Advancing Applications of Quantum Computers in Quantum Simulation, Optimization, Learning, and Topological Data Analysis

Thesis by
Robbie King

In Partial Fulfillment of the Requirements for the
Degree of
PhD

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2025
Defended April 29, 2025

Robbie King
ORCID: 0000-0002-4174-0801

ACKNOWLEDGEMENTS

I am deeply grateful to my advisor, Thomas Vidick, for his kindness, patience, and unwavering belief in me. At Caltech, Leonard Schulman has been a constant source of support, always making me feel at home. I am grateful to John Preskill for his visionary leadership and for fostering an environment at Caltech where students can truly realize their potential. I owe a special thanks to Ryan Babbush, who has given me invaluable guidance and opportunities which significantly shaped my research goals.

I draw continuous inspiration from my colleagues. Many of them have left a lasting impact on me: Amira Abbas, Atul Arora, Eric Anschuetz, Joao Basso, Thiago Bergamaschi, Matthias Caro, Ulysse Chabaud, Anthony Chen, Chris Chen, Andrea Coladangelo, Abhinav Deshpande, Matthew Ding, Dar Gilboa, Jonathan Gross, Bailey Gu, Casper Gyurik, Dominik Hangleiter, Matt Harrigan, Oscar Higgott, Robert Huang, Bill Huggins, Alex Jahn, Jiaqing Jiang, Bobak Kiani, Tanuj Khattar, Tamara Kohler, Laura Lewis, Saeed Mehraban, Tony Metger, Chinmay Nirkhe, Quynh Nguyen, Chris Pattison, Alex Poremba, Daniel Ranard, Chaithanya Rayudu, Sascha Schmidhuber, Tommy Schuster, Noah Shutty, Joe Slote, Jun Takahashi, Eugene Tang, Kevin Thompson, Yu Tong, Ben Villalonga, Agi Villanyi, James Watson, Tina Zhang, Leo Zhou, and Alex Zlokapa.

My journey into this field began thanks to Richard Jozsa and Sergii Strelchuk, who introduced me to quantum computing as an undergraduate at Cambridge. Since then, I have been fortunate to learn from many more exceptional senior researchers: Sergio Boixo, Toby Cubitt, Vedran Dunjko, Eddie Farhi, David Gosset, Zhang Jiang, Stephen Jordan, Robin Kothari, Guang Hao Low, Jarrod McClean, Tom O'Brien, Ojas Parekh, Pedram Roushan, Nick Rubin, Rolando Somma, Nathan Wiebe, and John Wright.

Finally, I am endlessly grateful for my parents Chieko and Marshall, whose love and sacrifices gave me the best possible start in life. I have also had incredible support from my brother Alex, my old friends back home in the UK, and the new friends I have made in California.

ABSTRACT

This thesis investigates novel directions for harnessing the potential of quantum computers in future applications. It is structured into three sections.

Quantum Simulation. We address two key questions: what systems exhibit quantum advantage in predicting ground state properties, and how can we reduce the cost of quantum simulations? For the former, we find that strongly interacting fermionic systems have promising characteristics for quantum advantage. For the latter, we develop an improved method for compiling block encodings using sum-of-squares optimization.

Learning with Entangled Measurements. We explore the benefits of leveraging entangled measurements on quantum states stored in quantum memory. These learning algorithms can be applied to the readout stage of quantum simulations, or to learn from quantum data from nature.

Topological Data Analysis. Using complexity-theoretic insights, we demonstrate that certain problems in topological data analysis possess a quantum mechanical structure, suggesting opportunities for quantum algorithms in this area.

PUBLISHED CONTENT AND CONTRIBUTIONS

[Kin+25a] Robbie King et al. “Quantum simulation with sum-of-squares spectral amplification”. In: *arXiv preprint arXiv:2505.01528* (2025). URL: https://arxiv.org/abs/2505.01528.
I conceived the project, established the key method, and contributed to the writing.

[Kin+25b] Robbie King et al. “Triply efficient shadow tomography”. In: *Proceedings of the 2025 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2025. DOI: 10.1137/1.9781611978322.27.
I conceived the project, contributed to the main results, and contributed to the writing.

[KK24] Robbie King and Tamara Kohler. “Gapped Clique Homology on Weighted Graphs is QMA 1-Hard and Contained in QMA”. In: *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2024, pp. 493–504. DOI: 10.1109/FOCS61266.2024.00039.
I conceived the project, established the key method, and was the main writer.

[KWM24] Robbie King, Kianna Wan, and Jarrod R. McClean. “Exponential Learning Advantages with Conjugate States and Minimal Quantum Memory”. In: *PRX Quantum* 5 (4 Oct. 2024), p. 040301. DOI: 10.1103/PRXQuantum.5.040301.
I conceived the project, proved the main results, and was the main writer.

[Kin+24] Robbie King et al. “Strongly interacting fermions are non-trivial yet non-glassy”. In: *arXiv preprint arXiv:2408.15699* (2024). URL: https://arxiv.org/abs/2408.15699.
I conceived the project, proved the key result, and contributed to the writing.

TABLE OF CONTENTS

Acknowledgements . . . iii
Abstract . . . iv
Published Content and Contributions . . . v
Table of Contents . . . v
Chapter I: Overview . . . 1
Chapter II: Strongly Interacting Fermions . . . 4
2.1 Commutation index and Lovász theta function . . . 12
2.2 Circuit lower bound . . . 25
2.3 Annealed approximation . . . 29
Chapter III: Sum of Squares Spectral Amplification . . . 40
3.1 Quantum simulation using block-encodings and linear combination of unitaries . . . 50
3.2 Quantum simulation by spectral amplification . . . 53
3.3 SOS optimization and example SOSSA block-encoding . . . 74
3.4 Application to SYK model . . . 77
Chapter IV: Learning Fermionic Observables . . . 83
4.1 Commutation, learning, and coloring . . . 97
4.2 Coloring commutation graphs with bounded clique number . . . 104
4.3 Learning local fermionic operators . . . 107
Chapter V: Learning Bosonic Observables . . . 115
5.1 Background on bosonic systems . . . 122
5.2 Learning displacement amplitudes . . . 125
5.3 Sample complexity lower bounds . . . 131
Chapter VI: Topological Data Analysis . . . 143
6.1 Preliminaries . . . 160
6.2 Hamiltonian to homology gadgets . . . 175
6.3 Spectral sequences . . . 179
6.4 Combining gadgets . . . 199
6.5 Postponed proofs . . . 208
Bibliography . . . 215

Chapter 1

OVERVIEW

Quantum computing represents the most powerful form of computation according to our current understanding of physics, with quantum technologies offering unprecedented control over fundamental reality—the storage and manipulation of quantum information itself. While significant progress is being made toward error-corrected quantum computers [23; Blu+24], a clear understanding of their most concrete and impactful future applications remains elusive [Dal+23a]. If quantum computers are not widespread within the next few decades, the reason may well be due to lack of incentive rather than an engineering obstacle. That said, there are reasons to be optimistic. Quantum computing holds the potential to unlock groundbreaking discoveries in physics and chemistry, ultimately leading to transformative technologies in fields such as drug development and material science. Quantum computing also has the potential to provide exponential computational speedups for certain problems, which could revolutionize processes at the industrial scale.

Simulation of quantum chemistry is one of the most widely heralded applications of quantum computing. This is because quantum chemistry governs the properties of drugs and materials, yet is hard to accurately model using classical algorithms. Two outstanding challenges remain for quantum algorithms in quantum chemistry. Firstly, we would like a better understanding of which systems are expected to exhibit the largest quantum advantages, for example in predicting groundstate properties. Can we characterize such systems, and understand *why* they are classically hard yet quantumly easy? Secondly, quantum algorithms for quantum chemistry remain expensive, and it is important to reduce the costs as much as possible.

We tackle the first problem in Chapter 2, where we study the SYK model as a model for strongly interacting fermions. Here we were able to establish two new results [Kin+24]. Firstly, low energy states of the SYK model must have large polynomial circuit depth; that is, they are very entangled. Secondly, the SYK model has an annealed free energy at inverse polynomial temperatures, and there is no glassy phase. The first result suggests that low energy physics of SYK is classically hard, whilst the second result suggests that it is quantumly easy, since there is no obstruction to quantum thermalization. Remarkably, both results follow from the

	Fermionic optimization and learning	
Quantum simulation	**Chapter I** **Strongly Interacting Fermions** – Strongly interacting fermionic systems have promising properties for quantum advantage. – Eigenstates are highly complex. – Annealed and non-glassy at low temperatures.	**Chapter II** **Sum of Squares Spectral Amplification** – Order of magnitude improvement over state of art quantum chemistry algorithms. – Preprocess Hamiltonian using sum of squares optimization and amplify low energy spectrum. – Asymptotic improvements in SYK toy model.
Learning with entangled measurements	**Chapter III** **Learning Fermionic Observables** – Aim to measure fermionic k-RDM. – Necessary in readout stage of quantum chemistry algorithms. – Exponential sample complexity improvements using entangled measurements on two copies of the state.	**Chapter IV** **Learning Bosonic Observables** – Aim to measure expectation values of bosonic displacement operators. – Exponential sample complexity improvements using measurements on $\rho \otimes \rho^*$. – Proof that $\rho \otimes \rho^*$ is necessary.

Table 1.1: An illustration of the themes in Chapters I-IV

same quantity which we call the commutation index, which measures the extent to which the terms in the Hamiltonian are non-commuting. For quantum spin glasses the commutation index is constant, whilst for random fermionic interactions it decays polynomially with system size. This provides a hint that strongly interacting fermionic systems behave more like random matrices than like spin glasses and are a better target for quantum algorithms.

Next in Chapter 3 we develop a new method for compiling more efficient block encodings of quantum chemistry Hamiltonians for use in quantum algorithms. We refer to our method as *sum of squares spectrum amplification* (SOSSA), and it has already been used in Ref. [Low+25] to achieve an order of magnitude improvement in the total runtime over state of art quantum chemistry algorithms for large molecules of industrial interest, such as FeMo-co. The key innovation lies in performing spectral amplification on the low-energy spectrum of the Hamiltonian. This approach typically works only when the Hamiltonian is frustration-free. However, by harnessing sum-of-squares optimization, we can classically preprocess the Hamiltonian to express it in a low-frustration form.

