, so as to perform complicated quantum computations. Relevant examples of such computations include finding prime-factorizations of n-bit integers (with poly(n) effort)[1] [1] and measuring the state of a quantum system after it has evolved over time [2]. These results are encouraging in the search for useful applications of quantum computers. This search, however, is a surprisingly delicate task and requires a careful treatment.

When identifying applications of quantum computers, one naturally encounters questions such as:

Q1 *Which computing tasks can be performed efficiently using quantum computers?*

Q2 *For which tasks do quantum computers provide a speedup over their classical counterparts?*

Q3 *Which tasks — pertaining to quantum mechanics — can be performed efficiently using classical computers?*

Q4 *Which tasks are likely to be intractable on quantum computers?*

In other words, *what are quantum computers good at?*, *when should you use one?*, *when are classical computers fine?* and *what should you not even hope to be able to do with quantum computers?* Questions like these sparked research in the field of quantum complexity theory (see e.g. [3]), which formally classifies the hardness of certain quantum computational tasks.

When designing quantum algorithms, Q1 and Q2 play a central role[2]. We wish to design quantum algorithms for solving interesting problems, where – ideally – solving them on a classical computer would take "much" longer. A particularly appealing use of quantum algorithms in that respect is performing quantum simulation, exploiting

[1] For completeness, let us note that a robust superpolynomial speedup compared to classical computation has not (yet) been established for this task. That is to say, it is not known whether this problem can be solved with polynomial effort classically.

[2] Q2 is often understudied in comparison to Q1.

the natural ability of quantum computers to emulate quantum physics. Applications of such quantum simulation techniques include estimating spectral properties of quantum many-body systems [4, 5] — which we will get back to in Section 1.2 below.

When dealing with Q1 (i.e., what are quantum computers good at?), one often assumes the quantum computer to be error-corrected, so that it can perform large-scale computations. Although significant experimental progress towards realizing such error-corrected machines is being made [6], answering Q1 with reference to *current* quantum hardware hinges on the clever design of algorithms that minimize the strain on these early quantum computers. Namely, some hardware limitations might yield running some asymptotically efficient quantum algorithms intractable on current machines, but algorithmic improvements could turn that around to some extent.

Answering Q3 (i.e., when are classical computers fine?) is important not just because it provides insights for answers to Q2 (i.e., when do quantum computers provide a speedup?), but also because it leads to the development of efficient classical algorithms for seemingly purely quantum computational tasks. Examples of the latter include calculating certain properties of non-interacting quantum particles (see e.g. [7]) and solving a quantum version of the classical 2-SAT problem [8–10].

The relation between Q2 and Q3 also has a more subtle side: Suppose a task can be efficiently solvable on a classical computer, in the sense that the classical effort scales polynomially in the problem size. In some cases, quantum computers can *even then* provide a speedup. Examples of such tasks are simulating the time dynamics of coupled classical oscillators [11] and the simulation of free fermion physics [12].

Q4 (i.e., what should you not even hope to be able to do with quantum computers?) deserves some more explanation. Some computational tasks are likely to be intractable on quantum computers – just like NP-hard problems are likely to be intractable on classical computers. Such computational tasks include calculating the ground energy and ground state of local quantum systems [3, 13, 14], which is of interest in physics and chemistry. An active direction of research for those problems is to design *classical* approximation algorithms to obtain provably "good" approximations to the true ground state – touching upon Q3.

It is worth mentioning that the majority of this thesis focuses on computational tasks related to *fermionic* quantum systems. The fermionic nature of the problems brings along some opportunities and intricacies, which are emphasized in each chapter.

The contents of this thesis relate to each of the questions and their interplay as discussed in this preface. In Section 1.4, we briefly discuss the chapters of this thesis and on which of the questions Q1-Q4 they touch.

First, we will briefly introduce quantum algorithms (in particular in the setting of quantum simulation) in Section 1.2 and fermionic optimization problems in Section 1.3, as both topics are important themes in this thesis.

1.2. Quantum algorithms

Quantum algorithms are run on quantum computers, the basic building blocks of which are qubits. The *states* of such a quantum computer (consisting of, say, n qubits) are normalized vectors $|\psi\rangle$ in a Hilbert space $(\mathbb{C}^2)^{\otimes n}$. When running a quantum algorithm,

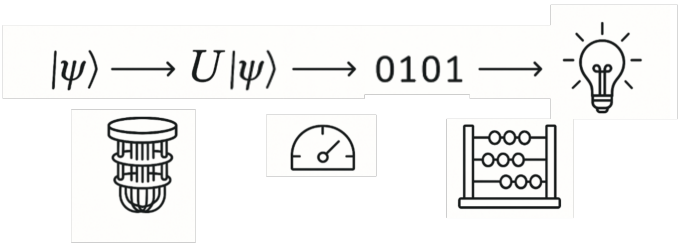

Figure 1.1: Illustration of a quantum algorithm, which typically involves a unitary evolution stage, a measurement stage and a classical post-processing stage.

an initial state $|\psi\rangle$ is processed by some unitary matrix $U \in \mathbb{C}^{2^n \times 2^n}$ to obtain a state $U|\psi\rangle$. These unitaries are products of a poly(n) number of smaller unitaries (*gates*), each acting on few qubits at a time. Having obtained $U|\psi\rangle$, one performs (projective) *measurements* on it (in a local qubit basis) to obtain some quantity of interest — typically after classical post-processing of the measurement outcomes. The idea is that there exist such quantities of interest that can be efficiently obtained using quantum algorithms, but cannot be obtained via purely classical computation. A schematic depiction of a quantum algorithm is given in Figure 1.1.

1.2.1. Extracting spectral features of Hamiltonians

To illustrate how quantum algorithms can be used to learn physically relevant information, let us consider the following algorithm. The goal of the algorithm is to estimate spectral properties of an n-qubit *local Hamiltonian* H. The Hamiltonian $H = \sum_i H_i$ is a Hermitian operator on $(\mathbb{C}^2)^{\otimes n}$, where each H_i acts non-trivially on few qubits. One approach to learning some desired properties of the spectrum of H is to run the following $(n+1)$-qubit quantum algorithm, which makes use of the fact that states can be efficiently time-evolved under local Hamiltonians (i.e., $|\psi\rangle \to \exp(-iHt)|\psi\rangle$) on a quantum computer [2]. This algorithm will be an important theme in Chapters 6 and 7 of this thesis.

1. Prepare $|\psi\rangle_n \otimes |0\rangle$.
2. Apply a so-called Hadamard gate to the $(n+1)$th qubit to obtain $|\psi\rangle_n \otimes 1/\sqrt{2}\big(|0\rangle + |1\rangle\big)$.
3. Apply the n-qubit unitary[3] $U = \exp(-iHk)$ (with k an integer) *conditioned* on the $(n+1)$th qubit to obtain $1/\sqrt{2}\big(|\psi\rangle_n \otimes |0\rangle + U|\psi\rangle_n \otimes |1\rangle\big)$.
4. Apply another a Hadamard gate to the $(n+1)$th qubit to obtain $1/2\big(\mathbb{1} + U\big)|\psi\rangle_n \otimes |0\rangle + 1/2\big(\mathbb{1} - U\big)|\psi\rangle_n \otimes |1\rangle$.
5. Measure the $(n+1)$th qubit in the standard basis $\to 0/1$.

[3] Note that U is not (yet) a product of unitaries that each act on few qubits, which we need it to be to run the algorithm. Using a method called *Trotterization* [2, 15], such a decomposition of U can be approximately realized, where one uses the fact that H is a sum of local terms.

The probability of measuring 0 in step 5 is $\frac{1}{2} + \frac{1}{2}\mathrm{Re}\left[\langle\psi|_n U |\psi\rangle_n\right]$. If we apply a $\exp\left(-i\frac{\pi}{4}Z\right)$-gate[4] to the $(n+1)$th qubit prior to step 4, the probability of measuring 0 is $\frac{1}{2} - \frac{1}{2}\mathrm{Im}\left[\langle\psi|_n U |\psi\rangle_n\right]$. Hence by repeating these experiments many times, we can estimate $f[k] := \langle\psi|_n U |\psi\rangle_n = \langle\psi|_n \exp(-iHk) |\psi\rangle_n \in \mathbb{C}$, where the sampling error scales inverse polynomially with the number of experiments. Repeating the estimation of $f[k]$ for several $k = 0, 1, 2, \ldots$, one learns the time-evolution "signal" of a state $|\psi\rangle$ under a Hamiltonian H. The oscillation frequencies of this signal correspond to the eigenvalues of the Hamiltonian H and thus contain valuable physical information.

Estimating the smallest eigenvalue of general q-local Hamiltonians H up to $1/\mathrm{poly}(n)$ is QMA-complete for $q \geq 2$ [3, 16], implying it cannot be done with $\mathrm{poly}(n)$ effort on a classical or quantum computer (unless BPP=QMA or BQP=QMA, respectively). If we attempted anyway, then what would go wrong when running the algorithm above? The difficulty lies in the preparation of $|\psi\rangle_n$ (one can only learn those eigenvalues that are at least $1/\mathrm{poly}(n)$ supported on $|\psi\rangle_n$) and in the ability to classically extract eigenvalues from the time-evolution signal. Hence we should pick a less ambitious goal and instead try to extract *some* spectral features of H from the time-evolution signal $f[k]$. One such strategy is proposed in Ref. [4], where one learns the spectral density of H with respect to $|\psi\rangle$. Other strategies consist of assuming a guarantee on the overlap between the input state $|\psi\rangle_n$ and the eigenstates whose eigenvalues one tries to estimate [5, 17].

1.2.2. Quantum simulation beyond time dynamics

Quantum simulation does not just encompass implementing time evolution of states under local Hamiltonians [2]. Other types of quantum simulation include preparing Gibbs (or thermal) states $\propto \exp(-\beta H)$, i.e., the state of a system in thermal equilibrium at inverse temperature β (modern references include [18–20]), and estimating out-of-time-order correlators (OTOCs) [21] which can be used to characterize many-body systems.

1.3. Fermionic optimization problems

1.3.1. Fermionic Hamiltonian optimzation

Let us leave the discussion of quantum algorithms for now and switch gears to the problem of (approximately) *optimizing* local Hamiltonians — and more specifically local *fermionic* Hamiltonians. The k-local *qubit* Hamiltonians $H = \sum_i H_i$ discussed so far are defined such that each H_i acts non-trivially on $k = O(1)$ qubits. Instead, let us consider Hamiltonians H that are local in terms of other degrees of freedom; fermionic modes — as these Hamiltonians are ubiquitous in quantum many-body physics. By locality, we mean that each term in H involves $O(1)$ fermionic modes[5]. Such local fermionic Hamiltonians model various physical systems, such as electrons in condensed matter and quantum chemistry — prime targets for quantum simulation. A problem of particular

[4] With Z denoting the Pauli matrix $\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$.

[5] Note that this notion of locality does *not* imply geometric locality or sparsity of the Hamiltonian. *Sparse* fermionic Hamiltonians are defined in Chapter 4.

interest in physics is to find/approximate the ground state of such Hamiltonians. Despite its practical and conceptual relevance, the general problem of approximating interacting fermionic ground states is currently less well understood than its qubit counterpart.

Instead of calculating ground energies/states of these local fermionic or local qubit Hamiltonians – which is QMA-complete [3, 14, 16] even when allowing $1/\mathrm{poly}(n)$ error – one might attempt the less ambitious but more realistic goal of computing a state that achieves a *constant* fraction of the ground energy. When setting up *classical* methods for doing so, one restricts to certain classes of simple states. For qubit Hamiltonians the natural choice are product states, while for fermionic Hamiltonians they are Gaussian states. Note that Gaussian states play a prominent role in fermionic optimization problems using the mean-field Hartree-Fock method, see e.g. Ref. [22]. This naturally leads to the questions *do product states achieve a constant approximation ratio for general local qubit Hamiltonians?* and *do Gaussian states achieve a constant approximation ratio for general local fermionic Hamiltonians?*

To align with computer science conventions, we, w.l.o.g., look to *optimize H* and look at the energy that some ansatz achieves compared to the *maximum* energy, as a proxy for the approximation quality. The above questions then reduce to

$$r_{\mathrm{qubit}} := \max_{\rho_{\mathrm{product}}} \frac{\mathrm{tr}\big(\rho_{\mathrm{product}} H_{\mathrm{qubit}}\big)}{\lambda_{\max}(H_{\mathrm{qubit}})} \stackrel{?}{=} \Omega(1), \tag{1.1a}$$

$$r_{\mathrm{fermion}} := \max_{\rho_{\mathrm{Gaussian}}} \frac{\mathrm{tr}\big(\rho_{\mathrm{Gaussian}} H_{\mathrm{fermion}}\big)}{\lambda_{\max}(H_{\mathrm{fermion}})} \stackrel{?}{=} \Omega(1), \tag{1.1b}$$

for *any* local and *traceless* qubit Hamiltonian H_{qubit} and fermionic Hamiltonian H_{fermion}.

Lieb [23] proved that equation (1.1a) is true for 2-local qubit Hamiltonians, with $r_{\mathrm{qubit}} = 1/9$. In sharp contrast, it was shown in Ref. [24] that equation (1.1b) is *not* true, i.e., there exist local fermionic Hamiltonians for which Gaussian states only ever achieve a vanishing fraction of the maximum energy. This raises the questions (1) *if* $r_{\mathrm{fermion}} \neq \Omega(1)$, *then what can one say about its scaling?* and (2) *can we restrict* H_{fermion} *in some physically motivated way such that equation* (1.1b) *becomes true?*

Before addressing these questions, let us note that they pertain to the *existence* of a product state or Gaussian state that achieves a certain approximation ratio. A stronger demand is to output the description of such a state using a polynomial-time (classical) algorithm. In fact, *finding* a constant-ratio approximation in polynomial-time can be ruled out under mild complexity-theory assumptions [25, 26], since both H_{qubit} and H_{fermion} contain the traceless Max Cut problem as a particular instance. Of course, if one assumes more structure, such a constant-ratio algorithm might become efficient.

Going back to questions (1) and (2), it was shown in Ref. [27] that there is a polynomial-time algorithm outputting a fermionic Gaussian state that achieves approximation ratio $1/\big(n\log(n)\big)$ with high probability. In Ref. [28] (included in this thesis), it was shown that if one assumes H_{fermion} to be sparse, then equation (1.1b) becomes true and such states can be efficiently obtained. Ref. [29] (also included in this thesis) shows that assuming that all interactions in H_{fermion} have a particular yet ubiquitous structure, namely that they are of density-density type, is also sufficient for equation (1.1b) to become true. These results shed some light on the interplay between

physically motivated constraints on Hamiltonians and approximability by simple classes of states such as fermionic Gaussian states.

1.3.2. The fermionic satisfiability problem

Instead of (approximately) optimizing a sum of k-local Hamiltonian terms, one could also consider the problem of determining whether there exists a state that is in the simultaneous ground space of *each* local term. Such quantum generalizations of the classical k-SAT problem are more mild complexity-wise than the problem of optimizing local Hamiltonians. In fact, both the qubit and fermionic versions of this problem are efficiently solvable classically for $k = 2$ (see respectively Refs. [8] and [30], with the latter being part of this thesis). For larger k, these problems become QMA-hard.

1.4. This thesis

The remainder of this introductory chapter will be devoted to an introduction to each chapter of this thesis and to a discussion of preliminaries on fermions in Section 1.5.

Chapter 2

We study the fermionic satisfiability problem, FERMIONIC k-SAT. This is the problem of determining whether there is a state in the null-space of some fermionic parity-preserving projectors, each involving k modes at a time. We prove several results related to this problem. First, we show that if each projector involves at most 2 modes (i.e., $k = 2$), the problem can be solved efficiently classically. Then, we investigate the complexity of adding *global* constraints to the $k = 2$ problem. Adding a constraint that also fixes the particle number *parity* of the state leaves the problem classically solvable. On the other hand, when fixing the particle number itself, the problem becomes NP-complete. In addition, we show that FERMIONIC 9-SAT is QMA_1-hard.

Chapter 3

We consider the optimization problem (ground energy search) of fermionic Hamiltonians with *classical* interactions. This QMA-hard problem is motivated by the Coulomb electron-electron interaction being diagonal in the position basis, a fundamental fact that underpins electronic-structure Hamiltonians in quantum chemistry and condensed matter. We prove that fermionic Gaussian states achieve an approximation ratio of at least 1/3 for such Hamiltonians, independent of sparsity. This shows that classical interactions are sufficient to prevent the vanishing Gaussian approximation ratio observed in SYK-type models. We also give efficient semi-definite programming algorithms for Gaussian approximations to several families of traceless and positive-semidefinite classically interacting Hamiltonians, with the ability to enforce a fixed particle number. The technical core of our results is the concept of a Gaussian blend, a construction for Gaussian states via mixtures of covariance matrices.

Chapter 4

We consider the problem of approximating the ground state energy of a fermionic Hamiltonian using a Gaussian state. In sharp contrast to the dense case [24], we prove that strictly q-local *sparse* fermionic Hamiltonians have a constant Gaussian approximation ratio; the result holds for any connectivity and interaction strengths. Sparsity means that each fermion participates in a bounded number of interactions, and strictly q-local means that each term involves exactly q fermionic (Majorana) operators. We extend our proof to give a constant Gaussian approximation ratio for sparse fermionic Hamiltonians with both quartic and quadratic terms. In each setting we show that the Gaussian state can be efficiently determined. Finally, we prove that the $O(n^{-1/2})$ Gaussian approximation ratio for the normal (dense) SYK-4 model extends to SYK-q for even and constant $q > 4$, with an approximation ratio of $O(n^{1/2-q/4})$. Our results identify non-sparseness as the prime reason that the SYK-4 model can fail to have a constant approximation ratio [24].

Chapter 5

Simulating non-interacting fermion systems is a common task in computational many-body physics. In absence of translational symmetries, modeling free fermions on N modes usually requires poly(N) computational resources. While often moderate, these costs can be prohibitive in practice when large systems are considered. We present several free-fermion problems that can be solved by a quantum algorithm with substantially reduced computational costs. The memory costs are exponentially improved, poly log(N). The runtime improvement, compared to the best known classical algorithms, is either exponential or significantly polynomial, depending on the geometry of the problem. The simulation of free-fermion dynamics is BQP-hard. This implies (under standard assumptions) that our algorithm yields an exponential speedup for any classical algorithm at least for some geometries. The key technique in our algorithm is the block-encoding of objects such as correlation matrices and Green's functions into a unitary. We demonstrate how such unitaries can be efficiently realized as quantum circuits, in the context of dynamics and thermal states of tight-binding Hamiltonians. The special cases of disordered and inhomogeneous lattices, as well as large non-lattice graphs, are presented in detail. Finally, we show that our simulation algorithm generalizes to other promising targets, including free boson systems.

Chapter 6

We consider the task of spectral estimation of local quantum Hamiltonians. The spectral estimation is performed by estimating the oscillation frequencies or decay rates of signals representing the time evolution of states. We present a classical Monte Carlo (MC) scheme which efficiently estimates an imaginary-time, decaying signal for stoquastic (i.e. sign-problem-free) local Hamiltonians. The decay rates in this signal correspond to Hamiltonian eigenvalues (with associated eigenstates present in an input state) and can be extracted using a classical signal processing method like ESPRIT. We compare the efficiency of this MC scheme to its quantum counterpart in which one extracts eigenvalues of a general local Hamiltonian from a real-time, oscillatory signal obtained

through quantum phase estimation circuits, again using the ESPRIT method. We prove that the ESPRIT method can resolve $S = \text{poly}(n)$ eigenvalues, assuming a $1/\text{poly}(n)$ gap between them, with $\text{poly}(n)$ quantum and classical effort through the quantum phase estimation circuits, assuming efficient preparation of the input state. We prove that our Monte Carlo scheme plus the ESPRIT method can resolve $S = O(1)$ eigenvalues, assuming a $1/\text{poly}(n)$ gap between them, with $\text{poly}(n)$ purely classical effort for stoquastic Hamiltonians, requiring some access structure to the input state. However, we also show that under these assumptions, i.e. $S = O(1)$ eigenvalues, assuming a $1/\text{poly}(n)$ gap between them and some access structure to the input state, one can achieve this with $\text{poly}(n)$ purely classical effort for *general* local Hamiltonians. These results thus quantify some opportunities and limitations of Monte Carlo methods for spectral estimation of Hamiltonians.

Chapter 7

We demonstrate the use of adapted classical phase retrieval algorithms to perform control-free quantum phase estimation. We eliminate the costly controlled time evolution and Hadamard test commonly required to access the complex time-series needed to reconstruct the spectrum. This significant reduction of the number of coherent controlled-operations lowers the circuit depth and considerably simplifies the implementation of statistical quantum phase estimation in near-term devices. This seemingly impossible task can be achieved by extending the problem that one wishes to solve to one with a larger set of input signals while exploiting natural constraints on the signal and/or the spectrum. We leverage well-established algorithms that are widely used in the context of classical signal processing, demonstrating two complementary methods to do this, vectorial phase retrieval and two-dimensional phase retrieval. We numerically investigate the feasibility of both approaches for estimating the spectrum of the Fermi-Hubbard model and discuss their resilience to inherent statistical noise.

	Q1	Q2	Q3	Q4
Chapter 2			✓	✓
Chapter 3			✓	✓
Chapter 4			✓	
Chapter 5	✓	✓	✓	
Chapter 6	✓		✓	
Chapter 7	✓			

Table 1.1: Questions from Q1-Q4 (as discussed in the Preface) that are touched upon in each chapter.

1.5. A crash course on fermions

In this thesis, we take the perhaps unorthodox (but arguably more sensible) route of first discussing fermions, and afterwards discussing qubits and their relation to fermions.

We consider a system of n fermionic *modes*, which are degrees of freedom which can be *filled* or *empty*. Let $\mathcal{H}$ denote a 2^n-dimensional Hilbert space and let $a_j \in \mathbb{C}^{2^n \times 2^n}$ for $j \in [n]$ – i.e., one for each mode – denote operators on $\mathcal{H}$. We will refer to these operators as *annihilation* operators, and to their Hermitian conjugates $a_j^\dagger$ as *creation* operators. These operators are defined to satisfy the anti-commutation relations

$$\{a_j, a_k^\dagger\} := a_j a_k^\dagger + a_k^\dagger a_j = \delta_{j,k} \mathbb{1} \text{ and } \{a_j, a_k\} = 0. \tag{1.2}$$

We will refer to vectors in $\mathcal{H}$ with unit l^2-norm as *states*. Let the *vacuum* $|\text{vac}\rangle \in \mathcal{H}$ be a state such that $a_j |\text{vac}\rangle = 0$ for all $j \in [n]$.

Clearly, the 2^n states $a_S^\dagger |\text{vac}\rangle \in \mathcal{H}$ with $a_S^\dagger := a_{j_1}^\dagger a_{j_2}^\dagger \dots a_{j_{|S|}}^\dagger$ $(j_1 < j_2 < \dots < j_{|S|})$ for all subsets of *filled* modes $S \subseteq [n]$ form a basis of $\mathcal{H}$. Any state $|\psi\rangle \in \mathcal{H}$ can therefore be written as

$$|\psi\rangle = \sum_{S \subseteq [n]} \alpha_S a_S^\dagger |\text{vac}\rangle, \tag{1.3}$$

with $\alpha_S \in \mathbb{C}$ and $\sum_{S\subseteq[n]} |\alpha_S|^2 = 1$. Note that $|\text{vac}\rangle$ is unique (up to a phase factor) for the following reason. If $|\text{vac'}\rangle = \sum_{S\subseteq[n]} \alpha_S a_S^\dagger |\text{vac}\rangle$ were another vacuum state, then – by definition – $a_S |\text{vac'}\rangle = 0$ for all non-empty $S \subseteq [n]$. It can be shown straightforwardly that this implies $\alpha_S = 0$ for all non-empty $S \subseteq [n]$, hence $|\alpha_\emptyset| = 1$. The states $|\text{vac}\rangle$ and $|\text{vac'}\rangle$ are thus related by a phase factor: $\langle \text{vac'}|\text{vac}\rangle = \sum_{S\subseteq[n]} \alpha_S^* \langle \text{vac}| a_S |\text{vac}\rangle = \alpha_\emptyset^* \langle \text{vac}|\text{vac}\rangle = \alpha_\emptyset^*$.

It will be useful to define the following *canonical transformation* of the creation and annihilation operators. Given a collection of operators $\{a_j\}$ (and $\{a_j^\dagger\}$) and a unitary matrix $U \in \mathbb{C}^{n\times n}$, then $a_j \to \tilde{a}_j = \sum_{k=1}^n U_{j,k} a_k$ (for all $j \in [n]$) gives another collection of operators satisfying the same anti-commutation relations as in equation (1.2). We will refer to a given collection of operators $\{a_j\}$ as a *mode basis*, so that canonical transformations transform between mode bases.

Let $n_j := a_j^\dagger a_j$ denote the particle number operator for mode j and $\hat{N} := \sum_{j=1}^n n_j$ the total particle number operator. Since $n_j^2 = n_j$, the particle number expectation in each mode is in $[0,1]$. Note that $n_j a_S^\dagger |\text{vac}\rangle = a_S^\dagger |\text{vac}\rangle$ if $j \in S$ (i.e., if j is filled) and zero otherwise. Furthermore, for any canonical transformation $U \in \mathbb{C}^{n\times n}$ transforming $\{a_j\}$ to $\{\tilde{a}_j\}$, we have that $\hat{\tilde{N}} = \hat{N}$.

In addition to the Dirac operators a_j (and $a_j^\dagger$) for $j \in [n]$, it will be convenient to define the (Hermitian) Majorana operators $c_j \in \mathbb{C}^{2^n \times 2^n}$ for $j \in [2n]$ as follows.

$$c_{2p-1} := a_p + a_p^\dagger, \quad c_{2p} := i(a_p - a_p^\dagger). \tag{1.4}$$

These operators satisfy $\{c_j, c_k\} = 2\delta_{j,k}\mathbb{1}$ for $j,k \in [2n]$. We note that any transformation $R \in SO(2n)$ of these operators s.t. $\tilde{c}_j = \sum_k R_{j,k} c_k$ (for all $j \in [2n]$) preserves these anti-commutation relations and thus gives rise to a new set of Majorana operators $\{\tilde{c}_j\}_{j=1}^{2n}$. Like for the Dirac operators, we say that R transforms between mode bases. Note that

any grouping of $2n$ Majorana operators $\{c_j\}_{j=1}^{2n}$ into pairs defines *n fermionic modes* via equation (1.4), since permutations are in $SO(2n)$.

For later reference, let us define Hermitian Majorana monomials

$$C_I := i^{\binom{q}{2}} c_{j_1} c_{j_2} \dots c_{j_q}, \tag{1.5}$$

where I is an ordered subset of $[2n]$. Note that $C_I^2 = \mathbb{1}$ so that the eigenvalues of C_I are all in $\{\pm 1\}$. It can be shown straightforwardly that these C_I's are traceless and are mutually orthogonal under the trace inner product. Furthermore, for $I, J \subseteq [2n]$, they obey

$$C_I C_J = (-1)^{|I||J|-|I\cap J|} C_J C_I. \tag{1.6}$$

1.5.1. Even fermionic operators

Let us discuss so-called *even* Hermitian operators on $\mathcal{H}$. These operators are relevant in a physical context, since any observable corresponds to an even Hermitian operator[6]. Any such operator X on $\mathcal{H}$ can be expressed as

$$X = \alpha_\emptyset \mathbb{1} + \sum_{I\subseteq[2n] \text{ s.t. } |I|=\text{even}} \alpha_I C_I, \tag{1.7}$$

with α_I's real. Indeed, the dimensionality of the even operator space (over the reals) is 2^{2n-1}, which equals the number of even Majorana monomials $|\{I \subseteq [2n] : |I| \text{ even}\}| = 2^{2n-1}$ (which are mutually orthogonal under the trace inner product).

An example of such even Hermitian operators is a *Hamiltonian*, which is a low-degree operator in the sense that it only contains C_I's for small constant $|I|$. Hamiltonians for which $|I| \leq 2$ are free-fermion/quadratic Hamiltonians (see also Section 1.5.2 below) and those that include terms with $|I| = 4$ or higher are *interacting* Hamiltonians. This definition of fermionic Hamiltonians is equivalent to the one given in Section 1.3, where terms in the Hamiltonian *involve at most a constant number of fermionic modes*.

Importantly, any even Hermitian operator X commutes with the particle number parity operator

$$\hat{P} := C_{[2n]} = i^{\binom{2n}{2}} c_1 c_2 \dots c_{2n} = i^n c_1 c_2 \dots c_{2n}. \tag{1.8}$$

Moreover, it can be straightforwardly shown that any Hermitian operator that commutes with $\hat{P}$ is an even Hermitian operator.

Density matrices ρ are Hermitian operators on $\mathcal{H}$ s.t. $\rho \succeq 0$ and $\mathrm{tr}(\rho) = 1$. Therefore, they can be decomposed as probabilistic mixtures $\rho = \sum_j p_j |\psi_j\rangle\langle\psi_j|$ of orthogonal

[6] To illustrate this, let us consider two disjoint subsystems A and B. Let X be an observable on A and Y be an observable on B, so that X and Y do not overlap on any Majorana operator. For simplicity, let us take them to be Hermitian Majorana *monomials*. Let us consider binary projective measurements expressed as POVMs – i.e., $(P, \mathbb{1} - P)$ with P a projector. In particular, we can construct POVMs for measurements associated with the ± 1 eigenvalues of X and Y as follows: $(X_{+1} = (\mathbb{1} + X)/2,\ X_{-1} = (\mathbb{1} - X)/2)$ on A and $(Y_{+1} = (\mathbb{1} + Y)/2,\ Y_{-1} = (\mathbb{1} - Y)/2)$ on B. Since A and B are disjoint subsystems, we should be able to perform joint order-independent measurements on them. That is to say, there should be a joint POVM $(R_{+1,+1}, R_{+1,-1}, R_{-1,+1}, R_{-1,-1})$ s.t. $X_{+1} = R_{+1,+1} + R_{+1,-1}$ and $Y_{+1} = R_{+1,+1} + R_{-1,+1})$. Due to orthogonality of POVM elements, this implies $X_{+1} Y_{+1} = R_{+1,+1} = Y_{+1} X_{+1}$. However, if X and Y are odd Majorana monomials, X_{+1} and Y_{+1} do *not* commute, leading to a contradiction.

states $|\psi_j\rangle \in \mathcal{H}$. They are *pure* iff $\mathrm{tr}(\rho^2) = 1$ (meaning that there is just one state in the mixture), else they are *mixed*. Next, let us argue that density matrices ρ are wlog also *even* Hermitian operators. To that end, we introduce the notion of a quantum channel $\rho \to \mathcal{E}(\rho)$ which describes any evolution of a quantum state – including natural time evolution and measurements.

Definition 1.5.1 (Quantum channel)**.** *A quantum channel $\mathcal{E}(\rho) = \sum_{j=1}^m K_j \rho K_j^\dagger$ is a collection of operators K_j on $\mathcal{H}$ s.t. $\sum_{j=1}^m K_j^\dagger K_j = \mathbb{1}$, known as Kraus operators.*

The "density matrices are even" argument is based on the two statements; (1) the Kraus operators of a quantum channel w.l.o.g. are even fermionic operators and (2) if two channels reproduce the same measurement statistics, they are operationally equivalent.

In all generality, measurement outcome probabilities on ρ after it has undergone some quantum evolution are given by $p = \mathrm{tr}(\rho \sum_{j\subset[m]} K_j^\dagger K_j)$, where the sum now runs over a subset of $\{K_j^\dagger K_j\}_{j=1}^m$ in general. Note that the measurement operator $X := \sum_{j\subset[m]} K_j^\dagger K_j$ is an even Hermitian operator. Since any $|\psi_j\rangle$ in $\rho = \sum_j p_j |\psi_j\rangle\langle\psi_j|$ can be decomposed into an even particle number and an odd particle number contribution $\alpha_j|\psi_j^{\mathrm{even}}\rangle + \beta_j|\psi_j^{\mathrm{odd}}\rangle$ s.t. $\hat{P}|\psi_j^{\mathrm{even}}\rangle = +1|\psi_j^{\mathrm{even}}\rangle$ and $\hat{P}|\psi_j^{\mathrm{odd}}\rangle = -1|\psi_j^{\mathrm{odd}}\rangle$, it is clear that

$$\langle\psi_j| X |\psi_j\rangle = |\alpha_j|^2 \langle\psi_j^{\mathrm{even}}|X|\psi_j^{\mathrm{even}}\rangle + |\beta_j|^2 \langle\psi_j^{\mathrm{odd}}|X|\psi_j^{\mathrm{odd}}\rangle, \tag{1.9}$$

for any even Hermitian operator X since $[X, \hat{P}] = 0$. Therefore,

$$\mathrm{tr}(\rho X) = \mathrm{tr}(\rho_{\mathrm{even}} X) + \mathrm{tr}(\rho_{\mathrm{odd}} X), \tag{1.10}$$

with $\rho_{\mathrm{even}} := \sum_j p_j |\alpha_j|^2 |\psi_j^{\mathrm{even}}\rangle\langle\psi_j^{\mathrm{even}}|$ and $\rho_{\mathrm{odd}} := \sum_j p_j |\beta_j|^2 |\psi_j^{\mathrm{odd}}\rangle\langle\psi_j^{\mathrm{odd}}|$. So, in terms of measurement statistics, there is no interference between odd particle number and even particle number contributions and so w.l.o.g. we can take density matrices to be commuting with $\hat{P}$ and thus to be even operators.

Definition 1.5.2 (Density matrix)**.** *A density matrix ρ is an even Hermitian operator on $\mathcal{H}$ s.t. $\rho \succeq 0$ and $\mathrm{tr}(\rho) = 1$.*

The time dynamics of a closed quantum system is fully determined by its initial density matrix $\rho(t=0)$ and its Hamiltonian H: $\rho(t=0) \to \rho(t) = e^{-iHt}\rho(t=0)e^{+iHt}$. In that case, the corresponding quantum channel (see Definition 1.5.1) is given by a single Kraus operator $K = e^{-iHt}$.

Ground states of Hamiltonians are defined as $\mathrm{argmin}_\rho \mathrm{tr}(\rho H)$, i.e., they are the solutions of minimizing low-degree even Hermitian operators over density matrices – where the minimum is wlog achieved at a pure density matrix. *Thermal states* of Hamiltonians are more practically relevant and are density matrices $\rho \propto \exp(-\beta H)$ (at some inverse temperature $\beta \in \mathbb{R}_+$). Note that zero-temperature thermal states are ground states of H.

In the context of physics, we are typically interested in low-energy states and – in particular – a poly(n) classical description of such states. Since $\rho \in \mathbb{C}^{2^n \times 2^n}$, such a classical description generally requires superpolynomial memory. To make classical treatments tractable, we thus typically restrict to states that do allow an efficient description, such as *fermionic Gaussian states*.

1.5.2. Fermionic Gaussian states

A class of Hamiltonians that is particularly simple is that of quadratic (i.e., free-fermion) Hamiltonians. The eigenstates and thermal states of these Hamiltonians define a simple class of density matrices — fermionic Gaussian states. Fermionic Gaussian states can be represented with poly(n) classical memory, making them particularly useful for classical characterizations of fermionic systems.

Quadratic Hamiltonians are degree-2 polynomials in $\{c_j\}_{j=1}^{2n}$:

$$H = i \sum_{j \neq k}^{2n} h_{j,k} c_j c_k, \tag{1.11}$$

with h a real-valued anti-symmetric matrix and $i c_j c_k$ are Hermitian Majorana monomials C_I with $|I| = 2$. Note that since h is an anti-symmetric matrix, it can be brought to block-diagonal form by $R \in SO(2n)$ as follows.

$$h = R^T \bigoplus_{j=1}^{n} \begin{pmatrix} 0 & -b_j \\ b_j & 0 \end{pmatrix} R, \tag{1.12}$$

with $b_j \in \mathbb{R}$.

Definition 1.5.3 (Fermionic Gaussian states)**.** *Given $2n$ Majorana operators $\{c_j\}_{j=1}^{2n}$. A fermionic Gaussian state is a density matrix of the form*

$$\rho_{\text{Gauss}} = \frac{1}{A} \exp\Big(- i \sum_{j \neq k}^{2n} h_{j,k} c_j c_k \Big), \tag{1.13}$$

where h is a real-valued anti-symmetric matrix as in equation (1.12) *and $A = 2^n \prod_{j=1}^{n} \cosh\big(2 b_j\big)$.*

Using equation (1.12), equation (1.13) can be written as

$$\rho_{\text{Gauss}} = \frac{1}{2^n} \prod_{j=1}^{n} \big(\mathbb{1} + i \lambda_j \tilde{c}_{2j-1} \tilde{c}_{2j} \big), \tag{1.14}$$

where $\tilde{c}_j = \sum_k R_{j,k} c_k$ (with R as in equation (1.12)) and $\lambda_j = \tanh\big(2 b_j\big) \in [-1, +1]$. Iff ρ_{Gauss} is a *pure* fermionic Gaussian state, we have $\lambda_j = \pm 1 \ \forall j \in [n]$ since only then $\text{Tr}(\rho_{\text{Gauss}}^2) = 1$.

Remark 1.5.4. *Any mixed fermionic Gaussian state is a mixture of pure fermionic Gaussian states. To see this, consider equation* (1.14)*, with for some j's, $-1 < \lambda_j < +1$ in the decomposition. For each such j, one can write $\big(\mathbb{1} + i \lambda_j \tilde{c}_{2j-1} \tilde{c}_{2j} \big) = p_j \big(\mathbb{1} + i \tilde{c}_{2j-1} \tilde{c}_{2j} \big) + (1 - p_j) \big(\mathbb{1} - i \tilde{c}_{2j-1} \tilde{c}_{2j} \big)$ with $\lambda_j = 2 p_j - 1$, resulting in a mixture over pure fermionic Gaussian states.*

Fact 1.5.5. *Given a quadratic fermionic Hamiltonian $H = \sum_{j \neq k}^{2n} h_{j,k}\, i c_j c_k$, with h a real-valued, anti-symmetric matrix. The eigenstates of H are fermionic Gaussian states $\rho_{Gauss} = \frac{1}{2^n} \prod_{j=1}^{n} \big(\mathbb{1} + i \lambda_j \tilde{c}_{2j-1} \tilde{c}_{2j} \big)$, with $\lambda_j = \pm 1 \ \forall j \in [n]$ and $\tilde{c}_j = \sum_k R_{j,k} c_k$, where $R \in SO(2n)$ block-diagonalizes h as in equation* (1.12)*. Hence the eigenstates can be obtained with polynomial effort in n.*

1.5.2.1. Particle number conserving setting

A particular type of *pure* fermionic Gaussian state is a Slater determinant state. These states are the eigenstates of particle number conserving free-fermion Hamiltonians, i.e., of free-fermion Hamiltonians H s.t. $[H,\hat{N}]=0$. These Hamiltonians are w.l.o.g. of the form

$$H=\sum_{j,k=1}^{n} h_{j,k}a_k^\dagger a_j, \tag{1.15}$$

where $h\in\mathbb{C}^{n\times n}$ is now a Hermitian matrix.

Definition 1.5.6 (Slater determinant and classical states)**.** *A pure Slater determinant state is a state of the form*

$$\left|\psi_{\text{Slater}}\right\rangle=\tilde{a}_1^\dagger\tilde{a}_2^\dagger\dots\tilde{a}_N^\dagger\left|\text{vac}\right\rangle, \tag{1.16}$$

where $\tilde{a}_j=\sum_{k=1}^n U_{j,k}a_k$ *with* $U\in\mathbb{C}^{n\times n}$ *a unitary matrix and* $N\in\{0,1,\dots,n\}$*. A particular type of Slater determinant state is a classical state which is of the form*

$$\left|\boldsymbol{x}\right\rangle\hat{=}\,a_{j_1}^\dagger a_{j_2}^\dagger\dots a_{j_N}^\dagger\left|\text{vac}\right\rangle, \tag{1.17}$$

where $j_1<j_2<\dots<j_N\in[n]$*,* $N\in\{0,1,\dots,n\}$ *and* $\boldsymbol{x}=(x_1,\dots,x_n)$ *with* $x_{j_1}=\dots=x_{j_N}=1$ *and all other* x_j*'s equal to* 0*.*

1.5.3. Covariance matrices

For any density matrix ρ, we can define a covariance matrix $\Gamma\in\mathbb{R}^{2n\times 2n}$ with entries

$$\Gamma_{j,k}:=\frac{i}{2}\operatorname{tr}\big(\rho[c_j,c_k]\big). \tag{1.18}$$

By its definition $\Gamma_{j,k}\in[-1,1]$, $\Gamma_{j,k}=-\Gamma_{k,j}$, $\Gamma_{k,k}=0$. An anti-symmetric real matrix such as Γ can be block-diagonalized so that

$$\Gamma=R^T\bigoplus_{j=1}^{n}\begin{pmatrix}0&\lambda_j\\-\lambda_j&0\end{pmatrix}R, \tag{1.19}$$

with $R\in\mathrm{SO}(2n)$, $\lambda_j\in\mathbb{R}$. Since $(R\Gamma R^T)_{j,k}=\frac{i}{2}\operatorname{tr}\big(\rho[\tilde{c}_j,\tilde{c}_k]\big)$ (with $\tilde{c}_j=\sum_k R_{j,k}c_k$) and $||C_I||_2\le 1$ for any Hermitian Majorana monomial, we have that $|\lambda_j|\le 1$ for all j. Hence for general fermionic density matrices, Gaussian or non-Gaussian, we have $\Gamma^T\Gamma\le\mathbb{1}$. On the other hand, one can show that ρ is a (pure) fermionic Gaussian state iff $\Gamma\Gamma^T=\mathbb{1}$ [31].

The following basic result about fermionic Gaussian states and their covariance matrices can be derived easily.

Proposition 1.5.7. *Any matrix* $\Gamma\in\mathbb{R}^{2n\times 2n}$ *that is anti-symmetric and has no eigenvalues outside of* $[-i,+i]$ *corresponds to the covariance matrix of a fermionic Gaussian state.*

Proof. One block-diagonalizes the anti-symmetric matrix Γ as in equation (1.19) and the fact that the eigenvalues of Γ are within $[-i,+i]$ implies that $\forall j$, $|\lambda_j|\le 1$. Hence the (unnormalized) fermionic Gaussian state associated with Γ in equation (1.19), is given in equation (1.13) with a block-diagonalized h in equation (1.12) with $b_j=\operatorname{arctanh}(\lambda_j)/2$. □

An essential property of fermionic Gaussian states is Wick's theorem, i.e. expectation values of quartic or higher-order correlations can be expressed in terms of entries of the covariance matrix of the fermionic Gaussian state (see e.g. [31]).

Proposition 1.5.8 (Wick's Theorem)**.** *Let ρ_{Gauss} be a fermionic Gaussian state with covariance matrix Γ. Then*[7]

$$\operatorname{tr}\left(\rho_{\text{Gauss}}\, C_I\right) = \operatorname{Pf}(\Gamma_I), \tag{1.20}$$

where Γ_I is the $|I|\times|I|$ sub-matrix of Γ restricted to the rows and columns in the ordered subset $I\subseteq[2n]$ and $\operatorname{Pf}(\cdot)$ denotes the Pfaffian. In particular, for $|I|=4$:

$$\operatorname{tr}\left(\rho_{\text{Gauss}}\, c_i c_j c_k c_l\right) = -\Gamma_{i,j}\Gamma_{k,l} + \Gamma_{i,k}\Gamma_{j,l} - \Gamma_{i,l}\Gamma_{j,k}, \tag{1.21}$$

with $i<j<k<l$.

From Proposition 1.5.8, we can derive the following simple statement about the relation between fermionic Gaussian states and their covariance matrices, leading to a poly(n)-sized classical description of fermionic Gaussian states.

Proposition 1.5.9. *A fermionic Gaussian state $\rho_{Gauss}\in\mathbb{C}^{2^n\times 2^n}$ is fully specified by its covariance matrix $\Gamma\in\mathbb{R}^{2n\times 2n}$.*

Proof. Since ρ_{Gauss} is a density matrix, it is an even Hermitian matrix (see Definition 1.5.2) and can thus be expanded as in equation (1.7). From Proposition 1.5.8, the fact that $C_I^2=\mathbb{1}$ and the fact that the C_I's are mutually orthogonal, we conclude that all coefficients α_I in the decomposition of ρ_{Gauss} in equation (1.7) are determined by Pfaffians of sub-matrices of Γ. □

1.5.3.1. Particle number conserving setting

An object closely related to the covariance matrix – and particularly relevant in particle number conserving settings – is the 1-RDM M of a density matrix ρ. M is an $n\times n$ Hermitian matrix with entries $M_{j,k}=\operatorname{tr}\left(\rho\, a_j^\dagger a_k\right)$, with $a_j^{(\dagger)}$ the annihilation (creation) operator of mode j. Since M is diagonalized by a unitary $U\in\mathbb{C}^{n\times n}$ – and $\tilde{a}_j=\sum_k U_{j,k}a_k$ is a canonical transformation – the eigenvalues of M are given by $\operatorname{tr}\left(\rho\,\tilde{a}_j^\dagger\tilde{a}_j\right)=\operatorname{tr}\left(\rho\,\tilde{n}_j\right)$ so that $0\preceq M\preceq\mathbb{1}$. M is particularly relevant in particle number conserving settings since pure Slater determinants (i.e., eigenstates of particle number conserving free-fermion Hamiltonians, see Definition 1.5.6) are fully specified by their 1-RDM. For completeness, let us prove this simple statement.

Proposition 1.5.10. *A pure Slater determinant state $\left|\psi_{\text{Slater}}\right\rangle$ (see Definition 1.5.6) is fully specified by its 1-RDM $M\in\mathbb{C}^{n\times n}$.*

Proof. Pure Slater determinant states have the form $\left|\psi_{\text{Slater}}\right\rangle=\tilde{a}_1^\dagger\tilde{a}_2^\dagger\ldots\tilde{a}_N^\dagger|\text{vac}\rangle$, where

$$\tilde{a}_j=\sum_{k=1}^{n}U_{j,k}a_k, \tag{1.22}$$

[7] The Pfaffian of a $2k\times 2k$ matrix A is defined as $\operatorname{Pf}(A)=\frac{1}{2^k k!}\sum_{\pi\in S_{2k}}\operatorname{sign}(\pi)\prod_{i=1}^{k}A_{\pi(2i-1),\pi(2i)}$.

with $U \in \mathbb{C}^{n\times n}$ a unitary matrix and $N \in \{0,1,\ldots,n\}$. They are thus fully determined by a unitary $U \in \mathbb{C}^{n\times n}$ and an occupation number N (w.l.o.g. the first N modes are occupied, as permutations can be incorporated into U). Let M be any 1-RDM associated with a pure Slater determinant state, and let it be diagonalized by $U' \in \mathbb{C}^{n\times n}$. Then U' is the unitary matrix that – combined with the number of eigenvalues that equal 1, i.e., some N' – defines the pure Slater determinant state, since the eigenvalues correspond to occupations in the mode basis defined by U'. Therefore, M fully specifies a pure Slater determinant state. □

1.5.4. Free-fermion time dynamics

Given a fermionic Gaussian state $\rho(0)$ and a free-fermion Hamiltonian $H = i\sum_{j\neq k}^{2n} h_{j,k}c_j c_k$, with h as in equation (1.12). If $\rho(0)$ is left to evolve under H (in a closed manner), then the state of the system at a later time t is given by $\rho(t) = e^{-iHt}\rho(0)e^{iHt}$. As we will show next, $\rho(t)$ is a fermionic Gaussian state and its covariance matrix $\Gamma(t)$ can be obtained with poly(n) classical effort.

As we have the freedom to choose a mode-basis to work in, we can take $\rho(0) = \frac{1}{2^n}\prod_{j=1}^{n}(\mathbb{1} + i\lambda_j c_{2j-1}c_{2j})$, in which case:

$$\rho(t) = \frac{1}{2^n}\prod_{j=1}^{n}\Big[\mathbb{1} + i\lambda_j\big[\sum_p \big(e^{-4ht}\big)_{2j-1,p}c_p\big]\big[\sum_q \big(e^{-4ht}\big)_{2j,q}c_q\big]\Big] = \frac{1}{2^n}\prod_{j=1}^{n}\Big[\mathbb{1} + i\lambda_j \tilde{c}_{2j-1}\tilde{c}_{2j}\Big], \tag{1.23}$$

where we have used that $e^{-4ht} \in SO(2n)$ (since $\det(e^{-4ht}) = e^{\mathrm{tr}(-4ht)} = +1$ because h is traceless). Clearly, $\rho(t)$ is a fermionic Gaussian state – fully specified by its covariance matrix (see Proposition 1.5.9) – and the time dynamics simply induces a canonical transformation of the Majorana operators. The covariance matrix of $\rho(t)$ is given by

$$\Gamma(t)_{j,k} = \frac{i}{2}\,\mathrm{tr}\big(\rho(t)[c_j, c_k]\big) = \big[e^{+4ht}\Gamma(0)e^{-4ht}\big]_{j,k}, \tag{1.24}$$

which can (clearly) be obtained with poly(n) classical effort and can be used to evaluate e.g. expectations on Majorana monomials using Wick's Theorem.

Remark 1.5.11. *When we introduced Majorana operators in equation* (1.4)*, we specified that any transformation $R \in SO(2n)$ s.t. $\tilde{c}_j = \sum_k R_{j,k}c_k$ (for all $j \in [2n]$) preserves the right anti-commutation relations and thus gives rise to a new set of Majorana operators $\{\tilde{c}_j\}_{j=1}^{2n}$. In principle, $R \in O(2n)$ also leads to a valid new set of Majorana operators that obey the right anti-commutation relations. However, these transformations are not induced by free-fermion dynamics (as discussed above, where we explicitly have* $\det(R) = +1$*) and so we typically restrict to $R \in SO(2n)$.*

A more detailed discussion of the classical simulation of free-fermion dynamics can be found in Refs. [7, 32], including for instance the classical simulation of free-fermion evolution where part of the evolution is conditioned on occupations of some of the modes.

Although BQP-hard in general, some fermionic time evolution beyond *free* fermions can also be simulated classically. This is typically done by considering simulations that

are "slightly" non-Gaussian; Gaussian evolution of slightly non-Gaussian initial states (see e.g. [33]) or slightly non-Gaussian evolution of Gaussian initial states (see e.g. [34]).

1.5.4.1. Particle number conserving setting

Although the time evolution of Slater determinant states under particle number conserving free-fermion Hamiltonians (i.e., H in equation (1.15)) is a particular case of the discussion above[8], let us consider it in some more detail to introduce some terminology that will be used in later chapters of this thesis.

Given an initial pure Slater determinant state $\left|\psi(t=0)\right\rangle = a_1^\dagger a_2^\dagger \dots a_N^\dagger |\text{vac}\rangle$ – which we take to be a classical state as we have the freedom to choose a mode basis to work in. The (pure) state at time t is given by $\left|\psi(t)\right\rangle = e^{-iHt}\left|\psi(t=0)\right\rangle$, with H as in equation (1.15), which can be straightforwardly shown to be equal to:

$$\begin{aligned}\left|\psi(t)\right\rangle = \Big[\sum_{j_1=1}^{n} (e^{-ith})_{1,j_1} a_{j_1}^\dagger\Big]\Big[\sum_{j_2=1}^{n} (e^{-ith})_{2,j_2} a_{j_2}^\dagger\Big]\dots\Big[\sum_{j_N=1}^{n} (e^{-ith})_{N,j_N} a_{j_N}^\dagger\Big]|\text{vac}\rangle \\ = \tilde{a}_1^\dagger \tilde{a}_2^\dagger \dots \tilde{a}_N^\dagger |\text{vac}\rangle\,, \quad (1.25)\end{aligned}$$

where for the second equality we have used that $e^{+ith^*} \in U(n)$ is a canonical transformation s.t. $\tilde{a}_j = \sum_{k=1}^n \left(e^{+ith^*}\right)_{j,k} a_k$ (for all $j \in [n]$). Clearly, $\left|\psi(t)\right\rangle$ is a Slater determinant state – fully specified by its 1-RDM (see Proposition 1.5.10). The 1-RDM of $\left|\psi(t)\right\rangle$ is given by

$$M(t)_{j,k} = \left\langle\psi(t)\right| a_j^\dagger a_k \left|\psi(t)\right\rangle = \left[e^{+ith} M(0) e^{-ith}\right]_{j,k}, \quad (1.26)$$

which, again, can clearly be obtained with poly(n) classical effort.

1.5.5. Fermions and qubits

Since quantum algorithms are executed on quantum computers, whose basic degrees of freedom are qubits, we will briefly discuss qubit systems and in particular their relation to fermions. The n-mode fermionic Hilbert space $\mathcal{H}$ (we will refer to it as $\mathcal{H}_{\text{fermion}}$) is isomorphic to the n-qubit Hilbert space $\mathcal{H}_{\text{qubit}} = (\mathbb{C}^2)^{\otimes n}$. Indeed, the states $|\mathbf{x}\rangle = |x_1 x_2 \dots x_n\rangle$ with $x_j \in \{0,1\}$ for $j \in [n]$ form a basis for $\mathcal{H}_{\text{qubit}}$ and are related to fermionic basis states $a_S^\dagger|\text{vac}\rangle$ (defined at the start of Section 1.5) by the following unitary:

$$U|x_1 x_2 \dots x_n\rangle = a_S^\dagger|\text{vac}\rangle\,, \text{ with } j \in S \text{ if } x_j = 1 \text{ and } j \notin S \text{ if } x_j = 0. \quad (1.27)$$

Let us define the Pauli matrices for $\{X_j, Y_j, Z_j\}_{j\in[n]}$, which are (in the standard basis) given by:

$$X_j = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y_j = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z_j = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (1.28)$$

These are operators on $\mathcal{H}_{\text{qubit}}$[9] that commute if they act on different qubits, and that anti-commute when acting on the same qubit.

[8] In case H in equation (1.15) is taken to be traceless.

[9] When we write $P_j \in \mathbb{C}^{2\times 2}$ for a Pauli matrix on qubit j, we actually mean $\mathbb{1}_1 \otimes \mathbb{1}_2 \otimes \dots P_j \dots \otimes \mathbb{1}_n$.

Any Hermitian operator on $\mathcal{H}_{\text{qubit}}$ can be written as

$$X = \sum_{\mathbf{z}\in\{0,1,2,3\}^n} \alpha_{\mathbf{z}} \prod_{j=1}^{n} P_j^{(z_j)}, \tag{1.29}$$

where $P_j^{(0)} = \mathbb{1}_j$, $P_j^{(1)} = X_j$, $P_j^{(2)} = Y_j$, $P_j^{(3)} = Z_j$ and $\alpha_{\mathbf{z}} \in \mathbb{R}$. Indeed, Hermitian Pauli monomials $P_{\mathbf{z}} := \prod_{j=1}^{n} P_j^{(z_j)}$ square to identity and are orthogonal under the trace inner product, and the number 4^n of such monomials equals the dimensionality of the Hermitian operator space (over the reals)[10].

Qubit density matrices ρ are Hermitian operators s.t. $\rho \succeq 0$ and $\text{tr}(\rho) = 1$ and (local) qubit Hamiltonians H are Hermitian operators that consist only of Pauli monomials $P_{\mathbf{z}}$ of constant weight (i.e., that involve only a constant number of qubits). Just like for fermions, a closed qubit system with Hamiltonian H evolves as follows: $\rho(t=0) \to \rho(t) = e^{-iHt}\rho(t=0)e^{+iHt}$. *Product states* are qubit density matrices that factorize as $\rho = \rho_1 \otimes \rho_2 \otimes \ldots \otimes \rho_n$, where ρ_j is a single-qubit density matrix on qubit j. Hamiltonians that only contain Pauli monomials of weight one are classically diagonalizable (like free-fermion Hamiltonians in the fermionic setting) and their eigenstates are product states.

The unitary in equation (1.27) can naturally be used to transform Hermitian qubit operators to Hermitian fermionic operators and vice versa. This transformation is known as the Jordan-Wigner transformation [35] and can be explicitly stated as follows[11]

$$Uc_{2k-1}U^\dagger = \Big[\prod_{j<k} Z_j\Big] X_k, \quad Uc_{2k}U^\dagger = \Big[\prod_{j<k} Z_j\Big] Y_k. \tag{1.30}$$

These Pauli monomials naturally reproduce the Majorana anti-commutation relations $\{c_p, c_q\} = 2\delta_{p,q}\mathbb{1}$. Via equation (1.30), Hermitian Majorana monomials map onto Hermitian Pauli monomials and vice versa. Importantly, the weight of a Majorana monomial C_I is not preserved when mapped onto a Pauli monomial $P_{\mathbf{z}}$ and hence local Fermionic Hamiltonians do not map onto local qubit Hamiltonians under the Jordan-Wigner transform (the locality of the latter can even become $\Theta(n)$).

The Jordan-Wigner transform is not the only manner in which fermionic operators can be mapped onto qubit operators. Other mappings arise when picking different unitaries than the one in equation (1.27) to identify basis states for $\mathcal{H}_{\text{qubit}}$ with basis states for $\mathcal{H}_{\text{fermion}}$. Clever choices for these unitaries can lead to improved properties of the mapping: the Bravyi-Kitaev transformation [36] for instance maps constant-locality fermionic Hamiltonians onto $\log(n)$-local qubit Hamiltonians.

Other ways of mapping fermions onto qubits include the encoding of n-mode fermionic states/operators into a larger $m > n$-qubit Hilbert space, rather than via Hilbert space isomorphisms, in such a way that the fermionic algebra is reproduced. These approaches [36, 37] bring along some advantages, like the fact that *sparse* constant-locality fermionic Hamiltonians get mapped onto *sparse* constant-locality

[10] Note that the dimensionality is double that of the space of even Hermitian fermionic operators, and indeed some Hermitian qubit operators map onto odd Hermitian fermionic operators.

[11] Note that there is some freedom in choosing this transformation, such as adding phase factors to the Z_j operators.

qubit Hamiltonians. The disadvantage – from a more practical perspective – is that one has to impose certain stabilizer constraints on the qubit Hilbert space, which have more-than-constant locality in general. Especially for lattice models, however, these constraints and their imposition simplify, see Refs. [38, 39].

References

[1] P. 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.

[2] S. Lloyd. 'Universal Quantum Simulators'. In: *Science* 273.5278 (1996), pp. 1073–1078. DOI: 10.1126/science.273.5278.1073.

[3] A. Y. Kitaev, A. H. Shen and M. N. Vyalyi. *Classical and Quantum Computation*. USA: American Mathematical Society, 2002. ISBN: 0-8218-3229-8.

[4] R. D. Somma. 'Quantum eigenvalue estimation via time series analysis'. In: *New Journal of Physics* 21.12 (Dec. 2019), p. 123025. DOI: 10.1088/1367-2630/ab5c60.

[5] T. E. O'Brien, B. Tarasinski and B. M. Terhal. 'Quantum phase estimation of multiple eigenvalues for small-scale (noisy) experiments'. In: *New Journal of Physics* 21.2 (Feb. 2019), p. 023022. DOI: 10.1088/1367-2630/aafb8e.

[6] R. Acharya *et al.* 'Quantum error correction below the surface code threshold'. In: *Nature* 638.8052 (2025), pp. 920–926. DOI: 10.1038/s41586-024-08449-y.

[7] B. M. Terhal and D. P. DiVincenzo. 'Classical simulation of noninteracting-fermion quantum circuits'. In: *Physical Review A* 65.3 (Mar. 2002). ISSN: 1094-1622. DOI: 10.1103/physreva.65.032325.

[8] S. Bravyi. 'Efficient algorithm for a quantum analogue of 2-SAT'. In: *Contemporary Mathematics*. Vol. 536. American Mathematical Society, 2011. eprint: quant-ph/0602108.

[9] N. de Beaudrap and S. Gharibian. 'A Linear Time Algorithm for Quantum 2-SAT'. In: *31st Conference on Computational Complexity (CCC 2016)*. Vol. 50. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016, 27:1–27:21. ISBN: 978-3-95977-008-8. DOI: 10.4230/LIPIcs.CCC.2016.27.

[10] I. Arad, M. Santha, A. Sundaram and S. Zhang. 'Linear-Time Algorithm for Quantum 2-SAT'. In: *Theory of Computing* 14 (2018), pp. 1–27. URL: https://theoryofcomputing.org/articles/v014a001/.

[11] R. Babbush, D. W. Berry, R. Kothari, R. D. Somma and N. Wiebe. 'Exponential Quantum Speedup in Simulating Coupled Classical Oscillators'. In: *Phys. Rev. X* 13 (4 Dec. 2023), p. 041041. DOI: 10.1103/PhysRevX.13.041041.

[12] M. Stroeks, D. Lenterman, B. M. Terhal and Y. Herasymenko. 'Solving free fermion problems on a quantum computer'. In: *Phys. Rev. Res.* 7 (4 Nov. 2025), p. 043176. DOI: 10.1103/zmwm-gdmw.

[13] T. Cubitt and A. Montanaro. 'Complexity Classification of Local Hamiltonian Problems'. In: *SIAM Journal on Computing* 45.2 (2016), pp. 268–316. DOI: 10.1137/140998287.

[14] B. O'Gorman, S. Irani, J. Whitfield and B. Fefferman. 'Intractability of Electronic Structure in a Fixed Basis'. In: *PRX Quantum* 3 (2 May 2022), p. 020322. DOI: 10.1103/PRXQuantum.3.020322.

[15] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe and S. Zhu. 'Theory of Trotter Error with Commutator Scaling'. In: *Phys. Rev. X* 11 (1 Feb. 2021), p. 011020. DOI: 10.1103/PhysRevX.11.011020.

[16] J. Kempe, A. Kitaev and O. Regev. 'The Complexity of the Local Hamiltonian Problem'. In: *SIAM Journal on Computing* 35.5 (2006), pp. 1070–1097. DOI: 10.1137/S0097539704445226.

[17] M. E. Stroeks, J. Helsen and B. M. Terhal. 'Spectral estimation for Hamiltonians: a comparison between classical imaginary-time evolution and quantum real-time evolution'. In: *New Journal of Physics* 24.10 (Oct. 2022), p. 103024. DOI: 10.1088/1367-2630/ac919c.

[18] C. Rouzé, D. S. França and Á. M. Alhambra. *Optimal quantum algorithm for Gibbs state preparation*. 2024. arXiv: 2411.04885 [quant-ph]. URL: https://arxiv.org/abs/2411.04885.

[19] C.-F. Chen, M. J. Kastoryano, F. G. S. L. Brandão and A. Gilyén. *Quantum Thermal State Preparation*. 2023. arXiv: 2303.18224 [quant-ph]. URL: https://arxiv.org/abs/2303.18224.

[20] Š. Šmíd, R. Meister, M. Berta and R. Bondesan. *Polynomial Time Quantum Gibbs Sampling for Fermi-Hubbard Model at any Temperature*. 2025. arXiv: 2501.01412 [quant-ph]. URL: https://arxiv.org/abs/2501.01412.

[21] D. A. Abanin *et al.* 'Observation of constructive interference at the edge of quantum ergodicity'. In: *Nature* 646.8086 (2025), pp. 825–830. DOI: 10.1038/s41586-025-09526-6.

[22] C. V. Kraus and J. I. Cirac. 'Generalized Hartree–Fock theory for interacting fermions in lattices: numerical methods'. In: *New Journal of Physics* 12.11 (2010), p. 113004. URL: https://doi.org/10.1088/1367-2630/12/11/113004.

[23] E. H. Lieb. 'The classical limit of quantum spin systems'. In: *Communications in Mathematical Physics* 31.4 (1973), pp. 327–340. DOI: 10.1007/BF01646493.

[24] M. B. Hastings and R. O'Donnell. 'Optimizing Strongly Interacting Fermionic Hamiltonians'. In: *Proc. of STOC 2022*. New York, NY, USA: ACM, 2022, pp. 776–789. DOI: https://doi.org/10.1145/3519935.3519960.

[25] S. Arora, E. Berger, H. Elad, G. Kindler and M. Safra. 'On non-approximability for quadratic programs'. In: *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*. 2005, pp. 206–215. DOI: 10.1109/SFCS.2005.57.

[26] S. Khot and R. O'Donnell. 'SDP gaps and UGC-hardness for Max-Cut-Gain'. In: *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*. 2006, pp. 217–226. DOI: 10.1109/FOCS.2006.67.

[27] S. Bravyi, D. Gosset, R. König and K. Temme. 'Approximation algorithms for quantum many-body problems'. In: *Journal of Mathematical Physics* 60.3 (2019), p. 032203. eprint: https://doi.org/10.1063/1.5085428. URL: https://doi.org/10.1063/1.5085428.

[28] Y. Herasymenko, M. Stroeks, J. Helsen and B. Terhal. 'Optimizing sparse fermionic Hamiltonians'. In: *Quantum* 7 (Aug. 2023), p. 1081. ISSN: 2521-327X. DOI: 10.22331/q-2023-08-10-1081.

[29] M. Stroeks, B. M. Terhal and Y. Herasymenko. *Optimizing fermionic Hamiltonians with classical interactions*. 2025. arXiv: 2510.02122 [quant-ph]. URL: https://arxiv.org/abs/2510.02122.

[30] M. Stroeks and B. M. Terhal. 'Complexity of Fermionic 2-SAT'. In: *Quantum* 9 (Oct. 2025), p. 1900. ISSN: 2521-327X. DOI: 10.22331/q-2025-10-31-1900.

[31] F. de Melo, P. Ćwikliński and B. M. Terhal. 'The power of noisy fermionic quantum computation'. In: *New Journal of Physics* 15.1 (Jan. 2013), p. 013015. DOI: 10.1088/1367-2630/15/1/013015.

[32] S. Bravyi. 'Lagrangian representation for fermionic linear optics'. In: *Quantum Info. Comput.* 5.3 (May 2005), pp. 216–238. ISSN: 1533-7146.

[33] B. Dias and R. Koenig. 'Classical simulation of non-Gaussian fermionic circuits'. In: *Quantum* 8 (May 2024), p. 1350. ISSN: 2521-327X. DOI: 10.22331/q-2024-05-21-1350.

[34] O. Reardon-Smith, M. Oszmaniec and K. Korzekwa. 'Improved simulation of quantum circuits dominated by free fermionic operations'. In: *Quantum* 8 (Dec. 2024), p. 1549. ISSN: 2521-327X. DOI: 10.22331/q-2024-12-04-1549.

[35] P. Jordan and E. Wigner. 'Über das Paulische Äquivalenzverbot'. In: *Zeitschrift für Physik* 47.9 (1928), pp. 631–651. DOI: 10.1007/BF01331938.

[36] S. B. Bravyi and A. Y. Kitaev. 'Fermionic Quantum Computation'. In: *Annals of Physics* 298.1 (2002), pp. 210–226. URL: https://doi.org/10.1006/aphy.2002.6254.

[37] K. Setia, S. Bravyi, A. Mezzacapo and J. D. Whitfield. 'Superfast encodings for fermionic quantum simulation'. In: *Phys. Rev. Res.* 1 (3 Oct. 2019), p. 033033. DOI: 10.1103/PhysRevResearch.1.033033.

[38] C. Derby, J. Klassen, J. Bausch and T. Cubitt. 'Compact fermion to qubit mappings'. In: *Phys. Rev. B* 104 (3 July 2021), p. 035118. DOI: 10.1103/PhysRevB.104.035118.

[39] C. Derby and J. Klassen. *A Compact Fermion to Qubit Mapping Part 2: Alternative Lattice Geometries*. 2021. arXiv: 2101.10735 [quant-ph]. URL: https://arxiv.org/abs/2101.10735.

2

Complexity of Fermionic 2-SAT

Tell me where the story ends

Lust for Life (2025), Courting

2.1. Introduction

Classical k-satisfiability (k-SAT) is the problem of deciding whether there exists a configuration of n Boolean variables which satisfies a collection of k-variable Boolean clauses. While 2-SAT is solvable in linear time [1], ($k \geq 3$)-SAT is NP-complete. Bravyi introduced the problem QUANTUM k-SAT [2]: this is the problem of deciding whether there exists an n-qubit state which is in the null-space of some collection of k-qubit projectors. Bravyi showed that QUANTUM 2-SAT can be decided classically efficiently, and later it was shown it can even be done in linear time [3, 4]. For $k \geq 3$, it has been shown that QUANTUM k-SAT is QMA_1-complete [2, 5].

Here, we introduce the *fermionic* satisfiability problem. An n-mode fermionic system corresponds to set of annihilation operators $\{a_j\}_{j=1}^n$ (with Hermitian conjugate creation operators $\{a_j^\dagger\}_{j=1}^n$) which satisfy $\{a_i^\dagger, a_j\} := a_i^\dagger a_j + a_j a_i^\dagger = \delta_{ij} I$, $\{a_i, a_j\} = 0$. In addition, there is a state $|\mathrm{vac}\rangle$ —called the vacuum, or empty, zero-particle state— for which $a_i |\mathrm{vac}\rangle = 0$, $\forall i$ (see also Section 1.5).

FERMIONIC k-SAT is the problem of deciding whether there exists an n-mode fermionic state which is in the null-space of a collection of k-mode fermionic projectors. Each k-mode projector is a polynomial in a_j and $a_j^\dagger$ with $j \in S$, where $S \subseteq [n]$ and $|S| = k$. Furthermore, the k-mode projector commutes with the parity operator $\hat{P} := (-1)^{\sum_{i=1}^n a_i^\dagger a_i}$. As a special case, we also study PARTICLE-NUMBER-CONSERVING (PNC) FERMIONIC k-SAT, where each projector additionally commutes with the particle-number operator $\hat{N} := \sum_{i=1}^n a_i^\dagger a_i$. In addition to studying the satisfiability of these problems, we investigate the complexity of PNC FERMIONIC 2-SAT and FERMIONIC 2-SAT when, respectively, the particle number is fixed to some given value $\hat{N} = N \in \{0, 1 \ldots, n\}$, or the parity $\hat{P} = P \in \{+1, -1\}$ is fixed.

2.1.1. Main results

Problem 2.1.1 (FERMIONIC k-SAT)**.** *Given is a set of projectors $\{\Pi_i\}$ on n fermionic modes, where each Π_i is a polynomial in the operators $a_j^\dagger$ and a_j with index j contained in a subset of at most k fermionic modes. Furthermore, $[\Pi_i, \hat{P}] = 0$ for all projectors $\{\Pi_i\}$ with $\hat{P} = (-1)^{\sum_{i=1}^n a_i^\dagger a_i}$. Decide whether there exists a $|\psi\rangle$ s.t. $\forall i, \Pi_i |\psi\rangle = 0$ (YES), or for all $|\psi\rangle$, $\sum_i \langle\psi|\Pi_i|\psi\rangle > \frac{1}{\text{poly}(n)}$ (NO).*

To be precise in our definition of FERMIONIC k-SAT one should, in principle, provide the accuracy with which the projectors Π_i are specified, i.e. with respect to what elementary gate set, see e.g. [3, 6] for how this is handled for QUANTUM k-SAT. Since this is separate from our results, we leave such precise definition to future work.

We also consider FERMIONIC k-SAT with global additional constraints, i.e. one asks for an assignment with given particle number or parity:

Problem 2.1.2 (FERMIONIC k-SAT with fixed parity P)**.** *Given n fermionic modes, $P \in \{-1, +1\}$, and a set of projectors $\{\Pi_i\}$ where each projector Π_i is a polynomial in the operators $a_j^\dagger$ and a_j with index j contained in a subset of at most k fermionic modes. Furthermore, $\forall i, [\Pi_i, \hat{P}] = 0$. Decide whether there exists $|\psi\rangle$ s.t. $\forall i$, $\Pi_i |\psi\rangle = 0$ and $\hat{P}|\psi\rangle = P|\psi\rangle$ (YES), or $\forall |\psi\rangle$ s.t. $\hat{P}|\psi\rangle = P|\psi\rangle$, $\sum_i \langle\psi|\Pi_i|\psi\rangle > \frac{1}{\text{poly}(n)}$ (NO).*

A special case of Problem 2.1.1 is the case when each projector commutes with the particle number, i.e. $\forall i$, $[\Pi_i, \hat{N}] = 0$, which we will refer to as PARTICLE-NUMBER-CONSERVING (PNC) FERMIONIC k-SAT. We also consider the following problem:

Problem 2.1.3 (PNC FERMIONIC k-SAT with fixed particle number N)**.** *Given n fermionic modes, an integer $N \in \{0, 1, \ldots, n\}$ and a set of projectors $\{\Pi_i\}$ where each projector Π_i is a polynomial in the operators $a_j^\dagger$ and a_j with index j contained in a subset of at most k fermionic modes. Furthermore, $\forall i, [\Pi_i, \hat{N}] = 0$. Decide whether there exists a state $|\psi\rangle$ s.t. $\forall i$, $\Pi_i |\psi\rangle = 0$ and $\hat{N}|\psi\rangle = N|\psi\rangle$ (YES), or $\forall |\psi\rangle$ s.t. $\hat{N}|\psi\rangle = N|\psi\rangle$, $\sum_i \langle\psi|\Pi_i|\psi\rangle > \frac{1}{\text{poly}(n)}$ (NO).*

Our main results are:

Theorem 2.1.4. FERMIONIC 2-SAT $\in$ P, *and can be decided in time $O(n+m)$, where $m = |\{\Pi_i\}|$ denotes the number of projectors.*

Theorem 2.1.5. FERMIONIC 2-SAT with fixed parity $\in$ P, *and can be decided in time $O(nm)$, where $m = |\{\Pi_i\}|$ denotes the number of projectors.*

Theorem 2.1.6. PNC FERMIONIC 2-SAT with fixed particle number *is* NP*-complete.*

As part of the proof of Theorem 2.1.5, we show that classical 2-SAT with fixed (Hamming weight) parity can be solved in $O(nm)$ time, see Theorem 2.C.1 in Appendix 2.C for details. We became aware that this classical fixed-parity result was also derived in [7].

To provide further context for our results, we also provide a straightforward proof that FERMIONIC k-SAT is QMA_1-hard for $k = 9$ in Section 2.6.

2.1.2. Overview and main ideas

FERMIONIC 2-SAT instances are defined on a graph G whose vertices are fermionic modes and whose edges correspond to projectors involving pairs of modes. These projectors are of rank at most three (if there is to be a satisfying assignment) in the two fermionic mode subspace, consisting of at most three rank-1 projectors. These rank-1 projectors — we will refer to them as *clauses* — can be either genuinely quantum clauses, or classical clauses which exclude a particular occupation of the two modes. When *restricted to a two-qubit space*, a genuinely quantum clause can exclude support on a state such as $\alpha|00\rangle + \beta|11\rangle$ (with $|\alpha|, |\beta| \neq 0$) which we call a $\Pi_e^{02,q}$ clause, or exclude support on a state such as $\alpha|01\rangle + \beta|10\rangle$ (with $|\alpha|, |\beta| \neq 0$) which we call a $\Pi_e^{1,q}$ clause. Since projectors in FERMIONIC 2-SAT are parity conserving by definition, these are the only allowed quantum clauses besides the classical clauses. Note that we mention the caveat *when restricted to a two-qubit space* since fermionic operators when viewed as acting on an n-qubit space are generally non-local due to Jordan-Wigner strings.

Let us discuss the ideas behind the proofs of our main results of Section 2.1.1. Given the graph G, we first remove all classical clauses, leaving us possibly with a set of disconnected graphs which we call quantum clusters, see e.g. Figure 2.1 in Section 2.2.3. We will refer to modes that are not involved in any quantum clauses as classical modes. It turns out to be useful to distinguish at least two types of quantum clusters, namely ones that are so-called hidden particle-number-conserving (hPNC) and ones which are not (non-hPNC). The hPNC clusters can be brought to a form in which they only contain $\Pi_e^{1,q}$ quantum clauses by a particle-hole transformation defined in Section 2.2.2: after this transformation, every clause commutes with the cluster particle number operator, hence the terminology. In Section 2.2 we introduce and discuss these various concepts.

In Section 2.3 we start by examining the effect of 'constraint propagation' by $\Pi_e^{1,q}$ and $\Pi_e^{02,q}$ quantum clauses, see Lemma 2.3.1 and its Corollary 2.3.2. This leads up to a characterization of the satisfying assignments: wlog, they are of — what we call — *cluster-product form*, see Proposition 2.3.3. Compared to QUANTUM 2-SAT which, except for on some rank-3 projectors, always has a product satisfying assignment (if any) [3], this form is more general, but also restricted in an interesting way. For assignments with or without overall fixed parity, we will show that one can limit oneself to cluster-product states which are products of operators (acting on the vacuum) which create (1) classical occupations on so-called classical modes, (2) Gaussian states on hPNC clusters, (3) Gaussian states on non-hPNC clusters which are lines or loops, and (4) non-Gaussian states on certain 4-fermion non-hPNC clusters. To obtain this result, we proceed in Section 2.3.3 to exclude certain quantum numbers, —these are hidden cluster particle number or cluster parity which uniquely fix a cluster satisfying assignment—, on quantum clusters. Crucial in this is Lemma 2.3.4 which provides a uniquely fermionic simplification, which leaves us to analyze quantum clusters which are lines and loops in Section 2.3.3.1. Fermions on a line do map directly onto qubits, hence we can use some proof techniques borrowed from QUANTUM 2-SAT investigations. Importantly, these results allow one to efficiently (with effort linear in the number of modes and clauses) verify which quantum numbers, i.e. hidden cluster particle number and cluster parity, are allowed on quantum clusters in the cluster-product form. This is expressed in Lemma

2.3.10.

To prove Theorems 2.1.4 and 2.1.5 in Section 2.3, we bring back all classical clauses in the graph G and reduce the problem of deciding FERMIONIC 2-SAT after several preprocessing steps to deciding a classical 2-SAT instance (thereby proving Theorem 2.1.4). For Theorem 2.1.5, there is the additional constraint of asking for a satisfying assignment of fixed parity. In Appendix 2.C, we provide an efficient algorithm for deciding classical 2-SAT where one asks for a solution of fixed Hamming weight parity (Theorem 2.C.1), which is used as a subroutine in the algorithm for the fermionic problem (Theorem 2.1.5). This result on parity-constrained classical 2-SAT may not be surprising, but we are unaware of any previous algorithm or proof. Crucial in the reductions in the proofs of Theorems 2.1.4 and 2.1.5 is the fact that non-classical assignments on quantum clusters heavily constrain how classical clauses which straddle the quantum cluster and a classical mode can be satisfied, namely the clause has to be true by only the choice of the classical mode, since the non-classical assignment requires that each mode inside the cluster should be allowed to be both empty (0) *and* filled (1). The nature of the hPNC and non-hPNC clusters asks for a different treatment in the proofs of these theorems.

In Section 2.5, we examine the complexity of particle-number-conserving FERMIONIC 2-SAT where we ask for an assignment with given fixed particle number, and prove Theorem 2.1.6. Given the efficiently checkable cluster-product form, it may not be surprising that this problem is contained in NP. It can be proved to be NP-complete since the class of problems contains the weighted 2-SAT problem which asks for a 2-SAT satisfying assignment with fixed Hamming weight, as a special case. In Section 2.6, we prove a quantum hardness result for FERMIONIC k-SAT, still leaving a considerable gap between our efficient classical algorithms for $k = 2$ and this hardness result for $k = 9$. We discuss open problems in the Discussion Section 2.7.

2.1.2.1. Differences with Quantum 2-SAT

We take a paragraph to reflect on the difference between our FERMIONIC 2-SAT findings and QUANTUM 2-SAT.

For QUANTUM k-SAT and FERMIONIC k-SAT, the projectors act nontrivially on k qubits resp. involve k fermionic modes. Fermion-to-qubit mappings allow for the projectors of a given instance of FERMIONIC k-SAT to be mapped to qubit projectors. However, each fermionic projector involving k modes gets mapped —in a general setting— onto a q-qubit projector, where $q \geq k$ [8]. Therefore, one cannot deduce the complexity of FERMIONIC k-SAT directly from the complexity of QUANTUM k-SAT.

Central to the solution of QUANTUM 2-SAT is the fact that a satisfying assignment, if it exists, can be taken to be a product state (except for two-qubit projectors whose null-space is a single entangled state), see e.g. Theorem 4 in [3]. We do not have such a statement for FERMIONIC 2-SAT. In fact, we find that satisfying assignments on non-hPNC clusters can include small 4-fermion non-Gaussian states. Furthermore, for FERMIONIC 2-SAT, quantum clusters with vertices with degree at least 3 obey specific constraints which relate to uniquely fermionic signs, see Lemma 2.3.4 in Section 2.3.3, excluding many particle numbers (but leaving only Gaussian states for hPNC clusters). In contrast, for quantum 2-SAT, product state assignments are generally built from superpositions of all particle numbers. Indirectly, Lemma 2.3.4 also limits the use

of any possible non-Gaussian assignment: non-Gaussian states are only *sometimes* needed on non-hPNC clusters of at most 4 fermionic modes (and any state of three or fewer fermionic modes is always Gaussian [9]). In Appendix 2.B we work through this 4-fermion non-Gaussian example explicitly, showing how the departure from product state assignments for QUANTUM 2-SAT comes about.

2.2. Preliminaries

In this section we briefly review fermionic language and concepts (see Section 1.5 for a more detailed discussion), and introduce several tools for the treatment of assignments on quantum clusters.

For a subset $S \subseteq [n]$ and $|S| = k$, we define the operator $a_S^\dagger := a_{i_1}^\dagger a_{i_2}^\dagger \dots a_{i_k}^\dagger$, where $i_1 < i_2 < \dots < i_k$. This ordering of the fermionic modes will later be conveniently chosen depending on the FERMIONIC 2-SAT instance in Section 2.2.3. We can write any n-mode fermionic state as

$$|\psi\rangle = \sum_{S\subseteq[n]} \alpha_S |S\rangle\,, \; |S\rangle := a_S^\dagger |\text{vac}\rangle\,, \tag{2.1}$$

where $\alpha_S \in \mathbb{C}$ is the amplitude for state $|S\rangle$. Each state $|S\rangle$ corresponds to a configuration of mode *occupations* $\{x_i\}_{i=1}^n$, i.e. the bit $x_i = 1$ (resp. $x_i = 0$) signifies that mode i is occupied, so $i \in S$, (resp. unoccupied, so $i \notin S$). When we speak of a *classical assignment* (see also Definition 1.5.6) in the remainder of this chapter, we mean a state of type $|S\rangle = a_S^\dagger |\text{vac}\rangle$.

One can map an (ordered) set of fermionic creation and annihilation operators via the Jordan-Wigner (JW) transformation onto a set of qubit operators (see also equation (1.30)):

$$a_i^\dagger \overset{J.W.}{=} Z_1 \dots Z_{i-1} \sigma_i^+, \tag{2.2}$$

and the fermionic vacuum state $|\text{vac}\rangle \to |00\dots0\rangle$, using the definition $\sigma_i^+ = |1\rangle\langle 0|_i$ (and similarly $\sigma_i^- = |0\rangle\langle 1|_i$).

Let us note that the n-mode states with particle number $\hat{N} = 0$ (vacuum), $\hat{N} = 1$ (single-particle "Slater determinant" state), $\hat{N} = n-1$ (single-hole state) and $\hat{N} = n$ (all-filled) are Gaussian states (see Definitions 1.5.3 and 1.5.6). The classical states $|S\rangle$ are — of course — also examples of Gaussian states.

2.2.1. Characterization of the Fermionic 2-SAT projectors

Let $G = (V, E)$ be a graph such that the vertices $v \in V$ label the modes ($|V| = n$) and each edge is associated with a projector $\{\Pi_e\}_{e\in E}$. Since the projectors in FERMIONIC 2-SAT are parity conserving, any projector $\Pi_{e=(j,k)}$ is a polynomial in $a_j, a_j^\dagger, a_k$ and $a_k^\dagger$ of even degree. Obviously, since each projector Π_e on modes j and k commutes with the overall parity operator $\hat{P}$, it also commutes with the edge parity $P_{jk} = (-1)^{a_j^\dagger a_j + a_k^\dagger a_k}$. Similarly, if Π_e commutes with $\hat{N}$, it commutes with the edge particle number $\hat{N}_{jk} = \hat{N}_j + \hat{N}_k$. Hence, if we consider a two-mode projector in isolation, it can be represented as a QUANTUM

2-SAT projector on two qubits which is parity-conserving (or, in special cases, also particle number-conserving). This means that if the projector, say of rank-1 in the 2-qubit space, projects onto a product state, the product state can only be a computational basis state. And if it projects onto an entangled state, then in the number-conserving case, it can only project onto states such like $\alpha|01\rangle + \beta|10\rangle$. In the merely parity-conserving case, it can also project onto a state $\alpha|00\rangle + \beta|11\rangle$ for some α, β. This picture is entirely preserved when we consider the action of the fermionic projectors onto n modes. In principle, one can apply a Jordan-Wigner transformation, see equation (2.2), and represent these fermionic projectors as 2-qubit projectors to which strings of Pauli Z's are appended, so that the qubit projectors are not 2-local, but this does not affect the number- or parity-conserving properties. In some sense, we thus consider a simplified form of QUANTUM 2-SAT due to number- and parity-conserving constraints, but the fermionic nature of the problem, i.e. the effect of the Pauli Z-strings through the Jordan-Wigner transformation, makes the problem genuinely different from QUANTUM 2-SAT.

Let's now discuss the nature of rank-1 projectors[1] in fermionic language. With the two-qubit picture in mind, we have two different types of rank-1 projectors for an edge $e = (j,k)$ (with convention $j < k$). The only possible rank-1 projector onto an odd-parity state is

$$\Pi_e^1 = |\beta_e|^2 a_j^\dagger a_j a_k a_k^\dagger + \beta_e \gamma_e^* a_j^\dagger a_k - \gamma_e \beta_e^* a_j a_k^\dagger + |\gamma_e|^2 a_j a_j^\dagger a_k^\dagger a_k \\ \rightarrow \text{ "project onto/exclude } \beta_e|10\rangle + \gamma_e|01\rangle\text{"}, \quad (2.3)$$

with $|\beta_e|^2 + |\gamma_e|^2 = 1$. When Π_e^1 does not project onto just "01" or "10", we call this a genuinely quantum clause and write $\Pi_e^{1,q}$, otherwise we refer to them as classical clauses $\Pi_e^{1,c}$. The only type of rank-1 projector onto an even-parity state is Π_e^{02}:

$$\Pi_e^{02} = |\alpha_e|^2 a_j a_j^\dagger a_k a_k^\dagger - \alpha_e \delta_e^* a_j a_k + \delta_e \alpha_e^* a_j^\dagger a_k^\dagger + |\delta_e|^2 a_j^\dagger a_j a_k^\dagger a_k \\ \rightarrow \text{ "project onto/exclude } \alpha_e|00\rangle + \delta_e|11\rangle\text{"}, \quad (2.4)$$

with $|\alpha_e|^2 + |\delta_e|^2 = 1$. Again, when Π_e^{02} does not project onto just "00" or "11", it is a genuinely quantum clause and we write $\Pi_e^{02,q}$, otherwise it is a classical projector $\Pi_e^{02,c}$.

For convenience, we introduce the following shorthand for a given quantum clause on an edge e:

$$u_e := \gamma_e^*/\beta_e^* \in \mathbb{C}, \text{ for } \Pi_e^{1,q}, \\ v_e := \alpha_e^*/\delta_e^* \in \mathbb{C}, \text{ for } \Pi_e^{02,q}, \quad (2.5)$$

which we shall only use later on when $\beta_e, \delta_e \neq 0$, i.e. in Lemmas 2.3.1, 2.3.4, 2.3.7 and 2.3.8.

[1] They have rank-1 only in the 2-mode subspace, not in the n-mode fermionic space, but for convenience we use this term throughout this paper. Same for statements on higher rank projectors which have a rank $r \leq 4$ in the 2-mode subspace.

For PNC FERMIONIC 2-SAT, we only have projectors of type Π_e^1 in equation (2.3), and only classical projectors

$$\begin{aligned}\Pi_e^0 = a_j a_j^\dagger a_k a_k^\dagger \quad &\rightarrow \text{"project onto/exclude } |00\rangle\text{"},\\ \Pi_e^2 = a_j^\dagger a_j a_k^\dagger a_k \quad &\rightarrow \text{"project onto/exclude } |11\rangle\text{"}.\end{aligned} \tag{2.6}$$

Proposition 2.A.1 in Appendix 2.A proves that these projectors are invariant under taking any unitary linear combination of the two modes j and k, but this understanding also follows from the two-qubit picture above.

2.2.1.1. Higher-rank projectors

The projectors appearing in FERMIONIC 2-SAT are of rank at most three and can be obtained by summing the unit-rank projectors above as long as these project onto orthogonal states. This is true since the 1-eigenvalue sub-space of a projector Π_e is spanned by eigenstates of the commuting edge parity operator $\hat{P}_{e=(j,k)}$, or, in the particle-number-conserving case, by eigenstates of the commuting edge number operator $\hat{N}_{e=(j,k)}$. By viewing the action of the rank-1 projectors in the two-qubit space, it is clear that there exist two orthogonal projectors in the $\{|01,10\rangle\}$ space and two orthogonal projectors in the $\{|00\rangle, |11\rangle\}$ space.

Observe that a rank-2 projector which is the sum of two orthogonal $\Pi_e^{02,q}$ clauses is a rank-2 classical projector (a sum of two classical clauses, excluding both "00" and "11"), and similarly a rank-2 projector which is the sum of two orthogonal $\Pi_e^{1,q}$ clauses is classical (excluding both "01" and "10"). We will call an edge *a classical edge* ($\in E_c$) when its at-most-rank-3 projector is constructed from rank-1 classical clauses, otherwise we call it a quantum edge ($\in E_q$), thus the edgeset can be written as $E = E_c \cup E_q$. Hence, an edge in E_q consists of one — $\Pi_e^{02,q}$ *or* $\Pi_e^{1,q}$ — or two — $\Pi_e^{02,q}$ *and* $\Pi_e^{1,q}$ — quantum clauses. Otherwise, the edge belongs to E_c. It is worthwhile to note that a rank-2 projector $\Pi_e^{1,q} + \Pi_e^{02,q}$ only consists of quadratic terms in creation and annihilation operators, i.e. the quartic contributions cancel as the quartic contribution $\propto a_j^\dagger a_j a_k^\dagger a_k$ of each quantum clause is independent of the choice of $\alpha_e, \beta_e, \gamma_e, \delta_e$ and opposite in sign for a $\Pi_e^{02,q}$ clause versus a $\Pi_e^{1,q}$ clause as can be seen from equation (2.3) and equation (2.4). This fact will be used in Lemma 2.3.9.

2.2.2. Particle-hole transformation

We define the unitary particle-hole transformation K_S on a subset $S \subseteq [n]$ of the n modes. The transformation interchanges the creation of a particle in S with the annihilation of a particle, i.e. the creation of a hole or absence of a particle, in S:

$$\begin{aligned}\forall\, j \in S \quad \rightarrow \quad K_S a_j K_S^{-1} &= \Big[\prod_{\substack{k<j\\ \text{s.t. } k\in S}}(-1)\Big]\, a_j^\dagger, \qquad K_S a_j^\dagger K_S^{-1} = \Big[\prod_{\substack{k<j\\ \text{s.t. } k\in S}}(-1)\Big]\, a_j,\\ \forall\, j \notin S \quad \rightarrow \quad K_S a_j K_S^{-1} &= \Big[\prod_{\substack{k<j\\ \text{s.t. } k\in S}}(-1)\Big]\, a_j, \qquad K_S a_j^\dagger K_S^{-1} = \Big[\prod_{\substack{k<j\\ \text{s.t. } k\in S}}(-1)\Big]\, a_j^\dagger.\end{aligned} \tag{2.7}$$

K_S also transforms the vacuum state $|\text{vac}\rangle$, i.e. $|\text{vac}\rangle \to K_S|\text{vac}\rangle$, where the new vacuum state corresponds to the all-filled state for the modes in the set S, and the original vacuum for the modes which are not in the set S. The $\pm$ signs in the transformation K_S in equation (2.7) can be understood by performing a Jordan-Wigner transformation which leads to the manifestly-unitary transformation $K_S = \bigotimes_{k\in S} X_k$, with X_k denoting Pauli-X on qubit k. When $S = \emptyset$, $K_{S=\emptyset} = I$.

When we apply the particle-hole transformation K_S to the genuinely quantum projectors in equation (2.3) and equation (2.4), it can interchange $\Pi_e^{1,q}$ clauses and $\Pi_e^{02,q}$ clauses, depending on the edge e and the subset S. In Appendix 2.A we give the effect of the transformation in mathematical detail (which is not extremely insightful).

An additional property of the particle-hole transformation K_S is that it preserves Gaussianity of a state ρ. We can express the unitary transformation K_S on the Majorana operators $c_{2i-1} \overset{J.W.}{=} Z_1 \dots Z_{i-1} X_i, c_{2i} \overset{J.W.}{=} Z_1 \dots Z_{i-1} Y_i$ (consistent with equation (2.2) and equation (1.18)). Then K_S is some product of Majorana operators $\Pi_{i\in S_{\text{Maj}}} c_i$, using $X_k = \left[\prod_{1\le j<k}(-ic_{2j-1}c_{2j})\right]c_{2k-1}$. We have $K_S c_k K_S^{-1} = \pm c_k$, with the sign depending on whether $k \in S_{\text{Maj}}$ and $|S_{\text{Maj}}|$, inducing a (diagonal) orthogonal transformation in the $\{c_i\}_{i=1}^{2n}$ basis. Any orthogonal transformation $R \in O(2n)$ preserves the Gaussianity of a pure fermionic state as it preserves the property that the covariance matrix Γ still obeys $\Gamma^T\Gamma = I$ [9], hence K_S preserves Gaussianity. Note that $K_S\hat{P}K_S^{-1} = \pm\hat{P}$, switching the overall parity P when $|S_{\text{Maj}}|$ is odd.

2.2.3. Quantum clusters

The graph $G = (V, E)$ describing the fermionic 2-SAT problem (where an edge represents a rank-1, 2 or 3 projector) can contain zero, one or more *quantum clusters*, see Fig. 2.1, defined as follows.

Definition 2.2.1 (Quantum clusters)**.** *Take the subgraph G_{sub} of G obtained by removing all classical clauses in E_c, and removing all vertices that are not touching any $e \in E_q$. We define the quantum clusters $G_q^i = (V_q^i, E_q^i)$ (for $i = 1, \dots, Q$, and with $n_q^i := |V_q^i|$) to be the connected components of G_{sub}. The cluster particle number equals $\hat{N}_q^i = \sum_{j\in V_q^i} a_j^\dagger a_j$ and the cluster parity is $\hat{P}_q^i = (-1)^{\sum_{j\in V_q^i} a_j^\dagger a_j}$.*

If the quantum cluster contains only $\Pi_e^{1,q}$ clauses, we can call the cluster particle-number-conserving, and states with a fixed cluster particle number span the set of possible satisfying assignments in the cluster. Some clusters are not directly particle-number conserving, but can be brought to this form by a unitary particle-hole transformation defined in Section 2.2.2. We call these clusters *hidden particle number conserving* (hPNC), see also Fig. 2.2.

Definition 2.2.2 (Hidden Particle Number Conserving Cluster (hPNC))**.** *A quantum cluster G_q on the set of vertices V_q is hidden particle number conserving (hPNC) if we can divide the set of vertices into two sets, $V_q = A \cup B$, and all quantum clauses connecting vertices from A with vertices from B are $\Pi_e^{02,q}$ clauses, while all quantum clauses connecting vertices in either A or B internally, are $\Pi_e^{1,q}$ clauses. A special case is*

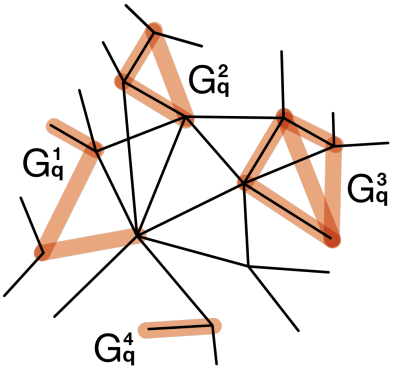

Figure 2.1: Several connected subgraphs G_q^i of G with genuinely quantum clauses (orange), the other clauses (black) are classical clauses. The graph G_{sub} is obtained by removing all classical edges and vertices not touching quantum edges (referred to as classical modes), leaving the disconnected quantum clusters G_q^i.

when there are no $\Pi_e^{02,q}$ clauses whatsoever in the quantum cluster —A is the set of all vertices, and $B = \emptyset$— and G_q is particle number conserving (PNC).

Observe that in a hPNC cluster G_q, one can assume wlog that all edges are single clauses (rank-1): if they are of rank-2 or higher, than either, they are classical edges (hence they had been removed from the graph G, hence cluster G_q), or they are a sum of a $\Pi_e^{02,q}$ clause and a $\Pi_e^{1,q}$ clause, and hence the cluster cannot be hPNC. Note that if we have the cluster parity operator $\hat{P}_q$ for n_q modes, then the particle-hole transformation in equation (2.7) for a subset $S \subseteq [n_q]$, transforms $K_S \hat{P}_q K_S^{-1} = \pm \hat{P}_q$: hence an eigenstate of cluster parity is still an eigenstate of cluster parity after a particle-hole transformation, but possibly with a flipped eigenvalue.

In order to efficiently determine whether G_q is hPNC and examine some other properties later, it will be useful to define a “maximal spanning" subgraph T_q which is hPNC as follows. The idea is that this (non-unique) subgraph T_q can be obtained via a greedy approach and can be used to efficiently decide whether a quantum cluster G_q is hPNC.

Definition 2.2.3 (Maximal spanning hPNC subgraph of a quantum cluster)**.** *Given a quantum cluster G_q (with n_q vertices and $|E_q|$ quantum clauses/edges), the graph T_q is a graph with the same vertices as G_q but a subset of quantum clauses of G_q such that (1) T_q is hPNC, (2) The quantum clauses in T_q generate a spanning tree for the graph G_q (i.e. all vertices in G_q are connected by some quantum clauses in T_q) and (3) T_q is maximal, i.e. it becomes non-hPNC when adding any more clauses from G_q. Such (generally non-unique) T_q can be constructed with effort $O(n_q + |E_q|)$ as follows.*

→ Start with T_q containing all n_q vertices, but no clauses. Pick a vertex and assign it to A.
→ For each vertex adjacent to the current one, add it to A if the connecting clause is of type

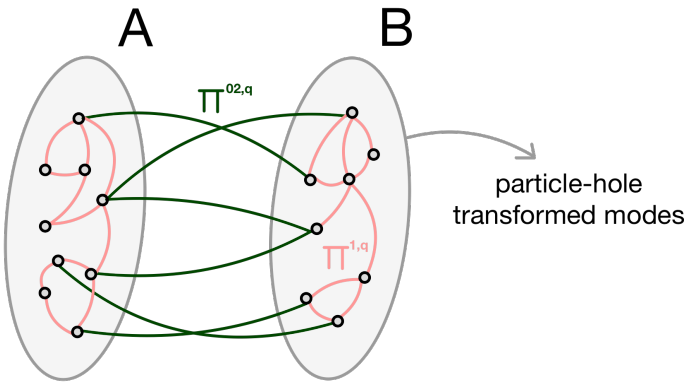

Figure 2.2: Example of a hidden particle number conserving quantum cluster G_q as given in Definition 2.2.2 where the green quantum edges are $\Pi_e^{02,q}$ clauses and the pink quantum edges are $\Pi_e^{1,q}$ clauses. If one does a particle-hole transformation K_B on all the fermionic modes in B (or K_A on A), the green $\Pi_e^{02,q}$ clauses become $\tilde{\Pi}_e^{1,q}$ clauses, while the pink $\Pi_e^{1,q}$ clauses become $\tilde{\Pi}_e^{1,q}$ clauses, see equation (2.30), thus staying of particle-conserving type. Hence, after the particle-hole transformation, all clauses in G_q are particle-number-conserving, and the space of satisfying assignments is then spanned by eigenstates $|\psi\rangle$ of the so-called *hidden cluster particle number* $\hat{N}_q$, obeying $\hat{N}_q(K_B|\psi\rangle) = N_q(K_B|\psi\rangle)$. Note that for a hPNC cluster with $B \neq \emptyset$, a satisfying assignment $|\psi\rangle$ will generally not be an eigenstate of cluster particle number in the original basis, but an eigenstate of the cluster parity.

$\Pi_e^{1,q}$, and add it to B if it is of type $\Pi_e^{02,q}$ (if it is the sum of $\Pi_e^{1,q}$ and $\Pi_e^{02,q}$, pick one of them at random). Add these connecting clauses to T_q and move to the neighbor vertex.
→ Repeat this procedure for all vertices adjacent to the previous vertices (if they have not yet been assigned to A or B) until all vertices are in the bipartition $A \cup B$.
→ Then, for all clauses in G_q which are not yet in T_q, do the following. Add the clause to T_q if it is a $\Pi_e^{02,q}$ clause and the vertices at its ends are in opposite sets (i.e., one in A and the other in B). Conversely, include the clause in T_q if it is a $\Pi_e^{1,q}$ clause and the vertices at its ends have equal assignments.

If after this efficient procedure, no clauses are left in G_q, then obviously $T_q = G_q$ and G_q is hPNC.

For later reference, we state a simple fact about a choice for picking the starting vertex in T_q:

Fact 2.2.4. *Suppose a quantum cluster G_q contains a vertex with degree at least 3, then a maximal spanning hPNC subgraph T_q can be efficiently constructed such that it also contains a vertex of degree at least 3, simply by choosing the initial vertex in the greedy approach in Definition 2.2.3 to be the vertex in G_q of degree at least 3.*

2.3. Characterizing the satisfying assignments of Fermionic 2-SAT

2.3.1. Properties of satisfying assignments

The following 'Partner' Lemma captures the idea of *propagation of a constraint* by genuinely quantum clauses: recall that in the qubit case, a rank-1 *entangled* projector is always 'propagating', see Definition 1 and Lemma 5 in [3]: the assignment of a product state for one qubit uniquely fixes the assignment of a product state for the other qubit, given that the state has to be orthogonal to the entangled projector. Here propagation is a more general concept as we can't assume a product state assignment. The Lemma roughly means that if a satisfying assignment for the whole graph G has support on a certain n-mode state, it also needs to have support on a 'partnered' n-mode state.

Lemma 2.3.1 (Partner Lemma)**.** *Given an edge $e=(j,k)$ (with $j<k$), with (1) clause $\Pi_e^{1,q}$ or (2) $\Pi_e^{02,q}$ (or both) in an instance of* FERMIONIC 2*-SAT. Suppose that $|\psi\rangle$ is an n-mode satisfying assignment.*

1. *$\Pi_e^{1,q}$ clause $\Rightarrow$ If $|\psi\rangle$ has support on a $|S\rangle = a_S^\dagger|\mathrm{vac}\rangle$ which is given by occupations $\{x_i\}_{i=1}^n$ such that $x_j \neq x_k$, then $|\psi\rangle$ must also have support on a 'partnered' state $|S'\rangle$ with occupations $\{x'_i\}_{i=1}^n$ with $x'_j=\overline{x_j}, x'_k=\overline{x_k}, x'_{i\neq j,k}=x_i$. In particular, we must have for the amplitudes α_S in the satisfying assignment:*

$$\begin{aligned}\alpha_S &= -sign(k,S')/sign(j,S)\, u_e\, \alpha_{S'},\ \text{if } x_j=1 \text{ and } x_k=0,\\ \alpha_{S'} &= -sign(k,S)/sign(j,S')\, u_e\, \alpha_S,\ \text{if } x_j=0 \text{ and } x_k=1,\end{aligned} \tag{2.8}$$

   where $sign(l,T)$ for $l\in T$ denotes the sign arising from the reordering of $a_T^\dagger$ to $a_l^\dagger a_{T\setminus l}^\dagger$ and u_e as in equation (2.5) *for $\Pi_e^{1,q}$.*

2. *$\Pi_e^{02,q}$ clause $\Rightarrow$ If $|\psi\rangle$ has support $|S\rangle$ with occupations $\{x_i\}_{i=1}^n$ such that $x_j=x_k$, then $|\psi\rangle$ must also have support on $|S'\rangle$ with $\{x'_i\}_{i=1}^n$ with $x'_j=\overline{x_j}, x'_k=\overline{x_k}, x'_{i\neq j,k}=x_i$. In particular, we must have for the amplitudes α_S in the satisfying assignment:*

$$\begin{aligned}\alpha_S &= -sign(jk,S)\, v_e\, \alpha_{S'},\ \text{if } x_j=x_k=1,\\ \alpha_{S'} &= -sign(jk,S')\, v_e\, \alpha_S,\ \text{if } x_j=x_k=0,\end{aligned} \tag{2.9}$$

   where $\mathrm{sign}(lm,T)$ *for $l,m\in T$ denotes the sign arising from reordering of $a_T^\dagger$ to $a_l^\dagger a_m^\dagger a_{T\setminus\{l,m\}}^\dagger$ and v_e as in equation* (2.5) *for $\Pi_e^{02,q}$.*

Proof. Consider case 1. Since $|\psi\rangle=\sum_{T\subseteq[n]}\alpha_T a_T^\dagger|\mathrm{vac}\rangle$ is a satisfying assignment, we must have $\langle\psi|\Pi_e^{1,q}|\psi\rangle=0$. Using equation (2.3) for $\Pi_{e=(j,k)}^{1,q}$ $(j<k)$, this implies

$$\langle\psi|\Pi_e^{1,q}|\psi\rangle = \sum_{T\subseteq[n]\ \text{s.t.}\ j\in T,\,k\notin T} \left|\,\mathrm{sign}(j,T)\,\beta_e^*\alpha_T + \mathrm{sign}(k,T')\,\gamma_e^*\alpha_{T'}\right|^2 = 0, \tag{2.10}$$

where for each subset $T=(x_1,\ldots,x_j=1,\ldots,x_k=0,\ldots,x_n)$, the partnered subset T' is defined as $(x_1,\ldots,x_j=0,\ldots,x_k=1,\ldots,x_n)$. The relation between α_S and $\alpha_{S'}$ follows from

equation (2.10). Consider case 2. To ensure $\langle\psi|\,\Pi_e^{02,q}\,|\psi\rangle = 0$ with $\Pi_e^{02,q}$ in equation (2.4), we must have that

$$\langle\psi|\,\Pi_e^{02,q}\,|\psi\rangle = \sum_{T\subset[n]\text{ s.t. }j,k\in T} \left|\,\alpha_e^*\alpha_{T'} + \text{sign}(jk,T)\,\delta_e^*\alpha_T\right|^2 = 0, \tag{2.11}$$

where for each subset $T = (x_1,\ldots,x_j=1,\ldots,x_k=1,\ldots,x_n)$, the partnered subset T' is defined as $(x_1,\ldots,x_j=0,\ldots,x_k=0,\ldots,x_n)$. The relation between α_S and $\alpha_{S'}$ follows from equation (2.11).

□

Lemma 2.3.1 strongly restricts the possible satisfying assignments on quantum clusters. The idea is that for a hPNC cluster G_q repeatedly applying the case 1 partner rule in Lemma 2.3.1 makes particles propagate on the connected quantum cluster graph G_q, so that all occupations for a fixed hidden cluster particle number are generated in the cluster. Then given a fixed hidden cluster particle number, there is at most 1 unique satisfying assignment.

If the cluster is non-hPNC, one can argue that there will be some additional $\Pi_e^{02,q}$ clauses, creating and annihilating pairs of particles via case 2, and in that case there is at most 1 unique satisfying assignment given fixed parity (and no satisfying assignments with fixed hidden cluster particle number).

In these arguments, care must be taken not to assume anything about the structure of global assignment, i.e. we derive the necessary structure of such global assignment based on how quantum clauses on each cluster can be satisfied. We capture these ideas in the following Corollary:

Corollary 2.3.2. *Suppose a quantum cluster G_q is hPNC with bipartition $A\cup B$, and $|\psi\rangle$ is a satisfying assignment for the entire graph G containing G_q. If $K_B|\psi\rangle$ (with $K_B = I$ in the strictly PNC case) has support on a n-mode state $|S\rangle$ with hidden cluster particle number N_q on G_q, i.e. $\hat{N}_q|S\rangle = N_q|S\rangle$, then it must have support on all $\binom{n_q}{N_q}$ n-mode states $|S'\rangle$ with particle number N_q on G_q where $|S'\rangle$ has equal occupation as $|S\rangle$ on all other modes in G. Moreover, there is at most one satisfying assignment per hidden cluster particle number $N_q\in\{0,1,\ldots,n_q\}$ on G_q alone. Thus for a hPNC cluster, N_q is the cluster particle number of the particle-hole transformed cluster, which we will refer to as the hidden cluster particle number.*

Suppose a quantum cluster G_q is non-hPNC and $|\psi\rangle$ is a satisfying assignment for G. If $|\psi\rangle$ has support on a n-mode state $|S\rangle$ with cluster parity P_q on G_q, i.e. $\hat{P}_q|S\rangle = P_q|S\rangle$, then it must have support on all $2^{n_q}/2$ n-mode states $|S'\rangle$ with parity P_q on G_q where $|S'\rangle$ has equal occupation as $|S\rangle$ on all other modes in G. There is at most one satisfying assignment per parity $P\in\{-1,+1\}$ on G_q alone.

Proof. The first part of the Corollary follows directly from repeatedly applying 'the propagation rule' of case 1 of Lemma 2.3.1 to an N_q-particle n-mode state $|S\rangle$ on which $K_B|\psi\rangle$ is supported, propagating the particles (only for $N_q = 0$ (all empty) and $N_q = n_q$ there is nothing to propagate and case 1 does not apply). The coefficient relation equation (2.8) implies that there is at most one satisfying assignment per hidden cluster

particle number $N_q \in \{0,1,\ldots,n_q\}$ on the vertices of G_q alone. As for the second part of the corollary, consider a non-hPNC cluster and construct the maximal spanning hPNC subgraph T_q in Definition 2.2.3 with bipartition $A \cup B$. If we were to apply the particle-hole transformation K_B on T_q, any satisfying assignment $|\psi\rangle$ for the projectors in this subgraph T_q would be such that $K_B|\psi\rangle$ is a unique state for fixed cluster particle number as per the first part of this corollary. Since G_q is non-hPNC, using this bipartition of T_q, there will be (a) at least one internal $\Pi_e^{02,q}$ clause inside either A or B, or (b) one $\Pi_e^{1,q}$ clause connecting A and B. If we apply the particle-hole transformation K_B, then in case (a) such $\Pi_e^{02,q}$ clause stays a $\tilde{\Pi}_e^{02,q}$ "pair-creating" clause, implying that any satisfying assignment $|\psi\rangle$ for the quantum clauses in G_q has the property that $K_B|\psi\rangle$ is a unique state given some cluster parity: this follows from case 2 in Lemma 2.3.1 and states of any particle number with that cluster parity must occur in the superposition (working in the particle-hole transformed basis). If $K_B|\psi\rangle$ is a state which involves all particle numbers with given parity, then so is the satisfying assignment $|\psi\rangle$, since K_B maps any state $|S\rangle$ onto some other state $|S'\rangle$. In case (b) upon the particle-hole transformation K_B, the connecting clause $\Pi_e^{1,q}$ transforms to "pair-creating" clause $\tilde{\Pi}_e^{02,q}$ and then the same arguments apply as in case (a). □

2.3.2. Cluster-product form of satisfying assignments

Let's examine the consequences of Corollary 2.3.2 for satisfying assignments for the entire graph G. Let hPNC be the set of hPNC quantum clusters, let non − hPNC be the set of non-hPNC quantum clusters and let Class be the collection of remaining modes not contained in any quantum cluster, to which we have referred as classical modes. Let's us fix an ordering of the fermionic modes in the graph G using these sets hPNC, non − hPNC and Class and some chosen internal ordering of modes inside the sets. Say, we first assign labels $1,2,\ldots,n_q^1$ to the modes in the first hPNC quantum cluster G_q^1, then labels $n_q^1+1, n_q^1+2,\ldots,n_q^1+n_q^2$ to the modes in the second hPNC quantum cluster G_q^2, etc., then do the same for the non-hPNC quantum clusters, and then finally for all the classical modes. What is relevant is that in this ordering, we order the modes in each cluster immediately following each other, so the satisfying assignment of the cluster itself can be used in the assignment for the whole graph in so-called cluster-product form. Given this ordering, Corollary 2.3.2 implies that *any* satisfying assignment $|\psi\rangle$ must be a superposition of states of ordered cluster-product form, i.e.

$$|\psi\rangle = \sum_{\{N_q^i\},\{P_q^j\},\{x_k\}} \beta_{\{N_q^i\},\{P_q^j\},\{x_k\}} \left(\prod_{k\in\text{Class}} O_k(x_k)\right)\left(\prod_{j\in\text{non-hPNC}} O_j(P_q^j)\right)\left(\prod_{i\in\text{hPNC}} O_i(N_q^i)\right)|\text{vac}\rangle\,, \tag{2.12}$$

Here the operator $O_i(N_q^i)$ creates the unique state with hidden cluster particle number N_q^i using the modes of hPNC quantum cluster i, the operator $O_j(P_q^j)$ creates the unique

state with cluster parity [2] $P_q^j = \pm 1$ using the modes in the non-hPNC cluster j, and the operator $O_k(x_k = 0) = I$, $O_k(x_k = 1) = a_k^\dagger$ for a classical mode $k \in \mathsf{Class}$. Note that the operators $O_k(x_k)$, $O_j(P_q^j)$, $O_i(N_q^i)$ do not generally commute. We note that for a hPNC cluster i, non-hPNC cluster j and $k \in \mathsf{Class}$, some values for respectively N_q^i, P_q^j and x_k in equation (2.12) may be excluded in order to be a satisfying assignment. For example, this can be due to further constraints on parity or particle number which we will examine in the next Section 2.3.3. We only consider the constraints due to classical clauses in the proofs of Theorems 2.1.4 and 2.1.5 in Section 2.4.