In applications of quantum simulation, after preparing a low energy state or simulating some dynamics on the quantum computer, we would like to extract meaningful physical information from the final quantum state. This readout stage of the quantum algorithm can be casted as a learning problem. In Chapter 4, we continue the theme of fermionic systems from Chapter 2 and study learning algorithms for fermionic observables which exploit entanglement measurements on multiple copies of the

unknown state. Specifically, we resolve a key bottleneck in shadow tomography for fermionic operators by reducing the quantum memory requirement from many copies to just two copies [Kin+25]. We do this by relating measurement strategies to colorings of a certain graph which we call the *commutation graph*.

In Chapter 5, we apply similar learning techniques from Chapter 4 to bosonic observables. In the bosonic setting, we uncover a new type of advantage where entangled measurements involving the complex conjugate of an unknown state provides exponential sample complexity advantages [KWM24]. We proved that the learning task is information-theoretically impossible without access to the complex conjugate state.

Chapter 6 is a final standalone chapter which looks at an application of quantum computers outside of quantum simulation. Can Hamiltonian simulation algorithms be applied to domains outside of the simulation of quantum systems themselves? One example is provided by the task of *topological data analysis* (TDA), for which quantum algorithms have been developed. The primary feature of TDA enabling quantum speedups is the ability to encode topological invariants in the ground space of an exponentially large sparse matrix, known as the combinatorial Laplacian [LGZ16; Ber+24a]. In Chapter 6 we show that deciding whether a topological space has a high-dimensional hole is QMA1-hard and contained in QMA [KK24]. This tells us that there is quantum structure in the TDA problem and provides evidence that the quantum TDA algorithm cannot be dequantized. Our proof technique crucially used a tool from algebraic topology known as spectral sequences, which allowed us to perform a version of perturbation theory on the Laplacian.

Chapter 2

STRONGLY INTERACTING FERMIONS

Simulating ground and thermal state properties of quantum systems is a key application of future quantum computers [Fey82; Llo96; McA+20; Lee+21a; Bur+21a; Bab+18b; Cha+20]. Nevertheless, the search for particular, favorable instances that are quantumly easy and classically hard is not clear-cut [Lee+22]. A challenge is that current quantum computers are limited in quality and size, requiring the community to rely on theoretical arguments to give computational separations. However, the ground states for standard few-body quantum spin models can be QMA-hard (as classical spin models are NP-hard) in the worst case [Kit+02; Aha+09; GI09]; in the average case, random classical and quantum spin models exhibit glassy physics where computational hardness may arise [Gam21; BS20]. To give an efficient quantum algorithm for low-temperature states, one must carefully avoid these instances.

Most chemical and condensed matter systems involve *fermionic* degrees of freedom, not only spins. Of particular importance in quantum chemistry is the *strongly interacting* regime, where Gaussian states do not give good approximations to the ground state and the Hartree–Fock method fails [SO12]. This has been proposed as a promising regime in which to apply quantum computers to achieve quantum advantage [McA+20]. The Sachdev–Ye–Kitaev (SYK) Hamiltonian provides a natural model for strongly interacting fermions [SY93; Kit15b; Kit15a]. As a counterpart to random spins, it is a random Hamiltonian consisting of all-to-all q-body Majorana fermions:

$$H_q^{\mathrm{SYK}} := i^{q/2}\binom{n}{q}^{-1/2} \sum_{j_1<\cdots<j_q} g_{j_1\ldots j_q}\gamma_{j_1}\cdots\gamma_{j_q}, \tag{2.1}$$

where q is assumed to be even, and the $g_{j_1\ldots j_q}$ are i.i.d. standard Gaussian random variables. The γ_j are the *Majorana operators*, which satisfy commutation relations $\gamma_i\gamma_j+\gamma_j\gamma_i = 2\delta_{ij}$. They arise from the fermionic creation and annihilation operators $a_j, a_j^\dagger$ by

$$\gamma_{2j-1} = a_j^\dagger + a_j \quad , \quad \gamma_{2j} = i(a_j^\dagger - a_j) \tag{2.2}$$

While the 4-body fermionic ground state problem can be just as hard as spin models in the worst case (NP-hard) [LCV07], average-case fermionic systems appear to

have qualitatively different physics and perhaps computational complexity than spin systems [HO22; MS16; BS20]. Extensive heuristic calculations (such as large-N expansions) together with numerical evidence indicate that the SYK model resembles a thermalizing chaotic system, not a frozen spin-glass as occurs with few-body quantum spin systems [BS20; Fac+19]. However, rigorous proofs that go beyond the physical arguments have been very limited [HO22; FTW19].

In this chapter, we study the strongly interacting SYK model and give quantitative evidence that random, all-to-all connected fermionic systems have a *classically non-trivial* yet *non-glassy* thermal state at constant temperatures. In contrast, these two properties are false for disordered spin systems [BS20; Bra+19]. Remarkably, the proofs of both main results rely on the same quantity, the *commutation index* [Kin+25]. To bound the commutation index of fermionic operators, we analyze the *Lovasz theta-function* [Knu93] of a certain graph encoding the fermionic commutation relations.

This quantity pinpoints a crucial and often overlooked distinction between fermionic and spin Hamiltonians: low-degree fermionic monomials have a very different commutation structure than low-weight Pauli operators. The commutation index captures this difference, quantifying the fundamental distinction in the physics of local spin systems and local fermionic systems. This disparity, we argue, is the origin of a potential quantum advantage in simulating strongly interacting fermionic systems. Although the SYK model is our primary example, we also show that versions of our results—depending only on the commutation index—apply to *all* models with i.i.d. Gaussian couplings. This highlights the surprising importance of quantum uncertainty in governing both glassiness and the entanglement of low-energy states in disordered many-body systems.

More precisely, we first show that all low-energy states (including constant-temperature thermal states) of the SYK model have high circuit complexity ('classically non-trivial')[1]. The following theorem is phrased in terms of states that maximize the energy $\mathrm{Tr}(\rho H_q^{\mathrm{SYK}})$. This is equivalent to states that minimize the energy, since the distribution of H_q^{SYK} is the same as that of $-H_q^{\mathrm{SYK}}$.

Theorem 2.1 (Low energy states are classically nontrivial)**.** *Consider the degree-q SYK model H_q^{SYK}. With high probability, the maximum energy is $\lambda_{\max}(H_q^{SYK}) \geq$*

[1] We here consider disordered models where the distribution of H is identical to that of $-H$, so the ground state energy is equivalent to the maximal energy.

$\Omega_q(\sqrt{n})$*, yet any state* ρ *such that*

$$\mathrm{Tr}\left(\rho H_q^{SYK}\right) \geq t\sqrt{n} \tag{2.3}$$

has circuit complexity

$$\tilde{\Omega}_q(n^{(q/2)+1}t^2). \tag{2.4}$$

The Ω_q *notations assume a fixed* q *and growing* n*.*

That is, low-energy states of the SYK model are highly entangled and require many parameters to describe; simple classical ansatzes, such as Gaussian states, must fail. In comparison, local quantum spin systems are known to have efficiently computable product state approximations to the ground state [Bra+19] and thus, in this sense, have 'trivial' states that achieve a constant-factor approximation of the ground state energy.

Second, we show that the quenched free energy of the SYK model agrees with the annealed free energy even at very low temperatures ('non-glassy'), formalizing and strengthening previous results of this nature [Gur17; BS20; Fac+19; GPS00].[2] Here, the free energy is normalized such that $\beta = O(1)$ corresponds to constant physical temperature.

Theorem 2.2 (Annealed at low temperatures)**.** *Consider the partition function of the degree-*q *SYK model* $Z_\beta := \mathrm{Tr}\exp(-\beta\sqrt{n}H_q^{SYK})$*. Then, we have:*

$$\frac{\mathbb{E}\log Z_\beta}{n} \leq \frac{\log \mathbb{E}Z_\beta}{n} \leq \frac{\mathbb{E}\log Z_\beta}{n} + O_q(\beta^2 n^{-q/2}). \tag{2.5}$$

The O_q *notations assume a fixed* q *and growing* n*.*

The quantitative agreement of the two free energies at (inverse-polynomially) low temperatures strikes a stark contrast with disordered spin systems: the SYK model does not experience a 'glass' phase transition in the sense of quenched-vs.-annealed free energy. For classical spin Hamiltonians, it is known that the annealed free energy $n^{-1}\mathbb{E}\log Z_\beta$ fails to agree with the quenched free energy $n^{-1}\log \mathbb{E}Z_\beta$ at constant temperatures where the Hamiltonian is in its glassy phase and algorithmic hardness arises; disordered quantum spin systems undergo a similar transition at constant temperature [BS20]. The lack of a glass transition for the SYK model

[2] In particular, Ref. [BS20] showed that the SYK model is *consistent* with an annealed approximation, and here we prove that the annealed approximation *holds*.

suggests that there may be no algorithmic obstructions to preparing low-temperature states of the model on a quantum computer, but we do not prove this claim. We leave finding such an efficient quantum algorithm for future work.

Although the SYK model is our primary example, our results apply more generally to a wider range of models. We will see shortly that the circuit lower bound in Theorem 2.1 and the bound on the quenched free energy in Theorem 2.2 depend only on the value of a single quantity which we call the *commutation index*. This quantity characterizes the quantum uncertainty in the model and provides a new tool to probe the chaotic properties of many-body systems.

Background and related work. The SYK model is a canonical instance of a chaotic Hamiltonian [SY93; Kit15b; Kit15a] with related models studied as far back as [FW70; BF71]. For even $q = o(\sqrt{n})$, the SYK model has a Gaussian spectrum [FTW19] and heuristics from physics indicate that the expected maximum energy of the SYK model scales as $\frac{\sqrt{2n}}{q}$ for even q [GJV18; GV16; HO22]. However, the only rigorous result we are aware of with explicit constants is an upper bound of $\sqrt{(\log 2)n}$ [FTW19]. Though Gaussian state approximation algorithms exist for fermionic systems [Bra+19; Her+23b], it is known that for the SYK model with $q \geq 4$, Gaussian states cannot achieve constant factor approximations to the maximum energy [HTS21a]. Separate from the SYK model, so-called no low-energy trivial states (NLTS) theorems rule out constant factor approximations to ground energies with low depth circuits in worst-case settings [ABN23; Her+23a]. For random nonlocal Hamiltonians, [Che+23] shows a circuit lower bound for sparse, sampled Pauli models using a similar technique as in the proof of Theorem 2.1, where the commutation index is much more straightforward to calculate.

The commutation index has connections to other areas of quantum information theory and Hamiltonian complexity. In [GHG23; XSW23], the commutation index (there termed the *generalized radius*) is used to study generalized Heisenberg uncertainty relations. Related to our work, [HO22] use the commutation index to analyze the performance of sum-of-squares relaxations of the SYK model and prove the $q = 4$ instance of Theorem 2.3, giving as well an algorithm verifying $\Omega(\sqrt{n})$ energy for $q = 4$. [AGK24] demonstrates that product states maximize the energy variance for random quantum spin Hamiltonians. Finally, the commutation index appears in quantum learning theory, where it provides a sample-complexity lower bound on how many copies of the state are required to learn the expectation values of a set of operators via shadow tomography [Che+22; Kin+25].

Commutation index

We define the *commutation index* $\Delta(\mathcal{S})$ of a given set of operators $\mathcal{S} = \{A_1, \ldots, A_m\}$ as follows.

Definition 2.1. *For a set $\mathcal{S}$ of Hermitian operators, define their* commutation index *by*

$$\Delta(\mathcal{S}) = \sup_{|\psi\rangle} \mathbb{E}_{A\in\mathcal{S}} \langle\psi|A|\psi\rangle^2. \tag{2.6}$$

When all $\|A_i\| \leq 1$ the commutation index takes values $0 < \Delta(\mathcal{S}) \leq 1$. Roughly, a more 'commuting' set of observables $\mathcal{S}$ gives a larger value of $\Delta(\mathcal{S})$. For example, if the operators are all mutually commuting and satisfy $A_i^2 = 1$, choosing $|\psi\rangle$ to be a simultaneous eigenstate gives $\Delta(\mathcal{S}) = 1$. The commutation index is related to the minimum uncertainty of the operators A_i, assuming they satisfy $A_i^2 = 1$:

$$\inf_{|\psi\rangle} \frac{1}{m} \sum_{i=1}^{m} \left(\Delta_{|\psi\rangle} A_i\right)^2 = 1 - \Delta(\mathcal{S}), \tag{2.7}$$

where $\left(\Delta_{|\psi\rangle} A_i\right)^2$ is the variance of A_i in the state $|\psi\rangle$:

$$\left(\Delta_{|\psi\rangle} A_i\right)^2 = \langle\psi| A_i^2 |\psi\rangle - \langle\psi| A_i |\psi\rangle^2 = 1 - \langle\psi| A_i |\psi\rangle^2 . \tag{2.8}$$

The commutation index has already been used in the context of quantum learning and state tomography [Kin+25]. In this paper, we will see that the commutation index has strong implications for the physics of the model $H = m^{-1/2} \sum_{i=1}^{m} g_i A_i$ with Gaussian coefficients g_i, and it will be the key quantity in our two main results. When the commutation index is small and the operators have high quantum uncertainty, two properties follow: the model has highly entangled low energy states, and the model is annealed at low temperatures.

Crucially, the commutation index controls the sensitivity of many physical properties when varying the couplings of the model. For instance, the norm of the energy gradient of a given state with respect to the disorder is bounded by:

$$\left\|\nabla_{\vec{g}} \langle\phi|H|\phi\rangle\right\|_2^2 = \frac{1}{m} \sum_{i=1}^{m} \langle\phi|A_i|\phi\rangle^2 \leq \Delta(\mathcal{S}). \tag{2.9}$$

Our key observation is that the commutation index of the set $\mathcal{S}_q^n$ of $\binom{n}{q}$ degree-q Majorana operators is very small:

Set $\mathcal{S}$	**Commutation index $\Delta(\mathcal{S})$**
Commuting	1
k-local Paulis	3^{-k} (Theorem 2.9)
Degree-q Majoranas	$\Theta_q(n^{-q/2})$ (Theorem 2.3)
All Paulis	2^{-n} [Che+22, Lemma 5.8]

Table 2.1: The commutation index $\Delta(\mathcal{S})$ characterizes how non-commuting a set $\mathcal{S}$ of operators is. The commutation index reveals a key distinction between local spin operators and local fermionic operators: in the fermionic case, the commutation index decays polynomially with system size, while it is constant in the case of spins. The Θ_q notation assumes a fixed q and growing n.

Theorem 2.3. *Let $\mathcal{S}_q^n$ be the set of degree-q Majorana operators on n fermionic modes. Then for any constant, even q:*

$$\Delta(\mathcal{S}_q^n) = \Theta_q(n^{-q/2}). \tag{2.10}$$

The decay with system size n is unique to the fermionic setting—for local Pauli operators, $\Delta(\mathcal{S})$ is constant with respect to n (see Table 2.1). This behaviour was first conjectured in [HO22] to our knowledge, and we establish the conjecture—including the setting when q scales with n—in Section 2.1.

The proof of Theorem 2.3 involves constructing the *commutation graph* $G(\mathcal{S})$ whose vertices correspond to operators $A_i \in \mathcal{S}$ with edges between operators if and only if they anti-commute. The commutation index can be upper bounded by $\Delta \leq \vartheta(G(\mathcal{S}))/|G|$, where $\vartheta(G(\mathcal{S}))$ is the so-called *Lovász theta function* of the commutation graph. The Lovász theta function can be efficiently computed via a semi-definite program [Knu93]. For the SYK Hamiltonian, $G(\mathcal{S}_q^n)$ is the graph of a certain Johnson association scheme [Del73].

In the course of writing our results we became aware of Ref. [Lin24], which also establishes the necessary results on the Lovász theta function of Johnson association schemes. Our results use different proof techniques and determine the explicit q-dependence of the constant in Equation (2.10), which was not derived in [Lin24].

Circuit lower bound

An almost direct consequence of a decaying commutation index is a lower bound on the complexity of any ansatz in constructing near-ground states, including the ansatz of quantum circuits.

Theorem 2.4 (Low energy states are classically non-trivial)**.** *Consider the random Hamiltonian $H = m^{-1/2} \sum_{i=1}^{m} g_i A_i$ with i.i.d. Gaussian coefficients g_i. Let Δ be the*

Ansatz	Circuit complexity*
Quantum circuit with G gates	$G \geq \tilde{\Omega}_q(n^{q/2+1}t^2)$
MPS with bond dimension χ	$\chi \geq \Omega_q(n^{q/4+1/2}t)$
Neural network with W parameters	$W \geq \Omega_q(n^{q/2+1}t^2)$

*min. complexity to achieve energy $t\lambda_{\max}(H_q^{\mathrm{SYK}})$ w.h.p.

Table 2.2: To achieve energy scaling as $t\lambda_{\max}(H_q^{\mathrm{SYK}})$ for the SYK Hamiltonian with high probability, ansatz complexity (e.g., circuit depth) must scale polynomially with n. See Section 2.2 for proofs. The Ω_q notations assume a fixed q and growing n.

commutation index of $\mathcal{S} = \{A_1, \ldots, A_m\}$ and let C be a set of fixed quantum states. Then the probability that C contains a low-energy state is exponentially small:

$$\mathbb{P}\left[\max_{|\psi\rangle \in C} \langle\psi| H |\psi\rangle \geq t\right] \leq \exp\left(\log|C| - \frac{t^2}{2\Delta}\right). \tag{2.11}$$

This theorem follows from a concentration argument and a union bound. The commutation index $\Delta(\{A_i\}_{i=1}^m)$ characterizes the maximum variance of the energy $\langle\psi|H|\psi\rangle$ for an arbitrary fixed state $|\psi\rangle$. Standard concentration bounds then imply that the probability a state $|\psi\rangle$ has energy t is bounded as $\exp(-\Omega(t^2/\Delta))$. This concentration is so strong that one can bound the maximum energy over extremely large sets of states (or ϵ-nets of infinite sets) $\mathcal{S}$ via a simple union bound argument with high probability over the disorder. In particular, we obtain a lower bound $|\mathcal{S}| = \exp(\Omega(t^2/\Delta))$ on the cardinality of the class of ansatzes needed to achieve a given energy t.

Specializing to the SYK model via Theorem 2.3, we summarize the implications of this result for various classes of states $\mathcal{S}$ in Table 2.2. For instance, we show that all states that achieve a constant (i.e., $t = \Theta(1)$) approximation ratio with the SYK ground state energy have a quantum circuit depth of $\Omega_q(n^{q/2})$. In contrast, product states give constant factor approximations to the ground state energy for any local spin Hamiltonian (see Section 2.2 for a short proof). Our argument also extends to classical ansatzes. For instance, tensor network methods require a bond dimension that grows polynomially with n to construct near-ground states [Sch11; Bañ23]. Similarly, popular methods based on neural quantum states [CT17; SH20; Sha+20; SSC22; NI21] need at least $\Omega(n^3)$ parameters to construct near-ground states for the standard $q = 4$ SYK model, implying a bounded depth fully connected network must have layer width that grows as $\Omega(n^{3/2})$.

Our circuit lower bound is related to the study of 'no low-energy trivial states' (NLTS) Hamiltonians, whose existence was conjectured in [FH13] and resolved

in [ABN23; Her+23a]. However, the settings are not strictly comparable: our instances are random (average-case), whereas NLTS is formalized for worst-case bounded interaction instances of Hamiltonians. The randomness allows us to prove stronger statements in two ways. First, our circuit lower bounds hold for states at *any* constant temperature, rather than for states below some energy threshold. Second, we can achieve arbitrary polynomial circuit depth lower bounds, whereas current constructions of NLTS only give a logarithmic depth lower bound. See Section 2.2 for more discussion.