Let us now argue that taking a superposition of cluster-product states as in Eq. (2.12) is unnecessary, i.e. individual cluster-product states suffice. For this, we first examine some basic properties of the cluster-product form. For any $i_0 \in \mathsf{hPNC}$, the cluster-product form is an eigenstate of the hidden cluster particle number $\hat{N}_q^{i_0}$ (where $K_B^{i_0}$ applies the particle-hole transformation for cluster i_0), i.e.

$$\begin{aligned}
&\hat{N}_q^{i_0}\left(K_B^{i_0}\left(\Pi_{k\in\mathsf{Class}}O_k(x_k)\right)\left(\Pi_{j\in\mathsf{non-hPNC}}O_j(P_q^j)\right)\left(\Pi_{i\in\mathsf{hPNC}}O_i(N_q^i)\right)|\mathrm{vac}\rangle\right) = \\
&\left(\Pi_{k\in\mathsf{Class}}O_k(x_k)\right)\left(\Pi_{j\in\mathsf{non-hPNC}}O_j(P_q^j)\right)\left(\Pi_{i\in\mathsf{hPNC}}[\hat{N}_q^i K_B^i]^{\delta_{i,i_0}}O_i(N_q^i)\right)|\mathrm{vac}\rangle = \\
&N_q^{i_0}\left(K_B^{i_0}\left(\Pi_{k\in\mathsf{Class}}O_k(x_k)\right)\left(\Pi_{j\in\mathsf{non-hPNC}}O_j(P_q^j)\right)\left(\Pi_{i\in\mathsf{hPNC}}O_i(N_q^i)\right)|\mathrm{vac}\rangle\right), \qquad (2.13)
\end{aligned}$$

using equation (2.7), i.e. we can commute these operators through the ordered product. Similarly, for any $j_0 \in \mathsf{non-hPNC}$

$$\begin{aligned}
&\hat{P}_q^{j_0}\left(\left(\Pi_{k\in\mathsf{Class}}O_k(x_k)\right)\left(\Pi_{j\in\mathsf{non-hPNC}}O_j(P_q^j)\right)\left(\Pi_{i\in\mathsf{hPNC}}O_i(N_q^i)\right)|\mathrm{vac}\rangle\right) = \\
&P_q^{j_0}\left(\left(\Pi_{k\in\mathsf{Class}}O_k(x_k)\right)\left(\Pi_{j\in\mathsf{non-hPNC}}O_j(P_q^j)\right)\left(\Pi_{i\in\mathsf{hPNC}}O_i(N_q^i)\right)|\mathrm{vac}\rangle\right), \qquad (2.14)
\end{aligned}$$

and for any mode $k_0 \in \mathsf{Class}$, we have

$$\begin{aligned}
&\hat{N}_q^{k_0}\left(\left(\Pi_{k\in\mathsf{Class}}O_k(x_k)\right)\left(\Pi_{j\in\mathsf{non-hPNC}}O_j(P_q^j)\right)\left(\Pi_{i\in\mathsf{hPNC}}O_i(N_q^i)\right)|\mathrm{vac}\rangle\right) = \\
&N_q^{k_0}\left(\left(\Pi_{k\in\mathsf{Class}}O_k(x_k)\right)\left(\Pi_{j\in\mathsf{non-hPNC}}O_j(P_q^j)\right)\left(\Pi_{i\in\mathsf{hPNC}}O_i(N_q^i)\right)|\mathrm{vac}\rangle\right), \qquad (2.15)
\end{aligned}$$

We will use the following shorthand for the cluster-product form:

$$\begin{aligned}
&\left|\phi\left(\{N_q^i\},\{P_q^j\},\{x_k\}\right)\right\rangle \equiv \\
&\qquad\left(\prod_{k\in\mathsf{Class}}O_k(x_k)\right)\left(\prod_{j\in\mathsf{non-hPNC}}O_j(P_q^j)\right)\left(\prod_{i\in\mathsf{hPNC}}O_i(N_q^i)\right)|\mathrm{vac}\rangle\,. \qquad (2.16)
\end{aligned}$$

These facts give us a final characterization of the satisfying assignments, see Fig. 2.3:

Proposition 2.3.3 (Cluster-product form of satisfying assignments)**.** *If there exists a satisfying assignment of a* FERMIONIC 2*-SAT problem, possibly with fixed particle number or fixed parity, then one can assume, with loss of generality, that it is of cluster-product form* $\left|\phi\left(\{N_q^i\},\{P_q^j\},\{x_k\}\right)\right\rangle$ *in equation* (2.16).

[2] Remember that cluster parity is defined in the original basis while *hidden* cluster particle number is defined in the particle-holed transformed basis in Definition 2.2.1 and Fig. 2.2.

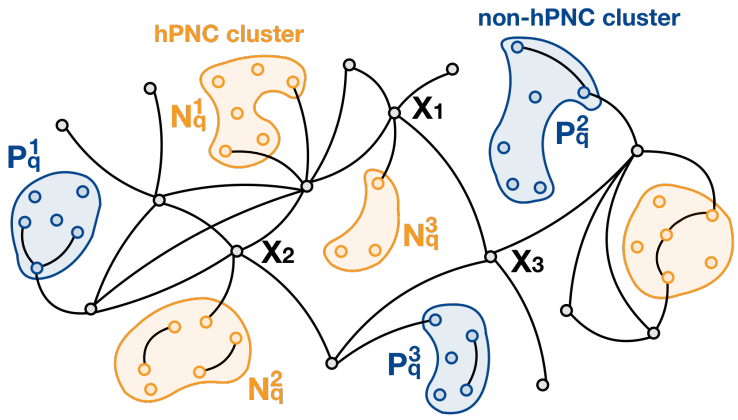

Figure 2.3: Satisfying assignments of FERMIONIC 2-SAT can be taken to be of cluster-product form: ordered products of operators which create satisfying assignments on quantum clusters and classical assignments on classical modes starting from the vacuum state, as in equation (2.16). The possibly-created states on a hPNC cluster $G_q^i \equiv i \in \mathsf{hPNC}$ (orange clusters) are completely specified by (hidden) cluster particle numbers $N_q^i = 0, 1, \ldots, n_q^i$. The possibly-created states on a non-hPNC cluster $j \in \mathsf{non-hPNC}$ (blue clusters) are completely specified by the cluster parity $P_q^j = \pm 1$. The possibly-created states for a classical mode $k \in \mathsf{Class}$ are completely specified by the occupied/unoccupied label $x_k = 0, 1$. The black edges correspond to classical clauses. Note that these classical clauses can also be internal to a quantum cluster or straddle a quantum cluster and a classical mode.

Proof. Given Corollary 2.3.2 the most general solution is of the form in equation (2.12). However, if any superposition state with different cluster parities $\{P_q^j\}$ for non-hPNC clusters is a satisfying assignment, then so is each individual term in the superposition since cross terms with distinct cluster parities must be zero as $\sum_i \Pi_i$ commutes with any cluster parity. Similarly, if any superposition state with different hidden particle numbers $\{N_q^i\}$ for hPNC clusters is a satisfying assignment, then so is each individual term in the superposition since (using manipulations in the particle-hole transformed basis) cross-terms must be 0. Similarly, if any superposition state with different values $\{x_k\}$ is a satisfying assignment, then so is each individual term in the superposition since cross-terms must be 0 as $\sum_i \Pi_i$ commutes with $\hat{N}_k$. Thus wlog we can restrict ourselveves to assignments of cluster-product form $\left|\phi\left(\{N_q^i\}, \{P_q^j\}, \{x_k\}\right)\right\rangle$. □

2.3.3. Excluding certain particle numbers or parities on quantum clusters

In this section we further examine the structure of satisfying assignments, namely the allowed values for (hidden) cluster particle numbers and parities in the cluster-product form in equation (2.16) which will restrict the search for an assignment. The following

Lemma provides an important simplification for quantum clusters with vertices with degree at least 3.

Lemma 2.3.4 (Fermionic Degree ≥ 3 Simplification)**.** *Let $G_q = (V_q, E_q)$ be a quantum cluster on n_q modes (=vertices) with at least one vertex of degree ≥ 3 and let T_q be its maximal spanning hPNC subgraph with partition $A \cup B$ in Definition 2.2.3. Then, for any satisfying n_q-mode assignment $|\psi\rangle$ for G_q, we must have $\langle S| K_B |\psi\rangle = 0$, for any $|S\rangle = a_S^\dagger |\text{vac}\rangle$ with hidden cluster particle number $\hat{N}_q = 2, 3, \ldots, n_q - 2$.*

Proof. Using Fact 2.2.4, the graph T_q contains a vertex j of degree ≥ 3 and let three of its neighbor vertices be labeled k, l, m. We may assume wlog that $j < k < l < m < i_1 < i_2 < \ldots$, with $i_1, i_2, \ldots$ labeling the other modes in the quantum cluster G_q. We write $e = (j,k)$, $e' = (j,l)$ and $e'' = (j,m)$. In the particle-hole transformed basis, all clauses adjacent to j are of type $\Pi_e^{1,q}$ by construction via Definition 2.2.3 and Fact 2.2.4. After the particle-hole transformation we write $\Pi^{1,q}_{e=(j,k)}$, $\Pi^{1,q}_{e'=(j,l)}$ and $\Pi^{1,q}_{e''=(j,m)}$ with non-zero coefficients u_e, $u_{e'}$ and $u_{e''}$ introduced in equation (2.5), uniquely characterizing the clauses.

Suppose that $K_B |\psi\rangle$ has support on a state $|S\rangle = a_S^\dagger |\text{vac}\rangle$ with particle number $N_q = 2, 3, \ldots, n_q - 2$. Lemma 2.3.1 implies that $K_B |\psi\rangle$ must have support on *all* $\binom{4}{2}$ states $|S'\rangle = a_{S'}^\dagger |\text{vac}\rangle$ s.t. in $|S'\rangle$ two modes from $\{j,k,l,m\}$ are occupied, and $|S'\rangle$ has occupation equal to $|S\rangle$ on all modes $V_q \backslash \{j,k,l,m\}$. Let the occupied modes in $V_q \backslash \{j,k,l,m\}$ be denoted by subset T (which depends on S). Let $K_B |\psi\rangle = \sum_S \alpha_S |S\rangle$, using, say, the fermionic mode ordering established at the start of Section 2.3.2, and assume wlog that in this ordering $T = T_{\text{before}} \cup T_{\text{after}}$ where T_{before} is a set of modes which come before j,k,l,m in the chosen fermionic order and T_{after} is a set of modes which come after and let $a_T^\dagger = a_{T_{\text{before}}}^\dagger a_{T_{\text{after}}}^\dagger$. Then $K_B |\psi\rangle$ must have support on each of the six states

$$a_j^\dagger a_k^\dagger a_T^\dagger |\text{vac}\rangle, a_j^\dagger a_l^\dagger a_T^\dagger |\text{vac}\rangle, a_j^\dagger a_m^\dagger a_T^\dagger |\text{vac}\rangle, \\ a_k^\dagger a_l^\dagger a_T^\dagger |\text{vac}\rangle, a_k^\dagger a_m^\dagger a_T^\dagger |\text{vac}\rangle, a_l^\dagger a_m^\dagger a_T^\dagger |\text{vac}\rangle, \quad (2.17)$$

with amplitudes $\alpha_{j,k,T}, \alpha_{j,l,T}, \alpha_{j,m,T}, \alpha_{k,l,T}, \alpha_{k,m,T}$ and $\alpha_{l,m,T}$. Lemma 2.3.1 implies the following relations between these amplitudes:

$$(1)\ \alpha_{j,k,T} = +u_{e'}\,\alpha_{k,l,T}, \quad (2)\ \alpha_{j,l,T} = -u_e\,\alpha_{k,l,T}, \quad (3)\ \alpha_{j,l,T} = +u_{e''}\,\alpha_{l,m,T}, \\ (4)\ \alpha_{j,k,T} = +u_{e''}\,\alpha_{k,m,T}, \quad (5)\ \alpha_{j,m,T} = -u_e\,\alpha_{k,m,T}, \quad (6)\ \alpha_{j,m,T} = -u_{e'}\,\alpha_{l,m,T}. \quad (2.18)$$

However, conditions (1-6) imply $+1 = -1$. Hence, $K_B |\psi\rangle$ cannot have support on n-mode states $|S\rangle = a_S^\dagger |\text{vac}\rangle$ with cluster particle number $\hat{N}_q = 2, 3, \ldots, n_q - 2$. □

We note that the previous Lemma is intrinsically fermionic, i.e. the anti-commutation of creation operators plays a role in the signs of equation (2.18) and there is no QUANTUM 2-SAT counterpart. The following corollary summarizes the consequences.

Corollary 2.3.5. *Consider an assignment of cluster-product form*

$$\left|\phi\left(\{N_q^i\}_{i\in\text{hPNC}}, \{P_q^j\}_{j\in\text{non-hPNC}}, \{x_k\}_{k\in\text{Class}}\right)\right\rangle, \quad (2.19)$$

in equation (2.16)*, which, per Proposition 2.3.3, exists if there is a satisfying assignment for G. We summarize some consequences for the quantum numbers N_q^i and P_q^i in the cluster-product form due to the quantum clauses on clusters.*

1. *For degree* ≥ 3 *hPNC clusters* $i \in \mathsf{hPNC}$*, the only hidden cluster particle numbers can be* $N_q^i = 0, 1, n_q^i - 1, n_q^i$ *which are Gaussian states, remaining Gaussian when undoing the particle-hole transformation as argued in Section 2.2.2.*

2. *For any hPNC cluster* $i \in \mathsf{hPNC}$ *with* n_q^i *modes, the hidden cluster particle number can* always *be* $N_q^i = 0$ *or* $N_q^i = n_q^i$*, since Lemma 2.3.1 imposes no constraints on the coefficients, i.e. particles don't propagate since there are either no particles or all modes are filled.*

3. *For a hPNC cluster* i *with* n_q^i *modes, let there be a cluster product assignment with hidden cluster particle number* $N_q^i \neq 0, N_q^i \neq n_q^i$*. Then the assignment will have support on states* $|S\rangle$ *in which any mode* j *in the cluster is empty* ($x_j = 0$) *and states for which it is filled* ($x_j = 1$)*. Hence in this case any classical clause straddling the cluster on mode* j *and a classical mode* $k \in \mathsf{Class}$ *or mode in another quantum cluster, will have to be satisfied (if it can) no matter the value for* $x_j = 0, 1$*. When* $n_q^i > 2$*, if the state with hidden cluster particle number* $N_q^i \neq 0, N_q^i \neq n_q^i$ *obeys any additional classical clause constraints inside the cluster, then either the all-empty* $N_q^i = 0$ *or all-filled assignment* $N_q^i = n_q^i$ *(or both) should also obey these constraints since on any pair of modes, at least 3 out of 4 occupations on these two modes will be allowed (thus always including either 00 or 11).*

4. *If G contains a degree* ≥ 3 *non-hPNC cluster* $j \in \mathsf{non-hPNC}$ *with a number of modes* $n_q^j \geq 5$*, then G has no satisfying assignment. Via Corollary 2.3.2 the satisfying assignment is characterized by total parity P and has support on states with any hidden cluster particle number with this parity. However, Lemma 2.3.4 implies that the assignment cannot have support on odd hidden cluster particle number* $N_q = 3$ *(in case of odd parity) or* $N_q = 2$ *(in case of even parity) when* $n_q \geq 5$*, hence there cannot be any satisfying assignment.*

5. *For a degree* ≥ 3 *non-hPNC cluster* $G_q \equiv j \in \mathsf{non-hPNC}$ *with* $n_q^j = 4$*,* $K_B |\psi\rangle$ *(with the set B defined through the graph* T_q *of the cluster* G_q*), can only have support on states with odd cluster parity. Undoing the particle hole transformation, this excludes one choice of parity P for this cluster in equation* (2.19) *(which one depends on whether* $|B|$ *is odd or even). In Appendix 2.B we show that the satisfying assignment for such cluster of four fermionic modes is a non-Gaussian fermionic state.*

6. *For a non-hPNC* G_q *cluster with* n_q *modes, let there be an cluster product assignment with some cluster parity* P_q*. Then this assignment will have support on states* $|S\rangle$ *in which any mode* j *in the cluster is empty and states for which it is filled: it has support on all occupations* x_j *for any* j *in the cluster, with the given fixed parity by Corollary 2.3.2. Hence any classical clause straddling the cluster on mode* j *and a classical mode* $k \in \mathsf{Class}$ *or mode in another quantum cluster, will have to be satisfied (if it can) no matter the value for* $x_j = 0, 1$*.*

The corollary means that if a vertex in the cluster is of degree at least 3, then, if the cluster is non-hPNC, we must have $n_q \leq 4$, and hence it is easy to verify whether an assignment exists. Otherwise, if it is hPNC and has a degree ≥ 3 vertex, its assignments are simple and Gaussian as they are very restricted in particle number (in the particle-hole transformed basis), collapsing the search space to being linearly dependent on the number of modes in the cluster. We capture this efficiency more precisely in Lemma 2.3.10 in the next section.

If all vertices in the cluster G_q are of degree less than 3, the graph G_q is a line or a loop. For a line, one can apply the Jordan-Wigner transformation and map the problem onto quantum 2-SAT on a line which can fully be solved, using existing methods. For a loop there is a small adaptation due to the Jordan-Wigner Pauli-Z strings. The cases of G_q being a line or a loop are dealt with in Section 2.3.3.1 below.

Corollary 2.3.5 also makes clear that the assignments of hidden particle number and parity on the cluster lead to strong restrictions on whether and how one can satisfy the remaining classical clauses in E_c, and this will be used to prove our main Theorems 2.1.4, 2.1.5 and 2.1.6.

2.3.3.1. Quantum clusters: lines and loops

To handle lines and loops, it is convenient to switch to qubit language as FERMIONIC 2-SAT projectors on a line map directly onto QUANTUM 2-SAT projectors on qubits via the Jordan-Wigner transformation in equation (2.2).

Let us first consider the taxonomy of graphs G_q that are lines or loops, shown in Figure 2.4. Remember that there is at most a single $\Pi_e^{02,q}$ or single $\Pi_e^{1,q}$ per edge e, otherwise one can reduce the pair to a classical edge, see the arguments in Section 2.2.1.1, and we do not yet consider any classical clauses in this section yet. If G_q is a line, its maximal spanning hPNC subgraph T_q (see Definition 2.2.3) is also a line, and G_q is either hPNC or non-hPNC as in Figure 2.4(a) and (b). If G_q is a loop, then T_q is either a loop, see Figure 2.4(c) and (d), or a line, Figure 2.4(e) and (f).

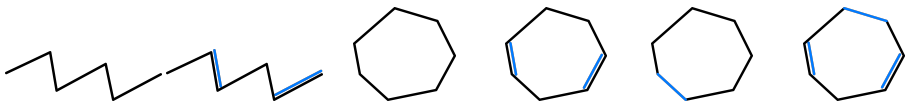

(a) Lemma 2.3.6 (hPNC). (b) Lemma 2.3.8 (non-hPNC). (c) Lemma 2.3.7 (hPNC). (d) Lemma 2.3.8 (non-hPNC). (e) Lemma 2.3.8 (non-hPNC). (f) Lemma 2.3.8 (non-hPNC).

Figure 2.4: Different types of line and loop clusters in the particle-hole-transformed basis. The black edges make up T_q (representing $\Pi_e^{1,q}$-type clauses) and the blue edges are $\Pi_e^{02,q}$-type clauses.

First, we consider the case of G_q being a line. The following Lemma is to be applied to T_q in the particle-holed transformed basis, thus containing only $\Pi^{1,q}$-type clauses. It is essentially a rephrasing of some results in [10].

Lemma 2.3.6. *Given a line with n_q fermionic modes and (rank-1) clauses $\Pi^{1,q}_{(i,i+1)}$ between modes i and $i+1$ for $i \in \{1,2,\ldots,n_q-1\}$. For each particle number $N_q \in \{0,1,\ldots,n_q\}$, there is a unique assignment satisfying the clauses $\Pi^{1,q}_{(i,i+1)}$.*

Proof. Via the Jordan-Wigner transformation in equation (2.2), each 2-mode rank-1 projector $\Pi^{1,q}_{(i,i+1)}$ becomes a projector onto a two-qubit quantum state $\Pi^{\text{qubit}}_{i,i+1} = \left|\psi\right\rangle\left\langle\psi\right|_{i,i+1}$ with $\left|\psi\right\rangle_{(i,i+1)} = \alpha|01\rangle + \beta|10\rangle$ (where $|\alpha|^2+|\beta|^2=1$). The entangled state $\left|\psi\right\rangle_{(i,i+1)}$ can generally be expressed as $B_i \otimes I|\Psi^-\rangle$ with a two-qubit singlet state $|\Psi^-\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$ and some diagonal, invertible B_i.

Let $H = I\otimes A_2 \ldots \otimes A_{n_q} H_{\text{FM-Heis}} I \otimes A_2^\dagger \ldots \otimes A_{n_q}^\dagger \geq 0$, with $H_{\text{FM-Heis}} = \sum_i |\Psi^-\rangle\langle\Psi^-|_{i,i+1}$, the 1D ferromagnetic Heisenberg model. Here A_2 is uniquely chosen such that $I\otimes A_2|\Psi^-\rangle_{1,2} = B_1\otimes I|\Psi^-\rangle_{1,2} = \left|\psi\right\rangle_{1,2}$. Given A_2, A_3 is chosen such that $A_2\otimes A_3|\Psi^-\rangle_{2,3} = B_2\otimes I|\Psi^-\rangle_{2,3} = \left|\psi\right\rangle_{2,3}$ etc. Importantly, the satisfying assignments of $\{\Pi^{\text{qubit}}_{(i,i+1)}\}$ are the zero-eigenvalue states of H. A zero-eigenvalue state with particle (or "excitation") number N_q can thus be obtained by applying $I\otimes A_2^{\dagger -1}\otimes\ldots\otimes A_{N_q}^{\dagger -1}$ to any state $\left|\psi^{N_q}_{\text{FM-Heis}}\right\rangle$ in the null-space of $H_{\text{FM-Heis}}$, with equal particle number N_q as the $A_i^{\dagger -1}$ matrices are diagonal. Hence, given a null-space solution with excitation number N_q for $H_{\text{FM-Heis}}$, one obtains a satisfying assignment for the fermionic problem on a line.

Due to the $U^{\otimes n_q}$-invariance of $H_{\text{FM-Heis}}$, the eigenspaces must be spanned by states of the form $\left|\chi^{\otimes n_q}\right\rangle$, hence symmetric subspaces of n_q qubits [11]. There is only one such state with a given particle number N_q, namely the fully permutation-symmetric state, which is in the null-space of $H_{\text{FM-Heis}}$ since $|\Psi^-\rangle\langle\Psi^-|_{i,i+1}\left|\psi^{N_q}_{\text{FM-Heis}}\right\rangle = |\Psi^-\rangle\langle\Psi^-|_{i,i+1}\pi(i,i+1)\left|\psi^{N_q}_{\text{FM-Heis}}\right\rangle = -|\Psi^-\rangle\langle\Psi^-|_{i,i+1}\left|\psi^{N_q}_{\text{FM-Heis}}\right\rangle = 0$ where $\pi(i,i+1)$ is the permutation between qubits i and $i+1$. Note that the unique satisfying assignment for the quantum 2-SAT clauses is thus $I\otimes A_2^{\dagger -1}\otimes\ldots\otimes A_{n_q}^{\dagger -1}\left|\psi^{N_q}_{\text{FM-Heis}}\right\rangle$. □

Let us now consider a hPNC cluster for which both G_q and T_q are a loop as in Fig. 2.4(c). Again the following Lemma is to be applied to T_q in the particle-hole-transformed basis in which it consists of $\Pi^{1,q}_e$ clauses only. For there to be any satisfying assignments on T_q, the clause closing the loop needs to be consistent with the other clauses. For FERMIONIC 2-SAT clauses this consistency condition also turns out to depend on the cluster parity:

Lemma 2.3.7. *Given a loop consisting of n_q fermionic modes with clauses $\Pi^{1,q}_{(i,i+1 \bmod n_q)}$ for $i\in\{1,2,\ldots,n_q\}$, with coefficients $u_{(i,i+1 \bmod n)}$ defined in equation (2.5). Besides the satisfying assignments at $N_q=0$ (all-empty) and $N_q=n_q$ (all-filled), there are satisfying assignments at a given N_q iff $(-1)^{N_q+n_q-1}\left[1/u_{(1,n_q)}\prod_{i=1}^{n_q-1}u_{(i,i+1)}\right]=1$. So, if $|1/u_{(1,n_q)}\prod_{i=1}^{n_q-1}u_{(i,i+1)}|=1$, there are only satisfying assignments at all even N_q (when n_q is odd) or odd N_q (when n_q is even).*

Proof. The proof of Lemma 2.3.7 is similar to the proof of Lemma 2.3.6. The only difference is that in the case of a loop, the clause $\Pi^{1,q}_{(1,n_q)}$, when Jordan-Wigner transformed, becomes

$$2\Pi^{\text{qubit}}_{(1,n_q)} = I_1 \otimes B_{n_q}\Big(|1\rangle_1 |0\rangle_{n_q} - Z_2 \ldots Z_{n_q-1} |0\rangle_1 |1\rangle_{n_q}\Big)\times \Big(\langle 1|_1 \langle 0|_{n_q} - Z_2 \ldots Z_{n_q-1} \langle 0|_1 \langle 1|_{n_q}\Big) I_1 \otimes B^\dagger_{n_q}, \quad (2.20)$$

with B_{n_q} diagonal and $(B_{n_q})_{11}/(B_{n_q})_{22} = -\beta_{(1,n_q)}/\gamma_{(1,n_q)} = -1/u^*_{(1,n_q)}$ (see equation (2.5)). The satisfying assignments of the loop problem correspond to zero-energy states of H in the proof of Lemma 2.3.6 that are also projected to zero by $\Pi^{\text{qubit}}_{(1,n_q)}$. We have that $(B_i)_{11}/(B_i)_{22} = -\gamma_{(i,i+1)}/\beta_{(i,i+1)} = -u^*_{(i,i+1)}$ for $i \in \{1,2,\ldots,n_q-1\}$, from which the entries of $A_2,\ldots,A_{n_q}$ follow (see the proof of Lemma 2.3.6). The relation $\Pi^{\text{qubit}}_{(1,n_q)} I \otimes A_2^{\dagger -1} \otimes \ldots \otimes A_{n_q}^{\dagger -1} \left|\psi^{N_q}_{\text{FM Heis}}\right\rangle = 0$ holds iff $u_{(1,n_q)} = (-1)^{N_q}(A_{n_q})^*_{22}/(A_{n_q})^*_{11}$ is obeyed (where we have used $(B_{n_q})_{11}/(B_{n_q})_{22} = -1/u^*_{(1,n_q)}$). Using that the entries of $A_2,\ldots,A_{n_q}$ can be expressed in terms of $\{u_{(i,i+1)}\}_{i=1}^{n_q-1}$, this condition is equivalent to the condition $(-1)^{N_q+n_q-1}\big[1/u_{(1,n_q)} \prod_{i=1}^{n_q-1} u_{(i,i+1)}\big] = 1$ from the lemma statement. □

2.3.3.2. non-hPNC lines and loops

We have thus far considered the hPNC line or loop graphs. For non-hPNC lines and loops, however, we must consider the effect of the $\Pi_e^{02,q}$ clauses (in the particle-hole-transformed basis) on the satisfiability of these clusters as in Fig. 2.4(b),(d), (e) and (f). These $\Pi_e^{02,q}$ clauses can be of two types: they can either be part of a rank-2 projector as in Fig. 2.4(d), or they can close a loop as in Figure 2.4(e). We will refer to the former as a *non-loop-closing* $\Pi_e^{02,q}$ clause and to the latter as a *loop-closing* $\Pi_e^{02,q}$ clause. Of course, any loop G_q contains at most one loop-closing $\Pi_e^{02,q}$ clause. If G_q is a line, then all $\Pi_e^{02,q}$ clauses are non-loop-closing as in Fig. 2.4(b). If G_q and its maximal spanning hPNC subgraph T_q are loops, all $\Pi_e^{02,q}$ clauses are also non-loop-closing (Fig. 2.4(d)).

Note that one can always apply a different particle-hole transformation (with a different T_q) such that the instance in Figure 2.4(f) becomes an instance in Figure 2.4(d), by ensuring that the loop is closed on a rank-2 edge. Therefore, if G_q contains a loop-closing $\Pi_e^{02,q}$ clause, then wlog it is the only $\Pi_e^{02,q}$ clause in G_q.

Lemma 2.3.8 below provides a detailed account of conditions under which a parity P_q is allowed on a quantum cluster G_q, provided that P_q is allowed on T_q. The important take-away is that for any non-hPNC line or loop, these conditions come down to straightforward relations between the u_e and v_e coefficients of respectively $\Pi_e^{1,q}$ and $\Pi_e^{02,q}$ clauses along the line or loop.

Lemma 2.3.8. *Given a non-hPNC cluster G_q which is a line or a loop on n_q modes. The following statements hold in the particle-hole-transformed basis in which T_q only has $\Pi_e^{1,q}$ clauses and P_q refers to the parity of the particle-hole-transformed assignment. Assume that T_q allows for satisfying assignments of parity P_q. If G_q contains a single $\Pi_e^{02,q}$ clause, then the same parities of T_q are allowed as satisfying assignments for G_q. If G_q contains at least two $\Pi_e^{02,q}$ clauses, then suppose:*

1. *G_q is a line (Figure 2.4(b)) then any P_q of T_q is allowed for G_q iff the following consistency condition is satisfied. For each pair of subsequent non-loop-closing $\Pi_e^{02,q}$*

clauses on edges $(j, j+1)$ *and* $(k, k+1)$ *(*$j < k$*),*

$$(\star)\quad v_{(j,j+1)} = \Big[\prod_{i=j}^{k-1} u_{(i,i+1)}\, u_{(i+1,i+2)}\Big]\, v_{(k,k+1)} \tag{2.21}$$

must hold.

2. *G_q is a loop and T_q a line (Figure 2.4(e). P_q is always allowed on G_q since there is only one $\Pi_e^{02,q}$ clause in G_q.*

3. *Both G_q and T_q are loops (Figure 2.4(d)). Any P_q of T_q is allowed for G_q, if condition $(\star)$ holds on $\Pi_e^{02,q}$ clauses along the loop, with the following caveat. If one of the non-loop-closing clauses is on edge $(1, n_q)$, then the one condition (on subsequent pairs) involving $\Pi_{(1,n_q)}^{02,q}$ has $v_{(1,n_q)}$ replaced by $-1/u_{(1,n_q)}^2\, v_{(1,n_q)}$ on the RHS of $(\star)$.*

Proof. By assumption, P_q is allowed on T_q. According to Corollary 2.3.2, if P_q is allowed on G_q, then a satisfying assignment has support on all states $|S\rangle = a_S^\dagger|\text{vac}\rangle$ with parity P_q. For P_q to be allowed on G_q, the $\Pi_e^{02,q}$ clauses on G_q should not lead to inconsistencies when connecting these states of parity P_q via Lemma 2.3.1. Such inconsistencies can come up when, starting with any $|S\rangle$, one repeatedly applies Lemma 2.3.1 to end up at state $|S\rangle$ again. If the associated coefficient relations in Lemma 2.3.1 are not consistent, P_q is excluded on G_q. For G_q lines and loops containing a single $\Pi_e^{02,q}$ clause, these inconsistencies are always avoided. For non-hPNC lines and loops with multiple $\Pi_e^{02,q}$ clauses, the consistency of these relations can be checked rather straightforwardly. To see this, we distinguish between cases (1), (2) and (3) in the lemma statement.

For case (1), all $\Pi_e^{02,q}$ clauses are non-loop-closing. The following process, realized by repeated application of Lemma 2.3.1, can lead to inconsistency. Start off with $|S\rangle$ with equal occupation on an edge with a $\Pi_e^{02,q}$ clause. One creates a pair of particles/holes on edge e, migrates them to a clause $\Pi_{e'}^{02,q}$ at edge e' via $\Pi_e^{1,q}$ clauses, and undoes the creation of the particle/hole pair through clause $\Pi_{e'}^{02,q}$ to end up at $|S\rangle$. Provided that the condition in the lemma statement holds, this can never lead to inconsistency. Ensuring that the condition holds on *subsequent* $\Pi_e^{02,q}$ clauses along G_q is sufficient to ensure the consistency between non-subsequent $\Pi_e^{02,q}$ clauses.

For case (2), the $\Pi_e^{02,q}$ clauses cannot lead to inconsistency, since there is only one such clause in G_q.

For case (3), the same process as for case (1) determines the consistency conditions. However, there can be a non-loop-closing $\Pi_{(1,n_q)}^{02,q}$ clause (i.e., on edge $(1, n_q)$). Again using Lemma 2.3.1, it can be shown straightforwardly that the altered $(\star)$ condition implies consistency between the $\Pi_{(1,n_q)}^{02,q}$ clause and its subsequent $\Pi_e^{02,q}$ clause along the loop. □

Importantly, we can show that satisfying assignments on non-hPNC clusters which are lines or loops are always fermionic Gaussian states of fixed parity. This is argued using the fact that a $\Pi_e^{1,q} + \Pi_e^{02,q}$ projector is quadratic, see Section 2.2.1.1, and arguing that one

can always add extra clauses onto edges in the cluster so that all edges become of this form $\Pi_e^{1,q} + \Pi_e^{02,q}$. Satisfying assignments must then lie in the null-space of a quadratic fermionic (positive semi-definite) Hamiltonian, hence be a Gaussian state.

Lemma 2.3.9. *Given a non-hPNC cluster G_q which is a line or loop. The satisfying assignments (if they exist) on this cluster are Gaussian states.*

Proof. Let us first consider line or loop clusters (in the particle-hole transformed basis) for which the only $\Pi_e^{02,q}$ clauses are non-loop-closing. There can be multiple such clauses (i.e., cases (1) and (3) in Lemma 2.3.8), or just a single one. In the former case, we assume equation (2.21) is satisfied (otherwise, the cluster would not be satisfiable). Let us add a non-loop-closing $\Pi_e^{02,q}$ clause to an edge not yet containing one. If this $\Pi_e^{02,q}$ clause also obeys equation (2.21), then the satisfying assignments of the original problem are still satisfying assignments. One can keep adding $\Pi_e^{02,q}$ clauses in this manner until *all* edges in the cluster are of type $\Pi_e^{1,q} + \Pi_e^{02,q}$. In that case, all cluster projectors are quadratic and hence the satisfying assignments are Gaussian states as per observations in Section 2.2.1.

What is left is to consider clusters whose only $\Pi_e^{02,q}$ clause is a loop-closing $\Pi_e^{02,q}$ clause (i.e., Figure 2.4(e) and case (2) in Lemma 2.3.8). Importantly, a quantum cluster with exactly the same satisfying assignments can be obtained by deleting the loop-closing $\Pi_e^{02,q}$ clause, and adding an appropriate non-loop-closing $\Pi_e^{02,q}$ clause anywhere along T_q. Hence we have obtained a cluster of type (1) in Lemma 2.3.8, for which we have already shown that the lemma statement holds. □

2.3.3.3. Efficient verification of allowed hidden particle number or parity on a quantum cluster

The following Lemma captures that verifying what cluster particle numbers and parities are allowed in the cluster-product form assignment, equation (2.19), already restricted by the results in the previous section, is classically efficient.

Lemma 2.3.10 (Cluster checks)**.** *Given a quantum cluster $G_q = (V_q, E_q)$ consisting of n_q modes and m_q quantum clauses. If G_q is an hPNC cluster, one can check in $O(n_q + m_q)$ time which hidden cluster particle numbers N_q (being the cluster particle number of $K_B |\psi\rangle$, with $|\psi\rangle$ a satisfying assignment) correspond to satisfying assignments. If G_q is a non-hPNC cluster, one can check in $O(n_q + m_q)$ time which cluster parities correspond to satisfying assignments.*

Proof. First, suppose G_q is a hPNC cluster. By Corollary 2.3.5 point 2, hidden cluster particle numbers $N_q = 0$ and $N_q = n_q$ are *always* permitted.

If the hPNC cluster G_q has a vertex of degree at least 3, then by Corollary 2.3.5 point 1, to see if a single-particle state with $N_q = 1$ and/or single-hole state with $N_q = n_q - 1$ are permitted, one simply lists all m_q partnering constraints (Lemma 2.3.1) between n_q coefficients in the potential satisfying assignment and checks the consistency of these constraints in time $O(n_q + m_q)$. On the other hand, suppose the hPNC cluster G_q has no vertex of degree at least 3. If G_q is a line, all hidden cluster particle numbers are allowed via Lemma 2.3.6, and the corresponding cluster particle states are unique. In case G_q is a

loop, then all cluster particle numbers for which the condition in Lemma 2.3.7 is obeyed, are permitted, with the caveat that $N_q = 0$ and $N_q = n_q$ are permitted in any case.

Now suppose that G_q is a non-hPNC cluster. If G_q has a vertex of degree at least 3, then by Corollary 2.3.5 any assignment is excluded for $n_q \geq 5$, one parity is excluded for $n_q = 4$, and obviously n_q cannot be ≤ 3. Clearly, one can efficiently verify the existence of an assignment for $n_q = 4$. If the non-hPNC cluster G_q does not have a vertex of degree at least 3, then G_q can be a line or a loop and which (if any) parity is allowed for those clusters can be checked in $O(n_q + m_q)$ time using the conditions in Lemmas 2.3.7 and 2.3.8. □

2.4. An efficient classical algorithm for Fermionic 2-SAT (with and without fixed parity)

2.4.1. Solving Fermionic 2-SAT: proof of Theorem 2.1.4

Let us now prove Theorem 2.1.4, which states that FERMIONIC 2-SAT can be solved in time $O(n+m)$. Solving FERMIONIC 2-SAT comes down to performing certain checks on properties of quantum clusters, and running a (classical) 2-SAT algorithm based on these constraints and the remaining classical clauses:

Proof of Theorem 2.1.4. Due to Proposition 2.3.3 we can restrict the assignment to cluster-product form. We will be constructing a classical 2-SAT instance, which will be used for deciding whether the FERMIONIC 2-SAT instance is satisfiable. This classical 2-SAT instance is denoted by $\mathcal{C}$. First, let us consider only the constraints from quantum edges.

1. We first consider all hPNC clusters. For these clusters, Corollary 2.3.5 point 2 implies that there are *always* two classical assignments that satisfy all quantum clauses in the cluster and that are each other's negations (see Figure 2.5 for an illustration). Now consider the set of classical clauses, either internal to the cluster or straddling between the cluster and some classical modes or between the cluster and some other quantum cluster. By Corollary 2.3.5 point 3 for $n_q > 2$, obeying any internal classical clauses on the cluster in any assignment with hidden particle number $N_q \neq 0, N_q \neq n_q$ would also allow for two classical assignments $N_q = 0, n_q$. Hence wlog we can assume that we have to choose from these classical assignments for any $n_q > 2$ cluster when constructing a global satisfying assignment. One can impose this using a set of 2-SAT clauses whose two unique satisfying assignments are these two classical assignments (which are each other's negations): note that we need to apply the particle-hole transformation K_B to obtain these 2-SAT clauses. These 2-SAT clauses, the classical clauses internal to the cluster and any straddling clauses are included in $\mathcal{C}$. The case $n_q = 2$ is dealt with separately for convenience. For $n_q = 2$, if there is a single clause and no internal classical clause, we do the same as for $n_q > 2$. If $n_q = 2$ with one $\Pi_e^{1,q}$ clause and one $\Pi_e^{02,q}$ clause, the cluster is non-hPNC and we deal with it in the next point 2. If $n_q = 2$ and the projector is rank-3, then it is hPNC, with 2 classical internal clauses and there is at most one satisfying Gaussian, non-classical assignment. For

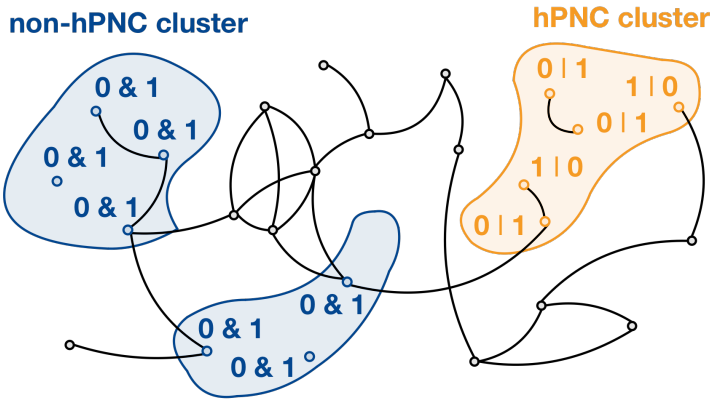

Figure 2.5: Satisfying assignments on non-hPNC quantum clusters are such that each mode has to be both empty *and* occupied, see Corollary 2.3.5 point 6. On hPNC clusters, one only has to assign either of two classical states to check whether the instance is satisfiable (apart from a special case when the cluster only has two modes, $n_q = 2$), when we do not aim to find an assignment of given parity. The black vertices and edges respectively denote modes that are not in any quantum cluster (i.e., classical modes) and classical clauses.

example, the two classical clauses exclude the assignments 00 and 11, allowing only the state $(\alpha a_1^\dagger + \beta a_2^\dagger)|\text{vac}\rangle$ orthogonal to some $\Pi^{1,q}_{(1,2)}$ clause. Hence, in this case only, one has to choose this non-classical (Gaussian) assignment and handle the possible classical clauses straddling this cluster via Corollary 2.3.5 point 3: a straddling clause on mode 1 in the cluster, i.e. $(x_1 \vee u)$ or $(\overline{x}_1 \vee u)$, gets replaced by clause (u) with literal u on some classical mode (and similarly for mode 2 in the cluster). These latter clauses are included in $\mathcal{C}$.

2. For all quantum clusters that are non-hPNC, we first check via Lemma 2.3.10 in $O(n+m)$ time whether there exists a satisfying assignment. Corollary 2.3.5 point 6 says that if there exists a satisfying assignment, there can be no additional internal classical clauses (as the state has support on all occupations with given parity). We handle any straddling classical clauses by replacing $(x_m \vee u_i)$ or $(\overline{x}_m \vee u_i)$, with mode m inside the non-hPNC cluster and classical literal u_i, by clauses (u_i). These latter clauses are included in $\mathcal{C}$. The case when the classical clause is straddling between two quantum clusters can be dealt with similarly.

After these preprocessing steps in which we may find that there can't be a satisfying assignment (then output NO), let the collection of classical clauses obtained via the preprocessing, as well as the remaining clauses on solely classical modes, be the classical 2-SAT instance $\mathcal{C}$. We then solve instance $\mathcal{C}$ in time $O(n+m)$ [1]. □

2.4.2. Solving Fermionic 2-SAT with fixed parity: proof of Theorem 2.1.5

Let us now address the question of whether fixing the parity changes the complexity of the FERMIONIC 2-SAT problem. We prove Theorem 2.1.5 in this section, which states that this problem can still be solved efficiently, in time $O(nm)$. Our algorithm for solving FERMIONIC 2-SAT with fixed parity $P \in \{-1,+1\}$ consists of running certain efficient checks on the allowed satisfying assignments on quantum clusters as a preprocessing step to construct a (classical) 2-SAT with fixed parity instance, on which one runs the $O(nm)$ time algorithm developed in Theorem 2.C.1 in Appendix 2.C.

Importantly, non-hPNC clusters are treated differently than hPNC clusters in this next proof. Namely, any assignment in a non-hPNC cluster satisfies the classical cluster-straddling clauses (or not) *independent* of its parity, while for hPNC clusters, assignments of different parity can lead to different ways of satisfying the straddling classical clauses.

Proof of Theorem 2.1.5.. By Proposition 2.3.3 an assignment of fixed parity is of cluster-product form. We proceed in three steps.

1. We run the $O(nm)$ time algorithm (developed in Appendix 2.C) for solving the (classical) 2-SAT instance $\mathcal{C}$ obtained in the proof of Theorem 2.1.4 via items 1 and 2, with fixed parity. If the algorithm outputs an assignment of the right parity, we are done and output YES. If there is no assignment, output NO. If there is an assignment but with the wrong parity (which the algorithm will tell you), let's call the assignment X_0, and proceed to the next step.

2. We consider whether it is possible to flip the parity of any non-hPNC cluster in X_0. No variables internal to a non-hPNC cluster enter the instance $\mathcal{C}$ due to how we handle the straddling clauses in item 2 in the proof of Theorem 2.1.4. Therefore, if we can flip the parity of a single cluster $G_q^j \in \mathsf{non-hPNC}$, the assignment on $G \backslash G_q^j$ of X_0 is still valid, and hence we have full assignment with the correct parity. According to Corollary 2.3.2, each non-hPNC cluster has at most one satisfying assignment per parity. Which parities (one or both) are allowed for each non-hPNC cluster can be checked in time $O(n+m)$ via Lemma 2.3.10. So we iterate over all non-hPNC clusters and if there is one whose parity can be flipped, we are done (output YES). If not, we proceed to the next step.

3. The next step pertains to all hPNC clusters, except the $n_q = 2$ clusters with a rank-3 edge which were treated differently in obtaining $\mathcal{C}$ in item 1 of the proof of Theorem 2.1.4: we will not attempt flipping parity on these clusters as there is only one possible assignment.

   Due to the failure of obtaining a flipped parity assignment in the previous two steps, we know that it is not possible to change the parity by flipping classical modes and/or flipping the hidden particle number $N_q = 0$ (all-empty) to $N_q = n_q$ (all-filled) assignment on any hPNC cluster (which are the possible assignments inside the hPNC clusters in $\mathcal{C}$ which are examined in step 1). This has the following consequence. We define the classical 2-SAT instances $\mathcal{C}^{+,\mathsf{M}}$ which is a superset of the classical clauses of $\mathcal{C}$ and some classical clauses which pertain to a subset of

non-hPNC clusters $\mathsf{M} \subseteq \mathsf{hPNC}$. That is, $\mathcal{C}^{+,\mathsf{M}}$ is obtained by adding to the collection of clauses in $\mathcal{C}$, for each cluster $G_q^j \in \mathsf{M}$, the following: replace each straddling clause between a mode $i \in G_q^j$ and classical mode, of the form $(x_i \vee u)$ or $(\overline{x}_i \vee u)$ by clause (u) with u a literal of the classical mode. One does the same for any straddling clause with another hPNC quantum cluster provided it is not in the set M (if it is non-hPNC or in the set M, there would be no satisfying assignment for $\mathcal{C}^{+,\mathsf{M}}$). Due to failure in the previous two steps, either an instance $\mathcal{C}^{+,\mathsf{M}}$ will have no satisfying assignment, or an assignment of the wrong parity, i.e. the same as X_0, for any M, since any satisfying assignment of $\mathcal{C}^{+,\mathsf{M}}$ is also a satisfying assignment for $\mathcal{C}$. Hence, the only way to still obtain a flipped parity assignment as compared to X_0 is to possibly flip the parity of some individual hPNC cluster by *not choosing* the hidden particle number $N_q = 0$ or $N_q = n_q$ assignment on it, as we did for $\mathcal{C}$. Note that if we do not constrain the overall parity, as in Theorem 2.1.4, there was never a reason to consider other possible assignments (see also Fig. 2.5).

Hence, we first consider for each hPNC cluster G_q^i whether one can flip its parity, and then if so, we decide whether an altered classical 2-SAT instance $\mathcal{C}_{G_q^i}$ has a satisfying assignment using a 2-SAT solver. To check whether one can flip its parity, observe that no internal classical clauses must be present as any assignment with hiddden particle numbers $N_q \neq 0, N_q \neq n_q$ has the property that each mode j in the cluster can be both 0 and 1. If classical clauses are present and/or the parity cannot be flipped for other reasons, we move onto the next hPNC cluster. As for checking on a hPNC cluster whether an assignment of flipped parity is allowed, given its quantum clauses, Lemma 2.3.10 shows that this can be done efficiently, in $O(n+m)$ time. For a hPNC cluster which is a line or a loop, in principle non-Gaussian assignments with hidden particle number $N_q \neq 0, 1, n_q - 1, n_q$ may be allowed. However, one can observe that, if a flipped parity assignment is possible, one can take this assignment actually to be a Gaussian state, i.e. one only with hidden particle number $N_q = 0, 1, n_q - 1, n_q$.

If the parity can be flipped for cluster G_q^i, we construct $\mathcal{C}_{G_q^i}$. To get from $\mathcal{C}$ to $\mathcal{C}_{G_q^i}$ we (1) replace all straddling clauses between a mode j in cluster G_q^i and a classical mode $(x_j \vee u)$ and $(\overline{x}_j \vee u)$ by clause (u). One does the same for a straddling clause with a mode in a hPNC quantum cluster (if there is a straddling clause with a non-hPNC quantum cluster, then there won't be a satisfying assignment when flipping cluster G_q^i). From $\mathcal{C}$ we remove all classical clauses which constrain the assignment on G_q^i to have hidden particle number $N_q = 0$ or $N_q = n_q$ as described in item 1 in the proof of Theorem 2.1.4: there are no more variables in $\mathcal{C}_{G_q^i}$ which pertain to $\mathcal{C}_{G_q^i}$. If $\mathcal{C}_{G_q^i}$ has an assignment, then it must now be an assignment of the right parity and we output YES. If not, for all hPNC clusters that we iterate over, we output NO, and this process is clearly efficient.

□

2.5. PNC Fermionic 2-SAT with fixed particle number is NP-complete

Let us now consider the complexity of solving Problem 2.1.3, which corresponds to deciding PNC FERMIONIC 2-SAT with particle number fixed to some $N \in \{0,1,\ldots,n\}$. In contrast to fixing the parity, we show that fixing the particle number to N makes FERMIONIC 2-SAT NP-complete (Theorem 2.1.6).

Proof of Theorem 2.1.6. Proposition 2.3.3 and Corollary 2.3.5 show that satisfying assignments of PNC FERMIONIC 2-SAT are of cluster-product form. In the PNC version of FERMIONIC 2-SAT, quantum clauses are only of type $\Pi_e^{1,q}$. As a consequence, all quantum clusters are PNC, even without performing a particle-hole transformation. Therefore, we can characterize satisfying assignments on quantum clusters by the particle number on that cluster. A witness W thus consists of a list of $O(n)$ integers $W = \left(\{N_q^i\}_{i\in\mathsf{hPNC}}, \{x_k\}_{k\in\mathsf{Class}}\right)$. The verifier algorithm verifies for each quantum cluster i, using Lemma 2.3.10 whether N_q^i is indeed an allowed particle number on that cluster. Note that the verifier does not need to construct or have the actual fermionic state to do this. Using the collection of particle numbers on quantum clusters N_q^i and Corollary 2.3.5, she infers for each mode in a quantum cluster whether it is empty, occupied or both. Combined with $\{x_k\}_{k\in\mathsf{Class}}$, it is then verified whether all classical clauses $e \in E_c$ are indeed satisfied. As a final step, the verifier checks if the total particle number of the witness W is indeed the given N. This verification algorithm is efficient. To show NP-hardness, let us consider an instance of FERMIONIC 2-SAT consisting of monotone classical clauses only, i.e. only clauses of type $(x_i \vee x_j)$. Deciding this problem with fixed Hamming weight N is equivalent to deciding whether there is vertex cover with N vertices, which is NP-complete. Indeed, it is known that classical 2-SAT with fixed Hamming weight, i.e. *the weighted 2-SAT problem* (W2SAT) is NP-complete, see `https://en.wikipedia.org/wiki/2-satisfiability`. □

Remark 2.5.1 (Classical Assignments)**.** *Part 1 of the proof of Theorem 2.1.4 for hPNC clusters and classical modes shows that PNC* FERMIONIC *2-SAT (which only has* $\Pi_e^{1,q}$ *quantum clauses and thus hPNC quantum clusters) without any particle number constraint always has a classical satisfying assignment, i.e., a satisfying assignment of the form* $a_S^\dagger|\text{vac}\rangle$*, with the exception of 2-mode quantum clusters with rank-3 projectors in Gaussian states. Note that this is similar to quantum 2-SAT always having a product assignment with the exception of rank-3 edges [3]. In contrast, when fixing particle number to some integer N as in Problem 2.1.3, there are instances with only non-Gaussian satisfying assignments, like the satisfying assignments on a line for* $1 < N_q < n_q - 1$ *in Lemma 2.3.6.*

2.6. Complexity of Fermionic k-SAT and related problems

Let us provide some background and perspective on the (known) complexity of related fermionic problems. Using qubit-to-fermion mappings, such as the one in Ref. [12], we can straightforwardly argue the following, see Appendix 2.D for the proof.

Lemma 2.6.1. FERMIONIC k-SAT $\in$ QMA *and* FERMIONIC k-SAT *is* QMA_1*-hard for* $k = 9$.

This Lemma mirrors the results for QMA_1-hardness of QUANTUM k-SAT for $k \geq 3$ [2, 5].

One may expect that FERMIONIC k-SAT $\in \mathrm{QMA}_1$ for all variants of the problem (particle or non-particle number conserving, with a fixed particle number or not), but one has to be cautious about what basic gates are used in the class QMA_1 (usually H, CNOT and T) and whether this set of gates allows for a zero-error acceptance in the YES case when the fermionic problem and its specification is mapped to qubits.

It may be possible to reduce $k = 9$ in Lemma 2.6.1 to a lower k by adapting the space-time circuit-to-Hamiltonian construction in Section 3.3 in [13] which proves QMA-completeness of a fermionic circuit Hamiltonian with projectors which involve at most 4 fermionic modes ($k = 4$), *under the restriction* that there is 1 particle per track on the 2D lattice. Thus we don't have an overall particle constraint, but several fixed particle sectors, which can possibly be shown to be equivalent to an overall constraint.

Let us also mention a result that is related to this work. It is known that the 'FERMIONIC MAX-2-SAT' problem *with* particle number constraint is QMA-complete, see the next Theorem 2.6.2. Note that this problem is fundamentally different from FERMIONIC 2-SAT, since the former is a ground-state energy estimation problem and not a question of whether a given Hamiltonian is frustrated or not.

Theorem 2.6.2 (Theorem 2 in [14])**.** *Determining the ground state energy with* $1/\mathrm{poly}(n)$ *precision for a class of Fermi-Hubbard Hamiltonians with n fermionic modes, with particle number* $N = n/2$ *(half-filling) is* QMA*-complete. The particle-number-conserving Fermi-Hubbard Hamiltonian on a graph* $G = (V, E)$ *is*

$$H_{\mathrm{FH}} = U \sum_{i \in V} n_{i,+} n_{i,-} + \sum_{(i,j) \in E} \sum_{\sigma = \pm} t_{i,j} (a^\dagger_{i,\sigma} a_{j,\sigma} + a^\dagger_{j,\sigma} a_{i,\sigma}), \tag{2.22}$$

with $n_{i,+} = a^\dagger_{i,+} a_{i,+}$, $U, t_{i,j} \in \mathbb{R}$. *For* QMA*-completeness, bounds are specified on the parameters* U *and* $\{t_{i,j}\}$.

2.7. Discussion

Interestingly, it is not clear whether there is an efficient classical algorithm to solve instances of QUANTUM 2-SAT (with only parity-conserving projectors) when we ask for an assignment with fixed parity: in this case there can be non-product satisfying assignments. In addition, we don't have the "natively-fermionic" Lemma 2.3.4 for QUANTUM 2-SAT. We conjecture that this problem is of different complexity than FERMIONIC 2-SAT with fixed parity which we have proved to be efficiently solvable in Theorem 2.1.5. This would mesh elegantly with the point of view that parity-conserving interactions are fundamental for fermions.

We have seen that the satisfying assignments of FERMIONIC 2-SAT are of cluster-product form with some modes with purely classical occupations, some modes in a possibly-large clusters in Gaussian states of fixed parity or hidden fixed particle number, and some modes in a 4-fermion cluster in a non-Gaussian state. It will be interesting to explore how we can use these results for FERMIONIC 2-SAT to develop approximation algorithms or heuristic strategies, quantum or classical, to solve FERMIONIC k-SAT for $k > 2$. For FERMIONIC $(k > 2)$-SAT one may expect that satisfying assignments can be genuinely

many-mode non-Gaussian states: due to its QMA_1-hardness, satisfying assignments for general FERMIONIC k-SAT problems are not expected to be classically efficiently describable. Thus for such problems, one can seek quantum heuristic strategies, i.e. a quantum equivalent of classical heuristic SAT solvers, which aim at constructing a satisfying state on a quantum computer: such strategies could build on the nature of satisfying assignments for the (FERMIONIC) 2-SAT problem. It is an open question whether there are interesting classical mathematical problems which can be formulated as a question about the existence of a satisfying assignment to a finite-size, quantum or FERMIONIC k-SAT problem,—perhaps the results in [15] can be useful here—. Such construction would be a quantum counterpart to a classical computer-assisted proof obtained through the use of a classical heuristic SAT solver [16].

As a general question, it might be interesting to consider FERMIONIC ($k > 2$)-SAT problems with additional fermionic symmetry, for example consider whether there are complexity-theoretic consequences of time-reversal, spatial-parity and charge-conjugation symmetries as used in the classification of non-interacting fermionic models [17].

References

[1] B. Aspvall, M. F. Plass and R. Tarjan. 'A linear-time algorithm for testing the truth of certain quantified Boolean formulas'. In: *Information Processing Letters* 8.3 (1979), pp. 121–123. ISSN: 0020-0190. DOI: https://doi.org/10.1016/0020-0190(79)90002-4.

[2] S. Bravyi. 'Efficient algorithm for a quantum analogue of 2-SAT'. In: *Contemporary Mathematics*. Vol. 536. American Mathematical Society, 2011. eprint: quant-ph/0602108.

[3] I. Arad, M. Santha, A. Sundaram and S. Zhang. 'Linear-Time Algorithm for Quantum 2-SAT'. In: *Theory of Computing* 14 (2018), pp. 1–27. URL: https://theoryofcomputing.org/articles/v014a001/.

[4] N. de Beaudrap and S. Gharibian. 'A Linear Time Algorithm for Quantum 2-SAT'. In: *31st Conference on Computational Complexity (CCC 2016)*. Vol. 50. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016, 27:1–27:21. ISBN: 978-3-95977-008-8. DOI: 10.4230/LIPIcs.CCC.2016.27.

[5] D. Gosset and D. Nagaj. 'Quantum 3-SAT is QMA_1-Complete'. In: *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. IEEE, Oct. 2013. DOI: 10.1109/focs.2013.86.

[6] D. Rudolph. *Towards a universal gateset for* QMA_1. 2024. arXiv: 2411.02681 [quant-ph]. URL: https://arxiv.org/abs/2411.02681.

[7] J. Brakensiek, S. Gopi and V. Guruswami. 'CSPs with global modular constraints: algorithms and hardness via polynomial representations'. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2019. Phoenix, AZ, USA: Association for Computing Machinery, 2019, pp. 590–601. ISBN: 978-1-4503-6705-9. DOI: 10.1145/3313276.3316401.

[8] S. Bravyi and A. Kitaev. 'Fermionic Quantum Computation'. In: *Annals of Physics* 298.1 (2002), pp. 210–226. URL: https://doi.org/10.1006%5C%2Faphy.2002.6254.

[9] F. de Melo, P. Ćwikliński and B. M. Terhal. 'The power of noisy fermionic quantum computation'. In: *New Journal of Physics* 15.1 (Jan. 2013), p. 013015. DOI: 10.1088/1367-2630/15/1/013015.

[10] Z. Ji, Z. Wei and B. Zeng. 'Complete characterization of the ground-space structure of two-body frustration-free Hamiltonians for qubits'. In: *Phys. Rev. A* 84 (4 Oct. 2011), p. 042338. DOI: 10.1103/PhysRevA.84.042338.

[11] A. W. Harrow. *The Church of the Symmetric Subspace*. 2013. arXiv: 1308.6595 [quant-ph]. URL: https://arxiv.org/abs/1308.6595.

[12] Y. Herasymenko, A. Anshu, B. M. Terhal and J. Helsen. 'Fermionic Hamiltonians without trivial low-energy states'. In: *Phys. Rev. A* 109 (5 May 2024), p. 052431. DOI: 10.1103/PhysRevA.109.052431.

[13] N. P. Breuckmann and B. M. Terhal. 'Space-time circuit-to-Hamiltonian construction and its applications'. In: *Journal of Physics A: Mathematical and Theoretical* 47.19 (Apr. 2014), p. 195304. DOI: 10.1088/1751-8113/47/19/195304.

[14] B. O'Gorman, S. Irani, J. Whitfield and B. Fefferman. 'Intractability of Electronic Structure in a Fixed Basis'. In: *PRX Quantum* 3 (2 May 2022), p. 020322. DOI: 10.1103/PRXQuantum.3.020322.

[15] R. King and T. Kohler. *Gapped clique homology on weighted graphs is QMA_1-hard and contained in QMA*. 2024. arXiv: 2311.17234 [quant-ph]. URL: https://arxiv.org/abs/2311.17234.

[16] M. J. H. Heule, O. Kullmann and V. W. Marek. 'Solving and Verifying the Boolean Pythagorean Triples Problem via Cube-and-Conquer'. In: *Theory and Applications of Satisfiability Testing – SAT 2016*. Ed. by N. Creignou and D. Le Berre. Cham: Springer International Publishing, 2016, pp. 228–245. ISBN: 978-3-319-40970-2.

[17] A. Kitaev, V. Lebedev and M. Feigel'man. 'Periodic table for topological insulators and superconductors'. In: *AIP Conference Proceedings*. AIP, 2009. DOI: 10.1063/1.3149495.

[18] R. Tarjan. 'Depth-First Search and Linear Graph Algorithms'. In: *SIAM Journal on Computing* 1.2 (1972), pp. 146–160. DOI: 10.1137/0201010.

[19] R. Tarjan. 'Edge-disjoint spanning trees and depth-first search'. In: *Acta Informatica* 6.2 (1976), pp. 171–185. DOI: 10.1007/BF00268499.

[20] T. Feder. 'Network flow and 2-satisfiability'. In: *Algorithmica* 11 (1994), pp. 291–319.

[21] T. Morimae, D. Nagaj and N. Schuch. 'Quantum proofs can be verified using only single-qubit measurements'. In: *Physical Review A* 93.2 (Feb. 2016). ISSN: 2469-9934. DOI: 10.1103/physreva.93.022326.

2.A. Mathematical micro-facts for Section 2.2

2.A.1. Invariance of Fermionic 2-SAT clauses that exclude all-empty or all-filled states

For an edge $e = (j,k)$ one can define rotated annihilation operators:

$$\begin{pmatrix} \tilde{a}_j \\ \tilde{a}_k \end{pmatrix} = \begin{pmatrix} U_{11} & U_{12} \\ U_{21} & U_{22} \end{pmatrix} \begin{pmatrix} a_j \\ a_k \end{pmatrix}, \tag{2.23}$$

with U a unitary matrix.

We can define

$$\Pi_e^0 = \big(I - \tilde{a}_j^\dagger \tilde{a}_j\big)\big(I - \tilde{a}_k^\dagger \tilde{a}_k\big), \tag{2.24}$$

and

$$\Pi_e^2 = \tilde{a}_j^\dagger \tilde{a}_j \tilde{a}_k^\dagger \tilde{a}_k. \tag{2.25}$$

The following holds

Proposition 2.A.1. *Projectors Π_e^0 and Π_e^2 in equation* (2.24) *and equation* (2.25) *are invariant under any transformation U, as in equation* (2.23)*, and can thus be viewed as classical clauses in the $\{a_j, a_k\}$ mode basis, with the classical Boolean variables corresponding to occupation numbers of the $\{a_j, a_k\}$ modes. That is, Π_e^2 becomes the clause $(\overline{x}_j \vee \overline{x}_k)$ (excluding 11) and Π_e^0 becomes $(x_j \vee x_k)$ (excluding 00) as in equation* (2.6).

Proof. Inserting the transformation U into Π_e^0 with $e = (j,k)$ gives

$$\begin{aligned} \Pi_e^0 = \; & \tilde{a}_j \tilde{a}_j^\dagger \tilde{a}_k \tilde{a}_k^\dagger \\ = \big[I - \big(U_{11}^* a_j^\dagger + U_{12}^* a_k^\dagger\big)\big(U_{11} a_j + U_{12} a_k\big)\big]\big[I - \big(U_{21}^* a_j^\dagger + U_{22}^* a_k^\dagger\big)\big(U_{21} a_j + U_{22} a_k\big)\big] \\ = I - a_j^\dagger a_j - a_k^\dagger a_k + |\det U|^2 a_j^\dagger a_j a_k^\dagger a_k = a_j a_j^\dagger a_k a_k^\dagger, \end{aligned} \tag{2.26}$$

where we have used unitarity of U: $U_{21}U_{22}^* = -U_{11}U_{12}^*$ and $|\det U| = 1$. Similarly, since $\Pi_e^2 = I - \tilde{a}_j^\dagger \tilde{a}_j - \tilde{a}_k^\dagger \tilde{a}_k + \Pi_e^0$, and $\tilde{a}_j^\dagger \tilde{a}_j + \tilde{a}_k^\dagger \tilde{a}_k$ is invariant under U, Π_e^2 is also invariant under U: $\Pi_e^2 = \tilde{a}_j^\dagger \tilde{a}_j \tilde{a}_k^\dagger \tilde{a}_k = a_j^\dagger a_j a_k^\dagger a_k$. □

2.A.2. Action of the particle-hole transformation K_S on $\Pi_e^{1,q}$ and $\Pi_e^{02,q}$ clauses

The transformation rules are

$$j\in S,\ k\notin S\text{:}\quad\begin{cases} K_S\Pi_e^{1,q}K_S^{-1}=\tilde{\Pi}_e^{02,q} \text{ with } \tilde{\alpha}_e=-\Big[\prod\limits_{\substack{j\le i<k\\ \text{s.t. } i\in S}}(-1)\Big]\,\beta_e \text{ and } \tilde{\delta}_e=\gamma_e \\ K_S\Pi_e^{02,q}K_S^{-1}=\tilde{\Pi}_e^{1,q} \text{ with } \tilde{\beta}_e=-\Big[\prod\limits_{\substack{j\le i<k\\ \text{s.t. } i\in S}}(-1)\Big]\,\alpha_e \text{ and } \tilde{\gamma}_e=\delta_e. \end{cases} \tag{2.27}$$

$$j\notin S,\ k\in S\text{:}\quad\begin{cases} K_S\Pi_e^{1,q}K_S^{-1}=\tilde{\Pi}_e^{02,q} \text{ with } \tilde{\alpha}_e=\Big[\prod\limits_{\substack{j\le i<k\\ \text{s.t. } i\in S}}(-1)\Big]\,\gamma_e \text{ and } \tilde{\delta}_e=\beta_e. \\ K_S\Pi_e^{02,q}K_S^{-1}=\tilde{\Pi}_e^{1,q} \text{ with } \tilde{\beta}_e=\Big[\prod\limits_{\substack{j\le i<k\\ \text{s.t. } i\in S}}(-1)\Big]\,\delta_e \text{ and } \tilde{\gamma}_e=\alpha_e. \end{cases} \tag{2.28}$$

$$j\in S,\ k\in S\text{:}\quad\begin{cases} K_S\Pi_e^{1,q}K_S^{-1}=\tilde{\Pi}_e^{1,q} \text{ with } \tilde{\beta}_e=-\Big[\prod\limits_{\substack{j\le i<k\\ \text{s.t. } i\in S}}(-1)\Big]\,\gamma_e \text{ and } \tilde{\gamma}_e=\beta_e. \\ K_S\Pi_e^{02,q}K_S^{-1}=\tilde{\Pi}_e^{02,q} \text{ with } \tilde{\alpha}_e=-\Big[\prod\limits_{\substack{j\le i<k\\ \text{s.t. } i\in S}}(-1)\Big]\,\delta_e \text{ and } \tilde{\delta}_e=\alpha_e. \end{cases} \tag{2.29}$$

$$j\notin S,\ k\notin S\text{:}\quad\begin{cases} K_S\Pi_e^{1,q}K_S^{-1}=\tilde{\Pi}_e^{1,q} \text{ with } \tilde{\beta}_e=\Big[\prod\limits_{\substack{j\le i<k\\ \text{s.t. } i\in S}}(-1)\Big]\,\beta_e \text{ and } \tilde{\gamma}_e=\gamma_e. \\ K_S\Pi_e^{02,q}K_S^{-1}=\tilde{\Pi}_e^{02,q} \text{ with } \tilde{\alpha}_e=\Big[\prod\limits_{\substack{j\le i<k\\ \text{s.t. } i\in S}}(-1)\Big]\,\alpha_e \text{ and } \tilde{\delta}_e=\delta_e. \end{cases} \tag{2.30}$$

2.B. Example of a unique 4-fermion non-Gaussian satisfying assignment

We work through explicitly how the satisfying assignments depend on fermionic parity for a simple illustrative 4-fermionic problem where one has a line of three $\Pi_e^{1,q}$ clauses on modes 1,2,3 and 4 as in Lemma 2.3.6, and one adds a single $\Pi_{(2,4)}^{02,q}$ clause between modes 2 and 4. This is an example of a non-hPNC cluster with vertex 2 having degree 3, i.e. an example of Fig. 2.4(d).

This example gives insight into why product state assignments suffice for QUANTUM 2-SAT (with the exception of cases involving rank-3 projectors) but 4-fermion non-product, non-Gaussian states are needed for FERMIONIC 2-SAT. On a separate note: from Lemma 2.3.6 it is clear that if we ask for fixed particle number for QUANTUM 2-SAT, one *can* have a unique satisfying assignment which is not a product state, i.e. a fully permutation-symmetric state of n_q qubits with a fixed number N_q of 1s (with N_q unequal to 0 or the maximum n_q).

The instance we consider here consists of 4 fermionic modes with projectors $\Pi_{(1,2)}^{1,q}$, $\Pi_{(2,3)}^{1,q}$, $\Pi_{(3,4)}^{1,q}$ which we give in qubit language. Specifically, we take $\Pi_{(1,2)}^{\text{qubit}}=|\Psi^-\rangle\langle\Psi^-|_{12}$,

$\Pi^{\text{qubit}}_{(2,3)} = |\Psi^-\rangle\langle\Psi^-|_{23}, \Pi^{\text{qubit}}_{(3,4)} = |\Psi^-\rangle\langle\Psi^-|_{34}$, with $|\Psi^-\rangle$ denoting a singlet state. For a fixed particle number $0 \le N \le 4$, the permutation symmetric state with N excitations is the unique satisfying assignment (i.e., ground state of the ferromagnetic Heisenberg model) via Lemma 2.3.6. For each N, there is thus a satisfying assignment, i.e.

$$\begin{aligned}
|\psi_{N=0}\rangle &= |0000\rangle \\
|\psi_{N=1}\rangle &= |0001\rangle + |0010\rangle + |0100\rangle + |1000\rangle\,, \\
|\psi_{N=2}\rangle &= |1100\rangle + |0011\rangle + |1010\rangle + |1001\rangle + |0110\rangle + |0101\rangle \\
|\psi_{N=3}\rangle &= |0111\rangle + |1011\rangle + |1101\rangle + |1110\rangle\,, \\
|\psi_{N=4}\rangle &= |1111\rangle\,.
\end{aligned} \tag{2.31}$$

Now image we wish to add a fourth fermionic $\Pi^{02,q}_{(2,4)}$-type projector with amplitudes $\alpha_{(2,4)} = \delta_{(2,4)} = 1/\sqrt{2}$ which, in qubit language —note the additional Z_3 due to the Jordan-Wigner transformation in equation (2.2)—, equals

$$\begin{aligned}
\Pi^{\text{qubit}}_{(2,4)} &= \frac{1}{2}\Big(\sigma_2^-\sigma_2^+\sigma_4^-\sigma_4^+ + \sigma_2^- Z_3\sigma_4^- + \sigma_2^+ Z_3\sigma_4^+ + \sigma_2^+\sigma_2^-\sigma_4^+\sigma_4^-\Big) \\
&= \frac{1}{2}\Big(|00\rangle\langle 00|_{24} + Z_3[|00\rangle\langle 11|_{24} + |11\rangle\langle 00|_{24}] + |11\rangle\langle 11|_{24}\Big).
\end{aligned} \tag{2.32}$$

From the expression for $\Pi^{\text{qubit}}_{(2,4)}$, it is clear that an assignment $|\psi\rangle$ can only be a satisfying assignment if it is of the form $|\psi_{\text{even}}\rangle = a_0|\psi_{N=0}\rangle + a_2|\psi_{N=2}\rangle + a_4|\psi_{N=4}\rangle$ (even parity) or $|\psi_{\text{odd}}\rangle = a_1|\psi_{N=1}\rangle + a_3|\psi_{N=3}\rangle$ (odd parity). Let us see which (if any) of these states is projected to zero by $\Pi^{\text{qubit}}_{(2,4)}$.

$$\begin{aligned}
\Pi^{\text{qubit}}_{(2,4)}|\psi_{\text{even}}\rangle = \frac{1}{2}\Big(&a_0(|0000\rangle + |0101\rangle) + a_2(|1010\rangle - |1111\rangle + |0000\rangle + |0101\rangle) \\
&+ a_4(-|1010\rangle + |1111\rangle)\Big), \\
\Pi^{\text{qubit}}_{(2,4)}|\psi_{\text{odd}}\rangle = \frac{1}{2}\Big(&a_1(|0010\rangle - |0111\rangle + |1000\rangle + |1101\rangle) \\
&+ a_3(-|0010\rangle + |0111\rangle + |1000\rangle + |1101\rangle)\Big),
\end{aligned} \tag{2.33}$$

where the different signs are caused by Z_3 in $\Pi^{\text{qubit}}_{(2,4)}$. Clearly, $\Pi^{\text{qubit}}_{(2,4)}|\psi_{\text{even}}\rangle = 0$ for $a_0 = -a_2$ and $a_2 = a_4$, and $\Pi^{\text{qubit}}_{(2,4)}|\psi_{\text{odd}}\rangle$ cannot be zero. So, $|\psi_{\text{even}}\rangle$ with $a_0 = -a_2$ and $a_2 = a_4$ is the unique satisfying assignment of $\{\Pi^{\text{qubit}}_{(1,2)}, \Pi^{\text{qubit}}_{(2,3)}, \Pi^{\text{qubit}}_{(3,4)}, \Pi^{\text{qubit}}_{(2,4)}\}$, which is clearly not a product state. Let us set $a_0 = 1/2\sqrt{2}$, $a_2 = -1/2\sqrt{2}$ and $a_4 = -1/2\sqrt{2}$ wlog (so that $|\psi_{\text{even}}\rangle$ is normalized).

Back in fermionic language, we thus have a satisfying assignment

$$|\psi_{\text{even,f}}\rangle = \frac{1}{2\sqrt{2}}\Big(I - (a_1^\dagger a_2^\dagger + a_3^\dagger a_4^\dagger + a_2^\dagger a_3^\dagger + a_1^\dagger a_4^\dagger + a_2^\dagger a_4^\dagger + a_1^\dagger a_3^\dagger) - a_1^\dagger a_2^\dagger a_3^\dagger a_4^\dagger\Big)|\text{vac}\rangle\,. \tag{2.34}$$

To prove that this is a non-Gaussian state, we argue as follows. A pure Gaussian state has a covariance matrix $\Gamma \in \mathbb{R}^{8\times 8}$ (with entries in Eq. (1.18)) with orthonormal columns as

$\Gamma^T\Gamma = I$. To show that this does not hold for our state, let's evaluate the first column of Γ ($\Gamma_{j,1}$ for $j = 1,2,\ldots,8$).

By definition, we have that $\Gamma_{1,1} = 0$. Furthermore, $\Gamma_{2,1} = 2\langle\psi_{\text{even,f}}| a_1^\dagger a_1 |\psi_{\text{even,f}}\rangle - 1$, which can be simply seen to equal 0. For the other entries, let us distinguish between

$$\begin{aligned}\Gamma_{2j-1,1} &= i\langle\psi_{\text{even,f}}|(a_j + a_j^\dagger)(a_1 + a_1^\dagger)|\psi_{\text{even,f}}\rangle, \text{ with odd index } 2j-1,\\ \Gamma_{2j,1} &= i\langle\psi_{\text{even,f}}| i(a_j - a_j^\dagger)(a_1 + a_1^\dagger)|\psi_{\text{even,f}}\rangle, \text{ with even index } 2j,\end{aligned} \tag{2.35}$$

for $j = 2,3,4$. For either type of entry, we have to evaluate the following expectation values, which can be done straightforwardly.

$$\langle\psi_{\text{even,f}}| a_j a_1^\dagger |\psi_{\text{even,f}}\rangle = \begin{cases} -1/4 & \text{if } j=2,\\ 0 & \text{if } j=3,\\ 1/4 & \text{if } j=4.\end{cases} \quad \langle\psi_{\text{even,f}}| a_j^\dagger a_1 |\psi_{\text{even,f}}\rangle = \begin{cases} 1/4 & \text{if } j=2,\\ 0 & \text{if } j=3,\\ -1/4 & \text{if } j=4.\end{cases}$$

$$\langle\psi_{\text{even,f}}| a_j^\dagger a_1^\dagger |\psi_{\text{even,f}}\rangle = \begin{cases} 0 & \text{if } j=2,\\ 1/4 & \text{if } j=3,\\ 0 & \text{if } j=4.\end{cases} \quad \langle\psi_{\text{even,f}}| a_j a_1 |\psi_{\text{even,f}}\rangle = \begin{cases} 0 & \text{if } j=2,\\ -1/4 & \text{if } j=3,\\ 0 & \text{if } j=4.\end{cases} \tag{2.36}$$

Using these expressions, we conclude the following.

$$\Gamma_{2j-1,1} = 0 \text{ for } j = 2,3,4 \quad \text{and} \quad \Gamma_{2j,1} = \begin{cases} 1/2 & \text{if } j=2,\\ 1/2 & \text{if } j=3,\\ -1/2 & \text{if } j=4.\end{cases} \tag{2.37}$$

Thus the first column of Γ is $(0,0,0,1/2,0,1/2,0,-1/2)^T$, which has 2-norm $\sqrt{3/4} < 1$, so $|\psi_{\text{even,f}}\rangle$ is not Gaussian.

Let's compare our findings briefly with an equivalent set of QUANTUM 2-SAT constraints: Will there be a product assignment here? The QUANTUM 2-SAT equivalent of the FERMIONIC 2-SAT instance considered here is $\left\{\Pi^{\text{qubit}}_{(1,2)} = |\Psi^-\rangle\langle\Psi^-|_{12}, \Pi^{\text{qubit}}_{(2,3)} = |\Psi^-\rangle\langle\Psi^-|_{23}, \Pi^{\text{qubit}}_{(3,4)} = |\Psi^-\rangle\langle\Psi^-|_{34}, \Pi^{\text{qubit}}_{(2,4)} = \frac{1}{2}(|00\rangle + |11\rangle)(\langle 00| + |11\rangle)_{24}\right\}$ (note the absence of the Jordan-Wigner Z_3). This QUANTUM 2-SAT instance indeed has a product state satisfying assignment, namely $\bigotimes_{j=1}^4 \frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle)_j$.

Our pedestrian findings here are captured by the statement for $n_q = 4$ in Corollary 2.3.5 point 5. The degree-3 vertex in this case is vertex 2, the particle-hole-transformation acts on modes $1,3,4$ and changes the parity of the non-Gaussian even parity state in equation (2.34) to an odd parity state, while an even parity state in the particle-hole transformed basis (hence odd particle here) is disallowed. The exclusion of one of the two parities for this FERMIONIC 2-SAT instance is essentially what causes the unique satisfying assignment to be non-product and non-Gaussian.

2.C. Classical 2-SAT with fixed parity

In this section, we prove that classical 2-SAT can be solved efficiently, even when constraining the Hamming weight parity of the assignment. This is not necessarily

a trivial problem as they are, for example, simple examples of 2-SAT instances with exponentially many solutions, *all* with the same parity. Consider for instance

$$\bigwedge_{j=0}^{n/4-1} (x_{j+1} \vee \overline{x}_{j+2}) \wedge (x_{j+2} \vee \overline{x}_{j+3}) \wedge (x_{j+3} \vee \overline{x}_{j+4}) \wedge (\overline{x}_{j+1} \vee x_{j+4}), \tag{2.38}$$

with n a multiple of four. This instance corresponds to disjoint units of four variables, where the only satisfying assignments on each unit are the all-zeros or the all-ones assignment. Clearly, there are exponentially many satisfying assignments, and they are all of even parity. In fact, there are even simple examples of *connected* 2-SAT instances with exponentially many satisfying assignments, all with the same parity. Thus even if one is guaranteed that a 2-SAT instance has (exponentially) many satisfying assignments, the question of whether there is a satisfying assignment with a given parity is a nontrivial one. We prove the following.

Theorem 2.C.1. *Given an instance of classical 2-SAT on n variables with m clauses, and a parity $P \in \{-1,+1\}$. Decide whether there exists a satisfying assignment x with Hamming weight parity P (YES), or there is no satisfying assignment with Hamming weight parity P (NO). This problem can be decided in time $O(nm)$.*

Proof. First, we find a satisfying assignment in time $O(n+m)$ [1] for the 2-SAT instance, if it is satisfiable. If this solution has Hamming weight parity P, then we are done (output YES), else we proceed. For convenience, let us redefine the 2-SAT instance such that the obtained solution is the all-zeros assignment. Then what is left is to check whether there exists an odd Hamming weight satisfying assignment for the redefined instance. The redefined instance wlog consists of clauses $(x_i \vee \overline{x}_j)$ (with $i < j$ or $i > j$) and $(\overline{x}_i \vee \overline{x}_j)$ (with $i < j$ wlog). Note that there can be multiple clauses per pair of variables i, j.

Let us consider the sub-graph G_{sub} consisting of just the clauses of type $(x_i \vee \overline{x}_j)$ (with $i < j$ or $i > j$). We associate a directed edge $i \to j$ with each clause $(x_i \vee \overline{x}_j)$ (i.e., edges point from "variable" i to "negated variable" j). Note that two variables i, j can simultaneously be connected by a $(x_i \vee \overline{x}_j)$ clause and a $(x_j \vee \overline{x}_i)$ clause, resulting in a $i \to j$ edge and a $j \to i$ edge in G_{sub}. Next, we identify the strongly-connected components (SCC's) of the directed graph G_{sub} in time $O(n+m)$ [18]. The only satisfying assignments on these SCC's are the all-zeros and the all-ones assignments, see Lemma 2.C.2 below. Hence we compress each SCC into a single Boolean variable. We label the new collection of variables by $\{y_j\}_{j=1}^{\leq n}$, where some were previously single variables and others are compressed SCC's. Since each directed cycle is (part of) an SCC, the compressed problem cannot contain any directed cycles, i.e., it is a directed acyclic graph. The graph with $(y_i \vee \overline{y}_j)$-type edges can be topologically sorted in time $O(n+m)$ [19]. The topologically-sorted graph now only includes edges of type $(y_i \vee \overline{y}_j)$ for which $i < j$. Next, let us add the $(\overline{x}_i \vee \overline{x}_j)$-type clauses, which after compression have become either $(\overline{y}_i \vee \overline{y}_j)$ clauses or self-edges $(\overline{y}_i \vee \overline{y}_i)$. The former can be added s.t. $i < j$ in the topologically sorted graph (since, obviously, $(\overline{y}_i \vee \overline{y}_j)$ and $(\overline{y}_j \vee \overline{y}_i)$ are equivalent) to construct the new graph $\tilde{G}$, see Fig. 2.C.2. The self-edges of type $(\overline{y}_i \vee \overline{y}_i)$ exclude $y_i = 1$. These variables with self-edges can be flagged in time $O(n+m)$.

With each variable j in the compressed problem, we associate a weight $w_j \in \{0,1\}$. The weight w_j corresponds to the number (modulo 2) of original variables in the SCC that has

been compressed into variable y_j. Note that for any uncompressed variable, we assign weight 1. To obtain a satisfying assignment of odd Hamming weight (if it exists), we need to flip an odd number of odd-weight variables. In order to do so, we might have to simultaneously flip some even-weight variables.

Importantly, if we flip a variable j (with some weight w_j) to $y_j = 1$, then, by construction, all clauses to variables in *later* layers of the topological sort are still satisfied. Clauses from variable j to variables in *earlier* layers can now become unsatisfied and thus some variables i s.t. $i < j$ might also have to be set to $y_i = 1$. Similarly, flipping such a variable i can only lead to unsatisfied clauses that connect i to variables in earlier layers of the topological sort, not to variables in later layers.

The algorithm runs by first identifying the earliest odd-weight variable j in the topological sort. For some instances, this *might* already be in the first layer. We set $y_j = 1$ and see whether the propagation to earlier layers of the topological sort does not lead to any contradiction. If it does not, then we have found a satisfying assignment of flipped parity (output YES), since all variables that are flipped in earlier layers due to propagation are even-weight variables by definition. If the propagation leads to contradiction, then clearly there is *no* satisfying assignment for which $y_j = 1$. Next, we reset to the all-zeros assignment and identify the next odd-weight variable k in the topological sort (which is in the same or in a later layer). Again, one checks whether setting $y_k = 1$ leads to contradiction in earlier layers. If it does not, then we have constructed a satisfying assignment with flipped parity (output YES), since y_k is the *only* odd-weight variable set to 1 in that assignment. Indeed, if y_j would also have to be set to 1, then there would be a contradiction. If setting $y_k = 1$ does lead to contradiction, then there is *no* satisfying assignment for which $y_k = 1$. We proceed again by resetting to the all-zeros assignment and identifying the next odd-weight variable in the topological sort. We iterate over all odd-weight variables in this manner and either obtain a satisfying assignment of flipped parity or conclude that no odd-weight variable can be set to 1 consistently and hence there is no satisfying assignment with flipped parity (output NO). Iterating over all odd-weight variables and checking for the consistency of flipping them takes time $O(nm)$. □

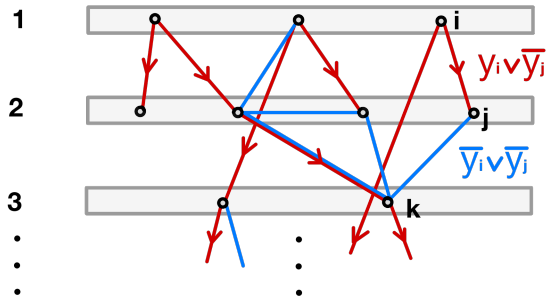

Figure 2.6: First three layers of a compressed topologically sorted graph $\tilde{G}$ derived from G, where each variable j has been given a weight w_j. The red clauses are of type $(y_i \vee \overline{y}_j)$ (with $i < j$) and the blue clauses are of type $(\overline{y}_i \vee \overline{y}_j)$.

Lemma 2.C.2. *Given an instance of classical 2-SAT with $(x_i \vee \overline{x}_j)$-type clauses only. Let us consider the directed graph with edges $i \to j$ for each clause $(x_i \vee \overline{x}_j)$ (i.e., edges pointing from "variable" i to "negated variable" j). If this graph is strongly connected, then only the all-zeros and all-ones assignments are satisfying assignments.*

Proof. If the directed graph is strongly connected, then for *each* pair of variables (i, j) there is path from i to j and a path from j to i. Let us set $x_i = 0$. Then clauses along the path of directed edges from variable i to variable j can only be satisfied if $x_j = 0$. Conversely, let us set $x_i = 1$ and let us now travel "upstream" to variable j. The clauses along the path of directed edges from variable j to variable i can only be satisfied if $x_j = 1$. Hence the lemma statement follows. □

Remark 2.C.3. *The algorithm in Theorem 2.C.1 cannot be used to solve the* NP*-complete problem of finding an assignment with fixed Hamming weight (particle number) since one can generally not appropriately change the particle number by flipping a polynomial set of designated variables. Similarly, the algorithm above cannot be used to count the number of solutions to a 2-SAT instance (which is a #P-complete problem), although there is a way of enumerating all solutions with effort growing with the number of solutions, which uses similar techniques as our algorithm [20].*

2.D. Proof of Lemma 2.6.1

Proof. Containment in QMA can be obtained by simply mapping the n-mode fermionic problem onto a n-qubit problem using e.g. the Jordan-Wigner transformation in equation (2.2). This means that each projector Π_i using k fermionic modes is mapped to a projector which is a sum of terms, each of which has a k-local part to which a string of Pauli Zs of some length is appended. Hence the FERMIONIC k-SAT problem maps onto a local Hamiltonian problem but some Pauli strings in the Hamiltonian have weight larger than k. However, one can still apply standard QMA-arguments for the proof verification as in Section IV in [21] since this method only relies on the fact that the number of Pauli terms in the Hamiltonian is poly(n).

To prove hardness, we can map the QMA_1-complete problem in [5] onto a FERMIONIC $k = 9$-SAT problem using the unitary "assimilation mapping" described in Lemma 15 in [12]. We take the 3-local circuit Hamiltonian, with its 3-local projectors, on some n qubits in the QMA_1-proof and view it as a Hamiltonian on n qubits plus $n/2$ fermionic modes which gets mapped on a space of $3n/2$ fermionic modes and the Pauli operators of one qubit get mapped on products of two (Majorana) fermion operators out of three fermionic modes. Thus each 3-local projector becomes a fermionic projector Π_i involving at most 9 fermionic modes which conserves fermionic parity. In the YES case, if there is a n-qubit satisfying assignment for the QUANTUM 3-SAT problem, it can be rewritten as a $3n/2$-fermionic state which is a satisfying assignment for the FERMIONIC 9-SAT problem. In the NO case, if there is no approximate satisfying assignment for QUANTUM 3-SAT, imagine that there is a fermionic state $|\Phi\rangle$ for the FERMIONIC 9-SAT problem for which $\sum_i \langle\Phi|\Pi_i|\Phi\rangle = \beta$. Then performing the inverse mapping and tracing over the additional fermionic registers in $|\Phi\rangle\langle\Phi|$ to get the qubit reduced density matrix ρ^{qubit}, this implies that $\sum_i \text{Tr}\,\Pi_i^{\text{qubit}}\rho^{\text{qubit}} = \beta$ which by assumption implies that $\beta \geq 1/\text{poly}(n)$. □

3

Optimizing fermionic Hamiltonians with classical interactions

The players tried for a forward pass, with the jester on the sidelines in a cast
Now the halftime air was sweet perfume, while sergeants played a marching tune
We all got up to dance, oh but we never got the chance
'Cause the players tried to take the field, the marching band refused to yield
Do you recall what was revealed, the day the music died

American Pie (1971), Don McLean

3.1. Introduction

In this paper we study energy optimization, or ground energy search, for fermionic Hamiltonians. Mathematically, it means finding the largest[1] eigenvalue of a 2^n-dimensional Hermitian matrix, which is a low-degree polynomial in fermionic creation and annihilation operators $\{a_j^\dagger, a_j\}_{j\in[n]}$:

$$a_j^\dagger a_k + a_k a_j^\dagger = \delta_{jk}, \quad a_j a_k + a_k a_j = 0, \quad n_j |\boldsymbol{x}\rangle = x_j |\boldsymbol{x}\rangle, \tag{3.1}$$

where $n_j := a_j^\dagger a_j$ and $|\boldsymbol{x}\rangle$ for $\boldsymbol{x} = (x_1, \ldots, x_n) \in \{0,1\}^n$ are the computational basis states. Energy optimization is one of the key computational problems in many-body physics and appears in a number of contexts in condensed matter physics and quantum chemistry. In general, it is a QMA-hard optimization task [1, 2]; in numerical practice, it is being accomplished with a number of approximation tools [3–7]. An interesting goal mathematically is to give rigorous performance guarantees for such approximate methods.

This text focuses on approximating the highest energy state with a Gaussian (i.e., free-fermionic) state [8–10]. Generally, Gaussian states are defined as the Gibbs states of

[1] We flip the sign convention of the Hamiltonian to align with that used in computer science.

Hamiltonians which are quadratic in $\{a_j^\dagger, a_j\}_{j\in[n]}$; they admit a classically efficient description in terms of a $2n$-sized *covariance matrix* (morally analogous to the stabilizer tableau for stabilizer states). In the domain of computational many-body physics, the standard method for finding Gaussian ground state approximations is generalized Hartree-Fock, which is heuristic [3, 4]. But in recent years, also *rigorous* guarantees for Gaussian ground state approximations (or lack thereof) have started to appear [11–18]. These use a common metric for any optimization method — *approximation ratio*, i.e., a guarantee on the ratio between the energy of the state obtained by a method, and the true ground energy[2]. It was discovered in [12, 13] that for *general* fermionic Hamiltonians, Gaussian states cannot yield ground energy even up to a constant approximation ratio. This dramatic effect – let us call it Gaussian (approximation) breakdown – was demonstrated for the Sachdev-Ye-Kitaev model. It was later extended to some other models which share the feature of having all-to-all, or at least non-sparse, fermion couplings [14, 16–18]. This breakdown can be viewed as a heuristic warning sign for optimization of general quantum chemistry Hamiltonians, as those are also strongly interacting and lack sparsity [17, 18]. On the other hand, it has been observed that the Hartree-Fock technique yields high approximation ratios in numerical practice [4, 12]. Rigorously speaking, it had not been settled if quantum chemistry Hamiltonians exhibit Gausisan breakdown.

3.2. Main results

Our work is motivated by a key difference between the Hamiltonians analyzed in Refs. [11–14, 16–18] and those arising in quantum chemistry. In particular, the quartic terms in real-space discretized chemistry Hamiltonians are not generic but *classical*, i.e., diagonal in computational basis [4, 6, 19–22]. This holds because chemistry interactions physically arise from Coulomb terms, built out of diagonal particle density operators $n(\boldsymbol{r})$. Using CIFH as an acronym for 'classically interacting fermionic Hamiltonians', we define

Problem 3.2.1 (TRACELESS CIFH OPTIMIZATION)**.** *Consider the Hamiltonian*

$$H = \sum_{(j,k)\in E} w_{j,k}(\mathbb{1}/4 - n_j n_k) + \sum_{j\in V} \mu_j\big(n_j - \mathbb{1}/2\big) + \sum_{(j,k)\in E'} w'_{j,k}\big(-a_j^\dagger a_k - a_k^\dagger a_j\big) \tag{3.2}$$

with $w_{j,k} \geq 0$, $w'_{j,k} \in \mathbb{R}$, and $\mu_j \in \mathbb{R}$, and vertex set V and edge sets E and E'. Compute $\lambda_{\max}(H) = \max_{\rho\in\mathbb{C}^{2^n\times 2^n}} \Big\{ \mathrm{tr}\big(\rho H\big) \text{ s.t. } \rho \succeq 0, \mathrm{tr}\big(\rho\big) = 1 \Big\}$.

Solving this problem with $1/\mathrm{poly}(n)$ precision is QMA-hard by adaptation of a result from [2], see Appendix 3.A. In the absence of the quadratic 'hopping' terms, $w'_{j,k} = 0$, this is a classical QUBO optimization problem [23]. As the main result of this work, we show, in Section 3.4, that

Theorem 3.2.2. *There exists a pure fermionic Gaussian state ρ that achieves an approximation ratio $\frac{1}{3}$ for* TRACELESS CIFH OPTIMIZATION *(Problem 3.2.1).*

[2] In the fermionic optimization literature, this ratio has been made well-defined by considering traceless Hamiltonians.

This implies that quantum chemistry, unlike general fermionic Hamiltonians, does not exhibit a Gaussian breakdown —even when the Hamiltonian is non-sparse (possibly dense).

Proof sketch. Split the Hamiltonian of equation (3.2) as $H = H_{\text{quad}} + H_{\text{class}}$, with the off-diagonal quadratic part H_{quad} and the diagonal (classical) part H_{class},

$$H_{\text{quad}} = \sum_{(j,k)\in E'} w'_{j,k}\big(-a_j^\dagger a_k - a_k^\dagger a_j\big), \tag{3.3}$$

$$H_{\text{class}} = \sum_{(j,k)\in E} w_{j,k}(\mathbb{1}/4 - n_j n_k) + \sum_{j\in V} \mu_j\big(n_j - \mathbb{1}/2\big). \tag{3.4}$$

Each of these Hamiltonians have Gaussian ground states, $\rho_{\max}(H_{\text{quad}})$ and $\rho_{\max}(H_{\text{class}})$, as the computational basis states are Gaussian. Therefore, if either H_{quad} or H_{class} is negligible in operator norm, some constant-ratio Gaussian solution can readily be obtained by choosing one of these states. To obtain the stronger Theorem 3.2.2, which guarantees the constant ratio of $\frac{1}{3}$ regardless of the relative size of H_{quad} and H_{class}, a few more steps are needed. One is to observe that $\rho_{\max}(H_{\text{class}})$ actually vanishes on H_{quad}, as we chose it to be off-diagonal. On the flip side, one can modify $\rho_{\max}(H_{\text{quad}})$, such that H_{class} only contributes to its energy non-negatively. This step is more technical; the key is to modify the covariance matrix of $\rho_{\max}(H_{\text{quad}})$ such that its first off-diagonal elements are removed. This can be done while preserving the validity —and Gaussianity— of the state. Choosing either thus modified solution $\rho_{\max}(H_{\text{quad}})$, or $\rho_{\max}(H_{\text{class}})$, allows to guarantee the approximation ratio of $\frac{1}{3}$ at the worst.

The Gaussian states which *exist* by Theorem 3.2.2 should not in general be efficiently constructable. In fact, finding *any* constant-ratio approximation in poly-time (for the classical problem with $H_{\text{quad}} = 0$) was ruled out under mild complexity-theory assumptions [24, 25]. But in structured cases this is achievable. Indeed, in Section 3.5.1 we show

Theorem 3.2.3. *There is a deterministic polynomial-time algorithm that outputs a fermionic Gaussian state ρ achieving approximation ratio $\frac{1}{3}$ for* TRACELESS CIFH OPTIMIZATION *(Problem 3.2.1), provided that the graph $G_{class} = (V,(w,E))$ is bipartite.*

A simple example of such bipartite interaction graph is a Fermi-Hubbard model with an onsite interaction between spin-up and spin-down electrons, so that the bi-partition is between spin-up and spin-down modes (note that the hopping Hamiltonian remains unconstrained). Problem 3.2.1 with a bipartite interaction graph stays QMA-hard, and does not need to be sparse.

The key to proving Theorem 3.2.3 is that the global optimum of H_{class} can be efficiently found, using a linear program which exploits the bipartite structure. This solution can be then used to give a constant-ratio approximating Gaussian, similarly to that of Theorem 3.2.2. We show that this Gaussian is a feasible solution for an $O(n)$-dimensional semi-definite program. This program accounts for H_{quad} in its objective and for H_{class} in its constraints. In practice it yields states with better approximation ratios, and is an interesting subject for future study. We note that if G_{class} had some other structure that allowed the optimum of H_{class} to be efficiently obtained, then Theorem 3.2.3 would carry over to those cases as well.

Tracelessness is one of two main conventions which make the approximation ratio well-defined. The other common option is to make every term of the Hamiltonian positive semi-definite, motivating

Problem 3.2.4 (POSITIVE SEMI-DEFINITE CIFH OPTIMIZATION)**.** *Consider the Hamiltonian*

$$H = \sum_{(j,k)\in E} w_{j,k}\big(\mathbb{1} - n_j n_k\big) + \sum_j \mu_j n_j + \sum_{(j,k)\in E'} w'_{j,k}\big(\mathbb{1} - a_j^\dagger a_k - a_k^\dagger a_j\big) \succeq 0, \tag{3.5}$$

with $w_{j,k} \geq 0$, $w'_{j,k} \in \mathbb{R}$, *and* $\mu_j \geq 0$ *and vertex set* V, *and edge sets* E *and* E'. *Compute* $\lambda_{\max}(H) = \max_{\rho\in\mathbb{C}^{2^n\times 2^n}} \Big\{ \mathrm{tr}\big(\rho H\big)$ *s.t.* $\rho \succeq 0, \mathrm{tr}\big(\rho\big) = 1 \Big\}$.

In case $\forall (j,k)$, $w'_{j,k} = 0$ and $\mu_j = 0$, Problem 3.2.4 is the weighted Max Cut problem. A special case of Problem 3.2.4 is FERMIONIC MAX CUT, see Section 3.5.3, which, when the graph is a line, coincides with (weighted) QUANTUM MAX CUT [26]. Unlike in the traceless case, here the goal is guaranteeing not just a constant approximation ratio, but one that is substantially better than that guaranteed by a fully mixed state (in this case $\frac{1}{2}$). In Section 3.5.2, we ask: can our methods give an interesting Gaussian approximation to this 'positive semidefinite' type of optimization? We find

Theorem 3.2.5. *There is a polynomial-time algorithm that with probability* $\Omega(1)$ *outputs a Gaussian state* ρ *that achieves an approximation ratio* 0.637 *for* PSD CIFH OPTIMIZATION *(Problem 3.2.4).*

This state can be found using a semi-definite program similar to that implied in Theorem 3.2.3, and the ratio is guaranteed by a similarly constructed feasible solution. For the classical part of the solution, given a lack of structure, we adapt the Goemans-Williamson approach [27]. In more structured settings, better algorithms for optimizing H_{class}, see e.g. [23], could improve the approximation ratio in Theorem 3.2.5.

The key technical contribution of our work is the concept of a *Gaussian blend*, introduced in Section 3.3.3. A Gaussian blend is defined as a Gaussian state whose covariance matrix is a weighted combination of covariance matrices of several input Gaussian states. The modification of the state $\rho_{\max}(H_{\text{quad}})$, hinted at in the proof sketch of Theorem 3.2.2, is in fact given as a Gaussian blend between two states. A more complex Gaussian blend is key to addressing the following problem,

Problem 3.2.6 (q-PARTICLE TRACELESS CIFH OPTIMIZATION)**.** *Consider the Hamiltonian*

$$H = \sum_{(j,k)\in E} w_{j,k}(\mathbb{1}/4 - n_j n_k) + \sum_{(j,k)\in E'} w'_{j,k}\big(-a_j^\dagger a_k - a_k^\dagger a_j\big) \tag{3.6}$$

with $w_{j,k} \geq 0$, $w'_{j,k} \in \mathbb{R}$, *and vertex set* V, *and edge sets* E *and* E'. *Compute*

$$\lambda_{\max,\langle q\rangle}(H) = \max_{\rho\in\mathbb{C}^{2^n\times 2^n}} \Big\{ \mathrm{tr}\big(\rho H\big) \text{ s.t. } \mathrm{tr}\big(\rho \hat{N}\big) = q,\ \rho \succeq 0,\ \mathrm{tr}\big(\rho\big) = 1 \Big\}, \tag{3.7}$$

with $q \in \{0, 1, \ldots, \lfloor n/2 \rfloor\}$.

In this problem, $\hat{N} := \sum_{j\in V} a_j^\dagger a_j$ is the total particle number operator. Essentially, this is Problem 3.2.1 with a constraint that the particle number is equal to q in expectation. This type of an optimization task is inspired by quantum chemistry and condensed matter theory: there the number of fermions is fundamentally a conserved quantity, which is often fixed by the physical setup.

In Section 3.6, we show

Theorem 3.2.7. *If $G_{class} = (V,(w,E))$ is bipartite and $q \le \lfloor n/2 \rfloor$, then there exists a fermionic Gaussian state ρ that achieves an approximation ratio $\frac{1}{2\left((n-2q)/n+3/2\right)}$ for Problem 3.2.6. Such a state can be obtained in polynomial time.*

Due to the constraining nature of the problem, the proof of Theorem 3.2.7 involves additional technicalities compared to that of Theorem 3.2.3. The semi-definite program which is used to produce the desired state, now includes the q-particle condition as a linear constraint; the provided feasible solution is a Gaussian blend involving $\rho_{\max}(H_{\text{class}})$, $\rho_{\max}(H_{\text{quad}})$, and a third, auxiliary state. Here, the particular condition of G_{class} being bipartite is more essential than in Theorem 3.2.3: in addition to being used in the efficient algorithm for the optimization of H_{class}, the bipartite structure is used (in a different way) in our proof that the constructed Gaussian state satisfies the q-particle constraint.

Besides proving Gaussian approximation ratios, we also give an argument in Appendix 3.B which shows that there are instances of traceless fermionic Hamiltonians with classical interactions where the Gaussian approximation ratio is upper-bounded away from 1 by a constant. Improving such upper-bounding techniques further is an interesting direction for future research.

3.3. Preliminaries

In this work, we will be using preliminaries presented in the current section in addition to those presented in Section 1.5.

The following fact on correlations in Slater determinant states (see Definition 1.5.6) will be used later on.

Lemma 3.3.1. *For the density matrix of a Slater determinant state $\rho = |\psi\rangle\langle\psi|$ one has*

$$\forall\, j,k \in [n], \quad i\operatorname{tr}\left(\rho c_{2j-1}c_{2k}\right) = -i\operatorname{tr}\left(\rho c_{2j}c_{2k-1}\right), i\operatorname{tr}\left(\rho c_{2j-1}c_{2k-1}\right) = i\operatorname{tr}\left(\rho c_{2j}c_{2k}\right). \quad (3.8)$$

Proof. Since Slater determinant states ρ are eigenstates of the particle number operator $\hat{N}$, $\operatorname{tr}\left(\rho a_j a_k\right) = \operatorname{tr}\left(\rho a_j^\dagger a_k^\dagger\right) = 0$. Using the definition of Majorana operators in equation (1.4), one has

$$\begin{aligned}
ic_{2j-1}c_{2k} &= -a_j a_k + a_j a_k^\dagger - a_j^\dagger a_k + a_j^\dagger a_k^\dagger,\\
ic_{2j}c_{2k-1} &= -a_j a_k - a_j a_k^\dagger + a_j^\dagger a_k + a_j^\dagger a_k^\dagger,\\
ic_{2j-1}c_{2k-1} &= i\left(a_j a_k + a_j a_k^\dagger + a_j^\dagger a_k + a_j^\dagger a_k^\dagger\right),\\
ic_{2j}c_{2k} &= i\left(-a_j a_k + a_j a_k^\dagger + a_j^\dagger a_k - a_j^\dagger a_k^\dagger\right),
\end{aligned} \quad (3.9)$$

from which the claim follows. □

3.3.1. Optimization over Gaussian states

Due to properties of the covariance matrix of a Gaussian fermionic state, the optimization of a general Hermitian (traceless) fermionic Hamiltonian H with quadratic and quartic terms in $\{c_i\}$ over the set of Gaussian fermionic states can be formulated as an optimization of the form [11]

$$F(\Gamma) = \max_{\Gamma^T\Gamma\leq\mathbb{1},\Gamma\in\mathcal{L}} \sum_{ijkl} W_{ijkl}\Gamma_{i,j}\Gamma_{k,l} + \sum_{ij} V_{ij}\Gamma_{i,j}, \tag{3.10}$$

with real fully anti-symmetric V_{ij} and W_{ijkl}. Here $\mathcal{L}$ is the space of real anti-symmetric $2n\times 2n$ matrices, obeying the condition $\Gamma^T\Gamma\leq\mathbb{1}$. Using the anti-symmetry of Γ, this is equivalent to $i\Gamma\leq\mathbb{1}$ as formulated in Proposition 1.5.7. It was shown in [11] that one can rewrite the linear term in $\Gamma_{i,j}$ as part of the quadratic term, we also use the classical version of this trick in the proof of Lemma 3.5.4. Hence, the general problem of optimizing over Gaussian fermionic states is that of a quadratic optimization problem over a convex set of covariance matrices, see also Lemma 3.3.4. Due to Remark 1.5.4 this optimum is achieved for a pure Gaussian state, one for which $\Gamma^T\Gamma=\mathbb{1}$, i.e. the eigenvalues of $i\Gamma$ are ±1. This quadratic optimization problem is generally hard to solve: Ref. [11] has considered efficient approximate optimizations via known results in the literature.

3.3.2. Optimization of quadratic Hamiltonians as a semi-definite program

One can efficiently optimize quadratic Hamiltonians over Gaussian fermion states which includes additional linear constraints on the covariance matrix of the Gaussian fermionic state. This essentially follows from formulating the optimization as a semi-definite program (SDP), i.e. the quadratic term in Γ in equation (3.10) is absent. First, we prove

Lemma 3.3.2. *Any Hermitian matrix $X\in\mathbb{C}^{2n\times 2n}\succeq 0$ with linear constraints $\forall i\neq j, X_{i,j}=-X_{j,i}$ and $\forall i, X_{i,i}=1$, can be written as $X=\mathbb{1}+i\Gamma$, with Γ the covariance matrix of a fermionic Gaussian state.*

Proof. Since X is anti-symmetric on the off-diagonal and equal to 1 along the diagonal, we have that X w.l.o.g. equals $X=\mathbb{1}+iB$ with B a real-valued anti-symmetric matrix. Now let us use the following two facts. (1) The eigenvalues of a real-valued anti-symmetric matrix B come in $\pm i\lambda_j$ pairs (with $j\in[n]$), with each λ_j real-valued. (2) $X=\mathbb{1}+iB\succeq 0$. Therefore, there are no eigenvalues of B outside of the interval $[-i,+i]$. So, B is an anti-symmetric real-valued matrix with no eigenvalues outside of the $[-i,+i]$ interval. Through Proposition 1.5.7, B corresponds to a valid covariance matrix Γ of a fermionic Gaussian state. □

The standard form of a semi-definite program is

$$\begin{aligned} \text{maximize} \quad & \mathrm{Tr}(CX) \\ \text{subject to} \quad & \mathrm{Tr}(A_iX)=b_i, \quad i=1,\ldots,m, \\ & X\succeq 0 \end{aligned}$$

where C, A_i and X are Hermitian matrices and $\mathbf{b}\in\mathbb{R}^m$ with $m=\mathrm{poly}(n)$. Clearly, one can choose the set of feasible solutions of a SDP to be of the form $X=\mathbb{1}+i\Gamma$ by appropriately

choosing the equality constraints given by $\{A_i, b_i\}$ to match those in Lemma 3.3.2 so that Γ is the covariance matrix of a fermionic Gaussian state. Thus we can show the following

Lemma 3.3.3. *Given a quadratic Hamiltonian on n modes $H = \sum_{j,k} h_{j,k}\, i c_j c_k$, with real anti-symmetric matrix h. The Gaussian state ρ_{Gauss} that maximize $\mathrm{tr}\big(\rho_{Gauss} H\big)$ can be obtained by solving a semi-definite program, hence in polynomial time in n, also in the presence of poly(n) additional linear constraints on the covariance matrix Γ of ρ_{Gauss}.*

Proof. We have $\mathrm{tr}\big(\rho_{\mathrm{Gauss}} H\big) = \mathrm{tr}\big(h^T \Gamma\big) = -\mathrm{tr}(h\Gamma) = \mathrm{tr}(i h X) = \mathrm{tr}(CX)$ where $X = \mathbb{1} + i\Gamma$ is a feasible solution of the SDP capturing the properties in Lemma 3.3.2, and the matrix $C = ih$ is Hermitian. A polynomial number of linear constraints on Γ and thus X can be freely added to define the feasible set. □

3.3.3. Blending Gaussian states

Given m Gaussian states $\rho_1, \rho_2, \ldots, \rho_m$, their mixture $\sum_{j=1}^m p_j \rho_j$ (with $\sum_j p_j = 1$) obviously does not need to be a Gaussian state as Wick's theorem in equation (1.21) does not apply to such mixture. In Ref. [28] such general mixtures were called convex-Gaussian states. Here, we define a Gaussian state obtained by mixing the covariance matrices instead, to which we refer as the *blended* Gaussian state. It is straightforward to prove the following as we can cast the (convex) feasible set of a semi-definite program as the set of covariance matrices:

Lemma 3.3.4. *Given covariance matrices $\Gamma^1, \Gamma^2, \ldots, \Gamma^m$, there exists a fermionic Gaussian state with covariance matrix $\Gamma = \sum_{i=1}^m p_i \Gamma^i$, for any probability distribution $\{p_i\}_{i=1}^m$.*

Proof. Since $\Gamma^1, \ldots, \Gamma^m$ are covariance matrices, $\Gamma = \sum_i p_i \Gamma^i$ is also real-valued and anti-symmetric. Consequently, its eigenvalues come in $\pm i\lambda$ pairs, with $\lambda \in [-1, +1]$. Since $i\Gamma^1, \ldots, i\Gamma^q$ only have eigenvalues in $[-1, +1]$, we have that $\lambda_{\max}(i\Gamma) \leq \sum_i p_i \lambda_{\max}(i\Gamma^i) \leq \sum_i p_i = 1$. Therefore, Γ has no eigenvalues outside $[-i, +i]$ and is thus a covariance matrix. Through Proposition 1.5.7, we can associate a fermionic Gaussian state with Γ. □

3.4. Proof of Theorem 3.2.2: existence of constant-ratio Gaussian approximations.

Proof of Theorem 3.2.2. Let us denote the maximum energy classical eigenstate of H_{class} in equation (3.4) by ρ_{class} and the maximum energy eigenstate of H_{quad} in equation (3.3) by ρ_{quad}, and their covariance matrices be Γ^{class} and Γ^{quad} respectively. Both of these states are fermionic Gaussian states. Since ρ_{class} is classical, $\Gamma^{\mathrm{class}} = \bigoplus_{j=1}^n \begin{pmatrix} 0 & \lambda_j \\ -\lambda_j & 0 \end{pmatrix}$ with $\lambda_j \in \{\pm 1\}$. Note that one does not necessarily have an efficient algorithm to compute Γ^{class}: this will be addressed in Section 3.5. We construct a blended Gaussian state with covariance matrix

$$\Gamma = p_{\mathrm{class}} \Gamma^{\mathrm{class}} + \frac{1 - p_{\mathrm{class}}}{2} \big(\Gamma^{\mathrm{mediator}} + \Gamma^{\mathrm{quad}}\big), \tag{3.11}$$

and let ρ_{Gauss} be the associated fermionic Gaussian state. Here, the covariance matrix Γ^{mediator} is defined as

$$\begin{aligned}
\Gamma^{\text{mediator}}_{2j-1,2j} &= -\Gamma^{\text{quad}}_{2j-1,2j},\ \forall j \in V,\\
\Gamma^{\text{mediator}}_{2j,2j-1} &= -\Gamma^{\text{quad}}_{2j,2j-1},\ \forall j \in V,\\
\Gamma^{\text{mediator}}_{j,k} &= 0,\ \text{elsewhere.}
\end{aligned} \tag{3.12}$$

Since $\left|\Gamma^{\text{quad}}_{2j-1,2j}\right| \leq 1$, the eigenvalues of the anti-symmetric matrix Γ^{mediator} lie in $[-i,+i]$ and thus Γ^{mediator} is the covariance matrix of some fermionic Gaussian state via Proposition 1.5.7. Via Lemma 3.3.4, Γ is thus a valid covariance matrix. In order to prove what minimum amount of energy it achieves, we consider the following SDP, which is an instance of the SDP discussed in Lemma 3.3.3 optimizing H_{quad} in equation (3.3). It takes as input the classical optimum $\Gamma^{\text{class}} \in \mathbb{R}^{2n\times 2n}$, the edge set E in H_{class} and the parameter $p_{\text{class}} \in [0,1]$, and the constraint matrix $C = ih$ corresponds to that of H_{quad} in equation (3.3):

$$\begin{aligned}
\max_{X\in\mathbb{C}^{2n\times 2n}} \quad & \operatorname{tr}(CX)\\
\text{subject to} \quad & X \succeq 0,\\
& \forall i, X_{i,i} = 1, \forall i \neq j, X_{i,j} = -X_{j,i}, \text{(anti-symmetry)}\\
& \text{and}\\
& X_{2j-1,2j} = ip_{\text{class}}\Gamma^{\text{class}}_{2j-1,2j}, \text{for all } j \in V,\\
& X_{2j-1,2k} = -X_{2j,2k-1} \text{ for all } (j,k) \in E,\\
& X_{2j-1,2k-1} = X_{2j,2k} \text{ for all } (j,k) \in E.
\end{aligned} \tag{3.13}$$

The covariance matrix Γ is constructed to also obey the additional equality constraints in this SDP (and thus corresponds to a feasible solution of the SDP), i.e. for $X = \mathbb{1} + i\Gamma$ one can argue:

1. Since $\forall j \in V$, $\Gamma^{\text{mediator}}_{2j,2j-1} + \Gamma^{\text{quad}}_{2j,2j-1} = -\Gamma^{\text{mediator}}_{2j,2j-1} - \Gamma^{\text{quad}}_{2j,2j-1} = 0$ we have $X_{2j-1,2j} = ip_{\text{class}}\Gamma^{\text{class}}_{2j-1,2j}$.

2. Since Γ^{class} and Γ^{mediator} are zero at all entries other than $(2j-1,2j)$ and $(2j,2j-1)$ $\forall j \in V$, we have $X = i\frac{1-p_{\text{class}}}{2}\Gamma^{\text{quad}}$ at all off-diagonal entries unequal to $(2j-1,2j)$ and $(2j,2j-1)$ $\forall j \in V$. Since ρ_{quad} is a Slater determinant state, Lemma 3.3.1 implies $X_{2j-1,2k} = -X_{2j,2k-1}$ and $X_{2j-1,2k-1} = X_{2j,2k}$ for all $(j,k) \in E$.

Since Γ^{class} and Γ^{mediator} are zero at all entries other than $(2j-1,2j)$ and $(2j,2j-1)$ $j \in V$, Γ achieves approximation ratio $(1-p_{\text{class}})/2$ on H_{quad}, i.e. $\operatorname{tr}\left(\rho_{\text{Gauss}}H_{\text{quad}}\right) \geq \frac{1-p_{\text{class}}}{2}\operatorname{tr}\left(\rho_{\text{quad}}H_{\text{quad}}\right) = \frac{1-p_{\text{class}}}{2}\lambda_{\max}(H_{\text{quad}})$.