Annealed approximation

The commutation index also has direct implications for the concentration of various physical properties of interest around their disordered expectation. One manifestation of this is in the relation between the *quenched* and *annealed* free energies:

$$\underbrace{\frac{1}{n}\mathbb{E}\log Z_\beta}_{\text{quenched}} \quad \leq \quad \underbrace{\frac{1}{n}\log \mathbb{E} Z_\beta}_{\text{annealed}}, \tag{2.12}$$

where Z_β is the partition function of the model $\sqrt{n}H$ at an inverse temperature β. The quenched free energy is physical but hard to calculate while the annealed free energy is much easier to calculate but nonphysical. The inequality above always holds due to Jensen's inequality, and the disagreement stems from fluctuations in $\log Z_\beta$ due to the disorder.

The quenched free energy assumes the disorder induced by the random couplings is fixed when averaging over thermal fluctuations; the annealed free energy treats these fluctuations on an equal footing. While the two quantities agree at high temperature, at low temperature the latter is incapable of accounting for frustration induced by the disorder of the random couplings which can induce a spin glass phase [Tal00; Par79]. Their disagreement is thus indicative of the presence of a spin glass phase. Motivated by this, our second main result bounds the difference in quenched and annealed free energies as a function of the temperature and the commutation index of the model.

Theorem 2.5 (Annealed at low temperatures)**.** *Consider the partition function* $Z_\beta := \operatorname{Tr}\exp(-\beta\sqrt{n}H)$ *of the random Hamiltonian* $H = m^{-1/2}\sum_{i=1}^m g_i A_i$ *with i.i.d. Gaussian coefficients* g_i*. Let* Δ *be the commutation index of* $\mathcal{S} = \{A_1, \ldots, A_m\}$*. Then the quenched and annealed free energies are bounded by:*

$$\frac{1}{n}\mathbb{E}\log Z_\beta \leq \frac{1}{n}\log \mathbb{E} Z_\beta \leq \frac{1}{n}\mathbb{E}\log Z_\beta + 4\beta^2\Delta. \tag{2.13}$$

Quantity f	**Rate** K
$\lambda_{\max}\left(H_q^{\mathrm{SYK}}\right)$	$\Omega_q\left(n^{q/2}\right)$
$\mathrm{Tr}\left(X\rho_\beta\right)$	$\Omega_q\left(\beta^{-2}n^{q/2-1}\right)$
$\mathrm{Tr}\left(H_q^{\mathrm{SYK}}\rho_\beta\right)$	$\Omega_q\left(\min\left(1,\beta^{-2}n^{-2}\right)n^{q/2}\right)^*$

*for t order of $\|H_q^{\mathrm{SYK}}\| = O(\sqrt{n})$

Table 2.3: Concentration bounds for functions f of the Hamiltonian around its mean, i.e., $\mathbb{P}\left[|f - \mathbb{E}\left[f\right]| \geq t\right] \leq 4\exp\left(-Kt^2\right)$. $\lambda_{\max}$ denotes the largest eigenvalue and X is an arbitrary bounded operator. ρ_β is the thermal state of $\sqrt{n}H$ at an inverse temperature β.

For the SYK model this directly implies Theorem 2.2. Informally, this bound is due to controlling the growth of the moment generating function of $\log\left(Z_\beta\right) - \mathbb{E}\left[\log\left(Z_\beta\right)\right]$ using the commutation index Δ. We formally prove Equation (2.13) in Section 2.3. We there also prove concentration bounds for observable expectations as well as two-point correlators, again following from bounding how sensitive these quantities are when varying the disorder. We summarize some of these results when applied to the SYK model in Table 2.3. We also emphasize that this general theorem applies to any model for which the commutation index is known. In the case of the k-body Pauli models, this indicates that the annealed approximation remains valid for β growing exponentially with the locality k as previously predicted [SW24].

2.1 Commutation index and Lovász theta function

Commutation index

In this section we introduce the *commutation index*, which quantifies the commutation structure of a set of operators. This allows us to study the commutation structure of local Majorana operators and how they differ from local Paulis.

The following result represents a kind of uncertainty principle, generalizing the familiar Bloch sphere constraint $\langle X\rangle^2 + \langle Y\rangle^2 + \langle Z\rangle^2 \leq 1$.

Lemma 2.6. *Let $A_1, \ldots, A_m$ be Hermitian operators which square to identity. If $A_1, \ldots, A_m$ pairwise anticommute, then for any state ρ*

$$\sum_j \mathrm{Tr}\left(A_j\rho\right)^2 \leq 1. \tag{2.14}$$

Versions of Lemma 2.6 appear in various papers, for example [Asa+16, Theorem 1]. Since the proof is simple, we reproduce it here.

Proof. Given ρ, let $a_j = \mathrm{Tr}\left(A_j\rho\right)$. We aim to show $\sum_j a_j^2 \leq 1$. Consider the observable

$$Q = \sum_j a_j A_j. \tag{2.15}$$

We will use the inequality $\mathrm{Tr}(Q^2\rho) - \mathrm{Tr}(Q\rho)^2 = \mathrm{Var}_\rho(Q) \geq 0$. Formally, this holds since ρ is positive semi-definite so $\mathrm{Tr}(O\rho) \geq 0$ for any positive semi-definite operator O and

$$\mathrm{Tr}\left(\left(Q - \mathrm{Tr}\left(Q\rho\right)\mathbb{1}\right)^2\rho\right) \geq 0 \implies \mathrm{Tr}\left(Q\rho\right)^2 \leq \mathrm{Tr}\left(Q^2\rho\right). \tag{2.16}$$

Due to anticommutativity of $A_1, \ldots, A_m$, we have

$$Q^2 = \sum_j a_j^2 A_j^2 + \sum_{j\neq l} a_j a_l A_j A_l = \sum_j a_j^2 \cdot \mathbb{1} + \frac{1}{2}\sum_{j\neq l} a_j a_l \{A_j, A_l\} = \sum_j a_j^2 \cdot \mathbb{1}. \tag{2.17}$$

Note $A_j^2 = \mathbb{1}$ since they are Hermitian unitaries. Thus $\mathrm{Tr}(Q^2\rho) = \sum_j a_j^2$. On the other hand $\mathrm{Tr}(Q\rho) = \sum_j a_j^2$, thus

$$\mathrm{Tr}\left(Q\rho\right)^2 \leq \mathrm{Tr}\left(Q^2\rho\right) \implies \left(\sum_j a_j^2\right)^2 \leq \sum_j a_j^2 \implies \sum_j a_j^2 \leq 1. \tag{2.18}$$

□

Lemma 2.6 tells us that pairwise anticommuting operators cannot all have large expected values on a quantum state: the sum of their squared expected values cannot exceed 1. This reminiscent of *Heisenberg's uncertainty principle*. More generally, in Definition 2.1 we defined the maximum sum of squares of a set of operators $\mathcal{S}$ to be its *commutation index*. For example, Lemma 2.6 shows that if all the operators in $\mathcal{S}$ anticommute, then $\Delta(\mathcal{S}) \leq 1/|\mathcal{S}|$.

We will see in Section 2.2 and Section 2.3 that the commutation index of the set $\mathcal{S}$ of operators has strong implications for the physics of the model consisting of the operators in $\mathcal{S}$ with Gaussian coefficients:

$$H = \frac{1}{\sqrt{|\mathcal{S}|}}\sum_{A\in\mathcal{S}} g_A A, \quad g_A \sim \mathcal{N}(0,1). \tag{2.19}$$

We will also see in Section 4.1 that the commutation index has connections to learning tasks on the set of operators $\mathcal{S}$ [Che+22].

Commutation graph and Lovasz theta function

Our computation of the commutation index will follow from studying the commutation structure of spin and fermionic operators, summarized by their *commutation graphs*.

Definition 2.2. (Commutation graph.) *The* commutation graph $G(\mathcal{S})$ *of a set* $\mathcal{S}$ *of Pauli or Majorana operators is defined as follows.*

- *The vertices of* $G(\mathcal{S})$ *correspond to operators* $A \in \mathcal{S}$.
- *We include an edge between any two vertices whose operators* anticommute.

We now introduce a key graph property which reveals the anticommutativity of the operators $\mathcal{S}$ through their commutation graph $G(\mathcal{S})$.

Definition 2.3. (Lovász theta function.) *Let* G *be a graph on* m *vertices. The* Lovász theta function $\vartheta(G)$ *is defined by the following semidefinite program of dimension* m*. Let* E *denote the edges in the graph* G*, and* $\mathbb{J}$ *the all-ones matrix.*

$$\begin{aligned} \max \{ \ & \mathrm{Tr}\,(\mathbb{J}X) \ , \ X \in \mathbb{R}^{m\times m} \\ & \textit{s.t. } X \succeq 0 \ , \ \mathrm{Tr}\,(X) = 1 \ , \ X_{jl} = 0 \ \forall (j,l) \in E \ \}, \end{aligned} \tag{2.20}$$

where E *denote the edges in the graph* G*;* $\mathbb{J}$ *the all-ones matrix;* $A \succeq B$ *denotes that* $A - B$ *is positive semidefinite;* $\mathrm{Tr}(X)$ *denotes the trace of* X*; and* X_{jl} *denotes entry* (j,l) *of* X*. It has dual*

$$\begin{aligned} \min \{ \ & \lambda \in \mathbb{R} \\ & \textit{s.t. } \exists\, Y \in \mathbb{R}^{m\times m} \ , \ Y_{jj} = 1 \ \forall j \ , \ Y_{jl} = 0 \ \forall (j,l) \notin E \ , \ \lambda Y \succeq \mathbb{J} \ \}. \end{aligned} \tag{2.21}$$

For any graph G, the following chain of inequalities is known:

$$I(G) \leq \vartheta(G) \leq \mathrm{chrom}(\overline{G}), \tag{2.22}$$

where $\overline{G}$ is the complement graph, $\mathrm{chrom}(\overline{G})$ is the chromatic number of $\overline{G}$, and $I(G)$ is the independence number of G. For example, see [Knu93].

The key reason for introducing the Lovasz theta function is that the Lovasz theta function of the commutation graph $G(\mathcal{S})$ upper bounds the commutation index $\Delta(\mathcal{S})$. This bound can be seen as a generalization of Lemma 2.6.

Lemma 2.7. *Let* $\mathcal{S}$ *be a set of Pauli operators*

$$\Delta(\mathcal{S}) \leq \frac{1}{|\mathcal{S}|}\vartheta(G(\mathcal{S})). \tag{2.23}$$

All together,

$$I(G(\mathcal{S})) \leq |\mathcal{S}| \cdot \Delta(\mathcal{S}) \leq \vartheta(G(\mathcal{S})) \leq \mathrm{chrom}(\overline{G}(\mathcal{S})). \tag{2.24}$$

Lemma 2.7 has appeared before in [GHG23; XSW23; HO22]; we reproduce a simplified proof here for the benefit of the reader.

Proof. The inequalities $I(G) \leq \vartheta(G) \leq \text{chrom}(\overline{G})$ always hold for any graph G [Knu93]. $\Delta(\mathcal{S}) \geq I(G(\mathcal{S}))/|\mathcal{S}|$ holds since we can choose ρ in the definition of $\Delta(\mathcal{S})$ to be in the simultaneous eigenbasis of the independent set of operators. It remains to establish $\Delta(\mathcal{S}) \leq \vartheta(G(\mathcal{S}))/|\mathcal{S}|$.

Denote $\mathcal{S} = \{A_1, \ldots, A_m\}$. Given ρ, let

$$a_j = \text{Tr}\left(A_j \rho\right). \tag{2.25}$$

We aim to show $\sum_j a_j^2 \leq \vartheta(G(\mathcal{S}))$.

Consider the observable

$$Q = \sum_j a_j A_j. \tag{2.26}$$

We have

$$Q^2 = \sum_{j,l} a_j a_l A_j A_l = \frac{1}{2} \sum_{j,l} a_j a_l \{A_j, A_l\}. \tag{2.27}$$

Note $A_j^2 = \mathbb{1}$ since they are Paulis.

Now take the trace with ρ. We get

$$\text{Tr}\left(Q^2 \rho\right) = \sum_{j,l} a_j a_l B_{jl} \leq \lambda_{\max}(B) \sum_j a_j^2, \tag{2.28}$$

where we defined the matrix

$$B_{jl} = \frac{1}{2} \text{Tr}\left(\{A_j, A_l\}\rho\right). \tag{2.29}$$

By positivity of the state ρ, we have $\text{Tr}\big(\big(Q - \text{Tr}\left(Q\rho\right)\mathbb{1}\big)^2 \rho\big) \geq 0$ and therefore

$$\text{Tr}\left(Q\rho\right)^2 \leq \text{Tr}\left(Q^2\rho\right) \tag{2.30}$$

$$\implies \Big(\sum_j a_j^2\Big)^2 \leq \lambda_{\max}(B) \sum_j a_j^2 \tag{2.31}$$

$$\implies \sum_j a_j^2 \leq \lambda_{\max}(B). \tag{2.32}$$

B satisfies $B_{jj} = 1 \; \forall j$ and $B_{jl} = 0$ for all edges (j, l). The latter holds since (j, l) is an edge precisely when $\{A_j, A_l\} = 0$. Positivity of the state ρ implies that B is

positive semidefinite, since for any vector $v \in \mathbb{R}^m$

$$v^T B v = \mathrm{Tr}\left(\Big(\sum_j v_j A_j\Big)^2 \rho\right) \geq 0. \tag{2.33}$$

Let's now take the supremum of the right-hand-side over all such B to get

$$\sum_j a_j^2 \leq \tilde{\vartheta}(G(\mathcal{S})), \tag{2.34}$$

where

$$\begin{aligned}\tilde{\vartheta}(G) = \max\ \{&\lambda_{\max}(B),\ B \in \mathbb{R}^{m\times m} \\ &\text{s.t.},\ B_{jj} = 1\ \forall j,\ B_{jl} = 0\ \forall (j,l) \in E,\ B \succeq 0\}.\end{aligned} \tag{2.35}$$

Lemma 2.8 completes the proof. □

Lemma 2.8. ([Knu93]) *The function* $\tilde{\vartheta}(G)$ *from Equation* (2.35) *satisfies* $\tilde{\vartheta}(G) \leq \vartheta(G)$.

Proof. We will use the dual description Equation (2.21). Let (λ, A) achieve the optimal dual value $\lambda = \vartheta(G)$. Define the $m \times (m+1)$ matrix

$$U = (\vec{1}, \sqrt{\lambda A - \mathbb{J}}), \tag{2.36}$$

where we padded with the all-ones column vector $\vec{1}$ on the left. (Recall $\mathbb{J}$ denotes the all-ones matrix.) Let B be any matrix feasible for $\tilde{\vartheta}(G)$. Decompose

$$B = Q^T D Q = V^T V \quad , \quad V = \sqrt{D} Q, \tag{2.37}$$

where Q is orthogonal and D is diagonal with $D_{11} = \lambda_{\max}(B)$. (The entries of D are the eigenvalues of B.) Now consider the collection of m matrices $\{Y^{(j)}\}$ of size $m \times (m+1)$ given by

$$Y^{(j)}_{ab} = V_{aj} U_{jb}. \tag{2.38}$$

We have

$$\mathrm{Tr}\left((Y^{(j)})^T Y^{(l)}\right) = \Big(\sum_a V_{aj} V_{al}\Big)\Big(\sum_b U_{jb} U_{lb}\Big) = \lambda B_{jl} A_{jl}. \tag{2.39}$$

If $j \neq l$, this is zero, since if (j,l) is an edge in G then $B_{jl} = 0$, and if not then $A_{jl} = 0$. If $j = l$, we get $\mathrm{Tr}\left((Y^{(j)})^T Y^{(j)}\right) = \lambda$. Thus $\{Y^{(j)}/\sqrt{\lambda}\}$ are orthonormal when viewed as vectors of dimension $m(m+1)$, and

$$1 \geq \sum_j \left(Y^{(j)}_{11}/\sqrt{\lambda}\right)^2 = \frac{D_{11}}{\lambda} \sum_j Q^2_{1j} = \frac{D_{11}}{\lambda} \implies D_{11} \leq \lambda. \tag{2.40}$$

□

It is in fact true that $\tilde{\vartheta}(G) = \vartheta(G)$, but we only need $\tilde{\vartheta}(G) \leq \vartheta(G)$ for our purposes.

Commutation index of local operators

The commutation index of the set of k-local Paulis is independent of system size n.

Theorem 2.9. *Let $\mathcal{P}_k^n$ be the set of k-local n-qubit Paulis. When $2n+1 \geq 3^k$, it holds:*

$$\Delta(\mathcal{P}_k^n) = 3^{-k}. \tag{2.41}$$

Moreover, the maximum is achieved by any product state.

Proof. First we aim to show $\mathbb{E}_{P\in\mathcal{P}_k^n}\langle\psi|P|\psi\rangle^2 = 3^{-k}$ for any product state of single-qubit states $|\psi\rangle = |\psi_1\rangle\otimes\cdots\otimes|\psi_n\rangle$. Denoting by $\mathcal{P}_k^S$ the set of Paulis on subsystem $S \subseteq [n]$, we have

$$\mathbb{E}_{P\in\mathcal{P}_k^n}\langle\psi|P|\psi\rangle^2 = \mathbb{E}_{S\subseteq[n],|S|=k}\mathbb{E}_{P\in\mathcal{P}_k^S}\langle\psi|P|\psi\rangle^2. \tag{2.42}$$

By tracing out $[n]\setminus S$, it is sufficient to show

$$\mathbb{E}_{P\in\mathcal{P}_k^k}\langle\psi|P|\psi\rangle^2 = 3^{-k} \tag{2.43}$$

for any product state $|\psi\rangle = |\psi_1\rangle\otimes\cdots\otimes|\psi_k\rangle$. Since $|\mathcal{P}_k^k| = 3^k$, this is equivalent to

$$\sum_{P\in\mathcal{P}_k^k}\langle\psi|P|\psi\rangle^2 = 1. \tag{2.44}$$

But this holds since

$$\sum_{P\in\mathcal{P}_k^k}\langle\psi|P|\psi\rangle^2 = \prod_{j=1}^{k}\Big(\sum_{P\in\{\sigma_X,\sigma_Y,\sigma_Z\}}\langle\psi_j|P|\psi_j\rangle^2\Big) = 1 \tag{2.45}$$

using that the single-qubit states $|\psi_j\rangle$ are pure.

For the upper bound, we will invoke Lemma 2.7. Recalling that $G(\mathcal{P}_k^n)$ is the commutation graph of k-local Paulis, it suffices to show that $\vartheta(G(\mathcal{P}_k^n)) \leq 3^{-k}\cdot|\mathcal{P}_k^n|$. For this purpose, we import a fact from [Knu93] using a proof technique similarly applied in [AGK24]. A graph G is *vertex-symmetric* if for any two vertices u, v, there is an automorphism of G taking u to v.

Fact 2.1 ([Knu93])**.** *If graph G is vertex-symmetric, then*

$$\vartheta(G)\cdot\vartheta(\overline{G}) = |G|. \tag{2.46}$$

where $\overline{G}$ is the complement graph and $|G|$ denotes the number of vertices in G.

The commutation graph $G(\mathcal{P}_k^n)$ of k-local Paulis is vertex-symmetric. Thus to establish the upper bound $\vartheta(G(\mathcal{P}_k^n)) \leq 3^{-k} \cdot |\mathcal{P}_k^n|$, it suffices to show $\vartheta(\overline{G}(\mathcal{P}_k^n)) \geq 3^k$. Using the independence number inequality of Lemma 2.7, it suffices to find an independent set in $\overline{G}(\mathcal{P}_k^n)$ of size at least 3^k. This is equivalent to a *clique* in $G(\mathcal{P}_k^n)$ of size at least 3^k. In other words, we must exhibit a set of 3^k mutually anticommuting k-local Paulis. This can be done using the ternary tree embedding of [Vla19; Jia+20a]. Let $2n + 1 = 3^k$. The ternary tree construction at depth k embeds $2n + 1$ mutually anticommuting operators into n qubits where each anticommuting operator has locality k. □

On the other hand, the commutation index of the set of degree-q Majorana operators decays polynomially with system size.