Next, we argue that the expectation on H_{class} in equation (3.4) for *any* feasible solution of the SDP in equation (3.13), and thus also for the optimum of the SDP, can be

lower-bounded as follows. Using Wick's theorem, Proposition 1.5.8, any feasible solution achieves expectation

$$\frac{1}{4}\sum_{(j,k)\in E} w_{j,k}\big(\Gamma_{2j-1,2j}+\Gamma_{2k-1,2k}-\Gamma_{2j-1,2j}\Gamma_{2k-1,2k}-\Gamma_{2j-1,2k}\Gamma_{2j,2k-1} \\ +\Gamma_{2j-1,2k-1}\Gamma_{2j,2k}\big)-\frac{1}{2}\sum_{j\in V}\mu_j\Gamma_{2j-1,2j}$$

$$\geq \frac{1}{4}\sum_{(j,k)\in E} w_{j,k}\big(p_{\text{class}}\Gamma^{\text{class}}_{2j-1,2j}+p_{\text{class}}\Gamma^{\text{class}}_{2k-1,2k}-p^2_{\text{class}}\Gamma^{\text{class}}_{2j-1,2j}\Gamma^{\text{class}}_{2k-1,2k}\big) \\ -\frac{1}{2}\sum_{j\in V}\mu_j p_{\text{class}}\Gamma^{\text{class}}_{2j-1,2j}, \qquad (3.14)$$

on H_{class}, where we have used the conditions in equation (3.13) in the inequality. Note that the final two conditions in equation (3.13) imply that $-\Gamma_{2j-1,2k}\Gamma_{2j,2k-1}\geq 0$ and $\Gamma_{2j-1,2k-1}\Gamma_{2j,2k}\geq 0$ for any feasible solution. The final step is to use equation (3.14) to prove that any feasible solution achieves at least expectation

$$p^2_{\text{class}}\lambda_{\max}\big(H_{\text{class}}\big), \qquad (3.15)$$

on H_{class}, for which we invoke Lemma 3.4.1, separately proved below with $\Gamma^{\text{class}}_{2j,2j-1}=-z_j$. The approximation ratio achieved by optimum of the SDP is thus

$$\geq \frac{p^2_{\text{class}}\lambda_{\max}(H_{\text{class}})+(1-p_{\text{class}})/2\,\lambda_{\max}(H_{\text{quad}})}{\lambda_{\max}(H)} \\ \geq \frac{p^2_{\text{class}}\beta+(1-p_{\text{class}})/2}{\beta+1}=f_\beta(p_{\text{class}}), \qquad (3.16)$$

with $\beta=\lambda_{\max}(H_{\text{class}})/\lambda_{\max}(H_{\text{quad}})\geq 0$ and $\lambda_{\max}(H)\leq\lambda_{\max}(H_{\text{class}})+\lambda_{\max}(H_{\text{quad}})$, and we have used the bounds derived above. $f_\beta(p_{\text{class}})$ is a convex function of the $p_{\text{class}}\in[0,1]$ and hence the optimal value given β is achieved at $p_{\text{class}}=0$ or $p_{\text{class}}=1$. At $\beta=1/2$, $f_\beta(p_{\text{class}}=0)=f_\beta(p_{\text{class}}=1)=1/3$ while at other values of β, $\max_{p_{\text{class}}=0,1} f_\beta(p_{\text{class}})=\max(\frac{\beta}{\beta+1},\frac{1/2}{\beta+1})\geq 1/3$, leading to the lower bound $1/3$ on the Gaussian approximation ratio. The state ρ_{Gauss} is not necessarily a pure fermionic Gaussian state. Through Remark 1.5.4, however, it is a mixture of pure fermionic Gaussian states. Therefore, at least one of the pure fermionic Gaussian states in the mixture achieves approximation ratio at least $1/3$. □

Lemma 3.4.1. *Given an interaction graph $G_{class}=((\mu,V),(w,E))$. Let $F(z_1,\ldots,z_n):=-\frac{1}{4}\sum_{j,k\in E}w_{j,k}(z_j+z_k+z_jz_k)+\frac{1}{2}\sum_{j\in V}\mu_j z_j$ with $\{z_j=\pm1\}_{j\in V}$. For any assignment $z_1,\ldots,z_n$, we have $\max\big(F(pz_1,\ldots,pz_n),F(-pz_1,\ldots,-pz_n)\big)\geq p^2\max\big(F(z_1,\ldots,z_n),F(-z_1,\ldots,-z_n)\big)$ for $p\in[0,1]$. Clearly, for the optimal assignment $y_1,\ldots,y_n$, we have that $F\big(p\,y_1,\ldots,p\,y_n\big)\geq p^2F\big(y_1,\ldots,y_n\big)$.*

Proof. We define

$$F_1(z_1,\dots,z_n) := -\frac{1}{4}\sum_{j,k\in E} w_{j,k}(z_j+z_k) + \frac{1}{2}\sum_{j\in V}\mu_j z_j,$$
$$F_2(z_1,\dots,z_n) := -\frac{1}{4}\sum_{j,k\in E} w_{j,k} z_j z_k, \tag{3.17}$$

so that $F(z_1,\dots,z_n) = F_1(z_1,\dots,z_n) + F_2(z_1,\dots,z_n)$, $F_1(-z_1,\dots,-z_n) = -F_1(z_1,\dots,z_n)$ and $F_2(-z_1,\dots,-z_n) = +F_2(z_1,\dots,z_n)$. Then, clearly, $\max\big(F_1(z_1,\dots,z_n), F_1(-z_1,\dots,-z_n)\big) \geq 0$, so $\max\big(F_1(p\,z_1,\dots,p\,z_n), F_1(-p\,z_1,\dots,-p\,z_n)\big) \geq p^2 \max\big(F_1(z_1,\dots,z_n), F_1(-z_1,\dots,-z_n)\big)$ for $p\in[0,1]$. Hence, for any assignment $z_1,\dots,z_n$

$$\max\big(F(p\,z_1,\dots,p\,z_n), F(-p\,z_1,\dots,-p\,z_n)\big) \geq \tag{3.18}$$
$$p^2\max\big(F_1(z_1,\dots,z_n), F_1(-z_1,\dots,-z_n)\big) + p^2 F_2(z_1,\dots,z_n) =$$
$$p^2\max\big(F(z_1,\dots,z_n), F(-z_1,\dots,-z_n)\big). \tag{3.19}$$

□

We can prove a small standalone corollary to Theorem 3.2.2 on the scaling of the maximum eigenvalue of the Hamiltonians in Problem 3.2.1. Namely, this scaling is fully determined by $\lambda_{\max}(H_{\text{class}})$ and $\lambda_{\max}(H_{\text{quad}})$, and there is little frustration between the two contributions H_{class} and H_{quad}.

Corollary 3.4.2. *For $H = H_{class} + H_{quad}$ as in Problem 3.2.1* TRACELESS CIFH OPTIMIZATION, *we can bound $\frac{\lambda_{\max}(H_{class})+\lambda_{\max}(H_{quad})}{3} \leq \lambda_{\max}(H) \leq \lambda_{\max}(H_{class}) + \lambda_{\max}(H_{quad})$, so that $\lambda_{\max}(H) = \Theta\big(\lambda_{\max}(H_{class}) + \lambda_{\max}(H_{quad})\big)$.*

Proof. The upper bound $\lambda_{\max}(H) \leq \lambda_{\max}(H_{\text{class}}) + \lambda_{\max}(H_{\text{quad}})$ simply follows from the triangle inequality. For the lower bound $\lambda_{\max}(H) \geq 1/3\big(\lambda_{\max}(H_{\text{class}}) + \lambda_{\max}(H_{\text{quad}})\big)$, note that the fermionic Gaussian state ρ_{Gauss} in the proof of Theorem 3.2.2 actually achieves $\operatorname{tr}\big(\rho_{\text{Gauss}} H\big) \geq 1/3\big(\lambda_{\max}(H_{\text{class}}) + \lambda_{\max}(H_{\text{quad}})\big)$ through equation (3.16), so that $\lambda_{\max}(H) \geq 1/3\big(\lambda_{\max}(H_{\text{class}}) + \lambda_{\max}(H_{\text{quad}})\big)$. □

3.4.1. Optimization procedure

In the proof of Theorem 3.2.2 the lower bound on the approximation ratio is achieved by taking either the pure state ρ_{class} or the Gaussian mixed state corresponding to $\frac{1}{2}(\Gamma_{\text{quad}} + \Gamma_{\text{mediator}})$. This does not mean that this solution will always be the optimal Gaussian state. In fact, one can run the semi-definite program in equation (3.13) —assuming access to Γ^{class} for the moment— and possibly get better solutions.

The input parameter p_{class} of the SDP can be chosen efficiently as follows. Run the SDP in equation (3.13) for $p_{\text{class}} = j/M$ for $j = 0,1,\dots,M = \text{poly}(n)$. For each j, obtain the optimum $X^{(j)}$ of the SDP in equation (3.13) and its associated covariance matrix $\Gamma^{(j)} = i\big(\mathbb{1} - X^{(j)}\big)$. Then, calculate the expectation that $\Gamma^{(j)}$ achieves on H (in Problem 3.2.1 TRACELESS CIFH OPTIMIZATION) and pick the $\Gamma^{(j)}$ that achieves the largest expectation. Since the values $p_{\text{class}} = 0,1$ are included in the sweep, the optimal $\Gamma^{(j)}$ will achieve

approximation ratio at least 1/3. In practice, we expect the optimal approximation ratio to be achieved at an intermediate value of p_{class} and to be larger than 1/3. To illustrate this, Figure 3.1 gives the approximation ratio achieved by the optimum of SDP in equation (3.13) on H as a function of p_{class}. Here H is taken to be a 3-site (i.e., 6-mode) traceless Fermi-Hubbard Hamiltonian on a triangle.

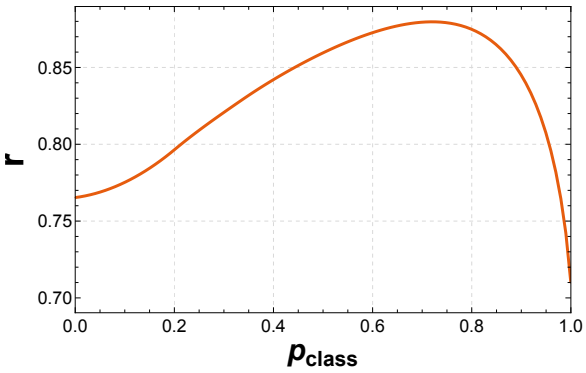

Figure 3.1: Approximation ratio r versus p_{class} for 3-site (i.e., 6-mode) traceless Fermi-Hubbard Hamiltonian. The maximum approximation ratio r_* (> 1/3) is achieved at an intermediate value of p_{class}.

Another input to the SDP in equation (3.13) is Γ^{class}. In the next section, we will discuss under what conditions (a sufficiently accurate approximation of) Γ^{class} can be obtained —i.e., under what conditions our method for constructing a fermionic Gaussian state with constant approximation ratio is constructive.

3.5. Efficient approximate constructions

In this section, we discuss under what conditions one can efficiently obtain Γ^{class} (the optimum of H_{class}) or an approximation of it – and use it to efficiently construct a Gaussian approximation using the SDP in equation (3.13). In particular, we show that if the *interaction* graph $G_{\text{class}} = (V, (w, E))$ in H_{class} in equation (3.4) is bipartite, then Γ^{class} can efficiently be found, leading to Theorem 3.2.3. In addition, for Problem 3.2.4 PSD CIFH OPTIMIZATION, we can efficiently obtain a provably accurate approximation of Γ^{class}, leading to Theorem 3.2.5.

In Section 3.5.3, we take a small detour and discuss how Gaussian approximations can be constructed using our methods for FERMIONIC MAX CUT —a fermionic version of the QUANTUM MAX CUT problem [26].

In all of our constructions, when claiming polynomial-time solvability, we rely on the fact that our SDPs are of dimensionality $O(n)$ and satisfy standard conditions for solvability of semi-definite programs in polynomial time (polynomial in dimensionality, logarithm of the error, and the number of digits of precision). In particular, one can rely on Section 5.3 of the textbook by Ben-Tal and Nemirovski [29], which demonstrates polynomial efficiency for semi-definite programs with polynomially bounded feasible sets

(see Theorem 5.3.1 in [29]). For the SDP we introduced in equation (3.13), the feasible set is polynomially bounded by bounding the Frobenius norm of the matrices of spectral radius 1 (which is required by one of the constraints). The exact same argument is sufficient for all semi-definite programs which will be introduced later in this Section.

3.5.1. Proof of Theorem 3.2.3: classical interactions on bipartite graphs

Lemma 3.5.1. *Let the Hamiltonian* $H_{class} = \sum_{(j,k)\in E} w_{j,k}\left(1/4 - x_j x_k\right) + \sum_{j\in V} \mu_j \left(x_j - 1/2\right)$ *with* $w_{j,k} \geq 0$*, be defined on a graph* $G_{class} = ((\mu, V), (w, E))$ *with binary variables* $x_j = 0, 1$*. If* G_{class} *is bipartite, then the classical state* ρ*, i.e. the vector* $\boldsymbol{x}$*, that optimizes* H_{class} *can be obtained in polynomial time.*

Proof. Our proof directly uses a known Theorem on quadratic binary optimization problems (QUBO). In particular, the maximization problem $\mathbf{x}^T Q\mathbf{x} + \mathbf{c}^T\mathbf{x}$ over the binary vector $\mathbf{x}$ with Q a real matrix with nonnegative off-diagonal entries (and $\mathbf{c}$ a real vector) can be efficiently solved: this is for example stated as Theorem 3.16 in [23]. We can introduce Ising spin variables $z_i = 1 - 2x_i = \pm 1$ and rewrite H_{class} in terms of these variables such that the quadratic term in H_{class} equals $-\frac{1}{4}\sum_{(j,k)\in E} w_{j,k} z_j z_k$. Since G_{class} is bipartite with bi-partition $V = V_A \cup V_B$, applying a spin-flip $z_i \to -z_i$ for all $i \in A$, will flip the sign so that the quadratic term becomes $+\frac{1}{4}\sum_{(j,k)\in E} w_{j,k} z_j z_k$. Switching back to the x_j variables thus gives nonnegative off-diagonal entries when constructing the matrix Q, while there are no constraints on the vector $\mathbf{c}$. Hence, we can efficiently obtain an optimal solution $\mathbf{x}$. Explicitly, the QUBO problem is rewritten as an integer linear program and relaxed to a linear program whose optimal solution can be shown to be achieved on $\{x_j \in \{0, 1\}\}$, see also [30]. □

The procedure used in Lemma 3.5.1 provides the optimum of H_{class} provided that G_{class} is a bipartite graph. For Fermi-Hubbard Hamiltonians G_{class} is not just bipartite, but it is simply a collection of disjoint edges. For such trivial graphs the optimum of H_{class} can be obtained in an extremely simple way:

Remark 3.5.2. *If* $G_{class} = ((\mu, V), (w, E))$ *consists of a collection of disjoint edges, then the optimum of* $H_{class} = \sum_{(j,k)\in E} w_{j,k}\left(\mathbb{1}/4 - x_j x_k\right) + \sum_{j\in V} \mu_j \left(x_j - \mathbb{1}/2\right)$ *on* G_{class} *can be obtained using the following simple procedure. Start from the vacuum state* $x_j = 0\ \forall\, j \in V$*. For each edge* $(j, k) \in E$*, set* $x_j = 1$ *if* $\mu_j \geq \mu_k$ *and set* $x_k = 1$ *if* $\mu_j < \mu_k$*. Then, for each edge* $(j, k) \in E$*, if* $\min\{\mu_j, \mu_k\} \geq w_{j,k}$*, also occupy the other mode* $\arg\min\{\mu_j, \mu_k\}$ *on the edge.*

Proof of Theorem 3.2.3. Theorem 3.2.3 thus follows immediately by noting that the SDP in equation (3.13), which leads to a state with approximation ratio 1/3, can be run in polynomial time, since its input Γ^{class} can be obtained in polynomial time. Therefore, the fermionic Gaussian state that achieves an approximation ratio at least 1/3 for Problem 3.2.1 can be obtained deterministically in polynomial time. □

3.5.2. Proof of Theorem 3.2.5: positive semi-definite classical interactions

In this section, we obtain an approximation of the optimum Γ^{class} of H_{class} in Problem 3.2.4. This is established in Lemma 3.5.4. In our proof, we make use of a result from [27], namely:

Lemma 3.5.3 (Theorem 3.2.1 in [27])**.** *Given a weighted graph $G=(V,(w,E))$ with non-negative weights $w_{j,k}\geq 0$ and a* $\mathrm{sign}(j,k)$ *for each edge $(j,k)\in E$. Consider the problem of computing*

$$\mathrm{MaxCut}_{\pm 1}=\max_{\{z_j=\{\pm 1\}\}_{j\in V}}\sum_{(j,k)\in E}w_{j,k}\big(1-\mathrm{sign}(j,k)z_jz_k\big). \tag{3.20}$$

An assignment $\{z_j\}_{j\in V}$ that (in expectation) achieves objective value $r_{GW}\mathrm{MaxCut}_{\pm 1}$ (with $r_{GW}=0.878$) for this problem can be obtained in polynomial time.

Lemma 3.5.4. *Let the Hamiltonian $H_{\mathrm{class}}=\sum_{(j,k)\in E}w_{j,k}\big(1-x_jx_k\big)+\sum_{j\in V}\mu_jx_j\geq 0$ be defined on a graph $G_{\mathrm{class}}=(V,(w,E))$, and $\mu_j\geq 0$. A classical state ρ, i.e. the binary vector $\boldsymbol{x}$, can be efficiently found which achieves expectation* $\mathrm{tr}\big(\rho H_{\mathrm{class}}\big)\geq r_{GW}\lambda_{\max}\big(H_{\mathrm{class}}\big)$.

Proof. We define $\tilde{\mu}_j=\mu_j/|E_j|\ \forall j\in V$, with $|E_j|$ denoting the number of edges adjacent to a vertex j and let $z_j=1-2x_j$ so that we have

$$H_{\mathrm{class}}(\{z_j\})=\sum_{(j,k)\in E}\Big[\Big(\frac{3}{4}w_{j,k}+\frac{1}{2}\tilde{\mu}_j+\frac{1}{2}\tilde{\mu}_k\Big)-\Big(\frac{1}{4}w_{j,k}-\frac{1}{2}\tilde{\mu}_j\Big)z_j-\Big(\frac{1}{4}w_{j,k}-\frac{1}{2}\tilde{\mu}_k\Big)z_k \\ -\frac{1}{4}w_{j,k}z_jz_k\Big]. \tag{3.21}$$

We can relate this optimization problem to a purely quadratic optimization problem by introducing a new variable $y=\pm 1$. Consider

$$H^{+}_{\mathrm{class}}(\{z_j\},y)=\sum_{(j,k)\in E}\Big[\Big(\frac{3}{4}w_{j,k}+\frac{1}{2}\tilde{\mu}_j+\frac{1}{2}\tilde{\mu}_k\Big)-\Big(\frac{1}{4}w_{j,k}-\frac{1}{2}\tilde{\mu}_j\Big)yz_j \\ -\Big(\frac{1}{4}w_{j,k}-\frac{1}{2}\tilde{\mu}_k\Big)yz_k-\frac{1}{4}w_{j,k}z_jz_k\Big]. \tag{3.22}$$

For an assignment $\big(y,\{z_j\}\big)$ achieving a value $r\max_{\{z_j\},y}H^{+}_{\mathrm{class}}$ (for some $r\leq 1$), the negated assignment $\big(\bar{y},\{\bar{z}_j\}\big)$ (with $\bar{z}$ denoting the spin-flip negation of z) clearly achieves the same value. Furthermore, either the assignment $\big(y,\{z_j\}\big)$ (if $y=+1$) or $\big(\bar{y},\{\bar{z}_j\}\big)$ (if $\bar{y}=+1$) achieves the value $r\max_{\{z_j\},y}H^{+}_{\mathrm{class}}\geq r\max_{\{z_j\}}H_{\mathrm{class}}$ for $H_{\mathrm{class}}(\{z_j\})$ in equation (3.21): either $\{z_j\}$ or $\{\bar{z}_j\}$ achieves the approximation ratio r for H_{class} in equation (3.21). Due to Lemma 3.5.3, we have an r_{GW}-approximation algorithm for the quadratic optimization problem in equation (3.22) which is of the form in equation (3.20) plus a constant, i.e.

$$H^{+}_{\mathrm{class}}=c+\sum_{(j,k)\in E}\Big[\Big|\frac{1}{4}w_{j,k}-\frac{1}{2}\tilde{\mu}_j\Big|\Big(1-\mathrm{sign}\Big(\frac{1}{4}w_{j,k}-\frac{1}{2}\tilde{\mu}_j\Big)yz_j\Big) \\ \Big|\frac{1}{4}w_{j,k}-\frac{1}{2}\tilde{\mu}_k\Big|\Big(1-\mathrm{sign}\Big(\frac{1}{4}w_{j,k}-\frac{1}{2}\tilde{\mu}_k\Big)yz_k\Big)+\frac{1}{4}w_{j,k}\big(1-z_jz_k\big)\Big], \tag{3.23}$$

with constant $c\geq 0$ when $\mu_j\geq 0$. Due to $c\geq cr_{GW}$ the approximation ratio for the problem in equation (3.21) is thus at least r_{GW}. □

Now we are ready to prove Theorem 3.2.5.

Proof of Theorem 3.2.5. The proof of Theorem 3.2.5 largely follows the same structure as that of Theorem 3.2.2, so we advise the reader to first read the latter. Through Lemma 3.5.4, we can efficiently obtain a classical state with covariance matrix $\Gamma^{\text{class}}_{GW}$ that achieves energy $r_{GW}\lambda_{\max}(H_{\text{class}})$ on H_{class} in Problem 3.2.4. We let $\tilde{\Gamma}^{\text{class}}$ be the state that achieves the largest expectation on H_{class} out of $\Gamma^{\text{class}}_{GW}$ and $-\Gamma^{\text{class}}_{GW}$.

Next, let us prove how $\tilde{\Gamma}^{\text{class}}$ can be used in combination with the SDP in equation (3.13) to obtain a fermionic Gaussian state that achieves approximation ratio 0.637 for Problem 3.2.4. We consider the SDP in equation (3.13), where we input $\tilde{\Gamma}^{\text{class}}$ instead of Γ^{class} and the cost function is shifted upwards by $\sum_{(j,k)\in E'} w'_{j,k}$. The blended Gaussian state with covariance matrix $\Gamma = p_{\text{class}}\Gamma^{\text{class}} + \frac{1-p_{\text{class}}}{2}(\Gamma^{\text{mediator}} + \Gamma^{\text{quad}})$ (with Γ^{mediator} defined in equation (3.12)) is a feasible solution of the SDP for the same reasons as in the proof of Theorem 3.2.2. Therefore, the optimum of the SDP achieves expectation $\sum_{(j,k)\in E'}\big(w'_{j,k} + \frac{1-p_{\text{class}}}{2}\frac{1}{2}(\Gamma^{\text{quad}}_{2j-1,2k} - \Gamma^{\text{quad}}_{2j,2k-1})\big)$ on H_{quad}. The optimum thus achieves expectation at least

$$\frac{1+\frac{1-p_{\text{class}}}{2}}{2}\cdot\lambda_{\max}(H_{\text{quad}}), \tag{3.24}$$

on H_{quad}, where we have used that $1+Cx \geq \frac{1+C}{2}(1+x)$ for $x\in[-1,+1]$ and $C\in[0,1]$, because $\frac{1+Cx}{1+x}$ decreases monotonically with x for any $C\in[0,1]$.

Next, let us argue that the approximation ratio on H_{class} can be lower bounded for any feasible solution of the SDP —and therefore also for the optimum— as follows. Using the conditions in the SDP in equation (3.13), with $\tilde{\Gamma}^{\text{class}}$ as input, we conclude that any feasible solution achieves expectation at least

$$\sum_{(j,k)\in E}\frac{1}{4}w_{j,k}\Big(3 + p_{\text{class}}\tilde{\Gamma}^{\text{class}}_{2j-1,2j} + p_{\text{class}}\tilde{\Gamma}^{\text{class}}_{2k-1,2k} - p^2_{\text{class}}\tilde{\Gamma}^{\text{class}}_{2j-1,2j}\tilde{\Gamma}^{\text{class}}_{2k-1,2k}\Big) + \frac{1}{2}\sum_{j\in V}\mu_j\big(1 - p_{\text{class}}\tilde{\Gamma}^{\text{class}}_{2j-1,2j}\big) \tag{3.25}$$

on H_{class}. Using reasoning similar as in Lemma 3.4.1, we argue the following.

$$\begin{aligned}&\sum_{(j,k)\in E}\frac{1}{4}w_{j,k}\Big(p_{\text{class}}\tilde{\Gamma}^{\text{class}}_{2j-1,2j} + p_{\text{class}}\tilde{\Gamma}^{\text{class}}_{2k-1,2k} - p^2_{\text{class}}\tilde{\Gamma}^{\text{class}}_{2j-1,2j}\tilde{\Gamma}^{\text{class}}_{2k-1,2k}\Big) - \frac{1}{2}\sum_{j\in V}\mu_j p_{\text{class}}\tilde{\Gamma}^{\text{class}}_{2j-1,2j}\\ &\geq p^2_{\text{class}}\Big[\sum_{(j,k)\in E}\frac{1}{4}\Big(\tilde{\Gamma}^{\text{class}}_{2j-1,2j} + \tilde{\Gamma}^{\text{class}}_{2k-1,2k} - \tilde{\Gamma}^{\text{class}}_{2j-1,2j}\tilde{\Gamma}^{\text{class}}_{2k-1,2k}\Big) - \frac{1}{2}\sum_{j\in V}\mu_j\tilde{\Gamma}^{\text{class}}_{2j-1,2j}\Big].\end{aligned} \tag{3.26}$$

This follows from the fact that the sum W_{lin} of the contributions *linear* in entries of $\tilde{\Gamma}^{\text{class}}$ is non-negative and so $p_{\text{class}}W_{\text{lin}} \geq p^2_{\text{class}}W_{\text{lin}}$. Else, $\tilde{\Gamma}^{\text{class}}$ would have been chosen to equal $\Gamma^{\text{class}}_{GW}$ with opposite sign. Hence any feasible solution achieves at least the following

expectation on H_{class}.

$$\sum_{(j,k)\in E} \frac{1}{4} w_{j,k} \Big(3 + p_{\text{class}}^2 \big(\tilde{\Gamma}^{\text{class}}_{2j-1,2j} + \tilde{\Gamma}^{\text{class}}_{2k-1,2k} - \tilde{\Gamma}^{\text{class}}_{2j-1,2j}\tilde{\Gamma}^{\text{class}}_{2k-1,2k}\big)\Big) + \frac{1}{2}\sum_{j\in V} \mu_j \big(1 - p_{\text{class}}^2 \tilde{\Gamma}^{\text{class}}_{2j-1,2j}\big)$$

$$\geq \frac{3+p_{\text{class}}^2}{4} \sum_{(j,k)\in E} \frac{1}{4} w_{j,k} \Big(3 + \big(\tilde{\Gamma}^{\text{class}}_{2j-1,2j} + \tilde{\Gamma}^{\text{class}}_{2k-1,2k} - \tilde{\Gamma}^{\text{class}}_{2j-1,2j}\tilde{\Gamma}^{\text{class}}_{2k-1,2k}\big)\Big) + \frac{1+p_{\text{class}}^2}{2} \frac{1}{2}\sum_{j\in V} \mu_j \big(1 - \tilde{\Gamma}^{\text{class}}_{2j-1,2j}\big)$$

$$\geq \frac{1+p_{\text{class}}^2}{2} r_{GW} \lambda_{\max}(H_{\text{class}}), \tag{3.27}$$

where we have used that $3 + Cx \geq \frac{3+C}{4}(3+x)$ for $x \in [-3,+1]$ and $C \in [0,1]$, that $1 - Cx \geq \frac{1+C}{2}(1-x)$ for $x \in [-1,+1]$ and $C \in [0,1]$, and that $\big(3 + p_{\text{class}}^2\big)/4 \geq \big(1 + p_{\text{class}}^2\big)/2$.

Therefore, the optimum of the SDP in equation (3.13) (with $\tilde{\Gamma}^{\text{class}}$ as input) achieves approximation ratio on H in Problem 3.2.4 which is at least

$$\frac{\big((1+p_{\text{class}}^2)/2\big) r_{GW}\beta + \big(1/2 + (1-p_{\text{class}})/4\big)}{\beta+1} = f_\beta(p_{\text{class}}), \tag{3.28}$$

with $\beta = \lambda_{\max}(H_{\text{class}})/\lambda_{\max}(H_{\text{quad}}) \geq 0$. The function $f_\beta(p_{\text{class}})$ is convex and for given β the optimum is achieved at $p_{\text{class}} = 0$ or $p_{\text{class}} = 1$. At $\beta = 2r_{GW}$, $f_\beta(p_{\text{class}} = 0) = f_\beta(p_{\text{class}} = 1) = r_{GW}/(r_{GW} + 1/2)$ while at other values of β, $\max_{p_{\text{class}}=0,1} f_\beta(p_{\text{class}}) \geq r_{GW}/(r_{GW} + 1/2)$. Hence, we obtain the lower bound of $r_{GW}/(r_{GW} + 1/2) \geq 0.637$ from the theorem statement. We note that to efficiently obtain a fermionic Gaussian state achieving at least this approximation ratio, one has to optimize over p_{class} as discussed in Section 3.4.1. □

3.5.3. Fermionic Max Cut

Inspired by QUANTUM MAX CUT [26], we introduce another model with positive semi-definite terms, namely:

Problem 3.5.5 (FERMIONIC MAX CUT)**.** *Consider the Hamiltonian*

$$H = \sum_{(j,k)\in E} \frac{1}{2} w_{j,k} \Big(-a_j^\dagger a_k - a_k^\dagger a_j + n_j + n_k - 2n_j n_k\Big) =$$

$$\sum_{(j,k)\in E} \frac{1}{4} w_{j,k} \big(\mathbb{1} + ic_{2j-1}c_{2k} - ic_{2j}c_{2k-1} + c_{2j-1}c_{2j}c_{2k-1}c_{2k}\big) =$$

$$\sum_{(j,k)\in E} w_{j,k} \left(\frac{\mathbb{1} + ic_{2j-1}c_{2k}}{2}\right)\left(\frac{\mathbb{1} - ic_{2j}c_{2k-1}}{2}\right) \geq 0, \tag{3.29}$$

with $w_{j,k} \geq 0$, and vertex set V and edge set E. Compute $\lambda_{\max}(H) = \max_{\rho\in\mathbb{C}^{2^n\times 2^n}} \Big\{ \operatorname{tr}\big(\rho H\big)$ s.t. $\rho \succeq 0$, $\operatorname{tr}\big(\rho\big) = 1 \Big\}$.

Note that H in equation (3.29) is manifestly positive semi-definite since each term is a projector onto a pure Gaussian 2-mode state with $\Gamma_{2j-1,2k} = +1$, $\Gamma_{2j,2k-1} = -1$. When put on a line, by the Jordan-Wigner transformation, this model is equivalent to QUANTUM MAX CUT, where the projector on each edge projects onto a 2-qubit singlet state. For more general graphs, FERMIONIC MAX CUT does not have the $U^{\otimes n}$ symmetry of QUANTUM MAX CUT (aka the anti-ferromagnetic Heisenberg model) which allows it to be more easily solvable/approximable [31]. We introduce FERMIONIC MAX CUT as a novel generalization of the classical Max Cut problem which may be more amenable to approximate optimization methods than the general Problem 3.2.4.

Indeed, observe that FERMIONIC MAX CUT is like Problem 3.2.4 with $\mu_j = \sum_{k\in E_j} \frac{1}{2} w_{j,k}$ $\forall\, j$ (with E_j denoting the subset of edges involving mode j), with the edge sets coinciding, i.e. $(w', E') = (\frac{1}{2}w, E)$, and only requiring that the *sum* of the classical interaction and hopping interactions is positive semi-definite. Note, however, that in that parameter regime, the trace of H in Problem 3.2.4 is larger than that of H in FERMIONIC MAX CUT by an amount $\sum_{j,k\in E} \frac{9}{8} w_{j,k} \operatorname{tr}(\mathbb{1})$. Hence, values for approximation ratios for Problem 3.2.4 correspond to different values of approximation ratios for FERMIONIC MAX CUT. Observe, for instance, that the maximally-mixed state achieves approximation ratio 1/4 for FERMIONIC MAX CUT. It is straightforward to prove the following proposition.

Proposition 3.5.6. *There exists a fermionic Gaussian state ρ that achieves approximation ratio $\frac{1}{2}$ for* FERMIONIC MAX CUT*, and this state can be obtained deterministically in polynomial time.*

Proof. Let us define $H_{\text{class}} = \sum_{(j,k)\in E} \frac{1}{4} w_{j,k}\big(\mathbb{1} + c_{2j-1}c_{2j}c_{2k-1}c_{2k}\big)$ and $H_{\text{quad}} = \sum_{(j,k)\in E} \frac{1}{4} w_{j,k}\big(i c_{2j-1}c_{2k} - i c_{2j}c_{2k-1}\big)$ and let Γ^{quad} denote the covariance matrix of the optimum of H_{quad}. Let Γ^{mediator} be defined as in equation (3.12). The Gaussian blend $\Gamma = \frac{1}{2}\big(\Gamma^{\text{mediator}} + \Gamma^{\text{quad}}\big)$ (see Lemma 3.3.4) with density matrix ρ_{Gauss} achieves expectation

$$
\begin{aligned}
&\operatorname{Tr}(\rho_{\text{Gauss}} H_{\text{class}}) = \\
&\sum_{j,k\in E} \frac{1}{4} w_{j,k}\big(1 - \Gamma_{2j-1,2j}\Gamma_{2k-1,2k} - \Gamma_{2j-1,2k}\Gamma_{2j,2k-1} + \Gamma_{2j-1,2k-1}\Gamma_{2j,2k}\big) = \\
&\sum_{j,k\in E} \frac{1}{4} w_{j,k}\Big(1 - \frac{1}{4}\Gamma^{\text{quad}}_{2j-1,2k}\Gamma^{\text{quad}}_{2j,2k-1} + \frac{1}{4}\Gamma^{\text{quad}}_{2j-1,2k-1}\Gamma^{\text{quad}}_{2j,2k}\Big) \geq \\
&\sum_{(j,k)\in E} \frac{1}{4} w_{j,k} \geq \frac{1}{2}\lambda_{\max}(H_{\text{class}}),
\end{aligned}
\tag{3.30}
$$

where we have used Proposition 1.5.8 and the fact that Γ^{quad} corresponds to a Slater determinant state, see Lemma 3.3.1. Since $\operatorname{Tr}(\rho_{\text{Gauss}} H_{\text{quad}}) = \frac{1}{2}\lambda_{\max}(H_{\text{quad}})$ on H_{quad}, the approximation of ρ_{Gauss} is thus at least $\operatorname{Tr}(\rho_{\text{Gauss}} H)/\lambda_{\max}(H) \geq 1/2$, using $\lambda_{\max}(H) \leq \lambda_{\max}(H_{\text{class}}) + \lambda_{\max}(H_{\text{quad}})$. Through Remark 1.5.4, there is a pure fermionic Gaussian state that achieves approximation ratio at least $\frac{1}{2}$. Since Γ^{quad} and Γ^{mediator} can be obtained in polynomial time, the Gaussian blend can be obtained efficiently. □

Fermionic Gaussian states thus achieve approximation ratio at least $\frac{1}{2}$ on FERMIONIC MAX CUT. This can be contrasted with the fact that product states achieve approximation ratio *at most* $\frac{1}{2}$ on QUANTUM MAX CUT [26].

3.6. Fermionic optimization in the presence of a particle constraint

Here, we consider approximation algorithms in case the optimization problem involves an average particle number constraint, see Problem 3.2.6. We prove that one can still obtain an approximation ratio of $1/\left[2\left((n-2q)/n+3/2\right)\right]$ (reducing to 1/3 at half-filling, i.e., for $q = n/2$) and an efficient algorithm to find such Gaussian state in case of a bipartite classical interaction graph when fixing the average particle number to q, see Theorem 3.2.7.

Let us denote the covariance matrix of the optimum of H_{class} in equation (3.4) (here with $\mu_j = 0\ \forall\, j$) at *exactly* q particles by Γ_q^{class}, and its associated expectation by $\lambda_{\max,q}(H_{\text{class}})$. We denote the covariance matrix of the *overall* optimum of H_{quad} by Γ^{quad}. To prove Theorem 3.2.7, let us first establish the following two simple lemmas. Lemma 3.6.1 says that the optimum of H_{class} at *average* particle number $q \in \{0,1,\ldots,\lfloor n/2\rfloor\}$ is in fact a classical state at particle number q and can be obtained efficiently. Lemma 3.6.2 establishes two facts about H_{class} that we will use in the proof of Theorem 3.2.7. We only prove these two lemmas for $q \leq \lfloor n/2\rfloor$, hence Theorem 3.2.7 is only proved for those q's.

Lemma 3.6.1. *Given a bipartite interaction graph $G_{class} = (V,(w,E))$. For each $q \in \{0,1,\ldots,\lfloor n/2\rfloor\}$, the optimal classical state Γ_q^{class} (see Definition 1.5.6) at particle number q achieves the average-q optimum $\lambda_{\max,\langle q\rangle}(H_{class}) = \max_{\rho\in\mathbb{C}^{2^n\times 2^n}} \left\{ \operatorname{tr}\left(\rho H_{class}\right) \text{ s.t. } \operatorname{tr}\left(\rho \hat{N}\right) = q,\ \rho \succeq 0,\ \operatorname{tr}\left(\rho\right) = 1 \right\}$. This state can be obtained in polynomial time.*

Proof. The graph $G_{\text{class}} = (V,(w,E))$ is bipartite w.r.t. a bi-partition $V = V_A \cup V_B$, where w.l.o.g. we take $|V_A| \geq |V_B|$ so that $|V_A| \geq \lfloor n/2\rfloor$. Therefore, the classical state $\prod_{j\in V_A} a_j^\dagger |\text{vac}\rangle$ (for which $n_j = 1$ for all $j \in V_A$ and $n_k = 0$ for all $k \in V_B$) is a state at particle number $|V_A| \geq q$ and achieves expectation $\frac{1}{4}\sum_{(j,k)\in E} w_{j,k}$. Clearly, one can annihilate particles in modes $j \in V_A$ until a classical state at particle number q is achieved, while preserving the expectation $\frac{1}{4}\sum_{(j,k)\in E} w_{j,k}$. Since the optimum $\lambda_{\max,\langle q\rangle}(H_{\text{class}})$ is at most $\frac{1}{4}\sum_{(j,k)\in E} w_{j,k}$, the lemma statement follows. □

Lemma 3.6.2. *Given a bipartite interaction graph $G_{class} = (V,(w,E))$. Let $F(z_1,\ldots,z_n) := -\frac{1}{4}\sum_{j,k\in E} w_{j,k}(z_j + z_k + z_j z_k)$ with $\{z_j = \pm 1\}_{j\in V}$. Let $z_1^q,\ldots,z_n^q$ denote the optimal assignment with exactly q variables set to $+1$, and $F_q = F(z_1^q,\ldots,z_n^q)$. Then, (1) $F_0 = F_1 = F_2 = \ldots = F_{\lfloor n/2\rfloor}$ and (2) $F(p\, z_1^q,\ldots,p\, z_n^q) \geq p^2 F(z_1^q,\ldots,z_n^q) = p^2 F_q$ for $p \in [0,1]$ and $0 \leq q \leq \lfloor n/2\rfloor$.*

Proof. As argued in the proof of Lemma 3.6.1, the optima at $0 \leq q \leq \lfloor n/2\rfloor$ are equal to $\frac{1}{4}\sum_{(j,k)\in E} w_{j,k}$ so that $F_1 = F_2 = \ldots = F_{\lfloor n/2\rfloor}$. The optimum at each $0 \leq q \leq \lfloor n/2\rfloor$ w.l.o.g. is such that for each edge $(j,k) \in E$, we have that $(z_j^q = -1, z_k^q = +1)$, $(z_j^q = +1, z_k^q = -1)$ or $(z_j^q = -1, z_k^q = -1)$. Let $F_1(z_1,\ldots,z_n) = -\frac{1}{4}\sum_{j,k\in E} w_{j,k}(z_j + z_k)$ and $F_2(z_1,\ldots,z_n) = -\frac{1}{4}\sum_{j,k\in E} w_{j,k} z_j z_k$. Then clearly, for each $0 \leq q \leq \lfloor n/2\rfloor$, $F_1(z_1^q,\ldots,z_n^q) \geq 0$. Therefore, $F_1(p\, z_1^q,\ldots,p\, z_n^q) \geq p^2 F_1(z_1^q,\ldots,z_n^q)$ for $p \in [0,1]$. Since $F_2(p\, z_1,\ldots,p\, z_n) = p^2 F_2(z_1,\ldots,z_n)$, we have that $F(p\, z_1^q,\ldots,p\, z_n^q) \geq p^2 F(z_1^q,\ldots,z_n^q)$ for $p \in [0,1]$ and $0 \leq q \leq \lfloor n/2\rfloor$. □

Having established these two facts, let us prove Theorem 3.2.7, which follows the same structure as the proof of Theorem 3.2.2. We advise the reader to first read the latter.

Proof of Theorem 3.2.7. Let us consider the SDP in equation (3.13), which we alter in two ways. We take as input $\Gamma^{\text{class}}_{q'}$ for some $0 \leq q' \leq q$ that will be specified later in this proof, and we add the linear constraint $\sum_{j=1}^{n}(-iX_{2j-1,2j}) = n-2q$ to the SDP. Clearly, the resulting SDP is still an instance of Lemma 3.3.3.

Let us define Γ^{mediator} as in equation (3.12), given Γ^{quad}. The Gaussian blend $-iX_{\text{blend}} = p_{\text{class}}\Gamma^{\text{class}}_{q'} + \frac{1-p_{\text{class}}}{2}(\Gamma^{\text{mediator}} + \Gamma^{\text{quad}})$ is a feasible solution to the SDP —provided that $p_{\text{class}} = \frac{n-2q}{n-2q'}$. To see this, note that

1. $-i(X_{\text{blend}})_{2j-1,2j} = p_{\text{class}}(\Gamma^{\text{class}}_{q'})_{2j-1,2j}$, $\forall\, j \in V$.

2. Γ^{quad} is a covariance matrix of a Slater determinant state. Therefore, we have $(X_{\text{blend}})_{2j-1,2k} = -(X_{\text{blend}})_{2j,2k-1}$ and $(X_{\text{blend}})_{2j-1,2k-1} = (X_{\text{blend}})_{2j,2k}$ for all $(j,k) \in E$ (see Lemma 3.3.1).

3. $\sum_{j=1}^{n}\left(-i(X_{\text{blend}})_{2j-1,2j}\right) = p_{\text{class}}\sum_{j=1}^{n}\left(\Gamma^{\text{class}}_{q'}\right)_{2j-1,2j} = p_{\text{class}}(n-2q') = n-2q$.

This feasible solution —and therefore the optimum— achieves approximation ratio at least $\frac{1-p_{\text{class}}}{2}$ of $\lambda_{\max}(H_{\text{quad}})$, and therefore also of $\lambda_{\max,\langle q\rangle}(H_{\text{quad}})$ (since $\lambda_{\max}(H_{\text{quad}}) \geq \lambda_{\max,\langle q\rangle}(H_{\text{quad}})$). What is left is to prove that all feasible solutions of the SDP —and so also its optimum— achieve expectation at least $p^2_{\text{class}}\lambda_{\max,\langle q\rangle}(H_{\text{class}})$ on H_{class}. Through Lemma 3.6.1, this is equivalent to showing that all feasible solutions achieve at least expectation $p^2_{\text{class}}\lambda_{\max,q}(H_{\text{class}})$ (i.e., the *exact* q-particle optimum) on H_{class}. Since we have $\lambda_{\max,q}(H_{\text{class}}) = \lambda_{\max,q'}(H_{\text{class}})$ for $q' \leq q \leq \lfloor n/2 \rfloor$ (see Lemma 3.6.2), it suffices to show that all feasible solutions achieve at least expectation $p^2_{\text{class}}\lambda_{\max,q'}(H_{\text{class}})$. This in turn follows from the SDP constraints $-i(X_{\text{blend}})_{2j-1,2j} = p_{\text{class}}\Gamma^{\text{class}}_{q'}$ $\forall\, j \in V$ in combination with Lemma 3.6.2.

Therefore, the optimum of the SDP achieves approximation ratio at least

$$\frac{p^2_{\text{class}}\lambda_{\max,\langle q\rangle}(H_{\text{class}}) + \frac{1-p_{\text{class}}}{2}\lambda_{\max,\langle q\rangle}(H_{\text{quad}})}{\lambda_{\max,\langle q\rangle}(H)} \geq \frac{p^2_{\text{class}}\beta + \frac{1-p_{\text{class}}}{2}}{\beta+1}$$
$$= \frac{\left(\frac{n-2q}{n-2q'}\right)^2\beta + \frac{1}{2} - \frac{n-2q}{2(n-2q')}}{\beta+1} = f_{\beta,q'}(p_{\text{class}}), \quad (3.31)$$

with $\beta := \lambda_{\max,\langle q\rangle}(H_{\text{class}})/\lambda_{\max,\langle q\rangle}(H_{\text{quad}})$ and we have used $\lambda_{\max,\langle q\rangle}(H) \leq \lambda_{\max,\langle q\rangle}(H_{\text{class}}) + \lambda_{\max,\langle q\rangle}(H_{\text{quad}})$. The function $f_{\beta,q'}(p_{\text{class}})$ is convex and so its optimum is achieved at the boundaries of its domain; at $p_{\text{class}} = \frac{n-2q}{n}$ (at $q'=0$) or at $p_{\text{class}} = 1$ (at $q' = q$). For each β, the approximation ratio can then be shown to be lower bounded by $1/\left[2\left((n-2q)/n+3/2\right)\right]$. Through Remark 1.5.4, there is a pure fermionic Gaussian state that achieves this approximation ratio. Using Lemma 3.6.1, we infer that a fermionic Gaussian state at that approximation ratio can also be efficiently constructed, because $\Gamma^{\text{class}}_{q'}$ can be efficiently obtained. □

Note that to obtain the approximation ratio $1/\left[2\left((n-2q)/n+3/2\right)\right]$ in this proof, we used a Gaussian blend feasible solution at $p_{\text{class}} = \frac{n-2q}{n}$ and $p_{\text{class}} = 1$, where the first choice for p_{class} constitutes a genuine three-component blend.

3.7. Discussion

This work deals with the optimization of classically interacting fermion Hamiltonians — a problem directly motivated by quantum chemistry and condensed matter theory. We give several guarantees on approximating ground energy of such Hamiltonians using Gaussian fermionic states. In particular, we show that traceless classically interacting Hamiltonians admit constant-ratio Gaussian approximations —ruling out the Gaussian breakdown scenario, previously found for SYK-like models. Furthermore, we provide efficient constructions for Gaussian approximations to several traceless and positive semi-definite fermionic Hamiltonians, allowing also the inclusion of a particle number constraint. Our results are derived using the new notion of a Gaussian blend, allowing to construct Gaussian states with desired properties using mixtures of covariance matrices. Another technical contribution is a semi-definite program for Gaussian optimization of classically interacting Hamiltonians, which may be an interesting object for further analysis. On the practical side, our results help to build a rigorous basis for the Hartree-Fock approach, a standard heuristic in computational quantum chemistry and materials science.

On an intuitive level, the reason behind classically interacting Hamiltonians avoiding the Gaussian breakdown of SYK-like models is the fact that the interactions here are commuting. The widespread non-commutation as the reason for Gaussian breakdown has been most directly pinpointed in [17], which gave circuit size lower bounds for SYK ground state approximations using the so-called *commutation index*. For non-classical particle number conserving Hamiltonians, the Gaussian states may not be guaranteed to yield a constant-ratio approximation. A more detailed analysis of SYK-like models with particle conservation, perhaps also employing the commutation index, is an interesting subject for future study.

To obtain the results in this work, we have used Gaussian blends that are *perfectly mediated* – i.e., Γ^{quad} and Γ^{mediator} are blended with equal weight. One might wonder whether our results can be improved by implementing *im*perfect mediation. If the weight of Γ^{quad} were larger than the weight of Γ^{mediator}, then the entries $\Gamma^{\text{quad}}_{2j-1,2j}$ would contribute to the expectation on H_{class}. One may generally have to assume that Γ^{quad} gives a contribution to the expectation on H_{class} that would scale like its minimum eigenvalue $\lambda_{\min}(H_{\text{class}})$. Interestingly, as we argue in Appendix 3.C for Problem 3.2.1 of traceless fermionic optimization, $|\lambda_{\min}(H_{\text{class}})|$ can scale as $n\lambda_{\max}(H_{\text{class}})$ in dense cases, making it difficult to obtain an improved lower bound on the approximation ratio using imperfect mediation. A regime in which imperfect mediation would be particularly useful is in the weakly interacting regime, since ideally one would obtain an approximation ratio equal to 1 in the limit of vanishing interactions. When implementing perfect mediation, however, the approximation ratio in that regime is at most $\frac{1}{2}$.

One may anticipate that these first values of approximation ratios can be improved by considering different optimization strategies, such as those that have been used for QUANTUM MAX CUT [26, 31–34]. In particular, in Appendix 3.D we provide the SDP relaxation hierarchy (Lasserre hierarchy) of the problem of optimizing fermionic Hamiltonians H over *general* fermionic states. It is an open question to what extent one can use the solution of such an SDP, which is not a physical state, to round to a Gaussian state with Γ approximately optimizing $F(\Gamma)$ in equation (3.10), for positive semi-definite Hamiltonians

with classical interactions (Problem 3.2.4) or specifically FERMIONIC MAX CUT in Section 3.5.3.

Another open question pertains to the H_{quad} contributions to the Hamiltonians in this work. Our proof techniques use that they are particle number conserving. Whether the same results can be obtained if H_{quad} is just parity preserving is currently not known to us.

References

[1] N. Schuch and F. Verstraete. 'Computational complexity of interacting electrons and fundamental limitations of density functional theory'. In: *Nature Physics* 5.10 (Aug. 2009), pp. 732–735. ISSN: 1745-2481. DOI: 10.1038/nphys1370.

[2] B. O'Gorman, S. Irani, J. Whitfield and B. Fefferman. 'Intractability of Electronic Structure in a Fixed Basis'. In: *PRX Quantum* 3 (2 May 2022), p. 020322. DOI: 10.1103/PRXQuantum.3.020322.

[3] A. Szabo and N. S. Ostlund. *Modern quantum chemistry: introduction to advanced electronic structure theory*. Courier Corporation, 1996.

[4] F. Jensen. *Introduction to computational chemistry*. John Wiley & Sons, 2017.

[5] R. J. Bartlett and M. Musiał. 'Coupled-cluster theory in quantum chemistry'. In: *Reviews of Modern Physics* 79 (1 Feb. 2007), pp. 291–352. DOI: 10.1103/RevModPhys.79.291.

[6] D. S. Sholl and J. A. Steckel. *Density functional theory: a practical introduction*. John Wiley & Sons, 2022.

[7] J. I. Cirac, D. Perez-Garcia, N. Schuch and F. Verstraete. 'Matrix product states and projected entangled pair states: Concepts, symmetries, theorems'. In: *Reviews of Modern Physics* 93 (4 Dec. 2021), p. 045003. DOI: 10.1103/RevModPhys.93.045003.

[8] B. M. Terhal and D. P. DiVincenzo. 'Classical simulation of noninteracting-fermion quantum circuits'. In: *Physical Review A* 65.3 (Mar. 2002). ISSN: 1094-1622. DOI: 10.1103/physreva.65.032325.

[9] B. Windt, A. Jahn, J. Eisert and L. Hackl. 'Local optimization on pure Gaussian state manifolds'. In: *SciPost Physics* 10.3 (Mar. 2021). DOI: 10.21468/scipostphys.10.3.066.

[10] J. Surace and L. Tagliacozzo. 'Fermionic Gaussian states: an introduction to numerical approaches'. In: *SciPost Physics Lecture Notes* (May 2022). ISSN: 2590-1990. DOI: 10.21468/scipostphyslectnotes.54.

[11] S. Bravyi, D. Gosset, R. König and K. Temme. 'Approximation algorithms for quantum many-body problems'. In: *Journal of Mathematical Physics* 60.3 (2019), p. 032203. eprint: https://doi.org/10.1063/1.5085428. URL: https://doi.org/10.1063/1.5085428.

[12] A. Haldar, O. Tavakol and T. Scaffidi. 'Variational wave functions for Sachdev-Ye-Kitaev models'. In: *Physical Review Research* 3.2 (2021). DOI: 10.1103/physrevresearch.3.023020.

[13] M. B. Hastings and R. O'Donnell. 'Optimizing Strongly Interacting Fermionic Hamiltonians'. In: *Proc. of STOC 2022*. New York, NY, USA: ACM, 2022, pp. 776–789. DOI: https://doi.org/10.1145/3519935.3519960.

[14] Y. Herasymenko, M. Stroeks, J. Helsen and B. Terhal. 'Optimizing sparse fermionic Hamiltonians'. In: *Quantum* 7 (Aug. 2023), p. 1081. ISSN: 2521-327X. DOI: 10.22331/q-2023-08-10-1081.

[15] D. Hothem, O. Parekh and K. Thompson. *An improved approximation for sparse fermionic Hamiltonians*. 2023. arXiv: 2301.04627 [quant-ph].

[16] J. Basso, C.-F. Chen and A. M. Dalzell. *Optimizing random local Hamiltonians by dissipation*. 2024. arXiv: 2411.02578 [quant-ph].

[17] E. R. Anschuetz, C.-F. Chen, B. T. Kiani and R. King. *Strongly interacting fermions are non-trivial yet non-glassy*. 2024. arXiv: 2408.15699 [quant-ph]. URL: https://arxiv.org/abs/2408.15699.

[18] M. Ding, R. King, B. T. Kiani and E. R. Anschuetz. *Optimizing Sparse SYK*. 2025. arXiv: 2506.09037 [quant-ph]. URL: https://arxiv.org/abs/2506.09037.

[19] J. Light, I. Hamilton and J. Lill. 'Generalized discrete variable approximation in quantum mechanics'. In: *The Journal of Chemical Physics* 82.3 (Feb. 1985), pp. 1400–1409. ISSN: 0021-9606. DOI: 10.1063/1.448462.

[20] E. Briggs, D. Sullivan and J. Bernholc. 'Real-space multigrid-based approach to large-scale electronic structure calculations'. In: *Physical Review B* 54 (20 Nov. 1996), pp. 14362–14375. DOI: 10.1103/PhysRevB.54.14362.

[21] R. Babbush, N. Wiebe, J. McClean, J. McClain, H. Neven and G. K.-L. Chan. 'Low-Depth Quantum Simulation of Materials'. In: *Physical Review X* 8 (1 Mar. 2018), p. 011044. DOI: 10.1103/PhysRevX.8.011044.

[22] J. Argüello-Luengo, A. González-Tudela, T. Shi, P. Zoller and J. I. Cirac. 'Analogue quantum chemistry simulation'. In: *Nature* 574.7777 (2019), pp. 215–218. DOI: 10.1038/s41586-019-1614-4.

[23] E. Çela and A. P. Punnen. 'The Quadratic Unconstrained Binary Optimization Problem'. In: ed. by A. Punnen. Springer, 2023. Chap. 3 Complexity and Polynomially Solvable Special Cases of QUBO. DOI: https://doi.org/10.1007/978-3-031-04520-2.

[24] S. Arora, E. Berger, H. Elad, G. Kindler and M. Safra. 'On non-approximability for quadratic programs'. In: *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*. 2005, pp. 206–215. DOI: 10.1109/SFCS.2005.57.

[25] S. Khot and R. O'Donnell. 'SDP gaps and UGC-hardness for Max-Cut-Gain'. In: *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*. 2006, pp. 217–226. DOI: 10.1109/FOCS.2006.67.

[26] S. Gharibian and O. Parekh. 'Almost optimal classical approximation algorithms for a quantum generalization of Max-Cut'. In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019)*. Vol. 145. Leibniz International Proceedings in Informatics (LIPIcs). 2019, 31:1–31:17. DOI: 10.4230/LIPICS.APPROX-RANDOM.2019.31.

[27] M. X. Goemans and D. P. Williamson. 'Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming'. In: *J. ACM* 42.6 (Nov. 1995), pp. 1115–1145. ISSN: 0004-5411. DOI: 10.1145/227683.227684.

[28] F. de Melo, P. Ćwikliński and B. M. Terhal. 'The power of noisy fermionic quantum computation'. In: *New Journal of Physics* 15.1 (Jan. 2013), p. 013015. DOI: 10.1088/1367-2630/15/1/013015.

[29] A. Ben-Tal and A. Nemirovski. *Lectures on modern convex optimization: analysis, algorithms, and engineering applications*. SIAM, 2001.

[30] A. Schrijver. *Theory of linear and integer programming*. USA: John Wiley & Sons, 1986. ISBN: 0-471-90854-1.

[31] F. Huber, K. Thompson, O. Parekh and S. Gharibian. *Second order cone relaxations for Quantum Max Cut*. 2024. arXiv: 2411.04120 [quant-ph]. URL: https://arxiv.org/abs/2411.04120.

[32] O. Parekh and K. Thompson. 'Application of the Level-2 Quantum Lasserre Hierarchy in Quantum Approximation Algorithms'. In: *Proc. of the International Colloquium on Automata, Languages, and Programming (ICALP)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. DOI: 10.4230/LIPICS.ICALP.2021.102.

[33] A. Anshu, D. Gosset and K. Morenz. 'Beyond Product State Approximations for a Quantum Analogue of Max Cut'. In: *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*. Vol. 158. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020, 7:1–7:15. ISBN: 978-3-95977-146-7. DOI: 10.4230/LIPIcs.TQC.2020.7.

[34] R. King. 'An improved approximation algorithm for Quantum Max-Cut on triangle-free graphs'. In: *Quantum* 7 (Nov. 2023), p. 1180. ISSN: 2521-327X. DOI: 10.22331/q-2023-11-09-1180.

[35] E. Lieb and D. Mattis. 'Ordering Energy Levels of Interacting Spin Systems'. In: *Journal of Mathematical Physics* 3.4 (July 1962), pp. 749–751. ISSN: 0022-2488. DOI: 10.1063/1.1724276.

[36] M. Hastings. *Improving Perturbation Theory with the Sum-of-Squares: Third Order*. 2024. arXiv: 2412.03564 [quant-ph]. URL: https://arxiv.org/abs/2412.03564.

3.A. QMA-hardness of optimizing classically interacting fermions

Adapting Theorem 3 in [2], one can show

Corollary 3.A.1. *Consider a system of $2n$ fermionic modes, with fermionic operators $\{a^\dagger_{i,\sigma}, a_{i,\sigma}\}_{i\in[n],\sigma\in\{\pm1\}}$. There exist constants $p > q > 0$ such that for all $u \geq n^{14+3p+2q}$, de-*

termining to precision n^{-q} the ground state energy[3] of a Hamiltonian

$$H = u \sum_{i\in[n]} n_{i,+1} n_{i,-1} + \sum_{i<j,\sigma\in\{\pm1\}} t_{i,j,\sigma} \left(a^\dagger_{i,\sigma} a_{j,\sigma} + a^\dagger_{j,\sigma} a_{i,\sigma} \right) - \mu \sum_{i\in[n],\sigma\in\{\pm1\}} n_{i,\sigma} \tag{3.32}$$

subject to $|t_{i,j,\sigma}| \leq \sqrt{n^p u}$, μ and $u/10 \geq \mu \geq 10 \cdot n^2 \cdot \sqrt{n^p u}$ is QMA-complete.

Proof. For convenience, let us split the Hamiltonian into terms $H = H_u + H_t + H_\mu$, defined in a straightforward way based on the form of equation (3.32).

The statement only differs from Theorem 3 in [2] by two points. First, we included an additional chemical potential term H_μ into the Hamiltonian. And second, we are interested in the ground energy of the Hamiltonian as a whole, while [2] considered the Hamiltonian projected onto the n-particle (Hamming weight n) subspace. Our goal will be to show that including the chemical potential term, the ground state of the Hamiltonian in equation (3.32) is guaranteed to have n particles. This would be sufficient for the statement to follow from [2] directly, because within the n-particle subspace where H_μ is constant, the search of the ground energy of $H = H_u + H_t$ is equivalent to that of $H = H_u + H_t + H_\mu$.

First, let us show that the ground state cannot have more than n particles. Consider a general state ρ with n particles, such that for all i, either $\mathrm{tr}(\rho\, n_{i,1}) = 1$ or $\mathrm{tr}(\rho\, n_{i,-1}) = 1$ (and thus $\mathrm{tr}(\rho H_u) = 0$). Any such state has lower energy than any state ρ' with $n' > n$ particles, because (using a triangle inequality on H_t)

$$\mathrm{tr}\left(\rho' H\right) \geq u(n'-n) - |\lambda_{\min}(H_t)| - \mu n' \geq u(n'-n) - 2n^2\sqrt{n^p u} - \mu n', \tag{3.33}$$

$$\mathrm{tr}\left(\rho H\right) \leq \lambda_{\max}(H_t) - \mu n \leq 2n^2\sqrt{n^p u} - \mu n, \tag{3.34}$$

and $u(n'-n) > \mu(n'-n) + 4n^2\sqrt{n^p u}$ for large enough n, given that $(n'-n) \geq 1$ and the assumptions on u and μ.

Secondly, the ground state cannot have less than n particles. Indeed, for any state ρ' with $n' < n$ particles, its energy is:

$$\mathrm{tr}\left(\rho' H\right) \geq -|\lambda_{\min}(H_t)| - \mu n' \geq -2n^2\sqrt{n^p u} - \mu n'. \tag{3.35}$$

Comparing to equation (3.34), we see that ρ' has energy greater than any state ρ, because $\mu(n-n') \geq 4n^2\sqrt{n^p u}$ due to $(n-n') \geq 1$ and the assumptions on μ. This concludes the proof. □

3.B. Upper bound on the Gaussian approximation ratio

We give a small $n = 4$ example of a fermionic Hamiltonian —mapping to the anti-ferromagnetic Heisenberg model on a line of 4 qubits— which has a unique non-Gaussian maximum eigenstate, allowing to bound the Gaussian approximation ratio in this instance. Note that for $n < 4$, all states are Gaussian [28]. To our knowledge, this result is the first rigorous upper bound for a Gaussian approximation ratio for any classically interacting fermionic Hamiltonian.

[3] In this section, in keeping with [2] and without loss of generality for our purposes, we refer to the *smallest* eigenvalue as the ground energy.

Proposition 3.B.1. *For any fermionic Gaussian state* ρ_{Gauss}*, one has* $\frac{\text{tr}(\rho_{\text{Gauss}}H)}{\lambda_{\max}(H)} \leq 0.99904$*, with* H *an instance of Problem 3.2.1* TRACELESS CIFH OPTIMIZATION *for* $n = 4$*.*

Proof. We consider the following 4-mode Hamiltonian on a line:

$$H = -\frac{1}{2}\sum_{i=1}^{3}\left(a_i^\dagger a_{i+1} + a_{i+1}^\dagger a_i\right) + \sum_{i=1}^{3}\left(\mathbb{1}/4 - n_i n_{i+1}\right) + \frac{1}{2}\left(n_1 - \mathbb{1}/2\right) + \sum_{i=2}^{3}\left(n_i - \mathbb{1}/2\right) + \frac{1}{2}\left(n_4 - \mathbb{1}/2\right). \quad (3.36)$$

This Hamiltonian maps onto $H = -\frac{1}{4}\sum_{i=1}^{3}\left(X_iX_{i+1} + Y_iY_{i+1} + Z_iZ_{i+1}\right)$ under the Jordan-Wigner transformation, hence its maximum eigenstate corresponds to the ground state of the anti-ferromagnetic Heisenberg model on a line. For the remainder of the argument, we need the known values of $\lambda_{\max}(H)$, the spectral gap $\Delta := \lambda_{\max}(H) - \lambda_{\max-1}(H)$, and the maximum energy eigenstate $|\psi_{\max}\rangle$. Due to the Lieb-Mattis Theorem [35] H has a *unique* eigenstate at $\lambda_{\max}(H)$. One can find that $\lambda_{\max}(H) = \frac{1}{4}(3+2\sqrt{3})$ and $\Delta = \frac{1}{2}(1+\sqrt{3}-\sqrt{2})$. For the unique maximum energy eigenstate $|\psi_{\max}\rangle$, one can show that its fermionic covariance matrix obeys $\Gamma_{\psi_{\max}}\Gamma_{\psi_{\max}}^T = s\mathbb{1}$ with $s = \frac{1}{9}(5+2\sqrt{3}) < 1$. This implies that $|\psi_{\max}\rangle$ is non-Gaussian, since it is a pure state. We can assume w.l.o.g. that the Gaussian state achieving the maximum Gaussian approximation ratio is pure as discussed in Section 3.3.1. Let us express any 4-mode state as

$$|\psi\rangle = \alpha\,|\psi_{\max}\rangle + \sqrt{1-|\alpha|^2}\,|\psi_\perp\rangle, \quad (3.37)$$

with $\alpha \in \mathbb{C}$ and where $|\psi_\perp\rangle$ is any state s.t. $\langle\psi_{\max}|\psi_\perp\rangle = 0$. Then,

$$\langle\psi|H|\psi\rangle = |\alpha|^2\lambda_{\max}(H) + (1-|\alpha|^2)\langle\psi_\perp|H|\psi_\perp\rangle \leq |\alpha|^2\lambda_{\max}(H) + (1-|\alpha|^2)(\lambda_{\max}(H) - \Delta). \quad (3.38)$$

Naturally, there is a value for $|\alpha|$ above which all states $|\psi\rangle$ in equation (3.37) are non-Gaussian, hence $|\alpha|$ should be below this value to ensure Gaussianity, thus upperbounding $\langle\psi|H|\psi\rangle$ achieved by any pure fermionic Gaussian state $|\psi\rangle$. To find such $|\alpha|$, we evaluate the entries of the covariance matrix of $|\psi\rangle$ in equation (3.37). Let $\{\tilde{c}_j\}_{j=1}^8$ be the Majorana basis in which $\Gamma_{\psi_{\max}}$ is (2×2) block-diagonal, so that in that basis $\Gamma_{\psi_{\max}} = \bigoplus_{i=1}^4 \begin{pmatrix} 0 & \lambda_i \\ -\lambda_i & 0 \end{pmatrix}$ with $\lambda_i \in [-1,+1]$. For the covariance matrix Γ_ψ of $|\psi\rangle$ in equation (3.37), we then have

for $j \neq k$

$$\begin{aligned}
\left(\Gamma_\psi\right)_{j,k} &= \langle\psi| i\tilde{c}_j\tilde{c}_k |\psi\rangle \\
&= |\alpha|^2 \langle\psi_{\max}| i\tilde{c}_j\tilde{c}_k |\psi_{\max}\rangle + 2\sqrt{1-|\alpha|^2}\,\mathrm{Re}\big(\alpha \langle\psi_{\max}| i\tilde{c}_j\tilde{c}_k |\psi_\perp\rangle\big) \\
&\qquad + (1-|\alpha|^2)\langle\psi_\perp| i\tilde{c}_j\tilde{c}_k |\psi_\perp\rangle \\
&= \begin{cases} \pm|\alpha|^2\lambda_i + (1-|\alpha|^2)\langle\psi_\perp| i\tilde{c}_j\tilde{c}_k |\psi_\perp\rangle, & \text{for } (j,k) = (2i-1,2i) \text{ or } (2i,2-i), \\ 2\sqrt{1-|\alpha|^2}\,\mathrm{Re}\big(\alpha \langle\psi_{\max}| i\tilde{c}_j\tilde{c}_k |\psi_\perp\rangle\big) & \\ \qquad + (1-|\alpha|^2)\langle\psi_\perp| i\tilde{c}_j\tilde{c}_k |\psi_\perp\rangle, & \text{elsewhere,} \end{cases}
\end{aligned} \tag{3.39}$$

where we have used that $|\psi_{\max}\rangle$ is an eigenstate of $i\tilde{c}_{2i-1}\tilde{c}_{2i}$ for $i = 1,2,3,4$. For any pure fermionic Gaussian state with covariance matrix Γ, we have that $\Gamma\Gamma^T = \mathbb{1}$ (see Section 3.3), so that $\sum_k \Gamma_{j,k}^2 = 1\ \forall j$. Using that $\langle\psi_\perp| i\tilde{c}_j\tilde{c}_k |\psi_\perp\rangle \in [-1,+1]$ and $\mathrm{Re}\big(\alpha \langle\psi_{\max}| i\tilde{c}_j\tilde{c}_k |\psi_\perp\rangle\big) \in [-|\alpha|,+|\alpha|]$, we find

$$\sum_k \Gamma_{j,k}^2 \leq \big(|\alpha|^2\sqrt{s} + (1-|\alpha|^2)\big)^2 + 6\big(2|\alpha|\sqrt{1-|\alpha|^2} + (1-|\alpha|^2)\big)^2 = g_s(|\alpha|), \tag{3.40}$$

where we have used that $|\lambda_i| = \sqrt{s}$ for any $i = 1,2,3,4$. We have $g_{s=\frac{1}{9}(5+2\sqrt{3})}(|\alpha_*|) = 1$ at $|\alpha_*| \approx 0.998818$, so that $|\psi\rangle$ in equation (3.37) is non-Gaussian for $|\alpha| > |\alpha_*|$ where $g_{s=\frac{1}{9}(5+2\sqrt{3})}(|\alpha|) < 1$. Therefore, the approximation ratio achievable by Gaussian states is upper bounded by

$$\max_{0\leq|\alpha|\leq|\alpha_*|} \big[|\alpha|^2\lambda_{\max}(H) + (1-|\alpha|^2)(\lambda_{\max}(H) - \Delta)\big]/\lambda_{\max}(H) < 0.99904, \tag{3.41}$$

where we have used $\lambda_{\max}(H) = \frac{1}{4}(3+2\sqrt{3})$ and $\Delta = \frac{1}{2}(1+\sqrt{3}-\sqrt{2})$. □

We note that this type of argument does not work to bound the Gaussian approximation ratio for a system of growing size n, as the gap scaling is small relative to the scaling of the maximum energy. However, the antiferromagnetic Heisenberg model in 1D, i.e. QUANTUM MAX CUT on a line, is a well studied model, solved via the Bethe ansatz, and we expect that when translated to fermions, its ground state is never Gaussian. Note that product states do not generally translate to fermionic Gaussian states via the Jordan-Wigner transformation, nor vice versa.

We note that Proposition 3.B.1 also gives a Gaussian upper bound for traceless FERMIONIC MAX CUT, i.e., Problem 3.5.5 made traceless.

Numerically, we find that there exists a fermionic Gaussian state that achieves approximation ratio *at least* 0.9788 for the (traceless) H in equation (3.36). How does this compare to a product state approximation ratio for this 1D QUANTUM MAX CUT model? Making H positive semi-definite like in FERMIONIC MAX CUT, this numerically obtained Gaussian state achieves the ratio ≈ 0.9855. Since product states achieve at most approximation ratio $\frac{1}{2}$ on a single QUANTUM MAX CUT edge $H_{i,i+1}$ and $\lambda_{\max}(H_{i,i+1}) = 1$, they achieve ratio at

most $\frac{3/2}{\lambda_{\max}(H)} \approx 0.6340$ for H the unit weight (positive semi-definite) QUANTUM MAX CUT Hamiltonian on a $n = 4$ line. Thus, the relatively high approximation ratio achieved by a fermionic Gaussian state for QUANTUM MAX CUT on a $n = 4$ line suggests that fermionic Gaussian states might be used in approximations for QUANTUM MAX CUT more generally (beyond 1D systems).

3.C. Example system with $|\lambda_{\min}(H_{\text{class}})|/\lambda_{\max}(H_{\text{class}}) = n$ for a traceless H_{class}

Consider the traceless classical Hamiltonian H_{class} in equation (3.4), where $G_{\text{class}} = ((\mu, V), (w, E))$ is the complete graph with $w_{j,k} = w\ \forall (j,k) \in E$ and $\mu_j = \mu\ \forall\, j \in V$ and denote $n_j = x_j \in \{0,1\}$, a binary vector $\mathbf{x}$ of length n to be optimized. Let $N = \sum_{j\in V} x_j$. For such instance, we have

$$H_{\text{class}} = \frac{1}{2} w \sum_{j\neq k} \big(1/4 - x_j x_k\big) + \mu \sum_{j\in V} \big(x_j - 1/2\big)$$
$$= \frac{1}{8} wn(n-1) - \frac{1}{2} w(N^2 - N) + \mu(N - n/2). \quad (3.42)$$

This cost function is a concave function of N, whose maximum is achieved at $N = \frac{1}{2} + \mu/w$ and whose minimum is achieved at either $N = 0$ or $N = n$. The associated values of the cost function are

$$\lambda_{\max}(H_{\text{class}}) = \frac{1}{8} w\big(1 + n(n-1)\big) + \frac{\mu^2}{2w} - \frac{\mu}{2}(n-1),$$
$$\lambda_{\min}(H_{\text{class}}) = \min\Big\{\frac{1}{8} wn(n-1) - \frac{\mu n}{2}, -\frac{3}{8} wn(n-1) + \frac{\mu n}{2}\Big\}. \quad (3.43)$$

Setting $\mu = wn/2$, this reduces to $\lambda_{\max}(H_{\text{class}}) = \frac{1}{8} w(n+1)$ and $\lambda_{\min}(H_{\text{class}}) = -\frac{1}{8} wn(n+1)$, so that $|\lambda_{\min}(H_{\text{class}})|/\lambda_{\max}(H_{\text{class}}) = n$.

In contrast, for any *sparse* traceless interactions H_{class} (with constant coefficients $w_{j,k}, \mu_j$), we have $\lambda_{\max}(H_{\text{class}}) = \Theta(n)$ —and therefore $|\lambda_{\min}(H_{\text{class}})| = \Theta(n)$ since $-H_{\text{class}}$ is also sparse— so that the ratio $|\lambda_{\min}(H_{\text{class}})|/\lambda_{\max}(H_{\text{class}})$ is a constant.

3.D. SDP relaxation and rounding?

Another potential direction of further research is the following. Similar to the approaches used to optimize QUANTUM MAX CUT, one can define an SDP hierarchy, see e.g. [13, 36], optimizing H over correlation matrices —expressing correlations between weight-k Majorana monomials— of poly(n) size for $k = O(1)$. These SDP's are relaxations of the eigenvalue problem s.t. $\text{SDP}_{k=1} \geq \text{SDP}_{k=2} \geq \ldots \geq \text{SDP}_{k=n} = \lambda_{\max}(H)$, with SDP_k denoting the optimum of the SDP at level k. The feasible solutions of such SDP's do not correspond to valid density matrices in general, let alone to fermionic Gaussian states. If, however, we could round the optimum $M^{(k)}$ of SDP_k to a fermionic Gaussian state ρ_{Gauss} s.t. $\text{tr}\big(\rho_{\text{Gauss}} H\big) = r\,\text{SDP}_k \geq r\,\lambda_{\max}(H)$ (with $0 \leq r \leq 1$), then we have an r approximation algorithm for Hamiltonian H. For traceless Hamiltonians, we know such a rounding

approach does not exist for constant r in general [24]. Therefore, any such rounding scheme should leverage specific properties of the Hamiltonian at hand, such as it being positive semi-definite. For completeness, let us briefly discuss the semi-definite program that computes SDP_k.

Let

$$C_I := i^{\binom{k}{2}} c_{i_1} \dots c_{i_k}, I = \{i_1 < i_2 < \dots < i_k\}, \tag{3.44}$$

be a weight-k Majorana monomial labeled by the ordered subset I, with $C_I^\dagger = C_I$, $C_I^2 = \mathbb{1}$ and $C_I C_J = (-1)^{|I||J|-|I\cap I'|} C_J C_I$. The spectral norm of operators C_I or $C_I C_J$ is at most 1.

The collection of monomials C_I of weight at least 1 and at most k is denoted by $\mathcal{C}_k \subseteq \mathcal{C}_{2n}$. Clearly, the number of monomials in $\mathcal{C}_k$ is $|\mathcal{C}_k| = \sum_{m=1}^k \binom{2n}{m} = \Theta(n^k)$. For any Hermitian matrix σ with $\mathrm{tr}\,\sigma = 1$ (not necessarily a density matrix), we define the weight-k correlation matrix

$$M^{(k)}_{I,J}(\sigma) := \mathrm{tr}\left(C_I^\dagger C_J \sigma\right) \in \mathbb{C},\ \forall I\ M^{(k)}_{I,I} = 1,\ |M^{(k)}_{I,J}| \le 1, \tag{3.45}$$

with $C_I, C_J \in \mathcal{C}_k$. $M^{(k)} \succeq 0$ by construction, since $\boldsymbol{\alpha}^\dagger M^{(k)} \boldsymbol{\alpha} = \sum_{I,J} \alpha_I^* \alpha_J M_{I,J} = \mathrm{tr}\left(E^\dagger E \sigma\right) \ge 0$, with $E = \sum_I \alpha_I C_I$, for any complex vector $\boldsymbol{\alpha}$. For $k = n$, $M^{(k)}$ can be associated with a valid density matrix σ.

The constraints on M are those of the feasible set of an SDP, i.e. $M^{(k)} \succeq 0$, obeying some linear equality constraints related to anti-commutation, or product rules of the C_I operators, as well the linear inequality constraints given in equation (3.45). This prescribed range of the entries is bounded appropriately by the constraints $M^{(k)}_{I,I} = 1$ and $M^{(k)} \succeq 0$ (since all principal minors of a positive semi-definite matrix are non-negative). Let us refer to this feasible set of positive semi-definite matrices $M^{(k)}$ as $\mathcal{L}_k$.

Given a quartic fermionic Hamiltonian $H = \sum_{I,J:\,|I|,|J|\le 2} h_{J,I} C_I^\dagger C_J$ (which has non-zero trace in general since $C_I^2 = \mathbb{1}$), with $h_{J,I} \in \mathbb{C}$ and h s.t. H contains only quadratic and quartic terms in the $\{c_i\}$, one can write, for any fermionic density matrix $\rho \succeq 0$

$$\mathrm{tr}\,\rho H = \sum_{I,J} h_{J,I} M^{(2)}_{I,J}(\rho) = \mathrm{tr}\left(h M^{(2)}(\rho)\right). \tag{3.46}$$

Hence, optimization over fermionic density matrices ρ can be relaxed to optimization over weight-k correlation matrices, i.e. one defines the hierarchy $\mathrm{SDP}_k = \sup_{M^{(k)}} \mathrm{tr}\left(h M^{(2)}\right)$ s.t. $M^{(k)} \in \mathcal{L}_k$ whose solution may correspond to $\sigma \not\succeq 0$ (sometimes called a pseudo-density matrix).

An approach used in QUANTUM MAX CUT [26] to obtain an (almost optimal) product state solution is to employ a randomized rounding scheme similar to GW rounding [27]. An analogous qubit SDP hierarchy is defined and the optimum weight-1 Pauli correlation matrix is Cholesky decomposed into $3n$-dimensional vectors. Then, these vectors are rounded to n 3-dimensional (normalized) vectors —which directly relate to the expectation with respect to a product state solution that achieves an approximation ratio close to the optimal one over product states. One might wonder whether a similar approach could be applied to FERMIONIC MAX CUT to obtain good Gaussian approximations. The fact that one rounds to Gaussian states instead of e.g. product states makes it essentially different. Namely, given a Cholesky decomposition $\{\mathbf{b}_I\}$ of an optimum $M^{(k)}$ (s.t. $M^{(k)}_{I,J} = \mathbf{b}_I^\dagger \mathbf{b}_J$) of the

fermionic SDP hierarchy, one needs to randomly project these vectors $\mathbf{b}_I$ to a valid Gaussian covariance matrix. Naive attempts at such rounding procedures do not necessarily yield valid Gaussian covariance matrices.

4

Optimizing sparse fermionic Hamiltonians and O(1)-weight SYK Hamiltonians

4.1. Introduction

Approximating the ground state energy of a local Hamiltonian is a central problem in both physics and computer science. In computer science it plays a key role in complexity theory [1], while in physics ground states capture the behaviour of systems at low energy. Two common families of Hamiltonians of interest are those defined on collections of qubits and those acting on fermionic degrees of freedom. Fermionic Hamiltonians model various physical systems, such as electrons in condensed matter and quantum chemistry — prime targets for quantum simulation. Despite its practical and conceptual relevance, the general problem of approximating fermionic ground state energies is currently less well understood than its qubit counterpart.

Some rigorous progress in studying this problem – both for qubits and for fermions – was made from the perspective of optimization. In this subfield of computer science, one of the central tasks is efficiently finding problem solutions that are provably close to optimal [2]. The closeness is usually quantified by an approximation ratio, i.e. the ratio between the value attained by an algorithm and the optimal value for a given problem. For the classical equivalent of the ground state energy finding – Constraint Satisfaction Problems (CSPs) – such approximation ratios have been extensively studied [3].

For quantum Hamiltonians, an interesting question is how well the ground state energy can be approximated using “classical" or “mean-field” states. For qubit Hamiltonians the natural choice of classical states are product states, while for fermionic Hamiltonians they are Gaussian states. Gaussian states play a prominent role in fermionic optimization problems using the mean-field Hartree-Fock method, see e.g. [4], or dynamical mean-field theory via solving impurity problems [5], or the simulation of free fermionic computation [6, 7].

Formal guarantees on approximation ratios characterize numerical simulation methods using classical states and outline their limitations compared to quantum

computing. For qubit Hamiltonians, it was first proved by Lieb [8] (see [9] for a simplified proof) that there always exists a product state which approximates the ground state energy of a traceless 2-local qubit Hamiltonian by a factor of 1/9. Many more results on approximating ground state energies of many-body systems by product states can be found in [10–15]. In [9] it was shown, through the Goemans-Williamson method, that for a 2-local traceless qubit Hamiltonian a product state can always be *efficiently* found with approximation ratio $O(1/\log(n))$ where n is the number of qubits. Ref. [9] also considered fermionic Hamiltonians with quadratic ($q = 2$) and quartic ($q = 4$) fermionic terms. They left as an open question whether all 4-local fermionic Hamiltonians have a constant approximation ratio with respect to Gaussian states (a Gaussian approximation ratio).

A surprising counterexample to this conjecture was recently presented in Refs. [16, 17] — the family of SYK-4 models (Sachdev-Ye-Kitaev models with quartic fermionic interactions, see equation (4.2)). It was shown that with high probability, SYK-4 Hamiltonians admits a Gaussian approximation ratio no better than $O(1/\sqrt{n})$ where n is the number of fermionic modes. Contrasting this result to Refs. [8, 11], it means that qubit and fermionic ground states strongly differ in their approximability by classical states. Moreover, this opens up the question of which fermionic Hamiltonians *do* allow finite Gaussian approximation ratios.

This is the question that we aim to answer here. We do this by considering *sparse* Hamiltonians, i.e. Hamiltonians where each fermionic mode participates in a bounded number of interactions. Sparsity holds for many physically relevant Hamiltonians, such as the Fermi-Hubbard model. It also holds for exotic Hamiltonians, such as those determined by constant-degree expander hypergraphs; notably, it does not hold for the SYK model. Sparsity of interactions has been considered in the classical CSP literature. It was shown in [18] that the MaxQP problem has an efficient constant approximation ratio algorithm on graphs of bounded chromatic number, in particular graphs with bounded degree. We show that a similar assumption of sparsity is enough to guarantee constant Gaussian approximation ratios for 4-local and strictly q-local Hamiltonians. Moreover, we show that a constant Gaussian approximation ratio can be achieved for the *sparse* SYK-4 model [19] (which has a logarithmically growing interaction participation and is thus not sparse by our definition). Finally, we consider in more detail the optimal approximation ratio for the *dense* SYK-q model for $q > 4$ (thus extending the work of [17]). We show that the shortfall of Gaussian states is even more pronounced in this setting.

To avoid confusion, we note that instead of the *ground state* energy, existing works often consider approximating the *maximal* eigenvalue of the Hamiltonian $\lambda_{\max}(H)$. These two optimization problems are equivalent if the family of Hamiltonians considered is invariant under a change of sign (e.g. traceless q-local Hamiltonians). For mathematical convenience and consistency with the literature, in the rest of the text, we will also be formulating our results in terms of approximating $\lambda_{\max}(H)$.

4.2. Statement of results

4.2.1. Preliminaries

Before surveying our results, we introduce the basic setup of fermionic Hamiltonians and q-locality. This subsection also defines the SYK-q model and spells out the previous

result of a vanishing Gaussian approximation ratio for SYK-4.

We consider a system of $2n$ traceless Majorana fermion operators (see also Section 1.5) c_i, $i = 1,\dots,2n$ with $c_i^2 = \mathbb{1}, c_i^\dagger = c_i$, forming a Clifford algebra, i.e., $\{c_j, c_k\} = 2\delta_{j,k}\mathbb{1}$ and representing n fermionic modes. We denote as I an ordered subset $I = \{i_1, i_2, ..i_q\} \subseteq [2n] \equiv \{1,\dots,2n\}$ where $i_1 < i_2 < \dots i_q$ with q even. We denote C_I as the Hermitian Majorana monomial

$$C_I := i^{\binom{q}{2}} c_{i_1}..c_{i_q}, \tag{4.1}$$

as in Section 1.5.

We can think about a subset I as corresponding to a term or interaction in a Hamiltonian. Indeed, it is natural to impose some form of locality:

Definition 4.2.1 (q-local fermionic Hamiltonian)**.** *Let H be a fermionic Hamiltonian on $2n$ Majorana operators. We say that H is q-local if H is a sum of Hermitian traceless terms C_I of weight at most q, i.e. each term is proportional to a product of at most q operators c_i. H is said to be strictly q-local when all terms have exactly weight q.*

A local traceless fermionic Hamiltonian $H = \sum_{I\in\mathcal{I}} J_I C_I$ is thus characterized by an interaction set $\mathcal{I}$ and the coefficients $J_I \in \mathbb{R}$. The maximum eigenvalue of H is denoted by $\lambda_{\max}(H) := \max_\rho \operatorname{Tr}(H\rho)$. Sometimes we will refer to a collection of sets I denoted as $\mathcal{I} = \{I_1, I_2, \dots\}$. The support of $\mathcal{I}$ is defined as $\operatorname{Sup}(\mathcal{I}) = \cup_i I_i$ and $\mathcal{I}' \subseteq \mathcal{I}$ implies that the sets in $\mathcal{I}'$ are also sets in $\mathcal{I}$.

Definition 4.2.2 (SYK-q Model)**.** *A q-local (with q even) SYK model on $2n$ Majoranas is defined as a family of Hamiltonians*

$$H = \binom{2n}{q}^{-1/2} \sum_{I\subseteq[2n],|I|=q} J_I C_I, \tag{4.2}$$

where each J_I is a Gaussian random variable (i.e., with zero mean and unit variance) and each C_I is the product of the q distinct Majorana operators as in equation (4.1)*. We normalize the model in expectation, i.e.,* $\mathbb{E}\left(\operatorname{Tr}(H^2)\right) = \binom{2n}{q}^{-1} \sum_{I\subseteq[2n],|I|=q} \mathbb{E}(J_I^2)\ \operatorname{Tr}(\mathbb{1}) = 2^n$.

In Ref. [16] it was shown that with high probability (over the draw of J_Is) for the SYK-4 model, one has

$$\max_{\text{Gaussian } \rho} \operatorname{Tr}(H\rho) = O(1).$$

In order to thus provide a counterexample to a constant Gaussian approximation ratio, one needs to prove a lower bound on $\lambda_{\max}(H)$ for the SYK-4 model, which holds with high probability, which was done in Ref. [17]:

Theorem 4.2.3. *[17] There is a* $\operatorname{poly}(n)$*-time quantum algorithm that, given any SYK-4 Hamiltonian H, returns a quantum state ρ. With probability* $1 - \exp\left(-\Omega(n)\right)$ *(over the draw of the J_Is), this state ρ has* $\operatorname{Tr}(H\rho) = \Omega(\sqrt{n})$.

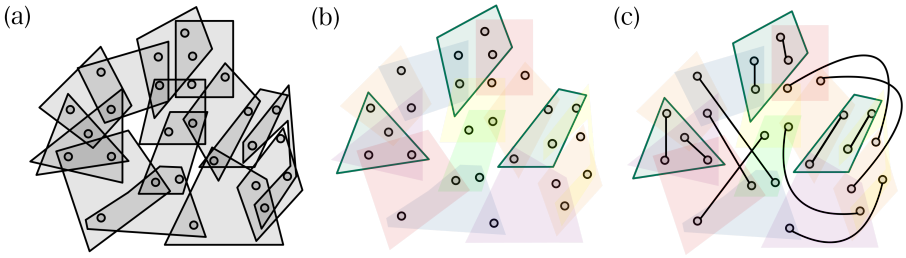

Figure 4.1: Illustrating the key idea of the proof of Theorem 4.2.5. An example of a strictly 4-local Hamiltonian is given in (a), vertices and faces representing Majorana operators and their interactions. The Hamiltonian is split into sets of terms – different colors in (b) – well separated from each other inside each set (so-called diffuse sets, see Definition 4.5.1). The next step is to match all Majorana operators, i.e., split the vertices into disjoint pairs, each connected by an edge (see panel (c)). We separately match the support of each term in one *targeted* set of terms (the color highlighted in (b) and (c)). The remaining vertices are matched in such a way that no two vertices connected by an edge belong to the same term. The Gaussian state is then created from the resulting matching, with only terms from the *targeted* set contributing to the energy. By optimizing the choice of the targeted set, a finite approximation ratio can be guaranteed.

4.2.2. Sparse fermionic Hamiltonians

Key to our work is the notion of a sparse Hamiltonian.

Definition 4.2.4. *Let H be a local traceless fermionic Hamiltonian of $2n$ Majorana operators. We say that H is k-sparse, for an integer k, if no Majorana operator c_i occurs in more than k terms of the Hamiltonian.*

Using graph theoretic terminology, one may say that interactions in a k-sparse Hamiltonian form a hypergraph of bounded degree k. This condition allows us to efficiently find Gaussian states with constant approximation ratio. We have the following theorem, which is the main result of our work:

Theorem 4.2.5. *Let H be a traceless fermionic Hamiltonian on $2n$ Majorana operators with maximal eigenvalue $\lambda_{\max}(H)$. If H is k-sparse and strictly q-local and $n > (q^2-1)k$, a Gaussian state ρ can be efficiently constructed such that*

$$\frac{\mathrm{Tr}(H\rho)}{\lambda_{\max}(H)} \geq \frac{1}{Q}, \tag{4.3}$$

for $Q = q(q-1)(k-1)^2 + q(k-1) + 2$.

The proof of this theorem is given in Section 4.5; its basic idea is explained in Figure 4.1.

We note that this proof only holds for Hamiltonians with terms of *exactly* weight q. Typical physical Hamiltonians, however, have quadratic (kinetic energy of the electrons)

and quartic terms (due e.g. to Coulomb interaction). Fortunately, we can also show that in the $q = 4$ case we can include $q = 2$ terms. For this we use a trick from [9] to lift such a 4-local Hamiltonian to a strictly 4-local Hamiltonian. This trick makes the Hamiltonian non-sparse. However, we show in Section 4.6 that, in this special case, we can circumvent the non-sparseness of the Hamiltonian and achieve a constant Gaussian approximation ratio.

Theorem 4.2.6. *Let H be a traceless fermionic Hamiltonian on $2n$ Majorana operators with maximal eigenvalue $\lambda_{\max}(H)$. If H is k-sparse with terms of weight 2 and 4 and $2n > 15k$, a Gaussian state ρ can be efficiently constructed, such that*

$$\frac{\mathrm{Tr}(H\rho)}{\lambda_{\max}(H)} \geq \frac{1}{2Q} \tag{4.4}$$

for $Q = 12(k-1)^2 + 4(k-1) + 2$.

4.2.3. The sparse $q = 4$ SYK model

In view of Theorem 4.2.5 it is worth revisiting the lack of a constant Gaussian approximation for the SYK model. The SYK-q model in Definition 4.2.2 is extremely non-sparse, in the sense that every Majorana operator occurs together with all other Majorana operators. This makes the SYK model somewhat unphysical, and several *sparse* versions of the model have been considered [19, 20]. Such sparse models intend to produce the same (low energy) physics, while being easier to simulate on both quantum and classical computers (see sections III and V in [19]). The sparse SYK model is generated by including terms by a Bernoulli trial with a certain probability p tuned such that the expected sparsity is bounded:

Definition 4.2.7. *The sparse SYK-4 or SSYK-4 model on $2n$ Majorana operators with expected sparsity $k = O(1)$ is given as*

$$H = \frac{1}{\sqrt{2kn}} \sum_{I \subset [2n], |I|=4} X_I J_I C_I \tag{4.5}$$

where the X_I are i.i.d. Bernoulli random variables with $p = \Pr(X_I = 1) = \frac{k}{\binom{2n-1}{3}}$ and the J_I are i.i.d. Gaussian random variables with mean 0 and variance 1.

Unlike the full SYK model with $\binom{2n}{4}$ terms in H, the sparse SYK model has a number of terms $\theta(n)$ in expectation. Note that the SSYK-4 model is only k-sparse in expectation, and with high probability there is a Majorana operator with degree $\Omega\big(\frac{\log(n)}{\log\log(n)}\big)$ (the degree distribution follows that of an Erdős-Renyi hypergraph. See Theorem 3.4 in [21] for a proof of the statement for Erdős-Renyi graphs. The hypergraph version follows by the same logic). This means that Theorem 4.2.5 does not directly apply. However, one can show, through a truncation argument, that almost all instantiations of SSYK-4 can be *sparsified*, giving rise to a constant approximation ratio result that holds with high probability.

Theorem 4.2.8. *Let H be a SSYK-4 Hamiltonian in equation* (4.5) *with expected degree $k = O(1)$, such that $n > 120(k+1)$. With probability at least $1 - 4\exp\left[-\frac{e^{-16(k+1)}k^3}{64(8k+7)}n\right]$, a Gaussian state ρ can be efficiently constructed such that*

$$\frac{\mathrm{Tr}(H\rho)}{\lambda_{\max}(H)} \geq \frac{1}{Q}, \tag{4.6}$$

where $Q = 1236 + 2752k + 1536k^2$.

Thus we arrive at the surprising conclusion that the SSYK-4 model has a constant Gaussian approximation ratio, while the dense SYK-4 model does not — even though SSYK-4 has similar physical properties as SYK-4.

Note that the proof of Theorem 4.2.8 is given in Ref. [22] and is not included in this thesis.

4.2.4. Higher-q SYK models

We investigate what Gaussian approximation ratios can be achieved for the *dense* SYK model of even weight $q > 4$, as this was left as an open question in [17]. We establish an upper bound on the largest Gaussian expectation value of SYK-q, which behaves rather dramatically for $q > 4$. We prove the following lemma employing a method similar to the one used in Ref. [16].

Lemma 4.2.9. *Let H be the* dense *SYK-q Hamiltonian (with even $q \geq 4$ and $q = O(1)$). With probability at least $1 - \exp\big(-\Omega(n)\big)$ over the draw of SYK-q Hamiltonians, the expectation value of every Gaussian state ρ is bounded, more precisely:*

$$\max_{\text{Gaussian } \rho} \mathrm{Tr}(H\rho) = O\big(n^{1-q/4}\big). \tag{4.7}$$

This lemma is proved in Section 4.7. Our second result establishes a lower bound on the largest eigenvalue for SYK-q, essentially generalizing what was established in [17] for $q = 4$. We prove the following Lemma (its proof can be found in Section 4.7):

Lemma 4.2.10. *Let H be the* dense *SYK-q Hamiltonian with even $q \geq 4$ (and $q = O(1)$). With probability at least $1 - \exp\big(-\Omega(n)\big)$ over the draw of SYK-q Hamiltonians, $\lambda_{\max}(H) = \Omega(\sqrt{n})$.*

As an immediate consequence of the previous results, we see that the Gaussian approximation ratio of the dense SYK-q model can be no better than $O\big(n^{1/2-q/4}\big)$:

Theorem 4.2.11. *Let H be the* dense *SYK-q Hamiltonian (with even $q \geq 4$ and $q = O(1)$). With probability at least $1 - \exp\big(-\Omega(n)\big)$ over the draw of SYK-q Hamiltonians, we have*

$$\max_{\rho \text{ Gaussian}} \frac{\mathrm{Tr}(H\rho)}{\lambda_{\max}(H)} = O\big(n^{1/2-q/4}\big). \tag{4.8}$$

Proof. Theorem 4.2.11 follows from combining Lemma 4.2.9 and Lemma 4.2.10 and applying the union bound. □

4.3. Discussion

The goal of this section is to place our results in a broader context and mention a few open questions.

First, let us discuss the relation between this work and the fermion-to-qubit mapping methods. As was shown in [23], one can map a *sparse* $O(1)$-local fermionic Hamiltonian onto a sparse $O(1)$-local qubit Hamiltonian (BK-superfast encoding). However, for this mapping one needs to enforce parity checks which are in general nonlocal; therefore, we cannot obtain our Theorem 4.2.5 in this way. There is also an additional obstacle: using the BK-superfast encoding, an approximating product state for the qubit Hamiltonian does not necessarily map back to a Gaussian fermionic state.

Ref. [23] also showed that one can map a general local fermionic Hamiltonian (like a SYK model) onto a qubit Hamiltonian with terms which are $O(\log n)$-local. Such a qubit Hamiltonian is generally not expected to have a constant approximation ratio by a product state due to its n-dependent locality. In fact, one can easily prove that a dense model like the SYK model can only be mapped onto a qubit Hamiltonian which is $\Omega(\log n)$-local. We give the argument in Appendix 4.A.

These observations suggest that approximation ratios by states such as Gaussian states or product states are likely to be affected by sparsity in the case of fermions, which is consistent with our new results.

Another question which is raised by our work and that of [17] and [12], is whether studying fermionic Hamiltonians can lead to new insights into the possibility of a quantum PCP theorem [24]. In this context it is important to mention that, besides the lower bound in Theorem 4.2.3, Ref. [17] also determined an upper bound on $\lambda_{\max}$ of the SYK-4 model showing that with high probability $\lambda_{\max} = \Theta(\sqrt{n})$. This shows that the SYK-4 model is extremely frustrated: the maximal average expected energy per term, *the energy density*, is only $\Theta(n^{-3/2})$. In contrast, our results for the sparse SYK model (see Theorem 4.2.8) show that the maximal average expected energy per term is $\Omega(1)$, which is the more 'natural' physical scaling. A simple fermionic toy model in which the maximal average energy per term decreases is a model in which an extensive set of Majorana operators is mutually anti-commuting, see Lemma 4.A.2 in Appendix 4.A. The presence of many such fully-anticommuting sets in the SYK model can be seen as one of the intuitive reasons why the maximal energy density achieved is so low.

For k-local qubit Hamiltonians researchers have looked at the hardness of approximating the maximal energy density with constant error ϵ: showing that this problem is QMA-complete would prove the quantum PCP theorem. For dense (non-sparse) k-local qubit Hamiltonians, it was proved in [12] (Theorem 13) that there is a polynomial-time classical algorithm to approximate the maximal energy density, using product state approximations. Ref. [25] generalized this result and formulated an efficient classical algorithm which approximately estimates the free energy of a 2-local dense qubit Hamiltonian.

One can similarly ask the question of approximating the maximal energy density for dense q-local fermionic Hamiltonians. Observe that the question is moot if the maximal energy density decreases as a function of n (as in the SYK model), since for large enough n (depending on ϵ) the classical algorithm could always output 0 and make an error less than ϵ. However, other dense $O(1)$-local fermionic Hamiltonians could exist for which

this question is nontrivial and not already covered by the dense qubit case.

As an open direction for further research, it would be interesting to provide a non-random family of fermionic Hamiltonians without a constant approximation ratio with respect to Gaussian states.

4.4. Additional background

In this section, we first provide some background on fermionic Gaussian states — in addition to that in Section 1.5 — that will be used throughout the remainder of this chapter.

Fermionic Gaussian states (see Definition 1.5.3) can always be brought to the standard form in equation (1.14). Such a standard form can be rephrased as a perfect matching M of Majorana operators. Such matching M is specified by n disjoint pairs (m_1, m_2) with $m_1 < m_2$. For each pair we have a coefficient $\lambda_{(m_1,m_2)} = \pm 1$, together forming the n-dimensional vector $\vec{\lambda}$. The class of states are of the form

$$\rho(M, \vec{\lambda}) = \frac{1}{2^n} \Pi_{(m_1,m_2)\in M} (\mathbb{1} + i\lambda_{(m_1,m_2)} c_{m_1} c_{m_2}). \tag{4.9}$$

It is useful to introduce a notion of *consistency* between this class of Gaussian states specified by a matching M and an interaction subset I.

Definition 4.4.1. *An (even) interaction subset $I \subseteq [2n]$ and a perfect matching M on $[2n]$ are called consistent if M contains a perfect matching of the elements of I. Given a set of interactions $\mathcal{I}$, we say that M is consistent (resp. inconsistent) with $\mathcal{I}$ if M is consistent (resp. inconsistent) with **each** interaction in $\mathcal{I}$.*

The following lemma is straightforward.

Lemma 4.4.2. *Consider a matching M and an interaction $I = \{i_1, i_2, .. i_q\}$.*

1. *If M is consistent with interaction I, let the perfect matching on the subset I be given by pairs $(i_{\pi(2l-1)}, i_{\pi(2l)})$ for $l = 1, \ldots, q/2$ and a permutation $\pi \in S_q$ where $i_{\pi(2l-1)} < i_{\pi(2l)}$. Then, the following holds:*

$$\mathrm{Tr}(C_I \rho(M, \vec{\lambda})) = \mathrm{sign}(\pi) \prod_{l\in\{1,\ldots,q/2\}} \lambda_{(i_{\pi(2l-1)}, i_{\pi(2l)})}.$$

   where $\mathrm{sign}(\pi) = \pm 1$.

2. *If M is inconsistent with I, then*

$$\mathrm{Tr}(C_I \rho(M, \vec{\lambda})) = 0. \tag{4.10}$$

Proof. In order for the trace to be nonzero, one needs to exactly match the Majorana operators in C_I with some in the expansion of $\rho(M, \vec{\lambda})$ since $\mathrm{Tr}(C_{I'}) = 0$ for any non-empty subset I'. If M is inconsistent, there is no term in the expansion of ρ which precisely

matches C_I, so the expectation vanishes. If M is consistent, we have

$$\begin{aligned}\mathrm{Tr}(C_I\rho(M,\vec{\lambda})) &= \frac{1}{2^n}\mathrm{Tr}\big(C_I\Pi_{(m_1,m_2)\in M}(\mathbb{1}+i\lambda_{(m_1,m_2)}c_{m_1}c_{m_2})\big)\\ &= \frac{1}{2^n}\mathrm{Tr}\big(C_I\Pi_{(m_1,m_2)\in M,(m_1,m_2)\cap I\neq\emptyset)}(i\lambda_{(m_1,m_2)}c_{m_1}c_{m_2})\big)\\ &= \mathrm{sign}(\pi)\prod_{l\in[1,..,q/2]}\lambda_{(i_{\pi(2l-1)},i_{\pi(2l)})}. \end{aligned} \tag{4.11}$$

Here we have used that one can first reorder C_I such that the pairs in the perfect matching are adjacent, i.e. $C_I = \mathrm{sign}(\pi)i^{q/2}c_{i_{\pi(1)}}c_{i_\pi(2)}\dots c_{i_{\pi(q)}}$, then one can commute through each pair to its matching pair in ρ and use $(c_ic_j)^2=-\mathbb{1}$, $i^q=(-1)^{q/2}$ and $\mathrm{tr}(\mathbb{1})=2^n$. □

4.5. Approximation ratios for sparse fermionic Hamiltonians

In this section we prove Theorem 4.2.5. We begin by setting up needed definitions and stating several technical lemmas (which are proved in the appendices of this chapter).

The key auxiliary notion in the proof of Theorem 4.2.5 is that of a *diffuse* subset of Hamiltonian terms. Intuitively, the terms in a diffuse subset are well separated from each other while covering only a limited part of the system. This idea is formalized as follows:

Definition 4.5.1. *Consider a set of q-local interactions $\mathcal{I}$ on $2n$ Majorana operators. A subset of these interactions $\mathcal{I}'\subset\mathcal{I}$ is* diffuse *with respect to $\mathcal{I}$, if the following three conditions apply:*

1. *$\forall I_1, I_2\in\mathcal{I}'$, I_1 and I_2 don't share any Majorana operators, i.e. $I_1\cap I_2=\emptyset$.*
2. *$\forall I_1, I_2\in\mathcal{I}'$, there exists no $I_3\in\mathcal{I}$ which shares Majorana operators with both I_1 and I_2 (if $I_3\cap I_1\neq\emptyset$ then $I_3\cap I_2=\emptyset$ and vice versa).*
3. *The size of support of $\mathcal{I}'$, i.e. $|\mathrm{Sup}(\mathcal{I}')|$, is smaller than $\frac{2qn}{q+1}$.*

In the setting of Theorem 4.2.5, diffuse sets of terms appear naturally due to the following Lemma.

Lemma 4.5.2. *Consider a k-sparse strictly q-local fermionic Hamiltonian H on $2n$ Majoranas. The interaction set $\mathcal{I}$ of H can be split into Q disjoint subsets $\mathcal{I}_\alpha$ ($\alpha\in[Q]$) all of which are diffuse with respect to $\mathcal{I}$ such that*

$$\mathcal{I}=\bigcup_{\alpha=1}^{Q}\mathcal{I}_\alpha. \tag{4.12}$$

The parameter Q is given as $Q=q(q-1)(k-1)^2+q(k-1)+2$ and does not depend on n. The construction of this splitting can be done efficiently, in time poly(n).

Lemma 4.5.2 is a special case of Lemma 4.6.1, which is proven in Appendix 4.B. The proof relies on a combinatorial argument on a graph that takes Hamiltonian terms as vertices and connects them with an edge if the pair violates conditions 1 or 2 of Definition 4.5.1. By the sparsity assumption, this graph has an efficiently constructable

coloring with a bounded number of colors, from which the split $\mathcal{I} = \bigcup_{\alpha=1}^{Q} \mathcal{I}_\alpha$ can be constructed.

The usefulness of diffuse sets comes from Lemma 4.6.2, see its proof in Appendix 4.C. Here we state its corollary, relevant to proving Theorem 4.2.5:

Lemma 4.5.3. *Let the interaction set $\mathcal{I}'$ be diffuse w.r.t. $\mathcal{I} \supset \mathcal{I}'$ ($\mathcal{I}'$ and $\mathcal{I}$ are strictly q-local and k-sparse). If $n > (q^2-1)k$, one can efficiently construct a matching M of the set $[2n]$ that is consistent with each interaction in $\mathcal{I}'$ and inconsistent with each interaction in $\mathcal{I}\backslash\mathcal{I}'$.*

With matchings introduced above, one can construct useful Gaussian states. The tool to do so is given by the following statement:

Lemma 4.5.4. *Let $H = \sum_{I\in\mathcal{I}} J_I C_I$ be strictly q-local and $\mathcal{I}'$ be a diffuse subset of $\mathcal{I}$. Let M be a matching of $[2n]$ as guaranteed by Lemma 4.5.3. One can efficiently construct a Gaussian state $\rho_{\mathcal{I}'}$ with the property:*

$$\mathrm{Tr}(H\rho_{\mathcal{I}'}) = \sum_{I\in\mathcal{I}'} |J_I|. \tag{4.13}$$

Lemma 4.5.4 is a specific case of a slightly more general Lemma 4.6.3, which is stated and proven in Appendix 4.D. We denote

$$\mathbf{J}\left(\mathcal{I}'\right) := \sum_{I\in\mathcal{I}'} |J_I|. \tag{4.14}$$

As shown below, Theorem 4.2.5 can be proven by constructing a diffuse $\mathcal{I}' \subset \mathcal{I}$ and a corresponding Gaussian state $\rho_{\mathcal{I}'}$ with large enough $\mathrm{Tr}(H\rho_{\mathcal{I}'}) = \mathbf{J}\left(\mathcal{I}'\right)$.

Theorem (Repetition of Theorem 4.2.5)**.** *Let H be a traceless fermionic Hamiltonian on $2n$ Majoranas with maximal eigenvalue $\lambda_{\max}(H)$. If H is k-sparse and strictly q-local and $n > (q^2-1)k$, a Gaussian state ρ can be efficiently constructed, such that*

$$\frac{\mathrm{Tr}(H\rho)}{\lambda_{\max}(H)} \geq \frac{1}{Q}, \tag{4.15}$$

for $Q = q(q-1)(k-1)^2 + q(k-1) + 2$.

Proof. For a Hamiltonian $H = \sum_{I\in\mathcal{I}} J_I C_I$, we construct the splitting of $\mathcal{I}$ into diffuse subsets $\mathcal{I} = \cup_\alpha \mathcal{I}_\alpha$ as guaranteed by Lemma 4.5.2. Next, find $\alpha = \mathrm{argmax}_{\alpha'} \mathbf{J}(\mathcal{I}_{\alpha'})$; since Q in Lemma 4.5.2 is constant, α can be found efficiently. Next, use Lemma 4.5.3 to construct a matching $M(\mathcal{I}_\alpha)$ (the condition $n > (q^2-1)k$ is satisfied by assumptions of Theorem 4.2.5). Since $\mathcal{I}_\alpha$ is diffuse with respect to $\mathcal{I}$, the Gaussian state $\rho_{\mathcal{I}_\alpha}$ can be efficiently constructed from $M(\mathcal{I}_\alpha)$ via Lemma 4.5.4. Using $\mathrm{Tr}(H\rho_{\mathcal{I}_\alpha}) = \mathbf{J}(\mathcal{I}_\alpha)$, the following inequality can be obtained for the resulting approximation ratio:

$$\frac{\mathrm{Tr}(H\rho_{\mathcal{I}_\alpha})}{\lambda_{\max}(H)} \geq \frac{\mathbf{J}(\mathcal{I}_\alpha)}{\sum_{\alpha'} \mathbf{J}(\mathcal{I}_{\alpha'})} \geq \frac{1}{Q}. \tag{4.16}$$

For the first inequality, note that $\lambda_{\max}(H) \leq \sum_{I\in\mathcal{I}} |J_I| = \sum_\alpha \mathbf{J}(\mathcal{I}_\alpha)$. The second inequality comes from a pigeonhole-type argument: if $\mathbf{J}(\mathcal{I}_\alpha) = \max_{\alpha'} \mathbf{J}(\mathcal{I}_{\alpha'})$, it directly follows that $\mathbf{J}(\mathcal{I}_\alpha) \geq \frac{1}{Q}\sum_{\alpha'} \mathbf{J}(\mathcal{I}_{\alpha'})$. The inequality in equation (4.16) concludes the proof, as it asserts the approximation ratio bound claimed in the Theorem. □

4.6. Sparse Hamiltonians with terms of weight 2 and 4

In this section we prove Theorem 4.2.6. We will again need to use the concept of diffuse subsets in Definition 4.5.1. The proof of Theorem 4.2.6 is similar in its basic idea to that of Theorem 4.2.5. The main obstacle in this case is the presence of terms of different weight, which does not allow one to use Lemmas 4.5.2-4.5.4 directly. This can be resolved by a slightly more elaborate construction and applying the more general Lemmas 4.6.1-4.6.3 which are proved in the appendices of this chapter and Lemmas 4.5.2-4.5.4 directly follow as special cases.

Lemma 4.6.1 (Generalization of Lemma 4.5.2)**.** *Let $\mathcal{I}$ be the interaction set of a k-sparse q-local Hamiltonian on the set of Majorana fermions $[2n]$. The set $\mathcal{I}$ can be split into $(qQ)/2$ disjoint, strictly $2q'$-local subsets $\mathcal{I}_\alpha^{(2q')}$ (with $\alpha \in [Q]$ and $q' \in [q/2]$) each of which is diffuse with respect to $\mathcal{I}$:*

$$\mathcal{I} = \bigcup_{q'=1}^{q/2} \bigcup_{\alpha=1}^{Q} \mathcal{I}_\alpha^{(2q')}. \tag{4.17}$$

The parameter $Q = q(q-1)(k-1)^2 + q(k-1) + 2$ does not grow with n. The construction of this splitting can be done efficiently, in time $\mathrm{poly}(n)$.

Lemma 4.6.2 (Generalization of Lemma 4.5.3)**.** *Let a strictly q'-local $\mathcal{I}'$ be diffuse w.r.t. q-local k-sparse $\mathcal{I}$ on $[2n]$, such that $n > (q^2-1)k$. One can efficiently construct a matching M of $[2n]$ that is consistent with $\mathcal{I}'$ and inconsistent with all interactions $I \in \mathcal{I}\backslash\mathcal{I}'$ such that (1) $|I| \geq q'$ or (2) $I \notin \mathrm{Sup}(\mathcal{I}')$.*

Lemma 4.6.3 (Generalization of Lemma 4.5.4)**.** *Let $H = \sum_{I\in\mathcal{I}} J_I C_I$ on $[2n']$ be q-local and $\mathcal{I}'$ be a diffuse subset of $\mathcal{I}$. Consider a matching M of $[2n']$. If M is consistent with $\mathcal{I}'$ and inconsistent with $\mathcal{I}\backslash\mathcal{I}'$, one can efficiently construct a Gaussian state $\rho_{\mathcal{I}'}$ with the property:*

$$\mathrm{Tr}(H\rho_{\mathcal{I}'}) = \sum_{I\in\mathcal{I}'} |J_I|. \tag{4.18}$$

In Lemma 4.6.3, we use n' instead of n to avoid confusion, as it will also be used for $n' \neq n$. The Lemmas above are proven in Appendices 4.B-4.D. With these in hand, we are ready to proceed with the proof of Theorem 4.2.6.

Theorem (Repetition of Theorem 4.2.6)**.** *Let H be a traceless fermionic Hamiltonian on $[2n]$ with maximal eigenvalue $\lambda_{\max}(H)$. If H is k-sparse with terms of weight 2 and 4 and $2n > 15k$, a Gaussian state ρ can be efficiently constructed, such that*

$$\frac{\mathrm{Tr}(H\rho)}{\lambda_{\max}(H)} \geq \frac{1}{2Q} \tag{4.19}$$

with $Q = 12(k-1)^2 + 4(k-1) + 2$.

Proof. We make use of the construction in Ref. [9] which relates a Hamiltonian with weights 2 and 4 on a set of fermionic modes $[2n]$, that is,

$$H = \sum_{I\in\mathcal{I}^{(2)}} J_I C_I + \sum_{I\in\mathcal{I}^{(4)}} J_I C_I, \tag{4.20}$$

to a strictly 4-local Hamiltonian $\tilde{H}$ on an extended set of fermions $[2n+2]$:

$$\tilde{H} = \sum_{I\in\mathcal{I}^{(2)}} (-ic_{2n+1}c_{2n+2})J_I C_I + \sum_{I\in\mathcal{I}^{(4)}} J_I C_I. \tag{4.21}$$

Introducing $\tilde{\mathcal{I}}^{(2)} \equiv \{(2n+1,2n+2)\cup I \,|\, I\in\mathcal{I}^{(2)}\}$, $\tilde{H}$ can be also written as:

$$\tilde{H} = -\sum_{I\in\tilde{\mathcal{I}}^{(2)}} J_I C_I + \sum_{I\in\mathcal{I}^{(4)}} J_I C_I. \tag{4.22}$$

The relation between $\tilde{H}$ and H is via the following property:

Lemma 4.6.4 (Lemma 6 of [9])**.** *For H and $\tilde{H}$ introduced above, $\lambda_{\max}(H) = \lambda_{\max}(\tilde{H})$. Moreover, for any Gaussian state $\tilde{\rho}$ of $2n+2$ Majorana modes, one can efficiently compute a Gaussian state ρ of $2n$ Majorana modes s.t.* $\mathrm{Tr}(H\rho) \geq \mathrm{Tr}(\tilde{H}\tilde{\rho})$.

Although strictly 4-local, Hamiltonian $\tilde{H}$ is no longer sparse since the operators c_{2n+1} and c_{2n+2} participate in $|\mathcal{I}^{(2)}|$ terms (which is generally $O(n)$). This prevents a direct application of Lemma 4.5.2 to $\tilde{H}$. We resolve the issue as follows.

Similarly to the proof of Theorem 4.2.5, we start by splitting each set of the original interactions $\mathcal{I}^{(2,4)}$ in H into subsets diffuse w.r.t. $\mathcal{I}^{(2)}\cup\mathcal{I}^{(4)}$: $\mathcal{I}^{(2)} = \cup_\alpha \mathcal{I}^{(2)}_\alpha$, $\mathcal{I}^{(4)} = \cup_\alpha \mathcal{I}^{(4)}_\alpha$. Each of the two splittings exists and can be done efficiently, as guaranteed by Lemma 4.5.2 (since the original H is sparse). Since $\mathcal{I}^{(2)}\cup\mathcal{I}^{(4)}$ is k-sparse and 4-local, we can bound $|\{\mathcal{I}^{(2)}_\alpha\}| < Q$, $|\{\mathcal{I}^{(4)}_\alpha\}| < Q$ for $Q = 12(k-1)^2 + 4(k-1) + 2$. In what follows, we will use the splittings $\mathcal{I} = \bigcup_\alpha \mathcal{I}^{(2)}_\alpha \cup \bigcup_\alpha \mathcal{I}^{(4)}_\alpha$ to construct two Gaussian states $\tilde{\rho}(\mathcal{I}^{(2)}_\alpha)$ and $\tilde{\rho}(\mathcal{I}^{(4)}_\alpha)$ on $[2n+2]$ with good properties relative to $\tilde{H}$, that is,

$$\mathrm{Tr}(\tilde{H}\tilde{\rho}(\mathcal{I}^{(2,4)}_\alpha)) = \sum_{I\in\mathcal{I}^{(2,4)}_\alpha} |J_I| =: \mathbf{J}(\mathcal{I}^{(2,4)}_\alpha). \tag{4.23}$$

With these Gaussian states, we will then show that the Gaussian state $\tilde{\rho}(\mathcal{I}^{(2,4)}_\alpha)$ for $q,\alpha = \mathrm{argmax}_{q',\alpha'}(\mathbf{J}(\mathcal{I}^{(q')}_{\alpha'}))$ is efficiently constructable and yields the desired approximation ratio for $\tilde{H}$. We will then apply Lemma 4.6.4 and extend the statement to the original Hamiltonian H, thus finishing the proof.

Following the outline above, we now move to construct the Gaussian state $\tilde{\rho}(\mathcal{I}^{(2)}_\alpha)$. Consider an ansatz of the form $\tilde{\rho} := \rho_{[2n]}\sigma_{\{2n+1,2n+2\}}$, where $\rho_{[2n]}$ is itself a Gaussian state of $[2n]$. To construct $\rho_{[2n]}$, note that each $\mathcal{I}^{(2)}_\alpha$ is 2−local and diffuse w.r.t. $\mathcal{I}^{(2)}\cup\mathcal{I}^{(4)}$ which is 4-local. Since $2n > 15k$ by assumptions of Theorem 4.2.6, we can apply Lemma 4.6.2 with $q=4$ to construct a matching $M(\mathcal{I}^{(2)}_\alpha)$ that is consistent with $\mathcal{I}^{(2)}_\alpha$. Since $\mathcal{I}^{(2)}_\alpha$ is 2-local, Lemma 4.6.2 also implies that the matching $M(\mathcal{I}^{(2)}_\alpha)$ is inconsistent with the entirety of $\mathcal{I}^{(4)}\cup\mathcal{I}^{(2)}\setminus\mathcal{I}^{(2)}_\alpha$. We then use $M(\mathcal{I}^{(2)}_\alpha)$ in Lemma 4.6.3 (substituting $n'=n$ for n' used in the Lemma) to construct $\rho_{[2n]}$ in $\tilde{\rho} = \rho_{[2n]}\sigma_{\{2n+1,2n+2\}}$. This implies the following expression (using equation (4.21) for $\tilde{H}$):

$$\mathrm{Tr}(\tilde{H}\tilde{\rho}) = \mathrm{Tr}(\sigma(-ic_{2n+1}c_{2n+2}))\sum_{I\in\mathcal{I}^{(2)}_\alpha} |J_I|. \tag{4.24}$$

By choosing σ to be the $+1$ eigenstate projector of operator $-ic_{2n+1}c_{2n+2}$, we arrive at the desired outcome:

$$\mathrm{Tr}(\tilde{H}\tilde{\rho}) = \sum_{I\in\mathcal{I}^{(2)}_\alpha} |J_I| = \mathbf{J}(\mathcal{I}^{(2)}_\alpha). \tag{4.25}$$

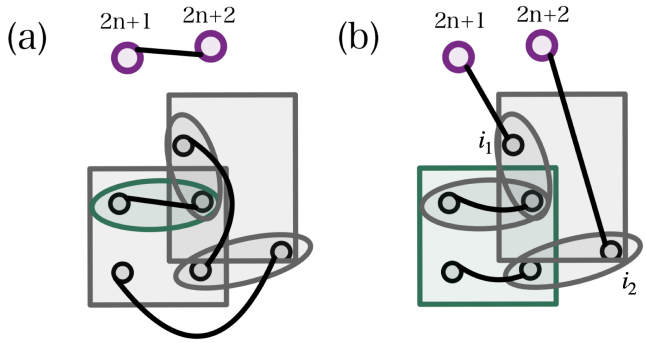

Figure 4.2: Demonstration of the method in the proof of Theorem 4.2.6. (a) Matching $M(\mathcal{I}_\alpha^{(2)})$ for $\mathcal{I}_\alpha^{(2)}$, here comprised of a single term (shown in green). To ensure consistency with $\mathcal{I}_\alpha^{(2)}$ in $\tilde{H}$, $M(\mathcal{I}_\alpha^{(2)})$ perfectly matches these terms and the pair $(2n+1, 2n+2)$. The rest of the vertices are matched so that each pair does not belong to the same term in $\mathcal{I}\backslash\mathcal{I}_\alpha^{(2)}$ (grey). (b) Matching $M(\mathcal{I}_\alpha^{(4)})$ for $\mathcal{I}_\alpha^{(4)}$ shown in green. Vertices i_1, i_2 are chosen not to belong to the same term in $\mathcal{I}^{(2)}$, ensuring no accidental consistency with a term in $\tilde{H}$. Note the special status of the term from $\mathcal{I}^{(2)}$ that is a subset of the $\mathcal{I}_\alpha^{(4)}$ term. From the perspective of $\tilde{H}$, it is not consistent with $M(\mathcal{I}_\alpha^{(4)})$ although it coincides with an edge from $M(\mathcal{I}_\alpha^{(4)})$. This is due to the intentional absence of the edge $(2n+1, 2n+2)$ in $M(\mathcal{I}_\alpha^{(4)})$.

The constructed Gaussian state $\tilde{\rho}$ we will denote as $\tilde{\rho}(\mathcal{I}_\alpha^{(2)})$.

For a diffuse $\mathcal{I}_\alpha^{(4)} \subset \mathcal{I}^{(2)} \cup \mathcal{I}^{(4)}$, we construct the Gaussian states $\tilde{\rho}(\mathcal{I}_\alpha^{(4)})$ in a different way. First we use Lemma 4.6.2 to construct a matching $M(\mathcal{I}_\alpha^{(4)})$ of $[2n]$. This matching is guaranteed to be consistent with $\mathcal{I}_\alpha^{(4)}$. However, since $\mathcal{I}^{(2)}$ is 2-local and $\mathcal{I}^{(4)}$ is 4-local, while in general $\mathrm{Sup}(\mathcal{I}^{(2)}) \cap \mathrm{Sup}(\mathcal{I}_\alpha^{(4)}) \neq \emptyset$, Lemma 4.6.2 implies that $M(\mathcal{I}_\alpha^{(4)})$ is inconsistent with $\mathcal{I}^{(4)}\backslash\mathcal{I}_\alpha^{(4)}$ but may be consistent with some terms in $\mathcal{I}^{(2)}$ (as those I don't obey the $|I| \geq q' = 4$ condition). At the same time, we aim to achieve $\mathrm{Tr}(\tilde{H}\tilde{\rho}(\mathcal{I}_\alpha^{(4)})) = \mathbf{J}(\mathcal{I}_\alpha^{(4)})$ which excludes contributions from $\mathcal{I}^{(2)}$. Thus we cannot extend $M(\mathcal{I}_\alpha^{(4)})$ to the extended set $[2n+2]$ directly, as it was done for $\mathcal{I}_\alpha^{(2)}$. Instead, we will create a matching of $[2n+2]$ using a reduced version of $M(\mathcal{I}_\alpha^{(4)})$ which inherits its beneficial properties, and then complete the matching by making it inconsistent with $\tilde{\mathcal{I}}^{(2)}$ – eliminating the difficulty described above.

To enable this, we find and mark an edge $(i_1, i_2) \in M(\mathcal{I}_\alpha^{(4)})$, such that $i_1 \notin \mathrm{Sup}(\mathcal{I}_\alpha^{(4)})$. This is always possible since $\mathcal{I}_\alpha^{(4)}$ is diffuse and thus $[2n]\backslash\mathrm{Sup}(\mathcal{I}_\alpha^{(4)})$ is non-empty (cf. Condition 3 in Definition 4.5.1). Note that $M(\mathcal{I}_\alpha^{(4)})$ is constructed via Lemma 4.6.2 and $\{i_1, i_2\} \notin \mathrm{Sup}(\mathcal{I}_\alpha^{(4)})$. This implies that as a two-fermion interaction, $\{i_1, i_2\}$ is guaranteed not to belong to $\mathcal{I}^{(2)}$. The latter statement is the key property of the marked edge (i_1, i_2) that we will employ momentarily.

We construct a matching $\tilde{M}(\mathcal{I}_\alpha^{(4)})$ of $[2n+2]$ in two stages. First we construct an intermediate matching $M'(\mathcal{I}_\alpha^{(4)})$ of $[2n+2]\backslash\{i_1, i_2, 2n+1, 2n+2\}$ by removing the edge

from $M(\mathcal{I}_\alpha^{(4)})$:

$$M'(\mathcal{I}_\alpha^{(4)}) = M(\mathcal{I}_\alpha^{(4)}) \backslash (i_1, i_2). \tag{4.26}$$

Since $\{i_1, i_2\} \not\subset \mathrm{Sup}(\mathcal{I}_\alpha^{(4)})$, we are guaranteed that $M'(\mathcal{I}_\alpha^{(4)})$ is consistent with $\mathcal{I}_\alpha^{(4)}$ and inconsistent with $\mathcal{I}^{(4)} \backslash \mathcal{I}_\alpha^{(4)}$ (from the construction of $M(\mathcal{I}_\alpha^{(4)})$). In the second stage, we complete $\tilde{M}(\mathcal{I}_\alpha^{(4)})$ to the entire set of $2n+2$ modes by adding two edges: $(i_1, 2n+1)$ and $(i_2, 2n+2)$:

$$\tilde{M}(\mathcal{I}_\alpha^{(4)}) = M(\mathcal{I}_\alpha^{(4)}) \cup \{(i_1, 2n+1), (i_2, 2n+2)\}. \tag{4.27}$$

These new edges render $\tilde{M}(\mathcal{I}_\alpha^{(4)})$ inconsistent with $\tilde{\mathcal{I}}^{(2)}$. To see it, note that all interactions in $\tilde{\mathcal{I}}^{(2)}$ take the form $I = \{j_1, j_2, 2n+1, 2n+2\}$ where $\{j_1, j_2\} \in \mathcal{I}^{(2)}$. By construction $\{i_1, i_2\} \notin \mathcal{I}^{(2)}$, thus we have $\{j_1, j_2\} \neq \{i_1, i_2\}$. As a result, matching $\tilde{M}(\mathcal{I}_\alpha^{(4)})$ of $[2n+2]$ is consistent with $\mathcal{I}_\alpha^{(4)}$ and inconsistent with $\tilde{\mathcal{I}}^{(2)} \cup \mathcal{I}^{(4)} \backslash \mathcal{I}_\alpha^{(4)}$. We continue by applying Lemma 4.6.3 to such $\tilde{M}(\mathcal{I}_\alpha^{(4)})$ and $\tilde{H}$ (substituting $n' = n+1$ for n' used in the Lemma). This efficiently constructs a Gaussian state $\tilde{\rho}(\mathcal{I}_\alpha^{(4)})$ that yields:

$$\mathrm{Tr}(\tilde{H}\tilde{\rho}(\mathcal{I}_\alpha^{(4)})) = \sum_{I \in \mathcal{I}_\alpha^{(4)}} |J_I| \equiv \mathbf{J}(\mathcal{I}_\alpha^{(4)}), \tag{4.28}$$

as desired.

The Gaussian state claimed in Theorem 4.2.6 is to be chosen among the states $\tilde{\rho}(\mathcal{I}_\alpha^{(2,4)})$ whose existence we've proven above. We make the choice by identifying the highest energy in the respective Gaussian state: $(q, \alpha) = \mathrm{argmax}_{(q,\alpha)} \mathbf{J}_\alpha^{(q)}$. As we showed, the respective Gaussian state $\tilde{\rho}(\mathcal{I}_\alpha^{(q)})$ can be efficiently constructed and the following is guaranteed:

$$\frac{\mathrm{Tr}(\tilde{H}\tilde{\rho}(\mathcal{I}_\alpha^{(q)}))}{\lambda_{\max}(\tilde{H})} \geq \frac{\mathbf{J}_\alpha^{(q)}}{\sum_{q',\alpha'} \mathbf{J}_{\alpha'}^{(q')}} \geq \frac{1}{2Q}. \tag{4.29}$$

Here we used that $\lambda_{\max}(\tilde{H}) \leq \sum_{q',\alpha'} \mathbf{J}_{\alpha'}^{(q')}$ and that $\mathbf{J}_\alpha^{(q)} = \max_{(q',\alpha')} \mathbf{J}_{\alpha'}^{(q')}$.

With the state $\tilde{\rho}(\mathcal{I}_\alpha^{(q)})$ on $[2n+2]$ fermions at hand, we finalize the proof by an application of Lemma 4.6.4. This relates $\lambda_{\max}(H)$ to $\lambda_{\max}(\tilde{H})$ and allows us to efficiently construct the Gaussian state $\rho(\mathcal{I}_\alpha^{(q)})$ of $[2n]$, with the desired property:

$$\frac{\mathrm{Tr}(H\rho(\mathcal{I}_\alpha^{(q)}))}{\lambda_{\max}(H)} \geq \frac{\mathrm{Tr}(\tilde{H}\tilde{\rho}(\mathcal{I}_\alpha^{(q)}))}{\lambda_{\max}(\tilde{H})} \geq \frac{1}{2Q}. \tag{4.30}$$

□

4.7. Upper bound on Gaussian approximation ratio for SYK-q Hamiltonians

4.7.1. Upper bound for Gaussian expectations on SYK-q models

We consider the expectation value of a SYK-q Hamiltonian H with respect to fermionic Gaussian states and we obtain an upper bound on its expectation value, with high probability over the random couplings J_I.

Lemma (Repetition of Lemma 4.2.9)**.** *Let H denote a Hamiltonian drawn from the q-local SYK Hamiltonians (with $q \geq 4$ even and $q = O(1)$), i.e. the coupling strengths J_I are drawn according to their distribution. With probability at least $1 - \exp(-\Omega(n))$, H has the property that, for any fermionic Gaussian state ρ*

$$\mathrm{Tr}(H\rho) \leq (q-1)!!\, 2^{1/2-q/4} q^{1/2+q/2} \sqrt{\log[q/\log(3/2)]}\, (2n)^{1-q/4}. \tag{4.31}$$

Proof. We first use Wick's theorem (Proposition 1.5.8) on the expectation of a product of Majorana operators w.r.t. a fermionic Gaussian state ρ characterized by a correlation matrix Γ. Note that the correlation matrix $\Gamma_{i<j}$ can be viewed as a real $d := (2n^2 - n)$-dimensional vector. We note that $\sum_{i<j} \Gamma_{ij}^2 = \frac{1}{2}\,\mathrm{Tr}(\Gamma^T\Gamma) \leq \frac{1}{2}\,\mathrm{Tr}(\mathbb{1}) = n$ so that $\|\Gamma\|_2 \leq n^{1/2}$.

Let $M(I)$ be a perfect matching of the indices in I ($|I|$ even), there are $(q-1)!!$ such matchings. We have

$$\mathrm{tr}(C_I\rho) = i^{q/2} \sum_{M(I)} \mathrm{sign}\big(M(I)\big)\, \mathrm{tr}\big(c_{i_1(M)} c_{i_2(M)}\rho\big)\mathrm{tr}(c_{i_3(M)} c_{i_4(M)}\rho)\ldots\mathrm{tr}(c_{i_{q-1}(M)} c_{i_q(M)}\rho). \tag{4.32}$$

Here we have assumed that for each matching $M(I)$; $i_1(M) < i_2(M)$, $i_3(M) < i_4(M)$, …, $i_{q-1}(M) < i_q(M)$, i.e. any sign arising from getting the expression to this form is absorbed in $\mathrm{sign}\big(M(I)\big)$.

The expectation of H in equation (4.2) w.r.t. fermionic Gaussian states ρ can be written as:

$$\begin{aligned}\mathrm{tr}(H\rho) &= \binom{2n}{q}^{-1/2} i^{q/2} \sum_{I\subseteq[2n],|I|=q} J_I \Big[\sum_{M(I)} \mathrm{sign}\big(M(I)\big) \prod_{t=1}^{q/2} \mathrm{tr}\big(c_{i_{2t-1}(M)} c_{i_{2t}(M)}\rho\big)\Big] \\ &= \binom{2n}{q}^{-1/2} \sum_{I\subseteq[2n],|I|=q} J_I \sum_{M(I)} \mathrm{sign}\big(M(I)\big) \prod_{t=1}^{q/2} \Gamma_{i_{2t-1}(M),i_{2t}(M)}.\end{aligned} \tag{4.33}$$

We note that we can view $\mathrm{Tr}(H\rho)$ as a sum of $(q-1)!!$ terms, one for each matching M of some subset of indices I, i.e. $\mathrm{Tr}(H\rho) = \binom{2n}{q}^{-1/2} \sum_M \mathrm{Tr}(H_M\rho)$ where $H_M = \sum_I \tilde{J}(M,I) \prod_{t=1}^{q/2} \Gamma_{i_{2t-1}(M),i_{2t}(M)}$. We have defined the $q/2$-way, $d \times d \times \ldots \times d$, tensor $\tilde{J}(M,I)$, whose entries are equal to either zero (when the indices coincide or are not ordered properly) or to a standard Gaussian random variable. Each J_I appears only once in $\mathrm{tr}(H_M\rho)$ and therefore all entries of $\tilde{J}(M,I)$ are statistically independent. We note that $\mathrm{sign}(M)$ does not depend on which (ordered) subset I one chooses. To bound each term $\mathrm{Tr}(H_M\rho)$, with high probability, we invoke the following lemma:

Lemma 4.7.1. *(Theorem 1 in [26].) Let A be a random K-way tensor $\in \mathbb{R}^{d_1\times d_2\times\ldots\times d_K}$ and w_i be vectors $\in \mathbb{R}^{d_i}$ and*

$$A(w_1,w_2,\ldots,w_K) := \sum_{k_1,\ldots,k_K} A_{k_1,\ldots,k_K}(w_1)_{k_1}\ldots(w_{q/2})_{k_K}.$$

If we have for each fixed unit vector $w_i/\|w_i\|$ ($i\in\{1,\ldots,K\}$):

$$\Pr\Big(|A(w_1/\|w_1\|,\ldots,w_K/\|w_K\|)|\geq t\Big)\leq 2\exp\Big(-t^2/(2\sigma^2)\Big), \tag{4.34}$$

then the spectral norm $|||A||| := \max_{w_1,\ldots,w_K} A\big(w_1/\|w_1\|,\ldots,w_K/\|w_K\|\big)$ (with $w_i\in\mathbb{R}^{d_i}$) can be bounded as follows:

$$|||A|||\leq\Big[8\sigma^2\Big[\Big(\sum_{i=1}^{K}d_i\Big)\log\big[2K/\log(3/2)\big]+\log\Big(\frac{2}{\delta}\Big)\Big]\Big]^{\frac{1}{2}},$$

with probability at least $1-\delta$.

To apply the Lemma, note that the vectors w_i correspond to $\Gamma_{i<j}$ viewing $i<j$ as a single index and we can use their norm $\|\Gamma\|\leq n^{1/2}$. In addition, for each entry in the tensor we have $\mathbb{E}\big[\exp\big(t\tilde{J}(M,I)_{k_1,\ldots,k_{q/2}}\big)\big]\leq\exp\big(t^2/2\big)$ (for $t\geq 0$) as the entry is zero or a Gaussian variable with variance 1 and mean zero. Using Chernoff's bound and the fact that all entries of $\tilde{J}(M,I)$ are statistically independent, we conclude that for any set of real vectors $w_1,\ldots,w_{q/2}$ one has

$$\Pr\Big[\Big|\sum_{k_1,\ldots,k_{q/2}}\tilde{J}(M,I)_{k_1,\ldots,k_{q/2}}\frac{(w_1)_{k_1}}{\|w_1\|}\ldots\frac{(w_{q/2})_{k_{q/2}}}{\|w_{q/2}\|}\Big|\geq t\Big]\leq 2\exp\big(-t^2/2\big). \tag{4.35}$$

Therefore, for each term H_M we can apply Lemma 4.7.1 and, using $K=q/2$ and $\sigma=1$, obtain

$$\Big|\Big|\Big|\tilde{J}(M,I)\Big|\Big|\Big|\leq\Big[4q(2n^2-n)\log\big[q/\log(3/2)\big]+8\log\big(2\delta^{-1}\big)\Big]^{1/2}, \tag{4.36}$$

with probability at least $1-\delta$. Then we can first bound

$$\begin{aligned}\max_{\text{Gaussian }\rho}\operatorname{tr}(H\rho)&\leq\binom{2n}{q}^{-1/2}\|\Gamma\|^{q/2}\sum_M\Big|\Big|\Big|\tilde{J}(M,I)\Big|\Big|\Big|\\&\leq\big(q/\sqrt{2}\big)^{q/2}(2n)^{-q/4}\sum_M\Big|\Big|\Big|\tilde{J}(M,I)\Big|\Big|\Big|,\end{aligned} \tag{4.37}$$

where we have used that $\binom{2n}{q}\geq(2n/q)^q$. We can now combine the upper bound in equation (4.36) and equation (4.37). Applying the union bound, we have with probability at least $1-(q-1)!!\,\delta$, that

$$\begin{aligned}\max_{\text{Gaussian }\rho}\operatorname{tr}(H\rho)\leq(q-1)!!\Big[&2^{1-q/2}q^{q+1}\big[(2n)^{2-q/2}-(2n)^{1-q/2}\big]\log\big[q/\log(3/2)\big]\\&+2^{3-q/2}q^q(2n)^{-q/2}\log\big(2\delta^{-1}\big)\Big]^{1/2}.\end{aligned} \tag{4.38}$$

Therefore, we can take $\delta = \exp\big(-\Omega(n)\big)$ such that, asymptotically, we have (assuming $q = O(1)$):

$$\max_{\text{Gaussian } \rho} \operatorname{tr}(H\rho) \leq (q-1)!!\, 2^{1/2-q/4} q^{1/2+q/2} \sqrt{\log[q/\log(3/2)]}\, (2n)^{1-q/4}, \tag{4.39}$$

with probability at least $1-\delta$. Note that in deriving this upper bound we only use the norm of the correlation matrix Γ, hence this upper bound is not necessarily achievable by a Gaussian state (or any state) as the constraint $\Gamma^T\Gamma \leq \mathbb{1}$ imposes more conditions on Γ than just an upper bound on its norm. □

4.7.2. Maximum eigenvalue lower bound for q-local SYK Hamiltonians

To show that fermionic Gaussian states cannot achieve a constant approximation ratio for $q \geq 4$ SYK models, we derive a lower bound on the maximum eigenvalue of the Hamiltonians H in equation (4.2):

Lemma (Repetition of Lemma 4.2.10)**.** *For the class of q-local SYK Hamiltonians (with even $q \geq 4$) in equation (4.2), $\lambda_{\max}(H) = \Omega(\sqrt{n})$ with probability at least $1-\exp\big(-\Omega(n)\big)$ over the draw of Hamiltonians.*

The remainder of this section will be devoted to proving this lemma. The techniques used are similar to those used in Section 6 of Ref. [17]. We note that throughout this section, we shall use C to denote a quantity that is constant in n or is bounded from above and below by a constant in n, and it will generally differ from appearance to appearance (for the sake of clarity). Importantly, C can contain factors of q (note that $q = O(1)$).

We start by obtaining a lower bound on the maximum eigenvalue of a so-called 2-colored SYK model and will use this to prove Lemma 4.2.10. The Hamiltonian of such a 2-colored SYK model is slightly different from the standard SYK model Hamiltonian in equation (4.2). We divide the $2n$ Majorana operators into two subsets, with sizes n_1 and n_2 ($n_2 \leq n_1$), and denote the operators in the first set by $\phi_1, \ldots, \phi_{n_1}$ and the ones in the second set by $\chi_1, \ldots, \chi_{n_2}$. The Hamiltonian is now given by[1]:

$$H^{(2)} = \frac{i}{\sqrt{n_2}} \sum_{j=1}^{n_2} \tau_j \chi_j, \tag{4.40}$$

where

$$\tau_j = \binom{n_1}{q-1}^{-1/2} i^{q/2-1} \sum_{\substack{S \subseteq [n_1] \\ |S| = q-1}} J_{S,j}\, \phi^S. \tag{4.41}$$

Here ϕ^S the product of $q-1$ of the ϕ Majorana operators in subset S, and $J_{S,j}$ are independent Gaussian random variables. The subset S labels an ordered subset of $q-1$ Majorana operators (note that these are different from the subsets I defined before that correspond to ordered subsets of q Majorana operators). We note that the (Hermitian) τ_j operators do not necessarily obey $\{\tau_j, \tau_k\} = 2\delta_{jk}\mathbb{1}$, but instead satisfy $\mathbb{E}(\{\tau_j, \tau_k\}) = -i^{q-2}\delta_{jk}\mathbb{1}$.

[1] We denote Hamiltonians from the class of 2-colored SYK Hamiltonians by $H^{(2)}$.

Lemma 4.7.2. *Let $\{\phi_i\}_{i=1}^{n_1}$ and $\{\chi_i\}_{i=1}^{n_2}$ be n_1+n_2 Majorana operators. For the class of q-local 2-colored SYK Hamiltonians (with even $q \geq 4$) in equation (4.40) defined in terms of these Majorana operators, the maximum eigenvalue of the Hamiltonian $\lambda_{\max}(H^{(2)})$ is lower bounded by $C\sqrt{n}$ (with C a constant) with probability at least $1-\exp\big(-\Omega(n)\big)$ over the draw of Hamiltonians.*

Proof. We introduce a new set of Majorana operators (again of size n_2) $\sigma_1,\ldots,\sigma_{n_2}$ (which do obey $\{\sigma_j,\sigma_k\}=2\delta_{jk}\mathbb{1}$) and we define the *quadratic* Hamiltonian H':

$$H' = \frac{i}{\sqrt{n_2}}\sum_{j=1}^{n_2}\sigma_j\chi_j. \tag{4.42}$$

This quadratic Hamiltonian H' is optimized by the fermionic Gaussian state $\rho_0 = \frac{1}{2^{n_2+n_1/2}}\prod_{j=1}^{n_2}\big(\mathbb{1}+i\sigma_j\chi_j\big)$, which achieves $\mathrm{tr}(H'\rho_0)=\sqrt{n_2}$. The idea is now to construct a new state ρ_θ obtained from ρ_0 by applying a unitary transformation to ρ_0, and to find a lower bound for the expectation value of $H^{(2)}$ w.r.t. ρ_θ.

$$\rho_\theta := e^{-\theta\zeta}\rho_0 e^{+\theta\zeta}, \quad \text{where } \zeta := \sum_{j=1}^{n_2}\tau_j\sigma_j \text{ and } \theta\in\mathbb{R}. \tag{4.43}$$

The expectation value of $H^{(2)}$ w.r.t. ρ_θ is:

$$\mathrm{tr}(H^{(2)}\rho_\theta) = \mathrm{tr}(H_\theta^{(2)}\rho_0), \tag{4.44}$$

where $H_\theta^{(2)} := e^{+\theta\zeta}H^{(2)}e^{-\theta\zeta}$. Using the BCH expansion of H_θ and $\mathrm{tr}(H^{(2)}\rho_0)=0$, we obtain:

$$\begin{aligned}
\mathrm{tr}(H^{(2)}\rho_\theta) &= \theta\,\mathrm{tr}([\zeta,H^{(2)}]\rho_0) + \theta^2\int_0^1(1-s)\mathrm{tr}([\zeta,[\zeta,H^{(2)}]]\rho_{s\theta})\,ds \\
&= \theta\,\mathrm{tr}([\zeta,H^{(2)}]\rho_0) + \theta^2\mathop{\mathbb{E}}_{s\sim[0,1]}\Big[(1-s)\mathrm{tr}([\zeta,[\zeta,H^{(2)}]]\rho_{s\theta})\Big] \\
&\geq \theta\,\mathrm{tr}([\zeta,H^{(2)}]\rho_0) - \theta^2\|[\zeta,[\zeta,H^{(2)}]]\|,
\end{aligned} \tag{4.45}$$

where we have used the triangle inequality and $\|\cdot\|$ denotes the spectral norm. To lower bound $\mathrm{tr}(H^{(2)}\rho_\theta)$, one now has to (i) lower bound $\theta\,\mathrm{tr}([\zeta,H^{(2)}]\rho_0)$ and (ii) upper bound $\theta^2\|[\zeta,[\zeta,H^{(2)}]]\|$. Note that this proof technique is similar in spirit to the proof in [27], although their proof is for qubit Hamiltonians with bounded-degree interactions.

First, we find a lower bound for $\theta\,\mathrm{tr}([\zeta,H^{(2)}]\rho_0)$ which holds with high probability:

$$\begin{aligned}
\mathrm{tr}([\zeta,H^{(2)}]\rho_0) = \frac{i}{\sqrt{n_2}}\sum_{j,k=1}^{n_2}\mathrm{tr}([\tau_j\sigma_j,\tau_k\chi_k]\rho_0) &= \frac{i}{\sqrt{n_2}}\sum_{j=1}^{n_2}\mathrm{tr}([\tau_j\sigma_j,\tau_j\chi_j]\rho_0) \\
&= \frac{2i}{\sqrt{n_2}}\sum_{j=1}^{n_2}\mathrm{tr}(\sigma_j\chi_j\tau_j^2\rho_0) \\
&= \frac{2}{\sqrt{n_2}}2^{-(n_2+n_1/2)}\sum_{j=1}^{n_2}\mathrm{tr}(\mathbb{1}_{n_2}\tau_j^2) \\
&= \frac{2(-1)^{q/2}}{\sqrt{n_2}\binom{n_1}{q-1}}\sum_{j=1}^{n_2}\sum_{\substack{S\subseteq[n_1]\\|S|=q-1}}\big(J_{S,j}\big)^2,
\end{aligned} \tag{4.46}$$

where we have used that $\mathrm{tr}([\tau_j\sigma_j,\tau_k\chi_k]\rho_0)$ is non-zero only for $j=k$, and the definition of τ_j. The quantity $\mathrm{tr}([\zeta,H^{(2)}]\rho_0)$ is thus a chi-squared random variable (up to normalization factors and potentially a sign) with $n_2\binom{n_1}{q-1}$ degrees of freedom and its expectation value is given by:

$$\mathbb{E}\big[\mathrm{tr}([\zeta,H^{(2)}]\rho_0)\big]=\frac{2(-1)^{q/2}}{\sqrt{n_2}\binom{n_1}{q-1}}\sum_{j=1}^{n_2}\sum_{\substack{S\subseteq[n_1]\\|S|=q-1}}\mathbb{E}\big[\big(J_{S,j}\big)^2\big]=2\sqrt{n_2}\,(-1)^{q/2}, \tag{4.47}$$

where we have used that $\mathbb{E}\big[\big(J_{S,j}\big)^2\big]=1$. We note that in order to obtain a positive first-order contribution to $\mathrm{tr}(H^{(2)}\rho_\theta)$, one should take θ positive for $q/2$ even, and one should take θ negative for $q/2$ odd. Since $\mathrm{tr}([\zeta,H^{(2)}]\rho_0)$ is a chi-squared random variable with $n_2\binom{n_1}{q-1}$ degrees of freedom, the following tail bounds can be obtained [28]:

$$\Pr\big[\mathrm{tr}([\zeta,H^{(2)}]\rho_0)\leq\sqrt{n_2}\big]\leq\exp\Big(-\Omega\big(n_2\,n_1^{q-1}\big)\Big), \tag{4.48}$$

for $q/2$ even, and

$$\Pr\big[\mathrm{tr}([\zeta,H^{(2)}]\rho_0)\geq-\sqrt{n_2}\big]\leq\exp\Big(-\Omega\big(n_2\,n_1^{q-1}\big)\Big), \tag{4.49}$$

for $q/2$ odd. The random variable $\mathrm{tr}([\zeta,H^{(2)}]\rho_0)$ is thus equal to $2\sqrt{n_2}(-1)^{q/2}$ in expectation and the probability that – for any even $q\geq 4$ – its norm is smaller than half the norm of this expectation is at most exponentially small in the system size.

In order to upper bound $\theta^2\parallel[\zeta,[\zeta,H^{(2)}]]\parallel$, we first evaluate $[\zeta,[\zeta,H^{(2)}]]$:

$$\begin{aligned}[\zeta,[\zeta,H^{(2)}]]&=\frac{i^{q/2}}{\sqrt{n_2\binom{n_1}{q-1}}}\sum_{j=1}^{n_2}\sum_{\substack{S\subseteq[n_1]\\|S|=q-1}}J_{S,j}[\zeta,[\zeta,\phi^S\chi_j]]\\&=\frac{i^{q/2}}{\sqrt{n_2\binom{n_1}{q-1}}}\sum_{j,k,l=1}^{n_2}\sum_{\substack{S\subseteq[n_1]\\|S|=q-1}}J_{S,j}[\tau_k\sigma_k,[\tau_l\sigma_l,\phi^S\chi_j]]\\&=\frac{i^{3q/2-2}}{\sqrt{n_2\binom{n_1}{q-1}^3}}\sum_{j,k,l=1}^{n_2}\sum_{S,S',S''}J_{S,j}J_{S',k}J_{S'',l}[\phi^{S'}\sigma_k,[\phi^{S''}\sigma_l,\phi^S\chi_j]],\end{aligned} \tag{4.50}$$

where the final sum over S,S',S'' is over all $S,S',S''\subseteq[n_1]$ with $|S|=|S'|=|S''|=q-1$ (all sums over S,S',S'' will implicitly have this constraint from now on). The nested commutator in this expression simplifies as follows (note that the product of $i^{3q/2-2}$ and the nested commutator is Hermitian):

$$i^{3q/2-2}[\phi^{S'}\sigma_k,[\phi^{S''}\sigma_l,\phi^S\chi_j]]=\begin{cases}C(\phi^K\sigma_k\sigma_l\chi_j)_H, & \text{if}\quad(|S''\cap S|\text{ is odd})\\ & \quad\wedge(|S'\cap(S''\triangle S)|+\delta_{k,l}\text{ is odd})\\ & \quad\wedge(|S|=|S'|=|S''|=q-1),\\ 0, & \text{otherwise,}\end{cases} \tag{4.51}$$

where $(\phi^K \sigma_k \sigma_l \chi_j)_H$ denotes a Hermitian version of $\phi^K \sigma_k \sigma_l \chi_j$ (i.e., $\phi^K \sigma_k \sigma_l \chi_j$ up to potential integer powers of i) and $K := (S \triangle S' \triangle S'') \cup (S \cap S' \cap S'')$ (note that $|K|$ is odd). We therefore have:

$$[\zeta,[\zeta,H^{(2)}]] = C\frac{1}{\sqrt{n_2}}\binom{n_1}{q-1}^{-3/2}\sum_{j,k,l=1}^{n_2}\sum_{S,S',S''} J_{S,j}J_{S',k}J_{S'',l}\Big((\phi^K\sigma_k\sigma_l\chi_j)_H f(S,S',S'',j,k,l)\Big), \tag{4.52}$$

where we have defined

$$f(S,S',S'',j,k,l) := \begin{cases} 1, & \text{if} \quad (|S''\cap S| \text{ is odd}) \\ & \quad \wedge (|S'\cap(S''\triangle S)| + \delta_{k,l} \text{ is odd}) \\ & \quad \wedge (|S| = |S'| = |S''| = q-1), \\ 0, & \text{otherwise.} \end{cases} \tag{4.53}$$

We now wish to find an upper bound on the expected value of the spectral norm of $[\zeta,[\zeta,H^{(2)}]]$. And in addition, we would like to show that the spectral norm exceeds twice the value of this upper bound with probability that is at most exponentially small in the system size. To establish this, we will have to show the following:

$$\mathbb{E}\Big(\| [\zeta,[\zeta,H^{(2)}]] \|^k\Big) \leq \alpha^k, \tag{4.54}$$

for *even* k proportional to the system size and for some α. Equation (4.54) implies two things: First, since $\mathbb{E}\big(\| [\zeta,[\zeta,H^{(2)}]] \|\big)^k \leq \mathbb{E}\big(\| [\zeta,[\zeta,H^{(2)}]] \|^k\big)$ (using Jensen's inequality), it implies $\mathbb{E}\big(\| [\zeta,[\zeta,H^{(2)}]] \|\big) \leq \alpha$ (i.e., α is the upper bound on the expected value of the spectral norm). Second, applying Markov's inequality to the random variable $\| [\zeta,[\zeta,H^{(2)}]] \|$ and using equation (4.54) yields

$$\Pr\Big[\| [\zeta,[\zeta,H^{(2)}]] \| \geq \alpha'\Big] = \Pr\Big[\| [\zeta,[\zeta,H^{(2)}]] \|^k \geq (\alpha')^k\Big] \leq \big(\alpha/\alpha'\big)^k, \tag{4.55}$$

with $\alpha' \geq \alpha$. So taking $\alpha' = 2\alpha$ and k equal to the system size $2n$ $(= 2n_2 + n_1)$ yields the desired result of the probability of the spectral norm exceeding twice the value of the upper bound being at most exponentially small in the system size.

For convenience, we define $A := [\zeta,[\zeta,H^{(2)}]]$. Since A is Hermitian (by direct calculation), the spectrum of A^2 is non-negative and therefore we have $\|A\|^k = \lambda_{\max}(A^2)^{k/2} \leq \mathrm{tr}(A^k)$ (for even k). Using equation (4.52), we express A as $C\sum_{\tilde{S}\subseteq[2n_2+n_1]} Q_{\tilde{S}} C_{\tilde{S}}$ for convenience, where C is a non-negative constant, $Q_{\tilde{S}}$ are real random variables, and $C_{\tilde{S}}$ denotes a Hermitian (even) Majorana monomial. In addition, we define the random variable (which is obtained by replacing Majorana monomials in A with 1)

$$A(1) := C\sum_{\tilde{S}\subseteq[2n_2+n_1]} Q_{\tilde{S}} = C\frac{1}{\sqrt{n_2}}\binom{n_1}{q-1}^{-3/2}\sum_{j,k,l=1}^{n_2}\sum_{S,S',S''} J_{S,j}J_{S',k}J_{S'',l} f(S,S',S'',j,k,l). \tag{4.56}$$

If we now *assume* that

$$\mathbb{E}\big(Q_{\tilde{S}_1}\dots Q_{\tilde{S}_k}\big)\geq 0 \quad \text{and} \quad \mathbb{E}\Big(A(1)^k\Big)/\alpha^k \leq 1/2^{n_2+n_1/2}, \tag{4.57}$$

both hold for some even k and some constant α (note that the first condition will automatically be satisfied since $\{J_{S,j}\}$ is a collection of independent standard Gaussian random variables), then for even k we can establish

$$\begin{aligned}\mathbb{E}\big(\|A\|\big)^k \leq \mathbb{E}\big(\|A\|^k\big) \leq \mathbb{E}\big(\mathrm{tr}(A^k)\big) &= C^k \sum_{\substack{\tilde{S}_1,\dots,\tilde{S}_k \\ \subseteq[2n_2+n_1]}} \mathbb{E}\big(Q_{\tilde{S}_1}\dots Q_{\tilde{S}_k}\big)\mathrm{Re}\big[\mathrm{tr}\big(C_{\tilde{S}_1}\dots C_{\tilde{S}_k}\big)\big] \\ &\leq 2^{n_2+n_1/2}\, C^k \sum_{\substack{\tilde{S}_1,\dots,\tilde{S}_k \\ \subseteq[2n_2+n_1]}} \mathbb{E}\big(Q_{\tilde{S}_1}\dots Q_{\tilde{S}_k}\big) \\ &= 2^{n_2+n_1/2}\,\mathbb{E}\big(A(1)^k\big) \leq \alpha^k,\end{aligned} \tag{4.58}$$

where the first inequality is again Jensen's inequality and we have also used that $\mathbb{E}\big(\mathrm{tr}(A^k)\big)$ is real (since A is Hermitian) and that $\mathrm{Re}\big[\mathrm{tr}\big(C_{\tilde{S}_1}\dots C_{\tilde{S}_k}\big)\big]$ is always at most $2^{n_2+n_1/2}$ (note that $\mathrm{tr}\big(C_{\tilde{S}_1}\dots C_{\tilde{S}_k}\big)$ equals $2^{n_2+n_1/2}$ up to integer powers of i, but imaginary contributions vanish in the sum). This establishes equation (4.54), and thereby the desired result. Therefore, what is left is to show that the second condition in equation (4.57) is satisfied.

From this point onward, we shall take n_1 and n_2 proportional to n, where $2n = 2n_2 + n_1$ denotes the total number of Majorana operators. We now show that the second condition in equation (4.57) is satisfied for $k = 2n$ and $\alpha = C\sqrt{n}$. In order to do so, we show that $\mathbb{E}\big(A(1)^{2n}\big) \leq (C\sqrt{n})^{2n}$ (where the factor of $2^{n_2+n_1/2}$ is absorbed in C^{2n}). To that end, we thus need to find an upper bound on the $(2n)$th moment of the random variable $A(1)$ in equation (4.56).

In Appendix 4.E, we derive this upper bound and indeed show that $\mathbb{E}\big(A(1)^{2n}\big) \leq (C\sqrt{n})^{2n}$. Therefore,

$$\mathbb{E}\big(\|\,[\zeta,[\zeta,H^{(2)}]]\,\|\big) \leq C\sqrt{n} \quad \text{and} \quad \Pr\Big[\|\,[\zeta,[\zeta,H^{(2)}]]\,\| \geq 2C\sqrt{n}\Big] \leq \exp\big(-\Omega(n)\big), \tag{4.59}$$

which is the desired result.

Combining equation (4.45), equation (4.48), equation (4.49) and equation (4.59), we conclude that there exists a $\theta = O(1)$ such that

$$\mathrm{tr}(H^{(2)}\rho_\theta) \geq C\sqrt{n}, \tag{4.60}$$

with probability at least $1-\exp\big(-\Omega(n)\big)$. □

What is left is to show that this result also holds for the standard SYK Hamiltonian. This translation from 2-colored SYK Hamiltonian to standard SYK Hamiltonian is given in Lemma 4.7.3 below, and its proof is given in Appendix 4.F.

Lemma 4.7.3. *For the class of q-local SYK Hamiltonians (with even $q \geq 4$) in equation* (4.2)*, ρ_θ (defined in equation* (4.43)*) achieves* $\mathrm{Tr}(H\rho_\theta) \geq C\sqrt{n}$ *with probability*

at least $1-\exp\big(-\Omega(n)\big)$ *over the draw of Hamiltonians, provided that* ρ_θ *achieves* $\mathrm{Tr}(H^{(2)}\rho_\theta) \geq C\sqrt{n}$ *(with* $H^{(2)}$ *the* 2*-coloured SYK Hamiltonian defined in equation* (4.40)*) with probability at least* $1-\exp\big(-\Omega(n)\big)$ *over the draw of* 2*-coloured Hamiltonians.*

This also concludes the proof of Lemma 4.2.10, i.e., that $\lambda_{\max}(H) = \Omega(\sqrt{n})$ with probability at least $1-\exp\big(-\Omega(n)\big)$ over the draw of standard SYK Hamiltonians.

References

[1] S. Gharibian, Y. Huang, Z. Landau and S. W. Shin. 'Quantum Hamiltonian Complexity'. In: *Foundations and Trends in Theoretical Computer Science* 10.3 (2015), pp. 159–282. URL: https://doi.org/10.1561%2F0400000066.

[2] S. Khanna, M. Sudan, L. Trevisan and D. P. Williamson. 'The approximability of constraint satisfaction problems'. In: *SIAM Journal on Computing* 30.6 (2001), pp. 1863–1920. URL: https://doi.org/10.1137/S0097539799349948.

[3] M. X. Goemans and D. P. Williamson. 'Improved Approximation Algorithms for Maximum Cut and Satisfiability Problems Using Semidefinite Programming'. In: *J. ACM* 42.6 (1995), pp. 1115–1145. URL: https://doi.org/10.1145/227683.227684.

[4] C. V. Kraus and J. I. Cirac. 'Generalized Hartree–Fock theory for interacting fermions in lattices: numerical methods'. In: *New Journal of Physics* 12.11 (2010), p. 113004. URL: https://doi.org/10.1088/1367-2630/12/11/113004.

[5] S. Bravyi and D. Gosset. 'Complexity of Quantum Impurity Problems'. In: *Communications in Mathematical Physics* 356.2 (Aug. 2017), pp. 451–500. ISSN: 1432-0916. DOI: 10.1007/s00220-017-2976-9.

[6] S. Bravyi and R. Koenig. 'Classical simulation of dissipative fermionic linear optics'. In: *Quant. Inf. Comp.* 12.11–12 (2012), pp. 925–943. URL: https://doi.org/10.48550/arXiv.1112.2184.

[7] F. d. Melo, P. Ćwikliński and B. M. Terhal. 'The power of noisy fermionic quantum computation'. In: *New Journal of Physics* 15.1 (Jan. 2013), p. 013015. ISSN: 1367-2630. DOI: 10.1088/1367-2630/15/1/013015.

[8] E. H. Lieb. 'The classical limit of quantum spin systems'. In: *Communications in Mathematical Physics* 31.4 (1973), pp. 327–340. URL: https://doi.org/10.1007/BF01646493.

[9] S. Bravyi, D. Gosset, R. König and K. Temme. 'Approximation algorithms for quantum many-body problems'. In: *Journal of Mathematical Physics* 60.3 (2019), p. 032203. eprint: https://doi.org/10.1063/1.5085428. URL: https://doi.org/10.1063/1.5085428.

[10] N. Bansal, S. Bravyi and B. M. Terhal. 'Classical approximation schemes for the ground-state energy of quantum and classical Ising spin Hamiltonians on planar graphs'. In: *Quantum Inf. Comp.* 9.7-8 (2009), pp. 701–720. URL: https://doi.org/10.48550/arXiv.0705.1115.

[11] S. Gharibian and J. Kempe. 'Approximation Algorithms for QMA-Complete Problems'. In: *SIAM Journal on Computing* 41.4 (2012), pp. 1028–1050. URL: https://doi.org/10.1137/110842272.

[12] F. G. Brandao and A. W. Harrow. 'Product-State Approximations to Quantum Ground States'. In: *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*. STOC '13. Palo Alto, California, USA: Association for Computing Machinery, 2013, pp. 871–880. ISBN: 978-1-4503-2029-0. URL: https://doi.org/10.1145/2488608.2488719.

[13] A. W. Harrow and A. Montanaro. 'Extremal eigenvalues of local Hamiltonians'. In: *Quantum* 1 (Apr. 2017), p. 6. URL: https://doi.org/10.22331/q-2017-04-25-6.

[14] T. Bergamaschi. *Improved Product-state Approximation Algorithms for Quantum Local Hamiltonians*. 2022. URL: https://doi.org/10.48550/arXiv.2210.08680.

[15] S. Gharibian and O. Parekh. 'Almost Optimal Classical Approximation Algorithms for a Quantum Generalization of Max-Cut'. en. In: (2019). URL: https://doi.org/10.4230/LIPICS.APPROX-RANDOM.2019.31.

[16] A. Haldar, O. Tavakol and T. Scaffidi. 'Variational wave functions for Sachdev-Ye-Kitaev models'. In: *Phys. Rev. Research* 3 (2 2021), p. 023020. URL: https://doi.org/10.1103/PhysRevResearch.3.023020.

[17] M. B. Hastings and R. O'Donnell. 'Optimizing Strongly Interacting Fermionic Hamiltonians'. In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2022. Rome, Italy: ACM, 2022, pp. 776–789. URL: https://doi.org/10.1145/3519935.3519960.

[18] N. Alon, K. Makarychev, Y. Makarychev and A. Naor. 'Quadratic forms on graphs'. In: *Inventiones mathematicae* 163.3 (2006), pp. 499–522. URL: https://doi.org/10.1007/s00222-005-0465-9.

[19] S. Xu, L. Susskind, Y. Su and B. Swingle. *A Sparse Model of Quantum Holography*. 2020. URL: https://doi.org/10.48550/arXiv.2008.02303.

[20] A. M. García-García, Y. Jia, D. Rosa and J. J. M. Verbaarschot. 'Sparse Sachdev-Ye-Kitaev model, quantum chaos, and gravity duals'. In: *Phys. Rev. D* 103 (10 2021), p. 106002. URL: https://doi.org/10.1103/PhysRevD.103.106002.

[21] A. Frieze and M. Karoński. *Introduction to random graphs*. Cambridge University Press, 2016. URL: https://doi.org/10.1017/CBO9781316339831.

[22] Y. Herasymenko, M. Stroeks, J. Helsen and B. Terhal. 'Optimizing sparse fermionic Hamiltonians'. In: *Quantum* 7 (Aug. 2023), p. 1081. ISSN: 2521-327X. DOI: 10.22331/q-2023-08-10-1081.

[23] S. B. Bravyi and A. Y. Kitaev. 'Fermionic Quantum Computation'. In: *Annals of Physics* 298.1 (2002), pp. 210–226. URL: https://doi.org/10.1006/aphy.2002.6254.

[24] D. Aharonov, I. Arad and T. Vidick. 'The Quantum PCP Conjecture'. In: *ACM SIGACT News* 44 (2 2013), pp. 47–79. URL: https://doi.org/10.48550/arXiv.1309.7495.

[25] S. Bravyi, A. Chowdhury, D. Gosset and P. Wocjan. *On the complexity of quantum partition functions*. 2021. URL: https://doi.org/10.48550/arXiv.2110.15466.

[26] R. Tomioka and T. Suzuki. 'Spectral norm of random tensors'. In: *arXiv* (2014). URL: https://doi.org/10.48550/arXiv.1407.1870.

[27] A. Anshu, D. Gosset, K. J. Morenz Korol and M. Soleimanifar. 'Improved Approximation Algorithms for Bounded-Degree Local Hamiltonians'. In: *Phys. Rev. Lett.* 127 (25 2021), p. 250502. URL: https://doi.org/10.1103/PhysRevLett.127.250502.

[28] B. Laurent and P. Massart. 'Adaptive estimation of a quadratic functional by model selection'. In: *The Annals of Statistics* 28.5 (2000), pp. 1302–1338. URL: https://doi.org/10.1214/aos/1015957395.

[29] B. Bollobás. *Modern Graph Theory*. Graduate Texts in Mathematics 184. Springer-Verlag New York, 1998. URL: https://doi.org/10.1007/978-1-4612-0619-4.

[30] G. A. Dirac. 'Some theorems on abstract graphs'. In: *Proceedings of the London Mathematical Society* 3.1 (1952), pp. 69–81. URL: https://doi.org/10.1112/plms/s3-2.1.69.

[31] R. Latala. 'Estimates of moments and tails of Gaussian chaoses'. In: *The Annals of Probability* 34.6 (Nov. 2006). URL: https://doi.org/10.1214%2F009117906000000421.

[32] V. H. de la Pena, S. J. Montgomery-Smith and J. Szulga. 'Contraction and Decoupling Inequalities for Multilinear Forms and U-Statistics'. In: *The Annals of Probability* 22.4 (1994), pp. 1745–1765. URL: https://doi.org/10.48550/arXiv.math/9406214 (visited on 06/09/2022).

[33] R. Adamczak and P. Wolff. 'Concentration inequalities for non-Lipschitz functions with bounded derivatives of higher order'. In: *Probability Theory and Related Fields* 162.3 (2015), pp. 531–586. URL: https://doi.org/10.1007/s00440-014-0579-3.

4.A. Extensive sets of all anti-commuting terms

One can easily prove that when one maps a dense, non-sparse, fermionic model such as the SYK model onto a qubit Hamiltonian, the locality of the resulting Hamiltonian has to grow as some function of n, due to the following Lemma:

Lemma 4.A.1. *Any set of all-mutually anti-commuting Pauli strings $\{Q_i\}_{i=0}^{m-1}$, each of weight at most k, on n qubits has cardinality m bounded as*

$$m \leq 3 \times 2^{k(3k-1)}, \tag{4.61}$$

assuming that $k(k-1) < n$.

Proof. Take Q_0 of weight at most k and let $m-1$ Paulis Q_i anticommute with it. We can represent each Pauli string as a $2n$-bit string y, say $Q_0 = y_x y_z$ where the Hamming weight $|y_x| \leq k, |y_z| \leq k$. Any other Q_i in the set has to anti-commute with Q_0 on the support of the string y. First, note that the set of strings of length at most $2k$ which have symplectic inner product equal to 1 (so anti-commute) to a given string of length $2k$ is at most 2^{2k-1}. Now we pick the largest subset $\mathcal{M}_1$ of the set of elements $Q_1, \dots Q_{m-1}$ such that all elements in the subset act identically on the support of Q_0, i.e. are represented by the same string of length at most $2k$ while differing beyond the support of Q_0. Let the cardinality of this set be $|\mathcal{M}_1| = m_1 \leq m-1$ and $m_1 \geq \frac{m-1}{2^{2k-1}} \geq \frac{m}{2^{2k}}$ as the largest set should at least be a fraction $1/2^{2k-1}$ of the total. So now we consider this set $\mathcal{M}_1$ and their action on the remaining $n-k$ qubits (outside the support of Q_0), where these elements all have to anti-commute. In addition, each element has Pauli weight at most $k-1$ (as we had to overlap with at least one Pauli with Q_0). We then reapply this argument on this set, leading to a new set $\mathcal{M}_2$ with $|\mathcal{M}_2| = m_2 \geq \frac{m_1-1}{2^{2(k-1)-1}}$ acting on $n-2k$ qubits and having weight $k-2$ etc. We can reiterate this process l times so that the remaining weight of the set of Pauli strings $\mathcal{M}_l$ has $k-l=1$. This implies that $\mathcal{M}_l$ can contain at most 3 elements since they all need to anti-commute on a single qubit (assuming that $n-kl>0$ or $n-k(k-1)>0$). So we have

$$3 \geq |\mathcal{M}_{l=k-1}| = m_{k-1} \geq \frac{m}{4^{k+(k-1)+\dots+l}} = \frac{m}{2^{k(3k-1)}}. \tag{4.62}$$

□

The SYK-4 model contains large (of size n) sets of mutually anti-commuting terms. An example is the set of all terms which only overlap on one fixed Majorana. Lemma 4.A.1 then shows that any fermion-to-qubit mapping (an encoding possibly using more qubits) will require the weight of some of the resulting Pauli terms to grow as a function of n. Note that the actual mapping by Bravyi and Kitaev [23] with $k = O(\log n)$ shows that the upper bound in equation (4.61) is not completely tight.

Another straightforward observation on the energy scaling of a model where all terms anti-commute is that $\lambda_{\max}$ does not necessarily scale with the number of terms, as captured by the following lemma.

Lemma 4.A.2. *Let $H = \sum_{i\in\mathcal{I}} J_I C_I$ where the $\{C_I\}$ are a set of all-mutually anti-commuting Majorana operators on $[2n]$ (each C_I has even support). Then*

$$\lambda_{\max}(H) = \sqrt{\sum_I J_I^2}. \tag{4.63}$$

Proof. We have $H = \sum_I J_I C_I = \sqrt{\sum_I J_I^2}\sum_I \beta_I C_I$ with $\sum_I \beta_I^2 = 1$. Take the state $\rho = \frac{1}{2^n}(\mathbb{1} + \sum_I \beta_I C_I)$ and thus $\mathrm{Tr}(H\rho) = \sqrt{\sum_I J_I^2}\sum_I \beta_I^2 = \sqrt{\sum_I J_I^2}$. This is the maximal eigenvalue that can be reached since one can map each c_I onto a single Majorana operator $c_{i(I)}$ as these sets form identical algebras. Then we can use the normalization of β_I to view $\sum_i \beta_I c_{i(I)} = \tilde{c}_1$ with single Majorana operator $\tilde{c}_1$ (this is an example of a canonical transformation as discussed in Section 1.5). A single Majorana $\tilde{c}_1$ has spectrum ± 1 and hence the (hugely degenerate) spectrum of H is simply $\pm\sqrt{\sum_I J_I^2}$. □

Thus, if all J_I are of similar strength, we observe that the overall maximal energy scales as $\sqrt{|\mathcal{I}|}$ rather than $|\mathcal{I}|$.

4.B. Splitting sparse Hamiltonians into diffuse interaction sets

Lemma (Repetition of Lemma 4.6.1)**.** *Let $\mathcal{I}$ be the interaction set of a k-sparse q-local Hamiltonian on the set of fermions $[2n]$. The set $\mathcal{I}$ can be split into $(qQ)/2$ disjoint, strictly $2q'$-local subsets $\mathcal{I}_\alpha^{(2q')}$ (with $\alpha \in [Q]$ and $q' \in [q/2]$) each of which is diffuse with respect to $\mathcal{I}$:*

$$\mathcal{I} = \bigcup_{q'=1}^{q/2} \bigcup_{\alpha=1}^{Q} \mathcal{I}_\alpha^{(2q')}. \tag{4.64}$$

The parameter $Q = q(q-1)(k-1)^2 + q(k-1) + 2$ does not grow with n. The construction of this splitting can be done efficiently, in time poly(n).

Proof. Consider a graph $\mathcal{G}$ with vertices corresponding to interaction sets $I \in \mathcal{I}$, where two interaction sets I_1, I_2 are connected with an edge if either 1. they share at least one Majorana operator or 2. I_1 and I_2 both share Majorana operators with another set $I' \neq I_1, I_2$. For a q-local k-sparse Hamiltonian, $\mathcal{G}$ has maximal degree Q' with $Q' = q(q-1)(k-1)^2 + q(k-1)$. Here $q(k-1)$ is the maximal number of interactions I_2 directly sharing a Majorana fermion with any given interaction I_1, and $q(q-1)(k-1)^2$ is the maximal number of interactions satisfying condition 2. Since a Q'-sparse graph is vertex-colorable by at most $(Q'+1)$ colors [29], we can split $\mathcal{I}$ into $(Q'+1)$ subsets $\mathcal{I}_\alpha$, s.t. any two interactions I_1, I_2 from a set $\mathcal{I}_\alpha$ are not connected by an edge in $\mathcal{G}$. By definition of $\mathcal{G}$, this amounts to sets $\mathcal{I}_\alpha$ satisfying the first two conditions of Definition 4.5.1. A greedy algorithm can be used to assign the vertices $\mathcal{G}$ with $(Q'+1)$ colors, so $\mathcal{I}_\alpha$ can be constructed efficiently.

Each interaction set $\mathcal{I}_\alpha$ can contain terms of different weight. For each value of α we define strictly $2q'$-local sets $\mathcal{I}_\alpha^{(2q')}$ (for $q' = 1, .., q/2$) by restricting to the strictly $2q'$-local part of $\mathcal{I}_\alpha$. This gives a splitting of $\mathcal{I}$ into efficiently constructable subsets $\mathcal{I}_\alpha^{(2q')}$:

$$\mathcal{I} = \bigcup_{q'=1}^{q/2} \bigcup_{\alpha=1}^{Q'+1} \mathcal{I}_\alpha^{(2q')}, \tag{4.65}$$

where all sets $\mathcal{I}_\alpha^{(2q')}$ satisfy conditions 1 and 2 in Definition 4.5.1.

The rest of the proof is concerned with the third condition of a diffuse set in Definition 4.5.1, for all sets $\mathcal{I}_\alpha^{(2q')}$. This means ensuring that for all values of α and q', the support size $|\mathrm{Sup}(\mathcal{I}_\alpha^{(2q')})|$ is smaller than $2n\frac{q}{q+1}$. Fix q' and consider sets $\mathcal{I}_\alpha^{(2q')}$ for $\forall \alpha \in [Q'+1]$.

Consider the case where $|\mathrm{Sup}(\mathcal{I}_\alpha^{(2q')})| < 2n\frac{q}{q+1}$ does not hold for at least one value of α, which we set to be $\alpha = Q'+1$ without loss of generality.

Let us prove that the violation $|\mathrm{Sup}(\mathcal{I}_\beta^{(2q')})| \geq 2n\frac{q}{q+1}$ cannot hold for any $\beta \neq Q'+1$. Firstly, no interaction I from $\mathcal{I}_\beta^{(2q')}$ can be a strict subset of an interaction in $\mathcal{I}_{Q'+1}^{(2q')}$ or

share Majoranas with two terms in $\mathcal{I}^{(2q')}_{Q'+1}$ simultaneously. The first scenario is excluded since $\mathcal{I}^{(2q')}_{Q'+1}$ and $\mathcal{I}^{(2q')}_{\beta}$ are both strictly $2q'$-local and the second scenario is excluded because $\mathcal{I}^{(2q')}_{Q'+1}$ satisfies condition 2 of Definition 4.5.1. From these two facts it follows that each interaction in $\mathcal{I}^{(2q')}_{\beta}$ must involve at least one Majorana from $[2n]\backslash\mathrm{Sup}(\mathcal{I}^{(2q')}_{Q'+1})$. This implies

$$|\mathrm{Sup}(\mathcal{I}^{(2q')}_{\beta})| \leq 2q'|[2n]\backslash\mathrm{Sup}(\mathcal{I}^{(2q')}_{Q'+1})|, \tag{4.66}$$

This can be further bounded as $|\mathrm{Sup}(\mathcal{I}^{(2q')}_{\beta})| \leq q|[2n]\backslash\mathrm{Sup}(\mathcal{I}^{(2q')}_{Q'+1})|$, because $2q' \leq q$. Since we assumed $|\mathrm{Sup}(\mathcal{I}^{(2q')}_{Q'+1})| \geq 2nq/(q+1)$ and thus $|[2n]\backslash\mathrm{Sup}(\mathcal{I}^{(2q')}_{Q'+1})| < 2n/(q+1)$, it follows that $|\mathrm{Sup}(\mathcal{I}^{(2q')}_{\beta})| \leq q|[2n]\backslash\mathrm{Sup}(\mathcal{I}^{(2q')}_{Q'+1})| < 2nq/(q+1)$. Thus we have shown that for a given q', the condition 3 of Definition 4.5.1 – indeed cannot be violated by more than one $\mathcal{I}^{(2q')}_{\alpha}$.

Consider all q' for which there exists a violation $|\mathrm{Sup}(\mathcal{I}^{(2q')}_{Q'+1})| \geq 2n\frac{q}{q+1}$. Since $\frac{q}{q+1} > \frac{1}{2}$ for any q, this violation can be fixed by splitting $\mathcal{I}^{(2q')}_{Q'+1}$ in half. Introduce non-overlapping sets $\tilde{\mathcal{I}}^{(2q')}_{Q'+1}$ and $\tilde{\tilde{\mathcal{I}}}^{(2q')}_{Q'+1}$ of sizes $\lfloor|\mathcal{I}^{(2q')}_{Q'+1}|/2\rfloor$ and $\lceil|\mathcal{I}^{(2q')}_{Q'+1}|/2\rceil$: $\mathcal{I}^{(2q')}_{Q'+1} = \tilde{\mathcal{I}}^{(2q')}_{Q'+1} \cup \tilde{\tilde{\mathcal{I}}}^{(2q')}_{Q'+1}$. By implication, $|\mathrm{Sup}(\tilde{\tilde{\mathcal{I}}}^{(2q')}_{Q'+1})| \leq 2n/2 \leq 2nq/(q+1)$ and similarly $|\mathrm{Sup}(\tilde{\mathcal{I}}^{(2q')}_{Q'+1})| \leq 2nq/(q+1)$. We conclude the construction by modifying the set $\mathcal{I}^{(2q')}_{\alpha}$ for the considered q': we redefine $\mathcal{I}^{(2q')}_{Q'+1} \equiv \tilde{\mathcal{I}}^{(2q')}_{Q'+1}$, and introduce one extra interaction set $\mathcal{I}^{(2q')}_{Q'+2} \equiv \tilde{\tilde{\mathcal{I}}}^{(2q')}_{Q'+1}$.

The proof can now be finalized. Performing the above procedure for all q' where a violation was present, and completing the $\{\mathcal{I}^{(2q')}_{\alpha}\}$ without such violations with $\mathcal{I}^{(2q')}_{\alpha=Q'+2} = \emptyset$, we arrive at the splitting

$$\mathcal{I} = \bigcup_{q'=1}^{q/2} \bigcup_{\alpha=1}^{Q} \mathcal{I}^{(2q')}_{\alpha}, \tag{4.67}$$

where $Q = Q'+2 = q(q-1)(k-1)^2 + q(k-1) + 2$. Interaction sets $\mathcal{I}^{(2q')}_{\alpha}$ are diffuse (satisfying all three conditions of Definition 4.5.1) with respect to $\mathcal{I}$ for all q' and α. The construction of $\mathcal{I}^{(2q')}_{\alpha}$ is efficient, because each step can be implemented in time poly(n). □

4.C. Majorana matchings from diffuse interaction sets

Lemma (Repetition of Lemma 4.6.2)**.** *Let a strictly q'-local $\mathcal{I}'$ be diffuse w.r.t. q-local k-sparse $\mathcal{I}$ on $[2n]$, such that $n > (q^2-1)k$. One can efficiently construct a matching M of $[2n]$ that is consistent with $\mathcal{I}'$ and inconsistent with all interactions $I \in \mathcal{I}\backslash\mathcal{I}'$ such that (1) $|I| \geq q'$ or (2) $I \not\subset \mathrm{Sup}(\mathcal{I}')$.*

Proof. We first note that for $I \in \mathcal{I}\backslash\mathcal{I}'$ the condition $|I| \geq q'$ implies $I \not\subset \mathrm{Sup}(\mathcal{I}')$. Indeed, there are two possible options for $I \in \mathcal{I}\backslash\mathcal{I}'$ such that $I \subset \mathrm{Sup}(\mathcal{I}')$. The first option is that I

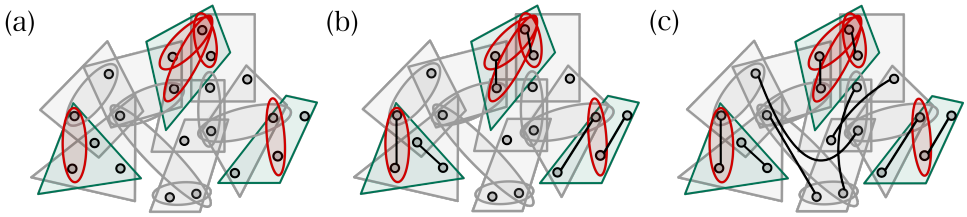

Figure 4.3: Example of the construction from the proof of Lemma 4.6.2. (a) The q-local set of interactions $\mathcal{I}$ ($q = 4$). Highlighted in green is diffuse and strictly q'-local $\mathcal{I}'$ ($q' = 4$), in red are the interactions in $\mathrm{Sup}(\mathcal{I}')$ with weight less than q', in grey are the rest of interactions in $\mathcal{I}$. The goal is to create a matching M consistent with green-colored terms, inconsistent with grey-colored terms, and with no guaranteed relation to the red-colored terms. (b) Matching M' on $\mathrm{Sup}(\mathcal{I}')$, consistent with $\mathcal{I}'$ by construction. Ensuring inconsistency with all red-colored terms is in general impossible. For example, consider the three overlapping red-colored terms at the top center. (c) Completing $M = M' \cup M''$ with a matching M'' on $[2n]\backslash\mathrm{Sup}(\mathcal{I}')$, ensuring inconsistency with all grey-colored terms. For this, the vertices are matched only if they belong to different interactions.

is a strict subset of a single interaction from $\mathcal{I}'$. However, this is not possible given $|I| \geq q'$, because $\mathcal{I}'$ is q'-local. The second option is for I to share Majorana modes with two or more interactions in $\mathcal{I}'$. This is ruled out because $\mathcal{I}'$ is diffuse with respect to $\mathcal{I}$ (cf. Condition 2 in Definition 4.5.1). The above implies that it is sufficient to construct the matching M that is consistent with $\mathcal{I}'$ and inconsistent with $\{I \in \mathcal{I}\backslash\mathcal{I}' | I \not\subset \mathrm{Sup}(\mathcal{I}')\}$.

We construct M in two steps. First we construct a matching M' of $\mathrm{Sup}(\mathcal{I}')$ (note $|\mathrm{Sup}(\mathcal{I}')|$ is always even). Next, we construct a matching M'' of the remaining Majorana modes $[2n]\backslash\mathrm{Sup}(\mathcal{I}')$. The desired matching of $[2n]$ is the union $M = M' \cup M''$.

To construct M', we match vertices of each $I \in \mathcal{I}'$ in an arbitrary way: for every such $I = \{i_1, .. i_{q'}\}$, $\{i_{2l-1}, i_{2l}\} \in M'$ for $l \in [1, .. q'/2]$. This matching is always possible, since $\mathcal{I}'$ is diffuse and thus different interactions from $\mathcal{I}'$ do not overlap. Thus constructed M' (and therefore also $M = M' \cup M''$) is explicitly consistent with all $I \in \mathcal{I}'$.

To construct a matching M'' of $[2n]\backslash\mathrm{Sup}(\mathcal{I}')$, we aim to ensure that no $(m_1, m_2) \in M''$ is a subset of any interaction in $\mathcal{I}$. For this, consider a 'permitted edge' graph $\mathcal{P}$ with vertices $[2n]\backslash\mathrm{Sup}(\mathcal{I}')$, and edges inserted between every pair (i_1, i_2) *unless* they belong to the same interaction in $\mathcal{I}$. We aim to construct M'' as a perfect matching of $\mathcal{P}$. Note that since $\mathcal{I}$ is q-local and k-sparse, the graph $\mathcal{P}$ has degree bounded from below as $|[2n]\backslash\mathrm{Sup}(\mathcal{I}')| - (q-1)k$. At the same time, since $\mathcal{I}'$ is diffuse, we're guaranteed by Condition 3 in Definition 4.5.1 that $|[2n]\backslash\mathrm{Sup}(\mathcal{I}')| \geq \frac{2n}{q+1}$. Therefore, since $n > (q^2-1)k$ by assumption, the degree of the vertices in $\mathcal{P}$ is lower bounded as $|[2n]\backslash\mathrm{Sup}(\mathcal{I}')| - (q-1)k \geq |[2n]\backslash\mathrm{Sup}(\mathcal{I}')|/2 + \frac{n}{q+1} - (q-1)k > |[2n]\backslash\mathrm{Sup}(\mathcal{I}')|/2$. Given this lower bound, we apply Dirac's theorem [30], which yields an efficiently constructable Hamiltonian cycle in the graph $\mathcal{P}$. Matching M'' is then obtained by pairing the sequential vertices in this cycle, making it a perfect matching of $\mathcal{P}$. By definition of

$\mathcal{P}$, M'' is guaranteed to contain at least one outgoing edge from every interaction in $\{I \in \mathcal{I}\backslash\mathcal{I}' | I \not\subset \mathrm{Sup}(\mathcal{I}')\}$. This makes $M = M' \cup M''$ inconsistent with $\{I \in \mathcal{I}\backslash\mathcal{I}' | I \not\subset \mathrm{Sup}(\mathcal{I}')\}$, as desired.

□

Lemma 4.5.3, which is used in the proof of Theorem 4.2.5, is a special case of Lemma 4.6.2. To obtain Lemma 4.5.3, one sets $q' = q$ and considers strictly q-local $\mathcal{I}$ instead of simply q-local. In this case all terms in $\mathcal{I}\backslash\mathcal{I}'$ satisfy the first condition of the Lemma, and therefore the constructed M is inconsistent with the entirety of $\mathcal{I}\backslash\mathcal{I}'$.

4.D. Matchings and Gaussian states

Lemma (Repetition of Lemma 4.6.3)**.** *Let $H = \sum_{I\in\mathcal{I}} J_I C_I$ on $[2n']$ be q-local and $\mathcal{I}'$ be a diffuse subset of $\mathcal{I}$. Consider a matching M of $[2n']$. If M is consistent with $\mathcal{I}'$ and inconsistent with $\mathcal{I}\backslash\mathcal{I}'$, one can efficiently construct a Gaussian state $\rho_{\mathcal{I}'}$ with the property:*

$$\mathrm{Tr}(H\rho(\mathcal{I}')) = \sum_{I\in\mathcal{I}'} |J_I|. \tag{4.68}$$

Proof. For the given matching M, consider its associated Gaussian state pure $\rho(M, \vec{\lambda})$ of the form:

$$\rho(M,\vec{\lambda}) = \frac{1}{2^n}\Pi_{\{m_1,m_2\}\in M}(\mathbb{1} + i\lambda_{(m_1,m_2)} c_{m_1} c_{m_2}). \tag{4.69}$$

Lemma 4.4.2 implies that the contribution to $\mathrm{Tr}(H\rho(M,\vec{\lambda}))$ from inconsistent interactions $\mathcal{I}\backslash\mathcal{I}'$ vanishes and contributions from $\mathcal{I}'$ yield:

$$\mathrm{Tr}(H\rho(M,\vec{\lambda})) = \sum_{I\in\mathcal{I}'} J_I \,\mathrm{sign}(\pi) \prod_{l\in\{1,..,|I|/2\}} \lambda_{(i_{\pi(2l-1)}, i_{\pi(2l)})}. \tag{4.70}$$

The proof is completed by choosing an appropriate value for $\vec{\lambda}$. Since $\mathcal{I}'$ is diffuse, by Condition 1 of Definition 4.5.1, distinct interactions from $\mathcal{I}'$ do not share Majorana fermions. This means that the values $\lambda_{(m_1,m_2)}$ for different I in equation (4.70) can be chosen independently. In particular, by picking appropriate $\lambda_{(m_1,m_2)} = \pm 1$, one can eliminate the sign of J_I sign(π) and achieve a contribution $|J_I|$ for each $I \in \mathcal{I}'$. Note that this procedure can be done efficiently, as it is simply a matter of choosing at most n ± 1 values by checking the sign of most $|\mathcal{I}'|$ terms. Denoting the thus chosen $\rho(M,\vec{\lambda})$ as $\rho(\mathcal{I}')$, this yields equation (4.68). □

A special case of Lemma 4.6.3 is Lemma 4.5.4 used in the proof of Theorem 4.2.5.

4.E. Moment bound for dense SYK-q

In this appendix, we establish the moment bound $\mathbb{E}\big(A(1)^{2n}\big) \le \big(C\sqrt{n}\big)^{2n}$, where $A(1)$ is defined as (in equation (4.56)):

$$A(1) = C\frac{1}{\sqrt{n}}\binom{n}{q-1}^{-3/2} \sum_{j,k,l=1}^{n} \sum_{S,S',S''} J_{S,j} J_{S',k} J_{S'',l}\, f(S,S',S'',j,k,l). \tag{4.71}$$

The function f in this expression is defined as (in equation (4.53)):

$$f(S,S',S'',j,k,l) = \begin{cases} 1, & \text{if} \quad (|S'' \cap S| \text{ is odd}) \\ & \quad \wedge\, (|S' \cap (S'' \triangle S)| + \delta_{k,l} \text{ is odd}) \\ & \quad\quad \wedge\, (|S| = |S'| = |S''| = q-1), \\ 0, & \text{otherwise.} \end{cases} \tag{4.72}$$

We classify the terms in the sum in equation (4.71) into five classes whose total contributions to the sum are denoted by D_0, D_1, D_2, D_3 and D_4. D_0 comprises of all terms for which the three J's are distinct. We shall therefore call the D_0 contribution the *diagonal-free* contribution. D_1 comprises of all terms for which the three J's are equal. D_2, D_3 and D_4 comprise of all terms for which exactly two out of three J's are equal.

Taking f into account, and thereby the terms that actually appear in $A(1)$, we conclude that the terms appearing in each class D_0, D_1, D_2, D_3 and D_4 correspond to the index sets given in Table 4.1. An illustration of *examples* of the index sets (S,j), (S',k) and (S'',l) associated with these different classes of contributions to $A(1)$ is given in Figure 4.4.

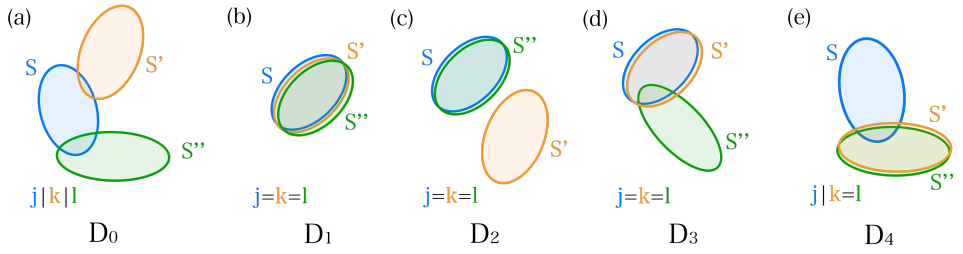

Figure 4.4: Illustration of *examples* of index sets (S,j), (S',k) and (S'',l) (corresponding to non-zero values of f in equation (4.72)) associated with the different classes of contributions to $A(1)$ in equation (4.71). The D_0 contribution in (a) is the diagonal-free contribution (i.e., (S,j), (S',k) and (S'',l) are unequal). The D_1, D_2, D_3 and D_4 contributions in resp. (b), (c), (d) and (e) are the diagonal contributions (i.e., at least two of (S,j), (S',k) and (S'',l) are equal).

To upper bound the $(2n)$th moment of $A(1)_{\min}$, we upper bound the rth moments (for even $r \leq 16 \cdot 2n$) of D_0, D_1, D_2, D_3, D_4 separately. In particular, if $\mathbb{E}\big((D_i)^r\big) \leq \big(C\sqrt{n}\big)^r$ for $i = 0,1,...,4$ and all even $r \leq 16 \cdot 2n$, then $\mathbb{E}\big(A(1)^{2n}\big) \leq \big(C\sqrt{n}\big)^{2n}$. Note that through the multinomial expansion and successive application of Cauchy-Schwarz inequality these former bounds indeed give an upper bound on the $(2n)$th moment of $A(1)$:

class	associated index sets of terms	associated index sets of terms in $A(1)$
D_0	$(S,j) \neq (S',k) \neq (S'',l) \neq (S,j)$	$(S,j) \neq (S',k) \neq (S'',l) \neq (S,j)$
D_1	$(S,j) = (S',k) = (S'',l)$	$(S,j) = (S',k) = (S'',l)$
D_2	$(S,j) = (S'',l) \neq (S',k)$	$S = S'' \neq S',\ j = l = k$
D_3	$(S,j) = (S',k) \neq (S'',l)$	$S = S' \neq S'',\ j = l = k$
D_4	$(S',k) = (S'',l) \neq (S,j)$	$(S',k) = (S'',l) \neq (S,j)$

Table 4.1: The index sets associated with each class of terms, and the index sets associated with each class of terms that appear in the expression for $A(1)$ (i.e., taking f into account).

$$\begin{aligned}\mathbb{E}\big(A(1)^{2n}\big) &= \mathbb{E}\Big((D_0+D_1+D_2+D_3+D_4)^{2n}\Big) = \sum_{k_0+\ldots+k_4=2n} \frac{2n!}{k_0!\ldots k_4!}\mathbb{E}\Big(D_0^{k_0}D_1^{k_1}\ldots D_4^{k_4}\Big)\\ &\leq \sum_{k_0+\ldots+k_4=2n} C^{2n}\,\mathbb{E}\big(D_0^{2k_0}\big)^{1/2}\mathbb{E}\big(D_1^{4k_1}\big)^{1/4}\mathbb{E}\big(D_2^{8k_2}\big)^{1/8}\mathbb{E}\big(D_3^{16k_3}\big)^{1/16}\mathbb{E}\big(D_4^{16k_4}\big)^{1/16}\\ &\leq \sum_{k_0+\ldots+k_4=2n} C^{2n}\left(C\sqrt{n}\right)^{k_0+\ldots+k_4} \leq \left(C\sqrt{n}\right)^{2n}, \end{aligned} \tag{4.73}$$

where we have used that the multinomial coefficient can be upper bounded by C^{2n} and that the number of 5-tuples of non-negative integers whose sum equals $2n$ is upper bounded by Cn^4 (which is smaller than C^{2n} for some constant C). Although clearly the rth moments of e.g. D_0 have to only be bounded for even $r \leq 2\cdot 2n$, we bound – for the sake of clarity – the rth moments for even $r \leq 16\cdot 2n$ for *all* D_i's. We first deal with the case of D_0, since the fact that this contribution is diagonal-free allows one to employ a decoupling technique. Afterwards, we will consider the D_1, D_2, D_3 and D_4 contributions.

First, we state the following lemma, which will be useful throughout this appendix.

Lemma 4.E.1. *Let P and P' be two polynomials of centered Gaussian random variables (i.e., the monomials are formed by products of elements from a sequence of independent centered Gaussian random variables, and each variable is allowed to appear in a monomial multiple times) with non-negative coefficients. Then, for any even r, $\mathbb{E}\big(|P+P'|^r\big) \geq \mathbb{E}\big(|P|^r\big)$.*

Proof. We have $\mathbb{E}\big(|P+P'|^r\big) = \mathbb{E}\big(|P|^r\big) + \sum_{k=1}^{r}\binom{r}{k}\mathbb{E}\big(P^{r-k}(P')^k\big)$, and $\mathbb{E}\big(P^{r-k}(P')^k\big)$ is non-negative (for any integers r,k) since P and P' have non-negative coefficients and all moments of centered Gaussian random variables are non-negative. □

4.E.1. Upper bound for moments of D_0 (diagonal-free contribution)

We start by noting that the function f takes on values 0 or 1, dependent on the index sets S,S',S'',j,k,l labeling the Majorana operators. We consider replacing f in each term of D_0 (equation (4.71)) with $\delta_{a,b}\delta_{c,d}$, where either

$$a \in (S'\cup k),\ b \neq c \in (S''\cup l),\ d \in (S\cup j), \qquad \text{(option 1)}$$

or

$$a \in (S' \cup k),\ b \neq c \in (S \cup j),\ d \in (S'' \cup l). \qquad \text{(option 2)}$$

We denote this modified sum as $D_{0,\delta\delta}$. By inspection, the index sets for which f is non-zero all correspond to a non-zero contribution for $\delta_{a,b}\delta_{c,d}$. Note that those index sets for which $\delta_{a,b}\delta_{c,d}$ is non-zero also include index sets for which f is zero. Hence, the terms associated with non-zero $\delta_{a,b}\delta_{c,d}$ (for the two options listed above) are a superset of the terms that correspond to non-zero values of f. Therefore, by Lemma 4.E.1, the upper bounds on even moments of D_0 can be obtained by upper bounding the even moments of $D_{0,\delta\delta}$.

We will denote the part of the sum $D_{0,\delta\delta}$ corresponding to option 1 as $D_{0,\min}$:

$$D_{0,\min} := C\frac{1}{\sqrt{n}}\binom{n}{q-1}^{-3/2} \sum_{\substack{j,k,l,\\ S,S',S'',\\ \text{s.t. } |(S'\cup k)\cap(S''\cup l)|\geq 1\\ \text{and } |(S''\cup l)\cap(S\cup j)|\geq 1}} J_{S,j}J_{S',k}J_{S'',l}, \tag{4.74}$$

where the sum is over indices such that $(S,j) \neq (S',k) \neq (S'',l) \neq (S,j)$ (by definition of D_0) and such that $(S' \cup k) \cap (S'' \cup l)$ and $(S'' \cup l) \cap (S \cup j)$ differ by at least one element. Any bound for all even moments of $D_{0,\min}$ also holds for $D_{0,\delta\delta} - D_{0,\min}$ which corresponds to option 2, due to the symmetry $(S,j) \leftrightarrow (S'',l)$ between the two options. An upper bound on all even moments of $D_{0,\delta\delta}$ (and, by implication, D_0) then follows from binomial expansion and application of the Cauchy-Schwarz inequality, similarly to equation (4.73). Thus it only remains to prove $\mathbb{E}\big(|D_{0,\min}|^r\big) < (C\sqrt{n})^r$ for all even r.

To upper bound the even moments of $D_{0,\min}$, we are going to employ a decoupling technique. To that end, we will study the even moments of a related decoupled quantity. This decoupled quantity is defined as $D_{0,\min}$ but with the standard Gaussian random variables $J_{S,j}$, $J_{S',k}$ and $J_{S'',l}$ (selected from a single sequence of standard Gaussian random variables) being replaced by their *decoupled* versions $J^{(1)}_{S,j}$, $J^{(2)}_{S',k}$ and $J^{(3)}_{S'',l}$ (selected from *three* independent sequences of standard Gaussian random variables). The related decoupled quantity is given by (where the sum is again over indices $(S,j) \neq (S',k) \neq (S'',l) \neq (S,j)$ and again such that $(S' \cup k) \cap (S'' \cup l)$ and $(S'' \cup l) \cap (S \cup j)$ differ by at least one element):

$$C\frac{1}{n^{3q/2-1}} \sum_{\substack{j,k,l,\\ S,S',S'',\\ \text{s.t. } |(S'\cup k)\cap(S''\cup l)|\geq 1\\ \text{and } |(S''\cup l)\cap(S\cup j)|\geq 1}} J^{(1)}_{S,j}J^{(2)}_{S',k}J^{(3)}_{S'',l}, \tag{4.75}$$

where we have additionally used that $(k/l)^l \leq \binom{k}{l} \leq (e\,k/l)^l$.

To upper bound the even moments of this decoupled quantity, we will make use of Lemma 4.E.2 below from [31]. The even moments of this decoupled sum are upper bounded by upper bounding the even moments of a decoupled sum whose terms are a superset of the terms in the sum in equation (4.75). Through Lemma 4.E.1, the

even moments of the latter sum are larger than those of the former sum. For each $J^{(i)}_{S,j}$, we introduce $q!-1$ additional independent standard Gaussian random variables associated with the $q!$ permutations of the indices in the subsets of size q. Furthermore, we introduce additional independent standard Gaussian random variables for which some or all of the q indices that label them are equal. We consider a sum over lists of $3q$ indices (which label the independent standard Gaussian random variables) $i^{(1)}_1,\ldots,i^{(1)}_q$, $i^{(2)}_1,\ldots,i^{(2)}_q$ and $i^{(3)}_1,\ldots,i^{(3)}_q$ (with each index in $[2n]$), instead of the sum over subsets of $[2n]$ in equation (4.75). Note that the sum over lists, by definition, can contain terms for which two (or three) of the Gaussian random variables have equal index sets.

The index lists $i^{(1)}_1,\ldots,i^{(1)}_q$ and $i^{(2)}_1,\ldots,i^{(2)}_q$ that are summed over each have *any* one index (denoted by resp. x and y) that is equal to an index in $i^{(3)}_1,\ldots,i^{(3)}_q$. If we additionally sum over all 'positions' of the x and y indices (where p_1, p_2, p_3 and $p_4 \in \{1,\ldots,q\}$ label these positions), we obtain the sum (see equation (4.76) below) whose terms are a superset of those in the sum in equation (4.75). Note that this sum in equation (4.76) contains all the contributions from a sum over lists of indices, and contains some terms multiple times that would occur only once in a sum over lists of indices: For example, in the hypothetical case $q=2$, one could have a contribution $J^{(1)}_{3,3}J^{(2)}_{6,7}J^{(3)}_{3,7}$ that would appear once in the sum over lists of indices but appears twice in the sum in equation (4.76) (once for $p_1=1$ and once for $p_1=2$). Through Lemma 4.E.1, the even moments of the sum in equation (4.76) will therefore be larger than those of the sum over lists of indices (and therefore larger than those of the sum in equation (4.75)), and it will thus suffice to upper bound the even moments of the sum in equation (4.76).

$$D^{\text{decoupled}}_{0,\min} := C\,\frac{1}{n^{3q/2-1}} \sum_{\substack{p_1,\ldots,p_4=1 \\ \text{s.t. } p_3\neq p_4}}^{q} \Bigg[\sum_{\substack{i^{(1)}_1,\ldots,i^{(1)}_{p_1-1},x, \\ i^{(1)}_{p_1+1},\ldots,i^{(1)}_q=1}}^{2n} \sum_{\substack{i^{(2)}_1,\ldots,i^{(2)}_{p_2-1},y, \\ i^{(2)}_{p_2+1},\ldots,i^{(2)}_q=1}}^{2n} \sum_{\substack{i^{(3)}_1,\ldots,i^{(3)}_{p_3-1},x,i^{(3)}_{p_3+1},\ldots, \\ i^{(3)}_{p_4-1},y,i^{(3)}_{p_4+1},\ldots,i^{(3)}_q=1}}^{2n} \left(J^{(1)}_{i^{(1)}_1,\ldots,i^{(1)}_{p_1-1},x,\, i^{(1)}_{p_1+1},\ldots,i^{(1)}_q} J^{(2)}_{i^{(2)}_1,\ldots,i^{(2)}_{p_2-1},y,\, i^{(2)}_{p_2+1},\ldots,i^{(2)}_q} J^{(3)}_{i^{(3)}_1,\ldots,i^{(3)}_{p_3-1},x,i^{(3)}_{p_3+1},\ldots,\, i^{(3)}_{p_4-1},y,i^{(3)}_{p_4+1},\ldots,i^{(3)}_q} \right) \Bigg]. \tag{4.76}$$

The free indices ($q-1$ indices of $i^{(1)}_1,\ldots,i^{(1)}_q$ and $i^{(2)}_1,\ldots,i^{(2)}_q$, and $q-2$ indices of $i^{(3)}_1,\ldots,i^{(3)}_q$) can be summed over to obtain new independent standard Gaussian random variables denoted by $K^{(1)}_{x,p_1}$, $K^{(2)}_{y,p_2}$ and $K^{(3)}_{x,p_3;y,p_4}$:

$$K^{(1)}_{x,p_1} := 1/\sqrt{(2n)^{q-1}} \sum_{i^{(1)}_1,\ldots,i^{(1)}_{p_1-1},i^{(1)}_{p_1+1},\ldots,i^{(1)}_q=1}^{2n} J^{(1)}_{i^{(1)}_1,\ldots,i^{(1)}_{p_1-1},x,i^{(1)}_{p_1+1},\ldots,i^{(1)}_q}, \tag{4.77a}$$

$$K^{(2)}_{y,p_2} := 1/\sqrt{(2n)^{q-1}} \sum_{i^{(2)}_1,\ldots,i^{(2)}_{p_2-1},i^{(2)}_{p_2+1},\ldots,i^{(2)}_q=1}^{2n} J^{(2)}_{i^{(2)}_1,\ldots,i^{(2)}_{p_2-1},y,i^{(2)}_{p_2+1},\ldots,i^{(2)}_q}, \tag{4.77b}$$

$$K^{(3)}_{x,p_3;y,p_4} := 1/\sqrt{(2n)^{q-2}} \sum_{\substack{i^{(3)}_1,\ldots,i^{(3)}_{p_3-1},i^{(3)}_{p_3+1},\ldots,\\ i^{(3)}_{p_4-1},i^{(3)}_{p_4+1},\ldots,i^{(3)}_q=1}}^{2n} J^{(3)}_{i^{(3)}_1,\ldots,i^{(3)}_{p_3-1},x,i^{(3)}_{p_3+1},\ldots,i^{(3)}_{p_4-1},y,i^{(3)}_{p_4+1},\ldots,i^{(3)}_q}, \tag{4.77c}$$

where we have used that the normalized sum $1/\sqrt{m}\sum_{i=1}^m J_i$ of a sequence of standard Gaussian random variables $J_1,\ldots,J_m$ is again a standard Gaussian random variable. We now obtain the following expression for $D^{\text{decoupled}}_{0,\min}$:

$$D^{\text{decoupled}}_{0,\min} := C \sum_{\substack{p_1,\ldots,p_4=1\\ \text{s.t. } p_3\neq p_4}}^{q} \Big[\frac{1}{n}\sum_{x,y=1}^{2n} K^{(1)}_{x,p_1} K^{(2)}_{y,p_2} K^{(3)}_{x,p_3;y,p_4}\Big]. \tag{4.78}$$

The sum over all free indices gives an extra total factor of $n^{3q/2-2}$, which partially cancels against $n^{3q/2-1}$ in equation (4.76). Importantly, we note that now the random variables $K^{(1)}_{x,p_1}$ and $K^{(1)}_{x',p_1}$ are independent for $x \neq x'$ (and equivalently for $K^{(2)}_{y,p_2}$ and $K^{(3)}_{x,p_3;y,p_4}$). We will apply Lemma 4.E.2 from [31] separately to each contribution to $D^{\text{decoupled}}_{0,\min}$ in equation (4.78) (with a contribution corresponding to one combination of p_i's).

Lemma 4.E.2 (Theorem 1 in [31])**.** *Let $Y \in \mathbb{R}^{N\times\ldots\times N}$ be a d-dimensional matrix and define:*

$$F\Big(\{K^{(j)}_1,\ldots,K^{(j)}_N\}_{j=1}^d\Big) := \sum_{i_1,\ldots,i_d=1}^{N} Y_{i_1,\ldots,i_d}\prod_{j=1}^{d} K^{(j)}_{i_j}, \tag{4.79}$$

where $\{K^{(j)}_1,\ldots,K^{(j)}_N\}_{j=1}^d$ are d independent sequences of N standard Gaussian random variables. Then for any integer $k \geq 2$:

$$\mathbb{E}\Big(\big|F\big(\{K^{(j)}_1,\ldots,K^{(j)}_N\}_{j=1}^d\big)\big|^k\Big) \leq \big(C\sum_{\mathcal{P}} k^{|\mathcal{P}|/2}\|Y\|_{\mathcal{P}}\big)^k \leq \big(C\max_{\mathcal{P}}\big[k^{|\mathcal{P}|/2}\|Y\|_{\mathcal{P}}\big]\big)^k, \tag{4.80}$$

where $\mathcal{P}$ are partitions of $[d]$ into non-empty parts $(P_1,\ldots,P_s)$. The second inequality holds because the number of partitions of $[d]$ into non-empty parts is constant in n (since d is constant in n). The quantity $\|Y\|_{\mathcal{P}}$ is defined as:

$$\|Y\|_{\mathcal{P}} = \|Y\|_{(P_1,\ldots,P_s)} := \max\Big\{\sum_{i_1,\ldots,i_d=1}^{N} Y_{i_1,\ldots,i_d}\, x^{(1)}_{i_{P_1}}\ldots x^{(s)}_{i_{P_s}} : \sum_{i_{P_1}}\big(x^{(1)}_{i_{P_1}}\big)^2 \leq 1,\ldots,\sum_{i_{P_k}}\big(x^{(k)}_{i_{P_k}}\big)^2 \leq 1\Big\}, \tag{4.81}$$

with each $x \in \mathbb{R}$.

Remark 4.E.3. *If F in equation (4.79) is diagonal-free (i.e., $Y_{i_1,\ldots,i_d} = 0$ if $i_j = i_k$ for any $j \neq k$) then the moments of the 'decoupled' F in equation (4.79) are (up to constants only depending on d) an upper bound for the moments of its 'coupled' counterpart $F'\big(K_1,\ldots,K_N\big) := \sum_{i_1,\ldots,i_d=1}^N Y_{i_1,\ldots,i_d}\prod_{j=1}^d K_{i_j}$ (i.e., where the random variables are all taken from the same sequence of N standard Gaussian random variables):*

$$\mathbb{E}\big((F')^r\big) \leq C\mathbb{E}\big((F)^r\big).$$

See e.g. Theorem 2.1 in [32].

The fact that this decoupling inequality only holds for diagonal-free polynomials is exactly the reason for differentiating between the diagonal-free contribution D_0 and the diagonal contributions D_1, D_2, D_3, D_4 to $A(1)$.

For $\sum_{x,y=1}^{2n} K^{(1)}_{x,p_1} K^{(2)}_{y,p_2} K^{(3)}_{x,p_3;y,p_4}$ in equation (4.78), we see that $d=3$ and hence the possible partitions $\mathcal{P}$ are $\{1,2,3\}$, $\{1\}\{2,3\}$, $\{2\}\{1,3\}$, $\{1,2\}\{3\}$, $\{1\}\{2\}\{3\}$. The associated $\|Y\|_{\mathcal{P}}$ values can be (straightforwardly) calculated and are given in Table 4.2. Using Table 4.2 and Lemma 4.E.2, we find the following upper bound on $\mathbb{E}\Big(\Big[C/n \sum_{x,y=1}^{2n} K^{(1)}_{x,p_1} K^{(2)}_{y,p_2} K^{(3)}_{x,p_3;y,p_4}\Big]^r\Big)$ (for all even r):

$$\Big(C/n \max\big(\sqrt{r}n, r\sqrt{n}, r\sqrt{n}, r, r^{3/2}\big)\Big)^r \leq \big(C\sqrt{n}\big)^r. \tag{4.82}$$

$\mathcal{P}$	$\lvert\mathcal{P}\rvert$	$\|Y\|_{\mathcal{P}}$
$\{1,2,3\}$	1	$2n$
$\{1\}\{2,3\}$	2	$\sqrt{2n}$
$\{2\}\{1,3\}$	2	$\sqrt{2n}$
$\{1,2\}\{3\}$	2	1
$\{1\}\{2\}\{3\}$	3	1

Table 4.2: The different partitions $\mathcal{P}$ of $[3]$ into non-empty parts, with the associated number of parts $|\mathcal{P}|$, and the associated $\|Y\|_{\mathcal{P}}$ for $\sum_{x,y=1}^{2n} K^{(1)}_{x,p_1} K^{(2)}_{y,p_2} K^{(3)}_{x,p_3;y,p_4}$ in equation (4.78). $\|Y\|_{\mathcal{P}}$ for the first four partitions can be straightforwardly evaluated by applying equation (4.81) to equation (4.75), and the fifth $\|Y\|_{\mathcal{P}}$ can be evaluated by additional application of the Cauchy-Schwarz inequality.

Note that $D^{\text{decoupled}}_{0,\min}$ in equation (4.78) consists of q^4 (with $q = O(1)$) contributions, each corresponding to a given combination of p_i's. We can again use the multinomial expansion and successive application of the Cauchy-Schwarz inequality (together with the fact that the multinomial coefficients can be upper bounded by C^r and that the number of q^4-tuples of non-negative integers whose sum equals r is upper bounded by C^r for some constant C) to conclude that the upper bounds of $(C\sqrt{n})^r$ for rth moments (for all even r) of these contributions imply an upper bound of $(C\sqrt{n})^r$ for rth moments (for all even r) of $D^{\text{decoupled}}_{0,\min}$.

We now employ the decoupling inequality from the above remark to obtain

$$\mathbb{E}\big(|D_{0,\min}|^r\big) \leq C\mathbb{E}\big(\big|D^{\text{decoupled}}_{0,\min}\big|^r\big) \leq \big(C\sqrt{n}\big)^r.$$

From the arguments given previously, this implies the desired bound $\mathbb{E}\big(|D_0|^r\big) \leq \big(C\sqrt{n}\big)^r$ for all even r, in particular for $r \leq 16 \cdot 2n$.

4.E.2. Upper bound for moments of D_1, D_2, D_3 and D_4

In the previous section we used a decoupling inequality to upper bound the rth moments (for even $r \leq 16 \cdot 2n$) of D_0. These decoupling inequalities hold for (Gaussian) polynomials for which each Gaussian monomial is a product of *distinct* Gaussian random variables, i.e., *diagonal-free* polynomials. This holds indeed – by definition – for the D_0 contribution

to $A(1)$, but not for contributions D_1, D_2, D_3 and D_4. For that reason, we cannot make use of the same decoupling inequality for the D_1, D_2, D_3 and D_4 contributions. Therefore, we have to resort to other methods to bound their rth moments (for even $r \leq 16 \cdot 2n$).

- The D_1 contribution can be written as:

$$D_1 := C\frac{1}{\sqrt{n}}\binom{n}{q-1}^{-3/2}\sum_{j,S}\left(J_{S,j}\right)^3. \tag{4.83}$$

The rth moment (with r even) of D_1 can be upper bounded as follows:

$$\begin{aligned}\mathbb{E}\left(|D_1|^r\right) &\leq C^r\left[\frac{1}{\sqrt{n}}\binom{n}{q-1}^{-3/2}\right]^r\mathbb{E}\left[\left(\sum_{S,j}\left(J_{S,j}\right)^3\right)^r\right]\\ &\leq C^r\left[\frac{1}{\sqrt{n}}\binom{n}{q-1}^{-3/2}\right]^r\left(\binom{n}{q-1}nr^{3/2}\right)^r\\ &\leq C^r n^{(5/2-q/2)r},\end{aligned} \tag{4.84}$$

where we have used that $\mathbb{E}\left[\left(\sum_{i=1}^m K_i\right)^r\right] = \sum_{k_1+\ldots+k_m=r}\frac{r!}{k_1!\ldots k_m!}\mathbb{E}\left(\left(K_1\right)^{k_1}\right)\ldots\mathbb{E}\left(\left(K_m\right)^{k_m}\right)$ (for $K_1,\ldots,K_m$ independent random variables), the fact that (S,j) can take on $2n\binom{2n}{q-1}$ values and the fact that the pth moment of a standard Gaussian random variable is equal to $(p-1)!!$ $(\leq p^{p/2})$. For even r, we therefore conclude that $\mathbb{E}\left(|D_1|^r\right) \leq \left(C\sqrt{n}\right)^r$.

- The D_2 and D_3 contributions are equivalent and can be written as:

$$D_2 = D_3 := C\frac{1}{\sqrt{n}}\binom{n}{q-1}^{-3/2}\sum_{j,S,S' \text{ s.t. } S\neq S'}\left(J_{S,j}\right)^2 J_{S',j}. \tag{4.85}$$

The rth moment (with r even) can be written as follows:

$$\mathbb{E}\left(|D_2|^r\right), \mathbb{E}\left(|D_3|^r\right) = C^r\left[\frac{1}{\sqrt{n}}\binom{n}{q-1}^{-3/2}\right]^r\mathbb{E}\left(\left(\sum_{j,S,S' \text{ s.t. } S\neq S'}\left(J_{S,j}\right)^2 J_{S',j}\right)^r\right). \tag{4.86}$$

We define

$$g := \sum_{j,S,S' \text{ s.t. } S\neq S'}\left(J_{S,j}\right)^2 J_{S',j}, \tag{4.87}$$

for which $\mathbb{E}(g) = 0$. We note that g is a homogeneous polynomial in standard Gaussian random variables of degree 3. To upper bound the moments of g, and thereby the moments of D_2 and D_3, we use the following result from [33]. This result is an extension of Lemma 4.E.2 from [31] to the setting where diagonal terms are allowed to appear in the polynomial. The extension also includes inhomogeneous polynomials, although in the current setting we are considering only homogeneous polynomials.

Lemma 4.E.4 (Theorem 1.3 in [33]). *Let $\mathbf{K} := K_1 \ldots, K_N$ denote a sequence of N independent standard Gaussian random variables and $g : \mathbb{R}^N \to \mathbb{R}$ a polynomial of degree D. Then, for all $r \geq 2$:*

$$\mathbb{E}\Big(\Big[g(\mathbf{K}) - \mathbb{E}\big(g(\mathbf{K})\big)\Big]^r\Big) \leq C^r \Big[\sum_{1\leq d\leq 3} \sum_{\mathcal{P}([d])} r^{|\mathcal{P}|/2} \big\|\mathbb{E}\big(\mathbf{D}^d g(\mathbf{K})\big)\big\|_{\mathcal{P}}\Big]^r, \tag{4.88}$$

where $\mathcal{P}$ are partitions of $[d]$ into non-empty parts, and $\|Y\|_{\mathcal{P}}$ (with Y a d-way tensor) is defined in equation (4.81). *$\mathbf{D}^d g(\mathbf{K})$ denotes the dth derivative of $g(\mathbf{K})$, which corresponds to a d-way tensor with entries equal to $\big[\mathbf{D}^d g(\mathbf{K})\big]_{i_1,\ldots,i_d} = \frac{\partial}{\partial K_{i_1}} \ldots \frac{\partial}{\partial K_{i_d}} g(\mathbf{K})$. For $d = D$, $\mathbf{D}^d g(\mathbf{K})$ is constant.*

For g in equation (4.87), we have that $N = n\binom{n}{q-1}$, since the sequence of Gaussian random variables corresponds to $\{J_{S,j}\}$. To find an upper bound for the rth moment of g using equation (4.88), we first calculate $\mathbf{D}^d g$ for $d = 1,2,3$. Then, for each d, we upper bound $\big\|\mathbb{E}\big(\mathbf{D}^d g\big)\big\|_{\mathcal{P}}$ for all partitions $\mathcal{P}$ of $[d]$. We will show that for all d and associated partitions $\mathcal{P}([d])$, $\big\|\mathbb{E}\big(\mathbf{D}^d g\big)\big\|_{\mathcal{P}}$ can be upper bounded in such a way that $\mathbb{E}\Big(|D_2|^r\Big), \mathbb{E}\Big(|D_3|^r\Big) \leq \big(C\sqrt{n}\big)^r$ for all even $2 \leq r \leq 16 \cdot 2n$. Finally, the 0th moment also (trivially) satisfies this upper bound, hence it holds for all even $r \leq 16 \cdot 2n$.

The derivatives of g are equal to:

$$Dg = \Big(\sum_{S' : S' \neq S} J_{S',j}^2 + 2J_{S,j} \sum_{S' : S' \neq S} J_{S',j}\Big)_{(S,j)} \Longrightarrow \mathbb{E}\big(Dg\big) = \binom{n}{q-1}_{(S,j)}, \tag{4.89}$$

$$D^2 g = \begin{pmatrix} 2\sum_{S' : S' \neq S} J_{S',j}, & \text{if } (S,j) = (T,k) \\ 2(J_{T,j} + J_{S,j}), & \text{if } S \neq T \text{ and } j = k \\ 0, & \text{if } j \neq k \end{pmatrix}_{(S,j),(T,k)} \Longrightarrow \mathbb{E}\big(D^2 g\big) = (0)_{(S,j),(T,k)}, \tag{4.90}$$

$$D^3 g = \begin{pmatrix} 2, & \text{if } (S = T \neq U \text{ or } S \neq T = U \text{ or } S = U \neq T) \text{ and } j = k = l \\ 0, & \text{if } (S = T = U \text{ or } S \neq T \neq U) \text{ and } j = k = l \\ 0, & \text{if } j,k,l \text{ are not all equal} \end{pmatrix}_{(S,j),(T,k),(U,l)} \Longrightarrow \mathbb{E}\big(D^3 g\big) = D^3 g. \tag{4.91}$$

In Table 4.3, we give the values of $\big\|\mathbb{E}\big(\mathbf{D}^d g\big)\big\|_{\mathcal{P}}$ for all partitions $\mathcal{P}([d])$ for $d = 1,2,3$. $\big\|\mathbb{E}\big(\mathbf{D}^d g\big)\big\|_{\mathcal{P}}$ for $d = 1$ can be straightforwardly evaluated using equation (4.81) and for $d = 2$ can be trivially evaluated by using $\mathbb{E}\big(\mathbf{D}^2 g\big) = 0$. For $d = 3$, $\big\|\mathbb{E}\big(\mathbf{D}^d g\big)\big\|_{\mathcal{P}}$ can be upper bounded using equation (4.81), and the triangle and Cauchy-Schwarz inequalities (for illustration purposes, we provide an example of the derivation of this upper bound for $\mathcal{P} = \{1,2\}\{3\}$ below).

Combining the upper bounds for $\big\|\mathbb{E}\big(\mathbf{D}^d g\big)\big\|_{\mathcal{P}}$ in Table 4.3 with the factor of $r^{|\mathcal{P}|/2} (\leq C n^{|\mathcal{P}|/2})$ in equation (4.88) and the normalization factor in equation (4.86), we find – using $\mathbb{E}(g) = 0$ – that indeed $\mathbb{E}\Big(|D_2|^r\Big), \mathbb{E}\Big(|D_3|^r\Big) \leq \big(C\sqrt{n}\big)^r$ for all even $r \leq 16 \cdot 2n$.

$\mathcal{P}$	$\lvert\mathcal{P}\rvert$	$\left\lVert\mathbb{E}\left(\mathbf{D}^d g\right)\right\rVert_{\mathcal{P}}$
$\{1\}$	1	$\binom{n}{q-1}^{3/2}$
$\{1,2\}$	1	0
$\{1\}\{2\}$	2	0
$\{1,2,3\}$	1	$C\sqrt{n}\binom{n}{q-1}$
$\{1\}\{2,3\}$	2	$Cn\binom{n}{q-1}$
$\{2\}\{1,3\}$	2	$Cn\binom{n}{q-1}$
$\{1,2\}\{3\}$	2	$Cn\binom{n}{q-1}$
$\{1\}\{2\}\{3\}$	3	$C\binom{n}{q-1}$

Table 4.3: The different partitions $\mathcal{P}$ of [3] into non-empty parts, with the associated number of parts $\lvert\mathcal{P}\rvert$, and (the upper bounds for) the associated $\left\lVert\mathbb{E}\left(\mathbf{D}^d g\right)\right\rVert_{\mathcal{P}}$ for g in equation (4.87).

Example: For illustration purposes, we give an explicit evaluation of $\left\lVert\mathbb{E}\left(\mathbf{D}^d g\right)\right\rVert_{\mathcal{P}}$ for $\mathcal{P} = \{1,2\}\{3\}$ (the evaluations for other $\mathcal{P}$'s follow using similar methods). By definition (equation (4.81)), we have:

$$\left\lVert\mathbb{E}\left(\mathbf{D}^3 g\right)\right\rVert_{\{1,2\}\{3\}} = \sup\Big\{ \sum_{(S,j),(T,k),(U,l)} \mathbb{E}\left(\mathbf{D}^3 g\right)_{(S,j),(T,k),(U,l)} x_{(S,j),(T,k)}\, y_{(U,l)} : \sum_{(S,j),(T,k)} x^2_{(S,j),(T,k)} \le 1, \sum_{(U,l)} y^2_{(U,l)} \le 1 \Big\}. \quad (4.92)$$

Using the expression obtained for $\mathbb{E}\left(\mathbf{D}^3 g\right)_{(S,j),(T,k),(U,l)}$, we obtain:

$$\begin{aligned}
&\left\lVert\mathbb{E}\left(\mathbf{D}^3 g\right)\right\rVert_{\{1,2\}\{3\}} \\
&\quad = \sup\Big\{ \sum_{j,S,T} 2\, x_{(S,j),(T,j)}\, y_{(T,j)} + \sum_{j,S,T} 2\, x_{(S,j),(T,j)}\, y_{(S,j)} + \sum_{j,S,U} 2\, x_{(S,j),(S,j)}\, y_{(U,j)} \\
&\qquad\qquad : \sum_{(S,j),(T,k)} x^2_{(S,j),(T,k)} \le 1, \sum_{(U,l)} y^2_{(U,l)} \le 1 \Big\} \\
&\quad \le \sup\Big\{ 2\sum_S \Big| \sum_{j,T} x_{(S,j),(T,j)}\, y_{(T,j)} \Big| + 2\sum_T \Big| \sum_{j,S} x_{(S,j),(T,j)}\, y_{(S,j)} \Big| \\
&\qquad\qquad + 2\sum_j \Big| \sum_S x_{(S,j),(S,j)} \Big| \Big| \sum_U y_{(U,j)} \Big| : \sum_{(S,j),(T,k)} x^2_{(S,j),(T,k)} \le 1, \sum_{(U,l)} y^2_{(U,l)} \le 1 \Big\} \\
&\quad \le \sup\Big\{ 2\sum_S \|x\|\,\|y\| + 2\sum_T \|x\|\,\|y\| \\
&\qquad\qquad + 2\sum_j \|x\| \sqrt{\binom{n}{q-1}}\, \|y\| \sqrt{\binom{n}{q-1}} : \sum_{(S,j),(T,k)} x^2_{(S,j),(T,k)} \le 1, \sum_{(U,l)} y^2_{(U,l)} \le 1 \Big\} \\
&\quad \le Cn\binom{n}{q-1}, \qquad (4.93)
\end{aligned}$$

where we have used the triangle inequality in the first inequality, and the Cauchy-Schwarz inequality for the second inequality (and we note that e.g. $\sum_U y_{(U,j)}$ is simply equal to the inner product of $y_{(j)} := (y_{(U_1,j)}, y_{(U_2,j)}, \ldots)$ with the all-ones vector).

- The D_4 contribution can be written as:

$$D_4 := C\frac{1}{\sqrt{n}}\binom{n}{q-1}^{-3/2} \sum_{\substack{j,k,S,S' \\ \text{s.t. } 0<|S\cap S'|<q-1 \\ \text{is odd}}} J_{S,j}\big(J_{S',k}\big)^2. \tag{4.94}$$

We note that the main difference with the D_2 and D_3 contributions is that, for D_4, the sum is over the double index j,k (instead of over the single index j), and over a restricted sum over sets S,S' (instead of over a free sum over sets S,S'). To bound the moments of D_4, we will employ a similar method as for D_2 and D_3. The rth moment (with r even) can be upper bounded as follows (where we drop the '$|S\cap S'|$ is odd' constraint using Lemma 4.E.1 and denote the collection of subsets S' such that $0<|S\cap S'|<q-1$ by $\sigma(S)$):

$$\mathbb{E}\Big(|D_4|^r\Big) \leq C^r\left[\frac{1}{\sqrt{n}}\binom{n}{q-1}^{-3/2}\right]^r \mathbb{E}\Big(\Big(\sum_{\substack{j,k,S, \\ S'\in\sigma(S)}} J_{S,j}\big(J_{S',k}\big)^2\Big)^r\Big). \tag{4.95}$$

We note that $|\sigma(S)|$ can be upper bounded and lower bounded by Cn^{q-2} (for some constants C). We define

$$h := \sum_{\substack{j,k,S, \\ S'\in\sigma(S)}} J_{S,j}\big(J_{S',k}\big)^2, \tag{4.96}$$

for which $\mathbb{E}(h)=0$. We note that h is a homogeneous polynomial in standard Gaussian random variables of degree 3. To upper bound the moments of g, and thus the moments of D_4, we again use Lemma 4.E.1 from [33]. We use equation (4.88) to find an upper bound for the rth moment of h. We first calculate $\mathbf{D}^d h$ for $d=1,2,3$. Then, for each d, we upper bound $\big\|\mathbb{E}\big(\mathbf{D}^d h\big)\big\|_{\mathcal{P}}$ for all partitions $\mathcal{P}$ of $[d]$. Thereby, we show that for all d and associated partitions $\mathcal{P}([d])$, $\big\|\mathbb{E}\big(\mathbf{D}^d h\big)\big\|_{\mathcal{P}}$ can be upper bounded such that $\mathbb{E}\Big(|D_4|^r\Big) \leq \big(C\sqrt{n}\big)^r$ for all even $2\leq r\leq 16\cdot 2n$. The 0th moment trivially satisfies this bound, and therefore it holds for all even $r\leq 16\cdot 2n$.