Definition 2.4. *The Majorana operators on n fermionic modes are defined abstractly as n operators $\{\gamma_1, \ldots, \gamma_n\}$ which satisfy the relations*

$$\gamma_a \gamma_b + \gamma_b \gamma_a = 2\delta_{ab} \mathbb{1}. \tag{2.47}$$

A degree-q Majorana operator is a degree-q monomial in the Majorana operators.

The following theorem is the key result which enables Theorem 2.1, Theorem 2.2, and Theorem 4.2.

Theorem 2.10. *Let $\mathcal{S}_q^n$ be the set of degree-q Majorana operators on n fermionic modes with q even. Then*

$$\left| \frac{\Delta(\mathcal{S}_q^n)}{\binom{n/2}{q/2} / \binom{n}{q}} - 1 \right| \leq O(n^{-1}) \tag{2.48}$$

for all n sufficiently large, for each q.

Proof. Let q and n be even. First we show the lower bound in Theorem 2.10. We can find a set $S \subseteq \mathcal{S}_q^n$ of mutually commuting degree-q Majoranas of size $\binom{n/2}{q/2}$ by taking $(q/2)$-wise products of $\{i\gamma_1\gamma_2, i\gamma_3\gamma_4, \ldots, i\gamma_{n-1}\gamma_n\}$. Let ρ be the state which is maximally mixed within the simultaneous +1-eigenspace of the operators in S. Then $\sum_{A \in \mathcal{S}_q^n} \langle\psi|A|\psi\rangle = |S| = \binom{n/2}{q/2}$ and $\Delta(\mathcal{S}_q^n) \geq \binom{n/2}{q/2} / \binom{n}{q}$. (Note the number of degree-q monomials on n Majoranas is $|\mathcal{S}_q^n| = \binom{n}{q}$.) The remainder of this section is devoted to showing the upper bound via the Lovász theta function. We aim to establish the following theorem.

Theorem 2.11. *Let $\mathcal{S}_q^n$ be the set of degree-q Majorana operators on n modes. Then*

$$\vartheta(G(\mathcal{S}_q^n)) \leq \binom{n/2}{q/2} + O(e^{O(q\log q)} n^{q/2-1}) \tag{2.49}$$

for all n sufficiently large, for each q.

Noting that $|\mathcal{S}_q^n| = \binom{n}{q}$, the upper bound in Theorem 2.10 follows from combining Theorem 2.11 and Lemma 2.7. Thus it suffices to establish Theorem 2.11. After completing our work, we became aware of Ref. [Lin24], in which they establish $\vartheta(G(\mathcal{S}_q^n)) \leq \binom{n/2}{q/2} + c(q)n^{q/2-1}$ for some function $c(q)$. Theorem 2.11 is stronger, since it specifies the asymptotic dependence $c(q) = O(e^{O(q\log q)})$. We give a self-contained proof of Theorem 2.11.

The *Johnson association scheme* $\mathcal{J}_d(n,q)$ is the graph whose vertices correspond to subsets $S \subseteq [n]$ of size $|S| = q$, and (S,T) forms an edge if $q - |S \cap T| = d$. Write $A_d^{n,q}$ for the adjacency matrix of the Johnson scheme $\mathcal{J}_d(n,q)$. The graph $G(\mathcal{S}_q^n)$ has adjacency matrix A equal to

$$A = A_1^{n,q} + A_3^{n,q} + \cdots + A_{q-1}^{n,q}. \tag{2.50}$$

We are interested in the Lovász theta function of $G(\mathcal{S}_q^n)$. The following three results advertised in [HO22] reduce $\vartheta(G(\mathcal{S}_q^n))$ to a linear program involving Hahn polynomials.

Lemma 2.12. *([Del73, p. 48]) The matrices $A_0^{n,q}, \ldots, A_q^{n,q}$ are simultaneously diagonalizable, with eigenvalues given by the* dual Hahn polynomials*:*

$$spec(A_d^{n,q}) = \{\tilde{H}_d^{n,q}(x) : x = 0, \ldots, q\} \tag{2.51}$$

$$\tilde{H}_d^{n,q}(x) = \sum_{j=0}^{d} (-1)^{d-j} \binom{q-j}{d-j} \binom{q-x}{j} \binom{n-q+j-x}{j}. \tag{2.52}$$

In particular, for all $d > 0$ the all-1's vector is an eigenvector of $A_d^{n,q}$ of multiplicity 1 with eigenvalue $\tilde{H}_d^{n,q}(0) = \binom{n-q}{d}\binom{q}{d}$.

Lemma 2.13 ([HO21], Lemma 4.25)**.** *In the dual formulation of the Lovász theta function in Definition 2.3 for the graph $G(\mathcal{S}_q^n)$ it suffices to minimize over matrices Y whose entries $Y(S,T)$ depend only on dist(S,T). Thus we can write*

$$\vartheta(G(\mathcal{S}_q^n)) = \\ \min\{\lambda \,:\, \exists\, a_1, a_3, \ldots, a_{q-1} \text{ s.t. } \lambda\big(\mathbb{1} + a_1 A_1^{n,q} + a_3 A_3^{n,q} + \cdots + a_{q-1} A_{q-1}^{n,q}\big) \succeq \mathbb{J}\}. \tag{2.53}$$

Corollary 2.14. ([HO21], Corollary 4.26)

$$\vartheta(G(\mathcal{S}_q^n)) \;=\; \min_{a_1,a_3,\ldots,a_{q-1}} \{ \binom{n}{q} \big/ \big(1+P(0)\big) \;:\; P(1),\ldots,P(q) \geq -1 \\ \text{where } P(x) = a_1 \tilde{H}_1^{n,q}(x) + a_3 \tilde{H}_3^{n,q}(x) + \cdots + a_{q-1}\tilde{H}_{q-1}^{n,q}(x)\}. \quad (2.54)$$

Our strategy to prove Theorem 2.11 will follow by finding a feasible polynomial P^* for the linear program (LP) in Corollary 2.14, and showing that $P^*(0)$ has the correct scaling in the large n limit:

$$\binom{n}{q} \big/ \big(1+P^*(0)\big) = \frac{1}{(q/2)!}(n/2)^{q/2} + O(n^{q/2-1}). \quad (2.55)$$

Then all that remains is to control the error term for finite n. This will follow from an application of the polynomial method.

We first state and prove a lemma which finds an appropriate feasible polynomial P^*.

Lemma 2.15. *We can choose a feasible polynomial P^* for the LP in Corollary 2.14 that simultaneously satisfies:*

- $P^*(0) = \frac{2^{q/2}(q/2)!}{(q)!} n^{q/2} - c(q) n^{q/2-1}$ *for some $c(q)$.*
- $P^*(0)/n^{q/2}$ *is a polynomial in n^{-1} of degree $2q-2$.*
- $|P^*(0)/n^{q/2}| \leq e^{2q} q^{2q}$ *for all n.*

Proof. For constant q, $x = 1,\ldots,q$ and large n, the leading order term in the Hahn polynomial is

$$\tilde{H}_d^{n,q}(x) = \begin{cases} \Theta(n^d) + O(n^{d-1}) & d \leq q-x \\ (-1)^{d+x-q}\Theta(n^{q-x}) + O(n^{q-x-1}) & d > q-x. \end{cases} \quad (2.56)$$

It will be useful later to be more specific about the coefficients in the cases $d = q-x-1$ and $d = q-x+1$. To leading order in n with q constant we have

$$\tilde{H}_{q-x-1}^{n,q}(x) = h^{(q)}(x)\cdot n^{q-x-1} + O(n^{q-x-2}) \quad , \quad \tilde{H}_{q-x+1}^{n,q}(x) = -g^{(q)}(x)\cdot n^{q-x} + O(n^{q-x-1}), \quad (2.57)$$

where

$$h^{(q)}(x) = \frac{q-x}{(q-x-1)!} \quad , \quad g^{(q)}(x) = \frac{x}{(q-x)!}\,. \quad (2.58)$$

Now let us examine the LP in Corollary 2.14 in the large-n limit. Our strategy is to sequentially go through $x = q, \ldots, 1$ and ensure that $P(x) \geq -1$ for sufficiently large n for each x. If we choose $a_1, a_3, \ldots, a_{q-1}$ to all be eventually positive, then automatically $P(x) \geq 0$ eventually for odd x. This is because the leading order term for large n in Equation (2.56) is always positive if x is odd (recall that q is even and d is always odd). Ensuring that $P(x) \geq -1$ eventually for even x will require choosing a_{q-x+1} as a function of a_{q-x-1}.

Let us first consider $P(q)$. It can be seen from the definition Equation (2.52) that $\tilde{H}_d^{n,q}(q) = -\binom{q}{d}$ is negative for all odd d and independent of n. If we set

$$a_1 = \frac{1 - C_1 n^{-1}}{g^{(q)}(q)} \tag{2.59}$$

for some sufficiently large constant C_1, and all other $a_d = O(n^{-1})$ for $d > 1$, then the constraint $P(q) \geq -1$ is satisfied.

Now consider $P(q-2)$. If we set

$$\begin{aligned} a_3 &= (1 - C_3 n^{-1}) \cdot \frac{h^{(q)}(q-2)}{g^{(q)}(q-2)} \cdot n^{-1} \cdot a_1 \\ &= (1 - C_3 n^{-1})(1 - C_1 n^{-1}) \cdot \frac{h^{(q)}(q-2)}{g^{(q)}(q-2)} \cdot \frac{1}{g^{(q)}(q)} \cdot n^{-1} \end{aligned} \tag{2.60}$$

for some sufficiently large constant C_3, and all other $a_d = O(n^{-3})$ for $d > 3$, then the constraint $P(q-2) \geq -1$ is satisfied.