The derivatives of h are equal to:

$$Dh = \Big(\sum_{S'\in\sigma(S),p} J^2_{S',p} + 2J_{S,j}\sum_{S'\in\sigma(S),p} J_{S',p}\Big)_{(S,j)} \Longrightarrow \mathbb{E}\big(Dh\big) \leq \big(Cn^{q-1}\big)_{(S,j)}, \tag{4.97}$$

where the sum over p runs from 0 to n and we have used the bounds on $|\sigma(S)|$. Note that this is a pointwise upper bound on the entries of the vector $\mathbb{E}\big(Dh\big)$, which

will be enough to bound the corresponding norm.

$$D^2 h = \begin{pmatrix} 2J_{T,k} + 2J_{S,j}, & \text{if } T \in \sigma(S) \text{ and } \forall k \\ 2\sum_{S' \in \sigma(S),p} J_{S',p}, & \text{if } (T,k) = (S,j) \\ 0, & \text{otherwise} \end{pmatrix}_{(S,j),(T,k)} \Longrightarrow \mathbb{E}\big(D^2 h\big) = (0)_{(S,j),(T,k)}, \tag{4.98}$$

$$D^3 h = \begin{pmatrix} 2, & \text{if } ((U,l) = (T,k)\,(U,T \in \sigma(S))) \text{ or} \\ & ((U,l) = (S,j) \text{ and } T \in \sigma(S) \text{ and } \forall k) \text{ or} \\ & ((T,k) = (S,j) \text{ and } U \in \sigma(S) \text{ and } \forall l) \\ 0, & \text{otherwise} \end{pmatrix}_{(S,j),(T,k),(U,l)} \Longrightarrow \mathbb{E}\big(D^3 h\big) = D^3 h. \tag{4.99}$$

In Table 4.4, we give the values of $\left\|\mathbb{E}\big(\mathbf{D}^d h\big)\right\|_{\mathcal{P}}$ for all partitions $\mathcal{P}([d])$ for $d = 1,2,3$. $\left\|\mathbb{E}\big(\mathbf{D}^d h\right)\|_{\mathcal{P}}$ for $d = 1$ can be straightforwardly evaluated using equation (4.81) and for $d = 2$ can be trivially evaluated by using $\mathbb{E}\big(\mathbf{D}^2 h\big) = 0$. For $d = 3$, $\left\|\mathbb{E}\big(\mathbf{D}^d h\right)\|_{\mathcal{P}}$ can be upper bounded using equation (4.81), and the triangle and Cauchy-Schwarz inequalities. To obtain these upper bounds, we have again used the bounds on $|\sigma(S)|$.

$\mathcal{P}$	$\lvert\mathcal{P}\rvert$	$\left\|\mathbb{E}\big(\mathbf{D}^d h\big)\right\|_{\mathcal{P}}$
$\{1\}$	1	$C\big(n^{q-1}\big)^{3/2}$
$\{1,2\}$	1	0
$\{1\}\{2\}$	2	0
$\{1,2,3\}$	1	$Cn^{q-1/2}$
$\{1\}\{2,3\}$	2	Cn^q
$\{2\}\{1,3\}$	2	Cn^q
$\{1,2\}\{3\}$	2	Cn^q
$\{1\}\{2\}\{3\}$	3	$Cn^{q/2}$

Table 4.4: The different partitions $\mathcal{P}$ of $[3]$ into non-empty parts, with the associated number of parts $|\mathcal{P}|$, and (the upper bounds for) the associated $\left\|\mathbb{E}\big(\mathbf{D}^d h\big)\right\|_{\mathcal{P}}$ for h in equation (4.96).

Combining the upper bounds for $\left\|\mathbb{E}\big(\mathbf{D}^d h\big)\right\|_{\mathcal{P}}$ in Table 4.4 with the factor of $r^{|\mathcal{P}|/2} (\leq Cn^{|\mathcal{P}|/2})$ in equation (4.88) and the normalization factor in equation (4.95), we find – using $\mathbb{E}(h) = 0$ – that indeed $\mathbb{E}\Big(|D_4|^r\Big) \leq \big(C\sqrt{n}\big)^r$ for all even $r \leq 16 \cdot 2n$.

In conclusion, we have shown that the rth moments (for even $r \leq 16 \cdot 2n$) of D_0, D_1, D_2, D_3 and D_4 can be upper bounded by $\big(C\sqrt{n}\big)^r$, and hence, by equation (4.73), the $(2n)$th moment of $A(1)$ can be upper bounded by $\big(C\sqrt{n}\big)^{2n}$. Thereby, we have also established that the second condition in equation (4.57) is satisfied.

4.F. Two-colored SYK to standard SYK

In this Appendix, we give the proof of Lemma 4.7.3.

Proof. To establish Lemma 4.7.3, we now show that the state ρ that achieves $\mathrm{tr}(H^{(2)}\rho) \geq C\sqrt{n}$ (with $H^{(2)}$ defined in equation (4.40)), with probability at least $1-\exp\big(-\Omega(n)\big)$ also achieves $\mathrm{tr}(H\rho) \geq C\sqrt{n}$ for the standard SYK Hamiltonian. To that end, we consider a standard SYK model with $2n$ Majorana operators and partition these Majorana operators into a subset of size $\frac{2n(q-1)}{q}$ and a complementary subset of size $\frac{2n}{q}$. The standard SYK model Hamiltonian H, see equation (4.2), consists of $\binom{2n}{q}$ terms. These terms are labeled by all ordered subsets $\{j_1 < \ldots < j_q\}$, and $\mathcal{I}$ denotes the collection of these subsets. The terms in H for which $q-1$ Majorana operators are in the first subset, and the other Majorana operator is in the complementary subset, are labelled by ordered subsets $\{j_1 < \ldots < j_q : j_1 < \ldots < j_{q-1} < \frac{q-1}{q} \leq j_q\}$. We denote the collection of these subsets by $\mathcal{T}$. The collection of other subsets is denoted by $\mathcal{T}' = \mathcal{I}\backslash\mathcal{T}$. $\mathcal{T}$ and $\mathcal{T}'$ thus correspond to collections of terms in the Hamiltonian. We denote the Hamiltonian consisting of the collection $\mathcal{T}$ by $H_{\mathcal{T}}$ and the Hamiltonian consisting of terms $\mathcal{T}'$ by $H_{\mathcal{T}'}$, hence $H = H_{\mathcal{T}} + H_{\mathcal{T}'}$. $H_{\mathcal{T}}$ corresponds exactly to the 2-colored Hamiltonian in equation (4.40) when multiplied by

$$e^{-(q-1)/2} \leq \frac{\sqrt{\binom{2n}{q}}}{\sqrt{\frac{2n}{q}}\sqrt{\binom{\frac{q-1}{q}2n}{q-1}}} \leq e^{q/2}, \tag{4.100}$$

which, importantly, is lower bounded and upper bounded by a constant in n. Note that the sizes of the two subsets into which the Majorana operators are partitioned can in fact be *any* $c2n$ and $(1-c)2n$ for $c = O(1)$ (instead of just $\frac{2n(q-1)}{q}$ and $\frac{2n}{q}$). The factor in equation (4.100) is lower bounded and upper bounded by a constant in n as well for all of these other partitions. Hence n is not constrained to be a multiple of q.

For any state ρ, $\mathbb{E}\big(\mathrm{tr}(H\rho)\big) = 0$, where the expectation value is w.r.t. the couplings in H since the couplings are random variables with zero mean. The state ρ_θ defined in Eq. equation (**??**) is able to achieve $\mathrm{tr}(H^{(2)}\rho_\theta) \geq C\sqrt{n}$ (with high probability) since it is constructed using a circuit that itself *depends* on the random couplings J_I ($I \in \mathcal{T}$) appearing in $H^{(2)}$. Since ρ_θ does *not* depend on the couplings J_I with $I \in \mathcal{T}'$, we have $\mathrm{tr}(H_{\mathcal{T}'}\rho_\theta) = 0$. Since: (i) $|\mathrm{tr}(C_I\rho)| \leq 1$ (for any ρ) for $I \in \mathcal{T}'$, (ii) that each J_I is a standard Gaussian random variable, and (iii) that $|\mathcal{T}'| \leq \binom{2n}{q}$, the quantity

$$\mathrm{tr}\big(H_{\mathcal{T}'}\rho\big) = \binom{2n}{q}^{-1/2} \sum_{I\in\mathcal{T}'} J_I \mathrm{tr}(C_I\rho) \tag{4.101}$$

is a Gaussian random variable with zero mean and variance at most one, for any ρ. Then, $\mathbb{E}\big[\exp\big(t\,\mathrm{tr}(H_{\mathcal{T}'}\rho)\big)\big] \leq \exp\big(\frac{1}{2}t^2\big)$ for all $t \geq 0$. Applying Chernoff's bound to $\mathrm{tr}(H_{\mathcal{T}'}\rho)$, and choosing $t = C\sqrt{n}$, we obtain:

$$\Pr\Big[|\mathrm{tr}(H_{\mathcal{T}'}\rho)| \geq C\sqrt{n}\Big] \leq 2\exp\big(-\Omega(n)\big), \tag{4.102}$$

for any constant C.

Using equation (4.102) and $\mathrm{tr}(H\rho) = \mathrm{tr}(H_{\mathcal{T}}\rho) + \mathrm{tr}(H_{\mathcal{T}'}\rho)$, we conclude that the state ρ_θ which achieves $\mathrm{tr}(H^{(2)}\rho_\theta) \geq C\sqrt{n}$ (i.e., for the 2-colored SYK Hamiltonian) with probability at least $1 - \exp\big(-\Omega(n)\big)$, also achieves $\mathrm{tr}(H\rho_\theta) \geq C\sqrt{n}$ (where H is the standard SYK Hamiltonian in equation (4.2)) with probability at least $1 - \exp\big(-\Omega(n)\big)$. Therefore, $\lambda_{\max}(H) \geq C\sqrt{n}$ with probability at least $1 - \exp\big(-\Omega(n)\big)$. □

5

Solving free fermion problems on a quantum computer

5.1. Introduction

Quantum many-body dynamics can be naturally simulated by a quantum computer [1], enabling its applications in condensed matter and quantum chemistry. For a system of size N, standard quantum algorithms use poly(N) resources for such simulations. It implies an exponential advantage over classical methods, when dealing with a generic many-body system. Such a general advantage may not hold in special cases of interest, such as the modeling of free fermions, where the best classical algorithms also have poly(N) cost [2–4]. This classical efficiency has been key to many successes of computational physics, as free fermions model a variety of systems in condensed matter and quantum chemistry; they have also been used in computational strategies for solving interacting fermion systems, using mean-field (Hartree-Fock), perturbative methods or dynamical mean-field theory. Nonetheless, in the practical simulations of noninteracting fermions, even the most efficient numerical methods become too expensive for large systems. This motivates the key question of this work: can a quantum computer boost free-fermion simulations beyond what can be done classically? We answer this question in the affirmative, presenting quantum algorithms with an exponential speedup and memory compression for several free-fermion problems.

To appreciate the value of such an exponential reduction, consider numerical simulations of free-fermion models of materials and interfaces for quantum transport [5, 6]. These can become prohibitive when involving more than $N = 10^9$ modes, which is of practical interest when simulating 3-dimensional lattice models. Upon compression, a system of 10^9 sites can be described by $n = 30$ qubits. Larger systems of practical interest could still be accessed with moderately sized quantum computers. Indeed, even simulating one mole ($N \simeq 10^{24}$) of fermionic modes requires fewer than $n = 80$ qubits in compressed form. This opens the door to modeling free fermions near the thermodynamic limit — a desired but often challenging goal.

Our result is based on an understanding of the reduced classical complexity of free-fermion systems. As an inspiration, we used the fact that the matchgate computations and the dynamics of free fermion problems on $N = 2^n$ modes can be

simulated in compressed form, using $O(n)$ *space* on a quantum computer [7–11]. In this work, we go beyond these memory compression results to identify free fermion problems that also permit an exponentially improved, poly(n) quantum *run-time*. Our key idea is to represent the relevant 2^n-sized object—such as the correlation matrix or a Green's function of a free-fermion state—as a block of an n-qubit unitary. This unitary can be given as an efficient quantum circuit; we provide explicit construction methods by leveraging the modern quantum algorithm toolbox of block-encoding manipulations [12–16]. In particular, we show how to construct the desired unitary for free-fermion states arising from time dynamics or thermal equilibrium. Given block-encodings of the aforementioned objects into a circuit, we show how to accurately extract various physical quantities for a state, including the occupation number on a given site, or energy density across the entire system. We analyze the application of our methods to free-fermion models on d-dimensional lattices and expander graphs. For the particular case of d-dimensional lattices, we argue that a polynomial runtime improvement can be expected, based on the comparison with best available classical algorithms. For expander graphs, the same analysis suggests an *exponential* speedup. On a more general geometry, the problem of single-particle time dynamics is BQP-hard [17] — as hard as any problem that can be efficiently solved by a quantum computer. This rigorously proves that our approach offers an exponential quantum speedup at least for some geometries (as long as quantum computers can offer exponential speedups in principle). Finally, we outline the generalization of our approach to systems beyond free fermions.

Our work can be viewed as a fermionic counterpart to [17], which shows how the time-dynamics of a system of coupled oscillators can be solved exponentially faster on a quantum versus a classical computer — with further applications in [18]. Compared to the alternative and recent work [19] which focuses on encoding a correlation matrix into a state, our method using block-encodings has an exponential advantage in signal strength for the extraction of local observables (see Appendix 5.A for more details).

We note that quantum algorithms for compressed simulation of *interacting* fermionic models have also been considered in e.g. [20], where a Fermi-Hubbard model is simulated in the $O(1)$-particle sub-space. By contrast, we consider N-mode systems with as many as $\Theta(N)$ particles.

5.1.1. Preliminaries

Throughout this work, we set $N = 2^n$. A particle-conserving free fermion Hamiltonian H can be written as

$$H = \sum_{i=0,j=0}^{N-1,N-1} h_{ij} a_j^\dagger a_i, \tag{5.1}$$

with Hermitian matrix h, which we will assume to be $O(1)$-sparse (i.e., there are at most a constant number of non-zero entries in each row) and $|h_{ij}| \leq 1$. Here $\{a_i^\dagger, a_j\} = \delta_{ij}, \{a_i, a_j\} = \{a_i^\dagger, a_j^\dagger\} = 0$. We denote the fermionic particle number operator as $\hat{N} = \sum_{i=0}^{N-1} a_i^\dagger a_i$, and we restrict ourselves to Hamiltonians which preserve particle number[1]. We allow for states ρ with an arbitrary number of particles $\mathrm{Tr}(\hat{N}\rho)$, which

[1] There are straightforward generalizations, using Majorana fermion language, to just parity-conserving free fermion Hamiltonians.

in general may scale with $N = 2^n$. Observe that in the case of single-particle dynamics $\mathrm{Tr}(\hat{N}\rho) = 1$, the fermionic nature of the system does not come into play and bosonic or fermionic dynamics are equivalent.

The Hermitian correlation matrix M of a fermionic state ρ on N modes is defined as

$$M_{ij} = \mathrm{Tr}(a_i^\dagger a_j \rho) \in \mathbb{C}, \tag{5.2}$$

and obeys $0 \leq M \leq I$, and $\mathrm{Tr}(M) = \langle \hat{N} \rangle$. M contains observable information about the fermionic state ρ: for example, M_{jj} is the mean fermion occupation number of a state ρ in the mode j. Furthermore, an expectation value of a free fermion Hamiltonian (equation (5.1)) can be expressed as $\mathrm{Tr}\big(H\rho\big) = \sum_{i,j} h_{ij} M_{ji}$. If ρ is itself free-fermionic, expectation values of *interacting* Hamiltonians can also be obtained from M, using Wick's theorem (see Proposition 1.5.8).

Throughout this work, we will use $[N = 2^n]$ in a non-traditional way, namely offset by 1: $[N] \equiv \{0, \ldots, N-1\}$. We also use the standard notation $f(x) = O(g(x))$ if a function is asymptotically upper bounded by $\mathrm{const} \cdot g(x)$, $f(x) = \Omega(g(x))$ if lower bounded, and $f(x) = \Theta(g(x))$ if both (i.e., scaling in the same way as $\mathrm{const} \cdot g(x)$).

5.1.2. Outline

In Section 5.2, we detail our objects of interest: correlation matrices for the time-evolved and thermal equilibrium states, as well as the Green's function matrix. Each of these objects carries physically meaningful information about the system, and has a form $F(h)$ — an explicit matrix function of h.

In Sections 5.3-5.4 we demonstrate, how the information contained in these matrices can be efficiently extracted from a quantum computer, using the framework of so-called *block-encodings*. Figure 5.1 provides a visual scheme, illustrating the structure of our approach. Section 5.3 explains the block-encoding framework, namely how any $N \times N$ sized matrix A can be encoded into a block of a unitary U_A on $O(n) = O(\log N)$ qubits. We also review the basic tools to produce and manipulate such unitaries U_A, which were previously established in the literature. Given the matrix functions $F(h)$ of our interest, we will aim to produce the block-encodings $U_{F(h)}$ as compact quantum circuits.

The starting point of our circuit construction are smaller unitaries which encode h itself; these unitaries are called sparse access oracles (as the matrix h is required to be sparse). In Section 5.3.1, we show how to to implement the sparse access oracles as quantum circuits of size $\mathrm{polylog}\, N$. Such implementations are specific to the model of interest: we discuss the cases of d-dimensional lattice models and some expander graph geometries; we also demonstrate that quenched disorder can be introduced efficiently.

In the following Section 5.3.2, having implemented the sparse access oracles for h, we move to the second step of the construction — realizing the block-encoding of matrix functions $F(h)$ of our interest. We detail how this can be done with quantum circuits of size that scales polynomially in $\log N$, as well as polynomially in parameters of the respective function, such as the evolution time t, the inverse temperature β, or the Green's function regularization parameter δ^{-1}.

Being able to run a circuit which realizes the block-encoding of the matrix $F(h)$ is not sufficient for a successful simulation: one still needs an efficient method to extract

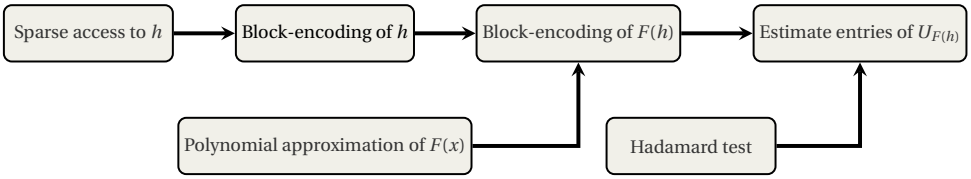

Figure 5.1: Overview of the proposed quantum computational method to extract properties of free fermionic systems such as the entries of matrices listed in Section 5.2. The elements of the construction illustrated here are described in detail in Sections 5.3-5.4.

physically relevant information from $F(h)$. Section 5.4 shows how this can indeed be done, using a Hadamard test and basic sampling techniques. In particular, we demonstrate that the local observables and global densities (such as the total energy density) can be accurately extracted from a block encoding of $F(h)$, while maintaining the poly$\log N$ complexity of the algorithm.

Sections 5.5 and 5.6 deal with a crucial question: does our approach provide a significant speedup compared to a classical computation? This question can be answered in the affirmative from two perspectives. In Section 5.5, we take a complexity theory perspective and observe that simulating free-fermionic time dynamics using $\log N$ qubits is BQP-hard. In other words, for a classical computer it is strictly as hard as simulating a general quantum computation on $\log N$ qubits — which is widely assumed to be exponentially hard in the number of qubits. This establishes that our approach yields an exponential quantum speedup for at least some system geometries. In Section 5.6, we take a more practical perspective, and focus on the geometries of direct physical interest (such as those given in Section 5.3.1). For these models, we compare the performance of our algorithm with the best classical algorithms which are currently available. In particular, we find that the quantum algorithm yields a power $(d+1)$ polynomial speedup when simulating the time dynamics of d-dimensional lattice models. For simulations of the expander models, we demonstrate an exponential quantum speedup.

We close the main text with the Section 5.7, where we sketch how our approach can be generalized to other systems, such as free fermions with pairing terms ($\sim \Delta a_j a_k$) and free bosons with particle conservation. In Section 5.8, we discuss the future directions.

5.2. Objects of interest

We consider three kinds of target objects — matrix functions of h, whose entries encode the physically relevant information.

- Correlation matrices $M^{(\beta)}$ of thermal states $\rho_\beta = e^{-\beta H}/\mathrm{Tr}(e^{-\beta H})$ associated with free-fermion Hamiltonians H:

$$M^{(\beta)} = \frac{I}{I + e^{\beta h}}. \tag{5.3}$$

  The eigenvalues $n_\beta(\epsilon_i) = (1 + e^{\beta\epsilon_i})^{-1}$ of $M^{(\beta)}$ correspond to the Fermi-Dirac

distribution, with ϵ_i the eigen-energies of h, and $\langle \hat{N} \rangle_\beta = \sum_i n_\beta(\epsilon_i)$. Note that h here includes a chemical potential term $-\mu\mathbb{I}$, if needed.

- Correlation matrices $M(t)$ of time-evolved states $\rho(t)$ (where the time evolution of $\rho(0)$ is under a free-fermion Hamiltonian H):

$$M(t) = e^{iht} M e^{-iht}, \tag{5.4}$$

with M denoting the correlation matrix of $\rho(0)$.

In fact, we will consider a slightly more general object:

$$M(t_1, t_2) = e^{iht_1} M e^{-iht_2}, \tag{5.5}$$

the entries of which correspond to

$$M_{ij}(t_1, t_2) = \mathrm{Tr}\,(a_i^\dagger(t_1) a_j(t_2) \rho), \tag{5.6}$$

with Heisenberg operators $a_i^\dagger(t), a_j(t)$ w.r.t. the free-fermion Hamiltonian H.

Note that for a Hamiltonian $H = H_0 + V$ with free-fermionic H_0 and interacting perturbation V, after applying $U(t) = e^{-iHt}$ to an initial free-fermionic state ρ, observables involving creation and annihilation operators can be obtained from $M(t_1, t_2)$ in equation (5.5). This can be done via a perturbative expansion of $U(t) = e^{-iHt}$ and using Wick's theorem.

- The Green's function (in the frequency domain) w.r.t. a thermal state ρ_β of a free-fermion Hamiltonian:

$$G^{(\delta,\beta,\omega)}(h) = \frac{\delta}{2}\Big[\Big(1 - \frac{1}{1+\exp(\beta h)}\Big)\frac{1}{i\delta - (h+\omega)} + \Big(\frac{1}{1+\exp(\beta h)}\Big)\frac{-1}{i\delta + (h+\omega)}\Big], \tag{5.7}$$

with $\delta > 0$ a regularization parameter.

$G^{(\delta,\beta,\omega)}(h)$ is a Fourier transform of the time-domain Green's function, the entries of which are given by (here we use time-ordering unlike in equation (5.6)):

$$\begin{aligned} G_{ij}(t_1, t_2) &= \begin{cases} i\mathrm{Tr}\big(a_i^\dagger(t_1) a_j(t_2) \rho_\beta\big), & \text{for } t_1 \geq t_2, \\ -i\mathrm{Tr}\big(a_j(t_2) a_i^\dagger(t_1) \rho_\beta\big), & \text{for } t_1 < t_2, \end{cases} \\ &= \begin{cases} \big(i e^{ih(t_1-t_2)} \frac{1}{1+\exp(\beta h)}\big)_{ij}, & \text{for } t_1 \geq t_2, \\ \big(-i e^{ih(t_1-t_2)}\big(1 - \frac{1}{1+\exp(\beta h)}\big)\big)_{ij}, & \text{for } t_1 < t_2. \end{cases} \end{aligned} \tag{5.8}$$

The regularization parameter δ in equation (5.7) ensures that the Fourier transform converges in the case of an isolated system, but can also model interactions with a bath at finite temperature [21].

5.3. Block-encodings

Let us consider encoding a Hermitian $(N \times N)$-dimensional matrix A into a block of an $n+m$ qubit unitary U_A. In general, an n-qubit matrix A is said to be block-encoded into U_A if it is equal to the block of U_A where m qubits are in a trivial state, with some coefficient α

$$A_{ij} = \alpha \langle i|_n \langle 0|_m U_A \left|j\right\rangle_n |0\rangle_m . \tag{5.9}$$

Here, the matrix indices $i, j \in [N]$ are interpreted as bit-strings of length n. The coefficient $\alpha \geq 1$ arises from the fact that $\|U_A\| = 1$ while A is arbitrary. If $\|A\| \leq 1$, we can take $\alpha = 1$. For a useful block-encoding, the coefficient α should not blow up beyond $\mathrm{poly}\log N$. Fortunately, in the applications considered in this work, α will remain a small constant. For the same reasons of maintaining efficiency, we will limit the number of ancillary qubits m to $O(\log N)$.

We will also allow block-encoding with error ε, the deviation in operator norm between A and $\alpha \langle 0|_m U_A |0\rangle_m$.

Definition 5.3.1. *For a matrix A on n qubits and $\alpha, \varepsilon \in \mathbb{R}_+$, an $(m+n)$-qubit unitary U_A is an (α, m, ε)-block-encoding of A, if*

$$\|A - \alpha(\langle 0|^{\otimes m} \otimes \mathbb{1}) U (|0\rangle^{\otimes m} \otimes \mathbb{1})\| \leq \varepsilon. \tag{5.10}$$

where $||\cdot||$ is the spectral norm.

The quantum circuits that approximately block-encode the matrix functions $F(h)$ are built using elementary circuits U_h that block-encode h. These latter block-encodings U_h, in turn, contain unitaries which realize so-called *sparse query access* to h. To access an s-sparse matrix h, i.e. a matrix which has up to $s = O(1)$ nonzero entries in any row, we will use 'oracle' unitaries O_r and O_a which produce the entries of h. The 'row' oracle O_r returns, for a given row i, all column indices where the matrix h has nonzero entries. The 'matrix entry' oracle O_a returns the value of h (given with n_a bits) for a given row and column index. This way, entries of h can be retrieved without explicit access to the $\Theta(2^n)$ nonzero entries of matrix h. Let us formally define the *oracle tuple* $\mathcal{O}_h$ of a sparse matrix h containing the row and matrix entry oracles, and also their inverses and controlled versions as follows.

Definition 5.3.2 (Sparse Access Oracle Tuple $\mathcal{O}_h$)**.** *Sparse access for an s-sparse $2^n \times 2^n$ matrix h is defined as*

$$\begin{aligned} O_r |i\rangle |0\rangle^{\otimes s(n+1)} &= |i\rangle |r(i,1)\rangle |r(i,2)\rangle \ldots |r(i,s)\rangle , \\ &\qquad\qquad \forall i \in [2^n], \\ O_a |i\rangle \left|j\right\rangle |0\rangle^{\otimes n_a} &= |i\rangle \left|j\right\rangle \left|h_{ij}\right\rangle , \quad \forall i, j \in [2^n], \end{aligned} \tag{5.11}$$

where $r(i,k)$ is the index for the kth nonzero entry of the ith row of h. Let us now cover a few technicalities. O_r is a matrix acting on $(s+1)(n+1)$ qubits, and so the first qubit of $|i\rangle$ is in $|0\rangle$. To accommodate rows with less than s non-zero entries, one uses the following. If the ith row contains $s' < s$ non-zero entries, then the last $(s-s')(n+1)$ qubits are put in the

state $|1\rangle|k\rangle$. *Note that for states* $|r(i,1)\rangle \ldots |r(i,s')\rangle$, *the first qubit is in* $|0\rangle$. h_{ij} *is the value of the* (i,j)*th entry of* h, *described by a bitstring with* n_a *binary digits (we will assume this representation to be exact).* O_a *is a matrix acting on* $2n+n_a$*-qubits.*

Furthermore, we define the controlled version of the above sparse access, consisting of

$$\begin{aligned} C\text{-}O_r &= O_r \otimes |1\rangle\langle 1|_a + \mathbb{1} \otimes |0\rangle\langle 0|_a, \\ C\text{-}O_a &= O_a \otimes |1\rangle\langle 1|_a + \mathbb{1} \otimes |0\rangle\langle 0|_a, \end{aligned} \tag{5.12}$$

where each matrix now acts on an additional (ancillary) qubit a. *We call the collection of six oracles* $(O_r, O_a, C\text{-}O_r, C\text{-}O_a, O_r^{-1}, O_a^{-1}, C\text{-}O_r^{-1}, C\text{-}O_a^{-1})$ *the sparse access oracle tuple* $\mathcal{O}_h$ *of* h.

The relation between this definition of the oracle tuple and another common definition is discussed in Appendix 5.B for completeness.

Let us now present the following statements, relating the construction of the block-encoding of h and that of polynomials of h. Note that these block-encoding constructions contain calls to oracles from the oracle tuple $\mathcal{O}_h$ in Definition 5.3.2. We shall use these results when constructing the block-encodings of our desired matrix functions. The following statements use Definition 5.3.1.

Proposition 5.3.3 (Lemma 48 in [22])**.** *A* $(s, n+3, \varepsilon_{BE_h})$*-block-encoding of* h, U_h *(and its controlled version) consists of* $O(1)$ *calls to oracles from* $\mathcal{O}_h$ *tuple,* $O\big(n+\log^{5/2}(s^2/\varepsilon_{BE_h})\big)$ *elementary gates and* $O\big(sn+n_a+\log^{5/2}(s^2/\varepsilon_{BE_h})\big)$ *ancillary qubits. Here* n_a *denotes the number of bits with which the entries of* h *are specified.*

Proposition 5.3.4 (Theorem 31 in [14])**.** *Let* $p_d(x)$ *denote a degree-*d *polynomial s.t.* $|p_d(x)| \leq 1/2$ *for* $x \in [-1,+1]$. *Then, a* $(1, n+5, 4d\sqrt{\varepsilon_{BE_h}/s}+\delta)$*-block-encoding of* $p_d(h/s)$, $U_{p_d(h/s)}$, *consists of* $O((n+4)d)$ *elementary gates, and at most* d *calls to unitaries* U_h, U_h^{-1} *or controlled-*U_h. *The classical description of this circuit can be obtained classically in* $poly(d, \log(1/\delta))$ *time.*

5.3.1. Sparse-access realization for physical systems

The starting point for our method is to realize the sparse access tuple $\mathcal{O}_h$ for the system Hamiltonian h, using efficient quantum circuits. In particular, we need circuit realizations for unitaries O_r and O_a (equation (5.11)); these in fact can be given as (reversible) classical circuits, as no entanglement generation is required. Then the controlled and inverse unitaries from $\mathcal{O}_h$ can also be obtained as efficient circuits (with a constant factor overhead), controlling or inverting the circuits for O_r and O_a gate-by-gate. Please note that 'efficient' in our case means $\text{poly}\log N$ gate complexity, i.e., polynomial in the number of qubits rather than the size of h. In other words, simply looking up the entries of the $N \times N$ matrix h would not suffice, as that takes time which is exponentially longer than desired. Despite this difficulty, the requirement of the efficient implementation of $\mathcal{O}_h$ can be satisfied for a variety of h of interest.

A large family of free-fermionic models for which the sparse access to h can be efficiently realized are d-dimensional tight-binding models. Consider a d-dimensional square lattice $\mathcal{L}$ with $L_1 \times L_2 \times .. \times L_d = N_s$ sites, with either periodic or open boundaries.

For each site $\vec{x}$, let there be up to $N_0 = O(1)$ onsite degrees of freedom such as spin, or local orbital degrees of freedom. We can thus represent each fermionic mode using $n = \left(\Pi_{i=1}^{d} \lceil \log_2 L_i \rceil\right) \times \lceil \log_2 N_0 \rceil$ qubits as $|\vec{x} = (x_1, \ldots, x_d), o\rangle$ where $N_s = \Theta(2^n)$. Inside the lattice, let there be $O(1)$ non-overlapping rectangular domains, modeling different physical regions such as leads versus bulk regions, where parameters in H can be different. We thus consider Hamiltonians of the following form:

$$H = \sum_{o_1, o_2} \sum_{\vec{x} \in \mathcal{L}, |\vec{t}|_{\mathrm{M}} \leq l} h_{\vec{x}, o_1, \vec{x}+\vec{t}, o_2} \, a^{\dagger}_{\vec{x}+\vec{t}, o_2} a_{\vec{x}, o_1} + \text{h.c.}, \tag{5.13}$$

where it is understood (but notationally awkward) that the sum over $\vec{x} \in \mathcal{L}, |\vec{t}|_{\mathrm{M}} \leq l$ only counts each possible hopping term once. In addition, we have

$$\begin{aligned} &h_{\vec{x}, o_1, \vec{x}+\vec{t}, o_2} = g\left(o_1, o_2, D(\vec{x}), D(\vec{x}+\vec{t}), \vec{t}\right), \\ &|h_{\vec{x}, o_1, \vec{x}+\vec{t}, o_2}| \leq 1. \end{aligned} \tag{5.14}$$

Here $|.|_{\mathrm{M}}$ means Manhattan distance in the lattice; the maximal range of the interaction is posited to be constant — $l = O(1)$. The function $D(\vec{x})$ returns the domain to which $\vec{x}$ belongs: since the domains are rectangular regions, $D(\vec{x})$ can be efficiently computed using standard reversible artihmetic circuits. If $\vec{x}$ or $\vec{x}+\vec{t}$ does not belong to any domain (for example, $\vec{x}+\vec{t}$ is beyond the boundaries of the lattice), the coefficient $h_{\vec{x}, o_1, \vec{x}+\vec{t}, o_2} = 0$. Thus, the function g only takes in $O(1)$ information and all $O(1)$ possible nonzero outputs of g can be stored classically, using, say, $O(n_a)$ bits. To realize the oracles $\mathcal{O}_h$ from Definition 5.3.2 as poly(n)-sized quantum circuits, observe that one can efficiently generate the $O(1)$ input to g and lookup the relevant information.

Going beyond local d-dimensional models, we give an example of a model on an *expander graph* which has sparse query access. These graphs have the important property that the number of vertices that lie a distance d away from a given vertex scales exponentially in d. Free-fermionic models on such graphs have been a subject of recent interest, especially in the studies of Anderson localization on random regular graphs [23–26]. In Appendix 5.C, we provide details of the realization of $\mathcal{O}_h$ as poly(n)-sized quantum circuits for a simple example: the Margulis expander graph.

So far, we have proposed models with efficient sparse access where there was only a limited number of possible options for the hopping parameters, and they were input 'by hand'. This is in line with a necessary limitation — even though the system has size N, we should be unable to assign every mode an independent value of the hopping parameter.

However, this restriction can be somewhat relaxed. In particular, one can show that local quenched disorder can also be incorporated into h. This has the significance for physics application, as it allows to study Anderson localization. For simplicity, let us focus on realizing onsite disorder in a single domain D^* of a tight-binding model. This means that we introduce a single change to the Hamiltonian of equation (5.13) and equation (5.14). Namely, if $D(\vec{x}) = D^*$ and $\vec{t} = 0$ (both equalities are efficiently checkable), the value of $h_{\vec{x}, o_1, \vec{x}+\vec{t}, o_2}$ will be replaced by

$$h_{\vec{x}, o_1, \vec{x}+\vec{t}, o_2} = \delta_{o_1, o_2} \, \mathrm{PRF}(\vec{x}), \tag{5.15}$$

where $\delta_{a,b}$ is the Kronecker symbol and PRF is a pseudo-random function of the lattice site coordinate $\vec{x}$. Note that a pseudo-random function can be realized as an efficient

classical circuit [27, 28]. Other models of local disorder can be realized similarly. We note that an independent work [29] discusses the application of simulating disordered free fermions in more detail.

5.3.2. Block-encodings of relevant matrix functions

Given the poly(n)-effort sparse access tuple $\mathcal{O}_h$, we now aim to realize a block-encoding of the desired matrix functions of h (Section 5.2) with an efficient quantum circuit. We will approximate these functions with polynomials of sufficiently low degree, enabling us to use standard methods of block-encodings manipulation (Proposition 5.3.4).

To construct the polynomial approximations, let us first establish the following. Proposition 5.3.4 prescribes how degree-d polynomials $p_d(x)$ with $x = h/s$ can be block-encoded, with s the sparsity of h. We thus require a polynomial approximation $p_d(x)$ to our functions of interest $F(h = sx)$ to be sufficiently accurate in the domain $x \in [-\|h\|/s, +\|h\|/s]$. It can be argued straightforwardly that this domain is at most $[-1, +1]$ by bounding the spectral norm of h:

Proposition 5.3.5. *Let h denote an $s = O(1)$-sparse Hermitian $N \times N$ matrix with $|h_{ij}| \leq 1$, $\forall i, j$. The spectral norm $||h||/s \leq 1$ by the Gershgorin circle lemma (which says that every eigenvalue of h lies within at least one of the N discs $D_i = \{z \in \mathbb{C} : |z - h_{ii}| \leq \sum_{j \neq i} |h_{ij}|\}$).*

To block-encode the thermal correlation matrix in equation (5.3) and the thermal Green's function in equation (5.7), we need to approximate the functions

$$f^{(\beta)}(x) := \frac{1}{4} \frac{1}{1 + \exp\big(\beta s x\big)} \tag{5.16}$$

and

$$g^{(\delta,\beta,\omega)}(x) := \frac{1}{4}\frac{\delta}{2}\Big[\Big(1 - \frac{1}{1+\exp\big(\beta s x\big)}\Big)\frac{1}{i\delta - (sx+\omega)} + \Big(\frac{1}{1+\exp\big(\beta s x\big)}\Big)\frac{-1}{i\delta + (sx+\omega)}\Big] \tag{5.17}$$

in the domain $x \in [-1, +1]$. These functions have poles in the complex plane at $z = (2k+1)i\pi/\beta$ (with $k \in \mathbb{Z}$), and at $z = (\pm i\delta - \omega)/s$, respectively. Since these poles might lie in the unit circle for general β and δ, we have to resort to polynomial approximation techniques beyond Taylor approximations to obtain a sufficiently accurate approximation for $x \in [-1, +1]$. In particular, we will employ Bernstein's Theorem:

Lemma 5.3.6 ([30])**.** *Let $f(x)$ be analytic on $[-1, +1]$ and analytically continuable to the interior of an ellipse defined by $E_r = \{\frac{1}{2}(z + z^{-1}) : |z| = r\}$ (for some real-valued $r \geq 1$). Furthermore, let $|f(z)| \leq C$ for $z \in E_r$. The error w.r.t. their polynomial approximation $p_d(x)$ (Chebyshev expansion truncated at degree d) can be bounded as*

$$\max_{x \in [-1,+1]} |f(x) - p_d(x)| \leq \frac{2Cr^{-d}}{r-1}. \tag{5.18}$$

Using this result, we derive the following error bounds for the polynomial approximations of equation (5.16) and equation (5.17). The proofs of Lemmas 5.3.7 and 5.3.8 are given in Appendices 5.D and 5.F.

Lemma 5.3.7 (Simplified version of Lemma 5.D.1 in Appendix 5.D)**.** *For the function $f^{(\beta)}(x)$ in Eq. (5.16) (with $\beta, s \geq 0$), one can efficiently construct a polynomial $p_d(x)$ of degree d such that*

$$\max_{x\in[-1,+1]} |f^{(\beta)}(x) - p_d(x)| \leq \text{poly}(\beta s)/d. \tag{5.19}$$

Lemma 5.3.8 (Simplified version of Lemma 5.F.1 in Appendix 5.F)**.** *For the function $g^{(\delta,\beta,\omega)}(x)$ in equation (5.17) (with $\beta, \delta, s > 0$), one can efficiently construct a polynomial $p_d(x)$ of (even) degree d such that*

$$\max_{x\in[-1,+1]} |g^{(\delta,\beta,\omega)}(x) - p_d(x)| \leq \big(\text{poly}(\beta s) + \text{poly}(s/\delta)\big)/d. \tag{5.20}$$

Combining Lemmas 5.3.7 and 5.3.8 with Propositions 5.3.3 and 5.3.4, we directly obtain Theorems 5.3.9 and 5.3.10 below. The detailed proofs are given in Appendices 5.D and 5.F. Note that – crucially, because of the factors $\frac{1}{4}$ in equation (5.16) and equation (5.17) – the polynomials $p_d(x)$ that are block-encoded obey $|p_d(x)| \leq 1/2$ for $x \in [-1,+1]$, provided that the error of the polynomial approximation is $O(1)$. The size of the circuits that block-encode $M^{(\beta)}(h)$ in equation (5.3) and $G^{(\delta,\beta,\omega)}(h)$ in equation (5.7) is poly(n), provided that $\beta, 1/\delta = \text{poly}(n)$ and when the oracles from $\mathcal{O}_h$ are poly(n)-sized circuits (such as those in Section 5.3.1).

Theorem 5.3.9 (Block-encoding of the thermal correlation matrix equation (5.3))**.** *For an s-sparse Hamiltonian h on n qubits, assume access to the oracle tuple $\mathcal{O}_h$. We denote the controlled $(1, n+5, \varepsilon)$-block-encoding of $\frac{1}{4}M^{(\beta)} = \frac{1}{4}I/(I + \exp(\beta h))$ by $C\text{-}U_{M^{(\beta)}}$. The implementation of this block-encoding for $\beta = poly(n)$ requires $poly(n)/\varepsilon$ calls to oracles from the oracle tuple $\mathcal{O}_h$, $O(n) + n_a + \log^{5/2}(poly(n)/\varepsilon^4)$ ancillary qubits and $O(n) + poly(n)/\varepsilon + \log^{5/2}(poly(n)/\varepsilon^4)$ additional elementary gates. To implement this block-encoding, an additional classical computing time of* poly$\big(n/\varepsilon, \log(1/\varepsilon)\big)$ *is required.*

Theorem 5.3.10 (Block-encoding of the thermal Green's function equation (5.7))**.** *For an s-sparse Hamiltonian h on n qubits, assume access to the oracle tuple $\mathcal{O}_h$. We denote the controlled $(1, n+5, \varepsilon)$-block-encoding of $\frac{1}{4}G^{(\delta,\beta,\omega)}(h)$ in equation (5.7) by $C\text{-}U_{G^{(\delta,\beta,\omega)}}$. The implementation of this block-encoding for $\beta, 1/\delta = poly(n)$ requires $poly(n)/\varepsilon$ calls to oracles from the oracle tuple $\mathcal{O}_h$, $O(n) + n_a + \log^{5/2}\big(poly(n)/\varepsilon^4\big)$ ancillary qubits and $O(n) + poly(n)/\varepsilon + \log^{5/2}\big(poly(n)/\varepsilon^4\big)$ additional elementary gates. To implement this block-encoding, an additional classical computing time of* poly$\big(n/\varepsilon, \log(1/\varepsilon)\big)$ *is required.*

Next, let us focus on block-encoding the time-evolved correlation matrix $M(t)$ in equation (5.4). To block-encode it, we will use a block-encoding of $\exp(iht)$ as a sub-routine. The construction of this latter block-encoding through polynomial approximations is already considered in [13, 22], and we will use this construction from [22] directly. We construct a block-encoding of $M(t)$ using the product of block-encodings of $\exp(iht)$, an initial correlation matrix M and $\exp(-iht)$. A detailed proof of Theorem 5.3.11 is given in Appendix 5.E. There, we in fact consider a block-encoding of the more general object $M(t_1, t_2)$ in equation (5.5).

Theorem 5.3.11 (Simplified version of Theorem 5.E.1 in Appendix 5.E: Block-encoding of the time-evolved correlation matrix in equation (5.4))**.** *For an s-sparse Hamiltonian h on N fermionic modes, assume access to the oracle tuple $\mathcal{O}_h$. In addition, assume access to the $(\alpha, m, \varepsilon_M)$-block-encoding U_M of a correlation matrix M of a fermionic state on N modes. The $\big(\alpha, 2n+m+10, \varepsilon+\varepsilon_M\big)$-block-encoding $U_{M(t)}$ of $M(t) = e^{iht}Me^{-iht}$ can be produced using $D(\alpha,\varepsilon,t) = O\big(|t| + \log(\alpha/\varepsilon)\big)$ calls to oracles from the tuple $\mathcal{O}_h$, and a single use of the block-encoding U_M. Moreover, one uses $O\big(n|t| + \log(\alpha/\varepsilon) + D(\alpha,\varepsilon,t)\big(n + \log^{5/2}(\alpha|t|/\varepsilon)\big)\big)$ elementary gates and $O\big(n_a + \log^{5/2}(\alpha|t|/\varepsilon)\big)$ ancillary qubits.*

5.4. Extracting observables

Having explicitly constructed (α, m, ε)-block-encodings $U_{F(h)}$ of our objects of interest $F(h)$, let us detail how to extract relevant observables from such block-encoding unitaries. If $U_{F(h)}$ is given as a poly(n)-sized quantum circuit, the real and imaginary parts of $F(h)_{ij}$ can be extracted efficiently using the so-called Hadamard test using an ancillary-qubit-controlled-$U_{F(h)}$. Note that the circuit size required to implement controlled-$U_{F(h)}$ scales the same as $U_{F(h)}$, up to a constant factor overhead. We can extract $F(h)_{ij}$ with an accuracy specified in the next Lemma 5.4.1. This Lemma is stated for a general block-encoding unitary and is proved in Appendix 5.G. From Lemma 5.4.1 it is clear that the error up to which $F(h)_{ij}$ can be estimated is $1/\text{poly}(n)$, since we allow for at most poly(n) calls to the block-encoding unitaries.

Lemma 5.4.1. *Given an n-qubit matrix A. Let $C\text{-}U_A$ (acting on $n+m+1$ qubits) denote the controlled version of the (α, m, ε)-block-encoding U_A of A. An estimate $\hat{A}_{ij}$ of entry A_{ij} can be obtained s.t. $\big|\hat{A}_{ij} - A_{ij}\big| \leq \varepsilon + \alpha\tilde{\varepsilon}$ with probability at least $1-\delta$, using poly(n)-sized circuits and at most $D(\tilde{\varepsilon},\delta) = \Theta\big(\tilde{\varepsilon}^{-2}\log\big(4\delta^{-1}\big)\big)$ calls to $C\text{-}U_A$.*

Combining Lemma 5.4.1 with Theorems 5.3.9, 5.3.10 and 5.3.11, we can respectively estimate entries of $M^{(\beta)}$ in equation (5.3), $G^{\delta,\beta,\omega}$ in equation (5.7), and $M(t)$ in equation (5.4), up to $1/\text{poly}(n)$ error with poly(n) effort. Note that — asymptotically — the circuit implementing the controlled block-encodings (which is required for the Hadamard test in Lemma 5.4.1) is of the same size as the block-encoding circuits themselves.

We note that in case when $F(h) = M$ is a correlation matrix and H corresponds to a lattice model, one can also obtain correlation matrix entries in momentum space — by using U_M and the efficient Quantum Fourier Transform circuit [31].

Going beyond individual matrix elements, for any local fermionic Hamiltonian term H_x in H, for example $H_x = \left(h_{ij}a_j^\dagger a_i + h_{ij}^* a_i^\dagger a_j\right)$ (with $|h_{ij}| \leq 1$) or $H_x = \left(V_{ijkl}a_i^\dagger a_j^\dagger a_k a_l + V_{ijkl}^* a_l^\dagger a_k^\dagger a_j a_i\right)$ (with $|V_{ijkl}| \leq 1$), the expectation of that term w.r.t. a state ρ can be efficiently extracted from the block-encoding of its correlation matrix U_M [2]. In this way one can also obtain the total energy density of ρ relative to a system Hamiltonian H. To do so, one can sample from the Hamiltonian terms uniformly at random and evaluate the expectation value of individual terms as mentioned above. For H being a free-fermion Hamiltonian, this sampling can be implemented using the sparse

[2] From this point onwards, all considered states are free-fermionic, unless stated otherwise.

access model discussed below; this method of sampling can be extended to interacting Hamiltonians. We can obtain the following concentration bound on this evaluated energy density e, assuming, for simplicity, that the expectation of an individual term is learned from U_M without error. By assumption, we have that $\left|\mathrm{Tr}(H_x\rho)\right| \leq 1$ for each Hamiltonian term H_x. This allows us to infer the Chernoff bound, which says that for sample size $S = \Theta\left(\varepsilon^{-2}\log\left(\delta^{-1}\right)\right)$, we have

$$\mathbb{P}\left(\left|e - \mathrm{Tr}\left(H\rho\right)/K\right| \leq \varepsilon\right) \geq 1-\delta, \tag{5.21}$$

where $K = \Theta(2^n)$ is the number of terms in the Hamiltonian H. Similarly, densities of other Hermitian operators can be learned through sampling, such as the particle density $\langle\hat{N}\rangle/2^n = \mathrm{Tr}(M)/2^n$.

5.5. Complexity

We have presented a method for simulating free-fermionic systems on $N = 2^n$ modes with polynomial resources in n in a variety of settings. The naive classical treatment of 2^n fermionic modes, on the other hand, requires exponential resources. Therefore, the naive speedup of our quantum method is exponential. However, our approach comes with manifest qualifications, namely the requirement for the oracle tuple $\mathcal{O}_h$ to be implementable using poly(n)-sized quantum circuits, time dynamics being simulable only for time $t = \mathrm{poly}\,(n)$, thermal states for $\beta = \mathrm{poly}\,(n)$ and Green's functions for $\beta, 1/\delta = \mathrm{poly}\,(n)$. Competing classical approaches could hypothetically exploit this structure of our setting. To settle this issue, one can readily argue that our method generally yields an exponential quantum speedup, by showing that it solves a BQP-complete problem. Roughly speaking, BQP-complete problems are the hardest problems which can be efficiently solved by a quantum computer [32]. Since for single-particle dynamics, the character of the particle, –be it a boson, fermion or distinguishable particle– is not relevant, BQP-hardness of time-dynamics follows in principle from Theorem 3 in [17], using techniques such as those developed in Ref. [33]. For completeness, we provide a slightly different proof for the complexity of the evolution of a multi-particle fermionic state in Appendix 5.H.

Theorem 5.5.1. *Let ρ_0 be a (multi-particle) fermionic state on 2^n modes, such that its correlation matrix M_0 is sparse, and the access oracle tuple $\mathcal{O}_{M_0}$ can be implemented as a poly(n)-sized quantum circuit. Given a quadratic Hamiltonian H on 2^n modes, let h be as in equation* (5.1) *and sparse, and we assume that the oracle tuple $\mathcal{O}_h$ is implemented as a* poly(n)*-sized quantum circuit. For $t = \mathrm{poly}(n)$, the problem is to decide whether, for some given mode j, $n_j(t) = \mathrm{Tr}\,(a_j^\dagger a_j e^{-iHt}\rho_0 e^{iHt}) \geq 1/p(\sqrt{n})$ (with p a polynomial) or $\leq \exp\left(-\sqrt{n}\right)$, given a promise that either one is the case. This problem is BQP-complete.*

5.6. Quantum speed-up in a variety of settings

We have established that our algorithms in principle provide an exponential speed-up, at least in the setting of time evolution. In this section, we argue what the speed-up is for several models of physical importance. To that end, let us first argue that for

$d = O(1)$-dimensional lattice models, entries of our matrix functions of interest (see Section 5.2) can be estimated *classically* with poly(n) effort for $\beta, 1/\delta, t = \text{poly}(n)$.

Lieb-Robinson bounds [34–36] imply that the time evolution of observables such as the occupation number of a mode i at some position (starting from a product state with some modes occupied and others unoccupied) is only affected by $O(t^d) = \text{poly}(n)$ sites in a ball of radius proportional to t around that position. Similarly, Ref. [34] shows that, for a given mode i, the thermal correlation matrix entries $|M^{(\beta)}_{ij}|$ decay exponentially with the distance between modes i and j, with a characteristic length $O(\beta)$. Mode i is therefore only non-trivially correlated with $O(\beta^d) = \text{poly}(n)$ modes in a ball of radius $O(\beta)$ around it. This latter fact suggests that an entry $M^{(\beta)}_{ij}$ can be classically evaluated with poly(n) effort, provided that $\beta = \text{poly}(n)$. Let us formalize this as follows.

Lemma 5.6.1. *Let $h \in \mathbb{C}^{2^n \times 2^n}$ be an $s = O(1)$-sparse matrix that corresponds to a $d = O(1)$-dimensional lattice model, cf. equation* (5.1) *with entries as in equation* (5.13) *and equation* (5.14)*. Assume* poly(n)*-effort classical access to the oracles O_r and O_a (see Definition 5.3.2) for h. Let $F(h)$ be a matrix function of h. If $\max_{x\in[-1,+1]} |F(x) - p_K(x)| \leq \text{poly}(n)/K$ with $p_K(x)$ a degree-K polynomial, then an entry $F(h)_{ij}$ can be estimated with that same error using* poly(K) × poly(n) *classical effort. For some $K = \text{poly}(n)$, the error thus becomes* $1/\text{poly}(n)$ *with* poly(n) *classical effort.*

Proof. If one is able to estimate $\langle i|\, h^k \,|j\rangle$ for any $k \in \{0, 1, \ldots, K\}$ with effort E, then $\langle i|\, p_K(h) \,|j\rangle = \sum_{k=0}^{K} \alpha_k \langle i|\, h^k \,|j\rangle$ can be evaluated with effort $K \times E$. By assumption, $\langle i|\, F(h) \,|j\rangle$ can then be classically approximated up to $\text{poly}(n)/K$ error with $K \times E$ effort. Since h corresponds to a $d = O(1)$-dimensional lattice model, $h^k \,|j\rangle$ is only supported on $O(k^d) = \text{poly}(k)$ $|i\rangle$'s. We can thus evaluate each $\langle i|\, h^k \,|j\rangle$ for $k \in \{0, 1, \ldots, K\}$ using poly(k) calls to the oracles and with a total $E = \text{poly}(k) \times \text{poly}(n)$ computational effort. Therefore, $\langle i|\, F(h) \,|j\rangle$ can be approximated classically with $\text{poly}(n)/K$ error with $K \times \text{poly}(k) \times \text{poly}(n) = \text{poly}(K) \times \text{poly}(n)$ effort. Clearly, there is a $K = \text{poly}(n)$ so that the error becomes $1/\text{poly}(n)$ and which yields a poly(n) classical effort. □

Combined with Lemmas 5.3.7 and 5.3.8, Lemma 5.6.1 implies the following for $d = O(1)$-dimensional lattice models. In the parameter regimes of Theorems 5.3.9 and 5.3.10, entries of the thermal correlation matrix in equation (5.3) and of the thermal Green's function in equation (5.7) can be estimated up to $1/\text{poly}(n)$ error with poly(n) classical effort.

Using similar reasoning, entries of the time-evolved correlation matrix $M(t)$ in equation (5.4) can be evaluated classically with poly(n) effort for $t = \text{poly}(n)$. In fact, assuming exact classical access to entries $\langle k|\, M \,|l\rangle$ of an initial correlation matrix M for given (k, l), one can obtain entries $M(t)_{ij}$ with $1/\exp(n)$ error. The improved error scaling comes from the fact that the polynomial approximation error of $\exp(ith)$ can be bounded by $1/\exp(n)$ even for degree $K = \text{poly}(n)$, provided that $t = \text{poly}(n)$. A detailed treatment is given in Appendix 5.I. Note that if we apply the time evolution to $M'^{(\beta)}$ (where $M'^{(\beta)}$ is the thermal correlation matrix corresponding to some $h' \neq h$), the accuracy reduces to $1/\text{poly}(n)$ due to the error in estimating entries of $M'^{(\beta)}$.

Despite losing the exponential speed-up for $d = O(1)$-dimensional lattice models, let us argue that we retain a power-$(d+1)$ polynomial speed-up for such models. Let us

	$d = O(1)$-dim. lattice models	expander graphs	general sparse models
Classical algorithms *	$r_{\text{prep},C} + N \cdot t$		
Quantum algorithms	$r_{\text{prep},Q} + \text{polylog}(N) \cdot t$		
Lieb-Robinson time	$N^{1/d}$	$\log(N)$	-
Speedup	power-$(d+1)$ polynomial **	exponential **	exponential ***

Table 5.1: Asymptotic run-times for evaluating entries of time-evolved correlation matrices (with $1/\text{poly}(n)$ error) for three different system types: lattice models, expander graphs and general sparse models. For the former two, we start from a thermal correlation matrix at $\beta = O(1)$ (of some h', different from h used for time evolution). For the latter, we start from a Slater determinant (free fermion pure) state. The third row gives the Lieb-Robinson time (only denoted for lattice models and expander graphs), which corresponds to the time it takes the Lieb-Robinson light cone to contain the entire system. The run-times of classical algorithms (for evolutions over a time interval which is at least the Lieb-Robinson time) and our quantum algorithms are given. In addition, we provide the associated speedups for the lattice models and expander graphs at the Lieb-Robinson time, and the speedup for general sparse models at $t = \text{poly}(n)$. The run-times required to prepare the starting state are denoted by $r_{\text{prep},C}$ and $r_{\text{prep},Q}$, for respectively the classical and quantum algorithms. Note that $r_{\text{prep},Q} = \text{poly}(n)$ in all three scenarios. *Run-times of —to the best of our knowledge— the best classical algorithm for these applications [37]. **Speedups compared to the aforementioned classical algorithms. ***Speedup assuming that it takes exponential (in n) time to solve BQP-complete problems classically.

focus on the task of estimating entries of the time-evolved correlation matrix from equation (5.4). In particular, let us focus on the task of time-evolution for t proportional to the Lieb-Robinson time t_{LR}, which is the time it takes for a Lieb-Robinson light cone to contain the entire system. For lattice models, $t_{LR} = N^{1/d}$. To then compute an entry of the correlation matrix $M(t) = e^{iht} M e^{-iht}$, known classical algorithms require $\Omega(Nt) = \Omega(t^{d+1})$ run-time [37]. Given the $\text{polylog}(N) \cdot t$ runtime of our quantum algorithm, we obtain a power-$d+1$ polynomial speedup. In particular, this yields a cubic speedup for $d = 2$ lattices and quartic speedup for $d = 3$ —which can be of interest in early fault-tolerant devices [38].

Crucially, our method can also be applied to settings other than lattice models, and the exponential speedup for those settings can be maintained. Let us consider tight-binding models on expander graphs, such as the Margulis graph considered in Section 5.3.1. The Lieb-Robinson time, due to the expansion property of the graph, will be logarithmic in the number of modes N: $t_{LR} = \log(N)$. We note that light cones also grow rapidly

in other graphs with *log-sized diameter*, such as the hyperbolic lattices (see [39] for recent studies of such tight-binding models). We expect to recover the full exponential quantum speedup for their simulation because at $t_{LR} = \log(N)$, the quantum run-time is $\mathrm{poly}\log(N)$ while known classical algorithms have run-time $\Omega(Nt) = \Omega(N\log(N))$ [37]. This speed-up can be of particular interest for, e.g., the study of Anderson localization on expander graphs [23–26].

To summarize the quantum advantage in different problem settings, Table 5.1 gives an overview of the asymptotic run-times of classical algorithms and our quantum algorithms, and associated quantum speedups for the problem of time-evolution.

5.7. Generalizations

The time-evolution framework presented in this paper can be made more general and applied to systems beyond free fermions. In a general quantum system described by a Hamiltonian H, one can consider a N-sized set of operators $\{O_j\}$ such that $[H, O_k] = \sum_{j=1}^{N} h_{jk} O_j$. This is sufficient for a matrix $M_{jk} = \mathrm{Tr}\big(\rho O_j^\dagger O_k\big)$ to transform as $M \mapsto e^{-iht} M e^{iht}$ under time evolution. Further assuming that h is a hermitian matrix, this allows treatment of M as a block-encoding of the type considered in this work. Beyond the free-fermionic systems on which we focused in this work, this general framework admits fermionic H which include pairing ($\Delta a_j a_k + h.c.$) terms. In this case the relevant set $\{O_j\}$ would include not just annihilation but also creation operators. Another example is a system of 2^n free bosons with particle conservation, in which case $\{O_j\}$ should be chosen as bosonic annihilation operators. Beyond $\mathrm{Tr}\big(\rho O_j^\dagger O_k\big)$, one can consider $M_{j_1,..,j_l;k_1,..,k_{l'}} = \mathrm{Tr}\big(\rho O_{j_1}^\dagger .. O_{j_l}^\dagger O_{k_1} .. O_{k_{l'}}\big)$, which can be considered as a rectangular matrix acting on $n \cdot \max(l, l')$ qubits, and block-encoded accordingly. The time evolution of these objects is defined similarly to that of M_{jk}, and therefore can be easily found as a block-encoding, given the block-encoding of the initial state. The flexibility of this general block-encoding framework is comparable to the one based on 'shadow' states, presented in Ref. [19] (see Appendix 5.A for a discussion of the differences).

5.8. Discussion

In this work, we develop quantum algorithms that solve several free fermion problems. We discuss in detail what type of speedup is achieved over classical algorithms and present generalizations of our approach.

One obvious avenue for future research is to apply our method to other matrix functions of h. For example, one should be able to estimate the free energy density of a 2^n-mode free-fermion system $\frac{F}{2^n} = -(\beta 2^n)^{-1} \log \mathrm{Tr}\,(e^{-\beta H}) = -(\beta 2^n)^{-1} \mathrm{Tr}\,(\log\big(I + e^{-\beta h}\big))$ with error ε, using a polynomial approximation of the function $\log\big(I + e^{-\beta h}\big)$ for $\beta = \mathrm{poly}(n)$, the block-encoding of h, and sampling entries to model the trace function. Using an estimate of the free energy density $F/2^n = (\langle H \rangle_\beta - \beta^{-1} S(\rho_\beta))/2^n$, one can in turn estimate an entropy density, given an energy density estimate, or a derivative of $F/2^n$ with respect to β such as the specific heat. Another possible generalization of our

work is a poly(n)-efficient estimation of matrix elements or observable expectations due to free-fermionic *dissipative* dynamics, which was shown to be classically simulatable in $O(2^{3n})$ time in [40].

One could also consider how block-encoding techniques fare when applied to estimating entries of a free-bosonic thermal correlation matrix $M^{(\beta)} = I/(e^{\beta h} - I)$ of Bose-Einstein form. A block-encoding of a polynomial approximation as developed in Lemma 5.3.7 and Theorem 5.3.9 in Section 5.3.2 requires a poly(n) bound on the mode occupation number (so that the matrix function be block-encoded), which can however grow as large as the number of particles for a Bose-Einstein condensate. Mathematically, the Bose-Einstein distribution with $\epsilon_i \geq 0$ has a singularity at $\epsilon_i = 0$ which has to be avoided (by choosing a small enough chemical potential μ) in order to place any bound. Note that similar points about only algebraic speed-ups for local lattice models (Lemma 5.6.1) were made for bosonic/oscillator systems in a more recent work [41].

Another outstanding open direction is to compute and optimize the precise implementation overhead and circuit depth for our proposed algorithms, as applied to simulation problems of practical interest.

Let us point out an open question in the setting of time-dynamics on 2^n fermionic modes (cf. equation (5.4)). One task that can be performed with poly(2^n) classical effort [2] is computing the overlap

$$|\langle S_1|\exp(-itH)|S_2\rangle|^2 = \mathrm{Tr}\big[\underbrace{\exp(-itH)|S_2\rangle\langle S_2|\exp(itH)}_{|S_3\rangle\langle S_3|}|S_1\rangle\langle S_1|\big], \tag{5.22}$$

with $|S_1\rangle$ and $|S_2\rangle$ single-Slater determinant states and H a free fermion Hamiltonian as in equation (5.1), and therefore $|S_3\rangle$ is also a Slater determinant state. If $|S_1\rangle\langle S_1|$ (for simplicity) is a standard-mode-basis Slater determinant state, then it can be expressed as a product of 2^{n+1} creation and annihilation operators. Using Wick's theorem, evaluating this weight-2^{n+1} correlator in equation (5.22) requires evaluating products of 2^n entries of the correlation matrix (cf. equation (5.2)) of state $|S_3\rangle$. This task – at least with naive attempts – cannot be performed using our methods with poly(n) quantum effort, since we can only evaluate poly(n) entries of the time-evolved correlation matrix, although approximate sampling methods could come into play.

References

[1] S. Lloyd. 'Universal Quantum Simulators'. In: *Science* 273.5278 (1996), pp. 1073–1078. DOI: 10.1126/science.273.5278.1073.

[2] B. M. Terhal and D. P. DiVincenzo. 'Classical simulation of noninteracting-fermion quantum circuits'. In: *Phys. Rev. A* 65 (3 Mar. 2002), p. 032325. DOI: 10.1103/PhysRevA.65.032325.

[3] E. Knill. *Fermionic Linear Optics and Matchgates*. 2001. arXiv: quant-ph/0108033 [quant-ph]. URL: https://arxiv.org/abs/quant-ph/0108033.

[4] S. Bravyi. 'Lagrangian representation for fermionic linear optics'. In: *Quantum Info. Comput.* 5.3 (May 2005), pp. 216–238. ISSN: 1533-7146.

[5] C. W. Groth, M. Wimmer, A. R. Akhmerov and X. Waintal. 'Kwant: a software package for quantum transport'. In: *New Journal of Physics* 16.6 (June 2014), p. 063065. DOI: 10.1088/1367-2630/16/6/063065.

[6] T. Kloss, J. Weston, B. Gaury, B. Rossignol, C. Groth and X. Waintal. 'Tkwant: a software package for time-dependent quantum transport'. In: *New Journal of Physics* 23.2 (Feb. 2021), p. 023025. DOI: 10.1088/1367-2630/abddf7.

[7] R. Jozsa, B. Kraus, A. Miyake and J. Watrous. 'Matchgate and space-bounded quantum computations are equivalent'. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 466.2115 (2009), pp. 809–830. DOI: 10.1098/rspa.2009.0433.

[8] B. Kraus. 'Compressed Quantum Simulation of the Ising Model'. In: *Physical Review Letters* 107.25 (Dec. 2011). ISSN: 1079-7114. DOI: 10.1103/physrevlett.107.250503.

[9] G. Blázquez-Cruz and P.-L. Dallaire-Demers. *Quantum supremacy regime for compressed fermionic models*. 2022. arXiv: 2110.09550 [quant-ph]. URL: https://arxiv.org/abs/2110.09550.

[10] B. Senjean, S. Yalouz and M. Saubanère. 'Toward density functional theory on quantum computers?' In: *SciPost Physics* 14.3 (2023), p. 055.

[11] A. Barthe, M. Cerezo, A. T. Sornborger, M. Larocca and D. García-Martín. *Gate-based quantum simulation of Gaussian bosonic circuits on exponentially many modes*. 2024. arXiv: 2407.06290 [quant-ph]. URL: https://arxiv.org/abs/2407.06290.

[12] D. W. Berry, A. M. Childs and R. Kothari. 'Hamiltonian simulation with nearly optimal dependence on all parameters'. In: *2015 IEEE 56th annual symposium on foundations of computer science*. IEEE. 2015, pp. 792–809.

[13] G. H. Low and I. L. Chuang. 'Hamiltonian Simulation by Qubitization'. In: *Quantum* 3 (July 2019), p. 163. ISSN: 2521-327X. DOI: 10.22331/q-2019-07-12-163.

[14] A. Gilyén, Y. Su, G. H. Low and N. Wiebe. 'Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics'. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. STOC '19. ACM, June 2019. DOI: 10.1145/3313276.3316366.

[15] L. Lin. *Lecture Notes on Quantum Algorithms for Scientific Computation*. arXiv.org:2201.08309. 2022. URL: https://math.berkeley.edu/~linlin/qasc/.

[16] P. Rall. 'Quantum algorithms for estimating physical quantities using block encodings'. In: *Physical Review A* 102.2 (Aug. 2020). ISSN: 2469-9934. DOI: 10.1103/physreva.102.022408.

[17] R. Babbush, D. W. Berry, R. Kothari, R. D. Somma and N. Wiebe. 'Exponential Quantum Speedup in Simulating Coupled Classical Oscillators'. In: *Phys. Rev. X* 13 (4 Dec. 2023), p. 041041. DOI: 10.1103/PhysRevX.13.041041.

[18] S. Danz, M. Berta, S. Schröder, P. Kienast, F. K. Wilhelm and A. Ciani. *Calculating response functions of coupled oscillators using quantum phase estimation*. 2024. arXiv: 2405.08694 [quant-ph]. URL: https://arxiv.org/abs/2405.08694.

[19] R. D. Somma, R. King, R. Kothari, T. O'Brien and R. Babbush. *Shadow Hamiltonian Simulation*. 2024. arXiv: 2407.21775 [quant-ph]. URL: https://arxiv.org/abs/2407.21775.

[20] A. Montanaro and S. Stanisic. *Compressed variational quantum eigensolver for the Fermi-Hubbard model*. 2020. arXiv: 2006.01179 [quant-ph]. URL: https://arxiv.org/abs/2006.01179.

[21] A. Altland and B. D. Simons. *Condensed Matter Field Theory*. 2nd ed. Cambridge University Press, Apr. 2010. ISBN: 0-521-76975-2.

[22] A. Gilyén, Y. Su, G. H. Low and N. Wiebe. *Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics*. 2018. arXiv: 1806.01838 [quant-ph]. URL: https://arxiv.org/abs/1806.01838.

[23] K. S. Tikhonov, A. D. Mirlin and M. A. Skvortsov. 'Anderson localization and ergodicity on random regular graphs'. In: *Physical Review B* 94.22 (2016), p. 220203.

[24] C. Vanoni, B. L. Altshuler, V. E. Kravtsov and A. Scardicchio. 'Renormalization group analysis of the Anderson model on random regular graphs'. In: *Proceedings of the National Academy of Sciences* 121.29 (2024). DOI: 10.1073/pnas.2401955121.

[25] I. García-Mata, O. Giraud, B. Georgeot, J. Martin, R. Dubertrand and G. Lemarié. 'Scaling Theory of the Anderson Transition in Random Graphs: Ergodicity and Universality'. In: *Phys. Rev. Lett.* 118 (16 Apr. 2017), p. 166801. DOI: 10.1103/PhysRevLett.118.166801.

[26] I. García-Mata, J. Martin, R. Dubertrand, O. Giraud, B. Georgeot and G. Lemarié. 'Two critical localization lengths in the Anderson transition on random graphs'. In: *Phys. Rev. Res.* 2 (1 Jan. 2020), p. 012020. DOI: 10.1103/PhysRevResearch.2.012020.

[27] O. Goldreich, S. Goldwasser and S. Micali. 'How to construct random functions'. In: *Journal of the ACM (JACM)* 33.4 (1986), pp. 792–807.

[28] A. Banerjee, C. Peikert and A. Rosen. 'Pseudorandom functions and lattices'. In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2012, pp. 719–737.

[29] J. Chen and G. K.-L. Chan. *Quantum linear algebra for disordered electrons*. 2025. arXiv: 2411.00434 [quant-ph]. URL: https://arxiv.org/abs/2411.00434.

[30] L. N. Trefethen. 'Chapter 8. Convergence for Analytic Functions'. In: *Approximation Theory and Approximation Practice, Extended Edition*. SIAM, 2013, pp. 55–62. DOI: 10.1137/1.9781611975949.ch8.

[31] M. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge Series on Information and the Natural Sciences. Cambridge, U.K.: Cambridge University Press, 2000. ISBN: 978-0-521-63503-5.

[32] E. Bernstein and U. Vazirani. 'Quantum Complexity Theory'. In: *SIAM Journal on Computing* 26.5 (1997), pp. 1411–1473. DOI: 10.1137/S0097539796300921.

[33] D. Nagaj. *Local Hamiltonians in Quantum Computation*. PhD thesis MIT. 2008. arXiv: 0808.2117 [quant-ph]. URL: https://arxiv.org/abs/0808.2117.

[34] M. B. Hastings. 'Decay of Correlations in Fermi Systems at Nonzero Temperature'. In: *Phys. Rev. Lett.* 93 (12 Sept. 2004), p. 126402. DOI: 10.1103/PhysRevLett.93.126402.

[35] C.-F. (Chen, A. Lucas and C. Yin. 'Speed limits and locality in many-body quantum dynamics'. In: *Reports on Progress in Physics* 86.11 (Sept. 2023), p. 116001. DOI: 10.1088/1361-6633/acfaae.

[36] M. C. Tran, C.-F. Chen, A. Ehrenberg, A. Y. Guo, A. Deshpande, Y. Hong, Z.-X. Gong, A. V. Gorshkov and A. Lucas. 'Hierarchy of Linear Light Cones with Long-Range Interactions'. In: *Phys. Rev. X* 10 (3 July 2020), p. 031009. DOI: 10.1103/PhysRevX.10.031009.

[37] P. C. S. Costa, S. Jordan and A. Ostrander. 'Quantum algorithm for simulating the wave equation'. In: *Physical Review A* 99.1 (Jan. 2019). ISSN: 2469-9934. DOI: 10.1103/physreva.99.012323.

[38] R. Babbush, J. R. McClean, M. Newman, C. Gidney, S. Boixo and H. Neven. 'Focus beyond Quadratic Speedups for Error-Corrected Quantum Advantage'. In: *PRX Quantum* 2 (1 Mar. 2021), p. 010103. DOI: 10.1103/PRXQuantum.2.010103.

[39] A. Kollár, M. Fitzpatrick, P. Sarnak and A. Houck. 'Line-Graph Lattices: Euclidean and Non-Euclidean Flat Bands, and Implementations in Circuit Quantum Electrodynamics'. In: *Commun. Math. Phys.* 376 (2020), pp. 1909–1956. DOI: 10.1007/s00220-019-03645-8.

[40] S. Bravyi and R. Koenig. 'Classical simulation of dissipative fermionic linear optics'. In: *Quantum Inf. Comput.* 12.11&12 (2012), pp. 925–943.

[41] K. Sakamoto and K. Fujii. *On the quantum computational complexity of classical linear dynamics with geometrically local interactions: Dequantization and universality*. 2025. arXiv: 2505.10445 [quant-ph]. URL: https://arxiv.org/abs/2505.10445.

[42] K. Temme, T. J. Osborne, K. G. Vollbrecht, D. Poulin and F. Verstraete. 'Quantum Metropolis sampling'. In: *Nature* 471.7336 (Mar. 2011), pp. 87–90. ISSN: 1476-4687. DOI: 10.1038/nature09770.

[43] J. Jiang and S. Irani. *Quantum Metropolis Sampling via Weak Measurement*. 2024. arXiv: 2406.16023 [quant-ph]. URL: https://arxiv.org/abs/2406.16023.

[44] H.-Y. Huang, R. Kueng and J. Preskill. 'Predicting Many Properties of a Quantum System from Very Few Measurements'. In: *Nature Physics* 16 (Oct. 2020). DOI: 10.1038/s41567-020-0932-7.

[45] A. Peres. 'Reversible logic and quantum computers'. In: *Phys. Rev. A* 32 (6 Dec. 1985), pp. 3266–3276. DOI: 10.1103/PhysRevA.32.3266.

[46] S. Boixo, E. Knill and R. Somma. 'Eigenpath traversal by phase randomization'. In: *Quantum Info. Comput.* 9.9 (Sept. 2009), pp. 833–855. ISSN: 1533-7146.

5.A. Alternative Encodings

In this section we describe alternative ways of representing a fermionic correlation matrix using qubits and their potential drawbacks.

A compressed representation of free-fermionic states on 2^n modes, as well as their dynamics, is readily obtained by using a (mixed) quantum state $\sigma = M/\mathrm{Tr}(M)$ of n qubits to represent the normalized correlation matrix of ρ. One then computes, —evolves and measures—, with σ to learn properties of ρ or its time-dynamics. For pure single-particle free-fermionic states ρ, σ is a rank-1 projector, and σ projects onto the bitstring $|i\rangle$ when ρ corresponds to $a_i^\dagger |\mathrm{vac}\rangle$, $i = 0, \ldots, N-1$ where $|\mathrm{vac}\rangle$ is the fermionic vacuum state. Once a state σ is prepared, its time-evolution can readily be simulated: when ρ evolves via e^{-iHt} with free-fermion Hamiltonian H, $\sigma \to e^{iht}\sigma e^{-iht}$. Sparse oracle access to h —see Definition 5.3.2— then allows for the efficient implementation of time-evolution in terms of its dependence on t and calls to the oracle [13, 15], starting from some easy-to-prepare initial state. For example, the initial state could be a set of fermions in a subset S of 2^m modes $|i\rangle$ (such that an efficient classical circuit can map S onto the set of m-bitstrings), or a subset of modes in the Fourier-transformed basis (as the QFT is an efficient quantum circuit). One can also adapt the heuristic quantum Metropolis-Hastings algorithm [42, 43] to the Fermi-Dirac distribution and sparse Hamiltonians h, since the algorithm uses quantum phase estimation for e^{iht} at its core. Even though the algorithm converges to the thermal state $\sigma_\beta = M^{(\beta)}/\mathrm{Tr}(M^{(\beta)})$, poly($n$)-efficiency is not guaranteed and unlikely for low-enough temperature.

Given a state σ, one can apply any learning algorithm for n-qubit states. For example, one can use shadow tomography [44] to estimate the expectation of L observables, such as $O_k = |k\rangle\langle k|, O^R_{lk} = |l\rangle\langle k| + |k\rangle\langle l|, O^L_{lk} = i(|l\rangle\langle k| - |k\rangle\langle l|)$, with computational effort $O(\log(L))$ using random Clifford circuits of poly(n) size.

There are a few disadvantages to this simple and direct method of representing the state via its correlation matrix. It is not immediately obvious how to estimate a time-dependent correlation function as in equation (5.6) as it relates to measurements on $e^{iht_1}\sigma e^{-iht_2}$ which is not a state. Second, and more crucially, any learning of a linear function of σ with accuracy ε, leads to learning with accuracy $\varepsilon\,\mathrm{Tr}(M) = \varepsilon\langle\hat{N}\rangle$ for the correlation matrix M itself. Therefore one expects poor accuracy for large particle number $\langle\hat{N}\rangle$; this in particular makes it impractical to extract individual matrix elements.

Thus in the main text of this paper we choose not to directly encode a correlation matrix as a quantum state, but rather apply quantum computational block-encoding techniques.

Recently, Ref. [19] introduced a general quantum simulation framework with compressed 'shadow' quantum states with applications to free bosons and free fermion systems. We note that the results in Ref. [19] use yet a different encoding than the encoding described above, or the block-encoding in the main text. Like for the encoding in the previous paragraph, the normalization of the shadow state in Ref. [19] can lead to a loss of efficiency if one wishes to estimate only few entries of the correlation matrix (this loss of efficiency is avoided in our block-encoding method). In particular, the normalization of the shadow state is a, which is bounded as $\sqrt{\sum_j(\langle\hat{N}_j\rangle - 1/2)^2} \le a \le \exp(n)$, where $\langle\hat{N}_j\rangle$ is the occupation number in the mode j of the represented state ρ. On the other hand, when estimating densities, for example the

energy density, our methods use sampling to estimate $\mathrm{Tr}(H\rho)/K$ (with $K = \Theta(2^n)$, the number of terms in H) with some error ε, while Ref. [19] estimates $\mathrm{Tr}(H\rho)/O(2^{n/2}a)$, which, depending on the value a, can be more efficient.

The precise relation between the shadow state approach [19] and the block-encoding framework presented in this work is currently unclear. A plausible hypothesis is that the latter is strictly more powerful, due to the signal strength difference discussed above. A concrete interesting question is whether a shadow state corresponding to $M_{jk} = \mathrm{Tr}\left(\rho a_j^\dagger a_k\right)$ (or, more generally, $\mathrm{Tr}\left(\rho O_j^\dagger O_k\right)$) can always be produced using a block-encoding U_M of M. In the 'typical' case $\mathrm{Tr}\left(M^\dagger M\right) = \Theta(2^n)$, this can be done simply by acting with U_M on the maximally entangled state between j and k registers, and postselecting on the zero value of ancillary qubits. This 'Choi–Jamiołkowski' strategy, however, does not give a constant success rate when $\mathrm{Tr}\left(M^\dagger M\right) = o(2^n)$, and should be adapted.

5.B. Remarks on oracle conventions

In this work, we define row and matrix entry oracles as in Definition 5.3.2. An alternative definition of a row oracle, used in, for instance, Ref. [14], is

$$O_r^{\mathrm{alt}} |i\rangle |k, 0^{(n+1)-\lceil \log(s) \rceil}\rangle = |i\rangle |r(i,k)\rangle, \quad \forall i \in [2^n], k \in [s], \tag{5.23}$$

with O_r^{alt} acting on $2(n+1)$ qubits. Again, if row i contains $s' < s$ non-zero entries, then the last $n+1$ qubits are set to $|1\rangle |k\rangle$. We note that having access to O_r in equation (5.11) implies access O_r^{alt} and vice versa.

In Ref. [14], O_r^{alt} and O_a are used to block-encode a sparse matrix h. In principle, this block-encoding scheme requires another (column) oracle O_c^{alt} when it is used to block-encode *general* sparse matrices h. If h is also Hermitian, which is the case for all applications considered in this work, this block-encoding can be implemented with just O_r^{alt} and O_a, since O_c^{alt} can be realized using O_r^{alt} and some SWAP gates.

5.C. Margulis Expander Graphs

In the main text, we have provided an example of a d-dimensional model which has sparse query access. Going beyond these models, in this appendix we consider an example of a model on an *expander graph* which has sparse query access. Expander graphs are bounded-degree graphs, which have the so-called *expansion* property. In particular, when counting the vertices away from a given vertex by a distance d, one obtains a number that scales exponentially with d. We will focus on realizing sparse access for a particular simple example, which is the Margulis expander graph.

A Margulis graph $\mathcal{G}_M$ of size N^2 has vertices v labeled by tuples $v = (v_1, v_2) \in [N] \times [N]$; an edge between two vertices u and v is placed if $u = t_l(v)$ where the functions t_l for $l \in [4]$ are defined as $t_0((v_1, v_2)) = (v_1 + 1 \bmod N, v_2)$, $t_1((v_1, v_2)) = (v_1, v_2 + 1 \bmod N)$, $t_2((v_1, v_2)) = (v_1 + v_2 \bmod N, v_2)$, and $t_3((v_1, v_2)) = (v_1, v_2 + v_1 \bmod N)$. In other words, the first two types of edges are simple nearest-neighbour links along the vertical and horizontal directions, with periodic boundary conditions. From this perspective, the edges t_2 and t_3 are geometrically non-local, and are the source of the expansion property

of the graph. We define our tight-binding Hamiltonian on the Margulis graph as follows. Each fermionic mode is labeled by the vertex of the graph, so the total number of modes is N^2. The Hamiltonian takes the form

$$H_{\text{Marg}} = \sum_{l\in[4]} \sum_{v\in[N]\times[N]} \left(a_v^\dagger a_{t_l(v)} + a_{t_l(v)}^\dagger a_v \right). \tag{5.24}$$

For a given v, modular addition circuits allow to efficiently generate a list of $u = t_l^{\pm 1}(v)$. This list can be used to construct an oracle O_r; to ensure distinct outputs, if some of 8 values of u coincide, one stores only one of the colliding outputs. The oracle O_a then represents collisions with an increased matrix element h_{vu}, realized by counting the times u occurs in the list of $t_l^{\pm 1}(v)$. We expect that more models on expander graphs can be implemented in a similar way – especially in the family of constant degree Ramanujan Cayley graphs, of which the Margulis graph is an example.

5.D. Block-encoding the thermal correlation matrix

In this appendix, we prove Theorem 5.3.9 from the main text. In particular, we prove a more detailed version of it, namely Theorem 5.D.2 below. In its proof we use Propositions 5.3.3 and 5.3.4 on the block-encoding of polynomials of sparse matrices, and Proposition 5.3.5 and Lemma 5.D.1 (of which Lemma 5.3.7 is a simplified version) on constructing a polynomial approximation to our desired matrix function $M^{(\beta)}$ in equation (5.3). We will first prove Theorem 5.D.2 and then Lemma 5.D.1.

As was argued in Section 5.3.2 of the main text using Proposition 5.3.5, we wish to construct accurate polynomial approximations of $F(sx)$ for $x \in [-1,+1]$. Let us state Lemma 5.D.1, which will be proved at the end of this section.

Lemma 5.D.1. *For a function $f(x) = \frac{1}{4}\frac{1}{1+exp\beta sx}$ (with $\beta s > 0, x \in [-1,+1]$), one can efficiently construct a polynomial $p_d(x)$ of degree d such that*

$$max_{x\in[-1,+1]}|f(x) - p_d(x)| \leq \begin{cases} \frac{3}{d}\left(\frac{\beta s}{\pi}\right)^4, & if \frac{\beta s}{2\pi} \geq 1, \\ \frac{10}{d}\left(\frac{\beta s}{\pi}\right)^2, & if \frac{\beta s}{2\pi} < 1. \end{cases} \tag{5.25}$$

Theorem 5.D.2. *For an s-sparse Hamiltonian h on n qubits, assume access to the oracle tuple $\mathcal{O}_h$. We denote the controlled $(1, n+5, \varepsilon_{PA} + \varepsilon_{p(h)} + \delta)$-block-encoding of $M^{(\beta)} = \frac{1}{4}\frac{1}{1+\exp(\beta h)}$ by $C\text{-}U_{M^{(\beta)}}$. The implementation of this block-encoding requires*

$$\begin{cases} \Theta\left(\frac{\beta^4 s^4}{\varepsilon_{PA}}\right), & if \frac{\beta s}{2\pi} \geq 1, \\ \Theta\left(\frac{\beta^2 s^2}{\varepsilon_{PA}}\right), & if \frac{\beta s}{2\pi} < 1, \end{cases} \tag{5.26}$$

calls to oracles from the oracle tuple $\mathcal{O}_h$,

$$O\left(sn + n_a + \log^{5/2}(16 s^9 \beta^8 / (\varepsilon_{PA}^2 \varepsilon_{p(h)}^2))\right) \tag{5.27}$$

ancillary qubits, and

$$O\left(n + (n+4)\beta^4 s^4 / \varepsilon_{PA} + \log^{5/2}(16 s^9 \beta^8 / (\varepsilon_{PA}^2 \varepsilon_{p(h)}^2))\right) \tag{5.28}$$

additional one-qubit and two-qubit gates. To implement this block-encoding, an additional classical computing time of $\text{poly}\left(\beta^4 s^4 / \varepsilon_{PA}, \log(1/\delta)\right)$ *is required.*

Proof. It follows from Proposition 5.3.3 (from [14]) that with $O(1)$ calls to the oracle tuple $\mathcal{O}_h$, one can construct a $(s, n+3, \varepsilon_{\mathrm{BE}_h})$-block-encoding U_h of h and its controlled version. For a given $\varepsilon_{\mathrm{BE}_h}$, the required number of ancillary qubits and (additional) elementary gates are given in Proposition 5.3.3.

Let $p_d(x)$ denote the degree-d polynomial approximation of the function $\frac{1}{4}\frac{1}{1+\exp(\beta s x)}$ as in Lemma 5.D.1. It follows from Lemma 5.D.1 that one can efficiently construct p_d such that

$$\| p_d(h/s) - 1/4\, M^{(\beta)} \| \leq \begin{cases} \frac{3}{d}\left(\frac{\beta s}{\pi}\right)^4, & \text{if } \frac{\beta s}{2\pi} \geq 1, \\ \frac{10}{d}\left(\frac{\beta s}{\pi}\right)^2, & \text{if } \frac{\beta s}{2\pi} < 1. \end{cases} \tag{5.29}$$

Taking $d = \Omega\left(\frac{\beta^4 s^4}{\varepsilon_{\mathrm{PA}}}\right)$ if $\frac{\beta s}{2\pi} \geq 1$ and $d = \Omega\left(\frac{\beta^2 s^2}{\varepsilon_{\mathrm{PA}}}\right)$ if $\frac{\beta s}{2\pi} < 1$, we achieve $\| p_d(h/s) - 1/4\, M^{(\beta)} \| \leq \varepsilon_{\mathrm{PA}}$.

For $\varepsilon_{\mathrm{PA}} < \frac{1}{4}$, we note that $|p_d(x)| \leq 1/2$ for $x \in [-1, +1]$. Therefore, we can apply Proposition 5.3.4 (from [14]): A $(1, n+5, 4d\sqrt{\varepsilon_{\mathrm{BE}_h}/s} + \delta)$-block-encoding of $p_d(h/s)$ consists of at most d uses of unitaries U_h, $U_h^\dagger$ or controlled-U_h and $O((n+4)d)$ elementary gates. In addition, it requires a classical computation with run-time as stated in Proposition 5.3.4. We take $\varepsilon_{p(h)} := 4d\sqrt{\varepsilon_{\mathrm{BE}_h}/s}$ so that for a given $\varepsilon_{p(h)}$, we should ensure that $\varepsilon_{\mathrm{BE}_h} = s\varepsilon_{p(h)}^2/(16d^2)$.

Let the $(1, n+5, \varepsilon_{p(h)} + \delta)$-block-encoding of $p_d(h/s)$ be denoted by $U_{p_d(h/s)}$. We can bound how well the block-encoding of $p_d(h/s)$ approximates the block-encoding of $1/4\, M^{(\beta)}$ as

$$\begin{aligned} \varepsilon_{\mathrm{Tot}} = || 1/4\, M^{(\beta)} - \langle 0|^{\otimes a} \otimes \mathbb{1} U_{p_d(h/s)} |0\rangle^{\otimes a} \otimes \mathbb{1} || \leq \\ || 1/4\, M - p_d(h/s) || + \\ || p_d(h/s) - \langle 0|^{\otimes a} \otimes \mathbb{1} U_{p_d(h/s)} |0\rangle^{\otimes a} \otimes \mathbb{1} || \leq \\ \varepsilon_{\mathrm{PA}} + \varepsilon_{p(h)} + \delta. \end{aligned} \tag{5.30}$$

We have thus constructed a $(1, n+5, \varepsilon_{\mathrm{PA}} + \varepsilon_{p(h)} + \delta)$-block-encoding of $1/4\, M^{(\beta)}$. To implement this block-encoding, we require a number of calls to oracles from the tuple $\mathcal{O}_h$, a number of ancillary qubits, and a number of one-qubit and two-qubit gates as in the lemma statement. □

Let us now give the proof of Lemma 5.D.1.

Proof. For the proof of this lemma, we will employ Bernstein's theorem (Lemma 5.3.6) which bounds the error of Chebyshev approximations. Such a Chebyshev approximation of degree d is of the form $p_d(x) = \sum_{k=0}^{d} a_k T_k(x)$, where $T_k(\cos(\theta)) := \cos(k\theta)$ is the degree k Chebyshev polynomial of the first kind. The coefficients a_k can be obtained by evaluating

$$a_k = \frac{2}{\pi} \int_{-1}^{+1} \frac{f(x)\, T_k(x)}{\sqrt{1-x^2}}\, dx, \tag{5.31}$$

with $\frac{2}{\pi}$ replaced by $\frac{1}{\pi}$ for $k = 0$. Each a_k can be evaluated classically with $\mathrm{poly}(\beta s k)$ resources for $f(x)$ in the lemma statement.

Note that the function $f(z = x + iy) = \frac{1}{1+\exp(\beta s z)}$ for $\beta s > 0$ is analytic for $|y| \leq \pi/\beta s$. Hence we can pick the ellipse $E_r = \{\frac{1}{2}(z + z^{-1}) : |z| = r\}$ with $r = \frac{1}{2}\sqrt{(2\pi/\beta s)^2 + 4}$ on which $f(z)$ is analytic, since within this ellipse $|y| \leq \frac{\pi}{2\beta s}$. We have $|f(z)| \leq C = 1$ for $z \in E_r$ since for $|y| \leq \frac{\pi}{2\beta s}$, we have

$$|1 + \exp(\beta s z)| \geq |1 + \exp(\beta s x)\cos(\beta s y)| \geq 1. \tag{5.32}$$

Using Lemma 5.3.6, we can thus bound $\max_{x\in[-1,+1]}|f(x) - p_d(x)|$ as

$$\max_{x\in[-1,+1]}|f(x) - p_d(x)| \leq \frac{2\left((\pi/\beta s)^2 + 1\right)^{-d/2}}{\frac{1}{2}\sqrt{(2\pi/\beta s)^2 + 4} - 1}. \tag{5.33}$$

Let us distinguish between scenario (1) $\beta s \geq 2\pi$ and scenario (2) $\beta s < 2\pi$. For scenario (1), we can bound

$$\frac{1}{2}\sqrt{(2\pi/\beta s)^2 + 4} - 1 \geq \frac{1}{12}(2\pi/\beta s)^2. \tag{5.34}$$

Furthermore, in both scenarios (1) and (2), we have that

$$\left((\pi/\beta s)^2 + 1\right)^{-d/2} \leq 1/\left((\pi/\beta s)^2 d/2 + 1\right) \leq 1/\left((\pi/\beta s)^2 d/2\right). \tag{5.35}$$

Combining these two facts lead to the following bound in scenario (1)

$$\max_{x\in[-1,+1]}|f(x) - p_d(x)| \leq \frac{12}{d}\left(\frac{\beta s}{\pi}\right)^4. \tag{5.36}$$

In scenario (2), we can simply bound the denominator in equation (5.33) by

$$\frac{1}{2}\sqrt{(2\pi/\beta s)^2 + 4} - 1 \geq \frac{1}{2}\sqrt{5} - 1 \geq 1/10. \tag{5.37}$$

Combining this with the upper bound above for the numerator in equation (5.33) (which holds in both scenarios), we obtain the following upper bound in scenario (2).

$$\max_{x\in[-1,+1]}|f(x) - p_d(x)| \leq \frac{40}{d}\left(\frac{\beta s}{\pi}\right)^2. \tag{5.38}$$

□

5.E. Block-encoding the time-evolved correlation matrix

In this appendix, we prove Theorem 5.E.1 below, which is a generalization of Theorem 5.3.11 for block-encoding $M(t_1, t_2)$ in equation (5.5). We will use a result from Ref. [22] on block-encoding $\exp(iht)$ using a block-encoding of h. Note that the error of the block-encoding of $M(t_1, t_2)$ in the theorem statement accounts for potential errors in the block-encoding of the initial correlation matrix as well.

Theorem 5.E.1. *For an s-sparse Hamiltonian h on 2^n fermionic modes, assume access to the oracle tuple $\mathcal{O}_h$. Also assume access to the $(\alpha, m, \varepsilon_M)$-block-encoding U_M of a correlation matrix M of a fermionic state on 2^n modes. The $\big(\alpha, 2n+m+10, \varepsilon+\varepsilon_M\big)$-block-encoding $U_{M(t_1,t_2)}$ of*

$$M(t_1,t_2) = e^{iht_1} M e^{-iht_2}, \tag{5.39}$$

can be produced using

$$D(\alpha,\varepsilon,t_1,t_2) = O\Big(s(|t_1|+|t_2|) + \log(12\alpha(|t_1|+|t_2|)/(|t_1|\varepsilon)) + \log(12\alpha(|t_1|+|t_2|)/(|t_2|\varepsilon))\Big) \tag{5.40}$$

calls to oracles from the tuple $\mathcal{O}_h$, and a single use of the block-encoding U_M. Moreover, one uses $O\big((n+3)(s(|t_1|+|t_2|) + \log(2\alpha(|t_1|+|t_2|)/(|t_1|\varepsilon)) + \log(2\alpha(|t_1|+|t_2|)/(|t_2|\varepsilon)) + D(\alpha,\varepsilon,t_1,t_2)(n+\log^{5/2}(2\alpha s^2(|t_1|+|t_2|)/\varepsilon))\big)$ one-qubit and two-qubit gates, and $O\big(n_a + \log^{5/2}(2\alpha s^2(|t_1|+|t_2|)/\varepsilon)\big)$ ancillary qubits (with n_a denoting the number of bits with which the entries of h are specified).

Proof. A block-encoding $U_{M(t_1,t_2)}$ of $M(t_1,t_2)$ can be constructed using products of block-encodings $U_{\exp(ith)}$ of $\exp(ith)$ (for times t_1 and $-t_2$) and U_M of M (where the latter is a $(\alpha, m, \varepsilon_M)$-block-encoding by assumption).

To construct a block-encoding of $\exp(iht)$, we employ a block-encoding of h. It follows from Proposition 5.3.3 (from [14]) that with $O(1)$ calls to the oracle tuple $\mathcal{O}_h$, one can construct a $(s, n+3, \varepsilon_{\mathrm{BE}_h})$-block-encoding U_h of h and its controlled version. For a given $\varepsilon_{\mathrm{BE}_h}$, the required number of ancillary qubits and (additional) elementary gates are given in Proposition 5.3.3.

Corollary 62 in [22] states that to implement a $(1, n+5, |2t|\varepsilon_{\mathrm{BE}_h})$-block-encoding of $\exp(ith)$, one is required to implement U_h or $U_h^\dagger$ a total of $6s|t| + 9\log\big((6/(|t|\varepsilon_{\mathrm{BE}_h})\big)$ times, and controlled-U_h or controlled-$U_h^\dagger$ three times. In addition, one has to use $O\big((n+3)(s|t| + \log(2/\varepsilon_{\mathrm{BE}_h})\big)$ two-qubit gates and $O(1)$ ancillary qubits. So to implement the $(1, n+5, |2t|\varepsilon_{\mathrm{BE}_h})$-block-encoding of $\exp(ith)$, one is required to call $\mathcal{O}_h$ a total of $O\big(s|t| + \log\big(6/(|t|\varepsilon_{\mathrm{BE}_h})\big)\big)$ times.

Using Lemma 30 in [14], the block-encoding $U_{M(t_1,t_2)}$ of $M(t_1,t_2)$ can be constructed using the product $U_{M(t_1,t_2)} = (\mathbb{1}_{n+5+m} \otimes U_{\exp(iht_1)})(\mathbb{1}_{2n+10} \otimes U_M)((\mathbb{1}_{n+5+m} \otimes U_{\exp(-iht_2)})$, such that $U_{M(t_1,t_2)}$ is a $(\alpha, 2n+m+10, 2\alpha\varepsilon_{\mathrm{BE}_h}(|t_1|+|t_2|) + \varepsilon_M)$-block-encoding. To implement this product, one is thus required to make

$$D(\varepsilon_{\mathrm{BE}_h}, t_1, t_2) = O\big(s(|t_1|+|t_2|) + \log\big(6/(|t_1|\varepsilon_{\mathrm{BE}_h})\big) + \log\big(6/(|t_2|\varepsilon_{\mathrm{BE}_h})\big)\big) \tag{5.41}$$

calls to oracles from the tuple $\mathcal{O}_h$. In addition, one has to use a total of $O\big((n+3)(s(|t_1|+|t_2|) + \log\big(1/(|t_1|\varepsilon_{\mathrm{BE}_h})\big) + \log\big(1/(|t_2|\varepsilon_{\mathrm{BE}_h})\big) + D(\varepsilon_{\mathrm{BE}_h}, t_1, t_2)(n + \log^{5/2}(s^2/\varepsilon_{\mathrm{BE}_h})\big)$ one-qubit and two-qubit gates, and $O\big(n_a + \log^{5/2}(s^2/\varepsilon_{\mathrm{BE}_h})\big)$ ancillary qubits.

We stress that a controlled version $C\text{-}U_{M(t_1,t_2)}$ of the block-encoding of $U_{M(t_1,t_2)}$ can be implemented with equivalent resources. □

5.F. Block-encoding the thermal Green's function

In this appendix, we prove Theorem 5.F.2, which is a more detailed version of Theorem 5.3.10. In its proof we will again use Propositions 5.3.3 and 5.3.4 on the block-encoding of

polynomials of sparse matrices. In addition, we will use Proposition 5.3.5 and Lemma 5.F.1 (of which Lemma 5.3.8 is a simplified version) on constructing a polynomial approximation to our desired matrix function $G^{(\delta,\beta,\omega)}$ in equation (5.7).

As was argued in Section 5.3.2 using Proposition 5.3.5, we would like to construct accurate polynomial approximations of $F(sx)$ for $x \in [-1,+1]$. The function to be approximated for block-encoding $G^{(\delta,\beta,\omega)}$ is

$$g^{(\delta,\beta,\omega)}(x) := \frac{\delta}{8}\Big[\Big(1-\frac{1}{1+\exp\big(\beta sx\big)}\Big)\frac{1}{i\delta-(sx+\omega)}+\Big(\frac{1}{1+\exp\big(\beta sx\big)}\Big)\frac{-1}{i\delta+(sx+\omega)}\Big]. \quad (5.42)$$

Note that $g^{(\delta,\beta,\omega)}(z)$ $(z\in\mathbb{C})$ has poles at $z=\frac{i\delta-\omega}{s}$ and $z=\frac{-i\delta-\omega}{s}$; the regularization parameter δ ensures that these poles lie off the real axis. For convenience, we define the functions $g_1^{(\delta,\omega)}(z)=1/(i\delta-(sz+\omega))$ and $g_2^{(\delta,\omega)}(z)=-1/(i\delta+(sz+\omega))$. Due to the poles, $|g_{1,2}^{(\delta,\omega)}(x)|$ can still grow as $1/\delta$. To be able to apply Proposition 5.3.4 for block-encoding polynomials, we have to ensure that the polynomial that approximates $g^{(\delta,\beta,\omega)}(x)$ has absolute value at most $1/2$ for $x\in[-1,+1]$. That is the reason for including a factor of $\delta/8$ in $g^{(\delta,\beta,\omega)}(x)$ (so that its absolute value is at most $1/4$, and that of its polynomial approximation at most $1/2$ for approximation error less than $1/4$).

Let us first state the following lemma, the proof of which will be provided at the end of this section, which will be used in the proof of Theorem 5.F.2 (and thus Theorem 5.3.10) on the block encoding of the matrix $G^{(\delta,\beta,\omega)}(h)$.

Lemma 5.F.1. *For a function $g^{(\delta,\beta,\omega)}(x)$ as in equation (5.42) (with $\beta,\delta,s>0$ and $x\in[-1,+1]$), one can efficiently construct a polynomial $p_d(x)$ of (even) degree d such that*

$$\max_{x\in[-1,+1]}|g^{(\delta,\beta,\omega)}(x)-p_d(x)| \le \begin{cases}\frac{12}{d}\big(\frac{\beta s}{\pi}\big)^4, & \text{if } \frac{\beta s}{2\pi}\ge 1,\\ \frac{40}{d}\big(\frac{\beta s}{\pi}\big)^2, & \text{if } \frac{\beta s}{2\pi}<1.\end{cases} + \begin{cases}\frac{128}{d}\big(\frac{s}{\delta}\big)^4, & \text{if } \frac{2s}{\delta}\ge 1,\\ \frac{32}{d}\big(\frac{s}{\delta}\big)^2, & \text{if } \frac{2s}{\delta}<1.\end{cases} \quad (5.43)$$

Theorem 5.F.2. *For an s-sparse Hamiltonian h on n qubits, assume access to the oracle tuple $\mathcal{O}_h$. We denote the controlled $(1,n+5,\varepsilon_{PA}+\varepsilon_{p(h)}+\delta_{class})$-block-encoding of $\frac{1}{4}G^{(\delta,\beta,\omega)}(h)$ in equation (5.7) by $C\text{-}U_{G^{(\delta,\beta,\omega)}}$. The implementation of this block-encoding requires*

$$\begin{cases}\Theta\Big(\frac{(\beta s)^4}{\varepsilon_{PA}}\Big), & \text{if } \frac{\beta s}{2\pi}\ge 1,\\ \Theta\Big(\frac{(\beta s)^2}{\varepsilon_{PA}}\Big), & \text{if } \frac{\beta s}{2\pi}<1.\end{cases} + \begin{cases}\Theta\Big(\frac{s^4}{\delta^4\varepsilon_{PA}}\Big), & \text{if } \frac{2s}{\delta}\ge 1,\\ \Theta\Big(\frac{s^2}{\delta^2\varepsilon_{PA}}\Big), & \text{if } \frac{2s}{\delta}<1.\end{cases} \quad (5.44)$$

calls to oracles from the oracle tuple $\mathcal{O}_h$,

$$O\big(sn+n_a+\log^{5/2}\big(16s^9(\beta^4+1/\delta^4)^2/(\varepsilon_{PA}^2\varepsilon_{p(h)}^2)\big)\big) \quad (5.45)$$

ancillary qubits, and

$$O\big(n+(n+4)(\beta^4s^4+s^4/\delta^4)/\varepsilon_{PA}+\log^{5/2}\big(16s^9(\beta^4+1/\delta^4)^2/(\varepsilon_{PA}^2\varepsilon_{p(h)}^2)\big)\big) \quad (5.46)$$

additional one-qubit and two-qubit gates. To implement this block-encoding, an additional classical computing time of $\mathrm{poly}\big((\beta^4s^4+s^4/\delta^4)/\varepsilon_{PA},\log(1/\delta_{class})\big)$ *is required.*

Proof. Like in the proof of Theorem 5.D.2, we employ Proposition 5.3.3 (from [22]) to construct a $(s, n+3, \varepsilon_{\mathrm{BE}_h})$-block-encoding U_h of h. Using this block encoding, we construct a block encoding of a polynomial approximation of $G^{(\delta,\beta,\omega)}(h)$. Let $p_d(x)$ denote the degree-d polynomial approximation of the function $g^{(\delta,\beta,\omega)}(x)$ from Lemma 5.F.1. It follows from Lemma 5.F.1 that one can efficiently construct p_d such that

$$\left|\left|p_d(h/s) - G^{(\delta,\beta)}(\omega, h)\right|\right| \tag{5.47}$$

is upper bounded by the RHS of the inequality in equation (5.43). Hence, taking

$$d \leq \begin{cases} \Theta\left(\frac{(\beta s)^4}{\varepsilon_{\mathrm{PA}}}\right), & \text{if } \frac{\beta s}{2\pi} \geq 1, \\ \Theta\left(\frac{(\beta s)^2}{\varepsilon_{\mathrm{PA}}}\right), & \text{if } \frac{\beta s}{2\pi} < 1. \end{cases} + \begin{cases} \Theta\left(\frac{s^4}{\delta^4 \varepsilon_{\mathrm{PA}}}\right), & \text{if } \frac{2s}{\delta} \geq 1, \\ \Theta\left(\frac{s^2}{\delta^2 \varepsilon_{\mathrm{PA}}}\right), & \text{if } \frac{2s}{\delta} < 1. \end{cases} \tag{5.48}$$

we obtain $\left|\left|p_d(h/s) - \delta/8\, G^{(\delta,\beta)}(\omega, h)\right|\right| \leq \varepsilon_{\mathrm{PA}}$.

For $\varepsilon_{\mathrm{PA}} \leq \frac{1}{4}$, we note that $|p_d(x)| \leq 1/2$ for $x \in [-1, +1]$, allowing us to apply Proposition 5.3.4 (from [14]). A $(1, n+5, 4d\sqrt{\varepsilon_{\mathrm{BE}_h}/s} + \delta)$-block-encoding of $p_d(h/s)$ consists of a circuit with $O((n+4)d)$ one-qubit and two-qubit gates, and at most d calls to unitaries U_h, $U_h^\dagger$ or controlled-U_h. The classical description of this circuit can be classically computed in $O\big(\mathrm{poly}(d, \log(1/\delta_{\mathrm{class}}))\big)$ time. We define $\varepsilon_{p(h)} := 4d\sqrt{\varepsilon_{\mathrm{BE}_h}/s}$ so that for a given $\varepsilon_{p(h)}$, we should ensure that $\varepsilon_{\mathrm{BE}_h} = s\varepsilon_{p(h)}^2/(16d^2)$.

Let the $(1, n+5, \varepsilon_{p(h)} + \delta_{\mathrm{classical}})$-block-encoding of $p_d(h/s)$ be denoted by $U_{p_d(h/s)}$. Like in the proof of Theorem 5.D.2, we have that $\varepsilon_{\mathrm{Tot}} = \|G^{(\delta,\beta,\omega)}(h) - \langle 0|^{\otimes a} \otimes \mathbb{1} U_{p_d(h/s)} |0\rangle^{\otimes a} \otimes \mathbb{1}\| \leq \varepsilon_{\mathrm{PA}} + \varepsilon_{p(h)} + \delta_{\mathrm{class}}$. We have thus constructed a $(1, n+5, \varepsilon_{\mathrm{PA}} + \varepsilon_{p(h)} + \delta_{\mathrm{class}})$-block-encoding of $G^{(\delta,\beta,\omega)}(h)$. To implement this block-encoding, we require a number of calls to oracles from the tuple $\mathcal{O}_h$, a number of ancillary qubits, and a number of one-qubit and two-qubit gates as in the lemma statement. □

Let us now give the proof of Lemma 5.F.1.

Proof. We wish to approximate $g^{(\delta,\beta,\omega)}(x)$ in equation (5.43) by a polynomial of degree d. Let us first express $g^{(\delta,\beta,\omega)}(x)$ as

$$\delta/8\Big((1 - f^{(\beta)}(x)) g_1^{(\delta,\omega)}(x) + f^{(\beta)}(x) g_2^{(\delta,\omega)}(x)\Big), \tag{5.49}$$

and its degree-d polynomial approximation $p_d(x)$ by

$$\delta/8\Big(\big(1 - f_{d/2}^{(\beta)}(x)\big) g_{1,d/2}^{(\delta,\omega)}(x) + f_{d/2}^{(\beta)}(x) g_{2,d/2}^{(\delta,\omega)}(x)\Big). \tag{5.50}$$

Note that

$$\begin{aligned} |\delta/8\, g^{(\delta,\beta,\omega)}(x) - p_d(x)| \leq \delta/8\Big(&|g_1^{(\delta,\omega)}(x) - g_{1,d/2}^{(\delta,\omega)}(x)| \\ &+ |g_2^{(\delta,\omega)}(x) - g_{2,d/2}^{(\delta,\omega)}(x)|\Big) + 1/2 |f^{(\beta)}(x) - f_{d/2}^{(\beta)}(x)|, \end{aligned} \tag{5.51}$$

where we have used that $|g_{1,d/2}^{(\delta,\omega)}(x)|, |g_{2,d/2}^{(\delta,\omega)}(x)| \leq 2/\delta$ for sufficiently large d (note that $|g_1^{(\delta,\omega)}(x)|, |g_2^{(\delta,\omega)}(x)| \leq 1/\delta$). Using the bound on $\max_{x\in[-1,+1]} |f^{(\beta)}(x) - f_{d/2}^{(\beta)}(x)|$ from Lemma 5.D.1, and applying Bernstein's theorem [30] (i.e., Lemma 5.3.6) to the functions $g_1^{(\delta,\omega)}(x)$ and $g_2^{(\delta,\omega)}(x)$ (with a Bernstein ellipse E_r with $r = \sqrt{(\delta/(2s))^2 + 1}$), we obtain the upper bound on $\max_{x\in[-1,+1]} |g^{(\delta,\beta,\omega)}(x) - p_d(x)|$ in the lemma statement. □

5.G. Proof of Lemma 5.4.1

Proof of Lemma 5.4.1. By assumption, we have that $\big|\langle i|\,A\,|j\rangle - \alpha\,\langle 0|^{\otimes m}\,\langle i|\,U_A\,|0\rangle^{\otimes m}\,|j\rangle\big| \leq \varepsilon$, where U_A acts on $n+m$ qubits. Let us consider estimating $\langle 0|^{\otimes m}\,\langle i|\,U_A\,|0\rangle^{\otimes m}\,|j\rangle$, which can alternatively be expressed as

$$\langle 0|^{\otimes m}\,\langle 0|^{\otimes n}\,(\mathbb{1}\otimes U_i^\dagger)\,U_A\,(\mathbb{1}\otimes U_j)\,|0\rangle^{\otimes m}\,|0\rangle^{\otimes n}, \tag{5.52}$$

where U_i, U_j are depth-1 circuits which prepare bit-strings i and j. We denote the estimate of $\langle 0|^{\otimes m}\,\langle i|\,U_A\,|0\rangle^{\otimes m}\,|j\rangle$ by $\widetilde{\langle i|\,A\,|j\rangle}$, so that if $\big|\langle 0|^{\otimes m}\,\langle i|\,U_A\,|0\rangle^{\otimes m}\,|j\rangle - \widetilde{\langle i|\,A\,|j\rangle}\big| \leq \tilde{\varepsilon}$, then $\big|\langle i|\,A\,|j\rangle - \alpha\widetilde{\langle i|\,A\,|j\rangle}\big| \leq \varepsilon + \alpha\tilde{\varepsilon}$.

One can obtain the estimate $\widetilde{\langle i|\,A\,|j\rangle}$ by running a series of Hadamard test circuits on $n+m+1$ qubits. These circuits correspond to running

$$\big(\mathbb{1}\otimes[H\,R_z(\theta)]_{\mathrm{a}}\big)\big(\mathbb{1}\otimes|0\rangle\langle 0|_{\mathrm{a}} + U\otimes|1\rangle\langle 1|_{\mathrm{a}}\big)\big(\mathbb{1}\otimes H_{\mathrm{a}}\big), \tag{5.53}$$

where $U = (U_i^\dagger\otimes\mathbb{1})\,U_A\,(U_j\otimes\mathbb{1})$, on the state $|0\rangle^{\otimes m}\,|0\rangle^{\otimes n}\,|0\rangle_{\mathrm{a}}$ (with the final qubit being an ancillary qubit). The output state of the ancillary qubit is measured a total of $D(\tilde{\varepsilon},\delta)$ times, half of the times for $\theta = 0$ and half of the times for $\theta = \pi/2$. The fractions of output-0 measurements for $\theta = 0$ and $\theta = \pi/2$ provide estimates of $\frac{1}{2} + \frac{1}{2}\mathrm{Re}\big(\langle 0|^{\otimes m}\,\langle i|\,U_A\,|0\rangle^{\otimes m}\,|j\rangle\big)$ and $\frac{1}{2} - \frac{1}{2}\mathrm{Im}\big(\langle 0|^{\otimes m}\,\langle i|\,U_A\,|0\rangle^{\otimes m}\,|j\rangle\big)$, respectively. Using a Chernoff concentration bound, one can show that $\big|\widetilde{\langle i|\,A\,|j\rangle} - \langle 0|^{\otimes m}\,\langle i|\,U_A\,|0\rangle^{\otimes m}\,|j\rangle\big| \leq \tilde{\varepsilon}$ with probability at least $1-\delta$ for $D(\tilde{\varepsilon},\delta) = \Theta\big(\tilde{\varepsilon}^{-2}\log\big(4\delta^{-1}\big)\big)$.

One can thus obtain an estimate of $\langle i|\,A\,|j\rangle$ (given by $\alpha\widetilde{\langle i|\,A\,|j\rangle}$) up to error $\varepsilon + \alpha\tilde{\varepsilon}$ with probability $1-\delta$, using $D(\tilde{\varepsilon},\delta) = \Theta\big(\tilde{\varepsilon}^{-2}\log\big(4\delta^{-1}\big)\big)$ calls to $C\text{-}U_A$.

□

5.H. BQP-completeness

Here we prove Theorem 5.5.1 in the main text, using the next Lemma 5.H.1 as a small tool:

Proof of Theorem 5.5.1. It is straightforward to see that evaluating the matrix element $M_{jj}(t)$ of the correlation matrix $M(t) = e^{iht}M_0e^{-iht}$ at $t = \mathrm{poly}(n)$ is a problem in BQP, given the promise. By Lemmas 5.4.1 and 5.E.1, given access to $\mathcal{O}_{M_0}$ and $\mathcal{O}_h$ as $\mathrm{poly}(n)$-sized quantum circuits, the problem is solved with $\mathrm{poly}(n)$ quantum effort.

To show BQP-hardness of our problem, we use the fact that for any promise problem in BQP of problem size m, we have the following property [32]: the problem can be decided by acting on an $k = \mathrm{poly}(m)$-qubit input $|00\ldots0\rangle$ with (a uniform family of) $\mathrm{poly}(k) = \mathrm{poly}(m)$-sized quantum circuits, outputing 1 (on the first qubit) with probability at least $2/3$ in case YES, and 1 with probability at most $1/3$ in case NO. In addition, one can boost the success and failure probabilities $2/3 \to 1-\exp(-\Theta(k))$ and $1/3 \to \exp(-\Theta(k))$, by running k instances of the $\mathrm{poly}(k)$-sized circuits in parallel and taking a majority vote on the first qubit of the output state for each instance (and copying the answer onto an ancillary qubit). The circuit corresponding to this boosted scenario acts on $q = k^2$ qubits, and its success and failure probabilities are respectively

$1-\exp\big(-\Theta(\sqrt{q})\big)$ and $\exp\big(-\Theta(\sqrt{q})\big)$. Let the quantum circuit for this problem with boosted probabilities be

$$U = W_L \dots W_1, \tag{5.54}$$

where W_l are elementary one-qubit and two-qubit gates and $L = \text{poly}(k) = \text{poly}(\sqrt{q})$. We represent this decision problem using time-evolution with a sparse circuit Hamiltonian. The circuit Hamiltonian, acting on a $q_{\text{clock}} = \log_2(L+1)$-qubit clock space (we assume wlog that $\log_2(L+1)$ is an integer) and the q-qubit space is given by

$$h = \sum_{l=1}^{L} \Big(|l+1\rangle\langle l|_{\text{clock}} \otimes W_l + |l\rangle\langle l+1|_{\text{clock}} \otimes W_l^\dagger \Big). \tag{5.55}$$

We take $n = q_{\text{clock}} + q$ and note that $q_{\text{clock}} < q$ for sufficiently large q, so that $n/2 \le q \le n$. The matrices W_l have at most 4 non-zero entries in a given row/column. Therefore, h is at most 8-sparse. Since $\{W_l\}_{l=1}^{L}$ are unitary matrices, the entries of h are $O(1)$ in absolute value.

Consider the evolution $\big|\psi(t)\big\rangle = e^{-iht}|1\rangle_{\text{clock}}|00\dots0\rangle$ with the Hamiltonian h from equation (5.55). This state can be decomposed as

$$\big|\psi(t)\big\rangle = \sum_{l=1}^{L+1} \alpha_{l,t}\,|l\rangle_{\text{clock}} \otimes \prod_{l'=1}^{l-1} W_{l'}\,|00\dots0\rangle \tag{5.56}$$

with coefficients $\alpha_{l,t}$ given by

$$\sum_{l=1}^{L+1} \alpha_{l,t}\,|l\rangle \equiv e^{-iJt}\,|1\rangle_{\text{clock}}, \tag{5.57}$$

where J is a Hamiltonian on the clock register

$$J = \sum_{l=1}^{L} \left(|l+1\rangle\langle l|_{\text{clock}} + |l\rangle\langle l+1|_{\text{clock}} \right). \tag{5.58}$$

Given the encoding of the clock register, one can write the probability of measuring $|L+1\rangle_{\text{clock}}$ on the clock and measuring $|1\rangle$ on the first of the q qubits as

$$p \equiv \Big| \big(\langle L+1|_{\text{clock}} \otimes \langle 1|_1\big)\,|\psi(t)\rangle \Big|^2 = \langle 1|_{\text{clock}}\langle 00\dots0|\,e^{iht} M_0 e^{-iht}\,|1\rangle_{\text{clock}}|00\dots0\rangle, \tag{5.59}$$

with $M_0 = \frac{1}{2^{q_{\text{clock}}+1}} \prod_{j=1}^{q_{\text{clock}}} (\mathbb{1} - Z_{\text{clock},j})(\mathbb{1} - Z_{\text{qubit},1})$. Hence, when the state $U|00\dots0\rangle$ outputs 1 on the first qubit with probability at least $1-\exp\big(-\sqrt{q}\big)$ (YES), it follows through Lemma 5.H.1 that $p = \Omega(1/\text{poly}(\sqrt{q})) = \Omega(1/\text{poly}(\sqrt{n}))$. When the state $U|00\dots0\rangle$ outputs 1 on the first qubit with probability at most $\exp\big(-\sqrt{q}\big)$ (NO), then $p \le \exp\big(-\sqrt{q}\big) \le \exp\big(-\sqrt{n/2}\big)$ through Lemma 5.H.1. Now, observe that M_0 is a valid and sparse correlation matrix of a multi-particle free-fermionic state on 2^n modes (in particular, a fraction $\Theta\big(1/\text{poly}(\sqrt{n})\big)$ of the modes is occupied), which is evolved in time $t = \text{poly}(\sqrt{n})$ by the sparse Hamiltonian h, after which one wishes to estimate a particular matrix element (labeled, say, by $j = 1_{\text{clock}}, 00\dots0$) of the time-evolved matrix, which is the problem stated in Theorem 5.5.1. The only thing left to argue is that given the description of $\{W_l\}$, one can implement $\mathcal{O}_h$ in Definition 5.3.2 as a poly(n)-sized circuit.

Oracle implementation: The oracle O_r from Definition 5.3.2, acting on $(s+1)(q_{\text{clock}}+q+1)$ qubits, can be implemented as follows. For convenience, we label the first $(q_{\text{clock}}+q+1)$ qubits by A and the last s $(q_{\text{clock}}+q+1)$-qubit registers by $B_1,\ldots,B_s$. For simplicity and wlog, we assume that all W_l are two-qubit gates and all entries of W_l in their two-qubit sub-spaces are non-zero. Note that for each $l\in\{1,2,\ldots,L\}$, we have access to the labels $Q_1^{(l)}$ and $Q_2^{(l)}$ (with $Q_1^{(l)}<Q_2^{(l)}$) of the qubits on which W_l acts non-trivially. The structure of h is such that each row contains 8 non-zero entries (apart from the rows associated with clock states $|1\rangle_{\text{clock}}$ and $|L+1\rangle_{\text{clock}}$), with a row $|i\rangle=|l\rangle_{\text{clock}}|x\rangle$ having four non-zero entries associated with clock register state $|l-1\rangle_{\text{clock}}$ and four non-zero entries associated with clock register state $|l+1\rangle_{\text{clock}}$. These entries correspond to the entries $\left\langle x_{Q_1^{l-1}},x_{Q_2^{l-1}}\right|W_{l-1}\left|y_1,y_2\right\rangle$ and $\left\langle x_{Q_1^{l}},x_{Q_2^{l}}\right|W_{l}\left|y_1,y_2\right\rangle$ (for $y\in\{0,1\}^2$), respectively. The rows associated with clock states $|1\rangle_{\text{clock}}$ and $|L+1\rangle_{\text{clock}}$ are 4-sparse.

We take workspace in the form of $2(L+1)$ additional $(q_{\text{clock}}+q)$-qubit registers (initialized in $|00\ldots0\rangle$), denoted by $C_1,\ldots,C_{2(L+1)}$. For each $j\in\{1,2,\ldots,L+1\}$, we transform the first $(L+1)$ qubits on registers C_{2j-1} and C_{2j} to $\left|j\right\rangle_{\text{clock}}$. Then, for each $j\in\{2,3,\ldots,L\}$ (so excluding 1 and $L+1$), we flip qubits $q_{\text{clock}}+Q_1^{j-1}$ and $q_{\text{clock}}+Q_2^{j-1}$ on register C_{2j-1} and qubits $q_{\text{clock}}+Q_1^{j}$ and $q_{\text{clock}}+Q_2^{j}$ on register C_{2j} to $|1\rangle$. In addition, we flip qubits $q_{\text{clock}}+Q_1^{1}$ and $q_{\text{clock}}+Q_2^{1}$ on register C_2 and $q_{\text{clock}}+Q_1^{L}$ and $q_{\text{clock}}+Q_2^{L}$ on register C_{2L-1} to $|1\rangle$.

Controlled on the clock state on register A being $|l\rangle_{\text{clock}}$, we set the clock state to $|l-1\rangle_{\text{clock}}$ on registers $B_1,\ldots,B_4$ (provided that $l>1$) and to $|l+1\rangle_{\text{clock}}$ on register $B_5,\ldots,B_8$ (provided that $l<L+1$). Controlled on the last q qubits of register A being in state $|x\rangle$, we copy $|x\rangle$ onto the final q qubits of $B_1,\ldots,B_4$, excluding qubits $q_{\text{clock}}+Q_1^{l-1}$ and $q_{\text{clock}}+Q_2^{l-1}$. These latter two qubits are transformed to $|00\rangle$, $|01\rangle$, $|10\rangle$ and $|11\rangle$ on registers $B_1,\ldots,B_4$, respectively. Similarly, we copy $|x\rangle$ onto the final q qubits of $B_5,\ldots,B_8$, apart from qubits $q_{\text{clock}}+Q_1^{l}$ and $q_{\text{clock}}+Q_2^{l}$, which are respectively transformed to $|00\rangle$, $|01\rangle$, $|10\rangle$ and $|11\rangle$. These operations make use of the states in the workspace registers $C_1,\ldots,C_{2(L+1)}$, which are uncomputed at the end of the protocol. In accordance with Definition 5.3.2, we need to account for rows of h having less than 8 non-zero entries. Since the rows of h associated with clock states $|1\rangle_{\text{clock}}$ and $|L+1\rangle_{\text{clock}}$ are 4-sparse, registers $B_1,\ldots,B_4$ are set to resp. $|1\rangle\otimes|5\rangle_{q_{\text{clock}}+q},\ldots,|1\rangle\otimes|8\rangle_{q_{\text{clock}}+q}$ controlled on the A clock state being $|1\rangle_{\text{clock}}$ (after which registers $(B_1,\ldots,B_4)$ and $(B_5,\ldots,B_8)$ are swapped), and registers $B_5,\ldots,B_8$ are set to resp. $|1\rangle\otimes|5\rangle_{q_{\text{clock}}+q},\ldots,|1\rangle\otimes|8\rangle_{q_{\text{clock}}+q}$ controlled on the A clock state being $|L+1\rangle_{\text{clock}}$. The size of the circuit implementing O_r is poly(n).

To implement oracle O_a, let us note that wlog the entries of W_l are 0, $\pm1/\sqrt{2}$ or 1, so that the entries can be encoded into a three bit string. By employing additional poly(n)-sized workspace (note that $L=\text{poly}(\sqrt{q})$ and each W_l has 16 entries), the oracle O_a can be implemented (by a poly(n)-sized circuit).

□

Remark: Like in [17], we could have adapted the BQP-verification circuit to output the state $|0\rangle_a\otimes|00\ldots0\rangle$ (so all qubits back to their initial state and an additional ancilla qubit a to 0) with high probability in the NO case, and with low probability in the YES case. This is done by simply copying the answer of the BQP-circuit onto an additional ancilla qubit a and applying the gates $W_L\ldots W_1$ in reverse on the other qubits. If we use

this cleaned-up circuit, it means that we are interested in estimating the probability for a specific output state — all qubits in $|0\rangle$ and clock state in $|L+1\rangle_{\text{clock}}$ — and this corresponds to estimating an entry of a time-evolved rank-1 projector $\tilde{M}_0$, corresponding to a single-particle state. Hence not surprisingly, time-evolution of single-particle states is also BQP-complete, as was shown in Theorem 3 in [17] (where more work was done to bring h in sign-free form to directly correspond to a sum of kinetic and potential energy).

The following lemma, which is used in the proof of Theorem 5.5.1, mainly follows the approach of [17]. Instead of employing this lemma, one could also adapt the coefficients in the hopping Hamiltonian h in equation (5.55) to allow for a perfect 1D state transfer from $|1\rangle_{\text{clock}} \to |L+1\rangle_{\text{clock}}$, using an idea first suggested by Peres [45], see also [17]: such adaptation requires extra ancilla qubit overhead in realizing the time-dynamics of h, hence we omit it.

Lemma 5.H.1. *For a Hamiltonian $J = \sum_{l=1}^{L}(|l\rangle\langle l+1| + |l+1\rangle\langle l|)$ on a $(L+1)$-dim Hilbert space with basis states $|l\rangle$, $l \in \{1,\ldots,L+1\}$, there exists a $t = O(L^2 \log L)$ such that*

$$|\langle L+1|\, e^{-iJt}\, |1\rangle| = \Omega(1/\sqrt{L}). \tag{5.60}$$

Proof. The Hamiltonian J has eigenstates

$$|\psi_k\rangle = \sum_{j=1}^{L+1} \alpha_j^{(k)} |j\rangle, \text{ with } \alpha_j^{(k)} = \sqrt{\frac{2}{L+2}} \sin\Big(\frac{\pi jk}{L+2}\Big), \tag{5.61}$$

and eigenvalues

$$\epsilon_k = 2\cos\Big(\frac{\pi k}{L+2}\Big), \tag{5.62}$$

with $k = 1 \ldots L+1$. We note that the gap between any two eigenvalues is at most 4. To prove a lower bound on $|\langle L+1|\, e^{-iJt}\, |1\rangle|$, we will derive a lower bound on the gaps $\Delta_m := |\epsilon_{m+1} - \epsilon_m|$ (for $m = 1, 2 \ldots L$) between the eigenvalues of J:

$$\Delta_m = |\epsilon_{m+1} - \epsilon_m| \geq \frac{\pi}{L+2} \min_{x \in \left[\frac{m\pi}{L+2}, \frac{(m+1)\pi}{L+2}\right]} \Big|\frac{d\, 2\cos(x)}{dx}\Big| \geq \frac{2\pi}{L+2} \sin\Big(\frac{\pi}{L+2}\Big) = \Omega(1/(L+2)^2). \tag{5.63}$$

Using the eigendecomposition of J, we infer that

$$\langle L+1|\, e^{-iJt}\, |1\rangle = \frac{2}{L+2} \sum_{k=1}^{L+1} e^{-i\epsilon_k t} (-1)^{k-1} \sin^2\Big(\frac{\pi k}{L+2}\Big), \tag{5.64}$$

so that

$$|\langle L+1|\, e^{-iJt}\, |1\rangle|^2 = \Big(\frac{2}{L+2}\Big)^2 \sum_{k,k'=1}^{L+1} e^{-i(\epsilon_k - \epsilon_{k'})t} (-1)^{k+k'} \sin^2\Big(\frac{\pi k}{L+2}\Big) \sin^2\Big(\frac{\pi k'}{L+2}\Big). \tag{5.65}$$

To show that there must be a time t for which $|\langle L+1|\, e^{-iJt}\, |1\rangle|^2 = \Omega(1/L)$, we use the fact that a probabilistically chosen time in a sufficiently large interval will give high success probability [33], and hence there must exist a specific time which works

sufficiently well. More precisely, for $k \neq k'$, there must exist a probability distribution $\{p(t)\}_{t=0}^{T} \geq 0$, $\sum_{t=0}^{T} p(t) = 1$, such that

$$\Big|\sum_{t=0}^{T} p(t) e^{-i(\epsilon_k - \epsilon_{k'})t}\Big| \leq \varepsilon, \tag{5.66}$$

provided that $\Delta = \Omega\big(1/(L+2)^2\big)$ and $T = O\big((L+2)^2 \log(1/\varepsilon)\big)$. Examples of probability distributions for which this is true are given in Ref. [46].

Therefore, for those $\{p(t)\}$'s we have that

$$\Big|\sum_{k\neq k'}\sum_{t=0}^{T} p(t) e^{-i(\epsilon_k-\epsilon_{k'})t}(-1)^{k+k'} \sin^2\Big(\frac{\pi k}{L+2}\Big)\sin^2\Big(\frac{\pi k'}{L+2}\Big)\Big| \leq$$
$$\varepsilon \sum_{k\neq k'} \sin^2\Big(\frac{\pi k}{L+2}\Big)\sin^2\Big(\frac{\pi k'}{L+2}\Big) = \varepsilon\Big(\frac{(L+2)^2}{4} - \frac{3(L+2)}{8}\Big) \leq \varepsilon\frac{(L+2)^2}{4},$$

where the equality follows from direct computation. We thus conclude that

$$\Big|\sum_{t=0}^{T} p(t)|\langle L+1|e^{-iJt}|1\rangle|^2 - \sum_{t=0}^{T} p(t)\Big(\frac{2}{L+2}\Big)^2 \sum_{k=1}^{L+1}\sin^4\Big(\frac{\pi k}{L+2}\Big)\Big| \leq \varepsilon. \tag{5.67}$$

The term $\sum_{t=0}^{T} p(t)\big(\frac{2}{L+2}\big)^2 \sum_{k=1}^{L+1}\sin^4\big(\frac{\pi k}{L+2}\big)$ can be evaluated to be $\frac{3}{2(L+2)}$. So choosing, for instance, $\varepsilon = \frac{1}{2(L+2)}$, we know that $\sum_{t=0}^{T} p(t)|\langle L+1|e^{-iJt}|1\rangle|^2 = \Omega\big(\frac{1}{L+2}\big)$. For $T = O\big((L+2)^2\log(2(L+2))\big)$, we conclude that there must be a $t = O(L^2 \log L)$ for which $|\langle L+1|e^{-iJt}|1\rangle|^2 = \Omega(1/L)$. □

5.I. Classically estimating entries of the time-evolved correlation matrix on lattice models

In this appendix we briefly argue the following. For $t = \text{poly}(n)$ and assuming classical access to entries $\langle k|M|l\rangle$ of an initial correlation matrix M for given (k,l), one can obtain entries $M(t)_{ij}$ with $1/\exp(n)$ error with $\text{poly}(n)$ classical effort. To see this, note that $\max_{x\in[-1,+1]}|p_K(x) - \exp(itsx)| = O\big((t/\sqrt{K})^{K+1}\big)$, with $p_K(x)$ a degree-K Taylor approximation. This implies

$$\Big|\big(p_K(h/s)Mp_K(-h/s)\big)_{ij} - M(t)_{ij}\Big| = O\big((t/\sqrt{K})^{K+1}\big), \tag{5.68}$$

where we have used that $\|M\| \leq 1$. Note that this error can be bounded by $1/\exp(n)$ for some $K = \text{poly}(n)$. Using the same reasoning as in the proof of Lemma 5.6.1, we can obtain $\langle i|h^{k_1}Mh^{k_2}|j\rangle$ for all $k_1, k_2 \leq K = \text{poly}(n)$, giving an estimate of $\big(p_K(h/s)Mp_K(-h/s)\big)_{ij}$. So for sufficiently large $K = \text{poly}(n)$, we obtain an estimate of $\big(e^{+ith}M_0e^{-ith}\big)_{ij}$ with $1/\exp(n)$ error.

6

Spectral estimation for Hamiltonians: Classical imaginary-time evolution versus quantum real-time evolution

Wil je blijven?

Ben ik te min (1967), Armand

6.1. Introduction

In general, it is a computationally intractable task to obtain, by classical or quantum means, the eigenvalues of a Hamiltonian H associated with a many-body quantum system. However, as discussed in Section 1.2, more restricted tasks related to estimating the spectrum of H can be executed on a quantum computer by means of quantum phase estimation algorithms [1–4], using the ability to simulate the real-time dynamics $e^{-iHt/\hbar}$ efficiently on a quantum computer via Trotterization [5].

Classical alternatives are provided by quantum Monte Carlo methods [6, 7]. The efficiency of quantum Monte Carlo methods when used to simulate many-body systems is generally limited by the *sign problem*. This can cause the variance of the estimator in the Monte Carlo algorithm to grow exponentially in the system size n, necessitating an exponential number of runs of the Monte Carlo algorithm.

A (ubiquitous) class of Hamiltonians that is sign-problem-free has been formalized under the name *stoquastic* Hamiltonians [8]. Roughly speaking, a (real-valued) Hamiltonian is stoquastic (in a particular basis $\mathcal{B}$) if its off-diagonal elements are non-positive: $\langle x| H \big| y\rangle \leq 0$, for $x \neq y$ (with $|x\rangle$, $\big|y\rangle$ being elements of $\mathcal{B}$). As a consequence, its associated Gibbs density matrix $e^{-\tau H}$ is an element-wise non-negative matrix (for $\tau \in \mathbb{R}_+$). This property makes it particularly suitable for Monte Carlo sampling as complexity results [8–10] and various algorithmic results [11–14] have demonstrated.

Since stoquastic Hamiltonians are sign-problem-free, it is of interest to see if one can indeed prove that (part of its) spectrum can be efficiently estimated through classical Monte Carlo methods. Conversely, can quantum algorithms, *even* for stoquastic Hamiltonians, provide an advantage over Monte Carlo algorithms in carrying out this task? In this work, we address these questions by making a direct comparison between the task of estimating the spectral content of a stoquastic local Hamiltonian in an input state via a quantum circuit versus via a classical Monte Carlo scheme. In addition, we investigate to what extent this task can be efficiently carried out classically for a general local Hamiltonian.

Central in our study is, first of all, the real-time *signal*

$$g_R(k) = \langle\Phi| e^{-iHk\Delta t} |\Phi\rangle = \sum_{j=1}^{2^n} \left|\langle\psi_j|\Phi\rangle\right|^2 \left(e^{-iE_j\Delta t}\right)^k, \tag{6.1}$$

for $k = 0, 1 \ldots, K$ and where $|\Phi\rangle$ is some pre-specified n-qubit input state. The estimation of $g_R(k)$ for various k is a crucial step in the quantum phase estimation algorithm (QPE). In what follows we will fix Δt so that the eigenstates $|\psi_j\rangle$ with nonzero or substantial overlap $|\langle\psi_j|\Phi\rangle|^2 > 0$, showing up in the signal, have the property that $E_j\Delta t \in [0, 2\pi)$. Thus, from now on, we assume that these E_j are shifted and rescaled to lie in $[0, 2\pi)$. We will assume that there are at most S eigenvectors with nonzero $|\langle\psi_j|\Phi\rangle|^2$, where S is desired to be poly(n) or less for overall efficiency. Identifying a state $|\Phi\rangle$ which has non-zero overlap on only a few ($S = \text{poly}(n)$ or $S = O(1)$) eigenstates and which obeys the assumptions in the following Theorems is not so simple, and can be considered one of the bottlenecks in using quantum phase estimation or other Monte Carlo methods to determine spectral information of the Hamiltonian.

Besides the real-time signal, one can define the imaginary-time *signal*

$$g_I(k) = \langle\Phi| e^{-Hk} |\Phi\rangle = \sum_{j=1}^{2^n} \left|\langle\psi_j|\Phi\rangle\right|^2 \left(e^{-E_j}\right)^k, \tag{6.2}$$

where again we can assume that $E_j \in [0, 2\pi)$. We will prove, *for local stoquastic Hamiltonians*, that the quantum cost of estimating $g_R(k)$ and the classical Monte Carlo cost of estimating $g_I(k)$ within error ϵ are *approximately identical*, although the assumptions on our knowledge/preparation costs of $|\Phi\rangle$ are slightly different in the two cases. We present the Monte Carlo scheme that estimates the signal in equation (6.2) in Section 6.2. We stress that this Monte Carlo scheme does *not* rely on the Metropolis-Hastings algorithm.

The two statements are as follows (the proof of these statements is given in Section 6.2):

Theorem 6.1.1. *For a local Hamiltonian acting on n qubits, one can estimate $g_R(k)$ in equation* (6.1) *with probability at least $1-\delta$ with sampling error ϵ and Trotter error ϵ_{trot} (and total error $\epsilon_{\text{tot}} = \epsilon + \epsilon_{trot}$), using quantum circuits acting on $n+1$ qubits, where the depth of the quantum circuit scales as $\mathcal{O}\left(k^{1+o(1)}\right)\mathcal{O}\left(\epsilon_{trot}^{-o(1)}\right) \times \text{poly}(n)$ and the number of times one executes the circuit is $\Theta(\epsilon^{-2}\log\left(4\delta^{-1}\right))$, under the assumption that $|\Phi\rangle$ is a state of n qubits which can be generated by a* poly(n)*-size quantum circuit. Hence to obtain*

$g_R(k)$ for $k = 0,\ldots,K$, with error at most $\epsilon_{\text{tot}} = \epsilon + \epsilon_{\text{trot}}$ for all k, with probability $1-\delta$ requires using quantum circuits for $k = 0,\ldots,K$, each acting on $n+1$ qubits, where the depth of the quantum circuit scales as $\mathcal{O}\left(k^{1+o(1)}\right)\mathcal{O}\left(\epsilon_{trot}^{-o(1)}\right) \times \text{poly}(n)$ and each circuit is repeated $\Theta(\epsilon^{-2}\left[\log\left(4\delta^{-1}\right) + \log(K)\right])$ times.

Theorem 6.1.2. *For a local stoquastic Hamiltonian acting on n qubits, one can estimate $g_I(k)$ in equation (6.2) with probability at least $1-\delta$ with total error $\epsilon_{\text{tot}} = \epsilon + \epsilon_{\text{trot}}$, using a classical MC algorithm on n-bit strings where the depth of the algorithm scales as $\mathcal{O}\left(k^{1+o(1)}\right)\mathcal{O}\left(\epsilon_{trot}^{-o(1)}\right) \times \text{poly}(n)$ and the number of times one runs the algorithm is $\Theta(\epsilon^{-2}\log(\delta^{-1}))$, under the assumption that $|\Phi\rangle = \sum_{x=1}^{2^n}\Phi(x)|x\rangle$ is a normalized state of n qubits such that (1) $\frac{\Phi(y)}{\Phi(x)}$ can be efficiently (*poly(n)*) calculated for a* given *x and y and (2) we can efficiently draw samples from the probability distribution $P(x) = \left|\Phi(x)\right|^2$. Hence to obtain $g_I(k)$ for all $k = 0,\ldots,K$, with error at most ϵ_{tot} for each k, with probability $1-\delta$ requires using a classical MC algorithm on n-bit strings for $k = 0,\ldots,K$, where the depth of each algorithm scales as $\mathcal{O}\left(k^{1+o(1)}\right)\mathcal{O}\left(\epsilon_{trot}^{-o(1)}\right) \times \text{poly}(n)$ and the number of times one runs the algorithm (for each k) is $\Theta(\epsilon^{-2}\left[\log(\delta^{-1}) + \log(K)\right])$.*

We then ask, given knowledge of either the real-time signal $g_R(k)$ or imaginary-time signal $g_I(k)$, what can be learnt about those eigenvalues E_j, whose associated eigenstates have nonzero overlap with the input state $|\Phi\rangle$? The signals $g_I(k)$ and $g_R(k)$ respectively correspond to a probabilistic sum of decaying components and a sum of oscillating components with decay rates and oscillation frequencies E_j as a function of discrete 'time' $k = 0,\ldots,K$. Hence a method which extracts those decay and oscillation rates from knowing $g_I(k)$ or $g_R(k)$ at various k is needed. A method of choice which has already been used in quantum information theory is the matrix pencil method [15–17] (with equivalent methods known as ESPRIT and MUSIC). This method has been used for processing randomized benchmarking data [18, 19], quantum phase estimation [3], spectral tomography of superoperators [20], for processing experimental time-series data to identify Hamiltonian parameters [21] or generally in processing discretely-sampled decaying Ramsey signals. In this work, we employ specifically the ESPRIT method and investigate its ability to extract the E_j's from the signals in equation (6.1) and equation (6.2). More details on the ESPRIT method are given in Section 6.3 and the algorithm implementing the ESPRIT method is given explicitly in Algorithm 6.2.

Using this method, it is known that if *either* $g_R(k)$ or $g_I(k)$ is known *exactly* for $k = 0,\ldots,K$ where $K+1 \geq 2S$, one can learn those eigenvalues E_j and probabilities $|\langle\psi_j|\Phi\rangle|^2$ *exactly*. However, in the presence of sampling and Trotter noise, the resolving power also depends on the gap between the eigenvalues E_j, the number S of eigenvalues and whether we extract them from an oscillating or decaying signal. Our work is thus focused on understanding whether there are fundamental advantages in learning $g_R(k)$ with noise versus learning $g_I(k)$ with noise, as this quantifies the benefit of a quantum algorithm versus a classical algorithm for spectral estimation of (stoquastic) Hamiltonians.

Not surprisingly, there are drawbacks to processing data from the imaginary-time evolution. As the signal decays exponentially, k cannot be chosen too large otherwise the

signal becomes smaller than the noise. Our goal is to quantify this precisely and show that, at least theoretically, a regime exists in which the Monte Carlo method may be competitive.

The first statement we make can be viewed as a summary of previous work, namely it combines Lemma 6.2.1 via Theorem 6.1.1 (which state the computational effort of estimating the real-time evolution signal up to a given error and with a given confidence) and the performance of the ESPRIT method in Theorem 6.3.1 (which gives an error bound for the eigenvalue estimates obtained from application of the ESPRIT method to the real-time evolution signal) in the presence of a gap:

Theorem 6.1.3. *Given a local Hamiltonian on n qubits. Let the number of eigenvectors supported in some (efficient-to-prepare) input state* $|\Phi\rangle$ *be* $S = p_1(n)$ *(with* $p_1(n)$ *some polynomial in n), and each occurs with nonzero probability at least* $1/\mathrm{poly}(n)$*. Furthermore, assume that the S eigenvalues* $\{E_i\}$ *with* $E_i \in [0, 2\pi)$ *are sufficiently well-separated, i.e. at least by a gap* $\Delta \geq C/K$ *with constant C and* $K = \Theta(p_1(n))$*. Then using Hadamard test (QPE) quantum circuits plus signal post-processing via ESPRIT, each requiring a* $\mathrm{poly}(n)$ *effort, one can resolve the eigenvalues* $\{E_j\}$ *with distance* $d(\{E_i\}, \{\tilde{E}_j\})$ *(defined in equation* (6.32)*) at most* $1/\mathrm{poly}(n)$*.*

For local stoquastic Hamiltonians the combination of Lemma 6.2.3 via Theorem 6.1.2 (which state the computational effort of estimating the imaginary-time evolution signal up to a given error and with a given confidence) and the performance of the ESPRIT method in Theorem 6.3.3 (which gives an error bound for the eigenvalue estimates obtained from application of the ESPRIT method to the imaginary-time evolution signal) in the presence of a gap leads to:

Theorem 6.1.4. *Given a local stoquastic Hamiltonian on n qubits. Let the number of eigenvectors supported in some efficient-to-sample (i.e. with* $\mathrm{poly}(n)$ *effort) input state* $|\Phi\rangle$ *be* $S = O(1)$*, and each occurs with nonzero probability at least* $1/\mathrm{poly}(n)$*. In addition, assume that for a fixed x, y it is efficient to compute* $\frac{\Phi(y)}{\Phi(x)}$*. Furthermore, assume that the S eigenvalues* $\{E_i\}, E_i \in [0, 2\pi)$ *are sufficiently well-separated, i.e. at least by* $\Delta \geq 1/\mathrm{poly}(n)$ *with some* $\mathrm{poly}(n)$*. Then using a Monte Carlo algorithm plus signal post-processing via ESPRIT, each requiring (some)* $\mathrm{poly}(n)$ *effort, one can resolve the eigenvalues* $\{E_j\}$ *with distance* $d(\{E_i\}, \{\tilde{E}_j\})$ *at most* $1/\mathrm{poly}(n)$*.*

Theorem 6.1.4 immediately begs the question whether such a result could hold for general local Hamiltonians as well: the assumptions that there are only $S = O(1)$ eigenstates in the initial state, as well as the assumption of efficient access to the initial state appear rather strong. To address this question, we define another real-valued, decaying signal as

$$g_D(k) = \langle\Phi| (I - H/2\pi)^k |\Phi\rangle = \sum_{j=1}^{2^n} \left|\langle\psi_j|\Phi\rangle\right|^2 \left(1 - E_j/2\pi\right)^k. \tag{6.3}$$

If $S = O(1)$ and if $g_D(k)$ can be estimated with some accuracy for $k = 1, \ldots, K = O(1)$, we can also apply the ESPRIT method to extract these S eigenvalues. We note that this requires that the eigenvalues E_j are bounded away from 2π. Hence if we use $g_D(k)$ we assume that we have shifted and rescaled the eigenvalues so that, say, the E_j's lie in $[0, \pi]$.

One can prove that for general local Hamiltonians, assuming $S = O(1)$ eigenvalues in $|\Phi\rangle$, one can estimate $g_D(k)$ with ϵ accuracy, under an assumption about the access to $|\Phi\rangle$ which is identical to the Monte Carlo case for stoquastic Hamiltonians (Theorem 6.1.4). In fact, this result shows that Theorem 6.1.4 is not particular to local stoquastic Hamiltonians at all, if we only care about 'nominally poly(n)' algorithms. However, the computational cost of estimating $g_D(k)$ for general local Hamiltonians is significantly higher in practice compared to the Monte Carlo method for stoquastic Hamiltonians. The result expressed in Lemma 6.2.6 can be viewed as 'dequantization' as it is similar in spirit to the Singular Value Transformation (SVT) tool (Theorem 3 in [22]). Theorem 3 in [22] is used to construct an algorithm that estimates the ground state energy of a Hamiltonian to $O(1)$ (in n) precision, given an initial state with only some constant overlap with the ground state.

Applying the ESPRIT analysis to Lemma 6.2.6, we will obtain the following Theorem:

Theorem 6.1.5. *Given a local Hamiltonian on n qubits. Let the number of eigenvectors supported in some efficient-to-sample (*poly(n) *effort) input state* $|\Phi\rangle$ *be* $S = O(1)$*, and each occurs with nonzero probability at least* $1/\text{poly}(n)$*. In addition, assume that for a fixed* x, y *it is efficient to compute* $\frac{\Phi(y)}{\Phi(x)}$*. Furthermore, assume that the S eigenvalues* $\{E_i\}, E_i \in [0, \pi]$ *are sufficiently well-separated, i.e. at least by* $\Delta \geq 1/\text{poly}(n)$ *with some* poly(n)*. Then using Lemma 6.2.6 plus signal post-processing via ESPRIT, each requiring (some)* poly(n) *classical effort, one can resolve the eigenvalues* $\{E_j\}$ *with distance* $d(\{E_i\}, \{\tilde{E}_j\})$ *at most* $1/\text{poly}(n)$*.*

In [22], it was additionally shown that estimating the smallest eigenvalue of a local Hamiltonian (with some restrictions on its locality) with inverse polynomial precision, even when provided that the guiding (input) state (which is a state with access structure similar to the one considered in this work) has close to a constant $(1/2 - \Omega(1/\text{poly}(n)))$ overlap with the ground state, is BQP-complete. Improving on this result, the authors of [23] have shown that this problem is BQP-complete even when the overlap of the guiding state with the ground state is $1 - \Omega(1/\text{poly}(n))$ (a result which was also obtained in [24]) and that equivalent results hold for the task of estimating excited state eigenvalues of the Hamiltonian. As also mentioned in [23], the ability to classically estimate $S = O(1)$ eigenvalues up to inverse polynomial error (provided an at least inverse polynomial spectral gap) – as described in Theorems 6.1.4 and 6.1.5 – thus depends strongly on the number of eigenvectors supported in the input state being *only* $S = O(1)$. Note that in Theorems 6.1.4 and 6.1.5 we furthermore require that the $S = O(1)$ eigenvalues are separated by an at least $1/\text{poly}(n)$ gap.

To investigate practical aspects of the MC scheme for stoquastic Hamiltonians and compare it to the quantum scheme, we numerically study the one-dimensional Ising

chain in a transverse field g [25] in a *proof-of-principle* setting. We numerically study, amongst several other aspects, the recovery of the ground-state and first-excited-state eigenvalues in the $(g > 1)$-regime from the signals $g_R(k)$ and $g_I(k)$ (in the presence of sampling noise and Trotter error) using the ESPRIT method.

An overview of the paper is as follows. In Section 6.2, we review the *Hadamard* or *overlap* quantum subroutine (Lemma 6.2.1) and we present the Monte Carlo algorithm (Lemma 6.2.3) for stoquastic Hamiltonians with its proof, as well as stating a straightforward Lemma 6.2.6 on 'dequantization'. Section 6.3 reviews the ESPRIT method and has an extensive 6.D in which we prove the performance of the ESPRIT method for imaginary-time decaying signals using many lemmas also needed in the real-time signal case. The arguments for Theorem 6.1.5 are presented in Section 6.3 as well. In Section 6.4, we numerically compare the quantum scheme and the Monte Carlo scheme (for stoquastic Hamiltonians) for determining part of the spectrum of a transverse field Ising chain. In Section 6.5, we discuss our work and propose some directions for future study. Several appendices give additional background information and details.

We note that very extensive literature exists on the Monte Carlo power method [6] in which one applies a sequences of steps which gradually project an initial input state onto the ground state. In this method, unlike in our MC scheme of Lemma 6.2.3, one renormalizes the state after each iteration, so that the signal does not die out. In our approach, we do not renormalize, but study the decay rates themselves. In terms of other previous work, we note that in [12] the ground state energy of a stoquastic Hamiltonian was efficiently estimated by means of a projector Monte Carlo scheme, under an additional 'guiding state' promise. In [26] the authors consider the implementation of the imaginary-time evolution $\exp(-\tau H)$ on a quantum computer in order to prepare a ground state of any local Hamiltonian. Note that our goal is not to prepare any ground or excited state but rather only learn some eigenvalues.

In the remainder of this section, we will review a few definitions which are used in this paper.

Definition 6.1.6. ***Stoquastic Hamiltonians*** *A (real-valued) Hamiltonian H is (globally) stoquastic [8] in a basis $\mathcal{B}$ if all its off-diagonal elements are non-positive: $\langle x| H \left|y\right\rangle \leq 0$, for $x \neq y$ (and states $|x\rangle$, $\left|y\right\rangle$ being elements of basis $\mathcal{B}$).*

In this work, we are interested in Hamiltonians that are *local* and *stoquastic*:

Definition 6.1.7. ***Local Hamiltonians*** *A Hamiltonian H associated with a system consisting of n degrees of freedom (e.g. spins/qubits) is local if it admits a decomposition into a set of Hermitian operators $\{H_i\}$ – i.e. $\sum_i^N H_i$ – such that each H_i acts non-trivially on $O(1)$ (not growing with n) degrees of freedom of the system.*

We denote the maximum number of degrees of freedom on which each H_i acts non-trivially (i.e. its *locality*) by k and note that the number of terms in a local Hamiltonian is $N = O(n^k)$.

For local Hamiltonians there is a slightly stronger notion of stoquasticity, called termwise stoquasticity, which can differ from the definition of stoquasticity given above, see [8, 9].

Definition 6.1.8. ***Termwise stoquastic Hamiltonians*** *A (real-valued) k-local Hamiltonian H is m-termwise stoquastic in a basis $\mathcal{B}$ if it admits a decomposition into (real-valued) $m(\geq k)$-local terms $\{H_a\}$ such that each H_a is stoquastic: $\forall a$, $\langle x| H_a \left|y\right\rangle \leq 0$, for $x \neq y$ (and states $|x\rangle$, $\left|y\right\rangle$ being elements of basis $\mathcal{B}$).*

Most many-body Hamiltonians considered in physics which are stoquastic are $O(1)$-termwise stoquastic. The results in this paper apply to both termwise stoquastic as well as globally stoquastic Hamiltonians (using some small adaptions employing results in [9]), and we will refer to them simply as 'stoquastic'.

For a matrix X we will use the operator or spectral norm $\|X\| = \sqrt{\lambda_{\max}(X^\dagger X)} = \sigma_{\max}(X)$, where $\sigma_{\max}(X)$ is the largest singular value of X. We also refer to the Frobenius norm $\|X\|_F = \sqrt{\mathrm{Tr}(X^\dagger X)}$ and the induced$-\infty$ norm $\|X\|_\infty = \max_i \sum_j |X_{ij}|$. For an $m \times n$ matrix X, we use $\|X\| \leq \sqrt{m}\|X\|_\infty$ and $\|X\| \leq \|X\|_F$.

6.2. Quantum scheme versus Monte Carlo scheme for spectral estimation

In this section we show how to estimate $g_R(k)$ on a quantum computer, and $g_I(k)$ for stoquastic Hamiltonians via a Monte Carlo algorithm, as well as how to estimate $g_D(k)$ inefficiently (in k) via a classical algorithm for general local Hamiltonians.

Lemma 6.2.1 states a well-known quantum subroutine, namely the Hadamard or overlap test, while a new result, a Monte Carlo version of the routine, is proved in Lemma 6.2.3. After these Lemmas, the proofs of Theorems 6.1.1 and 6.1.2 are given. Then we give Lemma 6.2.6 for general local Hamiltonians, using similar tools as in Lemma 6.2.3.

We note that the overlap test is used in versions of quantum phase estimation which do not aim at preparing an energy eigenstate of the Hamiltonian, but rather only learn the spectral content in its input state, as in Refs. [3, 4, 27]. Here we basically follow this approach for the real-time quantum evolution, which can in addition be randomized to save on implementation costs, see [28].

Lemma 6.2.1 (Hadamard or Overlap Test)**.** *Let $\mathcal{F} \equiv \langle\Phi| G_1 G_2 \ldots G_L |\Phi\rangle$, where:*

1. *$|\Phi\rangle = \sum_{x=1}^{2^n} \Phi(x)|x\rangle$ is a state of n qubits which can be generated by a* poly(n)*-size quantum circuit.*
2. *Each G_l is a k-local unitary matrix.*

$\mathcal{F}$ can be estimated within error ϵ with probability at least $1-\delta$ with a quantum circuit with $\Theta(\epsilon^{-2}\log(4\delta^{-1})) \times \left[\Theta(L) + poly(n)\right]$ single and two-qubit gates.

Proof. Figure 6.1 depicts the quantum circuit which is used. It involves an n-qubit register and a single ancillary qubit. The state of the composite system can be tracked through the circuit and the final state can be found to be (where $R(\theta) \equiv e^{-i\theta Z/2}$):

$$\frac{1}{2}\Big(\left(e^{-i\theta/2} I + e^{i\theta/2} G_1 G_2 \ldots G_L\right)|0\rangle_a \otimes |\Phi\rangle + \left(e^{-i\theta/2} I - e^{i\theta/2} G_1 G_2 \ldots G_L\right)|1\rangle_a \otimes |\Phi\rangle\Big). \quad (6.4)$$

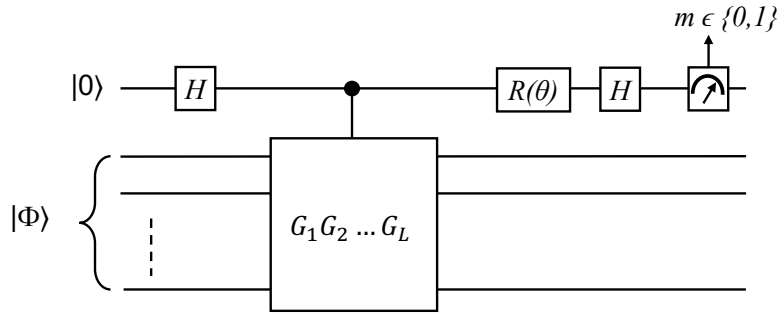

Figure 6.1: Basic circuit with a single ancillary qubit and an n-qubit register (initialized in state $|\Phi\rangle$).

A Z-measurement is now performed on the ancillary qubit, measuring either $|0\rangle$ or $|1\rangle$ with associated outcomes resp. $m=0$ or $m=1$. The probability to measure state $|0\rangle$ ($m=0$) on the ancillary qubit after application of the depicted gates is then given by:

$$\Pr\big(m=0\,|\,\theta\big)=\frac{1}{2}+\frac{1}{4}\Big(e^{i\theta}\,\langle\Phi|\,G_1G_2\ldots G_L\,|\Phi\rangle+e^{-i\theta}\big(\,\langle\Phi|\,G_1G_2\ldots G_L\,|\Phi\rangle\,\big)^*\Big)=\begin{cases}\frac{1}{2}+\frac{1}{2}\mathrm{Re}\Big(\,\langle\Phi|\,G_1G_2\ldots G_L\,|\Phi\rangle\Big), \text{ for } \theta=0,\\ \frac{1}{2}-\frac{1}{2}\mathrm{Im}\Big(\,\langle\Phi|\,G_1G_2\ldots G_L\,|\Phi\rangle\Big), \text{ for } \theta=\frac{\pi}{2}.\end{cases} \tag{6.5}$$

In the final expression, we have restricted ourselves to $\theta=0$ and $\theta=\frac{\pi}{2}$, which are the θ values of interest. Suppose that for $\theta=0$ and $\theta=\pi/2$, the quantum circuits are repeated $|\Sigma|$ times to obtain a set $2|\Sigma|$ of independent realizations of the ancillary-qubit state to be measured and let $\big|\Sigma_0^{\theta=0}\big|$ and $\big|\Sigma_0^{\theta=\pi/2}\big|$ be the number of times the ancilla measurement returns 0 so that

$$\tilde{\mathcal{F}}=\left(2\frac{\big|\Sigma_0^{\theta=0}\big|}{|\Sigma|}-1\right)-i\left(2\frac{\big|\Sigma_0^{\theta=\pi/2}\big|}{|\Sigma|}-1\right) \tag{6.6}$$

is our (unbiased) estimator, i.e. $\mathbb{E}(\tilde{\mathcal{F}}(t))=\mathcal{F}(t)$. Then by means of the Chernoff bound we have

$$\begin{aligned}\Pr\Big(\Big|\tilde{\mathcal{F}}-\mathcal{F}\Big|\le\epsilon\Big)&\ge\Pr\Big(\Big|\mathrm{Re}(\tilde{\mathcal{F}}-\mathcal{F})\Big|\le\epsilon/\sqrt{2}\Big)\Pr\Big(\Big|\mathrm{Im}(\tilde{\mathcal{F}}-\mathcal{F})\Big|\le\epsilon/\sqrt{2}\Big)\\&=\Big(1-\Pr\Big(\Big|\mathrm{Re}(\tilde{\mathcal{F}}-\mathcal{F})\Big|\le\epsilon/\sqrt{2}\Big)\Big)\Big(1-\Pr\Big(\Big|\mathrm{Im}(\tilde{\mathcal{F}}-\mathcal{F})\Big|\le\epsilon/\sqrt{2}\Big)\Big)\\&\ge\Big[\max\Big(0,\big(1-2\exp(-|\Sigma|\epsilon^2/4)\big)\Big)\Big]^2\\&\ge 1-4\exp\big(-|\Sigma|\epsilon^2/4\big)=1-\delta,\end{aligned} \tag{6.7}$$

where the number of samples is chosen as $|\Sigma|=\Theta(\epsilon^{-2}\log(4\delta^{-1}))$. □

Next, we will consider a classical Monte Carlo version of the quantum routine given above. The key result here is Lemma 6.2.3. However, before we can state it, we collect

a few facts about matrices G_i which will be useful in the proof of Lemma 6.2.3. The matrices G_i that we will consider now can be seen as analogous to the (unitary) local real-time propagation operators G_i considered earlier in Lemma 6.2.1, but are now local *imaginary-time* propagation operators. The G_i are no longer unitary but, for local stoquastic Hamiltonians, are elementwise nonnegative. They are of the form $G_i = e^{-a_i/M\,kH_i}$, where a_i/M is a positive parameter set by the Trotterization scheme and $k = 0, 1, \ldots, K$ denotes imaginary-time coordinate. We have the following proposition on further properties of these operators:

Proposition 6.2.2. *Let $G_i = e^{-a_i/M\,kH_i}$, where H_i is a stoquastic Hermitian matrix (a term in $H = \sum_i H_i$) which acts nontrivially on some subset of $O(1)$ qubits. Let the smallest eigenvalue of H_i be 0, i.e. $\lambda_{min}(H_i) = 0$. We have:*

- *The matrix G_i is an elementwise nonnegative and positive definite matrix, with eigenvalues in the interval $(0,1]$, acting nontrivially only on the same $O(1)$ qubits as H_i.*
- *If G_i is reducible, then we can write $G_i = \oplus_{b=1}^{B_i} G_i^b$ with B_i irreducible sub-matrices G_i^b. The set of bit string basis states on which the irreducible sub-matrix G_i^b acts is denoted by S_i^b, where $\cup_b S_i^b \subseteq \{0,1\}^n$.*
- *From the Perron-Frobenius Theorem (Theorem 8.4.4 in [29]) it follows that for each nonnegative and irreducible sub-matrix there exists a unique and strictly positive eigenstate associated with its largest eigenvalue, i.e.*

$$\left|\phi_i^b\right\rangle = \sum_{x \in S_i^b} \phi_i^b(x)\,|x\rangle\,, \quad G_i^b \left|\phi_i^b\right\rangle = \lambda_i^b \left|\phi_i^b\right\rangle, \tag{6.8}$$

  *where $\phi_i^b(x) > 0$, $\forall x \in S_i^b$. Since the spectrum of G_i is the union of spectra of the submatrices G_i^b, the spectrum of each G_i^b also lies in the interval $(0,1]$ and one of the blocks b will contain the largest eigenvalue of G_i equal to 1. In case G_i is irreducible itself, there is a largest nonnegative eigenvector as in Eq. equation (**??**) which has support $\phi_i(x) > 0$ for all x. In this case, the corresponding eigenvalue will be $\lambda_i = 1$.*
- *Naturally, since G_i acts nontrivially only on a subset of $O(1)$ qubits (and acts as I on other qubits) one can efficiently compute the blocks G_i^b, its largest eigenvalue λ_i^b and associated eigenstate $\left|\phi_i^b\right\rangle$ in each block b.*

We prove the following:

Lemma 6.2.3. *Let $\mathcal{F} \equiv \langle\Phi|\, G_1 G_2 \ldots G_L\, |\Phi\rangle$, where:*

1. *$|\Phi\rangle = \sum_{x=1}^{2^n} \Phi(x)\,|x\rangle$ is a normalized state of n qubits where $\Phi(x) \in \mathbb{C}$ $(\forall x)$ and $\sum_x \left|\Phi(x)\right|^2 = 1$. We assume that (1) $\frac{\Phi(y)}{\Phi(x)}$ can be efficiently (*poly(n)*) calculated for a* given *x and y and (2) we can efficiently draw samples from the probability distribution $P(x) = \left|\Phi(x)\right|^2$.*

2. *Each $G_l = G_l$ is a k-local, positive-definite, (elementwise) nonnegative matrix with eigenvalues in* $(0,1]$.

$\mathcal{F}$ can be estimated within error ϵ with probability at least $1-\delta$ with a classical MC algorithm with runtime $\Theta(\epsilon^{-2}log(\delta^{-1}))\times poly(n)\times\Theta(L)$.

Proof. The proof of Lemma 6.2.3 consists of two steps: To construct an estimator for $\mathcal{F}(\tau)$ and to show that the error of this estimator can be bounded according to the lemma.

We rewrite the quantity of interest $\mathcal{F}$ as follows (where $L-1$ complete sets of basis states are inserted in between the G_l operators in the final equality):

$$\mathcal{F} = \sum_{x_0,x_1,\ldots,x_L} |\Phi(x_0)|^2 \frac{\Phi(x_L)}{\Phi(x_0)} \langle x_0|G_1|x_1\rangle \langle x_1|G_2|x_2\rangle \ldots \langle x_{L-1}|G_L|x_L\rangle, \tag{6.9}$$

where we have set $|x\rangle = |x_0\rangle$ and $|y\rangle = |x_L\rangle$. $\mathcal{F}$ thus corresponds to the sum of an exponential number of products of (non-negative) matrix elements of $G_1,\ldots,G_L$, weighted by amplitudes in the state $|\Phi\rangle$. Evidently, only terms for which all the matrix elements in the product are non-zero contribute to the sum.

We now consider the string of basis states $|x_0\rangle,\ldots,|x_L\rangle$ and associate with each step $|x_{l-1}\rangle$ to $|x_l\rangle$ in this string a probability

$$P_l(x_{l-1}\to x_l) = \frac{1}{\lambda_l^b}\langle x_{l-1}|G_l|x_l\rangle \frac{\phi_l^b(x_l)}{\phi_l^b(x_{l-1})}, \tag{6.10}$$

where b labels the sub-block in $G_l = \oplus_b G_l^b$ which contains the strings x_{l-1} and x_l. Here $\phi_l^b(x) \equiv \langle x_{l-1}|\phi_l^b\rangle$ with $|\phi_l^b\rangle$ defined in Proposition 6.2.2.

The probability distribution P_l is thus non-negative as $\lambda_l^b \in (0,1]$, G_l is element-wise non-negative and $\phi_l^b(x_l) > 0$ and $\phi_l^b(x_{l-1}) > 0$. It can be shown to be normalized:

$$\begin{aligned}\sum_{x_l} P_l(x_{l-1}\to x_l) &= \sum_{x_l}\frac{1}{\lambda_l^b}\langle x_{l-1}|G_l|x_l\rangle\frac{\phi_l^b(x_l)}{\phi_l^b(x_{l-1})} = \sum_{x_l\in S_l^b}\frac{1}{\lambda_l^b}\langle x_{l-1}|G_l^b|x_l\rangle\frac{\phi_l^b(x_l)}{\phi_l^b(x_{l-1})}\\ &= \frac{1}{\phi_l^b(x_{l-1})}\frac{1}{\lambda_l^b}\langle x_{l-1}|G_l^b|\phi_l^b\rangle = \frac{1}{\phi_l^b(x_{l-1})}\frac{1}{\lambda_l^b}\langle x_{l-1}|\lambda_l^b|\phi_l^b\rangle = \frac{\langle x_{l-1}|\phi_l^b\rangle}{\phi_l^b(x_{l-1})} = 1.\end{aligned} \tag{6.11}$$

We use $P_l(x_{l-1}\to x_l)$ to rewrite $\mathcal{F}(\tau)$ as:

$$\begin{aligned}\mathcal{F}(\tau) = \sum_{x_0,x_1,\ldots,x_L} &\underbrace{|\Phi(x_0)|^2 P_1(x_0\to x_1)P_2(x_1\to x_2)\ldots P_L(x_{L-1}\to x_L)}_{\equiv\Pi(\boldsymbol{x}).}\\ &\times\underbrace{\frac{\Phi(x_L)}{\Phi(x_0)}\prod_{l=1}^{L}\lambda_l^{b(l)}\frac{\phi_l^{b(l)}(x_{l-1})}{\phi_l^{b(l)}(x_l)}}_{\equiv\mathcal{R}(\boldsymbol{x}).},\end{aligned} \tag{6.12}$$

where $\boldsymbol{x} \equiv (x_0, x_1, ..., x_L)$ and we have defined the quantities $\Pi(\boldsymbol{x})$ and $\mathcal{R}(\boldsymbol{x})$. Since $|\Phi(x_0)|^2$ and each P_l are probability distributions, $\Pi(\boldsymbol{x})$ is a probability distribution as well, i.e.

$$\sum_{x_0,x_1,...,x_L} \Pi(\boldsymbol{x}) = \sum_{x_0} \Big(|\Phi(x_0)|^2 \sum_{x_1} \Big(P_1(x_0 \to x_1) ... \sum_{x_L} \Big(P_L(x_{L-1} \to x_L) \Big) ... \Big) \Big) = 1. \quad (6.13)$$

Clearly, one can sample from $\Pi(\boldsymbol{x})$ by first sampling from $|\Phi(x_0)|^2$, then sampling from $P_1(x_0 \to x_1)$ to generate x_1 etc. until x_L.

By thus sampling from the probability distribution $\Pi(\boldsymbol{x})$ and obtaining a mean estimator for $\mathcal{F}(\tau)$ using the samples $\mathcal{R}(\boldsymbol{x})$, we can estimate $\mathcal{F}(\tau)$. We note that $\mathcal{F}(\tau) = \mathbb{E}\big(\mathcal{R}(\boldsymbol{x})\big)$. Since $\mathcal{R}(\boldsymbol{x}) \in \mathbb{C}$, a mean estimator over a finite number of samples will generally be complex-valued. Since $\mathcal{F}(\tau) \in \mathbb{R}$, we will instead obtain a mean estimator using samples $\mathrm{Re}(\mathcal{R}(\boldsymbol{x}))$. The mean estimator that we shall use to estimate $\mathcal{F}(\tau)$ is the median-of-means estimator [30]. Using a set Σ of samples $\{\boldsymbol{x}\}$ (distributed according to $\Pi(\boldsymbol{x})$), the median-of-means estimator is defined as follows: Divide the set Σ into q subsets $s_1, \ldots, s_q$ of size approximately $|\Sigma|/q$. Calculate the empirical mean of $\mathrm{Re}(\mathcal{R}(\boldsymbol{x}))$ over the samples in each subset: $f_j = \frac{1}{|s_j|} \sum_{\boldsymbol{x} \in s_j} \mathrm{Re}(\mathcal{R}(\boldsymbol{x}))$ for $j \in \{1, \ldots, q\}$ (each f_j is an unbiased estimator of $\mathcal{F}(\tau)$). Now the median-of-means estimator is given by the median of these empirical means: $\hat{\mathcal{F}} = \mathrm{M}(f_1, \ldots, f_q)$. See 6.C for more details.

The algorithm that efficiently produces $\hat{\mathcal{F}} = \mathrm{M}(f_1, \ldots, f_q)$ is explicitly given in Algorithm 6.1 below. Note that when $\Phi(x_0)$ is small for some x_0, the probability of drawing this x_0, $|\Phi(x_0)|^2$, is very small, but the ratio $\frac{\Phi(x_L)}{\Phi(x_0)}$ in the estimator could get very large.

Algorithm 6.1: Efficiently obtaining a median-of-means estimate of $\mathcal{F}(\tau)$ through sampling of the probability distribution $\Pi(\boldsymbol{x})$.

Input: Initial state $|\Phi\rangle$. Local propagation operators $\{G_l\}_{l=1}^L$. Sample size $|\Sigma|$. Number of subsets q.

Output: Median of means estimate of $\mathcal{F}(\tau)$.

for $\sigma \in \{1, 2, ..., |\Sigma|\}$ **do**

  Sample an initial basis state $|x_0\rangle$ from the probability distribution $|\Phi(x_0)|^2$. The state $|x_0\rangle$ is part of $S_1^{b(1)}$.

  for $l \in \{1, ..., L\}$ **do**

    Pick a state $|x_l\rangle \in S_l^{b(l)}$ with probability

    $P_l(x_{l-1} \to x_l) = \frac{1}{\lambda_l^{b(l)}} \langle x_{l-1}| G_l |x_l\rangle \frac{\phi_l^{b(l)}(x_l)}{\phi_l^{b(l)}(x_{l-1})}$. The state $|x_l\rangle$ is part of $S_{l+1}^{b(l+1)}$ (for $l < L$).

  end

  Given $\{x_l\}_{l=1}^L$ (sampled from $\Pi(\boldsymbol{x})$), calculate $\mathcal{R}_\sigma(\boldsymbol{x}) = \frac{\Phi(x_L)}{\Phi(x_0)} \prod_{l=1}^L \lambda_l^{b(l)} \frac{\phi_l^{b(l)}(x_{l-1})}{\phi_l^{b(l)}(x_l)}$.

end

Divide the $|\Sigma|$ samples into q subsets $s_1, \ldots, s_q$, such that $|s_j| \approx |\Sigma|/q$, $\forall j$.

for $j \in \{1, 2, \ldots, q\}$ **do**

  Calculate $f_j = \frac{1}{|s_j|} \sum_{\boldsymbol{x} \in s_j} \mathrm{Re}(\mathcal{R}(\boldsymbol{x}))$.

end

Output $\hat{\mathcal{F}} = \mathrm{M}(f_1, \ldots, f_q)$.

Algorithm 6.1 thus efficiently provides an estimate of $\mathcal{F}(\tau)$ (albeit biased). To complete the proof, we will show that the variance of $\mathrm{Re}(\mathcal{R}(\boldsymbol{x})) \in \mathbb{R}$ can be bounded which in turn is used to bound the number of samples to get an estimate close to the mean, leading to Lemma 6.2.3.

For a complex random variable $Z = \mathcal{R}(\boldsymbol{x})$, $\mathbb{E}(Z) \equiv \mathbb{E}\big(\mathrm{Re}(Z)\big) + i\mathbb{E}\big(\mathrm{Im}(Z)\big)$ and $\mathrm{Var}\big(Z\big) = \mathrm{Var}\big(\mathrm{Re}(Z)\big) + \mathrm{Var}\big(\mathrm{Im}(Z)\big) \geq \mathrm{Var}\big(\mathrm{Re}(Z)\big)$. Hence we can bound the variance of random variable $\mathrm{Re}(\mathcal{R}(\boldsymbol{x}))$ by bounding the variance of the random variable $\mathcal{R}(\boldsymbol{x})$. This variance is given by:

$$\mathrm{Var}(\mathcal{R}(\boldsymbol{x})) = \mathbb{E}\Big(\big|\mathcal{R}(\boldsymbol{x})\big|^2\Big) - \underbrace{\Big|\mathbb{E}\Big(\mathcal{R}(\boldsymbol{x})\Big)\Big|^2}_{=|\mathcal{F}(\tau)|^2=\mathcal{F}(\tau)^2.} \leq \mathbb{E}\Big(\big|\mathcal{R}(\boldsymbol{x})\big|^2\Big), \tag{6.14}$$

where the inequality holds because $\mathcal{F}^2 \geq 0$ (since $\mathcal{F} \in \mathbb{R}$). To obtain an upper bound on the variance, we shall investigate this expression in more detail:

$$\begin{aligned}
\mathbb{E}\Big(\big|\mathcal{R}(\boldsymbol{x})\big|^2\Big) &= \sum_{\boldsymbol{x}} \Pi(\boldsymbol{x})\big|\mathcal{R}(\boldsymbol{x})\big|^2 \\
&= \sum_{\boldsymbol{x}} |\Phi(x_L)|^2 \langle x_0|\, G_1\, |x_1\rangle \langle x_1|\, G_2\, |x_2\rangle \ldots \langle x_{L-1}|\, G_L\, |x_L\rangle \prod_{l=1}^{L} \lambda_l^{b(l)} \frac{\phi_l^{b(l)}(x_{l-1})}{\phi_l^{b(l)}(x_l)} \\
&= \sum_{\boldsymbol{x}} |\Phi(x_L)|^2\, Q_1(x_0,x_1) Q_2(x_1,x_2) \ldots Q_L(x_{L-1},x_L),
\end{aligned} \tag{6.15}$$

where in the last equality we defined the non-negative quantity $Q_l(x,y) \equiv \langle x|\, G_l\, \big|y\big\rangle \lambda_l^{b(l)} \frac{\phi_l^{b(l)}(x)}{\phi_l^{b(l)}(y)}$. Exploiting the Hermiticity of G_l^b, $Q_l(x,y)$ can be shown to have the following property:

$$\sum_x Q_l(x,y) = \sum_x \langle x|\, G_l\, \big|y\big\rangle \lambda_l^b \frac{\phi_l^b(x)}{\phi_l^b(y)} = \sum_{x \in S_l^b} \langle x|\, G_l^b\, \big|y\big\rangle \lambda_l^b \frac{\phi_l^b(x)}{\phi_l^b(y)} = \big(\lambda_l^b\big)^2 \frac{\langle \phi_l^b \big| y\rangle}{\phi_l^b(y)} = \big(\lambda_l^b\big)^2 \leq 1. \tag{6.16}$$

$Q_l(x,y)$ thus satisfies $0 \leq Q_l(x,y) \leq 1$, $\forall x,y$ and $\forall l \in \{1,2,...,L\}$. By consecutively exploiting the property in equation (6.16) for all Q_l's and the normalization property of state $|\Phi\rangle$ in the expression in equation (6.15), we obtain

$$\mathrm{Var}(\mathcal{R}(\mathbf{x})) \leq \mathbb{E}\Big(\big|\mathcal{R}(\boldsymbol{x})\big|^2\Big) \leq 1. \;\Rightarrow\; \mathrm{Var}\Big(\mathrm{Re}(\mathcal{R}(\mathbf{x}))\Big) \leq 1. \tag{6.17}$$

If we take the number of samples $|\Sigma|$, and divide them into q subsets $s_1,\ldots,s_q$ of size approximately $|\Sigma|/q$, then (by means of Chebyshev's inequality) each f_j obeys $|f_j - \mathcal{F}| \leq \sqrt{\mathrm{Var}\big(\mathrm{Re}(\mathcal{R}(\boldsymbol{x}))\big)}\sqrt{4q/|\Sigma|} \leq \sqrt{4q/|\Sigma|}$ with probability at least 3/4. Using Hoeffding's inequality and the definition of the mean, one can show that (see 6.C):

$$\Pr\Big(\big|\hat{\mathcal{F}} - \mathcal{F}\big| \leq \sqrt{4q/|\Sigma|}\Big) \geq 1 - e^{-q/8}. \tag{6.18}$$

Hence $\mathcal{F}$ can be estimated with error ϵ with probability at least $1-\delta$ (with $q = 8\log(\delta^{-1})$) for $|\Sigma| = \Theta(\log(\delta^{-1})\epsilon^{-2})$, where obtaining each sample takes a number of operations that scales linearly in L and poly(n). This completes the proof of Lemma 6.2.3. □

Remark 6.2.4. *Note that if one would have chosen the empirical mean* $\tilde{\mathcal{F}} = \frac{1}{|\Sigma|}\sum_{\boldsymbol{x}\in\Sigma} Re(\mathcal{R}(\boldsymbol{x}))$ *as a mean estimator for* $\mathcal{F}$ *(instead of the median-of-means estimator), then using equation* (6.17) *and Chebyshev's inequality, we obtain:*

$$Pr\Big(\big|\tilde{\mathcal{F}}-\mathcal{F}\big|\leq\epsilon\Big)\geq 1-\frac{\mathrm{Var}(Re(\tilde{\mathcal{F}}))}{\epsilon^2}\geq 1-\frac{1}{|\Sigma|\,\epsilon^2}. \tag{6.19}$$

Hence $\mathcal{F}$ *can be estimated using* $\tilde{\mathcal{F}}$ *with error* ϵ *with probability at least* $1-\delta$*, for* $|\Sigma|=\Theta(\delta^{-1}\epsilon^{-2})$*. Using the median-of-means estimator thus provides an exponential improvement in the required scaling of* $|\Sigma|$ *with* δ^{-1}*. Note that if we could upper and lower bound the range of* $\mathrm{Re}(\mathcal{R}(\mathbf{x}))$ *by some constants, then we could have used a Chernoff-Hoeffding bound for the empirical mean* $\tilde{\mathcal{F}}$ *which gives the aforementioned (exponentially) better dependence of the run-time of the algorithm with* δ^{-1} *(as in Lemma 6.2.1 where we do use a Chernoff-Hoeffding bound).*

Remark 6.2.5. *Note that the Lemma also applies to estimating* $\langle x|\,G_1G_2\ldots G_L\,|x'\rangle$ *(with* $1/\mathrm{poly}(n)$ *accuracy) as one simply starts the process at* $x_0=x$ *and* $\mathcal{R}(\mathbf{x})$ *is only nonzero when one arrives at* $x_L=x'$*. Similarly, one can estimate* $\langle\Phi_1|\,G_1G_2\ldots G_L\,|\Phi_2\rangle$ *with* $1/\mathrm{poly}(n)$ *accuracy, assuming one can sample from* $|\Phi_1(x)|^2$ *(or* $|\Phi_2(x)|^2$*)* and *compute for a given* x *and* y*, the ratio* $\frac{\Phi_2(y)}{\Phi_1(x)}$*. In addition, one can extend the Lemma to the case where the local propagation operators* G_l *are not Hermitian, but are still nonnegative matrices, see 6.B.*

We stress that Lemma 6.2.3 provides an efficient classical algorithm provided that: For a given $x, y\in\{0,1\}^n$, one can efficiently determine $\frac{\Phi(y)}{\Phi(x)}$ and the state $|\Phi\rangle$ is such that one can efficiently draw samples from $P(x)=\big|\Phi(x)\big|^2$. In many practical settings, $|\Phi\rangle$ is such that one can define a function $f:\{0,1\}^n\to\mathbb{C}$ which takes as input the n-bit string x, and efficiently outputs the corresponding coefficient $\Phi(x)$. This is e.g. the case for (matrix) product states or for other ansatz classes of states. Then, given x and y, the fraction $\frac{\Phi(y)}{\Phi(x)}$ can be efficiently obtained. Note that under this assumption one can set-up a Monte Carlo scheme based on the Metropolis algorithm to sample from $\big|\Phi(x)\big|^2$, although this scheme is only a heuristic strategy and its efficient convergence would have to be proved. A good class of states to which both Lemmas 6.2.3 and 6.2.1 apply are of course product states. Note that even when running the overlap test is too costly (as quantum circuits are noisy), but preparing the state $|\Phi\rangle$ is feasible, one could use this preparation to sample from $|\Phi(x)|^2$ for the application of the MC method. Of course the requirement of being able to compute $\frac{\Phi(y)}{\Phi(x)}$ remains. For the transverse field Ising model, an example of a $|\Phi\rangle$ which obeys these conditions will be given in Section 6.4.

Proof of Theorems 6.1.1 and 6.1.2: We require the Trotterization of e^{-ikH} resp. e^{-kH} into a string of local propagation operators G_i which are unitary (in Lemma 6.2.1) resp. Hermitian and non-negative (in Lemma 6.2.3). This non-unique decomposition of e^{-ikH} and e^{-kH} into an ordered string of local propagation operators depends on the Trotterization scheme and is discussed in 6.A (and more extensively in [31]). The Trotterization gives an error ϵ_{trot} (in addition to the sampling error ϵ in Lemmas 6.2.1 and 6.2.3) and the number of local propagation operators (for each sample) L in Lemmas 6.2.1 and 6.2.3 will be $L=\mathrm{poly}(n)\,\mathcal{O}\big(\Upsilon\,k^{1+1/p}\epsilon_{\mathrm{trot}}^{-1/p}\big)$ (for real time) and

$L = \text{poly}(n)\,\mathcal{O}\big(\Upsilon\, k^{1+1/p}\epsilon_{\text{trot}}^{-1/p}\big)$ (for imaginary time, provided that $M \geq 4\tau\Upsilon\big(\sum_\gamma \|H_\gamma\|\big)$, where M is the Trotter variable). Υ denotes the number of *stages* in the Trotterization scheme of order p, and typically scales exponentially in p (but p is chosen a constant). For given order $p = O(1)$ of the Trotterization scheme, L in Lemma 6.2.3 thus scales with the length of the time interval over which the system is simulated as $k^{1+o(1)}$ and $k^{1+o(1)}$, and with the imposed Trotter error as $\epsilon_{\text{trot}}^{-o(1)}$. Then, if we wish to estimate $g_R(k)$ and $g_I(k)$ at multiple $k = 0,\ldots,K$, we use that the probability that all K estimates are up to uncertainty ϵ equals unity minus the probability that at least one of the estimates is beyond ϵ (which, by the union bound, is at most $K\delta$).

Finally, before we move on to extracting eigenenergy estimates from the (real-time and imaginary-time) signals using the ESPRIT method, we prove the Lemma related to the signal $g_D(k)$ in equation (6.3).

Lemma 6.2.6. *Let $g_D(k)$ be defined as in equation* (6.3) *for a local n-qubit Hamiltonian H, with E_j in $[0,\pi]$, and assume that (1) one can efficiently (i.e. with* poly(n) *effort) sample from $|\Phi(x)|^2$, and (2) given x and y, one can compute $\Phi(y)/\Phi(x)$ efficiently. Then, $g_D(k)$ can be classically estimated within error ϵ with probability at least $1-\delta$ with $[\text{poly}(n)]^k \times \Theta(\epsilon^{-2}\log(\delta^{-1}))$ classical computational effort.*

Proof. By definition of $g_D(k)$, we can write

$$g_D(k) = \sum_{x,y} |\Phi(x)|^2 \frac{\Phi(y)}{\Phi(x)} \langle x|\,(I - H/2\pi)^k \,|y\rangle\,. \tag{6.20}$$

To estimate $g_D(k)$, one first draws an x from $P(x) = |\Phi(x)|^2$, and then one collects all y which are obtained after the application of $(I - H/2\pi)^k$ to $\langle x|$. Each application of $I - H/2\pi$ maps the input string onto at most poly(n) new output strings, hence one obtains at most $[\text{poly}(n)]^k$ such y's after k applications. Let $\boldsymbol{x} = (x_k = y, x_{k-1}, \ldots, x_1, x_0 = x)$ be a particular path of strings and let

$$\begin{aligned}\mathcal{R}(x) &\equiv \sum_y \frac{\Phi(y)}{\Phi(x)} \langle x|\,(I - H/2\pi)^k \,|y\rangle \\ &= \sum_{x_1,\ldots,x_{k-1},y} \frac{\Phi(y)}{\Phi(x)} \langle x|\,I - H/2\pi\,|x_1\rangle\langle x_1|\,I - H/2\pi\,|x_2\rangle \ldots \langle x_{k-1}|\,I - H/2\pi\,|y\rangle\,,\end{aligned} \tag{6.21}$$

so that $g_D(k) = \sum_x |\Phi(x)|^2\,\mathcal{R}(x)$. For each x that is sampled from $P(x)$, one thus computes and outputs $\text{Re}(\mathcal{R}(x))$ by summing over the contributions from *all* paths $\boldsymbol{x}$ that start at string x. As in Lemma 6.2.3, we need to establish how many samples $|\Sigma|$ we need to draw from $P(x)$ to obtain $g_D(k)$ within error ϵ with probability at least $1-\delta$. This analysis depends on the variance of the complex variable $\mathcal{R}(x)$ through equation (6.14), requiring

us to upper bound

$$\begin{aligned}\mathbb{E}\big(|\mathcal{R}(x)|^2\big) &= \sum_x |\Phi(x)|^2 \Big(\sum_y \frac{\Phi^*(y)}{\Phi^*(x)} \big\langle y\big| \left(I - H/2\pi\right)^k |x\rangle\Big)\Big(\sum_{y'} \frac{\Phi(y')}{\Phi(x)} \langle x| \left(I - H/2\pi\right)^k \big|y'\big\rangle\Big) \\ &= \sum_x \langle\Phi| \left(I - H/2\pi\right)^k |x\rangle \langle x| \left(I - H/2\pi\right)^k |\Phi\rangle \\ &= \langle\Phi| \left(I - H/2\pi\right)^{2k} |\Phi\rangle \leq 1, \end{aligned} \tag{6.22}$$

where in the final line we have used that the eigenvalues of H lie in $[0,\pi]$. As in the proof of Lemma 6.2.3, this establishes that $\mathrm{Var}\big(\mathrm{Re}(\mathcal{R}(x))\big) \leq 1$. Then we can use the median-of-means estimator as in the proof of Lemma 6.2.3 and 6.C to establish that with probability at least $1-\delta$, $g_D(k)$ can be estimated with error at most ϵ, taking $|\Sigma| = \Theta(\epsilon^{-2}\log(\delta^{-1}))$ samples from $P(x) = |\Phi(x)|^2$, and with $[\mathrm{poly}(n)]^k$ computational effort per sample. □

It is important to note that unlike in Lemma 6.2.3, here we only sample x and compute the rest as the estimator $\mathcal{R}(x)$, while in Lemma 6.2.3 we sample the whole path of length L. This is why the computational effort in Lemma 6.2.3 is efficient (linear) in L and thus polynomial in k, while in Lemma 6.2.6 the computational effort is exponential in k. This is thus the difference between the stoquastic Hamiltononian case versus the general Hamiltonian case. Note also that one can take each G_i in Lemma 6.2.3 to be $G = I - H/2\pi$ in principle, as it obeys condition (ii) when H is stoquastic.

We note that in [22] the sampling-access assumption is formulated slightly differently, that is, one gets access to $\Phi(x)$ for a given x, which can be stronger than only knowing the ratio $\Phi(x)/\Phi(y)$ for a given x and y. In addition, Ref. [22] allows an additional error in the sampling access whereas we gloss over this here and assume perfect sampling-access (similar to the exact assumptions in the other Lemmas).

6.3. Classically processing the signal: the ESPRIT method

We turn to discussing the ESPRIT method [32] which is a method like the matrix pencil method [15, 16, 33] for processing a signal as in equation (6.1) and equation (6.2) consisting of S components. Indeed, suppose a set of values for the signal $g(k)$,

$$g(k) = \sum_{j=1}^{S} c_j z_j^k, \tag{6.23}$$

where $|z_j| \leq 1$, for $k \in \{0,1,...,K\}$, K even. The goal is to determine the z_j and the coefficients $c_j > 0$ [1] using $g(k)$ for sufficiently many k. In case of the real-time signal $g_R(k)$, we have $z_j \equiv e^{-iE_j}$, in case of a purely-decaying imaginary-time signal $g_I(k)$, we have $z_j \equiv e^{-E_j} \in (e^{-2\pi}, 1]$ and for the purely-decaying signal $g_D(k)$ we have $z_j \equiv \big(1 - E_j/(2\pi)\big) \in [\frac{1}{2}, 1]$.

[1] Here we focus on determining the z_j, but given the z_j one can determine the c_j as well and methods for analyzing the performance also exist for this [34].

Due to sampling and Trotter noise, one is effectively given a noisy signal $y(k)$ (for $k \in \{0,1,...,K\}$), which is related to the original signal $g(k)$ by:

$$y(k) := g(k) + \eta(k) = \sum_{j=1}^{S} c_j z_j^k + \eta(k), \tag{6.24}$$

where $\eta(k)$ denotes e.g. the sampling and Trotter noise, and we have $|\eta(k)| \leq \epsilon_{\text{tot}}$ in Theorem 6.1.1 and 6.1.2 with high probability.

It is well-known that for a noiseless signal ($\eta(k) = 0$), the z_j's and the c_j's can be resolved perfectly via ESPRIT and the matrix pencil method if we take $K+1 \geq 2S$. Importantly, this result does not depend on whether the signal is oscillatory or decaying. For illustration, Figure 6.2 depicts the results of application of the matrix pencil method to a noiseless signal. We consider separately a decaying signal and an oscillating signal, and for both cases we depict respectively the estimates of the decay rates and oscillation frequencies as a function of K. When $K+1 \geq 2S$, the eigenvalues are indeed resolved both for the decaying and the oscillating signal. When $K+1 < 2S$, the eigenvalues are not resolved. We will see however, that in the presence of noise a decaying or oscillatory signal fares very differently.

Let us consider in more detail the task of obtaining the z_j's from the signal $y(k)$ in equation (6.24). The key object of study here is the Hankel matrix $H(y) := H(g) + H(\eta)$, containing all K data points of the noisy signal $y(k)$ and a positive integer 'matrix pencil' parameter L:

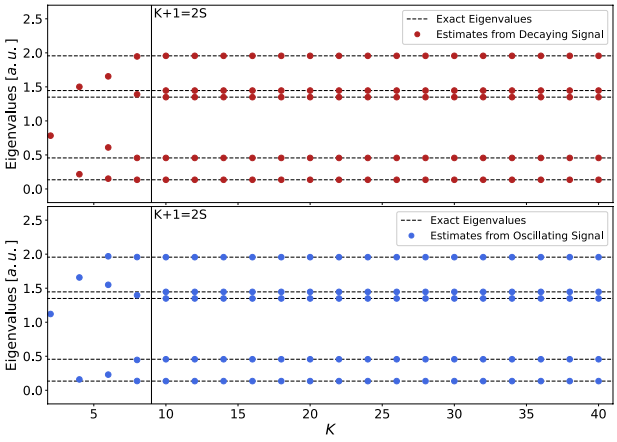

Figure 6.2: Estimates of the decay rates (of a decaying signal) and oscillation frequencies (of an oscillating signal) as a function of K. The estimates are obtained from applying the matrix pencil method [15, 16, 33] to the noiseless signals $g(k) = \sum_{j=1}^{S} c_j z_j^k$, where $z_j = e^{-E_j}$ (in the case of the decaying signal) and $z_j = e^{-iE_j}$ (in the case of the oscillating signal) for $k = 0,1,\ldots,K$ where $E_j \in [0,2\pi)$. All c_j's are set equal to $1/S$ and the E_j's have been randomly produced. The eigenvalues are recovered for $K+1 \geq 2S$.

$$H(y) = \begin{pmatrix} y(0) & y(1) & \dots & y(K-L) \\ y(1) & y(2) & \dots & y(K-L+1) \\ \vdots & \vdots & & \vdots \\ y(L) & y(L+1) & \dots & y(K) \end{pmatrix}_{(L+1)\times(K-L+1)}$$
$$= \sum_{j=1}^{S} c_j \begin{pmatrix} 1 & z_j & \dots & z_j^{K-L} \\ z_j & z_j^2 & \dots & z_j^{K-L+1} \\ \vdots & \vdots & & \vdots \\ z_j^L & z_j^{L+1} & \dots & z_j^K \end{pmatrix} + H(\eta), \quad (6.25)$$

where $H(\eta)$ is purely due to the noise and has norm $\|H(\eta)\|$. We can decompose the Hankel matrix $H(g)$ of the noiseless signal in terms of Vandermonde matrices V_L:

$$H(g) = \sum_{j=1}^{S} c_j \begin{pmatrix} 1 & z_j & \dots & z_j^{K-L} \\ z_j & z_j^2 & \dots & z_j^{K-L+1} \\ \vdots & \vdots & & \vdots \\ z_j^L & z_j^{L+1} & \dots & z_j^K \end{pmatrix}_{(L+1)\times(K-L+1)} = V_L C V_{K-L}^T, \quad (6.26)$$

where

$$C \equiv \mathrm{diag}(c_1, c_2, ..., c_S), \quad (6.27)$$

and V_L is

$$V_L = \begin{pmatrix} 1 & 1 & \dots & 1 \\ z_1 & z_2 & \dots & z_S \\ \vdots & \vdots & & \vdots \\ z_1^L & z_2^L & \dots & z_S^L \end{pmatrix}_{(L+1)\times S}. \quad (6.28)$$

In general, methods such as ESPRIT (see the ESPRIT Algorithm 6.2) rely on the parameter L and for convenience we will keep it general in some of the analysis (specifically in 6.D). Our results will, however, focus on the choice $L = K/2$. For $L = K/2$, we have

$$H(y) = \begin{pmatrix} y(0) & y(1) & \dots & y(K/2) \\ y(1) & y(2) & \dots & y(K/2+1) \\ \vdots & \vdots & & \vdots \\ y(K/2) & y(K/2+1) & \dots & y(K) \end{pmatrix}_{(K/2+1)\times(K/2+1)} \quad (6.29)$$

and

$$H(g) = V_{K/2} C V_{K/2}^T \in \mathbb{C}^{(K/2+1)\times(K/2+1)}. \quad (6.30)$$

Making contact with error bounds in the previous section, we see that (for $L = K/2$)

$$\forall k,\ \eta(k) \le \epsilon_{\mathrm{tot}} \Rightarrow \|H(\eta)\| = \sigma_{\max}(H(\eta)) \le \|H(\eta)\|_F \le K\epsilon_{\mathrm{tot}}. \quad (6.31)$$

From the 'Vandermonde decomposition' in equation (6.26) of the Hankel matrix encoding a real-time or imaginary-time signal, one can develop numerical algorithms to extract the the decay rates z_i. One such algorithm is ESPRIT (given in Algorithm 6.2), which specifically exploits the relation between the Vandermonde decomposition of $H(y)$ and its singular value decomposition.

Algorithm 6.2: ESPRIT algorithm.

Data: Time signal y, number of decay rates or oscillation frequencies S.
Result: List $\tilde{z}_1, \ldots, \tilde{z}_S$.
$K \leftarrow \text{length}(y)$; /* We will assume K is even for simplicity. */
$L \leftarrow K/2$; /* Not the most general choice, however it works well in practice. */
$H(y) \leftarrow$ Hankel matrix built from y;
$\tilde{U}, \tilde{\Sigma}, \tilde{W} \leftarrow \text{SVD}(H(y))$; /* Make sure $\tilde{\Sigma}$ is decreasingly ordered. */
$\tilde{U}_S \leftarrow$ First S columns of $\tilde{U}$; /* Remember $\tilde{U}$ is a $(L+1)\times(L+1)$ unitary matrix */
$\tilde{U}_0 \leftarrow$ First L rows of $\tilde{U}_S$;
$\tilde{U}_1 \leftarrow$ Last L rows of $\tilde{U}_S$;
$\tilde{\Psi} \leftarrow \tilde{U}_0^{+}\tilde{U}_1$; /* Make $S\times S$ signal matrix $\tilde{\Psi}$, $+$ denotes Moore-Penrose inverse. */
$\tilde{z}_1, \ldots, \tilde{z}_S \leftarrow$ eigenvalues of signal matrix $\tilde{\Psi}$.

We will see that this algorithm comes with recovery guarantees on the parameters $z_1, \ldots, z_S$, in both the real-time and imaginary-time signal case, provided the noise vector η is small enough. The strength of these guarantees differs significantly between the two types of signal, and we will discuss them separately in the next sections. From the $\tilde{z}_j$'s we can then (for both the real-time and imaginary-time signal) extract $\tilde{E}_j$'s, which denote the S estimates for $\{E_i \in [0, 2\pi)\}_{i=1}^{S}$ returned by the classical post-processing algorithm. The error in the energy estimates is set as the optimal matching distance [35]

$$d(\{E_i\}, \{\tilde{E}_j\}) = \frac{1}{2\pi} \min_{\pi \in \text{Perm}_S} \max_j |\tilde{E}_{\pi(j)} - E_j|, \tag{6.32}$$

i.e. the returned list is optimally matched with the actual eigenvalues and the error is set by the largest mismatch.

6.3.1. Real-time (oscillatory) signal

In this section we discuss the performance of ESPRIT on real-time (oscillatory) signals. This performance has been well studied in the signal processing literature. Here, we will follow the analysis of [32], which provides Theorem 6.3.1 relating $\|H(\eta)\|$ in equation (6.31) and the energy matching error defined in equation (6.32).

The performance of ESPRIT in the oscillatory signal case relies on lower bounding the smallest nonzero singular value of the Vandermonde matrix $V_{L=K/2}$ in equation (6.28), (or similarly upperbounding the condition number $\kappa(V_{K/2}) = \sigma_{\max}(V_{K/2})/\sigma_{\min}(V_{K/2})$). The smallest nonzero singular value of the Vandermonde matrix $V_{K/2}$ will depend on K, S and the location of the poles z_j. For the real-time signal, the z_j lie on the unit circle whereas

for the imaginary-time signal the z_j lie in the interval $(e^{-2\pi}, 1]$. Let the minimal gap between the E_i be defined as

$$\Delta = \frac{1}{2\pi} \min_{j \neq k} |E_j - E_k|. \tag{6.33}$$

It has been proved [34] for $z_j = e^{-iE_j}$ that

$$\Delta \geq \frac{C}{K} \Rightarrow \sigma^2_{\min}(V_{K/2}) \geq \frac{C-1}{C} K, \tag{6.34}$$

for some constant $C > 1$. Note that if there are S eigenvalues $E_j \in [0, 2\pi)$ in the signal, it is clear that the minimal gap $\Delta \leq 1/S$, hence one should at least take $K \geq CS$. Based on this bound, Theorem 4 in [32] says:

Theorem 6.3.1 ([32])**.** *Let $(g+\eta)(k)$ be a real-time signal with $k = 0, \ldots, K$, and with $g(k) = \sum_{i=1}^{S} c_i z_i^k$, $c_i > 0 \, \forall \, i$, $c_{\min} = \min_i c_i$ and $\eta(k)$ a small noise vector. Let $z_j = e^{-iE_j}$ with $j = 1, \ldots, S$ and $E_j \in [0, 2\pi) \ \forall \, j$, and $K \geq 2C/\Delta$ for some constant $C > 2$ with gap Δ, and $K + 1 \geq 2S$. If*

$$\left\| H(\eta) \right\| \leq c_{\min} K \, h_1(S, C, K), \tag{6.35}$$

with

$$h_1(S, C, K) = \frac{C-1}{8\sqrt{2S}C} \sqrt{1 - \frac{2CS}{(C-1)K}}, \tag{6.36}$$

then the ESPRIT algorithm outputs energy estimates $\{\tilde{E}_j\}$ with distance

$$d(\{E_i\}), \{\tilde{E}_j\}) \leq \left\| H(\eta) \right\| c_{\min}^{-1} K^{-1} \, h_2(S, C, K), \tag{6.37}$$

with

$$h_2(S, C, K) = 40\sqrt{2} S^2 \left(\frac{C}{C-1} \right)^{3/2} \left(1 - \frac{2CS}{(C-1)K} \right)^{-1}. \tag{6.38}$$

By equation (6.31) we have $\left\| H(\eta) \right\| \leq K \epsilon_{\text{tot}}$ and if we choose $K \sim S$, ϵ_{tot} can be chosen sufficiently small, inversely polynomial with S, such that at least equation (6.35) holds. Then $d(\{E_i\}), \{\tilde{E}_j\})$ will be $\Theta(\left\| H(\eta) \right\| S)$, hence decreasing like $S^2 \epsilon_{\text{tot}}$.

If we combine this Theorem with the quantum results of Theorem 6.1.1, then we obtain Theorem 6.1.3. These results thus form the theoretical underpinning of the ideas and numerical work in [3] in which quantum phase estimation was replaced by the repeated execution of a circuit applying controlled-U^k (conditioned on an ancilla qubit state) which gets Trotterized to the overlap test circuit in Fig. 6.1.

Remark 6.3.2. *It is noteworthy that even when the eigenvalues E_j are not well-separated but occur in 'clumps', results exist [32] which bound the performance of ESPRIT.*

6.3.2. Imaginary-time (decaying) signal

Let us now discuss what information can be extracted from the imaginary-time signal in the presence of sampling and Trotter noise and compare this to the known Theorem 6.3.1 for the real-time signal.

In 6.D we discuss in detail the recovery guarantees for ESPRIT for imaginary-time signals. This analysis is an adaptation of the work done in [32] for real-time signals, with the only true novelty being Lemma 6.D.7. However, since no rigorous analysis for imaginary-time signals exists in the literature we go through all the steps in considerable detail. The analysis will again depend on the condition number of the Vandermonde matrix $V_{L=K/2}$ in equation (6.28).

This condition number is much worse behaved, i.e. much larger, in case the z_i's all lie on the real axis –which is the case for the imaginary-time signal– but bounds on this condition number do exist [36]. Based on the work of Gautschi [37], we derive our own upper bounds on this condition number, which are asymptotically sub-optimal but have a clearer dependence on the choice of K and the given S than previous bounds in [36]. We then use the gap Δ to fill in the upper bound.

In analogy to Theorem 6.3.1, we then obtain the following:

Theorem 6.3.3. *Let $(g+\eta)(k)$ be an imaginary-time decaying signal with $k = 0,\dots,K$, and with $g(k) = \sum_{i=1}^{S} c_i z_i^k$, $c_i > 0$, $\forall i$, $c_{\min} = \min_i c_i$, and $\eta(k)$ a small noise vector. Let $z_i = e^{-E_i}$ with $E_i \in [0, 2\pi)$ and given eigenvalue gap $\Delta < 1$ in equation (6.33), and $\{\tilde{E}_i\}$ the energy estimates of ESPRIT with $L = K/2$. Let $K+1 \geq 2S$, K even and $K = TS$ for some positive integer T. If we have*

$$\left\| H(\eta) \right\| \leq \frac{c_{\min}}{\sqrt{K}} g_1(S,\Delta), \tag{6.39}$$

with

$$g_1(S,\Delta) = \frac{1}{32S^2} (e^{-2\pi}\pi\Delta)^{3(S-1)}, \tag{6.40}$$

then

$$d(\{\tilde{E}_i\},\{E_j\}) \leq \left\| H(\eta) \right\| c_{\min}^{-1} K\sqrt{K} g_2(S,\Delta), \tag{6.41}$$

with

$$g_2(S,\Delta) = e^{2\pi} 640\sqrt{2}\, S^{5.5} (e^{-2\pi}\pi\Delta)^{-5(S-1)}. \tag{6.42}$$

Since the dependence on S is exponential in equation (6.41), one cannot make the distance $d(\{\tilde{E}_i\},\{E_j\})$ small when the number of eigenvalues $S = \text{poly}(n)$, no matter what the gap. This is a crucial difference with the oscillatory real-time case. However, for $S = O(1)$, with sufficient, poly(n), effort one can make $\left\| H(\eta) \right\|$ sufficiently small to obey equation (6.39) and then reduce the error on the found eigenvalues to $1/\text{poly}(n)$. This assumes that the gap between the $O(1)$ rescaled eigenvalues present in the initial state is at least $1/\text{poly}(n)$ (and not exponentially small in n).

Furthermore, given that $\|H(\eta)\|$ should decrease at least as $\sim 1/\sqrt{K}$ through equation (6.39) but the upper bound in equation (6.41) scales as $\|H(\eta)\|K^{3/2}$, one obtains the optimal bound by choosing the *minimal* K, namely $K = 2S$, so that $L = K/2 = S$. In this case the Vandermonde matrix $V_{L-1} = V_{S-1}$ is square [2]. This expresses the intuitive fact that increasing K will not help beyond a point, as for larger K the signal simply dies out. This is unlike the oscillatory case of Theorem 6.3.1 in which the optimal K is required to grow with $1/\Delta$. Here the bound does not require that K grows with $1/\Delta$, so there is no 'super-resolution'. We note that the upper bounds may have a sub-optimal dependence on K and S, which is due to the proof techniques. Practically (roughly) speaking, whenever the condition number of the Vandermonde matrix $V_{L=K/2}$ grows by choosing a larger K, choosing that larger K can be beneficial.

For the other decaying signal $(g_D(k))$, a rather small change from $z_i = \exp(-E_i)$ to $z_i = 1 - E_i/2\pi$ gives:

Theorem 6.3.4. *Let $(g+\eta)(k)$ be a decaying signal with $k = 0,\dots,K$, and with $g(k) = \sum_{i=1}^{S} c_i z_i^k$, $c_i > 0$, $\forall i$, $c_{\min} = \min_i c_i$, and $\eta(k)$ a small noise vector. Let $z_i = 1 - E_i/2\pi$ with $E_i \in [0,\pi]$ and given eigenvalue gap $\Delta < 1$ in equation* (6.33)*, and $\{\tilde{E}_i\}$ the energy estimates of ESPRIT with $L = K/2$. Let $K+1 \geq 2S$, K even and $K = TS$ for some positive integer T. If we have*

$$\|H(\eta)\| \leq \frac{c_{\min}}{\sqrt{K}} \tilde{g}_1(S,\Delta), \tag{6.43}$$

with

$$\tilde{g}_1(S,\Delta) = \frac{1}{32S^2}\,\Delta^{3(S-1)}, \tag{6.44}$$

then

$$d(\{\tilde{E}_i\},\{E_j\}) \leq \|H(\eta)\|\, c_{\min}^{-1} K\sqrt{K}\,\tilde{g}_2(S,\Delta), \tag{6.45}$$

with

$$\tilde{g}_2(S,\Delta) = 640\sqrt{2}\,S^{5.5}\,\Delta^{-5(S-1)}. \tag{6.46}$$

Now to argue Theorem 6.1.5 from Theorem 6.3.4, we simply choose the minimal $K = 2S$, and since $S = O(1)$, it implies that the classical algorithm which estimates $g_D(k)$ for $k = 0,\dots,K(= O(1))$ within error ϵ using Lemma 6.2.6 requires poly(n) effort.

6.4. Spectral estimation for a transverse-field Ising chain

In this section, we numerically investigate the methods described thus far by applying them to an archetypal stoquastic Hamiltonian: The transverse field Ising chain. This system has been extensively studied [25] and will serve as a proof-of-principle test. The system consists of qubits on a one-dimensional lattice, which interact via an Ising

[2] Hence, strictly speaking Lemma 6.D.11 is not much of a help.

interaction and are exposed to an external magnetic field in the transverse direction. The Hamiltonian associated with this system is:

$$H = -J\Big(\sum_i Z_i Z_{i+1} + g\sum_i X_i\Big), \tag{6.47}$$

where X, Y, Z denote the Pauli matrices, $J > 0$ (for a ferromagnetic interaction) and $g \geq 0$, so that H is term-wise stoquastic in the standard basis. We take the field to be pointing in the x-direction without loss of generality [3].

The system exhibits an abrupt change in the ground state of the system as a function of g at $g = 1$ (for $n \to \infty$). On either side of the phase transition, one has:

- **Strong-coupling limit** ($g \gg 1$): In this limit, the Hamiltonian is dominated by the magnetic field terms and the ground state is given by $|\psi_0\rangle \approx |+\rangle^{\otimes n}$. The p-particle excitations correspond to states $|-\rangle_{q_1}|-\rangle_{q_2}...|-\rangle_{q_p}\prod_{i\neq q_1,q_2,...,q_p}|+\rangle_i$, i.e., the ground state with spin flips at p sites $q_1,...,q_p$ along the chain. These p-particle excited states are $\binom{n}{p}$-fold degenerate.

- **Weak-coupling limit** ($g \ll 1$): In this limit, the Hamiltonian is dominated by the Ising interaction terms and the (degenerate) ground state is given by either $|\psi_0\rangle \approx |0\rangle^{\otimes n}$ or $|\psi_0\rangle \approx |1\rangle^{\otimes n}$ (ferromagnetic phase). The excitations w.r.t. the ground state correspond to domain walls separating ferromagnetic regions of opposite spin.

To run the Monte Carlo scheme described in Lemma 6.2.3, the imaginary-time propagation operator e^{-kH} must be decomposed (by means of Trotterization) in terms of the local propagation operators $e^{-a_l k/M H_i}$ (where a_l and M are set by the Trotterization scheme) [4]. The local propagation operators acting on a subset of two qubits on the chain are given by:

$$e^{-\tilde{k}H_i} = \begin{pmatrix} \frac{\sinh(\lambda\tilde{k})}{\sqrt{1+g^2}} + \cosh(\lambda\tilde{k}) & 0 & \frac{g\sinh(\lambda\tilde{k})}{\sqrt{1+g^2}} & 0 \\ 0 & \frac{-\sinh(\lambda\tilde{k})}{\sqrt{1+g^2}} + \cosh(\lambda\tilde{k}) & 0 & \frac{g\sinh(\lambda\tilde{k})}{\sqrt{1+g^2}} \\ \frac{g\sinh(\lambda\tilde{k})}{\sqrt{1+g^2}} & 0 & \frac{-\sinh(\lambda\tilde{k})}{\sqrt{1+g^2}} + \cosh(\lambda\tilde{k}) & 0 \\ 0 & \frac{g\sinh(\lambda\tilde{k})}{\sqrt{1+g^2}} & 0 & \frac{\sinh(\lambda\tilde{k})}{\sqrt{1+g^2}} + \cosh(\lambda\tilde{k}) \end{pmatrix}, \tag{6.48}$$

where $\lambda = J\sqrt{1+g^2}$ and $\tilde{k} = a_l k/M$. This operator is element-wise non-negative and can be efficiently brought to bock-diagonal form (with each block being irreducible).

[3] The Hamiltonian can be transformed to $\tilde{H} = UHU^\dagger$ by the unitary transformation $U = \bigotimes_i \exp(\frac{i\theta Z_i}{2})$, which alters the direction of the field in the transverse plane while preserving the spectrum.

[4] We note that the numerical results presented in this section are obtained using a first-order Trotter decomposition.

Since the choice of $|\Phi\rangle$ directly governs which eigenvalues can be obtained from the real-time and imaginary-time evolution signals, it is a point of particular importance. In addition, the ability of ESPRIT to extract eigenvalues from the imaginary-time and real-time signals depends very strongly on the spectral gap between the eigenvalues in the signal. We consider a state $|\Phi\rangle$ which has considerable overlap with the ground state and the (n-fold degenerate) first excited state in the ($g > 1$)-regime. Since the gap between their associated eigenvalues increases monotonically as a function of g in this regime, this allows us to present the aforementioned gap dependence numerically. We shall call the state $\left|\Phi_{\text{optimal}}\right\rangle$ since in the ($g \gg 1$)-regime it optimally overlaps with the eigenstates of interest, i.e. $|\langle +^{\otimes n}|\psi_{p=0}\rangle|^2 = \sum_{q=1}^{n} |\langle +^{\otimes n}|\psi_{p=1,q}\rangle|^2 = \frac{1}{2}$. This state is given by:

$$\begin{aligned}\left|\Phi_{\text{optimal}}\right\rangle &= \frac{1}{\sqrt{2}}\Bigg(\underbrace{\prod_{i=1}^{n}|+\rangle_i}_{|\psi_{p=0}\rangle} + \sum_{q=1}^{n}\frac{1}{\sqrt{n}}\underbrace{|-\rangle_q\prod_{i\neq q}|+\rangle_i}_{|\psi_{p=1,q}\rangle}\Bigg) \\ &= \frac{1}{2^{(n+1)/2}}\sum_{q=1}^{n}\left(\left(\left(\frac{1}{n}+\frac{1}{\sqrt{n}}\right)|0\rangle_q + \left(\frac{1}{n}-\frac{1}{\sqrt{n}}\right)|1\rangle_q\right)\sum_{x\in\{0,1\}^{n-1}}|x\rangle\right),\end{aligned} \tag{6.49}$$

where $\sum_{x\in\{0,1\}^{n-1}}|x\rangle$ denotes an equal superposition of $(n-1)$-bit strings that exclude the bit in register q.

We note that for $\left|\Phi_{\text{optimal}}\right\rangle$, one can efficiently obtain $\frac{\Phi(y)}{\Phi(x)}$ for a given $x, y \in \{0,1\}^n$ and one can efficiently sample from $|\Phi(x)|^2$: From Eq. equation (**??**), one can infer a function $\Phi(x)$ ($\{0,1\}^n \to \mathbb{R}$) that (efficiently) gives the coefficient of the state $\left|\Phi_{\text{optimal}}\right\rangle$ associated with an n-bit string x: $\Phi(x) = 1/2^{(n+1)/2}\left(\left(\frac{1}{n}+\frac{1}{\sqrt{n}}\right)(n-|x|) + \left(\frac{1}{n}-\frac{1}{\sqrt{n}}\right)|x|\right)$, so $\Phi(x)$ only depends on the Hamming weight $|x|$ of bit string x, i.e. the quantity $\frac{\Phi(y)}{\Phi(x)}$ can be efficiently determined. Furthermore, since $\Phi(x)$ only depends on n and $|x|$, the distribution $|\Phi(x)|^2$ also depends solely on these quantities. This implies that one can indeed efficiently sample from this distribution: First, one draws a Hamming weight $|x|$ from the distribution $|\Phi(x)|^2 = |\Phi(|x|)|^2$. Then, given $|x|$, one constructs at random an n-bit string with this Hamming weight. This latter step can be efficiently implemented by starting from some n-bit string with Hamming weight $|x|$ (such as $\{1\}^{|x|}\{0\}^{n-|x|}$) and then applying a random permutation.

6.4.1. Numerical method and results

We briefly discuss the details of the numerical analysis that is used to obtain the results presented in this section. We use the Monte Carlo and quantum algorithms (where the latter is inefficiently implemented on a classical computer), which are presented in Section 6.2 and summarized in Theorems 6.1.2 and 6.1.1, to obtain resp. the imaginary-time and real-time evolution signals for the transverse-field Ising chain. We note that here we estimate the imaginary-time evolution signal using the empirical mean estimator, instead of the (asymptotically superior) median-of-means estimator. Having obtained these signals, we obtain estimates of the eigenvalues using the filtered ESPRIT method: This method corresponds to Algorithm 6.2 in combination with an

additional filtering step. This additional step is required since in principle the number of components in the signal S is not known a priori in the current setting. Therefore, we construct the matrix $\tilde{U}_S$ (in Algorithm 6.2) by taking the first S columns of $\tilde{U}$, where S is now the number of singular values in the SVD of the Hankel matrix $H(y)$ that exceed $\mathsf{TF}\,\sigma_{\max}$. TF denotes what we call a truncation factor, and $\sigma_{\max}$ denotes the largest singular value of $H(y)$. In this way, the number of components in the signal emerges from the analysis of its Hankel matrix, rather than being a quantity that is known beforehand. By implementing the remainder of Algorithm 6.2 as usual, we obtain estimates of the z_j's. From these estimates of the z_j's, we obtain the *spectral estimates* $\tilde{E}_j$ for the quantum algorithm and for the Monte Carlo algorithm.

Note that this approach of including a filtering step – which often resembles more closely the practically encountered scenario when running the algorithms from Lemmas 6.2.1 and 6.2.3 – differs from that considered in Theorems 6.3.1 and 6.3.3, where the number of components S in signals is known beforehand. Here, S is a quantity emerging in the analysis and it can even generally occur that components of the signal with very small coefficients – corresponding to eigenstates with very small overlap with $|\Phi\rangle$ – are filtered out.

In the results presented in this section, note that the real-time and imaginary-time increments have been chosen such that all E_j that are present in the signals lie in $[0, 2\pi)$. This does *not* mean that the whole spectrum of the Hamiltonian lies in $[0, 2\pi)$, as the majority of its eigenvalues will not be present in the signals.

We note that for the quantum algorithm, the parameters $\{z_j\}$ have unit norm. However, due to finite sampling, one determines a noisy version of the signal $g_R(k)$, resulting in estimated eigenvalues of the Trotterized unitary having norms that slightly deviate from unity. To ensure that the estimates $\tilde{E}_j$ are real-valued, we take them to be the real parts of $i\log(\tilde{z}_j)$.

The code that is used to obtain the numerical results presented in this work can be found at [38].

In Figure 6.3, the Monte Carlo signals $\langle\Phi|\, e^{-kH}\,|\Phi\rangle$ and the real and imaginary parts of the quantum algorithm signals $\langle\Phi|\, e^{-ikH}\,|\Phi\rangle$ for $|\Phi\rangle = |+\rangle^{\otimes n}$ and $\left|\Phi_{\text{optimal}}\right\rangle$ are depicted.

The upper three figures correspond to $|\Phi\rangle = |+\rangle^{\otimes n}$. For this choice of $|\Phi\rangle$, the signals are clearly dominated by a single eigenvalue (the ground state eigenvalue): The Monte Carlo signal decays with a single decay rate and the quantum algorithm signals oscillate with a single frequency. For the quantum algorithm signals, there are also higher-frequency components visible (due to $|+\rangle^{\otimes n}$ not having overlap with *only* the ground state).

For the lower three figures, we take $|\Phi\rangle = \left|\Phi_{\text{optimal}}\right\rangle$. For this choice of $|\Phi\rangle$, there are two eigenvalues present in the signals (the ground state and first excited state eigenvalues). For the Monte Carlo signal, the excited state eigenvalue can be seen to die out within a few units of time, after which only the ground state component is left. The quantum algorithm signals can be seen to be composed of a high-frequency (excited-state) component superposed on the ground-state component, where the excited-state component now obviously does *not* die out.

We now consider the spectral estimates that are obtained by applying ESPRIT to the evolution signals that are produced by the quantum algorithm (from Theorem 6.1.1) and

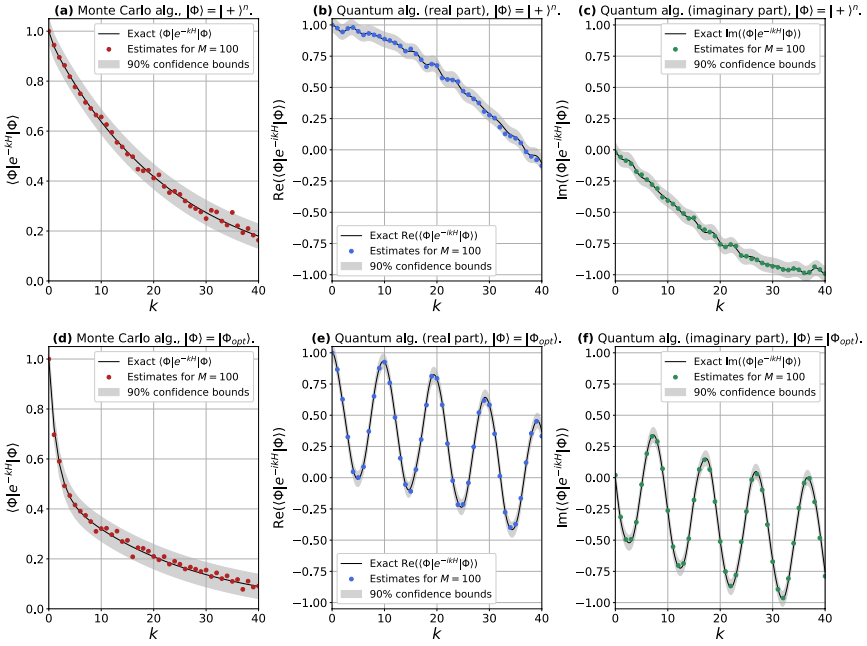

Figure 6.3: The evolution of the states $|+\rangle^{\otimes n}$ (in **(a)**, **(b)** and **(c)**, for which the ground state is the dominant component in the signal) and $\left|\Phi_{\text{optimal}}\right\rangle$ (in **(d)**, **(e)** and **(f)**, for which the ground state and first excited state are the dominant components in the signal) for $n = 7$ and $g = 4$ in imaginary time (in **(a)** and **(d)**) and in real time (in **(b)**, **(c)**, **(e)** and **(f)**). The signals in **(a)** and **(d)** are obtained through the Monte Carlo scheme of Theorem 6.1.2. The signals in **(b)**, **(c)**, **(e)** and **(f)** are obtained through the quantum algorithm of Theorem 6.1.1 (which is inefficiently implemented on a classical computer). The Trotter variable is taken to be $M = 100$ and $|\Sigma|$ is set to be 4200.

Monte Carlo algorithm (from Theorem 6.1.2). In particular, we determine both time evolution signals at a given total number of measurement points in real/imaginary time. We then determine the spectral estimates from both signals for increasing K, by including step-by-step more of the total number of measurement points in the analysis [5]. The truncation factor TF is taken to be equal to 0.02 throughout.

The top two plots in Figure 6.4 depict, for a given $|\Sigma|$, the eigenvalue estimates as a function of g and for several values of K. For both the quantum algorithm and Monte Carlo algorithm estimates, it is clear that a smaller spectral gap indeed requires a larger K for the eigenvalues to be obtained accurately. Furthermore, for a given $|\Sigma|$ and K, it is clear that the error of the estimate for the excited-state eigenvalue obtained from the imaginary-time signal is larger than that obtained from the real-time signal. We conclude furthermore that, in line with Theorems 6.3.1 and 6.3.3, increasing K beyond a certain

[5] For $K = 2$; $k = 0, 1, 2$. For $K = 4$; $k = 0, 1, 2, 3, 4$. Etc.

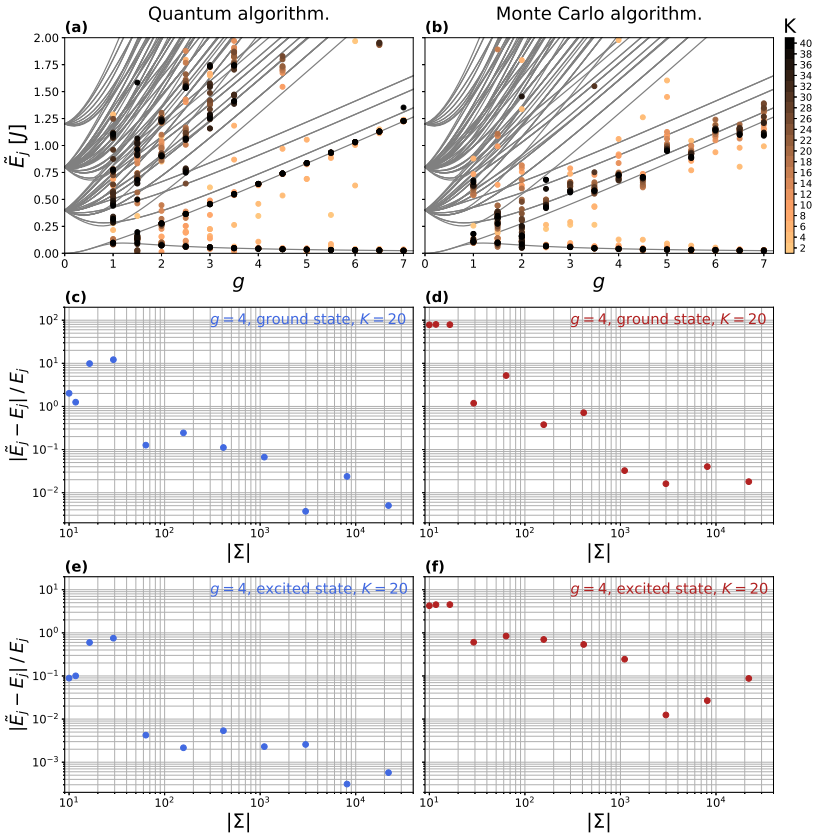

Figure 6.4: Spectral estimates of the ferromagnetic Ising chain in a transverse field (for $n = 7$) obtained through analysis of the evolution of $|\Phi_{\text{optimal}}\rangle$. Plots **(a)** and **(b)** depict the spectral estimates (together with the true spectrum) obtained through the quantum algorithm and the Monte Carlo algorithm for $|\Sigma| = 4200$ and $M = 100$ for several values of K. Plots **(c)**,**(d)** and plots **(e)**,**(f)** depict the relative error of the spectral estimates – i.e. $|\tilde{E}_j - E_j|/E_j$ – for the resp. ground state and excited state eigenvalues at $g = 4$, for $M = 400$ and as a function of $|\Sigma|$. The truncation factor is taken to be $\mathsf{TF} = 0.02$ throughout. The scaling of the error of the ground-state eigenvalue estimates is similar for both methods, while the error for excited-state eigenvalue is larger for the MC algorithm than for the quantum algorithm. The excited-state eigenvalue estimates also converge more quickly as a function of K for the quantum algorithm.

threshold does not necessarily reduce the error of the eigenvalue estimates.

It is apparent that as one approaches the $g = 1$ point, more higher-lying eigenvalues emerge from the ESPRIT analysis. This is especially true for the quantum algorithm (note that for the Monte Carlo signal, the larger the eigenvalues are, the quicker the associated

components in the signal die out). The appearance of these higher-lying eigenvalues can be attributed to the fact that (for finite n) the state $\left|\Phi_{\text{optimal}}\right\rangle$ starts to have significant overlap with states other than the two lowest-energy eigenstates in this regime.

The middle two and bottom two plots in Figure 6.4 depict the relative error of the spectral estimates – i.e. $|\tilde{E}_j - E_j|/E_j$ – for resp. the ground-state eigenvalue and excited-state eigenvalue (at fixed $g = 4$). We consider a range of values for $|\Sigma|$. For the ground-state eigenvalue, the scaling of the relative errors as a function of $|\Sigma|$ is similar for the quantum algorithm and the Monte Carlo algorithm. Clearly, the relative errors of the excited-state eigenvalue estimates for the quantum algorithm are smaller than those for the Monte Carlo algorithm.

We have also implemented the matrix pencil method in [20, 33] to estimate the eigenvalues from the real-time and imaginary-time signals. The only significant difference that was found between the estimates obtained through the ESPRIT method and through this matrix pencil method is that – in the $(K < 2S)$-regime – the matrix pencil method outputs estimates which resemble an average of the eigenvalues in the signal (as can be seen in Figure 6.2 in a noiseless setting), while this is not the case generally for the ESPRIT method.

6.5. Discussion

We have considered the problem of obtaining (some) eigenvalues of local stoquastic – i.e. sign-problem-free – Hamiltonians and general local Hamiltonians H by means of tracking the evolution of the system state, differentiating between the evolution of the system state in real time and imaginary time. In both cases, we examine the use of the matrix pencil ESPRIT method in extracting eigenvalues of H from the state evolution signal. The real-time (oscillating) evolution signal is obtained through running quantum circuits, while the imaginary-time (decaying) signal for local stoquastic Hamiltonians is obtained through a Monte Carlo scheme (developed in this work) that is implemented in a computationally tractable manner classically. Another type of decaying evolution signal – from which the ESPRIT method can extract eigenvalues of H – is obtained through a classical method for general local Hamiltonians that is similar in spirit to 'dequantization'.

We have invoked some known performance bounds of the ESPRIT method for the real-time signal and applied and extended bounds for the imaginary-time signal. Our bounds suggest that the ESPRIT method (or matrix pencil methods more generally) performs – not surprisingly – worse in extracting (multiple) eigenvalues from an imaginary-time decaying (MC algorithm) signal than from a real-time oscillating (quantum algorithm) signal in the presence of noise. However, we show that if the input state contains $S = O(1)$ eigenstates and the spectral gap is at least $1/\text{poly}(n)$, and the right access to the input state is available, the associated eigenvalues can be resolved efficiently (with $\text{poly}(n)$ classical effort) for local stoquastic as well as for general local Hamiltonians. Even though for $S = O(1)$, the classical effort for stoquastic as well as general Hamiltonians is $\text{poly}(n)$, the 'brute-force' algorithm for general Hamiltonians (in Lemma 6.2.6) incurs an exponential cost in k in estimating the signal $g_D(k)$, while for stoquastic Hamiltonians the cost is polynomial in k. Despite this difference in cost, the error bounds for the eigenvalue estimates obtained here through analysis of the ESPRIT

method applied to a decaying signal ($g_D(k)$ or $g_I(k)$) suggests that letting k grow as some function of n will generally not help.