Continue like this for $P(q-4), \ldots, P(2)$. For each of $x = q, q-2, \ldots, 2$, we will set

$$\begin{aligned} a_{q-x+1} &= (1 - C_{q-x+1} n^{-1}) \cdot \frac{h^{(q)}(x)}{g^{(q)}(x)} \cdot n^{-1} \cdot a_{q-x-1} \\ &= (1 - C_{q-x+1} n^{-1}) \ldots (1 - C_1 n^{-1}) \\ &\quad \cdot \frac{h^{(q)}(x) \cdot h^{(q)}(x+2) \cdot \cdots \cdot h^{(q)}(q-2)}{g^{(q)}(x) \cdot g^{(q)}(x+2) \cdot \cdots \cdot g^{(q)}(q-2)} \cdot \frac{1}{g^{(q)}(q)} \cdot n^{-q/2+(x/2)} \end{aligned} \tag{2.61}$$

for some sufficiently large constants $C_1, C_3, \ldots, C_{q-1}$. Notice that

$$\frac{h^{(q)}(t)}{g^{(q)}(t)} = \frac{(q-t)^2}{t} \tag{2.62}$$

and $g^{(q)}(q) = q$, so defining

$$\hat{a}_{q-x+1} := \frac{(q-x)^2 (q-x-2)^2 \ldots 2^2}{(q-2)(q-4)\ldots x} \cdot \frac{1}{q} \tag{2.63}$$

independent of n, we can write

$$a_{q-x+1} = (1 - C_{q-x+1}n^{-1}) \dots (1 - C_1 n^{-1}) \cdot \hat{a}_{q-x+1} \cdot n^{-q/2+(x/2)}. \tag{2.64}$$

Finally, let us look at $P(0)$. For large n, we have

$$\begin{aligned} P(0) &= a_{q-1} \cdot h^{(q)}(0) \cdot n^{q-1} + \dots \\ &= (1 - C_{q-x+1}n^{-1}) \dots (1 - C_1 n^{-1}) \cdot \frac{q}{(q-1)!} \cdot \hat{a}_{q-1} \cdot n^{q/2} - c(q) n^{q/2-1}, \end{aligned} \tag{2.65}$$

where $c(q)$ is some function of q (note $h^{(q)}(0) = q/(q-1)!$). The product in $\hat{a}_{q-1}$ telescopes to give

$$\hat{a}_{q-1} = \frac{1}{q}(q-2)(q-4)\dots 2, \tag{2.66}$$

so we get

$$P(0) = \frac{(1 - C_{q-x+1}n^{-1}) \dots (1 - C_1 n^{-1})}{(q-1)(q-3)\dots 1} \cdot n^{q/2} - c(q) n^{q/2-1}. \tag{2.67}$$

This establishes the first point of Lemma 2.15.

Let us now examine $P(0)/n^{q/2}$ in order to establish the second point of Lemma 2.15:

$$\begin{aligned} P(0)/n^{q/2} &= \frac{1}{n^{q/2}} \left(a_1 \tilde{H}_1^{n,q}(0) + a_3 \tilde{H}_3^{n,q}(0) + \dots + a_{q-1} \tilde{H}_{q-1}^{n,q}(0) \right) \\ &= \hat{a}_1 \binom{q}{1} \binom{n-q}{1} \cdot n^{-q/2} \cdot (1 - C_1 n^{-1}) \\ &\quad + \hat{a}_3 \binom{q}{3} \binom{n-q}{3} \cdot n^{-q/2-1} \cdot (1 - C_1 n^{-1})(1 - C_3 n^{-1}) \\ &\quad \dots + \hat{a}_{q-1} \binom{q}{q-1} \binom{n-q}{q-1} \cdot n^{-q+1} \cdot (1 - C_1 n^{-1}) \dots (1 - C_{q-1} n^{-1}), \end{aligned} \tag{2.68}$$

recalling $\tilde{H}_d^{n,q}(0) = \binom{q}{d}\binom{n-q}{d}$. Recall that the coefficients $\hat{a}_d$ are independent of n. From this expression, we can readily see that $P(0)/n^{q/2}$ is a polynomial in n^{-1} of degree $2q-2$, establishing the desired degree bound.

It remains to establish the third point of Lemma 2.15. For all $d = 1, 3, \dots, q-1$, $(1 - C_1 n^{-1}) \dots (1 - C_d n^{-1}) \leq 1$ eventually and $\hat{a}_d \leq \hat{a}_{q-1}$. Further, for all $d = 1, 3, \dots, q-1$

$$\binom{q}{d}\binom{n-q}{d} \leq \left(\frac{e^2 q(n-q)}{d^2} \right)^d \leq (e^2 q)^{q-1} n^d, \tag{2.69}$$

using the general bound $\binom{m}{r} \leq (em/r)^r$. Using these facts, we can bound

$$\begin{aligned} \left| P(0)/n^{q/2} \right| &\leq \hat{a}_{q-1} \cdot (e^2 q)^{q-1} \cdot \left(n^{-q/2+1} + n^{-q/2+2} + \dots + 1 \right) \\ &\leq q^{q/2-2} \cdot (e^2 q)^{q-1} \cdot (q/2) \leq e^{2q} q^{2q}, \end{aligned} \tag{2.70}$$

using $\hat{a}_{q-1} \leq q^{q/2-2}$ in the final step. □

The first bullet point in Lemma 2.15 gives the correct large-n limit for the Lovász theta function for constant q:

$$\binom{n}{q} / \big(1 + P^*(0)\big) = \left(\frac{1}{(q)!} \cdot n^q + O(n^{q-1})\right) \cdot \left(\frac{2^{q/2}(q/2)!}{(q)!} n^{q/2} - O(n^{q/2-1})\right)^{-1} \tag{2.71}$$

$$= \frac{1}{(q/2)!}(n/2)^{q/2} + O(n^{q/2-1}). \tag{2.72}$$

We are almost done. It remains to show that $c(q) \leq e^{O(q \log q)}$. We will do this using the second and third bullet points, combined with Markov's "other inequality" for bounded degree polynomials [BE95, Theorem 5.1.8]. We will need the following adaptation for functions defined on inverse-integer points $1/1, 1/2, \cdots, 1/n, \cdots$.

Lemma 2.16 (Markov's other inequality [BE95, Theorem 5.1.8], [Che+24a], [Che+24b])**.** *For any polynomial f of degree ℓ, there is an absolute constant c such that*

$$|f(1/n) - f(0)| \leq \frac{c\ell^4}{n} \sup_{n' \geq 1} |f(1/n')| \tag{2.73}$$

for each integer n.

Proof.

$$n|f(1/n) - f(0)| \leq \sup_{x \in [0,1/n]} |f'(x)| \qquad \text{(Fundamental theorem of calculus.)} \tag{2.74}$$

$$\leq \frac{2\ell^2}{a} \sup_{x \in [0,a]} |f(x)| \qquad \text{([Che+24a, Lemma 4.1])} \tag{2.75}$$

$$\leq \frac{4\ell^2}{a} \sup_{n' \geq 1/a} |f(n')| \quad \text{([Che+24a, Lemma 4.2]), } a = \frac{1}{4q^2}) \tag{2.76}$$

where the final line considers the supremum over inverse integer points, which concludes the proof. □

Let us view $P^*(0)/n^{q/2}$ as a polynomial in $1/n$ (note that we are fixing the input to the polynomial P^* to zero, and that we are concerned with the n-dependence of

$P^*(0)$). Using Lemma 2.15, we can apply **??** with $f(1/n) = P^*(0)/n^{q/2}$, $\ell = 2q-2$, and

$$\sup_{n' \geq 1} |f(1/n')| \leq e^{2q} q^{2q} = e^{O(q \log q)}. \tag{2.77}$$

This completes the proof of Theorem 2.11. □

Numerics on Lovász theta function of local Majorana operators

In this section we present some numerics on the Lovász theta function of the commutation graph of degree-q Majorana operators $\vartheta(G(\mathcal{S}_q^n))$.

	$\vartheta(G(\mathcal{S}_q^n))$					$\binom{n/2}{q/2}$				
n	$q=2$	$q=4$	$q=6$	$q=8$	$q=10$	$q=2$	$q=4$	$q=6$	$q=8$	$q=10$
2	1					1				
4	2	1				2	1			
6	3	3	1			3	3	1		
8	4	14	4	1		4	6	4	1	
10	5	14.57	14.57	5	1	5	10	10	5	1
12	6	15	52	15	6	6	15	20	15	6
14	7	21	57.34	57.34	21	7	21	35	35	21
16	8	28	64	198	64	8	28	56	70	56
18	9	36	100.13	218.34	218.34	9	36	84	126	126
20	10	45	153.11	251.22	787.17	10	45	120	210	252
22	11	55	195.13	429.91	885.15	11	55	165	330	462
24	12	66	236.42	759	982.84	12	66	220	495	792
26	13	78	286	990.80	1757.0	13	78	286	715	1287
28	14	91	364	1217.2	3260.2	14	91	364	1001	2002
30	15	105	455	1444.2	4643.9	15	105	455	1365	3003
32	16	120	560	1820.0	6040.7	16	120	560	1820	4368
34	17	136	680	2423.3	7240.0	17	136	680	2380	6188
36	18	153	816	3327.1	9269.4	18	153	816	3060	8568
38	19	171	969	4512.8	12552	19	171	969	3876	11628
40	20	190	1140	6022.1	17230	20	190	1140	4845	15504

Table 2.4: Numerical comparison of the Lovász theta function $\vartheta(G(\mathcal{S}_q^n))$ versus $\binom{n/2}{q/2}$. They are exactly equal for very small values of n, and also appear to be exactly equal for sufficiently large values of n for each q. For example at $q = 4$, which corresponds to the standard SYK-4 model, it appears that $\vartheta(G(\mathcal{S}_4^n)) = \binom{n/2}{2}$ for all even values of n apart from $n = 8$ and $n = 10$.