Even though our results show that for these Hamiltonians, for an input state supported on $S = O(1)$ eigenvalues (separated by an at least $1/\text{poly}(n)$ gap), these eigenvalues can be estimated with $\text{poly}(n)$ classical effort, it remains to be better understood how *practical* this MC method for stoquastic Hamiltonians or the 'dequantization' method in Lemma 6.2.6 are. The upper bounds for the errors on the eigenvalue estimates in Theorem 6.3.3 grow rather fast with S (and the computational effort grows fast with k in Lemma 6.2.6 for general local Hamiltonians), and it is not clear how much one can improve, say, the ESPRIT bounds.

Indeed, it would be interesting to show that the current bounds of ESPRIT for the imaginary-time decaying signal cannot be improved upon. There are definitely known negative results on the condition number of Vandermonde matrices [39], but there might be signal extraction algorithms that have better practical performance on decaying signals, or have looser requirements (such as the requirement that all data is evenly spaced). However, we suspect that the difficulty gap we observe between real-time and imaginary-time signal is universal. One possible way to argue this is through the Cramer-Rao bound (which has been analysed for real-time signals [40] but not for imaginary-time signals), which is a question we leave for further research.

In terms of numerical results, we find that: For a given spectral gap and sample size, the ability to distinguish between two eigenvalues indeed depends on the number of measurement points K at which the real-time and imaginary-time evolution signals are evaluated. The MC algorithm for stoquastic Hamiltonians and the quantum algorithm (in combination with the ESPRIT method) lead to a similar scaling of the relative error of the ground-state eigenvalue as a function of the sample size. However, for an excited-state eigenvalue, the quantum algorithm leads to significantly smaller relative errors than the MC algorithm. More extensive numerical studies, also of models other than the transverse-field Ising chain, may shed further light on whether the Monte Carlo + ESPRIT method is useful in practice. For frustrated stoquastic Hamiltonians, even the smallest eigenvalue may lead to a fast decaying signal, requiring small sampling error and Trotter error in practice.

As for other directions of further research, one can ask whether a hybrid approach in which imaginary-time data from an error-free Monte Carlo algorithm can strengthen the use of real-time data from a quantum algorithm obtained from a noisy quantum circuit. This approach requires combining the data where the poles/nodes $z_j = e^{-iE_j}$ on the unit circle each have a partner pole $z'_j = e^{-E_j}$ (or $z'_j = I - E_j/2\pi$) on the real axis. If the effect of noise can be modeled $z_j = e^{-iE_j} \to e^{iE_j - \gamma}$ [3], then the imaginary-time data may help in extracting the values for E_j. It may also be of interest to consider the case of sampling k for both the quantum circuit and Monte Carlo method at random (instead of picking $k = 0, 1, \ldots, K$). Another direction of further research is the following. Suppose the input state has overlap with S (here not necessarily $O(1)$) eigenstates of the Hamiltonian, one could asses how well the ESPRIT methods succeeds in extracting e.g. the ground-state eigenvalue by filtering out all other components in the real-time or imaginary-time evolution signals. Another rather different direction of further research is the estimation of low-lying eigenvalues of Laplacian matrices of graphs (which relate to

the connectivity of the graph), where we note that Laplacian matrices are stoquastic.

References

[1] M. Nielsen and I. Chuang. *Quantum computation and quantum information: 10th Anniversary Edition*. Cambridge University Press, 2010. DOI: 10.1017/CBO9780511976667.

[2] K. Svore, M. Hastings and M. Freedman. 'Faster phase estimation'. In: *Quantum Inf. Comput.* 14.3-4 (Mar. 2014), pp. 306–328. DOI: 10.26421/QIC14.3-4-7.

[3] T. O'Brien, B. Tarasinski and B. Terhal. 'Quantum phase estimation of multiple eigenvalues for small-scale (noisy) experiments'. In: *New Journal of Physics* 21.023022 (Feb. 2019). DOI: 10.1088/1367-2630/aafb8e/meta.

[4] R. Somma. 'Quantum eigenvalue estimation via time series analysis'. In: *New Journal of Physics* 21.12 (Dec. 2019), p. 123025. DOI: 10.1088/1367-2630/ab5c60.

[5] S. Lloyd. 'Universal quantum simulators'. In: *Science* 273.5278 (Aug. 1996), pp. 1073–1078. URL: http://www.jstor.org/stable/2899535.

[6] J. Gubernatis, N. Kawashima and P. Werner. *Quantum Monte Carlo Methods. Algorithms for lattice models*. Cambridge University Press, 2016. DOI: 10.1017/CBO9780511902581.

[7] W. Foulkes, L. Mitas, R. Needs and G. Rajagopal. 'Quantum Monte Carlo simulations of solids'. In: *Reviews of Modern Physics* 73.1 (Jan. 2001). DOI: 10.1103/RevModPhys.73.33.

[8] S. Bravyi, D. DiVincenzo, R. Oliveira and B. Terhal. 'The complexity of stoquastic local Hamiltonian problems'. In: *Quantum Inf. Comput.* 8.5 (May 2008), pp. 361–385. DOI: 10.5555/2011772.2011773.

[9] M. Ioannou, S. Piddock, M. Marvian, J. Klassen and B. Terhal. "Sign-curing local Hamiltonians: termwise versus global stoquasticity and the use of Clifford transformations". July 2020. arXiv: 2007.11964 [quant-ph].

[10] D. Aharonov and A. Bredariol Grilo. 'Stoquastic PCP vs. randomness'. In: *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*. 2019, pp. 1000–1023. DOI: 10.1109/FOCS.2019.00065.

[11] S. Bravyi and D. Gosset. 'Polynomial-Time classical simulation of quantum ferromagnets'. In: *Physical Review Letters* 119.10 (Sept. 2017). DOI: 10.1103/PhysRevLett.119.100503.

[12] S. Bravyi. 'Monte Carlo simulation of stoquastic Hamiltonians'. In: *Quantum Inf. Comput.* 15.13-14 (Oct. 2015), pp. 1122–1140. DOI: 10.5555/2871363.2871366.

[13] E. Crosson and A. Harrow. 'Rapid mixing of path integral Monte Carlo for 1D stoquastic Hamiltonians'. In: *Quantum* 5 (Feb. 2021), p. 395. ISSN: 2521-327X. DOI: 10.22331/q-2021-02-11-395.

[14] E. Crosson and S. Slezak. "Classical simulation of high temperature quantum Ising models". Feb. 2020. arXiv: 2002.02232v1 [quant-ph].

[15] T. Sarkar and O. Pereira. 'Using the matrix pencil method to estimate the parameters of a sum of complex exponentials'. In: *IEEE Antennas and Propagation Magazine* 37.1 (Feb. 1995), pp. 48–55. DOI: 10.1109/74.370583.

[16] Y. Hua and T. Sarkar. 'On SVD for estimating generalized eigenvalues of singular matrix pencil in noise'. In: *IEEE Transactions on Signal Processing* 39.4 (Apr. 1991), pp. 892–900. DOI: 10.1109/78.80911.

[17] D. Potts and M. Tasche. 'Parameter estimation for nonincreasing exponential sums by Prony-like methods'. In: *Linear Algebra and its Applications* 439.4 (Aug. 2013), pp. 1024–1039. DOI: 10.1016/j.laa.2012.10.036.

[18] E. Onorati, A. Werner and J. Eisert. 'Randomized benchmarking for individual quantum gates'. In: *Physical Review Letters* 123.6 (Aug. 2019). ISSN: 1079-7114. DOI: 10.1103/physrevlett.123.060501.

[19] J. Helsen, I. Roth, E. Onorati, A. Werner and J. Eisert. "A general framework for randomized benchmarking". Oct. 2020. arXiv: 2010.07974 [quant-ph].

[20] J. Helsen, F. Battistel and B. Terhal. 'Spectral quantum tomography'. In: *npj Quantum Information* 5.74 (Sept. 2019). DOI: 10.1038/s41534-019-0189-0.

[21] D. Hangleiter, I. Roth, J. Eisert and P. Roushan. "Precise Hamiltonian identification of a superconducting quantum processor". Aug. 2021. arXiv: 2108.08319 [quant-ph].

[22] S. Gharibian and F. Le Gall. "Dequantizing the Quantum Singular Value Transformation: Hardness and Applications to Quantum Chemistry and the Quantum PCP Conjecture". Nov. 2021. arXiv: 2111.09079 [quant-ph].

[23] C. Cade, M. Folkertsma and J. Weggemans. *Complexity of the Guided Local Hamiltonian Problem: Improved Parameters and Extension to Excited States*. July 2022. arXiv: 2207.10097 [quant-ph].

[24] S. Gharibian, R. Hayakawa, F. L. Gall and T. Morimae. *Improved Hardness Results for the Guided Local Hamiltonian Problem*. July 2022. arXiv: 2207.10250 [quant-ph].

[25] S. Sachdev. *Quantum phase transitions*. 2nd ed. Cambridge University Press, 2011. DOI: 10.1017/CBO9780511973765.

[26] M. Motta, S. Chong, A. Tan, M. O'Rourke, E. Ye, A. Minnich, F. Brandão and G. Chan. 'Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution'. In: *Nature Physics* 16.2 (Nov. 2019), pp. 205–210. ISSN: 1745-2481. DOI: 10.1038/s41567-019-0704-4.

[27] L. Lin and Y. Tong. 'Heisenberg-limited ground-state energy estimation for early fault-tolerant quantum computers'. In: *PRX Quantum* 3 (1 Feb. 2022), p. 010318. DOI: 10.1103/PRXQuantum.3.010318.

[28] K. Wan, M. Berta and E. Campbell. "A randomized quantum algorithm for statistical phase estimation". Oct. 2021. arXiv: 2110.12071 [quant-ph].

[29] R. Horn and C. Johnson. *Matrix analysis*. 2nd ed. Cambridge University Press, 2012. DOI: 10.1017/9781139020411.

[30] G. Lugosi and S. Mendelson. 'Mean estimation and regression under heavy-tailed distributions: a survey'. In: *Foundations of Computational Mathematics* 19 (Aug. 2019), pp. 1145–1190. DOI: 10.1007/s10208-019-09427-x.

[31] A. Childs, Y. Su, M. Tran, N. Wiebe and S. Zhu. 'Theory of Trotter error with commutator scaling'. In: *Phys. Rev. X* 11 (1 Feb. 2021), p. 011020. DOI: 10.1103/PhysRevX.11.011020.

[32] W. Li, W. Liao and A. Fannjiang. 'Super-resolution limit of the ESPRIT algorithm'. In: *IEEE Transactions on Information Theory* 66.7 (Feb. 2020), pp. 4593–4608. DOI: 10.1109/TIT.2020.2974174.

[33] Y. Hua and T. Sarkar. 'Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise'. In: *IEEE Transactions on Acoustics, Speech and Signal Processing* 38.5 (May 1990), pp. 814–824. DOI: 10.1109/29.56027.

[34] A. Moitra. 'Super-resolution, extremal functions and the condition number of Vandermonde matrices'. In: *Proceedings of the forty-seventh annual ACM symposium on Theory of Computing* (June 2015), pp. 821–830. DOI: 10.1145/2746539.2746561.

[35] R. Bhatia. *Graduate Texts in Mathematics: Matrix analysis*. Vol. 169. Springer Science & Business Media, 1997. DOI: 10.1007/978-1-4612-0653-8.

[36] F. Bazán. 'Conditioning of rectangular Vandermonde matrices with nodes in the unit disk'. In: *SIAM J. Matrix Anal. Appl.* 21.2 (Oct. 1999), pp. 679–693. ISSN: 0895-4798. DOI: 10.1137/S0895479898336021.

[37] W. Gautschi. 'On the inverses of Vandermonde and confluent Vandermonde matrices I.' In: *Numer. Math* 4 (1962), pp. 117–123. URL: https://www.cs.purdue.edu/homes/wxg/selected_works/section_01/016.pdf.

[38] M. Stroeks. "ClassQuantSimStoqHam". https://github.com/MStroeks/ClassQuantSimStoqHam. 2022.

[39] V. Pan. 'How Bad Are Vandermonde Matrices?' In: *SIAM Journal on Matrix Analysis and Applications* 37.2 (2016), pp. 676–694. DOI: 10.1137/15M1030170.

[40] P. Stoica and A. Nehorai. 'MUSIC, maximum likelihood, and Cramer-Rao bound'. In: *IEEE Transactions on Acoustics, speech, and signal processing* 37.5 (May 1989), pp. 720–741. DOI: 10.1109/29.17564.

[41] M. Suzuki. 'General theory of fractal path integrals with applications to many-body theories and statistical physics'. In: *Journal of Mathematical Physics* 32.2 (Feb. 1991), pp. 400–407. DOI: 10.1063/1.529425.

[42] T. Greville. 'Note on the generalized inverse of a matrix product'. In: *SIAM Review* 8.4 (Oct. 1966), pp. 518–521. DOI: 10.1137/1008107.

[43] G. W. Stewart. 'Perturbation theory for the singular value decomposition'. In: *SVD and signal processing, II: algorithms, analysis and applications* (1991), pp. 99–109. URL: https://users.math.msu.edu/users/iwenmark/Teaching/MTH995/Papers/SVD_Stewart.pdf.

[44] R. Li. 'Relative perturbation theory: II. Eigenspace and singular subspace variations'. In: *SIAM Journal on Matrix Analysis and Applications* 20.2 (1998), pp. 471–492. URL: https://doi.org/10.1137/S0895479896298506.

[45] P. Wedin. 'Perturbation theory for pseudo-inverses'. In: *BIT Numerical Mathematics* 13.2 (June 1973), pp. 217–232. URL: https://doi.org/10.1007/BF01933494.

6.A. Trotterization

Suppose $H = \sum_{i=1}^{N} H_i$ (where $N = \mathcal{O}(\text{poly}(n))$) represents a k-local Hamiltonian of a quantum system. $\{H_i\}_{i=1}^{N}$ is generally a set of non-commuting terms but can be divided into subsets, such that within each subset all terms commute. For a given set $\{H_i\}_{i=1}^{N}$, we denote the minimum possible number of these subsets by Γ. This number of subsets is at most N and equals 1 in the trivial case where all H_i's commute with each other. The Hamiltonian H can thus be decomposed as $H = \sum_{\gamma=1}^{\Gamma} H_\gamma$, where all H_γ do not commute with each other, but the terms of which each individual H_γ is composed do commute. Choosing a decomposition into the minimum number of subsets brings about an additional advantage of parallelizability when implementing the evolution of the systems in imaginary or real time.

The following Lemma (adaptation from [31]) upper bounds the errors of implementing imaginary-time and real-time state evolution through a first-order Trotter decomposition.

Lemma 6.A.1. ***First-Order Trotter Decomposition.*** *Given a k-local Hamiltonian $H = \sum_i^N H_i$. Furthermore, suppose the set $\{H_i\}_{i=1}^{N}$ can be divided into a minimum of Γ subsets $\{H_\gamma\}_{\gamma=1}^{\Gamma}$, such that within each individual subset all H_i's commute. Then the quantities $\left|\langle\Phi| e^{-itH} |\Phi\rangle - \langle\Phi| \left(\prod_\gamma e^{-itH_\gamma/M}\right)^M |\Phi\rangle\right|$ and $\left|\langle\Phi| e^{-\tau H} |\Phi\rangle - \langle\Phi| \left(\prod_\gamma e^{-\tau H_\gamma/M}\right)^M |\Phi\rangle\right|$ (where $|\Phi\rangle$ is a normalized state and $t, \tau \in \mathbb{R}_+$) are bounded as follows:*

$$\left|\langle\Phi| e^{-itH} |\Phi\rangle - \langle\Phi| \Big(\prod_{\gamma} e^{-itH_\gamma/M}\Big)^M |\Phi\rangle\right| \leq \sum_{\gamma'=1}^{\Gamma-1} \sum_{\gamma>\gamma'} ||\,[H_{\gamma'}, H_\gamma]\,|| \,\frac{t^2}{2M}, \tag{6.50a}$$

$$\left|\langle\Phi| e^{-\tau H} |\Phi\rangle - \langle\Phi| \Big(\prod_{\gamma} e^{-\tau H_\gamma/M}\Big)^M |\Phi\rangle\right| \leq 3e^2 \sum_{\gamma'=1}^{\Gamma-1} \sum_{\gamma>\gamma'} ||\,[H_{\gamma'}, H_\gamma]\,|| \,\frac{\tau^2}{2M}, \tag{6.50b}$$

where the second inequality holds provided that $\left|\left|e^{-\tau H/M}\right|\right| \leq 1$, $\left|\left|e^{-\tau H_\gamma/M}\right|\right| \leq 1$ $(\forall \gamma)$ and $\frac{\tau\left(\sum_\gamma ||H_\gamma||\right)}{M} \leq 1$, and M denotes the Trotter variable.

To obtain a better scaling of the errors as a function of the Trotter variable M, one can employ higher-order Trotter decompositions. We denote the pth-order approximants of $e^{-itH/M}$ and $e^{-\tau H/M}$ by $\mathcal{T}_M(p,t)$ and $\mathcal{T}_M(p,\tau)$, respectively. We denote $\left|\langle\Phi| e^{-itH} |\Phi\rangle - \langle\Phi| \mathcal{T}_M(p,t)^M |\Phi\rangle\right|$ and $\left|\langle\Phi| e^{-\tau H} |\Phi\rangle - \langle\Phi| \mathcal{T}_M(p,\tau)^M |\Phi\rangle\right|$ by ϵ_{trot}. In [31], it was shown that, for general p, ϵ_{trot} is upper bounded as follows:

$$\epsilon_{\text{trot}} \leq \mathcal{O}\Big(\alpha\, t^{p+1}/M^p\Big), \quad \text{for real-time evolution,} \tag{6.51a}$$

$$\epsilon_{\text{trot}} \leq \mathcal{O}\Big(\alpha\, \tau^{p+1}/M^p\Big), \quad \text{for imaginary-time evolution,} \tag{6.51b}$$

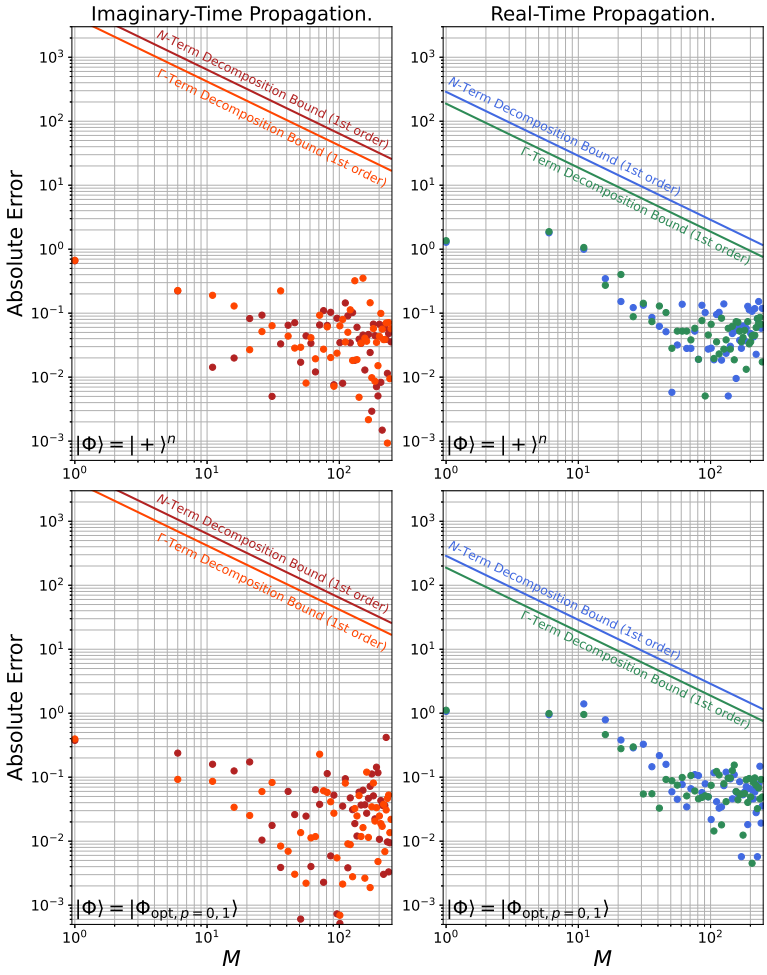

Figure 6.5: Absolute Trotter error (imposed on the signal estimate) as a function of the Trotter variable M for the imaginary-time and real-time signals. The noisy ($|\Sigma| = 200$) and Trotterized versions of $\langle\Phi| e^{-\tau H} |\Phi\rangle$ and $\langle\Phi| e^{-itH} |\Phi\rangle$ for a ferromagnetic Ising chain in a transverse field (for $g = 4$ and $n = 8$) are evaluated at $\tau = t = 3$ and several values of M. The Trotterization schemes are first-order N-term and first-order $\Gamma(= 2)$-term schemes. The associated error bounds are included in matching colors.

where $\alpha = \sum_{\gamma_1,\gamma_2,...,\gamma_{p+1}=1}^{\Gamma} || [H_{\gamma_{p+1}}, ..., [H_{\gamma_2}, H_{\gamma_1}]...] ||$ ($\alpha^{1/p}$ is typically poly(n)) and equation (6.51b) holds provided that $4\tau\Upsilon\left(\sum_\gamma \left\| H_\gamma \right\|\right)/M \leq 1$ (where Υ corresponds to the number of *stages* of the Trotter decomposition and typically scales exponentially in p) [6].

[6] In the remainder of this discussion it is assumed that this condition is satisfied.

In [41], a widely used scheme is discussed for constructing pth-order approximants.

It is important to consider the total number of k-local propagation operators L required to simulate e^{-itH} and $e^{-\tau H}$ (for a given order p and Trotter variable M). For the scheme in [41], the number of these k-local propagation operators required to be implemented for the simulation of e^{-itH} and $e^{-\tau H}$ for $p>1$ is $L=2MN5^{\frac{p}{2}-1}$ (and for $p=1$ is MN). If one wishes to obtain a given ϵ_{trot}, the number of k-local propagation operators into which the evolutions are decomposed scales as $L=\text{poly}(n)\,\mathcal{O}\left(\Upsilon\, t^{1+1/p}\epsilon_{\text{trot}}^{-1/p}\right)$ (for real time) and $L=\text{poly}(n)\,\mathcal{O}\left(\Upsilon\,\tau^{1+1/p}\epsilon_{\text{trot}}^{-1/p}\right)$ (for imaginary time). We thus conclude that for large p (i.e. high-order decompositions), L scales approximately linearly in the evolution time of the system under consideration (for real-time and imaginary-time evolution).

In Figure 6.5, we have depicted the absolute error of noisy MC (imaginary-time) and QPE (real-time) signals at fixed $\tau=t$ as a function of M, obtained through first-order N-term and $\Gamma(=2)$-term Trotterization schemes. We have included the first-order Trotter error bounds. We note that the apparent drastic increase in noise magnitude as a function of M is primarily due to the fact that the absolute error decreases as a function of M and is plotted on a logarithmic scale.

6.B. Extension to non-Hermitian propagation operators

In this Appendix we prove the following Lemma, extending Lemma 6.2.3:

Lemma 6.B.1. *Let $\mathcal{F}\equiv\langle\Phi|\,G_1G_2\ldots G_L\,|\Phi\rangle$, where:*

1. *$|\Phi\rangle=\sum_{x=1}^{2^n}\Phi(x)\,|x\rangle$ is a normalized state of n qubits where $\Phi(x)\in\mathbb{C}$ $(\forall\, x)$ and $\sum_x\left|\Phi(x)\right|^2=1$. We assume that (1) $\frac{\Phi(y)}{\Phi(x)}$ can be efficiently (*poly(n)*) calculated for a* given *x and y and (2) we can efficiently draw samples from the probability distribution $P(x)=\left|\Phi(x)\right|^2$.*
2. *Each G_l is a k-local (possibly non-Hermitian) element-wise nonnegative matrix with singular values in $(0,1]$.*

$\mathcal{F}$ can be estimated within error ϵ with probability at least $1-\delta$ with a classical MC algorithm with runtime poly$(n)\times\Theta(\epsilon^{-2}\delta^{-1})\times\Theta(L)$.

Proof. In addition to the n-qubit register, we exploit a single ancillary qubit. The matrices G_l are still element-wise non-negative. The state $|a\rangle$ denotes the state of the single ancillary qubit. By making use of the single ancillary qubit, the propagation operators can be symmetrized as follows:

$$F_l\equiv\begin{cases}G_l\otimes|0\rangle\langle1|+G_l^\dagger\otimes|1\rangle\langle0|\,, & \text{if } l \text{ is odd}\\ G_l\otimes|1\rangle\langle0|+G_l^\dagger\otimes|0\rangle\langle1|\,, & \text{if } l \text{ is even.}\end{cases}\tag{6.52}$$

In this form, F_l (the 'new' propagation operator) is element-wise non-negative, $k+1$-local and Hermitian and hence one can apply Lemma 6.2.3 to $\langle\Phi|\,F_1F_2\ldots F_L\,|\Phi\rangle$, provided that its eigenvalues lie in $(0,1]$. The eigenvalues λ of F_l (for l odd) can be found by solving:

$$\det\begin{pmatrix}-\lambda\mathbb{1} & G_l\\ G_l^\dagger & -\lambda\mathbb{1}\end{pmatrix}=\det\left(\lambda^2\mathbb{1}-G_lG_l^\dagger\right)=\det\left(G_lG_l^\dagger-\lambda^2\mathbb{1}\right)=0,\tag{6.53}$$

where we have used that $G_l^{(\dagger)}$ commutes with $\mathbb{1}$ and that $G_l^{(\dagger)}$ is of even dimensionality. The eigenvalues of the Hermitian and positive semi-definite matrix $G_l G_l^\dagger$ are thus equal to λ^2. Since the singular values of G_l are equal to the square root of the eigenvalues of $G_l G_l^\dagger$, the eigenvalues of F_l will lie in $(0,1]$ if the singular values of G_l lie in $(0,1]$. This can be similarly shown for l even and this statement thus holds for all l.

What is left to prove is that estimating the signal for the string of F_l's is equivalent to estimating the signal for the string of G_l's. Specifically, we want to prove the following identity: $G_1 G_2 ... G_L = \langle 0| F_1 F_2 ... F_L |L \bmod 2\rangle$, for $L \in \mathbb{Z}_+$. This is done below by means of induction.

- For $L = 1$:

$$\begin{aligned} \langle 0| F_1 |1\rangle &= \langle 0| \Big(G_1 \otimes |0\rangle\langle 1| + G_1^\dagger \otimes |1\rangle\langle 0| \Big) |1\rangle \\ &= G_1 \langle 0|0\rangle\langle 1|1\rangle + G_1^\dagger \langle 0|1\rangle\langle 0|1\rangle \\ &= G_1, \end{aligned} \tag{6.54}$$

- Assuming $G_1 G_2 ... G_L = \langle 0| F_1 F_2 ... F_L |L \bmod 2\rangle$ holds for L, it holds for $L+1$ as well: Making use of the definition in equation (6.52), we write F_{L+1} as follows:

$$F_{L+1} = G_{L+1} \otimes |L \bmod 2\rangle\langle L+1 \bmod 2| + G_{L+1}^\dagger \otimes |L+1 \bmod 2\rangle\langle L \bmod 2|. \tag{6.55}$$

  The quantity of interest – in the case of the length of the operator string being $L+1$ – can now be rewritten as follows:

$$\langle 0| F_1 F_2 ... F_L F_{L+1} |L+1 \bmod 2\rangle = \langle 0| F_1 F_2 ... F_L |L \bmod 2\rangle\, G_{L+1} = G_1 G_2 ... G_L G_{L+1}, \tag{6.56}$$

  which finishes the proof.

□

6.C. Median-of-means estimator

The MC scheme described in Section 6.2 produces a set of $|\Sigma|$ samples $\{\boldsymbol{x}\}$ which are distributed according to $\Pi(\boldsymbol{x})$. For each sample, $\mathrm{Re}(\mathcal{R}(\boldsymbol{x}))$ can be evaluated and subsequently an estimate of $\mathcal{F}$ can be obtained. Only the first and second moments of the random variable $\mathrm{Re}(\mathcal{R}(\boldsymbol{x}))$ can be upper bounded in general. Therefore, if one would use the empirical mean $\mathrm{Re}(\tilde{\mathcal{F}}) = \frac{1}{|\Sigma|}\sum_{\boldsymbol{x}\in\Sigma} \mathrm{Re}(\mathcal{R}(\boldsymbol{x}))$ as a mean estimator for $\mathcal{F}$, then the best achievable scaling of $|\Sigma|$ such that

$$\Pr\Big(\big| \mathrm{Re}(\tilde{\mathcal{F}}) - \mathcal{F} \big| \leq \epsilon \Big) \geq 1 - \delta, \tag{6.57}$$

is $|\Sigma| = \Theta(\epsilon^{-2}\delta^{-1})$ (by means of Chebyshev's inequality).

Taking the median-of-means estimator [30] as estimator (instead of the empirical mean), one can obtain a more convenient scaling of $|\Sigma|$ w.r.t. δ (despite the fact that only the first two moments of $\mathrm{Re}(\mathcal{R}(\boldsymbol{x}))$ can be upper bounded). The median-of-means estimator can be constructed as follows: Partition the set of MC samples Σ into q groups

$s_1, \ldots, s_q$ of size approximately $|\Sigma|/q$. One then computes the empirical mean of $\mathrm{Re}(\mathcal{R}(\boldsymbol{x}))$ over the samples in each group separately (giving q unbiased estimators of $\mathcal{F}$) and takes the median of these empirical means. We denote the empirical mean for each group by $f_j = \frac{1}{|s_j|}\sum_{\boldsymbol{x}\in s_j} \mathrm{Re}(\mathcal{R}(\boldsymbol{x}))$ (for $j \in \{1,\ldots,q\}$) and denote the median of these empirical means by $\hat{\mathcal{F}} = \mathrm{M}(f_1,\ldots,f_q)$. The estimator $\hat{\mathcal{F}}$ is the median-of-means estimator.

We define the median of q real numbers $a_1,\ldots,a_q$ as $\mathrm{M}(a_1,\ldots,a_q) = a_i$ with a_i such that

$$|\{j : a_j \le a_i\}| \ge q/2 \quad \wedge \quad |\{j : a_j \ge a_i\}| \ge q/2, \tag{6.58}$$

where we take the smallest i if multiple is obey this condition.

$\{\mathrm{Re}(\mathcal{R}(\boldsymbol{x}))\}$ are i.i.d. random variables with mean $\mathcal{F}$ and variance $\mathrm{Var}\big(\mathrm{Re}(\mathcal{R}(\boldsymbol{x}))\big) \le 1$. Let q and $|\Sigma|/q$ be positive integers, then

$$\Pr\Big(\big|\hat{\mathcal{F}} - \mathcal{F}\big| \le \sqrt{4q/|\Sigma|}\Big) \ge 1 - e^{-q/8}. \tag{6.59}$$

So for $q = 8\log(\delta^{-1})$ and $|\Sigma| = 4\,q\epsilon^{-2} = 32\log(\delta^{-1})\epsilon^{-2}$, we have:

$$\Pr\Big(\big|\hat{\mathcal{F}} - \mathcal{F}\big| \le \epsilon\Big) \ge 1 - \delta. \tag{6.60}$$

Note that the estimator $\hat{\mathcal{F}} = \mathrm{M}(f_1,\ldots,f_q)$ depends explicitly on the confidence since q scales with δ. Given that indeed $q = \Theta\big(\log(\delta^{-1})\big)$, the number of samples required to obtain equation (6.60) is $|\Sigma| = \Theta(\log(\delta^{-1})\epsilon^{-2})$ (which is an exponentially better scaling w.r.t. δ compared to that for the empirical mean estimator).

To see why equation (6.59) is true , see [30], note that one can apply Chebyshev's inequality to each of the empirical means f_j: with probability at least $3/4$, we have $\big|f_j - \mathcal{F}\big| \le \sqrt{4q/|\Sigma|}$. If $\big|\hat{\mathcal{F}} - \mathcal{F}\big| \ge \sqrt{4q/|\Sigma|}$, then, by definition of $\hat{\mathcal{F}}$, at least $q/2$ of the empirical means f_j satisfy $\big|f_j - \mathcal{F}\big| \ge \sqrt{4q/|\Sigma|}$. Hence the probability that $\big|\hat{\mathcal{F}} - \mathcal{F}\big| \ge \sqrt{4q/|\Sigma|}$ is upper bounded by the probability that a binomially distributed random variable with q draws and success probability $1/4$ exceeds $q/2$:

$$\begin{aligned}\Pr\Big(\big|\hat{\mathcal{F}} - \mathcal{F}\big| \ge \sqrt{4q/|\Sigma|}\Big) &\le \Pr\Big(\mathrm{Bin}(q,1/4) \ge q/2\Big) \\ &= \Pr\Big(\mathrm{Bin}(q,1/4) - \mathbb{E}\big(\mathrm{Bin}(q,1/4)\big) \ge q/4\Big) \le e^{-q/8},\end{aligned} \tag{6.61}$$

where we have used $\mathbb{E}\big(\mathrm{Bin}(q,1/4)\big) = q/4$ and Hoeffding's inequality.

6.D. Performance of ESPRIT on the imaginary-time (decaying) signal

In this section we prove a series of Lemmas that characterize the behaviour of the ESPRIT algorithm (Algorithm 6.2) on an imaginary-time signal obtained with finite error. They are direct generalisations of the work done in [32], which leads up to Theorem 6.3.1 for oscillatory signals, to signals composed of real exponential decays. We will see that the guarantees on the algorithm will be substantially weaker in this case. The end goal of this section is Theorem 6.3.3 in the main text.

The argument decomposes roughly into two halves. In the first half we argue that the behaviour of ESPRIT is controlled by the smallest non-zero singular value of the Vandermonde matrix V_L. In the second half we argue that that this smallest nonzero singular value can be controlled in terms of a gap condition on the energy eigenvalues of the imaginary-time signal.

We start by proving a short result on the smallest nonzero singular values of products of matrices.

Lemma 6.D.1. *Let the smallest* nonzero *singular value of a matrix X be $\sigma_{\min}(X)$. For any matrix, X we have $\sigma_{\min}(X) := \left\|X^+\right\|^{-1}$, where X^+ is the Moore-Penrose pseudo-inverse of X, i.e. through the SVD, we have $\sigma_{\min}^{-1}(X) = \left\|X^+\right\|$, where $\|X\|$ is the operator norm (the largest singular value). Let A, B be (non-square) matrices such that $(AB)^+ = B^+A^+$. Then we have that*

$$\sigma_{\min}(AB) \geq \sigma_{\min}(A)\sigma_{\min}(B). \tag{6.62}$$

Proof. By sub-multiplicativity of the operator norm, we have that

$$\sigma_{\min}(AB) = \left(\left\|(AB)^+\right\|\right)^{-1} = \left(\left\|B^+A^+\right\|\right)^{-1} \geq \left(\left\|B^+\right\|\left\|A^+\right\|\right)^{-1} = \sigma_{\min}(B)\sigma_{\min}(A). \tag{6.63}$$

□

We note that the product property on the Moore-Penrose pseudo-inverse does not hold for all matrices (unlike for the regular inverse). We will make use of the following sufficient condition:

Lemma 6.D.2 ([42])**.** *Let A, B be matrices and let A have full column rank, and B have full row rank. Then we have $(AB)^+ = B^+A^+$.*

Next, we argue that a small perturbation in the imaginary-time signal does not impact the space spanned by the the first S left singular vectors of the Hankel matrix $H(g)$ too strongly, see the ESPRIT Algorithm 6.2. It is a compressed version of Lemmas 4 and 5 in [32] (which are formulated for real-time signals only, but hold more generally). To state this Lemma we need to consider a freedom of choice in U_S and $\tilde{U}_S$ with $\tilde{U}_S$ as defined in the ESPRIT Algorithm 6.2 and U_S its noise-free version. It is possible that U_S and $\tilde{U}_S$ are far apart as operators, even if the spaces they span are close together.

We solve this by not considering U_S proper, but rather a rotated version of U_S. As we will see, this rotation will not impact the actual output of ESPRIT which are the eigenvalues of the signal matrix $\tilde{\Psi}$. The rotated version of U_S is given through the $S \times S$ unitary operator $(O_2O_1)^\dagger$, which is defined via the singular value decomposition of $U_S^\dagger \tilde{U}_S$, i.e.

$$O_1 U_S^\dagger \tilde{U}_S O_2 = D, \tag{6.64}$$

with $I_S \geq D \geq 0$ and D diagonal. The diagonal elements of the matrix D are cosines of the so-called *canonical angles*. We note that this internal rotation is performed implicitly in [32], whereas we have chosen to make it explicit at all times.

Lemma 6.D.3. *Let $(g+\eta)(k)$ be an imaginary-time signal with $g(k) = \sum_{i=1}^{S} c_i z_i^k$ and $\eta(k)$ a small noise vector. Consider the associated Hankel matrices $H(g)$ and $H(g+\eta)$, with singular value decompositions $H(g) = U\Sigma W$ and $H(g+\eta) = \tilde{U}\tilde{\Sigma}\tilde{W}$, and label the matrix of the first S columns of U (resp. $\tilde{U}$) as U_S (resp. $\tilde{U}_S$). Finally, let $O_1 U_S^\dagger \tilde{U}_S O_2 = D$ with $I_S \geq D \geq 0$ be the singular value decomposition of $U_S^\dagger \tilde{U}_S$. If*

$$\left\| H(\eta) \right\| \leq \sigma_{\min}(H(g))/2, \tag{6.65}$$

then

$$\left\| U_S (O_2 O_1)^\dagger - \tilde{U}_S \right\| \leq \frac{2\sqrt{2S}\left\| H(\eta) \right\|}{\sigma_{\min}(H(g))}. \tag{6.66}$$

Proof. First, we can observe that indeed $I_S \geq D$ as $\left\| O_1 U_S^\dagger \tilde{U}_S O_2 \right\| \leq \|O_1\| \left\| U_S^\dagger \right\| \left\| \tilde{U}_S \right\| \|O_2\| \leq 1$.

The proof follows from Wedin's $\sin\Theta$ theorem for perturbations of singular subspaces as well as Weyl's perturbation theorem for singular values, see e.g. [35, 43]. From this latter theorem we know that $|\sigma_i(H(g+\eta)) - \sigma_i(H(g))| \leq \left\| H(\eta) \right\| \leq \sigma_{\min}(H(g))/2$ where σ_i is the ith singular value (in order and some singular values can be zero). Let $\sigma_{\min}(H(g+\eta)) > 0$ be the kth singular value, and thus

$$\sigma_{\min}(H(g+\eta)) \geq \sigma_k(H(g)) - \sigma_{\min}(H(g))/2 \geq \sigma_{\min}(H(g))/2, \tag{6.67}$$

where the last inequality holds as $\sigma_k(H(g)) > 0$ and hence is at least $\sigma_{\min}(H(g))$. Hence we can use Wedin's theorem on singular values (Theorem 3.4 in [44], setting $\delta = \alpha = \sigma_{\min}(H(g))/2$) to conclude that

$$\left\| (U_S^\perp)^\dagger \tilde{U}_S \right\| \leq \frac{2\left\| H(\eta) \right\|}{\sigma_{\min}(H(g))}, \tag{6.68}$$

where $U_S^\perp$ is the matrix formed from the $L+1-S$ other (besides U_S) columns of the noiseless U. To connect this to $U_S(O_2O_1)^\dagger - \tilde{U}_S$ we can make the following long calculation:

$$\begin{aligned}
\left\| U_S (O_2 O_1)^\dagger - \tilde{U}_S \right\| &\leq \left\| U_S (O_2 O_1)^\dagger - \tilde{U}_S \right\|_F \\
&= \left[\mathrm{tr}\left(U_S (O_2 O_1)^\dagger (O_2 O_1) U_S^\dagger \right) + \mathrm{tr}\left(\tilde{U}_S \tilde{U}_S^\dagger \right) - \mathrm{tr}\left((U_S (O_2 O_1)^\dagger \tilde{U}_S^\dagger + \tilde{U}_S (O_2 O_1) U_S^\dagger) \right) \right]^{1/2} \\
&= \left[\mathrm{tr}\left(U_S U_S^\dagger \right) + \mathrm{tr}\left(\tilde{U}_S \tilde{U}_S^\dagger \right) - 2\,\mathrm{tr}(D) \right]^{1/2} \\
&\leq \left[2\,\mathrm{tr}\left(\tilde{U}_S \tilde{U}_S^\dagger \right) - 2\,\mathrm{tr}\left(D D^\dagger \right) \right]^{1/2} \\
&= \sqrt{2} \left[\mathrm{tr}\left(\tilde{U}_S \tilde{U}_S^\dagger \right) - \mathrm{tr}\left((O_1 U_S^\dagger \tilde{U}_S O_2)(O_1 U_S^\dagger \tilde{U}_S O_2)^\dagger \right) \right]^{1/2} \\
&= \sqrt{2} \left[\mathrm{tr}\left(\tilde{U}_S \tilde{U}_S^\dagger \right) - \mathrm{tr}\left(U_S^\dagger \tilde{U}_S \tilde{U}_S^\dagger U_S \right) \right]^{1/2} \\
&= \sqrt{2} \left[\mathrm{tr}\left(\tilde{U}_S \tilde{U}_S^\dagger \right) - \mathrm{tr}\left(U_S U_S^\dagger \tilde{U}_S \tilde{U}_S^\dagger \right) \right]^{1/2} \\
&= \sqrt{2} \left[\mathrm{tr}\left(\tilde{U}_S \tilde{U}_S^\dagger \right) - \mathrm{tr}\left(\tilde{U}_S \tilde{U}_S^\dagger \right) + \mathrm{tr}\left(U_S^\perp (U_S^\perp)^\dagger \tilde{U}_S \tilde{U}_S^\dagger \right) \right]^{1/2} \\
&= \sqrt{2} \left\| (U_S^\perp)^\dagger \tilde{U}_S \right\|_F \leq \sqrt{2S} \left\| (U_S^\perp)^\dagger \tilde{U}_S \right\|
\end{aligned} \tag{6.69}$$

In the second inequality we used that $\mathrm{tr}\left(U_S U_S^\dagger\right) = \mathrm{tr}\left(\tilde{U}_S \tilde{U}_S^\dagger\right) = S$ since U_S as well as $\tilde{U}_S$ consist of S orthonormal columns, and $D \geq D^2 = DD^\dagger$, since $I \geq D \geq 0$. In addition, at the end we use that $UU^\dagger = U_S U_S^\dagger + U_S^\perp (U_S^\perp)^\dagger = I$ as U is unitary. □

The next step is to bound the deviation of the ESPRIT signal matrix $\tilde{\Psi} = \tilde{U}_0^+ \tilde{U}_1$ from the rotated version of its noiseless variant $(O_2O_1)\Psi(O_2O_1)^\dagger = (O_2O_1)U_0^+ U_1 (O_2O_1)^\dagger$ in terms of $\left\| U_S(O_2O_1)^\dagger - \tilde{U}_S \right\|$. Recall that U_0 (resp. U_1) are constructed by removing respectively the first or last row from the matrix U_S. Note also that only the eigenvalues of the signal matrix Ψ matter in the ESPRIT Algorithm 6.2 and the additional unitary rotations O_2O_1 do not alter these eigenvalues. We first establish some intermediate result:

Lemma 6.D.4. *Let A, B be matrices such that* $\mathrm{rank}(A) = \mathrm{rank}(B)$. *If* $\|A-B\| \leq \sigma_{\min}(A)/2$ *then*

$$\left\|A^+ - B^+\right\| \leq \frac{1+\sqrt{5}}{2} \|A-B\| \left\|A^+\right\|^2 = \frac{1+\sqrt{5}}{2} \frac{\|A-B\|}{\sigma_{\min}^2(A)}. \tag{6.70}$$

Proof. From Theorem 4.1 in [45] we get that

$$\left\|A^+ - B^+\right\| \leq \frac{1+\sqrt{5}}{2} \|A-B\| \left\|A^+\right\| \left\|B^+\right\|. \tag{6.71}$$

Furthermore, since $\|A-B\| \leq \frac{\sigma_{\min}(A)}{2} < \frac{1}{\|A^+\|}$, we have by Lemma 3.1 in [45] that

$$\left\|B^+\right\| \leq \frac{\left\|A^+\right\|}{1 - \|A^+\| \|A-B\|} \leq \left\|A^+\right\|, \tag{6.72}$$

leading to the first inequality in equation (6.70) and the Lemma follows. □

The next Lemma establishes that if a (non-square) matrix is full rank, a sufficiently small perturbation does not *decrease* the rank (and hence rank is preserved). Note that full-rankness is really required, as an arbitrarily small perturbation can always *increase* the rank.

Lemma 6.D.5. *Let A be an $m \times n$ ($m \leq n$) matrix of rank m and let B an $m \times n$ matrix s.t.* $\|A-B\| \leq \sigma_{\min}(A)/2$. *Then* $\mathrm{rank}(A) = \mathrm{rank}(B)$.

Proof. By construction, we have $\mathrm{rank}(A) \geq \mathrm{rank}(B)$. Moreover we have that the smallest singular value of B is at least $\sigma_{\min}(A) - \|A-B\|$, by Weyl's singular value perturbation theorem [43]. Thus by construction the smallest singular value of B is at least $\sigma_{\min}(A)/2$ which is strictly larger than 0 as A is full rank and thus B is also full rank, and thus $\mathrm{rank}(A) = \mathrm{rank}(B)$. □

Finally, we will require a bound on the smallest nonzero singular value of U_0. This is the only Lemma where we deviate significantly from the work done in [32], where the corresponding result, Lemma 3 in [32], makes explicit use of the fact that in their scenario all poles z_j lie on the unit circle (what we call the real-time, oscillatory signal). The bound we present here is simpler and more general and thus applies to both imaginary (decaying) as well as real-time (oscillatory) signals, but leads to a suboptimal dependence

on the condition number of the Vandermonde matrix V_L defined in equation (6.28) (in particular $\sigma_{\min}(V_L)$). However, it is sufficient for our purpose. The Lemma will use some essential properties of how the ESPRIT Algorithm 6.2 works which we review first. Key to the functioning of ESPRIT is the fact that $H(g)$ has two decompositions

$$H(g) = U\Sigma W = V_L C V_{K-L}^T, \tag{6.73}$$

where V_L is the $(L+1)\times S$ Vandermonde matrix defined in equation (6.28) and the coefficient matrix C is given in Eq. equation (**??**). When $S \le L \le K+1-S$ (requiring $K+1 \ge 2S$), V_L and V_{K-L} are full rank. Then V_L and U_S have an image of the same dimension, which means there exists an invertible matrix A such that

$$U_S = V_L A, \tag{6.74}$$

and thus

$$U_0 = V_{L-1}A, U_1 = V_{L-1}ZA. \tag{6.75}$$

with $Z = \mathrm{diag}(z_1, \ldots, z_S)$. This implies that

$$\Psi = U_0^+ U_1 = A^{-1} V_{L-1}^+ V_{L-1} Z A = A^{-1} Z A. \tag{6.76}$$

and hence the eigenvalues of Ψ are the poles z_i.

Lemma 6.D.6. *Let U_0 be the $L\times S$ matrix obtained from U_S by removing the last row. If the associated Vandermonde matrix V_{L-1} is of (full) rank S then so is U_0, and moreover*

$$\sigma_{\min}(U_0) \ge \frac{\sigma_{\min}(V_{L-1})}{\|V_L\|}. \tag{6.77}$$

Proof. We have

$$I_S = U_S^\dagger U_S = (V_L A)^\dagger V_L A = A^\dagger V_L^\dagger V_L A, \tag{6.78}$$

which means that the singular values of A are precisely inverse to those of V_L, or equivalently that A^+ has the same singular spectrum as V_L. Moreover, by assumption V_{L-1} has full column rank, and A is invertible so

$$\sigma_{\min}^{-1}(U_0) = \left\|U_0^+\right\| = \left\|(V_{L-1}A)^+\right\| = \left\|A^{-1}V_{L-1}^+\right\| \le \left\|A^{-1}\right\|\left\|V_{L-1}^+\right\| = \frac{\|V_L\|}{\sigma_{\min}(V_{L-1})}, \tag{6.79}$$

which is the inverse of the Lemma statement. □

With Lemmas 6.D.4, 6.D.5 and 6.D.6 in hand, we can give a perturbation bound for the signal matrix $\tilde{\Psi}$. We will show that if $\tilde{U}_0$ does not deviate strongly from the rotated version of U_0, namely $U_0(O_2O_1)^\dagger$, then the noisy signal matrix is also close to the ideal (rotated) version.

Lemma 6.D.7. *Let $\Psi := U_0^+ U_1, \tilde{\Psi} = \tilde{U}_0^+ \tilde{U}_1$ be the ideal and perturbed version of the signal matrix, respectively. Now assume that $\left\|U_0(O_2O_1)^\dagger - \tilde{U}_0\right\| \le \sigma_{\min}(U_0)/2$, where O_2O_1 is defined through the singular value decomposition of $U_S^\dagger \tilde{U}_S$ (i.e. $O_1 U_S^\dagger \tilde{U}_S O_2 = D$). With this assumption we have*

$$\left\|(O_2O_1)\Psi(O_2O_1)^\dagger - \tilde{\Psi}\right\| \le 5\left\|U_S(O_2O_1)^\dagger - \tilde{U}_S\right\| \frac{\|V_L\|^2}{\sigma_{\min}^2(V_{L-1})}. \tag{6.80}$$

Proof. Following [32] we have

$$\left\|(O_2O_1)\Psi(O_2O_1)^\dagger-\tilde{\Psi}\right\| \le \left\|(O_2O_1)U_0^+\right\|\left\|U_1(O_2O_1)^\dagger-\tilde{U}_1\right\| + \left\|(O_2O_1)U_0^+-\tilde{U}_0^+\right\|\left\|\tilde{U}_1\right\|. \tag{6.81}$$

We have $\left\|\tilde{U}_1\right\| \le \left\|\tilde{U}_S\right\| = 1$, since $\tilde{U}_S^\dagger\tilde{U}_S = I_S$ and removing a row vector decreases the operator norm. Similarly we have $\left\|U_1(O_2O_1)^\dagger-\tilde{U}_1\right\| \le \left\|U_S(O_2O_1)^\dagger-\tilde{U}_S\right\|$. Now note that by our initial assumption and Lemma 6.D.5 (with $A = U_0(O_2O_1)^\dagger$ and $B = \tilde{U}_0$) we have $\mathrm{Rank}(U_0(O_2O_1)^\dagger) = \mathrm{Rank}(\tilde{U}_0)$. This means that we can use Lemma 6.D.4 to conclude that

$$\left\|(O_2O_1)U_0^+-\tilde{U}_0^+\right\| \le \frac{1+\sqrt{5}}{2}\frac{\left\|U_0(O_2O_1)^\dagger-\tilde{U}_0\right\|}{\sigma^2_{\min}(U_0)} \le \frac{1+\sqrt{5}}{2}\frac{\left\|U_S(O_2O_1)^\dagger-\tilde{U}_S\right\|}{\sigma^2_{\min}(U_0)}. \tag{6.82}$$

Hence we get

$$\begin{aligned}\left\|(O_2O_1)\Psi(O_2O_1)^\dagger-\tilde{\Psi}\right\| &\le \frac{\left\|U_S(O_2O_1)^\dagger-\tilde{U}_S\right\|}{\sigma_{\min}(U_0)} + \frac{1+\sqrt{5}}{2}\frac{\left\|U_S(O_2O_1)^\dagger-\tilde{U}_S\right\|}{\sigma^2_{\min}(U_0)} \\ &\le (\sigma_{\min}(U_0)+(1+\sqrt{5})/2)\frac{\left\|U_S(O_2O_1)^\dagger-\tilde{U}_S\right\|}{\sigma^2_{\min}(U_0)} \\ &\le \left(\frac{3}{2}+\sqrt{5}\right)\frac{\left\|U_S(O_2O_1)^\dagger-\tilde{U}_S\right\|}{\sigma^2_{\min}(U_0)},\end{aligned} \tag{6.83}$$

where we used $\sigma_{\min}(U_0) \le \|U_0\| \le 1$. Plugging in the lower bound on $\sigma_{\min}(U_0)$ (Lemma 6.D.6) and noting that $\frac{3}{2}+\sqrt{5} \le 5$ we obtain the Lemma statement. □

Combining all of this we get the following general theorem. From now on we restrict ourselves to the case $L = K/2$:

Theorem 6.D.8. *Let $(g+\eta)(k)$ be the signal with $g(k) = \sum_{i=1}^S c_i z_i^k$ and $\eta(k)$ a small noise vector to which we apply the ESPRIT algorithm 6.2. Consider then the associated Hankel matrices $H(g)$ and $H(g+\eta)$, with singular value decompositions $H(g) = U\Sigma W$ and $H(g+\eta) = \tilde{U}\tilde{\Sigma}\tilde{W}$, and label the matrix of the first S columns of U (resp. $\tilde{U}$) as U_S (resp. as $\tilde{U}_S$). Denote by U_0 (resp. U_1) the submatrix of U_S with the last row (resp. first row) removed and define the signal matrix $\Psi = U_0^+U_1$ (similarly for $\tilde{\Psi}$). Let $L = K/2$ and $K+1 \ge 2S$. Now, if*

$$4\left\|H(\eta)\right\| \le c_{\min}\sigma^2_{\min}(V_{K/2})\sigma_{\min}(V_{K/2-1})\|V_{K/2}\|^{-1}, \tag{6.84}$$

then

$$\left\|(O_2O_1)\Psi(O_2O_1)^\dagger-\tilde{\Psi}\right\| \le \frac{10\sqrt{2S}\left\|H(\eta)\right\|\|V_{K/2}\|^2}{\sigma^4_{\min}(V_{K/2-1})}c_{\min}^{-1}. \tag{6.85}$$

where O_2O_1 is defined through the singular value decomposition of $U_S^\dagger\tilde{U}_S$, equation (6.64).

Proof. We start from the requirement in Lemma 6.D.7 that $\left\|U_0(O_2O_1)^\dagger-\tilde{U}_0\right\| \le \left\|U_S(O_2O_1)^\dagger-\tilde{U}_S\right\| \le \sigma_{\min}(U_0)/2$. By Lemma 6.D.6, this is certainly satisfied if

$\left\|U_S(O_2O_1)^\dagger - \tilde{U}_S\right\| \leq \sigma_{\min}(V_{K/2-1})\|V_{K/2}\|^{-1}/2$. Moreover, from Lemma 6.D.3 we know that $\left\|U_S(O_2O_1)^\dagger - \tilde{U}_S\right\| \leq 2\sqrt{2S}\left\|H(\eta)\right\|\sigma^{-1}_{\min}(H(g))$ so now let's upperbound $\sigma^{-1}_{\min}(H(g))$ in terms of the smallest singular value of $V_{K/2}$. We have

$$\sigma^{-1}_{\min}(H(g)) = \left\|H(g)^+\right\| = \left\|(V^T_{K/2})^+ C^{-1} V^+_{K/2}\right\| \leq \sigma^{-2}_{\min}(V_{K/2})c^{-1}_{\min}, \tag{6.86}$$

where we used $c_{\min} = \min_i c_i$. This means that the condition

$$4\left\|H(\eta)\right\| \leq c_{\min}\sigma^2_{\min}(V_{K/2})\sigma_{\min}(V_{K/2-1})\|V_{K/2}\|^{-1}, \tag{6.87}$$

allows us to use Lemma 6.D.7

$$\begin{aligned}\left\|(O_2O_1)\Psi(O_2O_1)^\dagger - \tilde{\Psi}\right\| &\leq 10\sqrt{2S}\frac{\left\|H(\eta)\right\|\|V_{K/2}\|^2}{\sigma^2_{\min}(V_{K/2})\sigma^2_{\min}(V_{K/2-1})}c^{-1}_{\min} \\ &\leq \frac{10\sqrt{2S}\left\|H(\eta)\right\|\|V_{K/2}\|^2}{\sigma^4_{\min}(V_{K/2-1})}c^{-1}_{\min}, \end{aligned} \tag{6.88}$$

where we also used the general fact about Vandermonde matrices [36, Theorem 1] that $\sigma_{\min}(V_{K/2}) \geq \sigma_{\min}(V_{K/2-1})$ (i.e. the smallest non-zero singular value of V_L grows monotonically with L). □

We wish to translate the bound in Theorem 6.D.8 to a theorem on the distance between the inferred eigenvalues z_i and $\tilde{z}_i$. The argument is as follows. We know from the Bauer-Fike theorem and the fact that $(O_2O_1)\Psi(O_2O_1)^\dagger$ is diagonalizable (see [35, Exercise VIII.3.2]) that the operator norm bound on $(O_2O_1)\Psi(O_2O_1)^\dagger$ implies a matching bound distance on its eigenvalues z_i, using equation (6.76). That is, we have

$$d(\{z_i\},\{\tilde{z}_j\}) := \max_{\pi\in\mathrm{Perm}_S}\min_i |z_{\pi(i)} - \tilde{z}_i| \leq (2S-1)\kappa(A(O_2O_1)^\dagger)\left\|(O_2O_1)\Psi(O_2O_1)^\dagger - \tilde{\Psi}\right\|, \tag{6.89}$$

where $\kappa(A(O_2O_1)^\dagger) = \kappa(A) := \|A\|\left\|A^{-1}\right\|$ is the condition number of the invertible matrix A in equation (6.74). We have $A = V^+_{K/2}U_S$ and since U_S is an isometry we know that

$$\kappa(A) = \|A\|\left\|A^+\right\| = \|V_{K/2}\|\left\|V^+_{K/2}\right\| = \frac{\|V_{K/2}\|}{\sigma_{\min}(V_{K/2})}, \tag{6.90}$$

and hence we get a bound on the matching distance of the eigenvalues in terms of known quantities, as expressed in the following Theorem:

Theorem 6.D.9. *Let $y(k) = (g+\eta)(k)$ $(k = 0,\ldots,K)$ be the signal with $g(k) = \sum_{i=1}^S c_i z_i$, let $\eta(k)$ a small noise vector, and $K+1 \geq 2S$ $(L = K/2)$. Let $\tilde{z}_i$ be the output of the ESPRIT algorithm. Then under the noise condition:*

$$4\left\|H(\eta)\right\| \leq c_{\min}\sigma^2_{\min}(V_{K/2})\sigma_{\min}(V_{K/2-1})\|V_{K/2}\|^{-1}, \tag{6.91}$$

we have

$$d(\{z_i\},\{\tilde{z}_j\}) \leq (2S-1)\frac{10\sqrt{2S}\left\|H(\eta)\right\|\|V_{K/2}\|^3}{\sigma^5_{\min}(V_{K/2-1})}c^{-1}_{\min}. \tag{6.92}$$

This theorem thus holds for both the decaying as well as oscillatory signal.

In order to use the Theorem, one has to lower bound $\sigma_{\min}(V_{L-1})$ (and more trivially, upper bound $\|V_L\|$ as well) for $L = K/2$. For complex poles z_i on the unit circle one can get very good bounds, assuming a gap, see equation (6.34).

To lower bound $\sigma_{\min}(V_{L-1})$ for a purely decaying signal, we start with the following characterization of *square* Vandermonde matrices with real poles due to Gautschi:

Theorem 6.D.10 (Theorem 1 in [37])**.** *Let V_{S-1} be a square $S \times S$ Vandermonde matrix with S (unequal) real positive poles $z_1, \ldots, z_S$. Then ∞ norm of V_{S-1}^{-1} is*

$$\left\| V_{S-1}^{-1} \right\|_\infty := \max_i \sum_j \left| \left(V_{S-1}^{-1} \right)_{ij} \right| = \max_{i \in \{1,\ldots,S\}} \prod_{j=1, j\neq i}^{S} \frac{1+z_i}{|z_j - z_i|}. \tag{6.93}$$

Based on this Theorem we can work out a very similar statement for non-square Vandermonde matrices V_{L-1} where L is a multiple of S. Note that this Lemma does not depend on any gap.

Lemma 6.D.11. *Let V_{ST-1} be an $ST \times S$ Vandermonde matrix (where T is a positive integer) with S (unequal) real positive poles $z_1, \ldots, z_S \leq 1$. Then we have*

$$\left\| V_{ST-1}^{+} \right\|_\infty \leq 2 \left\| V_{S-1}^{-1} \right\|_\infty. \tag{6.94}$$

Proof. Note that

$$V_{ST-1} = \begin{pmatrix} V_{S-1}^T & Z^S V_{S-1}^T & Z^{2S} V_{S-1}^T & \cdots & Z^{(T-1)S} V_{S-1}^T \end{pmatrix}^T, \tag{6.95}$$

with $Z = \mathrm{diag}(z_1, \ldots, z_S)$, using $(V_{S-1})_{ij} = z_j^{i-1}$ and $(Z^S V_{S-1}^T)_{ij} = z_i^{s+j-1}$.

The pseudo-inverse V_{ST-1}^+ of size $S \times ST$ can be directly calculated as

$$V_{ST-1}^{+} = (I - Z^{2S})(I - Z^{2ST})^{-1} \begin{pmatrix} V_{S-1}^{-1} & Z^S V_{S-1}^{-1} & Z^{2S} V_{S-1}^{-1} & \cdots & Z^{(T-1)S} V_{S-1}^{-1} \end{pmatrix}, \tag{6.96}$$

using a geometric series. Hence $\left\| V_{ST-1}^{+} \right\|_\infty$ can be calculated to be

$$\begin{aligned} \max_{i \in \{1,\ldots,S\}} \frac{1-z_i^{2S}}{1-z_i^{2ST}} \sum_{p=0}^{T-1} z_i^{pS} \prod_{j=1, j\neq i}^{S} \frac{1+z_i}{|z_j - z_i|} \leq \max_i \left[\frac{1-z_i^{2S}}{1-z_i^{2ST}} \frac{1-z_i^{ST}}{1-z_i^{S}} \right] \left\| V_{S-1}^{-1} \right\|_\infty = \\ \max_i \left[\frac{1+z_i^S}{1+z_i^{ST}} \right] \left\| V_{S-1}^{-1} \right\|_\infty \leq 2 \left\| V_{S-1}^{-1} \right\|_\infty. \end{aligned} \tag{6.97}$$

which gives the Lemma statement. □

To apply this Lemma, we use that $\sigma_{\min}^{-1}(V_{ST-1}) = \left\| V_{ST-1}^{+} \right\| \leq \sqrt{S} \left\| V_{ST-1}^{+} \right\|_\infty$ so that

$$\sigma_{\min}(V_{ST-1}) \geq (2\sqrt{S})^{-1} \left\| V_{S-1}^{-1} \right\|_\infty^{-1}, \tag{6.98}$$

for any T. Unlike the lower bound for the real-time signal which explicitly uses the gap Δ in equation (6.34), this bound does not improve with T. We will now use the gap to upper bound $\left\| V_{S-1}^{-1} \right\|_\infty$ given in equation (6.93), thus lower bounding $\sigma_{\min}(V_{ST-1})$ for any T. This is done in the proof of our final Theorem 6.3.3 (and its slight adaptation Theorem 6.3.4) restated here:

Theorem. *Let $(g+\eta)(k)$ be an imaginary-time decaying signal (of length K) with $g(k) = \sum_{i=1}^{S} c_i z_i^k$, $c_i > 0\,\forall i$, $c_{\min} = \min_i c_i$ and $\eta(k)$ a small noise vector. Let $z_i = e^{-E_i}$ with $E_i \in [0, 2\pi)$ and given eigenvalue gap $\Delta < 1$ in equation (6.33), and $\{\tilde{E}_i\}$ the energy estimates of ESPRIT with $L = K/2$. Let $K+1 \geq 2S$, K even and $K = TS$ for some positive integer T. If we have*

$$\left\| H(\eta) \right\| \leq \frac{c_{\min}}{\sqrt{K}} g_1(S, \Delta), \tag{6.99}$$

with

$$g_1(S, \Delta) = \frac{1}{32 S^2} \left(e^{-2\pi} \pi \Delta\right)^{3(S-1)}, \tag{6.100}$$

then

$$d(\{\tilde{E}_i\}, \{E_j\}) \leq \left\| H(\eta) \right\| c_{\min}^{-1} K \sqrt{K} g_2(S, \Delta), \tag{6.101}$$

with

$$g_2(S, \Delta) = e^{2\pi} \sqrt{2} 640\, S^{5.5} \left(e^{-2\pi} \pi \Delta\right)^{-5(S-1)}. \tag{6.102}$$

Proof. First, we can lift the bound on the eigenvalue distance $d(\{z_i\}, \{\tilde{z}_j\})$ to one on energies $d(\{E_i\}, \{\tilde{E}_j\})$ defined in equation (6.32), with $\tilde{E}_i := -\log(\tilde{z}_i)$ and $E_i \in [0, 2\pi)$ by noting that

$$\begin{aligned} \frac{1}{2\pi} |E_{\pi(i)} - \tilde{E}_i| = \frac{1}{2\pi} |\log\left(z_{\pi(i)}\right) - \log(\tilde{z}_i)| = \frac{1}{2\pi} |\log\left(1 - (\tilde{z}_i - z_{\pi(i)})/\tilde{z}_i\right)| \\ \leq |(\tilde{z}_i - z_{\pi(i)})/\tilde{z}_i\big)| \leq e^{2\pi} |(\tilde{z}_i - z_{\pi(i)})|, \end{aligned} \tag{6.103}$$

using that $z_i \in (e^{-2\pi}, 1]$ and $\tilde{z}_i \in (e^{-2\pi}, 1]$. In particular, for the first inequality, let $x = (\tilde{z}_i - z_{\pi(i)})/\tilde{z}_i$. If $x < 0$, $|\log(1-x)| \leq |x|$. If $x > 0$, since $z_i \in (e^{-2\pi}, 1]$, we have $x \leq 1 - e^{-2\pi}$, so that $|\log(1-x)| \leq 2\pi |x|$.

Second, let us now use the gap condition $|E_i - E_j| \geq 2\pi\Delta$ in equation (6.33). This leads to a gap condition on the z_i themselves through (assuming w.l.o.g. that $z_i \geq z_j$):

$$2\pi\Delta \leq |E_i - E_j| = |\log\left(z_i / z_j\right)|, \tag{6.104}$$

and thus

$$e^{2\pi\Delta} \leq \frac{z_i}{z_j} = \frac{z_i - z_j}{z_j} + 1, \tag{6.105}$$

which gives

$$|z_i - z_j| = z_i - z_j \geq z_j (e^{2\pi\Delta} - 1) \geq e^{-2\pi} 2\pi\Delta. \tag{6.106}$$

This implies through equation (6.93) that

$$\left\| V_{S-1}^{-1} \right\|_\infty \leq (\pi\Delta)^{-(S-1)} e^{2\pi(S-1)}, \tag{6.107}$$

so that

$$\sigma_{\min}(V_{ST-1}) \geq (2\sqrt{S})^{-1} (\pi\Delta e^{-2\pi})^{(S-1)}, \tag{6.108}$$

for any integer T. We note that this lower bound on $\sigma_{\min}$ is exponentially small in S as $\Delta < 1$.

Third, we need an upper bound on $\|V_L\|$ in order to use Theorem 6.D.9. For $z_j \in (0,1]$, we have $\|V_L\| \leq \|V_L\|_F = \left(\sum_{i,j} z_j^{2(i-1)}\right)^{1/2} \leq S^{1/2}\left(\sum_{i=1}^{L+1} z_{\max}^{2(i-1)}\right)^{1/2} = S^{1/2}\left(\frac{1-z_{\max}^{2(L+1)}}{1-z_{\max}^2}\right)^{1/2}$ which can tend to $(S(L+1))^{1/2}$ when $z_{\max} \to 1$, so we use the simple upper bound $(S(L+1))^{1/2} \leq (SK)^{1/2}$. Putting all this together allows to translate equation (6.92) to equation (6.101). The condition on $\left\|H(\eta)\right\|$ in equation (6.91) then translates to the sufficient condition in equation (6.99) using that $\sigma_{\min}(V_{K/2}) \geq \sigma_{\min}(V_{K/2-1})$, the lower bound on $\sigma_{\min}(V_{K/2-1})$, and the upper bound on $\|V_{K/2}\|$. □

The adapted version, Theorem 6.3.4, is proved almost identically (but requires that all E_i are in principle bounded away from 2π):

Theorem. *Let $(g+\eta)(k)$ be an imaginary-time decaying signal (of length K) with $g(k) = \sum_{i=1}^{S} c_i z_i^k$, $c_i > 0$, $\forall i$, $c_{\min} = \min_i c_i$, and $\eta(k)$ a small noise vector. Let $z_i = 1 - E_i/2\pi$ with $E_i \in [0,\pi]$ and given eigenvalue gap $\Delta < 1$ in equation (6.33), and $\{\tilde{E}_i\}$ the energy estimates of ESPRIT with $L = K/2$. Let $K+1 \geq 2S$, K even and $K = TS$ for some positive integer T. If we have*

$$\left\|H(\eta)\right\| \leq \frac{c_{\min}}{\sqrt{K}}\tilde{g}_1(S,\Delta), \tag{6.109}$$

with

$$\tilde{g}_1(S,\Delta) = \frac{1}{32S^2}\Delta^{3(S-1)}, \tag{6.110}$$

then

$$d(\{\tilde{E}_i\},\{E_j\}) \leq \left\|H(\eta)\right\| c_{\min}^{-1} K\sqrt{K}\tilde{g}_2(S,\Delta), \tag{6.111}$$

with

$$\tilde{g}_2(S,\Delta) = 640\sqrt{2}\,S^{5.5}\,\Delta^{-5(S-1)}. \tag{6.112}$$

Proof. First, we convert the eigenvalue distance $d(\{z_i\},\{\tilde{z}_j\})$ to one on energies $d(\{E_i\},\{\tilde{E}_j\})$ defined in equation (6.32), with $\tilde{E}_i := 2\pi(1-\tilde{z}_i)$, so

$$\frac{1}{2\pi}|E_{\pi(i)} - \tilde{E}_i| = |z_{\pi(i)} - \tilde{z}_i|. \tag{6.113}$$

Second, let us now use the gap condition $|E_i - E_j| \geq 2\pi\Delta$ in equation (6.33). This leads to a gap condition on the z_i themselves through:

$$|z_i - z_j| \geq \Delta. \tag{6.114}$$

This implies through equation (6.93) that

$$\left\|V_{S-1}^{-1}\right\|_\infty \leq \Delta^{-(S-1)} \tag{6.115}$$

so that

$$\sigma_{\min}(V_{ST-1}) \geq (2\sqrt{S})^{-1}\Delta^{(S-1)}, \tag{6.116}$$

for any integer T. Following identical steps as in the proof of the previous Theorem then leads to the final statements. □

7

Quantum phase estimation without controlled unitaries

En ik kan niet
Ik kan er niet er omheen
Ik kan het niet
Ik kan het niet alleen

Ik kan het niet alleen (1989), De Dijk

7.1. Introduction

As discussed in Section 1.2, the evaluation of spectral properties of large many-body systems is a computationally intensive task, limited by numerical bottlenecks associated with classical simulation of time-evolution [1]. Quantum computers have the potential to make such simulations tractable, offering the opportunity to revolutionize material science and chemistry [2]. Algorithms based on quantum phase estimation offer a route towards obtaining such spectral properties, albeit requiring one to perform many costly time-evolution simulation steps *controlled* on one or several ancillary qubit(s), making it difficult to run on current *fault-prone* quantum computers. Therefore, eliminating or reducing the need for controlled time-evolution in quantum phase estimation has attracted a lot of attention, as it could render the prediction of spectral properties accessible to near-term devices.

A significant step in this direction was the proposal of statistical phase estimation (see e.g. [3–5]) which reduces quantum phase-estimation to the computation of a time series $f(t) = \langle\psi| e^{iHt} |\psi\rangle$ on a quantum device followed by classical post-processing to obtain an estimate of the spectrum of H with similar resolution as standard quantum phase estimation for most Hamiltonians. As shown in Figure 1.a, statistical phase estimation reduces the required quantum circuit to two variants of a Hadamard test circuit, giving the real and imaginary parts of $f(t)$. This result has spurred a plethora of theoretical work [6–8], followed by experimental demonstrations [9, 10], which can be adapted to estimation of more general spectral quantities such as linear response of quantum systems [11], operator-resolved density of states [12] or Green functions

[13], to give some examples. Even if statistical phase-estimation represents indisputable progress toward making quantum phase estimation more reachable for near-term devices, a single-qubit control time-evolution remains a highly non-trivial operation to implement on near-term devices. This makes it often impractical to implement statistical phase estimation on currently existing hardware (beyond trivial-size problems). Its difficulty arises from the significant numbers of qubits and non-trivial circuit depth time-evolution operations that need to be controlled by one (or few) qubit(s), making it extremely sensitive to decoherence, increasing the cost of each time-evolution step, and also significantly increasing the circuit depth of implementing all of the required controlled-time-evolution operations.

The significant limitation imposed by the *controlled* time-evolution operation has led to an effort to develop algorithmic techniques to retrieve spectra using control-free time evolution operations. Note that the quantity $|\langle\psi|\, e^{iHt}\, |\psi\rangle|^2$ can be obtained with the circuit shown in Figure 1.b, which does not make use of a controlled operation. Unfortunately, this only gives us access to the absolute values of the time-series $f(t)$, while its phases are key in recovering the spectrum. One way of exploiting the access to $|\langle\psi|\, e^{iHt}\, |\psi\rangle|^2$, is to compute its Fourier transform to retrieve a new spectrum composed of differences between energy levels, rather than obtaining the energies directly, as done in [14]. Alternatively, in [15] the authors use a complex-time evolution, and in [16] use the knowledge of one eigenvector and eigenenergies, to obtain information about the spectrum using control-free operations. As we detail in the next subsection, each technique has its own limitations and sometimes strong constraints, leaving room for significant improvements.

Retrieving a spectrum while not having (complete) access to the phases of its signal is a well-known problem in classical signal-processing called phase-retrieval [17], which has been extensively studied and offers a plethora of different solutions, each one having its own benefits and set of most appropriate applications. The time-series obtained on a digital quantum computer or quantum simulator can be understood as a classical signal and retrieving the spectrum of the many-body problem can therefore benefit from phase-retrieval ideas and algorithms.

In this work we show how the adaptation to quantum problems of two different techniques, *vectorial phase-retrieval* [18] and the *two-dimensional phase-retrieval* [19], allows us to remove the intricate controlled time-evolution currently limiting the scalability of phase-estimation algorithms. Vectorial phase-retrieval works by including measurements of absolute values of additional time series, and – crucially – measurements of the absolute values of interferences of the target time series and these additional time series. In addition, we show how to recast the problem as a two-dimensional one, which can be solved with the well known hybrid input-output algorithm [19]. We numerically investigate the feasibility of both approaches for estimating the spectrum of the Fermi-Hubbard model and discuss the resilience of the algorithms to inherent statistical noise.

7.1.1. Related work

There are a couple of techniques that allow to compute time-series without having to implement a controlled time evolution. Both require as condition the knowledge and

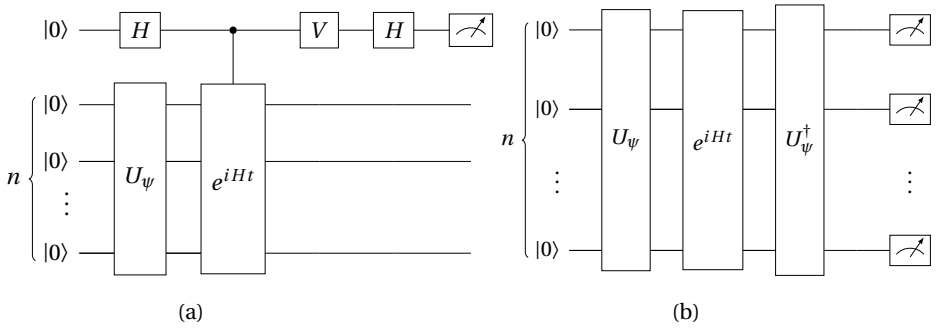

Figure 7.1: (a) Single-ancillary-qubit quantum phase estimation computing a time series $f(t) = \langle\psi| e^{iHt} |\psi\rangle$ for time t. The circuit U_ψ prepares the input state $|\psi\rangle$, which is followed by a circuit implementing e^{iHt} controlled by the single-ancillary-qubit. The gate V is an identity ($S^\dagger$ gate) when estimating the real (imaginary) part of the time series. (b) Ancillary-qubit-free circuit that allows one to estimate $|f(t)|^2 = |\langle\psi| e^{iHt} |\psi\rangle|^2$.

easy preparation of an eigenstate of the Hamiltonian of interest. If this can be easily satisfied for many physically motivated Hamiltonians, for example for all Hamiltonians having vacuum as an eigenstate, the techniques will fail for arbitrary Hamiltonians. The first approach, detailed in [20], requires doubling the number of qubits and the implementation of a pair of control-SWAP operations over the two register, where the time-evolution is applied locally on one of them and without the needs of control. It is worth pointing that for current QPU devices, doubling the number of qubits may turn unfeasible what could be an initial good candidate to quantum advantage. If doubling the number of qubits is not a big limitation, the approach has a limited additional cost given by the implementation pair of control-SWAP, which will result in a extra cost of $O(N)$ CNOT gates. The second approach [16, 21–23] requires the preparation of a superposition $|\psi\rangle + e^{i\theta} |R\rangle$, where $|R\rangle$ is an eigenvector of the Hamiltonian that will play the role of a phase reference. One can then exploit post-processing algorithms to extract the value of our time-series of interest. Unfortunately, even in the case where $|R\rangle$ is trivial to prepare, for example an all zeros state, creating the superposition would be in general cases non-trivial, as it will require a controlled implementation of $U_\psi^\dagger$. Again, removing of the control time-evolution shifts the burden to an extra cost that depends on the original gate count of $U_\psi^\dagger$, leading to a cost that would be of at least $O(N)$ CNOT gates but in many cases $O(N^2)$.

Last year, there were two works that removed those constraints on the input state requirements, making ancillary-free quantum phase-estimation closer to the range of applicability of traditional quantum phase-estimation. In [14], Chan et al. leverage the power of the recently proposed classical shadows technique to estimate energy gaps via the computation of expectation values of certain operators as a function of time. A downside of the technique is that it cannot provide the spectrum of $\langle\psi| e^{iHt} |\psi\rangle$, but rather the spectrum of $|\langle\psi| e^{iHt} |\psi\rangle|^2$, which leaves the challenge of recovering

the eigenvalues of the Hamiltonian from the set of all pairwise energy gaps. Almost at the same time, Yang et al. showed that one can remove the ancillary qubit by combining real-time and imaginary-time evolution [15], exploiting ideas from complex function analysis. Its main limitation is the need for an additional implementation of imaginary-time evolution on a quantum device itself, which currently works only for input states with finite correlation length and short evolution times, but improvement in this direction could render the technique more widely applicable. Finally, in the recent work [12] the authors propose a hybrid classical-quantum algorithm for calculating quantities of interest in many-body spectroscopy that exploits ideas inspired by classical shadows that do not require a controlled time-evolution due to particular properties of the spectral objects of interest. Nonetheless, the approach requires the application of a (potentially shallow) controlled unitary gate independent of the time evolution. In addition, despite the promise of delegating part of the computation to a classical computer, the constraints it imposes in terms of having access to families of states that can be classically emulated may limit its use for arbitrary Hamiltonians.

7.1.2. Comparison of controlled vs uncontrolled time evolution

There are two key advantages of removing the controlled time-evolution: (1) It takes out a significant source of errors resulting from the sensitivity of the original scheme to the decoherence of the single qubit controlling the dynamics; (2) It requires a smaller number of gates and less circuit depth to implement time evolution compared to its controlled counterpart. Here we summarize in Table 7.1 the reduction in circuit complexity achieved when removing the time-evolution for three simple examples that capture the essence of the discussion: i) the 1D Transverse-Field Ising Model (TFIM) implemented on hardware with all-to-all connectivity; ii) 1D TFIM with matching hardware connectivity; iii) the 2D Fermi-Hubbard model with square-lattice hardware connectivity. In each case we calculate the complexity of 1D vectorial phase retrieval. We do not take into account the cost of preparing or uncomputing the initial state $|\psi\rangle$, which we assume to be simple, nor the cost of preparing a superposition of $|\psi\rangle$ and another state[1]. However, we also make assumptions which significantly favor the controlled time evolution calculation which we use for comparison. We do not cost the extra control qubits used to reduce the circuit depth when implementing controlled time dynamics on 1D and 2D architectures. More details on the discussion and derivation can be found in Appendix 7.C. These calculations do not take into account any Hamiltonian specific techniques which could be used to optimize further the gate counts.

In the case of quantum hardware with all-to-all connectivity, a quantum circuit computing a unitary operator U can be converted into a quantum circuit for computing controlled-U by replacing every CNOT gate with a Toffoli gate (requiring 6 CNOT gates each), and every single-qubit unitary V with a controlled-V operator (requiring 2 CNOT gates each). Despite being only a constant factor gain, this can make a significant difference to the ability of near-term quantum computers to obtain meaningful results, by increasing the number of accessible Trotter steps.

[1] In our experiments, the superposition used was essentially a GHZ state, but we do not believe this is a fundamental requirement.

However, the difference in complexity between the two situations can be significantly greater if there are restrictions on hardware connectivity, as in 1D and 2D architectures. It is instructive to compare the bounds obtained in a realistic regime for near-term quantum algorithms. As an example, assume that we have a quantum computer with 100 qubits that can execute quantum circuits of CNOT depth at most 100 accurately. Then, based on the algorithmic complexities in Table 7.1, we could not even implement 1 Trotter step for the 1D TFIM under 1D connectivity if the standard method is used, but we could implement 25 Trotter steps using phase retrieval. The reductions in gate count are less significant than depth, but are still relatively substantial (e.g. approximately a factor of 3 for the 1D TFIM).

Model	H/w	CNOTs (PR)	CNOTs (no PR)	Depth (PR)	Depth (no PR)
1d TFIM	All-to-all	$(2n-2)k$	$(6n-4)k$	$4k$	$2\lceil \log_2 n \rceil + 10k$
1d TFIM	1d	$(2n-2)k$	$6(n-1)$ $+(6n-4)k$	$4k$	$6\lceil n/2 \rceil + 10k$
2d FH	2d	$32kn(n-1)/2$	$48kn(n-1)/2$ $+6\lceil (n-1)^2/2 \rceil$	$32k$	$48k+3(n-2)$

Table 7.1: The cost of implementing k Trotter layers for several Hamiltonians, an n qubits transverse field Ising model and an $n \times n$ spinless Fermi-Hubbard model. "Depth" is CNOT depth. Note that the algorithms achieving the minimal CNOT depth and CNOT count may be different. "2d FH" is Fermi-Hubbard without spin, where results are approximate. In each case "PR" refers to the 1D vectorial phase retrieval approach, in which we do not consider the cost of preparing the initial state.

A key aspect of making phase-retrieval work is to enlarge the problem that we wish to solve in a "cheap" way, either by enlarging the set of input states included in the procedure, or by implementing additional (comparatively simple) time dynamics. The saving in terms of circuit depth therefore comes at a price of a larger number of circuits to be implemented and required shots.

7.1.3. Structure of this chapter

Section 7.2 summarizes our results on the adaptation of both vectorial phase-retrieval and two-dimensional phase-retrieval to statistical quantum phase-estimation with a final comparison of the performance of both techniques. This is followed by a conclusion and discussion (Section 7.3). This chapter also contains two appendices on the technical preliminaries (Appendix 7.A), and the full details of our adaptation of vectorial phase-retrieval (Appendix 7.B) to statistical phase-estimation, the full details of our adaptation of two-dimensional phase-retrieval can be found in Appendix C of Ref. [24].

7.2. Summary of results

The classical post-processing in statistical phase-estimation is a standard signal processing problem: One is provided with a noisy version of some complex-valued signal $f(t)$, which is typically measured at integer multiples of some time increment Δt (resulting in a discrete time-series[2] $f[t]$). The goal is to approximately determine the Fourier transform $F(\omega)$ of the continuous signal, typically corresponding to some quantity of interest such as the spectrum of a Hamiltonian. This is a non-trivial task due to noise and to finite measurement time T. There exists extensive literature on this problem with applications to multiple areas of engineering, such as optical imaging, spectroscopy and audio processing [25–28]. We note that in the statistical phase-estimation scenario, the largest measurement time T is bounded by the largest circuit depth we can afford in the simulation of the *controlled* time evolution unitary e^{iHT}.

The goal of phase-retrieval as presented in classical signal processing literature is to reconstruct the signal from the absolute values of its discrete Fourier transform. At first sight, this is an impossible task, as the phases are key to achieve the reconstruction. As with any no-go result, the question is which rules one needs to change to circumvent it. This can be successfully achieved by extending the problem that one wishes to solve to one with a larger set of input signals and exploiting natural constraints on the signal and/or the spectrum. A key non-trivial insight is to understand when this generalized and constrained problem has a unique solution that recovers all phases of the signals. If such an insight can be used to develop an algorithm that solves the problem in the ideal noiseless case, a non-trivial aspect is to design algorithms that work efficiently in realistic regimes of noise and have reasonable computational cost. This work is devoted to designing these algorithms in the scenario of ancillary-free quantum phase-estimation (QPE). It is important to note that phase-retrieval literature defines *unique solution* (for the spectrum) up to unavoidable trivial ambiguities corresponding to a multiplication by a global phase, a translation, and reflection with complex-conjugation of the recovered spectrum $F[k]$. The first two have no real physical impact, as experimentally we only have access to energy differences while the final ambiguity can reflect the spectrum, an effect that can be mitigated if handled with proper care.

The adaptation of phase-retrieval algorithms to the quantum-phase-estimation scenario is unfortunately not a trivial “plug & play". First it requires a conceptual adaptation of the framework. The standard problem in classical phase-retrieval literature consists of reconstructing the phases $e^{i\theta[t]}$ of $f[t]$ while having access to $|F[\omega]|$, where in our case we know $|f[t]|$ and want to recover $e^{i\theta[t]}$ (and thereby $F[\omega]$) instead. Because of the symmetries of the standard classical signal processing problem where both signal and its spectrum are potentially complex functions, we can most of the time circumvent this slight difference by just exchanging the role of time and frequency in most phase-retrieval algorithms. Secondly, one needs to find an enlarged set of time-series, by enlarging the set input states or inducing additional dynamics, that combined with natural constraints on the spectrum, such as its non-negativity or the fact that it has a well-defined support, guarantee an efficient recovery of the phases $e^{i\theta[t]}$ (we shall come back to this point below and further details are provided in Section 7.B and in

[2] We use square brackets to stress that the function takes discrete input.

Appendix C of Ref. [24]). This will finally allow approximate reconstruction of our initial spectrum of interest, but – importantly – also of the phases of the time series $f[t]$. In this manuscript, we show how to re-design two well-known phase-retrieval techniques, vectorial phase-retrieval and two-dimensional phase-retrieval, to address the problem of ancillary-free quantum phase-estimation. Vectorial phase-retrieval [18] achieves its goal by including measurements of the absolute values of multiple well-behaved time-series, and – importantly – the time-series of their interferences. Where standard vectorial phase-retrieval makes use of two signals (the target signal and one additional signal) and their interferences, we generalize to multiple additional signals to improve the algorithm's resilience to noise in the quantum phase-estimation framework. The total number of required samples equals $(1+3R)$ times the number of samples taken for each individual times series, where R denotes the number of additional signals used in the analysis. Secondly, we show how to transform a one-dimensional time series (corresponding to time evolution of an input state under a Hamiltonian) into a two dimensional phase-retrieval problem by adding a dummy Hamiltonian H_D that commutes with H and produces non-trivial dynamics on the input state of interest. We then exploit the well-known *Hybrid Input Output* algorithm to retrieve the target spectra. The sampling cost then increases by a factor equal to the number of measurement points taken in this second "time" dimension.

It is important to point out that our objective is never to recover the individual eigenvalues in the spectrum of a given Hamiltonian, which is not practical even for a fault-tolerant computer due to the exponential increase of the number of spectral peaks with the size of the system, but to retrieve a realistic spectrum resulting from a time evolution of time T and time-steps ΔT. The unavoidable prescience of error in near-term devices affects both statistical phase-estimation and phase-retrieval techniques, and reconstructing the initial exact signal is a complex task beyond the scope of this chapter.

A fair goal for this work is to compare the phase-retrieval post-processing to its statistical phase-estimation counterpart, where both have the same realistic constraint on the time-dynamics simulation errors (Trotter errors for example) and choices of accessible time and size of time intervals. It is worth stressing we are not judging the quality of the output signal from a physicist, chemist or material scientist point of view, but rather given a time-series as reference, which quality can depend on the quality of time-dynamics simulation, how the trivial recovery by statistical phase-estimation compares to the non-trivial one using phase-retrieval techniques. We therefore focus on the errors resulting from the finite-size of samples that induces a shot noise in the evaluation of the expectation values of interest. Demonstrating a successful recovery by phase-retrieval algorithms in this regime is not trivial and is the focus of our work here. In a nutshell, the effect of Trotter errors or the effect of time-interval choices is to shift the ground truth for the pot-processing analysis. On the other hand, the gate errors can have slightly different impacts, as statistical phase-estimation implements a different circuit that the ones we propose here. Statistical phase-estimation needs a control time-evolution that has a higher gate counts than standard time-dynamics, while phase-retrieval may need a reference state preparation (vectorial phase-retrieval) or a slightly longer time-evolution (2D phase-retrieval). We leave a more refined effect of error in the quantum circuit for future work, as our primary focus here is the recovery

performance of our algorithm. However, errors due to quantum noise in the circuit can often be modeled as global depolarizing noise [29]. Under this model the time series will experience exponential decay as a function of time, this will become Lorentzian line broadening on the spectrum.

We stress that the goal of phase-retrieval is solely to reconstruct the phases of the time series and we could use different metrics to quantify the quality of its recovery. Because our ultimate aim is to use the time-series to later generate a quantum phase-estimation spectrum related to a given many-body problem, we will use the recovery of the spectrum as a good metric of performance of the phase-retrieval problem. More concretely we use the 1-norm between the reconstruction and the exact solution. In order to keep the analysis simple we just implement a recovery relying on a DFT of the time-series, however, we note that the discrete Fourier transform is not necessarily the most accurate reconstruction of the spectrum of a Hamiltonian. Typically, classical post-processing techniques like the one developed in [4] can be used to obtain a more useful description of the spectrum.

Let us stress that we work in the scenario in which the available circuit depth is limited to some finite value. In that scenario, the presented circuit-depth reduction might yield the estimation of some previously unavailable many-body spectra tractable. As outlined above, this comes at some sampling cost. We do not explicitly bound this sampling cost in terms of the accuracy with which the spectrum is recovered. However, our numerical investigations suggest that the number of samples to be taken remains reasonable, even for larger system sizes.

In the rest of the document we will use dimensionless time and frequency variables, i.e., we will use the notation $f[j]$, where $j \in \mathbb{N}$ and $t_j = j\Delta t$, and $F[k]$, where $k \in \mathbb{N}$ and $\omega_k = 2\pi k/T$, and we have defined $\Delta t = T/N$. We will be interested in instances of the Fermi-Hubbard model over a line or square lattice graph G, where V is the set of vertices and E the edges, where each lattice site, i.e., vertex, has two spin modes $\sigma \in \{\uparrow, \downarrow\}$. The Hamiltonian reads

$$H = -\tau \sum_{\langle i,j|i,j\rangle \in E,\sigma} \left(a^\dagger_{i\sigma} a_{j\sigma} + a^\dagger_{j\sigma} a_{i\sigma} \right) + U \sum_{v \in V} n_{v\uparrow} n_{v\downarrow},$$

where the first term of the Hamiltonian consists of hopping terms among modes of same spin while the second sum corresponds to interactions between particles of opposite spin at the same lattice site. Throughout this work we will take $\tau = 1$ and $U = 4$, which is an intermediate coupling regime where the model exhibits non-trivial behaviour [30]. When referring to states, we are – throughout this work – referring to qubit states, which are related to the fermionic states by a Jordan-Wigner transformation. For later reference, we note that a standard basis qubit state corresponds to a position-spin-basis Slater determinant.

7.2.1. Vectorial phase retrieval

Vectorial phase retrieval [18] is a one-dimensional phase retrieval technique. One-dimensional phase retrieval techniques aim to solve a problem that, clearly, is ambiguous in general. Namely, without any further constraints, *every* assignment of phases, denoted by $\{x_j\}_{j=0}^{N-1}$, constitutes a time series $\{|f[j]|x_j\}_{j=0}^{N-1}$ that is consistent with

the absolute value measurements $\{|f[j]|\}_{j=0}^{N-1}$. Therefore, in order to (approximately) solve the one-dimensional phase retrieval problem, one has to include additional constraints or measurements. In the vectorial phase retrieval framework, this is done by including absolute value measurements of interferences between the target time series f (which we shall denote by f_1 in the remainder of this section) and another (secondary) time series f_2. To be more precise, vectorial phase-retrieval resolves the ambiguity in the retrieval of the phases of the signal f_1 by measuring not only its absolute values $|f_1[j]|$, but also those of a secondary signal $|f_2[j]|$ and their interferences $|f_1[j]+f_2[j]|$ and $|f_1[j]+if_2[j]|$. The vectorial phase retrieval problem has a *unique* solution – up to trivial ambiguities – in the noiseless scenario, provided that (1) f_1 and f_2 are *spectrally independent* (i.e., their z-transforms have no common roots in the complex plane[3]), and (2) the discrete Fourier transforms of f_1 and f_2 (denoted resp. by F_1 and F_2) have no support outside some finite interval (of size at most $N/2$ or, equivalently, of size at most π in terms of frequency) [18].

In the setting of control-free quantum phase estimation, we take $f_1[j] = \langle\Phi|\exp\big(ijH\Delta t\big)|\Phi\rangle$ and $f_2[j] = \langle\Phi|\exp\big(ijH\Delta t\big)\big|\psi\big\rangle$ (for some state $\big|\psi\big\rangle$ different from $|\Phi\rangle$), so that $f_1[j]+f_2[j] = \langle\Phi|\exp\big(ijH\Delta t\big)\big(|\Phi\rangle+\big|\psi\big\rangle\big)$ and $f_1[j]+if_2[j] = \langle\Phi|\exp\big(ijH\Delta t\big)\big(|\Phi\rangle+i\big|\psi\big\rangle\big)$. The constraint of F_1 and F_2 only being supported on a frequency interval of size π can be satisfied by an appropriate choice of Δt. Note that for standard statistical phase estimation routines, such a constraint should also be satisfied, except that there the support should be of size 2π. We stress that the aforementioned condition (1) of f_1 and f_2 being spectrally independent does *not* mean that states $|\Phi\rangle$ and $\big|\psi\big\rangle$ can only have support on disjoint subsets of eigenvalues of H, but rather that the z-transforms of the time-series have no common roots in the complex plane.

Remark that in our scenario we are interested in retrieving the phases of the time series f_1, which can then – for instance – be used to obtain F_1. Two important factors that lead to one only obtaining an approximation of the correct assignment of phases $\{x_j\}_{j=0}^{N-1}$ in a practical setting are (1) the noise on the learned signals $|f_1|$ and $|f_2|$ and (2) the fact that F_1 and F_2 fail to have their support restricted to a well-defined finite interval. This second factor is caused by spectral leakage of F_1 and F_2 away from the eigenvalues of H, which may spread outside the expected restricted support. This can be particularly aggravated when the state $|\Phi\rangle$ and distribution of supported eigenvalues might be such that F_1 and F_2 decay slowly near either end of the $[0,\pi]$ interval.

We find that to ensure that the approximate solution obtained through vectorial phase retrieval is accurate, one has to ensure that the scenario is such that F_1 and F_2 have *approximately* well-defined support. This can generally be done by for instance increasing the total evolution time (at the expense of deeper circuits implementing the time evolution) or making informed choices for the reference states or for $|\Phi\rangle$ (when the problem we aim to solve allows for some freedom in the choice of the input state). Our numerical investigations suggest that vectorial phase retrieval performs relatively well even when F_1 and F_2 do not have well-defined support.

In what follows we summarize our implementation of the vectorial phase retrieval technique (and its application to control-free quantum phase estimation), more details

[3] Since we only work with relatively modest system sizes and total evolution times (since else we would have no interest in implementing *control-free* phase estimation), we note that the continuum limit behavior of these roots is not an immediate barrier for our method.

and derivations can be found in Section 7.B.

7.2.1.1. Vectorial phase retrieval as an optimization problem

The key aspect behind the vectorial phase-retrieval technique is the realization that the phases of $f_1[j]$ and $f_2[j]$ are related by the relation $e^{\theta_1[j]} = e^{\theta_2[j]}\, G[j]$, where[4]

$$G[j] := \frac{|f_1[j] + f_2[j]|^2 + i|f_1[j] + i f_2[j]|^2 - (1+i)\big(|f_1[j]|^2 + |f_2[j]|^2\big)}{2\,|f_1[j]|\,|f_2[j]|}, \tag{7.1}$$

is fully accessible from the measurement data. This allows us to define an optimization problem over the variable $\mathbf{y} \in \mathbb{C}^{2N}$ with unit absolute value entries, encoding the estimates of the phases $\{e^{\theta_1[j]}, e^{\theta_2[j]}\}$. The cost function to be optimized consists of two contributions. The first contribution quantifies how close we are from satisfying the relative phase constraints $e^{\theta_1[j]} = e^{\theta_2[j]}\, G[j]$ and reads,

$$Q_{\text{interference}}(\mathbf{y}) := \sum_{j=0}^{N-1} \big|\, y_j - y_{N+j}\, G[j] \,\big|^2. \tag{7.2}$$

As will be argued in Section 7.B (and in [18]), imposing these phase constraints is not sufficient for (approximate) recovery of the phases in general. Therefore, the cost function also contains a second contribution. This additional contribution quantifies how close we are from F_1 and F_2 being zero outside the support $\{0,1,\ldots,\sigma\}$ (for some integer $\sigma \in \{0,1,\ldots,\lfloor N/2 \rfloor - 1\}$, the *exact* value of which is unknown in a general setting). In an ideal noiseless scenario, this contribution is minimized if F_1 and F_2 are zero on the domain $\{\sigma+1, \sigma+2, \ldots, N-1\}$. In reality, one of the factors making this constraint not fully satisfied is the inevitable spectral leakage of the discrete Fourier transform, that will lead to F_1 and F_2 having support outside of the supported frequency range. Nonetheless, as we will demonstrate, the approach still recovers the phases relatively accurately provided that $|F_1[k]|$ and $|F_2[k]|$ decay quickly for $k > \sigma$. To impose that the solution (approximately) has the correct support size, we define a family of support cost functions, one for each $s \in \{0,1,\ldots,N-1\}$ (because, again, the true support size σ might be unknown, or not even well-defined), defined as

$$Q^{(s)}_{\text{support}}(\mathbf{y}) = \sum_{k=s}^{N-1} \left| \sum_{j=0}^{N-1} |f_1[j]|\, y_j \exp\big[-i2\pi jk/N\big] \right|^2 + \sum_{k=s}^{N-1} \left| \sum_{j=0}^{N-1} |f_2[j]|\, y_{N+j} \exp\big[-i2\pi jk/N\big] \right|^2, \tag{7.3}$$

quantifying how far the spectrum is from being zero on the region $\{s, s+1, \ldots, N-1\}$.

As we explain in more detail in section 7.B, one can define a family of total cost functions, combining the support component and the interference component, which can be expressed as a non-negative quadratic form

$$Q^{(s)}(\mathbf{y}) := Q^{(s)}_{\text{support}}(\mathbf{y}) + Q_{\text{interference}}(\mathbf{y}) = \mathbf{y}^\dagger A_s^\dagger A_s\, \mathbf{y}, \tag{7.4}$$

where $A_s \in \mathbb{C}^{(2(N-s)+N)\times 2N}$ entries are detailed in Section 7.B. Remark that $Q^{(s)}(\mathbf{y})$ is equal to zero for all $s \in \{\sigma+1, \sigma+2, \ldots, N-1\}$ if $\mathbf{y}$ is the unique solution of the ideal

[4] For $G[j]$ to be well-defined, the subsets of eigenstates of H with which $|\Phi\rangle$ and $|\psi\rangle$ have significant overlap should not be disjoint, as this would cause $|f_2[j]|$ to vanish.

noiseless vectorial phase-retrieval problem (i.e., for which $|F_1[k]|$ and $|F_2[k]|$ are indeed equal to zero for $k > \sigma$). When $|F_1[k]|$ and $|F_2[k]|$ are not exactly equal to zero for $k > \sigma$, the minimum does not equal zero for any $s \in \{0, 1, \ldots, N-1\}$. Typically, however, it will decay quickly as a function of s around the optimal choice of s. As presented in Figure 7.2, finding the solution to the optimization problems at different values of s typically allows one to obtain an estimate of an approximate σ, which in turn can be used to approximately retrieve the phases of f_1.

In any case, the optimal assignment of phases $\mathbf{y}$ for a given s can be found by minimizing the function $Q^{(s)}(\mathbf{y})$ over all assignments of *phases* $\mathbf{y}$, which is not a convex optimization problem due to the constraints on the entries of $\mathbf{y}$ (i.e., that each entry has unit absolute value). To approximate the optimal solution in a practical setting, one will have to relax the non-convex optimization over phase assignments. The relaxation that we will employ is the problem of optimizing $Q^{(s)}(\mathbf{y})$ over all vectors $\mathbf{y} \in \mathbb{C}^{2N}$ for which $||\mathbf{y}||_2 = \sqrt{2N}$, instead of optimization over $\mathbf{y}$ whose entries have unit absolute value.[5] Hence, we relax the problem to the problem determining the smallest eigenvalue of a matrix.

Empirical investigation, illustrated in Figure 7.2, has shown us that in order to make the phase-retrieval more resilient to noise, it is convenient to generalize the problem by including $R > 1$ *secondary* signals and their respective interferences with the target signal (as opposed to including just one such secondary signal). This leads to an equivalent expression for the cost function, now expressed in terms of $\tilde{A}_s \in \mathbb{C}^{((R+1)(N-s)+RN)\times(R+1)N}$ (where $\tilde{\cdot}$ is used to stress that noisy input data is used in the definition of $\tilde{A}_s$) and $\mathbf{y}_R \in \mathbb{C}^{(R+1)}$. Employing an equivalent relaxation as the one just discussed, we optimize over all vectors $\mathbf{y} \in \mathbb{C}^{(R+1)N}$ for which $||\mathbf{y}||_2 = \sqrt{(R+1)N}$, corresponding to the following eigenvalue problem.

$$\min_{\mathbf{y}\in\mathbb{C}^{2N} \text{ s.t. } ||\mathbf{y}||_2=\sqrt{(R+1)N}} \tilde{Q}^{(s)}(\mathbf{y}) = (R+1)N\,\lambda_{\min}\big(\tilde{A}_s^\dagger \tilde{A}_s\big). \tag{7.5}$$

Our numerical investigations suggest that for relatively small noise magnitudes, the vector $\mathbf{y}$ obtained by solving equation (7.5) will generally be close to the ideal solution. We do note that the entries of the optimal vector $\mathbf{y}$ are not necessarily phases (i.e., have unit absolute value), nonetheless, we approximate $\{f_1[j]\}_{j=0}^{N-1}$ by $\{|f_1[j]|y_j\}_{j=0}^{N-1}$ because our numerical investigations suggest that this generally gives a more accurate reconstruction compared to the solution that is obtained when rounding to a vector whose entries are phases. For a more detailed discussion of this latter fact, we refer to Section 7.B.3. We note that the sampling effort for running the $R > 1$ version of VPR scales linearly in R. How large R needs to be to obtain an accurate reconstruction depends on the problem instance and noise magnitude.

7.2.1.2. Implementation

To demonstrate the technique, we apply the vectorial phase retrieval framework to the state evolution signals generated by a (1×5) spinful Fermi-Hubbard Hamiltonian with $|U/\tau| = 4$. The target initial state is taken to be a (standard-basis) state at half filling,

[5] Remark that the relaxation we use is slightly different to the one implemented in [18].

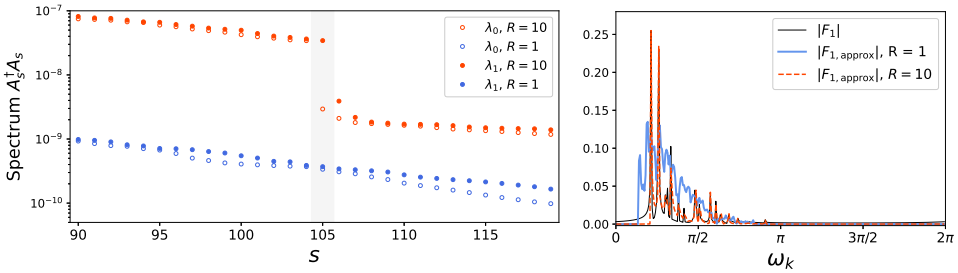

(a) $N_{\text{samples}} = 10^6$. (b) $N_{\text{samples}} = 10^6$, $s = 105$.

Figure 7.2: (a) The smallest two eigenvalues of $A_s^\dagger A_s$ for $R = 1$ (blue) and $R = 10$ (red) in the noisy setting corresponding to $N_{\text{samples}} = 10^6$ as a function of s; (b) Comparison between the reconstructed spectrum for $R = 1$ (blue) and $R = 10$ (red) versus the exact solution (black) (obtained by evaluating the smallest eigenvector of $A_s^\dagger A_s$ at $s = 105$). Note that in (a) it can be seen that, in the $R = 10$ scenario, the smallest eigenvalue drops rapidly around $s = 105$ while it is the opposite for $R = 1$. We show the spectrum on the range π to 2π, to validate there is no support between π to 2π, as required by condition (2) discussed at the start of Section 7.2.1.

where we map to qubits under a Jordan-Wigner transformation. We consider a scenario in which the number of secondary states is one (i.e., $R = 1$), and a scenario in which we include ten secondary states (i.e., $R = 10$). The *secondary* states are created as follows. Take the target state and select at most $p < n/2$ (distinct) pairs of qubits that are unequally occupied, then flip the states on these qubits. We obtain the time series at $N = 300$ points in time (with $N_{\text{samples}} = 10^6$ shots per point in time), with a time increment $\Delta t = 0.133$.

Note that superpositions of the target state and a given secondary state can be produced using similar circuit to those generating a multi-qubit GHZ state followed by local X gates, which requires $O(\log p)$ layers of two-qubit gates (given all-to-all connectivity). This is similar to a method that can be used to reduce the cost of controlled time-evolution on hardware with restricted connectivity, or the depth on hardware with all-to-all connectivity (see Table 7.1). While one should take this additional cost into account when comparing complexities, we believe that it is not a fundamental limitation of the method, as different secondary states could be used.

Figure 7.2a depicts the smallest two eigenvalues of the $\tilde{A}_s^\dagger \tilde{A}_s$ matrix, whose smallest eigenvector is used as an estimate for the assignment of phases of the time series entries in both scenarios. Our investigations suggest that:

1. The two smallest eigenvalues are a proxy for the quality of the retrieval of phases. That is, the smaller the smallest eigenvalue, the smaller the value of the two residuals Eq. (7.2) and (7.3) above.

2. A large gap opening between the two smallest eigenvalues suggests a larger overlap between the smallest eigenvector and the true assignment of phases. This being motivated by the fact that in the ideal noiseless case the solution is unique.

Typically, the smallest eigenvalue of $\tilde{A}_s^\dagger \tilde{A}_s$ drops rapidly as a function of s at a given value s^* (at $s = 105$ in Figure 7.2a), which our numerical investigations suggests is an indicator of an accurate retrieval of phases around that value of s. Note that both features mentioned above are present in the $R = 10$ scenario, leading to a good reconstruction, but not in the $R = 1$ scenario, which does not provide an accurate reconstruction, as one can see in Figure 7.2b (showing the associated estimate of the spectrum $|F_1|$ at $s = 105$).

More generally, our investigations suggest that the noise resilience of the vectorial phase retrieval algorithm is improved by taking $R > 1$. We note that increasing R increases linearly the amount of classical data to be processed and the number of quantum circuit shots required, but leads to a cubic scaling of the running time of the classical post-processing algorithm (eigenvalue problem). The trade-off between resources cost and quality of solution will therefore play a role in the final choice of R. In addition, it is important to cleverly pick the secondary quantum states in a general setting, so that the circuit generating them and their superposition with the target state is a shallow circuit, so that it does not significantly increase the depth of the quantum circuit needed for the implementation.

7.2.2. Two-dimensional phase-retrieval

It is well-known in the phase-retrieval literature that two dimensional problems are more resilient to non-trivial ambiguities of the solution [31, 32]. The non-trivial ambiguities in the 1D problem arise from the fact that the z-transform of the signal is a univariate reducible polynomial. In contrast, reducible multivariate polynomials have measure zero, which implies that the phase retrieval problem in the multi-dimensional case typically has a unique solution, unless disguising a single variable polynomial as a multivariable one. Nonetheless, this result strongly suggests that two-dimensional phase-retrieval should work in practice, an intuition confirmed by a broad range of applications in optical imaging and image processing that demonstrate its practical value [33, 34].

We embed our problem – which is naturally one dimensional – into a larger 2D problem so that we can exploit the benefits of two-dimensional phase-retrieval. The full details of this method can be found in Appendix C of Ref. [24], but in summary we introduce an additional dummy Hamiltonian H_D and a second independent time variable z. A new signal $f[j,l]$ can be defined from samples of

$$f(t,z) = \langle \psi | e^{itH} e^{izH_D} | \psi \rangle, \tag{7.6}$$

at the discrete set of times $t_j = j\Delta t$ and $z_l = l\Delta z$, where $\Delta t = T/N$ and $\Delta z = T'/M$, where for simplicity we choose $T' = T$ [6]. This lets us define a matrix $f[j,l]$ whose entries can roughly be thought of as samples of this function at these discrete times, up to a multiplication by a windowing function and a specific ordering of these elements to ensure the spectrum being real (see Appendix C of Ref. [24] for details).

In this work we apply one of the oldest and most widely used phase retrieval algorithms to this new 2D problem, the hybrid input-output algorithm (HIO) proposed by Fienup [19] in 1982. HIO is known to perform well when $F[k,m]$ – the DFT of $f[j,l]$ – is real and positive. This is because this knowledge can be used to drive the HIO algorithm. These

[6] Note that $T = T'$ is not a fundamental requirement and other choices are possible and potentially beneficial.

Algorithm 7.1: Hybrid

Input: $|f|$, β, L, F^1
for $i \leftarrow 1$ **to** L **do**
  $f^i \leftarrow \mathcal{DFT}^{-1}(F^i) = |f^i| e^{i \arg(f^i)}$
  $\tilde{f}^i \leftarrow |f| e^{i \arg(f^i)}$
  $\tilde{F}^i \leftarrow \mathrm{Re}(\mathcal{DFT}(\tilde{f}^i))$
  if $\tilde{F}^i[k,m] \leq 0$ **then**
    $F^{i+1}[k,m] \leftarrow F^i[k,m] - \beta \tilde{F}^i[k,m]$
  else
    $F^{i+1}[k,m] \leftarrow \tilde{F}^i[k,m]$

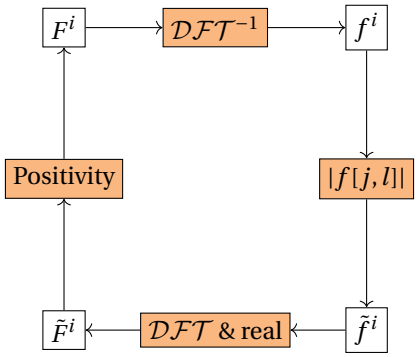

Figure 7.3: Pseudocode (left) and schematic (right) for Fienup's hybrid input-output algorithm in four steps. The i-th iteration starts with a candidate spectrum F^i: Step (1) transforms it into a time series f^i candidate via an inverse DFT; Step (2) generates a new $\tilde{f}^i$ that has same phase as f^i and satisfies $|\tilde{f}[j,l]| = |f[j,l]|$; Step (3) transforms it back to the Fourier domain and takes the real part; Step (4) implements the update rule in the schematic to impose a relaxed version of positivity of the spectrum.

constraints can be imposed on the problem very naturally, as we detail in Appendix C of Ref. [24]:

1. Positivity can be ensured by choosing a H_D which commutes with H, and by multiplying the time-series by a filter that has a positive discrete Fourier transform. For simplicity in our implementation we choose a triangular window and multiply by this to form a discrete array $f[j,l]$.

2. $F[k,l]$ is real if we can define $f[j,l]$ so that it satisfies $f[j,l] = f^*[N-j, M-l]$, as any signal that satisfies this condition has a real DFT. We achieve this by exploiting the fact that $f(t,z) = f^*(-t,-z)$ and including samples at both (t_j, z_l) and $(-t_j, -z_l)$ when defining $f[j,l]$ (see Appendix C of Ref. [24]).

We can then incorporate this additional knowledge about the solution into the phase retrieval algorithm. HIO also makes it relatively simple to add additional constraints exploiting symmetries of the problems or partial information we may have on the solution, as we show below for a specific example. In this work we have only explored the simplest windowing function which ensures the positivity constraint is satisfied. There is an extensive signal processing literature on filtering [35] and the optimum choice of windowing function remains an open question.

As schematized in Figure 7.3, Fienup's hybrid input-output algorithm (HIO) [19] involves transforming back and forth between the Fourier (left) and object domain (right) interspersed by projections that guarantee the satisfaction of the time-series values

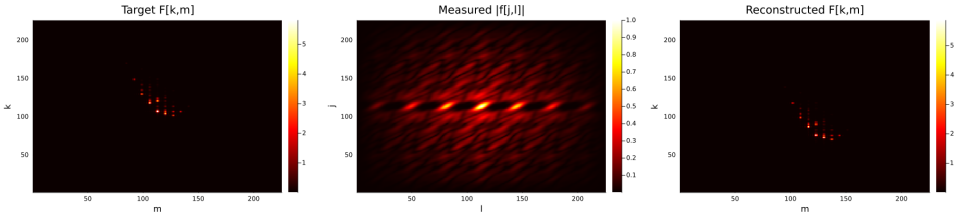

(a) Target spectrum. (b) Measured time-series. (c) Reconstructed spectrum.

Figure 7.4: These three plots show an example of 2D phase retrieval for a 2×2 Fermi-Hubbard Hamiltonian with $N \times M = 225 \times 225$. Here $T = 112$ and we take $|\psi\rangle$ to be a uniform super-position over computational basis states. H has been normalized so that its eigenvalues lie between 0 and π. We have chosen the H_D as the total number operator. This example excludes sampling noise. (7.4a) The target 2D discretized spectrum $F[k,m]$ is shown here as a heatmap; (7.4b) The absolute values $|f[j,l]|$ are shown here, these would be measured using only time-evolution and are the input to the HIO phase-retrieval algorithm; (7.4c) Here we show the result of using HIO on (7.4b), we recover this reconstruction. Observe that the recovered spectrum contains a *trivial* ambiguity, in the form of a shift.

$|f[j,l]|$ in the object domain and approximately guarantee the $F[k,m]$ constraints in the Fourier domain. The algorithm is non-linear and of an iterative nature, where the constraint on the spectrum defines a new element F^{i+1} with the update rule shown at the left side of Figure 7.3, where the parameter $0 \leq \beta \leq 1$ is carefully selected to optimize the performance. We will now summarise the details of a specific instance of this set-up for the Fermi-Hubbard model.

7.2.2.1. Implementation

We test this algorithm numerically using a 2D spin Fermi-Hubbard model, with $U = 4$ and $\tau = 1$. For the dummy Hamiltonian H_D we use the total particle number operator, which under the Jordan-Wigner encoding is simply a sum of single qubit Pauli-Z operators. We use 2×2 spin Fermi-Hubbard model and set $N = M = 225$, time $T = 112$ and choose $|\psi\rangle$ to be the uniform super-position over computational basis states. H has been normalized so that its eigenvalues lie between 0 and π. Additionally we incorporate knowledge of the phases of $f[0,l]$ as driving constraints, as they can be computed efficiently classically for our specific input state. These points correspond to evolution under H_D only, which can be performed with local gates in any potential experiment.

The first step is to use the HIO algorithm to reconstruct a 2D image, summarized in Figure 7.4. The left side of Figure 7.4 shows the target 2D-spectrum $F[k,m]$, which we obtain from a direct numerical calculation of the time-series $f[j,l]$ including its phase. In the centre we show the absolute values of $|f[j,l]|$, these would be measured using only time-evolution as sketched in Figure 7.1b, and form the input to the HIO algorithm. Observe that we also see the effect of the 2D triangular windowing function,

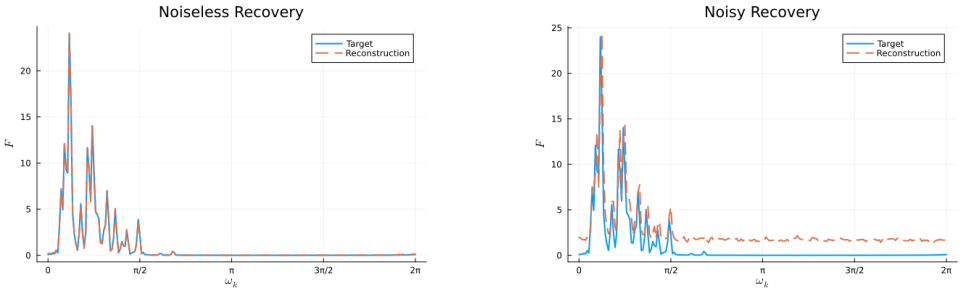

(a) Noiseless

(b) With sample noise.

Figure 7.5: This figure shows the recovered 1D spectrum following the image reconstruction in Figure 7.4. (7.5a) This shows $F[k] := \mathcal{DFT}[f[j,0]]$ where f is the inverse discrete Fourier transform of the reconstructed spectrum shown in (7.4c). This is in the noiseless scenario. Note that the ideal solution includes the triangular window by definition. (7.5a) This figure shows the effect of sample noise on this same problem. In this case we use only 1000 samples per time-series point and still recover the peaks but with a vertical offset in the reconstructed spectrum.

the magnitude of $|f[j,l]|$ decreases as we approach the edge of the plot[7]. Finally to the right we show the reconstruction of $F[k,m]$ using the the HIO algorithm with $\beta = 0.9$ and $L = 5000$. We can observe a trivial ambiguity in this reconstruction, the final spectrum is translated slightly. Nonetheless the reconstruction is correct up to trivial ambiguities which we correct for when plotting.

Finally, we recover the 1D information following this image reconstruction. To do this we apply a DFT to the recovered signal $f[j,0]$, which corresponds to the times-series of H without the dummy Hamiltonian. The result of this is shown in Figure 7.5. We also model the effect of sample noise on this algorithm and find that for a modest number of samples – 1,000 per calculation of $|f[j,l]|$ – the algorithm can still reconstruct the location of the peaks.

7.2.3. Comparative Overview of Proposed Methods

We are now ready to compare the performance of these two approaches in the context of spectral estimation of the Fermi-Hubbard model. We choose a 3×3 instance of the (spinful) Fermi-Hubbard model (with $|U/\tau| = 4$) and set up the following task. The target spectrum is the ideal solution obtained from the DFT of the ideal time series $f[j] = \langle\psi| e^{ijH\Delta t} |\psi\rangle$ for time step $\Delta t = 0.12$ and maximum time evolution of $T = 15$, where the time series is obtained using brute-force calculation of $e^{ijH\Delta t}$ (for varying j), with the input state being

[7] We have plotted the zero frequency in the centre.

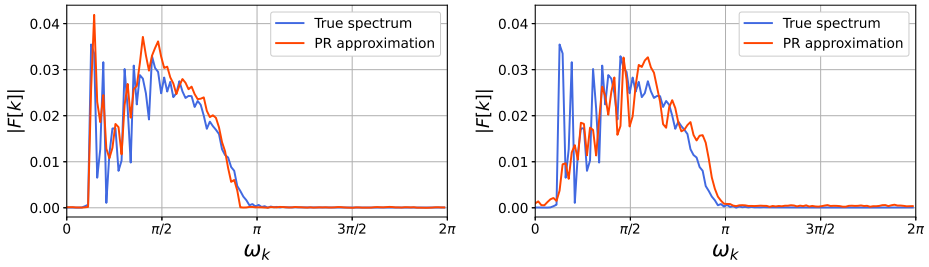

(a) Vectorial phase retrieval.

(b) 2D HIO.

Figure 7.6: This figure compares the two techniques on the same problem with the same sample budget of 5×10^9. We target the spectrum given the input state $|\psi\rangle$ in Eq. 7.7 for the Fermi-Hubbard Hamiltonian H on a lattice of size 3×3 with on-site interaction strength $U = 4$ and hopping strength $\tau = 1$, with time step $\Delta t = 0.12$ and maximum time evolution of $T = 15$. In both plots the ideal solution is obtained from $f[j] = \langle\psi| e^{ijH\Delta t} |\psi\rangle$ and shown in blue; (a) Here we show the estimated spectrum (red) for vectorial phase-retrieval with $R = 40$ and $6 \cdot 10^5$ shots per point in time j per signal (we have taken s, the support size estimate, to be 49); (b) Here we show the estimate spectrum (red) for two-dimensional phase-retrieval with $M = 25$, $\beta = 0.9$ and $L = 5000$ and $1.6 \cdot 10^6$ shots per $f[j,l]$.

$$|\psi\rangle = \frac{1}{\sqrt{3}}(|010101010\,010101010\rangle + |110101100\,010100110\rangle + |011101010\,000110111\rangle), \tag{7.7}$$

where the first (last) 9 bits are associated with the occupation of the spin-up (spin-down) modes. This state was chosen as it is suitable for our implementations of the vectorial phase retrieval algorithm and the two-dimensional phase retrieval algorithm. It does not provide an unfair advantage to either adaption and – importantly – it is not an eigenstate of H_D.

To get as close as possible to a fair comparison of both techniques we allocate a budget of $N_S = 5 \cdot 10^9$ circuit runs to our simulation of each phase retrieval technique. This directly limits the number of circuit runs required for each technique to retrieve each $|f[j]| = |\langle\psi| e^{ijH\Delta t} |\psi\rangle|$ which translates into shot noise on each value, which can later affect the performance of the phase-retrieval reconstruction. More concretely, we model the effect of shot noise induced by the finite number of samples during the statistical estimation of the time-series by sampling from a binomial distribution with success probability $|f[j]|^2$ (See Section 7.A.3).

We chose to set up a comparison in this fashion as both methods have a variety of parameters which can be altered to ensure quality performance in the presence of sample noise. By fixing only the total number of runs we can compare the performance of both methods on the same footing. For example, with vectorial phase-retrieval there is now a

trade-off between the number of additional secondary states R permitted and the total number of samples used for obtain an estimate of each $|f[j]|$. The same is true for the two-dimensional phase-retrieval implementation and the parameter M. The optimum setting for each implementation remains an open question and depends on the particular problem under consideration. In this work we tried to optimize the choices of parameters for both approaches trying to retrieve a spectrum as close to the target one as possible.

The phase-recovery is followed by a FFT to reconstruct the Fermi-Hubbard Hamiltonian on a lattice described above, depicted in Figure 7.6 for comparison.

For the 2D implementation we took $M = 25$ and used 1.6×10^6 samples per $f[j,l]$. In conjunction with $\beta = 0.9$ and $L = 5000$ iterations the HIO algorithm produces the spectrum shown in Figure 7.6. For the vectorial phase retrieval implementation, we used $R = 40$ secondary states, and $6 \cdot 10^5$ samples per signal per point in time j. Let us stress again that the *total* number of circuit runs $N_S = 5 \cdot 10^9$ is taken to be equal for both methods, to ensure a fair comparison. We observe that the 2D HIO method gives a poorer reconstruction in comparison to VPR. In particular, the HIO algorithm struggles to recover the amplitudes of the peaks. One can also check that the 1-norm (normalized by N) of the vector with entries $|F_{\text{reconstructed}}[k]| - |F_{\text{exact}}[k]|$ is 0.0018 for VPR and 0.0028 for 2DPR, confirming the intuition from a visual inspection of Figure 7.6. Another observation is that the peaks recovered using 2D HIO appear slightly broadened (which generally leads to a lower resolution on the recovered spectrum), due to the windowing function we use to enforce the positivity constraint we use to drive HIO.

The accurate retrieval of the spectrum using the vectorial phase retrieval algorithm can be attributed to the large number of secondary states employed in the algorithm, leading improved noise resilience (as is outlined in Section 7.2.1). Since the vectorial phase retrieval algorithm performs better as the support of the spectrum becomes well-defined, there might be instances of spectra with less well-defined support for which vectorial phase retrieval performs worse. We remark that when HIO is observed to perform well in the noiseless scenario – such as when the peaks are well defined as in Figure 7.5a – we observe that the corresponding noisy spectrum can typically be reconstructed well with a comparatively low number of samples as in Figure 7.5b. For the problem instance explored in Figure 7.6, HIO does not recover the spectrum as well as VPR in the noiseless setting and this gap in performance persists in the noisy setting as shown in Figure 7.6. We stress that the poorer performance of two-dimensional phase-retrieval in Figure 7.6 is not necessarily a fundamental fact and one could potentially adapt more advanced modern two-dimensional phase retrieval methods to quantum phase-estimation problem while also exploiting other commuting Hamiltonians H_D that could potentially lead to better results.

The restriction to an 18-mode Fermi-Hubbard instance is inherent to our numerical modeling demonstration, but we expect that the phase-retrieval algorithms can be applied to system sizes beyond classically tractable ones and can be applied to experimental data from actual quantum hardware. The two phase-retrieval techniques are presented in full detail in Section 7.B and in Appendix C of Ref. [24]. Finally, we believe further work could generalize the current framework to more general Loschmidt echo quantities, but also linear response of quantum systems, operator-resolved density of states or Green's functions.

7.3. Conclusion and Discussion

We have demonstrated how to design phase-retrieval algorithms for the problem of statistical phase estimation. These algorithms allow for the reconstruction of the phases of the (time evolution) time series of a quantum state $|\psi\rangle$ with respect to a given Hamiltonian H, while only having access to their absolute values. To achieve this seemingly impossible goal, one needs to address a more general problem involving the absolute values of a larger set of time series, resulting from increasing the set of input states or enlarging the range of time-evolutions available to us. To reach our goal we have adapted two well-known phase retrieval algorithms, vectorial phase retrieval and two-dimensional phase retrieval, which we demonstrate offer good performance in retrieving the spectrum of a Fermi-Hubbard model.

In the vectorial phase-retrieval scenario, we obtain absolute values of time series corresponding to time evolution of multiple input states, but – more importantly – time evolution of their superpositions. We also showed that despite quantum phase estimation being a one-dimensional signal processing problem, one can exploit the benefits of two-dimensional phase-retrieval by adding a quench evolution governed by a dummy Hamiltonian H_D commuting with H. Despite the two techniques being rather different in design and post-processing algorithms, there are some analogies in their adaptation to quantum phase estimation. One could see the time-evolution with a dummy Hamiltonian H_D as an alternative way of generating a non-trivial family of time series in addition to the one of the target input state $|\psi\rangle$, in a similar fashion as vectorial phase retrieval exploits a family of multiple input states and their superpositions. Our methods require some tailoring to the Hamiltonian at hand, such as a choice for the target and the secondary states for VPR such that the spectra have (approximately) well-defined support, or the choice of a dummy Hamiltonian for 2DPR. However, we have no reason to believe that our techniques only work for the Fermi-Hubbard model, and have merely chosen it because of its practical relevance.

We demonstrated that the performance of the algorithms is resilient to shot noise affecting the estimation of the time-series resulting from the limited number of accessible circuit runs. This is sufficient to analyses the non-trivial recovery performance of phase-retrieval. Further investigation of the subtle impact of circuit imperfections on the comparison between statistical phase-estimation and phase-retrieval is left to future work. Providing theoretical guarantees and asymptotic scaling of the different resource cost in realistic scenarios is an extremely complex task for the algorithms proposed here, specially for HIO. Nonetheless, it would be interesting to obtain heuristic asymptotic scalings, but this will need to use time-series resulting from time-evolutions that can be scaled numerically on classical computers, like free-fermions or 1D TFIM.

The fine tuning of the different phase-retrieval techniques to a specific spectrum of a given quantum state $|\psi\rangle$ and Hamiltoian H currently needs a handcrafted approach. It would be interesting to design techniques that can automatize the procedure for arbitrary states and Hamiltonians. We note that despite the Hybrid Input-Output algorithm used in two dimensional phase-retrieval working well in-practice, it is highly probable that relaxing the problem to a convex optimization would bring additional benefits, from convergence guarantees, to better running times, scaling and potentially better resilience to noise. We leave this question open for future work. Even further, one could dream

of having a unified framework that given a quantum state and Hamiltonian suggest the better phase-retrieval algorithm, ideally base on convex optimization tools that are known to be most resilient to noise.

We stress that phase retrieval algorithms recover the phases of the state-evolution time series, of which one learns just the absolute values in the scenario of control-free quantum phase estimation. In some portions of this work, for convenience, we compare the discrete Fourier transform of the retrieved time series to that of the true time series to judge the accuracy of the phase retrieval. However, we note that the discrete Fourier transform is not necessarily the best way to recover the spectrum of a Hamiltonian from its time-series. Typically, classical post-processing techniques like the one developed in [4] can be used to obtain a more accurate and useful description of the spectrum.

We observe that although our phase retrieval algorithms can achieve nontrivial reductions in quantum circuit complexity in many cases, whether this is useful in practice requires consideration of a series of tradeoff. Firstly, there is a subtle cost comparison to do between the additional cost of the controlled time-evolution in statistical phase-estimation and the additional GHZ-like state input state required for vectorial phase-retrieval or the additional time-dynamics of the dummy Hamiltonian in 2D phase-retrieval. In the particular case of vectorial phase-retrieval vs statistical phase-estimation on 2D architectures, both required to build a GHZ-resource state, the difference is that the second one needs to be as large as the system, where the first only a fraction of it, which implies a saving of resources. An interesting question we leave open is whether the GHZ resource required for vectorial phase-retrieval could be made of constant cost, not scaling with system size N. Secondly, there is a tradeoff between circuit complexity and number of shots required, since our algorithms generally increase the latter. It is an interesting question whether the increased number of shots required can be significantly reduced.

There is a wealth of advanced techniques in the classical phase retrieval literature that could be explored and evaluated to improve the techniques that we present in this work further. We believe further improvement in circuit depth, number of shots and resilience to error and recovery can be achieved by further adapting more advanced ideas in signal processing to quantum information theory.

References

[1] J. P. F. LeBlanc, A. E. Antipov, F. Becca, I. W. Bulik, G. K.-L. Chan, C.-M. Chung, Y. Deng, M. Ferrero, T. M. Henderson, C. A. Jiménez-Hoyos, E. Kozik, X.-W. Liu, A. J. Millis, N. V. Prokof'ev, M. Qin, G. E. Scuseria, H. Shi, B. V. Svistunov, L. F. Tocchio, I. S. Tupitsyn, S. R. White, S. Zhang, B.-X. Zheng, Z. Zhu and E. Gull. 'Solutions of the Two-Dimensional Hubbard Model: Benchmarks and Results from a Wide Range of Numerical Algorithms'. In: *Phys. Rev. X* 5 (4 Dec. 2015), p. 041041. DOI: 10.1103/PhysRevX.5.041041.

[2] B. Bauer, S. Bravyi, M. Motta and G. K.-L. Chan. 'Quantum Algorithms for Quantum Chemistry and Quantum Materials Science'. In: *Chemical Reviews* 120.22 (Oct. 2020), pp. 12685–12717. ISSN: 1520-6890. DOI: 10.1021/acs.chemrev.9b00829.

[3] 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.

[4] R. D. Somma. 'Quantum eigenvalue estimation via time series analysis'. In: *New Journal of Physics* 21.12 (Dec. 2019), p. 123025. DOI: 10.1088/1367-2630/ab5c60.

[5] T. E. O'Brien, B. Tarasinski and B. M. Terhal. 'Quantum phase estimation of multiple eigenvalues for small-scale (noisy) experiments'. In: *New Journal of Physics* 21.2 (Feb. 2019), p. 023022. DOI: 10.1088/1367-2630/aafb8e.

[6] L. Lin and Y. Tong. 'Heisenberg-Limited Ground-State Energy Estimation for Early Fault-Tolerant Quantum Computers'. In: *PRX Quantum* 3 (1 Feb. 2022), p. 010318. DOI: 10.1103/PRXQuantum.3.010318.

[7] K. Wan, M. Berta and E. T. Campbell. 'Randomized Quantum Algorithm for Statistical Phase Estimation'. In: *Phys. Rev. Lett.* 129 (3 July 2022), p. 030503. DOI: 10.1103/PhysRevLett.129.030503.

[8] G. Wang, D. S. França, R. Zhang, S. Zhu and P. D. Johnson. 'Quantum algorithm for ground state energy estimation using circuit depth with exponentially improved dependence on precision'. In: *Quantum* 7 (Nov. 2023), p. 1167. ISSN: 2521-327X. DOI: 10.22331/q-2023-11-06-1167.

[9] N. S. Blunt, L. Caune, R. Izsák, E. T. Campbell and N. 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.

[10] T. E. O'Brien, S. Polla, N. C. Rubin, W. J. Huggins, S. McArdle, S. Boixo, J. R. McClean and R. Babbush. 'Error Mitigation via Verified Phase Estimation'. In: *PRX Quantum* 2 (2 May 2021), p. 020317. DOI: 10.1103/PRXQuantum.2.020317.

[11] A. Baroni, J. Carlson, R. Gupta, A. 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.

[12] N. Maskara, S. Ostermann, J. Shee, M. Kalinowski, A. M. Gomez, R. A. Bravo, D. S. Wang, A. I. Krylov, N. Y. Yao, M. Head-Gordon, M. D. Lukin and S. F. Yelin. *Programmable Simulations of Molecules and Materials with Reconfigurable Quantum Processors*. 2023. arXiv: 2312.02265 [quant-ph].

[13] B. Bauer, D. Wecker, A. J. Millis, M. B. Hastings and M. Troyer. 'Hybrid Quantum-Classical Approach to Correlated Materials'. In: *Phys. Rev. X* 6 (3 Sept. 2016), p. 031045. DOI: 10.1103/PhysRevX.6.031045.

[14] H. H. S. Chan, R. Meister, M. L. Goh and B. Koczor. *Algorithmic Shadow Spectroscopy*. 2024. arXiv: 2212.11036 [quant-ph].

[15] Y. Yang, A. Christianen, M. C. Bañuls, D. S. Wild and J. I. Cirac. 'Phase-Sensitive Quantum Measurement without Controlled Operations'. In: *Phys. Rev. Lett.* 132 (22 May 2024), p. 220601. DOI: 10.1103/PhysRevLett.132.220601.

[16] A. E. Russo, K. M. Rudinger, B. C. A. Morrison and A. D. Baczewski. 'Evaluating Energy Differences on a Quantum Computer with Robust Phase Estimation'. In: *Phys. Rev. Lett.* 126 (21 May 2021), p. 210501. DOI: 10.1103/PhysRevLett.126.210501.

[17] A. Walther. 'The Question of Phase Retrieval in Optics'. In: *Optica Acta: International Journal of Optics* 10.1 (1963), pp. 41–49. DOI: 10.1080/713817747.

[18] O. Raz, N. Dudovich and B. Nadler. 'Vectorial Phase Retrieval of 1-D Signals'. In: *IEEE Transactions on Signal Processing* 61.7 (2013), pp. 1632–1643. DOI: 10.1109/TSP.2013.2239994.

[19] J. Fienup. 'Phase retrieval algorithms: a comparison'. In: *Applied optics* 21 (Aug. 1982), pp. 2758–69. DOI: 10.1364/AO.21.002758.

[20] W. J. Huggins, J. Lee, U. Baek, B. O'Gorman and K. B. Whaley. 'A non-orthogonal variational quantum eigensolver'. In: *New Journal of Physics* 22.7 (July 2020), p. 073009. DOI: 10.1088/1367-2630/ab867b.

[21] S. Lu, M. C. Bañuls and J. I. Cirac. 'Algorithms for Quantum Simulation at Finite Energies'. In: *PRX Quantum* 2 (2 May 2021), p. 020321. DOI: 10.1103/PRXQuantum.2.020321.

[22] A. Schuckert, A. Bohrdt, E. Crane and M. Knap. 'Probing finite-temperature observables in quantum simulators of spin systems with short-time dynamics'. In: *Phys. Rev. B* 107 (14 Apr. 2023), p. L140410. DOI: 10.1103/PhysRevB.107.L140410.

[23] K. Hémery, K. Ghanem, E. Crane, S. L. Campbell, J. M. Dreiling, C. Figgatt, C. Foltz, J. P. Gaebler, J. Johansen, M. Mills, S. A. Moses, J. M. Pino, A. Ransford, M. Rowe, P. Siegfried, R. P. Stutz, H. Dreyer, A. Schuckert and R. Nigmatullin. 'Measuring the Loschmidt Amplitude for Finite-Energy Properties of the Fermi-Hubbard Model on an Ion-Trap Quantum Computer'. In: *PRX Quantum* 5 (3 Aug. 2024), p. 030323. DOI: 10.1103/PRXQuantum.5.030323.

[24] L. Clinton, T. S. Cubitt, R. Garcia-Patron, A. Montanaro, S. Stanisic and M. Stroeks. 'Quantum phase estimation without controlled unitaries'. In: *PRX Quantum* (Jan. 2026), pp. –. DOI: 10.1103/7qcr-znl2.

[25] A. Oppenheim, A. Willsky and S. Nawab. *Signals & Systems*. Prentice-Hall signal processing series. Prentice-Hall International, 1997. ISBN: 978-0-13-651175-5. URL: https://books.google.co.uk/books?id=O9ZHSAAACAAJ.

[26] R. P. Millane. 'Phase retrieval in crystallography and optics'. In: *Journal of The Optical Society of America A-optics Image Science and Vision* 7.3 (1990), pp. 394–411. DOI: 10.1364/JOSAA.7.000394.

[27] J. Dainty and J. Fienup. 'Phase retrieval and image reconstruction for astronomy'. In: *Image Recovery: Theory Appl* 13 (Jan. 1987).

[28] A. Walther. 'The Question of Phase Retrieval in Optics'. In: *Journal of Modern Optics* 10 (1963), pp. 41–49. URL: https://api.semanticscholar.org/CorpusID:123494835.

[29] J. Vovrosh, K. E. Khosla, S. Greenaway, C. Self, M. S. Kim and J. Knolle. 'Simple mitigation of global depolarizing errors in quantum simulations'. In: *Physical Review E* 104.3 (Sept. 2021). ISSN: 2470-0053. DOI: 10.1103/physreve.104.035309.

[30] M. Qin, T. Schäfer, S. Andergassen, P. Corboz and E. Gull. 'The Hubbard Model: A Computational Perspective'. In: *Annual Review of Condensed Matter Physics* 13.1 (Mar. 2022), pp. 275–302. ISSN: 1947-5462. DOI: 10.1146/annurev-conmatphys-090921-033948.

[31] Y. Bruck and L. Sodin. 'On the ambiguity of the image reconstruction problem'. In: *Optics Communications* 30.3 (1979), pp. 304–308. ISSN: 0030-4018. DOI: https://doi.org/10.1016/0030-4018(79)90358-4.

[32] D. Kogan, Y. C. Eldar and D. Oron. 'On the 2D Phase Retrieval Problem'. In: *IEEE Transactions on Signal Processing* 65.4 (Feb. 2017), pp. 1058–1067. ISSN: 1941-0476. DOI: 10.1109/tsp.2016.2631455.

[33] T. Bendory, R. Beinert and Y. C. Eldar. 'Fourier Phase Retrieval: Uniqueness and Algorithms'. In: *Compressed Sensing and its Applications: Second International MATHEON Conference 2015*. Ed. by H. Boche, G. Caire, R. Calderbank, M. März, G. Kutyniok and R. Mathar. Cham: Springer International Publishing, 2017, pp. 55–91. ISBN: 978-3-319-69802-1. DOI: 10.1007/978-3-319-69802-1_2.

[34] E. Osherovich. 'Numerical methods for phase retrieval'. In: *ArXiv* abs/1203.4756 (2012). URL: https://api.semanticscholar.org/CorpusID:14566173.

[35] J. O. Smith. *Mathematics of the Discrete Fourier Transform (DFT)*. W3K Publishing, 2007. ISBN: 978-0-9745607-4-8. URL: http://www.w3k.org/books/.

[36] A. Walther. 'The Question of Phase Retrieval in Optics'. In: *Journal of Modern Optics* 10 (1963), pp. 41–49. URL: https://api.semanticscholar.org/CorpusID:123494835.

[37] D. P. Arovas, E. Berg, S. A. Kivelson and S. Raghu. 'The Hubbard Model'. In: *Annual Review of Condensed Matter Physics* 13.Volume 13, 2022 (2022), pp. 239–274. ISSN: 1947-5462. DOI: https://doi.org/10.1146/annurev-conmatphys-031620-102024.

[38] R. Beinert. 'One-Dimensional Phase Retrieval with Additional Interference Intensity Measurements'. In: *Results in Mathematics* 72.1 (2017), pp. 1–24. DOI: 10.1007/s00025-016-0633-9.

[39] C. Derby, J. Klassen, J. Bausch and T. Cubitt. 'Compact fermion to qubit mappings'. In: *Phys. Rev. B* 104 (3 July 2021), p. 035118. DOI: 10.1103/PhysRevB.104.035118.

[40] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.

7.A. Technical Preliminaries

In this appendix we first introduce basics concepts of statistical phase-estimation and signal processing that will be of particular interest to readers not familiar with

phase-estimation and its signal-processing material. Secondly, we give details on the Fermi-Hubbard Hamiltonian that we use as a benchmark and few technicalities related to our simulations.

7.A.1. Definitions

In what follows we adopt the following continuous Fourier transform convention

$$F(\omega) = \mathcal{FT}[f] := \int_{-\infty}^{\infty} f(t) e^{-i\omega t} dt \qquad f(t) = \mathcal{FT}^{-1}[F] := \frac{1}{2\pi} \int_{-\infty}^{\infty} F(\omega) e^{i\omega t} d\omega, \quad (7.8)$$

and the discrete Fourier Transform definition as

$$F[k] = \mathcal{DFT}[f] := \sum_{j=0}^{N-1} f[j] e^{-i2\pi \frac{kj}{N}} \qquad f[j] = \mathcal{DFT}^{-1}[F] := \frac{1}{N} \sum_{k=0}^{N-1} F[k] e^{i2\pi \frac{kj}{N}}. \quad (7.9)$$

For a spectrum $F(\omega)$ to be a real function it needs to satisfy $F(\omega) = F^*(\omega)$ for all ω. It is easy to check from the definitions above that this is only possible if

$$F(\omega) - F^*(\omega) = \int \left[f(t) - f^*(-t) \right] e^{i\omega t} dt = 0, \quad (7.10)$$

which holds if and only if $f(t) = f^*(-t)$.

Let's consider an Hamiltonian $H = \sum_i E_i \Pi_i$, where E_i is the i-th eigenvalue of the Hamiltonian and Π_i is the projector into its eigenspace. In the case of a rank 1 projector we have $\Pi_i = |E_i\rangle \langle E_i|$. Its time evolution reads $e^{iHt} = \sum_i e^{iE_i t} \Pi_i$ and $\mathrm{Tr}[\rho \Pi_i] = p_i$ is the amplitude of spectrum of ρ at energy E_i, where the full spectrum of ρ can be written in a compact form as $F(\omega) = 2\pi \sum_i p_i \delta(\omega - E_i)$, where the positivity of p_i implies that of the spectrum $F(\omega)$ The Hadamard test in Figure 7.1a used in statistical phase-estimation allows the computation of the complex amplitude

$$f(t) = \mathrm{Tr}[\rho e^{iHt}], \quad (7.11)$$

for different times t, i.e., time series. It is easy to see that $f(t) = \mathrm{Tr}[\rho e^{iHt}] = f^*(-t)$, due to ρ being hermitian and the cyclic property of the trace, which proves that the spectrum is real. We will now show that it is also positive in the ideal case.

7.A.1.1. Positivity of the spectrum

Applying the Fourier transform we obtain

$$\mathcal{FT}[f(t)] = \int \mathrm{Tr}[\rho \sum_i e^{iE_i t} \Pi_i] e^{-i\omega t} dt,$$

which using linearity of the trace and some rearrangements leads to

$$\mathcal{FT}[f(t)] = \sum_i \mathrm{Tr}[\rho \Pi_i] \int e^{-i(\omega - E_i)t} dt$$

which using the definition of the amplitudes of the spectrum and $\int e^{-i(\omega-E_i)t}dt = 2\pi\delta(\omega-E_i)$ leads to

$$\mathcal{FT}[f(t)] = 2\pi\sum_i p_i\delta(\omega-E_i) = F(\omega).$$

In what follows we assume the Hamiltonian eigenvalues E_i satisfy the condition $E_i \in [0,2\pi]$. Otherwise, assuming the Hamiltonian maximum and minimum eigenvalue energies, or providing a decent upper and lower bounds, one can always re-scale the Hamiltonian accordingly.

7.A.2. Basics concepts on signal processing and Fourier transforms

7.A.2.1. Time-window and resolution

In practice continuous signals $f(t)$ can only be probed inside a finite time window $[-T/2, T/2]$. The simplest windowing $\hat{f}(t) = f(t)\Pi_T(t)$ results from the multiplication of the original signal $f(t)$ and a time-window function

$$\Pi_T(t) = \begin{cases} 1, & |t| \leq T/2 \\ 0, & |t| > T/2 \end{cases} \tag{7.12}$$

Using the convolution theorem, relating the Fourier transform of a product to the convolution of its respective Fourier transforms ($\mathcal{FT}[A\cdot B] = \mathcal{FT}[A] * \mathcal{FT}[B]$), together with the Fourier transforms of the step function

$$\mathcal{FT}[\Pi_T(t)](\omega) = \frac{2\sin\left(\frac{T\omega}{2}\right)}{\omega} = T\,\mathrm{sinc}(T\omega/2) \tag{7.13}$$

where we define $\mathcal{FT}[f(t)](\omega) = F(\omega)$, leads to

$$\mathcal{FT}[\hat{f}(t)](\omega) = F(\omega) * T\,\mathrm{sinc}(T\omega/2). \tag{7.14}$$

We will use the notation $\hat{F}(\omega)$ for the convolution of $= F[\omega] * \mathrm{sinc}(T\omega)$, resulting from the broadening of the spectral peacks due to the finite-size time window.

Resolution The convolution $\hat{F}[\omega] = F[\omega] * \mathrm{sinc}(T\omega/2)$ generally leads to a broadening of the spectral peaks of $F(\omega)$. The main lobe width of $\mathrm{sinc}(T\omega/2)$ being $4\pi/T$, it is easy to see that a time window of size T can generally only resolve energy gaps of order $O(1/T)$. Therefore to have full resolution of the spectrum one needs to do time-evolution up to T inverse proportional to the smallest gap between eigenvalues in the spectrum.

The choice of T gives an upper-bound on how good the quality is of the recovered spectrum, i.e., even under ideal reconstruction conditions one can at best retrieve $= F[\omega] * \mathrm{sinc}(T\omega)$. For Hamiltonians with gaps scaling beyond $1/\mathrm{poly}(n)$, it is impossible to reconstruct the exact spectrum and retrieving the coarser version $\hat{F}[\omega]$ will be the only option.

Alternative time-windows Remark that despite the ideal spectrum $F(\omega)$ being positive the convolution with a sinc function, resulting from the finite-size time window, may lead to a spectrum that has some slightly negativities. In our 2D phase-retrieval algorithm this could be an issue for the convergence of our algorithm. We will therefore use a different time window that will ensure the target spectrum in our reconstruction being all positive, see Appendix C of Ref. [24] for more details. The simplest way to achieve this is via a triangle time-window, resulting from the convolution of two rectangular time-windows $\Pi_T(t)$, which has a squared sinc function as Fourier transform. Unfortunately, this doubles the width of the main lobe of its Fourier spectrum further widening the recovered spectrum. There is a whole literature on optimized filters that would allow to limit the widening of the spectrum while still guaranteeing its positivity, we leave that for future work.

7.A.2.2. Sampling and aliasing

In practice continuous signals $f(t)$ can only be probed inside a finite time window $[-T/2, T/2]$ and for a finite number N of time values. We will assume these samples are evenly spaced with time Δt. Note that in this discussion we always choose an odd number of samples N and set the spacing to be $\Delta t = T/N$. We have odd time values as $t = 0$ is also included and we are probing up to time $t_{N-1} = (N-1)/N \times T/2$. The discretised sampling over a finite range is modelled via the multiplication of the initial signal $\hat{f}(t)$ by a Dirac comb

$$\text{Ш}_{\Delta t}(t) = \sum_{j\in\mathbb{Z}} \delta(t - j\Delta t) \tag{7.15}$$

leading to the new discretised and finite signal

$$g(t) = f(t)\Pi_T(t)\text{Ш}_{\Delta t}(t). \tag{7.16}$$

Using the convolution theorem together with the Fourier transforms of the Dirac comb

$$\mathcal{FT}[\text{Ш}_{\Delta t}(t)](\omega) = \frac{2\pi}{\Delta t}\text{Ш}_{2\pi/\Delta t}(\omega) = \frac{2\pi}{\Delta t}\sum_{k\in\mathbb{Z}} \delta\left(\omega - k\frac{2\pi}{\Delta t}\right), \tag{7.17}$$

together with eq.(7.14) we obtain

$$G(\omega) \coloneqq \mathcal{FT}[g(t)](\omega) = 2\pi\frac{T}{\Delta t}\sum_{k\in\mathbb{Z}} F(\omega) * \text{sinc}(T\omega/2) * \delta\left(\omega - k\frac{2\pi}{\Delta t}\right). \tag{7.18}$$

One can simplify the previous equation using the properties of delta functions and their convolutions and using the definition $\hat{F}(\omega) = F(\omega) * T\text{sinc}(T\omega/2)$ to get

$$G(\omega) = \frac{2\pi}{\Delta t}\sum_{k\in\mathbb{Z}} \hat{F}\left(\omega - k\frac{2\pi}{\Delta t}\right). \tag{7.19}$$

Aliasing Due to the discrete-time nature of the sampling, the spectrum $\mathcal{FT}[g(t)](\omega)$ becomes periodic function of period $2\pi/\Delta t$. The celebrated Nyquist-Shannon theorem states that sampling at twice the bandwidth of the signal guarantees an unique recovery of the continuous function. If one samples at a lower rate, the weights of higher frequency

lines are shifted to the lower part of the spectrum, an effect called aliasing. In an ideal scenario (and unrealistic) scenario of an infinite time-window and the spectrum of H having largest absolute value frequency $\omega_{\max} = \max\{|\omega|\}$, it would be sufficient to guarantee the condition $2\Delta t \cdot \omega_{\max} < 1$ to have perfect reconstruction and zero aliasing. In practice, depending on the choice of time-window and the finite value of T we always observe some amount of aliasing, but a good choice of sampling rate, time-window shape and having access to large enough T will help make aliasing small or even negligible.

7.A.2.3. Connecting sampled time-series and discretized spectrum via a DFT

The algorithms we use exploit discretized version of time-series and spectrum connected via a DFT. In what follows we review this connection where the discretized spectrum can be understood from a periodic extension of the time-series.

We *choose* as set of discrete frequencies $\omega_k = 2\pi k/T$ where $k = 0, 1, \dots N-1$ and define a length N vector $F[k] := \mathcal{FT}[g(t)](\omega_k) = G(\omega_k)$, whose entries are samples of the continuous time Fourier transform of $g(t) = f(t)\Pi_T(t)\text{Ш}_{\Delta t}(t)$ at points ω_k. We begin by writing the Fourier transform of our sampled function as

$$F[k] := \mathcal{FT}[g(t)](\omega_k) \tag{7.20}$$

$$= \int_{-\infty}^{\infty} f(t)\Pi_T(t)\text{Ш}_{\Delta t}(t)e^{-i\omega t}dt, \tag{7.21}$$

where using the definition of the Dirac comb in eq.(7.15) and $\int f(t)\delta(x-a)dt = f(a)$ we obtain

$$F[k] = \sum_{j=-\infty}^{\infty} f(j\Delta t)\cdot\Pi(j\Delta t)e^{-i\omega_k j\Delta t} \tag{7.22}$$

which using the discretisation notation $\hat{f}[j] := f(j\Delta t)\Pi(j\Delta t)$ and the fact that $\omega_k j\Delta t = 2\pi jk/N$ leads to

$$F[k] = \sum_{j=-\infty}^{\infty} \hat{f}[j]e^{-2i\pi kj/N}. \tag{7.23}$$

The connection between the DFT and samples of the the continuous functions at ω_k and $t_j = j\Delta t$ is almost complete. We have obtained the conventional DFT frequencies, however we have an infinite sum in the above expression.

We deal with this as follows. One can replace the sum over integers by the following double sum

$$F[k] = \sum_{j'\in\mathbb{Z}}\sum_{j=0}^{N-1} \hat{f}[j-j'N]e^{-i(2\pi kj/N+2\pi j'k)} \tag{7.24}$$

$$= \sum_{j=0}^{N-1}\left(\sum_{j'\in\mathbb{Z}} \hat{f}[j-j'N]\right)e^{-i2\pi kj/N}, \tag{7.25}$$

$$\tag{7.26}$$

then using the definition

$$f[j] := \sum_{j'\in\mathbb{Z}} \hat{f}[j-j'N]. \tag{7.27}$$

leads to

$$F[k] = \sum_{j=0}^{N-1} f[j] e^{-i2\pi kj/N}. \tag{7.28}$$

This shows that $F[k]$ is actually the DFT of $f[j]$, which is simply an N periodic version of the windowed samples $\hat{f}[j]$.

Notice that $f[j]$ is a periodic sequence where the values within the window are repeated with period N. The previous exposition provides an interpretation of the discretized time-series linked to a discrete spectrum via the DFT, in a similar fashion as occurred for their continuous counterparts. Provided we properly choose T and Δt to avoid aliasing and limit leakage, we can safely work on a discrete setting using the DFT to move between the sampled time-series and a discretized version of the spectrum.

7.A.3. Fermi-Hubbard simulation and technicalities of our numerical experiments

In this manuscript, we benchmark our phase-retrieval algorithms for performing statistical QPE on spectrum recovery of the Fermi-Hubbard model.

Fermi Hubbard We will be interested in instances of the Fermi-Hubbard model over a square lattice graph G, where V is the set of vertices and E the edges, where each lattice site, i.e., vertex, has two spin modes $\sigma \in \{\uparrow, \downarrow\}$. The Hamiltonian reads

$$H = -\tau \sum_{\langle i,j|i,j\rangle \in E,\sigma} \left(a^\dagger_{i\sigma} a_{j\sigma} + a^\dagger_{j\sigma} a_{i\sigma} \right) + U \sum_{v \in V} n_{v\uparrow} n_{v\downarrow},$$

where the first term of the Hamiltonian consists of hopping terms among modes of same spin while the second sum corresponds to interactions between particles of opposite spin at the same lattice site. This model is of great interest, since despite its simple form it demonstrates interesting phenomena found in more complex materials and molecular Hamiltonians [30, 36, 37]. Throughout this work we will take $\tau = 1$ and $U = 4$, which s an intermediate coupling regime where the model exhibits non-trivial behaviour [30]. All simulations in this work are performed on the qubit level, by applying a Jordan-Wigner transformation to the fermionic model.

Time-evolution simulation In order to compute time-series of the form $|\langle\psi| e^{iHt} |\phi\rangle|^2$ (forming the input to our phase retrieval algorithms) and time-series $\langle\psi| e^{iHt} |\phi\rangle$ (which we use to compare to our phase-retrieved solution), one needs to simulate the time-evolution of the Fermi-Hubbard model e^{iHt}. We will be using numerical simulations to obtain the time-series, by brute-force calculation of e^{iHt_j} with $t_j = j\Delta t$. Then, for some states $|\psi\rangle$ and $|\phi\rangle$, we simply evaluate the expressions $f[j] = \langle\psi| e^{iHt_j} |\phi\rangle$ and $|f[j]|$. Due to the current limitation on near-term quantum computers, we will be interested to explore states that have shallow preparation circuits, as detailed in Sections 7.B and **??**.

Sampling Noise A crucial aspect of the practical implementation of phase-retrieval algorithms is their resilience to noise. We model the sampling noise on time series $|\langle\phi|\exp(ij\Delta tH)|\psi\rangle|^2$ (for some $|\psi\rangle$ and $|\phi\rangle$) by sampling from a binomial random variable. In an experimental setting, one can estimate $|\langle 0^n|U_\phi^\dagger \exp(ij\Delta tH)U_\psi|0^n\rangle|^2$ (where U_ψ and U_ϕ are state preparation unitaries s.t. $|\psi\rangle = U_\psi|0\rangle$ and $|\phi\rangle = U_\phi|0\rangle$) by acting with $U_\psi^\dagger \exp(ij\Delta tH)U_\psi$ on $|0^n\rangle$, and simply performing standard-basis projective measurements on each of the n qubits. The estimate is then given by the fraction of the number of times that output $|0\rangle^{\otimes n}$ is obtained and the total number of circuit executions M. Therefore, we model the sampling noise by estimating $|\langle\phi|\exp(ij\Delta tH)|\psi\rangle|^2$ by the number of successes M_0 (divided by M) obtained in a binomial random variable realization with success probability $|\langle\phi|\exp(ij\Delta tH)|\psi\rangle|^2$ and M trials.

Trivial Ambiguities We call operations that leave the absolute value of the time series unchanged trivial ambiguities. Throughout this chapter, we ignore these trivial ambiguities and therefore match the recovered signal to the true time-series $f \in \mathbb{C}^N$ up to: (1) a global phase $e^{i\phi}f$; (2) linear phase shifts $\{e^{-i2\pi jm/N}f[j]\}$; (3) complex conjugation f^*, or combinations thereof. These respectively induce the following operations on the DFT of f: (1) multiplication by a global phase; (2) translation by m; (2 combined with 3) reflection and complex conjugation. Note that a shift in the energies, provided that the support of the spectrum is bounded, can be ensured to not lead to confusion in the ordering of the recovered peaks. The shape of the spectrum will be correct up to a constant shift. The final ambiguity can reflect the spectrum, an effect that can be mitigated if handled with proper care.

7.B. Vectorial phase retrieval

In this section we give a detailed presentation of vectorial phase retrieval (v-PR) (see [18]) and its adaptation to the quantum scenario, which we discussed briefly in Section 7.2.1. This problem corresponds to a version of the one-dimensional phase retrieval problem, where one assumes access to particular additional input in the form of interference measurements.

7.B.1. Vectorial phase retrieval with a single interference signal

Let us first discuss under which circumstances the (single-interference-signal) v-PR problem has a unique solution. This understanding of the uniqueness of the v-PR problem was used in [18] to develop a method for obtaining estimates of the phases of the time series $\mathbf{f} \in \mathbb{C}^N$.

In the v-PR framework, the (non-trivial) ambiguity of the one-dimensional phase retrieval problem is reduced by not only obtaining absolute value measurements $\{|f_1[j]|\}$ of a time series $\mathbf{f}_1 \in \mathbb{C}^N$, but also measurements $\{|f_2(j)|\}$ of some other time series $\mathbf{f}_2 \in \mathbb{C}^N$, and – crucially – of the absolute values of the time series corresponding to sums of $\mathbf{f}_1$ and $\mathbf{f}_2$. It was shown in [38] that under some relatively mild conditions on the time series $\mathbf{f}_1$ and $\mathbf{f}_2$, and assuming *exact* access to the aforementioned absolute value measurements,

the resulting phase retrieval problem has a unique solution. The problem of *finding* this unique solution efficiently was addressed in [18].

Let us define the vectorial phase retrieval problem more exactly. Afterwards, we will discuss under which circumstances the problem has a unique solution, and how one can in practice obtain an accurate estimate of the full time series using the vectorial phase retrieval algorithm [18].

Problem 7.B.1. [Vectorial phase retrieval (v-PR)] Given measurements of the absolute values

$$|f_1[j]|^2, |f_2[j]|^2, |f_3[j]|^2 := |f_1[j] + f_2[j]|^2, |f_4[j]|^2 := |f_1[j] + i f_2[j]|^2, \tag{7.29}$$

at times $j = 0, 1, \ldots, N-1$, determine

$$F_1[k] := \sum_{j=0}^{N-1} f_1[j] \exp\big[-i2\pi jk/N\big], F_2[k] := \sum_{j=0}^{N-1} f_2[j] \exp\big[-i2\pi jk/N\big], \tag{7.30}$$

at $k = 0, 1, \ldots, N-1$.

In the setting of control-free quantum phase estimation, the input signals take the following form.

$$\begin{aligned} f_1[j] &= \langle\Phi| \exp\big(i\Delta t H j\big) |\Phi\rangle, \\ f_2[j] &= \langle\Phi| \exp\big(i\Delta t H j\big) \big|\psi\big\rangle, \\ f_3[j] &= \langle\Phi| \exp\big(i\Delta t H j\big) \big(|\Phi\rangle + \big|\psi\big\rangle\big), \\ f_4[j] &= \langle\Phi| \exp\big(i\Delta t H j\big) \big(|\Phi\rangle + i\big|\psi\big\rangle\big), \end{aligned} \tag{7.31}$$

for our states $|\Phi\rangle$ of interest and a secondary state $\big|\psi\big\rangle$. We will come back to which choices of $\big|\psi\big\rangle$ are appropriate, but we note that preparing their superposition ideally should not require execution of very deep circuits.

7.B.1.1. The noiseless scenario

Let us first address the setting in which one has access to the *noiseless* absolute value measurements as input to the v-PR problem. At the core of the vectorial phase retrieval algorithm is the understanding of when Problem 7.B.1 has a unique solution. To illustrate when uniqueness is guaranteed, let us introduce the following notions.

The DFTs $\mathbf{F}_1 \in \mathbb{C}^N$ and $\mathbf{F}_2 \in \mathbb{C}^N$ of the time-series $\mathbf{f}_1$ and $\mathbf{f}_2$ are spectrally independent if $P_1(z) := \sum_{k=0}^{N-1} F_1[k]\, z^k$ (with $z \in \mathbb{C}$) and $P_2(z) := \sum_{k=0}^{N-1} F_2[k]\, z^k$ (with $z \in \mathbb{C}$) have no common roots in the complex plane.

The DFTs $\mathbf{F}_1 \in \mathbb{C}^N$ and $\mathbf{F}_2 \in \mathbb{C}^N$ of the time-series $\mathbf{f}_1$ and $\mathbf{f}_2$ have well-defined support if they are such that there exist $k_{\min} < k_{\max} \in \{0, 1, \ldots, N-1\}$ s.t. $k_{\max} - k_{\min} = \sigma \leq \lfloor N/2 \rfloor$ and $|F_1[k]|, |F_2[k]| \neq 0$ for $k_{\min} \leq k \leq k_{\max}$ and $|F_1[k]| = |F_2[k]| = 0$ for $k < k_{\min}$ and $k > k_{\max}$. In other words, they have well-defined support if they are non-zero inside of some interval $(k_{\min}, \ldots, k_{\max})$ (of length $\sigma \leq \lfloor N/2 \rfloor$), and zero outside of that interval.

Especially the latter notion will play a central role in the vectorial phase retrieval algorithm, and in particular in our application of the algorithm to spectral estimation of Hamiltonians. The uniqueness of the solution of the vectorial phase retrieval problem is related to these notions in the following way.

The vectorial phase retrieval problem has a unique solution if and only if $\mathbf{F}_1$ *and* $\mathbf{F}_2$ *are spectrally independent and have well-defined support [18].*

See [18] for a proof of this statement. Note that Problem 7.B.1 has trivial ambiguities. The trivial ambiguities correspond in general to all signals related to $(\mathbf{f}_1, \mathbf{f}_2)$ by

$$\begin{aligned}
&\big\{\exp(i\alpha)\, F_1[k], \exp(i\alpha)\, F_2[k]\big\}_{k=0}^{N-1} \text{ (with } \alpha \in \mathbb{R}\text{)},\\
&\big\{F_1[k-k_0], F_2[k-k_0]\big\}_{k=0}^{N-1} \text{ (with } k_0 \in \mathbb{Z}\text{)},\\
&\big\{F_1^*[-k], F_2^*[-k]\big\}_{k=0}^{N-1},
\end{aligned} \tag{7.32}$$

or combinations thereof. Solutions to the vectorial phase retrieval problem will only ever be unique up to these trivial ambiguities.

The vectorial phase retrieval algorithm (presented in [18]) is based on the fact that the unique solution to Problem 7.B.1 can be obtained by solving a convex optimization problem in the ideal noiseless scenario. The optimization procedure optimizes over assignments of phases to the time series measurements $\{|f_1[j]|\}$ and $\{|f_2[j]|\}$ such that they are consistent with the interference measurements and such that $\mathbf{F}_1$ and $\mathbf{F}_2$ have the right support size.

The *input* to the v-PR algorithm consists of $|f_1[j]|^2$, $|f_2[j]|^2$, $|f_3[j]|^2$ and $|f_4[j]|^2$ (at $j = 0, 1, \ldots, N-1$), and of an estimate s of the support size σ of $\mathbf{F}_1$ and $\mathbf{F}_2$. Its *output* is given by a vector $\mathbf{y} \in \mathbb{C}^{2N}$, whose entries are the estimates of the phases of $f_1[j]$ and $f_2[j]$ at $j = 0, 1, \ldots, N-1$. The algorithm is an optimization procedure that minimizes a cost function over possible assignments of phases $\mathbf{y} \in \mathbb{C}^{2N}$. Let us denote the vector containing the correct phases of $\mathbf{f}_1$ and $\mathbf{f}_2$ at $j = 0, 1, \ldots, N-1$ by $\mathbf{x}$. The cost function has two components; one component ensures that the assignment of phases is consistent with the interference measurements, and the other component ensures that $\mathbf{F}_1$ and $\mathbf{F}_2$ are indeed zero outside of the interval of size s (which is an estimate of σ). These components are denoted by $Q^{(s)}_{\text{support}}(\mathbf{y})$ and $Q_{\text{interference}}(\mathbf{y})$, respectively. Note that due to the trivial ambiguities, minimization of $|F_1[k]|$ and $|F_2[k]|$ outside of their supported interval can be realized by simply minimizing $|F_1[k]|$ and $|F_2[k]|$ at $k = s, s+1, \ldots, N-1$ without loss of generality.

The expression for the support component of the cost function is as follows.

$$\begin{aligned}
Q^{(s)}_{\text{support}}(\mathbf{y}) &:= \sum_{k=s}^{N-1} \big|F_1[k]\big|^2 + \sum_{k=s}^{N-1} \big|F_2[k]\big|^2\\
&= \sum_{k=s}^{N-1} \Big| \sum_{j=0}^{N-1} |f_1[j]|\, y_j \exp\big[-i2\pi jk/N\big] \Big|^2\\
&\qquad + \sum_{k=s}^{N-1} \Big| \sum_{j=0}^{N-1} |f_2[j]|\, y_{N+j} \exp\big[-i2\pi jk/N\big] \Big|^2.
\end{aligned} \tag{7.33}$$

The expression for the interference component of the cost function is a slightly more involved. Because of the following relations derived from the definitions of $\mathbf{f}_3$ and $\mathbf{f}_4$ and involving the correct phases $\mathbf{x}$,

$$\begin{aligned}|f_3[j]|^2 &= |f_1[j]+f_2[j]|^2 = |f_1[j]|^2+|f_2[j]|^2+2|f_1[j]||f_2[j]|\,\mathrm{Re}\big(x_j x_{N+j}^*\big),\\ |f_4[j]|^2 &= |f_1[j]+i f_2[j]|^2 = |f_1[j]|^2+|f_2[j]|^2+2|f_1[j]||f_2[j]|\,\mathrm{Im}\big(x_j x_{N+j}^*\big),\end{aligned} \tag{7.34}$$

we have that

$$x_j = x_{N+j}\,G[j], \quad G[j] := \frac{|f_3[j]|^2 + i|f_4[j]|^2 - (1+i)\big(|f_1[j]|^2+|f_2[j]|^2\big)}{2|f_1[j]||f_2[j]|}. \tag{7.35}$$

We note that $|G[j]|=1$ for $j=0,1,\ldots,N-1$. The second component of the cost function is expressed as follows, and is equal to zero if $\mathbf{y}=\mathbf{x}$.

$$Q_{\text{interference}}(\mathbf{y}) := \sum_{j=0}^{N-1} \big|\, y_j - y_{N+j}\,G[j]\,\big|^2. \tag{7.36}$$

The total cost function is defined simply as the sum of the support component and the interference component. We note that this cost function is non-negative by definition.

$$Q^{(s)}(\mathbf{y}) := Q^{(s)}_{\text{support}}(\mathbf{y}) + Q_{\text{interference}}(\mathbf{y}). \tag{7.37}$$

Importantly, it can be expressed as

$$Q^{(s)}(\mathbf{y}) := \mathbf{y}^\dagger A_s^\dagger A_s\,\mathbf{y}, \tag{7.38}$$

where $A_s \in \mathbb{C}^{(2(N-s)+N)\times 2N}$ has entries

$$(A_s)_{p,q} = \begin{cases} |f_1[q]|\exp\big(i2\pi q(s+p)/N\big), & \text{if } q<N,\ p<N-s,\\ |f_2[q]|\exp\big(i2\pi(q-N)(s+p-(N-s))/N\big), & \text{if } N\le q<2N,\ (N-s)\le p<2(N-s),\\ 1, & \text{if } q<N,\ p=q+2(N-s),\\ -G[q-N], & \text{if } N\le q<2N,\ p=q-N+2(N-s),\\ 0, & \text{otherwise,}\end{cases} \tag{7.39}$$

with $p=0,1,\ldots,(2(N-s)+N)-1$ and $q=0,1,\ldots,2N-1$.

If there exists an assignment of phases $\mathbf{y}$ such that $Q^{(s)}(\mathbf{y})=0$, then we are guaranteed that for that assignment $F_1[k]=F_2[k]=0$ at $k=s,s+1,\ldots,N-1$, and $y_j = y_{N+j}\,G[j]$ at $j=0,1,\ldots,N-1$. The correct assignment of phases $\mathbf{x}$ is a minimizer of the cost function $Q^{(\sigma)}(\mathbf{y})$. In particular, since the correct assignment of phases attains zero cost on both cost functions (by definition), we have that

$$\phi_\sigma^{\text{noiseless}} := \min_{\substack{\mathbf{y}\in\mathbb{C}^{2N}\text{ s.t.}\\ |y_0|=|y_1|=\ldots=|y_{N-1}|=1}} Q^{(\sigma)}(\mathbf{y}) = Q^{(\sigma)}(\mathbf{x}) = 0, \tag{7.40}$$

where the optimization is over all vectors $\mathbf{y}\in\mathbb{C}^N$ whose entries are phase factors.

We are not, however, a priori guaranteed that $\mathbf{x}$ is the unique minimizer of $Q^{(\sigma)}(\mathbf{y})$, and therefore that solving equation (7.40) provides the correct assignment of phases $\mathbf{x}$. However, it was shown in [18] that $Q^{(\sigma)}(\mathbf{y}) = 0$ for *any* vector $\mathbf{y} \in \mathbb{C}^{2N}$ (with non-zero 2-norm) iff $\mathbf{y}$ is the exact assignment $\mathbf{x}$. From this fact we conclude that the correct assignment of phases $\mathbf{x}$ can also be obtained by solving the following relaxation of the non-convex optimization problem in equation (7.40).

$$\chi_\sigma^{\text{noiseless}} := \min_{\substack{\mathbf{y}\in\mathbb{C}^{2N} \text{ s.t.} \\ ||\mathbf{y}||_2=\sqrt{2N}}} Q^{(\sigma)}(\mathbf{y}) = (R+1)N\,\lambda_{\min}\left(A_\sigma^\dagger A_\sigma\right) = Q^{(\sigma)}(\mathbf{x}) = 0, \tag{7.41}$$

where the optimization is over *all* vectors with a 2-norm equal to $\sqrt{2N}$. In other words, in the noiseless scenario, the correct assignment of phases $\mathbf{x}$ can be obtained by determining the smallest eigenvector of $A_\sigma^\dagger A_\sigma$.

7.B.1.2. The noisy scenario

In practice, one only obtains the noisy versions of the absolute value measurements $|f_1[j]|^2$, $|f_2[j]|^2$, $|f_3[j]|^2$ and $|f_4[j]|^2$, which we respectively denote by $|\tilde{f}_1[j]|^2$, $|\tilde{f}_2[j]|^2$, $|\tilde{f}_3[j]|^2$ and $|\tilde{f}_4[j]|^2$. We denote the correct vector of phases of $\tilde{\mathbf{f}}_1$ and $\tilde{\mathbf{f}}_2$ by $\tilde{\mathbf{x}} \in \mathbb{C}^{2N}$. Let us now define the cost function in terms of these noisy absolute value measurements as follows.

$$\tilde{Q}^{(s)}(\mathbf{y}) := \tilde{Q}^{(s)}_{\text{support}}(\mathbf{y}) + \tilde{Q}_{\text{interference}}(\mathbf{y}), \tag{7.42}$$

with

$$\tilde{Q}^{(s)}_{\text{support}}(\mathbf{y}) := \sum_{k=s}^{N-1} \left| \frac{1}{N} \sum_{j=0}^{N-1} |\tilde{f}_1[j]|\, y_j \exp\left[-i2\pi jk/n\right] \right|^2 + \sum_{k=s}^{N-1} \left| \frac{1}{N} \sum_{j=0}^{N-1} |\tilde{f}_2[j]|\, y_{N+j} \exp\left[-i2\pi jk/n\right] \right|^2, \tag{7.43}$$

and

$$\tilde{Q}_{\text{interference}}(\mathbf{y}) := \sum_{j=0}^{N-1} \left| y_j - y_{N+j}\, \tilde{G}[j]/|\tilde{G}[j]| \right|^2, \quad \text{with } \tilde{G}[j] := \frac{|\tilde{f}_3[j]|^2 + i|\tilde{f}_4[j]|^2 - (1+i)\left(|\tilde{f}_1[j]|^2 + |\tilde{f}_2[j]|^2\right)}{2|\tilde{f}_1[j]|\,|\tilde{f}_2[j]|}, \tag{7.44}$$

where we stress that $\tilde{G}[j]$ does not necessarily have unit absolute value.

Again, the cost function can be expressed as

$$Q^{(s)}(\mathbf{y}) := \mathbf{y}^\dagger \tilde{A}_s^\dagger \tilde{A}_s \mathbf{y}, \tag{7.45}$$

where $\tilde{A}_s \in \mathbb{C}^{(2(N-s)+N)\times 2N}$ has entries

$$(\tilde{A}_s)_{p,q} = \begin{cases} |\tilde{f}_1[q]| \exp\big(i2\pi q(s+p)/N\big), & \text{if } q < N,\ p < N-s, \\ |\tilde{f}_2[q]| \exp\big(i2\pi(q-N)(s+p-(N-s))/N\big), & \text{if } N \le q < 2N,\ (N-s) \le p < 2(N-s), \\ 1, & \text{if } q < N,\ p = q + 2(N-s), \\ -\tilde{G}[q-N]/|\tilde{G}[q-N]|, & \text{if } N \le q < 2N,\ p = q - N + 2(N-s), \\ 0, & \text{otherwise,} \end{cases} \tag{7.46}$$

with $p = 0, 1, \ldots, (2(N-s)+N)-1$ and $q = 0, 1, \ldots, 2N-1$.

In the noisy setting, there is no guarantee that there is an assignment of phases $\mathbf{y}$ for which $\tilde{Q}^{(s)}(\mathbf{y}) = 0$, but the optimal assignment of phases could be found by solving the non-convex optimization problem

$$\phi_s^{\text{noisy}} := \min_{\substack{\mathbf{y}\in\mathbb{C}^{2N} \text{ s.t.} \\ |y_0|=|y_1|=\ldots=|y_{N-1}|=1}} \tilde{Q}^{(s)}(\mathbf{y}), \tag{7.47}$$

at $s = \sigma$. In general, we now have that $\phi_s^{\text{noisy}} > 0$. The correct assignment $\tilde{\mathbf{x}}$ is not guaranteed to be the unique minimizer of $\tilde{Q}^{(\sigma)}(\mathbf{y})$. However, for relatively small noise magnitudes, any assignment of phases for which $\tilde{Q}^{(\sigma)}(\mathbf{y})$ is minimized tends to be close to the correct assignment $\tilde{\mathbf{x}}$. Like in the noiseless scenario, we will consider the following relaxation of the problem in equation (7.47) to an eigenvalue problem.

$$\chi_\sigma^{\text{noisy}} := \min_{\substack{\mathbf{y}\in\mathbb{C}^{2N} \text{ s.t.} \\ ||\mathbf{y}||_2=\sqrt{2N}}} \tilde{Q}^{(\sigma)}(\mathbf{y}) = 2N\,\lambda_{\min}\big(\tilde{A}_\sigma^\dagger \tilde{A}_\sigma\big), \tag{7.48}$$

where we note that again, generally, $\chi_\sigma^{\text{noisy}} > 0$. For relatively small noise magnitudes, the optimal assignment obtained by solving equation (7.48) (which we denote by $\mathbf{y}_{\min}$) will be close to the correct assignment $\tilde{\mathbf{x}}$.

Before discussing the details of the algorithm, let us consider a generalization that we introduce in this work. This generalization consists of the inclusion of *multiple* time series that are used for interference, as opposed to just a single one. As we shall demonstrate, this leads to improved performance of the algorithm in the presence of noise.

7.B.2. Vectorial phase retrieval using multiple interference signals

In this work, we consider the scenario in which one employs *multiple* signals $\mathbf{f}_2^{(r)} \in \mathbb{C}^N$ (with $r = 1, 2, \ldots, R$) with which the target time series $\mathbf{f}_1$ interferes. As we will demonstrate in Section 7.B.3, we find that the performance of vectorial phase retrieval improves when multiple interference signals are employed in the noisy scenario. This multi-interference vectorial phase retrieval problem is formulated as follows.

Problem 7.B.2. [Multi-interference vectorial phase retrieval] Given measurements of the absolute values

$$|f_1[j]|^2,\ \Big\{|f_2^{(r)}[j]|^2, |f_3^{(r)}[j]|^2 := |f_1[j] + f_2^{(r)}[j]|^2, |f_4^{(r)}[j]|^2 := |f_1[j] + i f_2^{(r)}[j]|^2\Big\}_{r=1}^{R}, \tag{7.49}$$

at times $j = 0, 1, \ldots, N-1$, determine

$$F_1[k] := \sum_{j=0}^{N-1} f_1[j] \exp\big[-i2\pi jk/N\big], \ \Big\{F_2^{(r)}[k] := \sum_{j=0}^{N-1} f_2^{(r)}[j] \exp\big[-i2\pi jk/N\big]\Big\}_{r=1}^{R}, \quad (7.50)$$

at $k = 0, 1, \ldots, N-1$.

The trivial ambiguities of Problem 7.B.2 are equivalent to those of Problem 7.B.1, involving *all* $(R+1)$ signals $\mathbf{F}_1, \mathbf{F}_2^{(1)}, \ldots, \mathbf{F}_2^{(R)}$ (instead of just $\mathbf{F}_1, \mathbf{F}_2$).

We will discuss how the multi-interference v-PR algorithm can be implemented to (approximately) solve Problem 7.B.2. But let us first discuss how the control-free quantum phase estimation routine fits into the multi-interference v-PR framework. We take

$$f_1[j] = \langle\Phi| \exp\big(iH\Delta t\, j\big) |\Phi\rangle \text{ and } \Big\{f_2^{(r)}[j] = \langle\Phi| \exp\big(iH\Delta t\, j\big) \big|\psi_r\big\rangle\Big\}_{r=1}^{R}, \quad (7.51)$$

such that

$$\Big\{f_3^{(r)}[j] = f_1[j] + f_2^{(r)}[j] = \langle\Phi| \exp\big(iH\Delta t\, j\big)\big(|\Phi\rangle + \big|\psi_r\big\rangle\big) \text{ and}$$
$$f_4^{(r)}[j] = f_1[j] + i f_2^{(r)}[j] = \langle\Phi| \exp\big(iH\Delta t\, j\big)\big(|\Phi\rangle + i\big|\psi_r\big\rangle\big)\Big\}_{r=1}^{R}. \quad (7.52)$$

Hence by preparing states $|\Phi\rangle$ and $\big\{\big|\psi_r\big\rangle, 1/\sqrt{2}\big(|\Phi\rangle + \big|\psi_r\big\rangle\big), 1/\sqrt{2}\big(|\Phi\rangle + i\big|\psi_r\big\rangle\big)\big\}_{r=1}^{R}$, then time evolving *each* of the states under Hamiltonian H for a time $\Delta t\, j$, and measuring the overlap with $|\Phi\rangle$, we can estimate $|f_1[j]|$ and $\big\{|f_2^{(r)}[j]|, |f_3^{(r)}[j]|, |f_4^{(r)}[j]|\big\}_{r=1}^{R}$ at $j = 0, 1, \ldots, N-1$. Of course, we do not want the states $|\Phi\rangle$ and $\{\big|\psi_r\big\rangle\}_{r=1}^{R}$ to be such that preparing superpositions of $|\Phi\rangle$ with any $\big|\psi_r\big\rangle$ requires applying circuits of very large depth.

As discussed before, one only ever obtains noisy versions of the absolute value measurements $|f_1[j]|^2$, $\big\{|f_2^{(r)}[j]|^2, |f_3^{(r)}[j]|^2, |f_4^{(r)}[j]|^2\big\}_{r=1}^{R}$ in practice, which we denote by $|\tilde{f}_1[j]|^2$, $\big\{|\tilde{f}_2^{(r)}[j]|^2, |\tilde{f}_3^{(r)}[j]|^2, |\tilde{f}_4^{(r)}[j]|^2\big\}_{r=1}^{R}$. The multi-interference v-PR algorithm takes as *input* these noisy absolute value measurements, and an estimate s of σ (with σ denoting the support size of $\mathbf{F}_1, \mathbf{F}_2^{(1)}, \ldots, \mathbf{F}_2^{(R)}$). The *output* of the algorithm is a vector $\mathbf{y} \in \mathbb{C}^{(R+1)N}$ whose entries are estimates of the phases of $f_1[j]$, $f_2^{(1)}[j]$, $f_2^{(2)}[j]$, …, $f_2^{(R)}[j]$ at $j = 0, 1, \ldots, N-1$. The algorithm again consists of the minimization of a cost function over all these assignments of phases to the absolute value measurements $\{|\tilde{f}_1[j]|\}$, $\{|\tilde{f}_2^{(1)}[j]|\}$, …, $\{|\tilde{f}_2^{(R)}[j]|\}$. We denote by $\tilde{\mathbf{x}} \in \mathbb{C}^{(R+1)N}$ the correct phases of $f_1[j]$, $f_2^{(1)}[j]$, $f_2^{(2)}[j]$, …, $f_2^{(R)}[j]$ at $j = 0, 1, \ldots, N-1$.

The cost function again consists of a support component $\tilde{Q}_{\text{support}}^{(s)}(\mathbf{y})$ and an interference component $\tilde{Q}_{\text{interference}}(\mathbf{y})$. The support component of the cost function can be written as

$$\tilde{Q}_{\text{support}}^{(s)}(\mathbf{y}) := \sum_{k=s}^{N-1} \Big| \sum_{j=0}^{N-1} |\tilde{f}_1[j]|\, y_j \exp\big[-i2\pi jk/n\big]\Big|^2$$
$$+ \sum_{r=1}^{R} \Bigg[\sum_{k=s}^{N-1} \Big| \sum_{j=0}^{N-1} |\tilde{f}_2^{(r)}[j]|\, y_j^{(r)} \exp\big[-i2\pi jk/n\big]\Big|^2 \Bigg], \quad (7.53)$$

which is equivalent to the support component defined in the single-interference-signal framework, except that now it constrains the support of $R+1$ signals. The interference component $\tilde{Q}_{\text{interference}}(\mathbf{y})$, which is again an extension of the single-interference-signal interference component to $R+1$ signals, can be defined as follows.

$$\tilde{Q}_{\text{interference}}(\mathbf{y}) := \sum_{r=1}^{R} \Big[\sum_{j=0}^{N-1} \Big| y_j - y_j^{(r)} \tilde{G}_r[j]/|\tilde{G}_r[j]| \Big|^2 \Big], \tag{7.54}$$

where the quantities $\tilde{G}_r[j]$ are defined at $j = 0, 1, \ldots, N-1$ for $r = 1, 2, \ldots, R$ as

$$\tilde{G}_r[j] := \frac{|\tilde{f}_3^{(r)}[j]|^2 + i|\tilde{f}_4^{(r)}[j]|^2 - (1+i)\big(|\tilde{f}_1[j]|^2 + |\tilde{f}_2^{(r)}[j]|^2\big)}{2|\tilde{f}_1[j]|\,|\tilde{f}_2^{(r)}[j]|}. \tag{7.55}$$

The total cost function

$$\tilde{Q}^{(s)}(\mathbf{y}) := \tilde{Q}^{(s)}_{\text{support}}(\mathbf{y}) + \tilde{Q}_{\text{interference}}(\mathbf{y}) \tag{7.56}$$

can be expressed as

$$\tilde{Q}^{(s)}(\mathbf{y}) := \mathbf{y}^\dagger \tilde{A}_s^\dagger \tilde{A}_s \mathbf{y}, \tag{7.57}$$

where $\tilde{A}_{s,R} \in \mathbb{C}^{((R+1)(N-s)+RN)\times(R+1)N}$ has entries

$$(A_s)_{p,q} = \begin{cases} |\tilde{f}_1[q]| \exp\big(i2\pi q(s+p)/N\big), & \text{if } q < N,\ p < N-s, \\ |\tilde{f}_2[q]| \exp\big(i2\pi (q-N)(s+p-(N-s))/N\big), & \text{if } N \le q < 2N,\ (N-s) \le p < 2(N-s), \\ \vdots & \\ |\tilde{f}_{R+1}[q]| \exp\big(i2\pi (q-RN)(s+p-R(N-s))/N\big), & \text{if } RN \le q < (R+1)N,\ R(N-s) \le p < (R+1)(N-s), \\ 1, & \text{if } q < N,\ p = q + (R+1)(N-s), \\ -\tilde{G}_2[q-N]/|\tilde{G}_2[q-N]|, & \text{if } N \le q < 2N,\ p = q - N + (R+1)(N-s), \\ 1, & \text{if } q < N,\ p = q + (R+1)(N-s) + N, \\ -\tilde{G}_3[q-2N]/|\tilde{G}_3[q-2N]|, & \text{if } 2N \le q < 3N,\ p = q - N + (R+1)(N-s), \\ \vdots & \\ 1, & \text{if } q < N,\ p = q + (R+1)(N-s) + (R-1)N, \\ -\tilde{G}_{R+1}[q-RN]/|\tilde{G}_{R+1}[q-RN]|, & \text{if } RN \le q < (R+1)N,\ p = q - N + (R+1)(N-s), \\ 0, & \text{otherwise,} \end{cases} \tag{7.58}$$

with $p = 0, 1, \ldots, ((R+1)(N-s)+RN)-1$ and $q = 0, 1, \ldots, (R+1)N-1$.

The optimal assignment of phases could be found by solving the non-convex optimization problem

$$\Phi_\sigma^{\text{noisy}} := \min_{\substack{\mathbf{y}\in\mathbb{C}^{(R+1)N} \text{ s.t.} \\ |y_0|=|y_1|=\ldots=|y_{N-1}^{(R)}|=1}} \tilde{Q}^{(\sigma)}(\mathbf{y}). \tag{7.59}$$

In general, we have that $\Phi_\sigma^{\text{noisy}} > 0$. The correct assignment $\tilde{\mathbf{x}}$ is again not guaranteed to be a (unique) minimizer of $\tilde{Q}^{(\sigma)}(\mathbf{y})$. However, for relatively small noise magnitudes, any assignment of phases for which $\tilde{Q}^{(\sigma)}(\mathbf{y})$ is minimized tends to be close to the correct assignment $\tilde{\mathbf{x}}$.

We again relax the problem in equation (7.59) to an eigenvalue problem as follows.

$$X_\sigma^{\text{noisy}} := \min_{\substack{\mathbf{y}\in\mathbb{C}^{(R+1)N} \text{ s.t.} \\ ||\mathbf{y}||_2=\sqrt{(R+1)N}}} Q^{(\sigma)}(\mathbf{y}) = (R+1)N\,\lambda_{\min}\big(\tilde{A}_\sigma^\dagger \tilde{A}_\sigma\big), \tag{7.60}$$

where we note that again, generally, $X_\sigma^{\text{noisy}} > 0$. As we shall demonstrate in Section 7.B.3, for relatively small noise magnitudes, the optimal assignment $\mathbf{y}_{\min}$ obtained by solving equation (7.60) will be close to the correct assignment $\tilde{\mathbf{x}}$. Moreover, we will show that for a given noise magnitude, the assignment of phases obtained by solving equation (7.60) will typically become a better approximation for $\tilde{\mathbf{x}}$ as R increases.

We note that the optimal assignment $\mathbf{y}_{\min}$ obtained by solving equation (7.60) is not necessarily an assignment of phases. Although we ensure that $||\mathbf{y}_{\min}||_2 = \sqrt{(R+1)N}$, each entry can have absolute value different from one. One could consider rounding $\mathbf{y}_{\min}$ to an assignment of phases by dividing each entry by its absolute value. We do not perform this rounding as we find it does not *necessarily* improve the approximation to $\tilde{\mathbf{x}}$ in the scenarios that we consider.

In Section 7.B.3, we will demonstrate numerically how the v-PR algorithm can be used to perform spectral estimation for instances of the FH Hamiltonian, and discuss its performance in various settings.

7.B.3. Vectorial phase retrieval algorithm for spectral estimation of Hamiltonians

In this section, we numerically investigate the performance of the vectorial phase retrieval algorithm in the context of spectral estimation for Hamiltonians. The signals encountered in this context do not perfectly fit within the framework of v-PR. In particular, due to the fact that the (absolute values of) the time series are only measured for a finite measurement time, the signals $\mathbf{F}_1$, $\mathbf{F}_2^{(1)}$, …, $\mathbf{F}_2^{(R)}$ do not necessarily have well-defined support. Typically, $|F_1[k]|$, $|F_2^{(1)}[k]|$, …, $|F_2^{(R)}[k]|$ decay away from the supported eigenvalues, rather than being equal to zero exactly outside of some frequency interval. The more the signals suffer from this effect, the more ill-defined their support is becomes. As a result, a good performance of the vectorial phase retrieval algorithm cannot be guaranteed without further considerations.

To set the stage, we first provide an example for which the v-PR input data is obtained from time evolution of some state under the FH Hamiltonian, where we have artificially ensured that the input data fits exactly in the v-PR framework. Namely, we have altered the input data in such a way that $\mathbf{F}_1$, $\mathbf{F}_2^{(1)}$ (we take $R = 1$ here) are exactly equal to zero outside of an interval of size $\sigma = 25$. Note that to do this, one generally needs access to the full time series and not just their absolute values (so this cannot be done in practice). The data we have used here is noiseless data. Based on the discussion in Section 7.B.1.1, we expect the matrix $A_s^\dagger A_s$ to have a *single* zero eigenvalue at $s = \sigma$, and its smallest eigenvector to correspond to the exact assignment of phases. Figure 7.7 depicts the spectrum of $A_s^\dagger A_s$ as a function of s. Indeed, its smallest eigenvalue is non-zero for $s < \sigma$ and (numerically) zero at $s = \sigma$. For $s > \sigma$, there are several zero-eigenvalue eigenvectors, with the correct solution lying in the span of those eigenvectors. In the noisy setting, there will generally be *no* zero-eigenvalue solutions. However, for sufficiently small noise

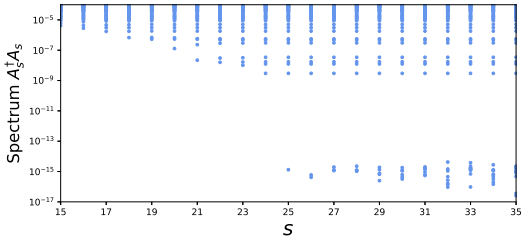

Figure 7.7: Spectrum of $A_s^\dagger A_s$ for $R = 1$ in the noiseless setting, as a function of s. The time series amplitudes used to construct $A_s^\dagger A_s$ correspond to evolution under a 1×5 (i.e., 10-mode) Fermi-Hubbard Hamiltonian for which $|\tau/U| = 1/4$. We have taken $N = 300$. The signals f_1 and f_2 were artificially processed to ensure that they have well-defined support (here, $\sigma = 25$).

magnitudes, there will still be a drop in the value of the smallest eigenvalue of $A_s^\dagger A_s$ around σ. Our investigations suggest that:

1. The two smallest eigenvalues are a proxy for the quality of the retrieval of phases. That is, the smaller the smallest eigenvalue, the smaller the value of the two residuals equations (7.2) and (7.3) above.

2. A large gap opening between the two smallest eigenvalues suggests a larger overlap between the smallest eigenvector and the true assignment of phases. This being motivated by the fact that in the ideal noiseless case the solution is unique.

In the remainder of this section, we will discuss the performance of vectorial phase retrieval for input data obtained from time evolution under FH Hamiltonians. The procedure that is implemented is given as a pseudocode below. As discussed, this input data does not perfectly fit into the v-PR framework. Nevertheless, the v-PR algorithm performs well in performing phase retrieval in the scenarios that we consider.

Algorithm 7.2: v-PR

Input: $|\tilde{f}_1[j]|$ for $j = 0, \ldots, N-1$;
$|\tilde{f}_2^{(r)}[j]|, |\tilde{f}_3^{(r)}[j]|, |\tilde{f}_4^{(r)}[j]|$ for $r = 1, \ldots, R$ and $j = 0, \ldots, N-1$
Output: $\{|f_1[j]|(y_{\min})_j\}_{j=0}^{N-1}$
for $s \leftarrow 0$ **to** $N-1$ **do**
  Evaluate $G_r[j]$ for $r = 1, \ldots, R$ and $j = 0, \ldots, N-1$
  Form $\tilde{A}_s^\dagger \tilde{A}_s \in \mathbb{C}^{(R+1)N \times (R+1)N}$
  Compute the two smallest eigenvalues of $\tilde{A}_s^\dagger \tilde{A}_s$ (denote the smallest by $\lambda_{\min}(\tilde{A}_s^\dagger \tilde{A}_s)$)
Choose s^* (heuristically)
Compute $\mathbf{y}_{\min} \in \mathbb{C}^{(R+1)N}$ such that $\mathbf{y}_{\min}^\dagger \tilde{A}_{s^*}^\dagger \tilde{A}_{s^*} \mathbf{y}_{\min} = (R+1)N\lambda_{\min}(\tilde{A}_{s^*}^\dagger \tilde{A}_{s^*})$
return $\{|f_1[j]|(y_{\min})_j\}_{j=0}^{N-1}$

Interplay between R and the noise resilience of the vectorial phase retrieval algorithm: Figure 7.8 illustrates the difference in noise resilience of the v-PR algorithm between the $R = 1$ and the $R = 10$ scenarios. This difference materializes as follows: In the noisy scenario, there is a sharp drop in the smallest eigenvalue of $\tilde{A}_s^\dagger \tilde{A}_s$ as a function of s for $R = 10$, which is absent for $R = 1$. This can be seen in Figure 7.8c. Our numerical investigations suggest that this sharp drop is a signature of an accurate retrieval of phases at that value of s. As can be seen in Figure 7.8d, indeed the estimate of $\mathbf{F}_1$ for $R = 10$ is more accurate than for $R = 1$. Figure 7.8b suggests that choosing larger R does not influence the accuracy significantly in the noiseless scenario.

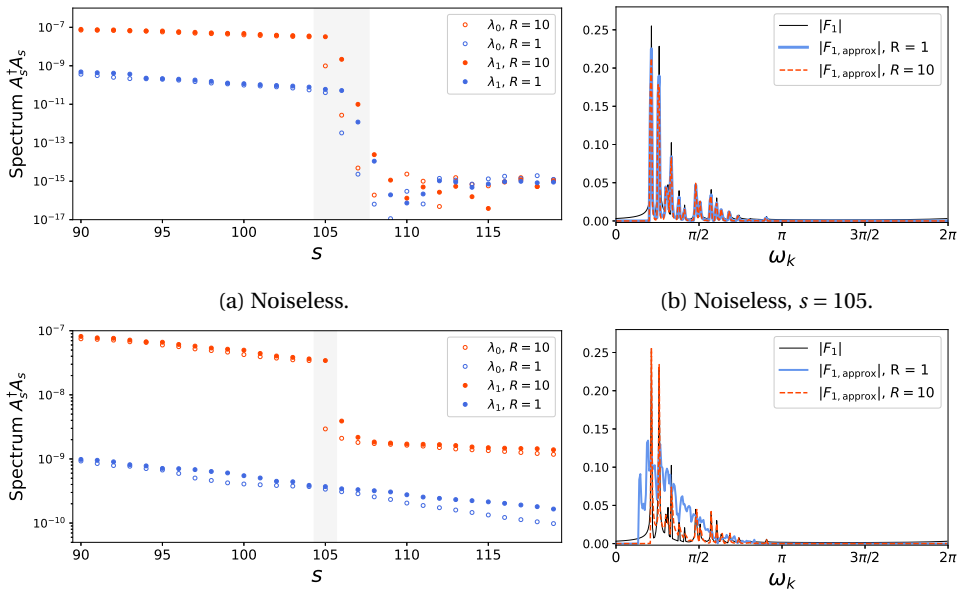

(a) Noiseless.

(b) Noiseless, $s = 105$.

(c) $N_{\text{samples}} = 10^6$.

(d) $N_{\text{samples}} = 10^6$, $s = 105$.

Figure 7.8: The smallest two eigenvalues of $A_s^\dagger A_s$ for $R = 1$ and $R = 10$ in the noiseless setting (see (a)), and for $N_{\text{samples}} = 10^6$ (see (c)). The eigenvalues are given as a function of s. Approximations to $|F_1|$ at fixed $s = 105$ for $R = 1$ and $R = 10$ are depicted in the noiseless setting in (b) and for $N_{\text{samples}} = 10^6$ in (d). The time series amplitudes used to construct $A_s^\dagger A_s$ correspond to evolution under a 1×5 (i.e., 10-mode) Fermi-Hubbard Hamiltonian for which $|\tau/U| = 1/4$.

Performance of the vectorial phase retrieval algorithm as a function of sample size: For completeness, we include a plot of the l^2-norm *error* of the recovered spectrum $\mathbf{F}_1$ in Figure 7.9. Let us note that the error floor appearing around 10^6 samples seems to be an artifact of the "choice" for measuring the error using the l^2 norm: As can be seen in Figure 7.8d, the peaks in the spectrum are accurately recovered for $R = 10$ at these noise magnitudes, but the recovery of the tails is poorer.

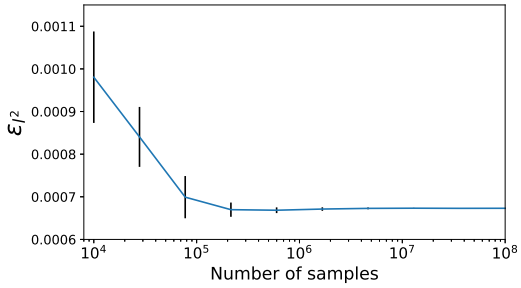

Figure 7.9: Spectrum error as a function of number of samples used to estimate each $|f[j]|$. Here, we consider the $R = 10$ instance from Figure 7.8 and we define the spectrum error as $\epsilon_{l^2} = \frac{1}{N}\sqrt{\sum_k |F[k] - F_{\text{true}}[k]|^2}$. At each fixed number of samples we average over 10 instances of the v-PR algorithm and include error bars indicating one standard deviation.

Performance of the vectorial phase retrieval algorithm for shallower time evolutions: The fact that $\mathbf{F}_1, \mathbf{F}_2^{(1)}, \ldots, \mathbf{F}_2^{(R)}$ do not have well-defined support is the most prominent reason that the v-PR input data in the current setting does not fit perfectly into the framework of v-PR. The support of these signals typically becomes more ill-defined (i.e., $\mathbf{F}_1, \mathbf{F}_2^{(1)}, \ldots, \mathbf{F}_2^{(R)}$ decay more slowly away from the supported frequencies) as the total time evolution becomes smaller (due to spectral leakage). Since the circuit depths are limited in our near-term scenario, it is of interest to see how the v-PR algorithm fares for shallower time evolutions.

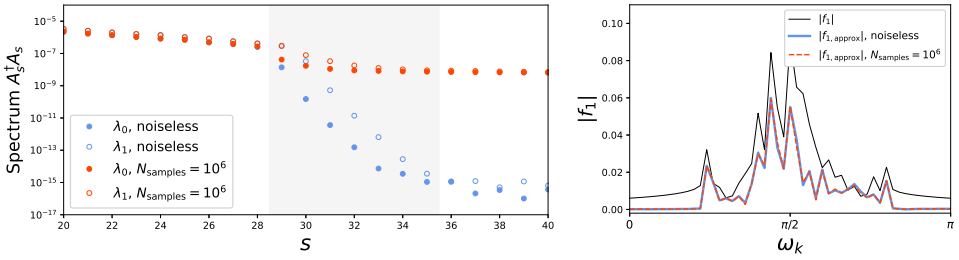

(a) $R = 10$. (b) $R = 10$, $s = 29$.

Figure 7.10: The smallest two eigenvalues of $A_s^\dagger A_s$ for $R = 10$ in the noiseless and noisy settings (see (a)). The eigenvalues are given as a function of s. Approximations to $|F_1|$ at fixed $s = 29$ for $R = 10$ are depicted in the noiseless setting and for $N_{\text{samples}} = 10^6$ in (b). The time series amplitudes used to construct $A_s^\dagger A_s$ correspond to evolution under a 1×5 (i.e., 10-mode) Fermi-Hubbard Hamiltonian for which $|\tau/U| = 1/4$, with a relatively small total time evolution T. This small total time evolution leads to ill-defined support of $|F_1|$. Note that the input state $|\Phi\rangle$ is chosen differently than in Figure 7.8.

Figure 7.11a depicts the smallest two eigenvalues $A_s^\dagger A_s$ as a function of s. Clearly, in the noiseless setting, $\lambda_{\min}\left(A_s^\dagger A_s\right)$ decays to zero comparatively slowly, which is a result the signals having ill-defined support. We find that the quality of the approximation $|F_{1,\text{approx}}|$ to $|F_1|$ is highest at the value of s at which $\lambda_{\min}\left(A_s^\dagger A_s\right)$ first starts to decay, i.e., $s = 29$ for Figure 7.10. Figure 7.11b depicts $|F_{1,\text{approx}}|$ and in the noiseless scenario, and for $N_{\text{samples}} = 10^6$. Remarkably, although the approximation of $|F_{1,\text{approx}}|$ to $|F_1|$ is worse (even in the noiseless scenario) than in Figure 7.8, the approximation retains its resilience against the addition of sampling noise.

To round, or not to round:
At the end of Section 7.B.2, we briefly commented on whether it is beneficial to round the smallest eigenvector $\mathbf{y}_{\min}$ (with 2-norm $\sqrt{(R+1)N}$) of the matrix $\tilde{A}_s^\dagger \tilde{A}_s$ to a vector whose entries are phase factors, by dividing each entry of $\mathbf{y}_{\min}$ by its absolute value. Our numerical investigations suggest that using the entries of this rounded vector to reconstruct the time series does not *necessarily* lead to a more accurate reconstruction. In terms of the reconstructed spectra F (see Figure 7.11 below, corresponding to the scenario in Figure 7.6a), we note the following differences. The unrounded spectrum recovers the peak locations relatively well. The rounded spectrum performs better at recovering the absolute values $|F|$. However, the latter also seems to recover some peaks for frequencies at which F is not actually supported. Whether rounding is beneficial thus depends on which features one wishes to recover. To quantify the error in the particular case of Figure 7.11, we note that the 1-norm (normalized by N) of the vector with entries $|F_{\text{reconstructed}}[k]| - |F_{\text{exact}}[k]|$ is 0.0016 without rounding and 0.0013 with rounding.

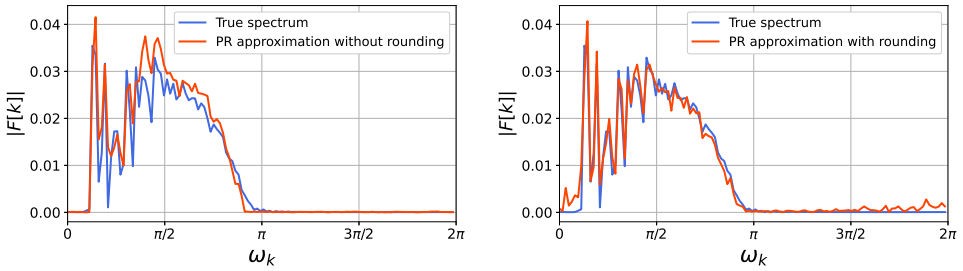

(a) Without rounding.

(b) With rounding.

Figure 7.11: Comparison between the unrounded (a) and rounded (b) solutions obtained through vectorial phase retrieval. The scenario used to obtain the figures is the same as that of Figure 7.6.

7.C. Comparison with the standard approach

There are two key advantages of phase retrieval compared with the standard approach to statistical phase estimation: each qubit no longer needs to interact with an ancilla "control" qubit, and there is no increase in circuit complexity caused by the use of controlled unitary operations. One can consider many metrics for circuit complexity;

here, we will focus on the number of CNOT gates required, and the quantum circuit depth.

To get a sense of how large the reduction in complexity can be, first observe that if the quantum hardware allows for all-to-all connectivity, any quantum circuit computing a unitary operator U can be converted into a quantum circuit for computing controlled-U by replacing every CNOT gate with a Toffoli gate, and every single-qubit unitary V with a controlled-V operator. A Toffoli gate can be implemented using 6 CNOT gates and controlled-V can be implemented using 2 CNOT gates for any single-qubit unitary V. So, if the original quantum circuit had c CNOT gates and s single-qubit gates, the controlled circuit will have at most $6c+2s$ CNOT gates – at most a constant factor increase in complexity, if $c \geq s$ (for example), but still relatively substantial for near-term applications. However, the difference in complexity between the two situations can be greater if there are restrictions on hardware connectivity, or if one considers circuit depth instead.

Here we approximately calculate the reduction in circuit complexity achieved by phase retrieval for several simple examples: the 1D Ising model with transverse field, both all-to-all and with matching hardware connectivity, and the spinless 2D Fermi-Hubbard model with square-lattice hardware connectivity. In each case, we consider the complexity of implementing k Trotter steps – i.e. k repetitions of an operator of the form $\prod_{j=1}^{m} e^{i\theta_j H_j}$ for a Hamiltonian of the form $H = \sum_j H_j$. Note that it is important to consider $k > 1$ steps because the cost of interacting with the control qubit can be amortised across multiple steps by preparing a GHZ state, as we will see below. This approach may add its own difficulties in terms of rendering the circuit more prone to decoherence, but we will ignore this issue for simplicity. We also stress that there is no guarantee that the implementations we describe here are optimal.

Model	H/w	CNOTs (PR)	CNOTs (no PR)	Depth (PR)	Depth (no PR)
1d TFIM	All-to-all	$(2n-2)k$	$(6n-4)k$	$4k$	$2\lceil \log_2 n \rceil + 10k$
1d TFIM	1d	$(2n-2)k$	$6(n-1)$ $+(6n-4)k$	$4k$	$6\lceil n/2 \rceil + 10k$
2d FH	2d	$32kn(n-1)/2$	$48kn(n-1)/2$ $+6\lceil (n-1)^2/2 \rceil$	$32k$	$48k+3(n-2)$

Table 7.2: The cost of implementing k Trotter layers for several Hamiltonians, an n qubits transverse field Ising model and an $n \times n$ spinless Fermi-Hubbard model. "Depth" is CNOT depth. Note that the algorithms achieving the minimal CNOT depth and CNOT count may be different. We consider the spinless Fermi-Hubbard model for simplicity, and these calculations are approximate.

1. **1D Ising model with transverse field.** This Hamiltonian is defined as $H = \sum_{i=1}^{n-1} Z_i Z_{i+1} + \sum_{j=1}^{n} X_j$. A Trotter step can be implemented using $2(n-1)$ CNOT gates and in CNOT depth 4. To implement a controlled Trotter step, we need to implement controlled-$e^{i\theta ZZ}$ and controlled-$e^{i\theta X}$ gates. These have circuits using four and two CNOT gates, respectively. With all-to-all connectivity, then,

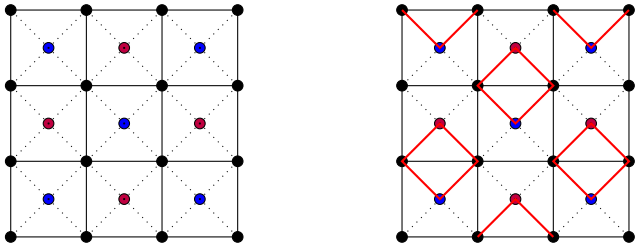

Figure 7.12: A possible layout of the spinless Fermi-Hubbard model on a lattice with 2D connectivity. The connectivity is indicated by dashed lines while the Fermi-Hubbard lattice interactions are indicated by solid lines. The face qubits in purple are needed for the compact encoding. The face qubits in blue are used to implement geometrically local control. The diagram on the right indicates one of the four Trotter layer groupings of interactions which can be implemented in parallel, in this case horizontal hopping interactions. With the control and encoding qubits these interactions are at most Pauli weight-4.

each Trotter layer can be implemented using $4(n-1)+2n = 6n-4$ CNOT gates. Implementing these in a straightforward way would lead to a high-depth circuit, as each controlled unitary acts on the same control qubit. However, we can reduce the depth by constructing a GHZ state of n qubits, where each gate that needs to be applied in parallel uses a different qubit of this state. The CNOT depth of producing such a state given all-to-all connectivity is at most $\lceil \log_2 n \rceil$. Following this, we can split the time-evolution steps into three groups (ZZ terms on (odd, even) qubits, ZZ terms on (even, odd) qubits, and X terms) and implement each group in parallel. At the end, the GHZ state is uncomputed, before measuring the control qubit as in the usual protocol. The total CNOT depth (with all-to-all connectivity) is then $2\lceil \log_2 n \rceil + 10k$ as each layer can be implemented in CNOT depth 10.

With 1D connectivity, the situation is more difficult as we need to create a distributed GHZ state. This can be done with a small extra cost in gate count, but a linear cost in depth. Now we include an ancilla qubit for controlled operations next to each qubit in the original 1D graph. Once these qubits have been prepared in a GHZ state, we can perform controlled operations across any pair of qubits in the original graph without the need for any swaps. The map $|000\rangle \mapsto |000\rangle$, $|100\rangle \mapsto |101\rangle$ can be carried out using three CNOT gates in 1d connectivity (CNOT_{12}, CNOT_{23}, CNOT_{12}), so the required state can be prepared using $3(n-1)$ CNOT gates and in CNOT depth $3\lceil n/2 \rceil$. Once this has been prepared, the costs for the time-evolution part are the same as for all-to-all connectivity. The total CNOT cost for k Trotter steps, including the cost of uncomputing the GHZ state, is then $6(n-1)+(6n-4)k$, and the CNOT depth is $6\lceil n/2 \rceil + 10k$.

2. **2D Fermi-Hubbard model.** Consider an $n \times n$ Fermi-Hubbard model without spin on an 2D lattice. Consider even n only and use the compact fermionic encoding [39]

and layout shown in Figure 7.12. The full Hamiltonian is made of horizontal $\sum_{\langle i,j\rangle}(X_iX_jY_{f(i,j)}+Y_iY_jY_{f(i,j)})/2$ and vertical $\sum_{\langle i,j\rangle}\pm(X_iX_jX_{f(i,j)}+Y_iY_jX_{f(i,j)})/2$ interactions, where the sum is over pairs of horizontally or vertically adjacent sites on an $n\times n$ lattice and $f(i,j)$ are face qubits placed on this lattice. The Hamiltonian can be split into four sets of $n(n-1)/2$ interactions. Within each group all interactions can be performed in parallel. Every interaction consists of two weight-3 (at most) Pauli interactions. Therefore each interaction can be done using 8 CNOT gates, meaning that a full Trotter step will have a depth of 32 CNOT gates and the total number of CNOT gates will be $32n(n-1)/2$. This calculation is approximate and simply treats edge interactions identically to those in the bulk. When performing controlled evolution under this Hamiltonian one possible layout of control ancilla qubits is shown in Figure 7.12. As in the case of the Ising model above, we consider creating a distributed GHZ state so that we can implement local control. This particular selection uses the remaining face qubits not used by the compact encoding. Now every "controlled" interaction can be performed by evolving under two weight-4 Pauli terms, where the interactions have been augmented so that they include a Pauli-Z on the local control ancilla. Therefore each interaction can be done using 12 CNOT gates, meaning that a full Trotter step will have a depth of 48 CNOT gates and the total number of CNOT gates will be $48n(n-1)/2$. (This all uses the standard approach to performing a controlled rotation by a weight-k Pauli operation [40]) The additional cost of creating the distributed GHZ state depends on the number of control ancillas $\lceil(n-1)^2/2\rceil$ and their connectivity. The total count of CNOT gates will be $3\times\lceil(n-1)^2/2\rceil$. The depth is approximately $3(n-2)/2$ (for even lattices with $n>2$). One can begin building the GHZ state from the central control qubit and move outward along the branches approximately in parallel. It takes 3 CNOT gates to spread the entanglement to the nearest control qubit, shown in blue. This is done in the same manner as described for the Ising model.

In all of the above costings, we do not include the cost of preparing (or uncomputing) the initial state $|\psi\rangle$ – which we assume to be simple – nor the cost of preparing a superposition of $|\psi\rangle$ with another state. In some situations, this may be fairly substantial, e.g. if this superposition is a GHZ state.

8

Discussion and outlook

In this thesis we have studied computational tasks that appear in quantum many-body physics, with a focus on fermionic systems. Most contributions are either complexity-theoretic classifications or – where possible – constructive algorithmic results. Next, let us discuss some of the most important aspects of this thesis and look ahead to potential new avenues of research.

In Chapter 2, we have established the complexity of the fermionic satisfiability problem, FERMIONIC 2-SAT. The bare problem is in P, and if we add the *classical* global constraints of fixing particle-number-parity or fixing the particle number itself, the problem is respectively in P and NP-complete. If we increase the locality of the problem to $k = 9$, then it becomes QMA_1-hard – in line with QUANTUM k-SAT [1] which is in P for $k = 2$ and QMA_1-hard for $k \geq 3$. What can we say about the complexity of FERMIONIC 2-SAT or QUANTUM 2-SAT instances – i.e., keeping the locality of the bare problems at $k = 2$ – plus generic $k > 2$ *classical* constraints (beyond the global constraints considered in Chapter 2)? Are there any collections of such $k > 2$ *classical* constraints that can make these problems QMA_1-hard or will they remain in NP? If they remain in NP, can one establish a dichotomy such that any instance of the problem can be classified as either being in P or being NP-complete? One of the reasons that these questions are interesting is that adding classical $k > 2$ constraints also has physical motivations; they can be used to fix particle number (as we have studied in Chapter 2) or magnetization in some regions of the system, or they can be used to exclude some global occupation/spin configurations.

In Chapters 3 and 4, we have investigated Gaussian approximations to interacting fermion ground states. We have shown that the so-called Gaussian breakdown that has been proved for general fermionic Hamiltonians [2, 3] – meaning that Gaussian states achieve at most a vanishing fraction of the ground energy – can be avoided by imposing certain structural constraints on fermionic Hamiltonians. Assuming (1) the Hamiltonian to be sparse (so that each mode is involved in only a constant number of Hamiltonian terms) or (2) assuming that all interactions are of a density-density type is sufficient. These structural assumptions are physically motivated as (1) for instance lattice models are sparse and (2) interactions in (real-space discretized) quantum chemistry are of density-density type. We have established that fermionic Gaussian states achieve a constant ratio of the ground energy for these types of fermionic Hamiltonians, but one might wonder how well such states reproduce the physical phenomena that these

systems exhibit. More generally, what is the relation between Gaussian approximability and the presence/absence of certain many-body phenomena?

An interesting direction in that respect is whether Gaussian witnesses can be used to determine whether a given fermionic system is in some physically relevant finite-temperature phase. For example, given a density-density interacting fermionic Hamiltonian (on some generic interaction graph and hopping graph), and given a temperature and potentially other ambient conditions, is the system in a superconducting phase? We know – through BCS theory [4] – that Gaussian states can account for many phenomena related to superconductivity. But is there always a Gaussian witness for a superconducting phase? If such a witness always exists, then one might wonder whether it can be used to efficiently test whether a given system is in a superconducting phase or not. More generally, it is of interest to figure out whether this problem is in P or BQP at all. Note that this is not ruled out a priori by uncomputability results for *zero-temperature* phase transitions [5]. To make progress towards showing that the problem is in BQP at all, one might attempt to use existing methods for preparing Gibbs states of fermionic systems at constant temperatures (see e.g. [6]) and combine them with a scheme for measuring certain constant-weight correlators from which it can be deduced whether a superconducting condensate is present or not.

Focusing on more pragmatic problems like *deciding superconductivity* is valuable for the field of quantum simulation. Beyond being intrinsically interesting, this kind of focus speaks to quantum simulation's broader utility promise: it should be pursued not *only* as a scientific program in its own right, but also to provide computational capabilities that are intended to be useful beyond the quantum-information community.

Going back to the contents of this thesis, we have presented a method for compressed simulation of sparse free-fermion physics using a quantum computer in Chapter 5. We established an exponential memory compression and – more importantly – an exponential or strong-polynomial speedup in runtime (depending on the problem geometry) for practically relevant free-fermion problems compared to classical techniques. The use of quantum computers to provide a speedup for problems that are classically asymptotically tractable holds great promise. Beyond applications that are currently known, which include Chapter 5 of this thesis and the simulation of systems of coupled classical oscillators presented in Ref. [7], let us propose another idea in the same spirit for obtaining good graph cuts using a quantum computer.

The Max Cut problem is an NP-hard classical discrete optimization problem with many real-world applications. Given a (possibly weighted) graph, it is the problem of determining a split of the vertex set (into two) that maximizes the total weight of the edges running *across* the split. Equivalently, the goal is to solve $\max_{z\in\{\pm1\}^{|V|}} \sum_{j,k\in E} \frac{1}{2} w_{j,k}(1 - z_j z_k)$. Classical techniques for solving Max Cut are essentially as good as they can be. That is to say, there is an efficient SDP-based algorithm (the Goemans-Williamson algorithm [8]) that achieves approximation ratio 0.878 of the optimum, and obtaining a ratio of $0.878 + \epsilon$ is NP-hard. Since it is unlikely that quantum algorithms can efficiently solve NP-hard problems, one should not look to improve this approximation ratio for general graphs. In recent work [9], the authors have looked at (among many other applications) improving Max Cut approximation ratios using quantum algorithms for *restricted* graph instances. Our idea is to use a quantum

computer *not* to improve the approximation guarantee (by focusing on particular instances), but to obtain a *speedup* for getting to classically attainable approximation ratios using compressed simulation techniques[1]. We consider *sparse* graphs $G = (V, E)$ for which $|V| = 2^n$ (note that sparsity does not kill the NP-hardness of Max Cut). What can one learn about *good* cuts on G using $\text{poly}(n) = \log(|V|)$ quantum and classical effort? Clearly, we cannot learn the whole cut $z \in \{\pm 1\}^{|V|}$, but at most $\text{poly}(n)$ of the z_j's. In other words, we can learn on what side of the cut a subset of the vertices lie, for some good cut. Note that even the fact that we have to learn the cut values of a few variables does not break the NP-hardness of the problem. Our idea entails block-encoding matrix functions (using techniques from Ref. [10]) of the adjacency matrix A of G (with $A_{j,k} = w_{j,k}$ for $(j,k) \in E$ and zero elsewhere). Inspired by classical techniques that use the smallest eigenvector of A to construct good cuts (typically underperforming compared to the GW algorithm, but still achieving non-trivial approximation ratios), one could block-encode matrix functions such as $\exp(-\beta A)$ and estimate its entries to gain insight into the cut values of individual vertices on such good cuts. To the best of our knowledge, there are two existing approaches for (provably) realizing a speedup to obtain good *classically attainable* graph cuts using quantum algorithms: (1) based on quantum SDP solvers [11, 12], one can obtain the aforementioned Goemans-Williamson solution with a quadratic speedup and (2) using Gibbs sampling methods, it was shown recently in Ref. [13] that the Goemans-Williamson approximation guarantee can be obtained with exponential speedup, provided that the input problem obeys several constraints which are stronger than assuming that the graph G is sparse (as is assumed here).

Chapters 6 and 7 are both concerned with the problem of estimating spectral properties of many-body Hamiltonians from time-domain data. Let us focus the discussion on Chapter 7. One way to obtain the required time-domain data – namely, a time-evolution signal of a state under a many-body Hamiltonian – is to run quantum phase estimation circuits on a quantum computer. A costly step in these circuits is implementing the time-evolution unitary *controlled* on a single qubit. Because this unitary is controlled, the resulting circuit is very deep – too deep to run on early quantum computers. Removing the control makes the circuits shallower (as discussed in detail in Chapter 7), but it also discards phase information from the time-domain data. Our solution is to collect time-dynamics data from many different initial states, with all of this data obtained by running such shallower control-free circuits. This enlarged set of phaseless time-domain data can then be used to approximately recover the phases of the target time-evolution signal, from which spectral features can be extracted. In other words, rather than running deep quantum circuits, we run shallower circuits many more times to achieve the same goal. We demonstrate the practical feasibility of this method by estimating spectra of (large-scale) Fermi–Hubbard model instances. We do not, however, explicitly bound the total number of samples (i.e., the total number of shallow-circuit runs) required to recover a spectrum; this is addressed in subsequent work [14], which employs a slightly different strategy. Note that this aspect is especially important because trading circuit depth for number of repetitions is not completely free of charge: one must be able to perform a very large number of circuit repetitions consistently – i.e., without substantial

[1] Note that running the (classical) GW algorithm can become prohibitive (despite being asymptotically efficient) for practically relevant problem sizes, even for sparse graphs.

drift in the device parameters over time, which can be in principle be a limiting factor.

References

[1] S. Bravyi. 'Efficient algorithm for a quantum analogue of 2-SAT'. In: *Contemporary Mathematics*. Vol. 536. American Mathematical Society, 2011. eprint: quant-ph/0602108.

[2] M. B. Hastings and R. O'Donnell. 'Optimizing Strongly Interacting Fermionic Hamiltonians'. In: *Proc. of STOC 2022*. New York, NY, USA: ACM, 2022, pp. 776–789. DOI: https://doi.org/10.1145/3519935.3519960.

[3] A. Haldar, O. Tavakol and T. Scaffidi. 'Variational wave functions for Sachdev-Ye-Kitaev models'. In: *Physical Review Research* 3.2 (2021). DOI: 10.1103/physrevresearch.3.023020.

[4] J. Bardeen, L. N. Cooper and J. R. Schrieffer. 'Theory of Superconductivity'. In: *Physical Review* 108.5 (Dec. 1957), pp. 1175–1204. DOI: 10.1103/PhysRev.108.1175.

[5] J. Bausch, T. S. Cubitt and J. D. Watson. 'Uncomputability of phase diagrams'. In: *Nature Communications* 12.1 (2021), p. 452. DOI: 10.1038/s41467-020-20504-6.

[6] Š. Šmíd, R. Meister, M. Berta and R. Bondesan. 'Polynomial-time quantum Gibbs sampling for the weak and strong coupling regime of the Fermi-Hubbard model at any temperature'. In: *Nature Communications* 16.1 (2025), p. 10736. DOI: 10.1038/s41467-025-65765-1.

[7] R. Babbush, D. W. Berry, R. Kothari, R. D. Somma and N. Wiebe. 'Exponential Quantum Speedup in Simulating Coupled Classical Oscillators'. In: *Phys. Rev. X* 13 (4 Dec. 2023), p. 041041. DOI: 10.1103/PhysRevX.13.041041.

[8] M. X. Goemans and D. P. Williamson. 'Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming'. In: *J. ACM* 42.6 (Nov. 1995), pp. 1115–1145. ISSN: 0004-5411. DOI: 10.1145/227683.227684.

[9] S. P. Jordan, N. Shutty, M. Wootters, A. Zalcman, A. Schmidhuber, R. King, S. V. Isakov, T. Khattar and R. Babbush. 'Optimization by decoded quantum interferometry'. In: *Nature* 646.8086 (2025), pp. 831–836. DOI: 10.1038/s41586-025-09527-5.

[10] A. Gilyén, Y. Su, G. H. Low and N. Wiebe. 'Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics'. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. STOC '19. ACM, June 2019. DOI: 10.1145/3313276.3316366.

[11] F. G. Brandao and K. M. Svore. ' Quantum Speed-Ups for Solving Semidefinite Programs '. In: *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. Los Alamitos, CA, USA: IEEE Computer Society, Oct. 2017, pp. 415–426. DOI: 10.1109/FOCS.2017.45.

[12] J. van Apeldoorn, A. Gilyén, S. Gribling and R. de Wolf. ‘Quantum SDP-Solvers: Better upper and lower bounds’. In: *Quantum* 4 (Feb. 2020), p. 230. ISSN: 2521-327X. DOI: 10.22331/q-2020-02-14-230.

[13] H. Yuan, D. S. França, I. Luchnikov, E. Tiunov, T. Haug and L. Aolita. *Exponential Speed-ups for Structured Goemans-Williamson relaxations via Quantum Gibbs States and Pauli Sparsity*. 2025. arXiv: 2510.08292 [quant-ph]. URL: https://arxiv.org/abs/2510.08292.

[14] B. F. Schiffer, D. S. Wild, N. Maskara, M. D. Lukin and J. I. Cirac. ‘Hardware-Efficient Quantum Phase Estimation via Local Control’. In: *PRX Quantum* 6 (4 Dec. 2025), p. 040348. DOI: 10.1103/zlp8-rz9n.

Acknowledgements

First and foremost, I would like to thank my advisor, Barbara Terhal, for all her kindness and guidance. You have been my supervisor for almost all of my academic life, and I am grateful for everything that you have taught me and for all the opportunities that you have given me. You have a particularly selfless way of supervising your students and you provide an environment that allows us to grow, both scientifically and personally. Your approach to science and life is very inspiring, and I hope to have inherited some of it for the future.

I also wish to thank Fabian Hassler for many interesting discussions especially in the earlier stages of my PhD, and for providing me with a second academic home in Aachen.

I am deeply grateful to Yaroslav Herasymenko. Thank you for all the science we have done together. You are a very pleasant person to be around and I look forward to working together in the future. If I'm ever half as good a physicist as you are, I would be very happy.

I also wish to thank Jonas Helsen for many fun collaborations. Your enthusiasm and positivity are infectious, making it very enjoyable to work with you. I appreciate all the times when you were willing to give your sharp thoughts and feedback on projects, even if you were not directly involved.

Beyond those already mentioned, I have been fortunate enough to have collaborated with many inspiring people during my PhD and prior to it. I especially wish to thank Stephanie Wehner and Miriam Blaauboer at Delft, Tom O'Brien and Alicja Dutkiewicz at Leiden, Jan Behrends and Benjamin Béri at Cambridge, and Laura Clinton, Raul Garcia-Patron, Stasja Stanisic, Ashley Montanaro, Toby Cubitt and Joel Klassen at Phasecraft.

Thank you to all the people that have been part of the Terhal group throughout the years. A special thank you to my office mates Mac Shaw and Marc Serra-Peralta for all the conversations on linguistics and sports, and for the endless opportunities to make a fool out of myself by asking dumb questions about quantum error correction. I also wish to thank the students that I supervised, Daan Lenterman and Koen Eggen. I have learned a lot from you both.

I appreciate the friends around me for their support and for occasional distractions from science. I am especially grateful for those that I have known from a very young age; Bram, Leopold, Ramon, Cid, Ingo and Jesper.

I wish to thank my parents, Alexander and Marion, and my siblings, Sophie and Maurits, for their endless support. Papa, you have raised us to dream big and to pursue those dreams. Mama, you have taught us (by example) to be positive and optimistic, especially in times of hardship. You all inspire me, each in your own way, and I could not have asked for a more loving family to grow up in.

Finally, I would like to thank Marla, without whom this thesis would not have been finished. I am grateful for all your warmth and support, and for the endless laughs that we share. You are a big inspiration to me, and I can't wait to see the things that you will achieve next.

Curriculum Vitæ

Maarten Eduard Hubertus Marie Stroeks

I was born in Heerlen in 1997 and grew up in Broekhem (Limburg). There, I attended primary school and later I attended secondary school in Meerssen.

Subsequently, I studied physics at TU Delft, receiving my bachelor of science in 2018 and my master of science in 2020. Afterwards, I obtained a master of advanced study (part III in physics) from the University of Cambridge in 2021.

After returning to Delft, I started my PhD in Barbara Terhal's group working on quantum and classical algorithms for problems in quantum many-body physics. During my PhD, I was an intern at Phasecraft, where I worked on designing quantum algorithms for near-term quantum computers. In the summer of 2025, I was a visiting graduate student at the Simons Institute for the Theory of Computing at UC Berkeley. My work was presented at several highly-selective conferences (including QIP, TQC and QCTiP) and workshops in the Netherlands, Germany, Belgium, Austria, United Kingdom, United States and Canada.

List of Publications

1. **Quantum phase estimation without controlled unitaries.** Laura Clinton, Toby S. Cubitt, Raul Garcia-Patron, Ashley Montanaro, Stasja Stanisic & Maarten Stroeks. Published in *PRX Quantum*, 2026. url: https://doi.org/10.1103/7qcr-znl2.

   [Chapter 7]

2. **Complexity of Fermionic 2-SAT.** Maarten Stroeks & Barbara M. Terhal. Published in *Quantum*, 2025. url: https://doi.org/10.22331/q-2025-10-31-1900.

   [Chapter 2]

3. **Solving free fermion problems on a quantum computer.** Maarten Stroeks, Daan Lenterman, Barbara M. Terhal & Yaroslav Herasymenko. Published in *Physical Review Research*, 2025. url: https://doi.org/10.1103/zmwm-gdmw.

   [Chapter 5]

4. **Optimizing fermionic Hamiltonians with classical interactions.** Maarten Stroeks, Barbara M. Terhal & Yaroslav Herasymenko. Pre-print url: https://arxiv.org/abs/2510.02122, 2025.

   [Chapter 3]

5. **Optimizing sparse fermionic Hamiltonians.** Yaroslav Herasymenko, Maarten Stroeks, Jonas Helsen & Barbara M. Terhal. Published in *Quantum*, 2023. url: https://doi.org/10.22331/q-2023-08-10-1081.

   [Chapter 4]

6. **Spectral estimation for Hamiltonians: a comparison between classical imaginary-time evolution and quantum real-time evolution.** Maarten Stroeks, Jonas Helsen & Barbara M. Terhal. Published in *New Journal of Physics*, 2022. url: https://doi.org/10.1088/1367-2630/ac919c.

   [Chapter 6]