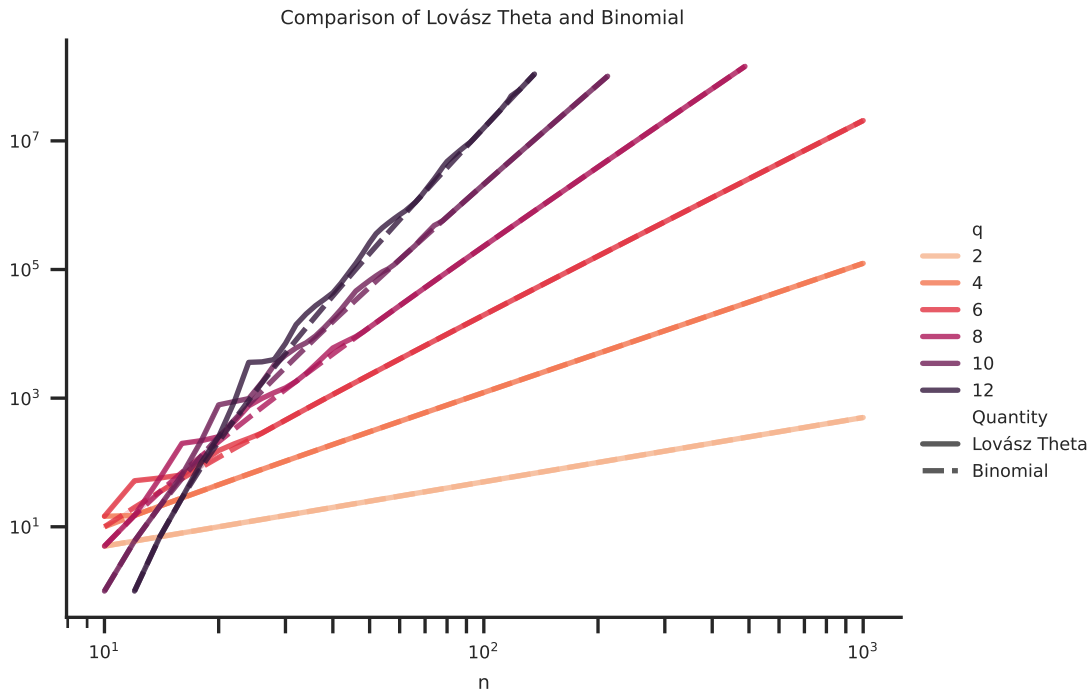

Figure 2.1: Log-log plot of $\vartheta(G(\mathcal{S}_q^n))$ versus $\binom{n/2}{q/2}$. $\vartheta(G(\mathcal{S}_q^n))$ fluctuates for small n, but for sufficiently large n it behaves the same as $\binom{n/2}{q/2}$.

2.2 Circuit lower bound

In this section we show that low-energy states of random Hamiltonians with small commutation index require ansatzes with high complexity. As a consequence, we get a circuit lower bound on the low energy states of the the SYK model. Our proof resembles the circuit lower bound of [Dal+23b, Appendix D].

Consider the random Hamiltonian

$$H = \frac{1}{\sqrt{|\mathcal{S}|}} \sum_{A \in \mathcal{S}} g_A A \quad , \quad g_A \sim_{i.i.d.} \mathcal{N}(0,1), \tag{2.78}$$

where $\mathcal{S}$ consists of Hermitian operators A_i satisfying $A_i^2 = \mathbb{1}$. Recall from Definition 2.1 that $\Delta(\mathcal{S})$ denotes the commutation index of $\mathcal{S}$. We first establish a concentration bound for the energy of a fixed state $|\psi\rangle$.

Lemma 2.17. *Fix any state $|\psi\rangle$. The energy $\langle\psi|H|\psi\rangle$ sharply concentrates:*

$$\mathbb{P}\left(\langle\psi|H_q^{SYK}|\psi\rangle \geq t\right) \leq \exp\left(-\frac{t^2}{2\Delta(\mathcal{S})}\right). \tag{2.79}$$

Proof. Since a sum of Gaussians is Gaussian, we have

$$\mathbb{P}\big(\langle\psi|H|\psi\rangle \geq t\big) \leq \exp\left(-\frac{t^2}{2\sigma^2}\right). \tag{2.80}$$

The variance is upper bounded by the commutation index:

$$\sigma^2 = \frac{1}{|\mathcal{S}|}\sum_{A\in\mathcal{S}}\langle\psi|A|\psi\rangle^2 \leq \Delta(\mathcal{S}). \tag{2.81}$$

□

The main result of this section now follows from a union bound.

Theorem 2.18 (Low energy states are classically non-trivial)**.** *Let C be a set of fixed quantum states. Then the probability that C contains a low-energy state is exponentially small.*

$$\mathbb{P}\left[\max_{|\psi\rangle\in C}\langle\psi|\,H\,|\psi\rangle \geq t\right] \leq \exp\left(\log\|C\| - \frac{t^2}{2\Delta}\right). \tag{2.82}$$

Proof. Applying a union bound to Lemma 2.17,

$$\mathbb{P}\left[\max_{|\psi\rangle\in C}\langle\psi|\,H\,|\psi\rangle \geq t\right] \leq |C|\exp\left(-\frac{t^2}{2\Delta(\mathcal{S})}\right) = \exp\left(\log\|C\| - \frac{t^2}{2\Delta}\right). \tag{2.83}$$

□

We now apply this to the circuit complexity of low temperature states of the SYK model. The SYK model was previously described in Equation (2.1), but we repeat its definition here for convenience.

Definition 2.5. *Let $\mathcal{S}_q^n$ denote the set of degree-q Majorana operators on n fermionic modes. The SYK$_q$ model is a random ensemble of Hamiltonians defined by*

$$H_q^{SYK} = \frac{1}{\sqrt{\binom{n}{q}}}\sum_{A\in\mathcal{S}_q^n} g_A A \quad , \quad g_A \sim_{i.i.d.} \mathcal{N}(0,1). \tag{2.84}$$

Theorem 2.19. (SYK model low-energy states have high circuit complexity.) *Let $circ(G)$ denote the set of unitaries generated by quantum circuits with at most G gates each taken from a finite universal set of 2-local unitary gates. Fix an arbitrary initial state $|\phi\rangle$. With high probability, for any even $q \geq 2$, it holds that the minimum circuit complexity to construct a state achieving at least $t\sqrt{n}$ on H_q^{SYK} is at least*

$$\min\left\{G :\ \exists U \in circ(G),\ \langle\phi|\,U^\dagger H_q^{SYK} U\,|\phi\rangle \geq t\sqrt{n}\right\} = \tilde{\Omega}_q(n^{(q/2)+1}t^2). \tag{2.85}$$

Proof. Theorem 2.10 tells us that

$$\Delta(\mathcal{S}_q^n) = \Omega_q(n^{-q/2}). \tag{2.86}$$

Let M be the number of gates in the universal gate set. Then the number of circuits in $\mathrm{circ}(G)$ is at most

$$|\mathrm{circ}(G)| \leq \left(M\binom{n}{2}\right)^G = \exp(O(G\log(n))). \tag{2.87}$$

Now apply Theorem 2.18 with $C = \mathrm{circ}(G)$ to complete the proof. □

The proof above can be extended to gates with continuous parameters by forming an ϵ-net over the gates. This comes at the cost of additional $\log(1/\epsilon)$ factors in the bound of Theorem 2.19.

Other notions of non-triviality

Though our focus so far has been on quantum circuit lower bounds, our results readily generalize to lower bounds for other classes of ansatzes via the construction of covering nets. For example, our argument shows that any state from the set of Gaussian states cannot be a near ground state for the SYK Hamiltonian for $q \geq 4$, since an ϵ-net over the set of Gaussian states has cardinality $\exp(\tilde{O}(n^2 + \mathrm{poly}\log(1/\epsilon)))$. This reproduces results from prior works [HTS21b; Her+23b].

Another popular classical ansatz are tensor network states, or matrix product states (MPSs) in particular. Implemented at any finite precision the number of configurations of a matrix product state on n sites grows with the *bond dimension* χ as $\left\|\{|\psi_j\rangle\}\right\| = \exp\left(\Theta\left(\chi^2 + \log n\right)\right)$. It is thus apparent from the same argument that the minimum bond dimension such that there is an MPS achieving an energy $t\sqrt{n}$ is

$$\chi = \Omega_q\left(n^{q/4+1/2}t\right) \tag{2.88}$$

with high probability. Similarly, a classical neural network representation of the state with W weights has a number of configurations growing as $\left\|\{|\psi_j\rangle\}\right\| = \exp\left(\Theta\left(W\right)\right)$, yielding the growth condition to achieve an energy $t\sqrt{n}$ with high probability:

$$W = \Omega_q\left(n^{q/2+1}t^2\right). \tag{2.89}$$

Relation to NLTS results

Our circuit lower bound is closely related to the study of 'no low-energy trivial states' (NLTS) Hamiltonians. Introduced in [FH13], a Hamiltonian $H = \sum_i g_i A_i$

has the NLTS property if there is no constant-depth circuit preparing a state whose energy is above the ground energy by less than some constant fraction of the ℓ_1 norm $\sum_i |g_i|$. Such Hamiltonians were first proven to exist in [ABN23] using quantum LDPC codes.

The circuit lower bounds we give are not quite comparable to the traditional notion of NLTS. This is because we compare the energy of our low-energy states to the Hamiltonian's maximum eigenvalue rather than the ℓ_1 norm of the coefficients. Unlike the quantum code Hamiltonian studied in [ABN23; Her+23a], the SYK model is highly frustrated and thus the operator norm and ℓ_1 norms have vastly different scalings: $\Theta(n^2)$ and $\Theta(\sqrt{n})$, respectively.

Despite these differences from the standard NLTS setting, the circuit lower bounds we can establish are much stronger in two ways when compared to current progress on NLTS [EH17; AB22; AGK23; Her+23a; ABN23]. First, our circuit lower bounds hold for states at *any* energy which is a constant fraction of the ground state energy, rather than for states below some constant-fraction energy threshold. Second, we can achieve arbitrary polynomial circuit depth lower bounds, whereas current constructions of NLTS only give a logarithmic depth lower bound.

Product state approximations for spin Hamiltonians

It is worth pointing out that there cannot be a k-local spin Hamiltonian with the property in Theorem 2.19. For any traceless k-local spin Hamiltonian H, there is a product state achieving energy at least $\lambda_{\max}(H)/3^k$. The argument is imported from [Bra+19, proof of Theorem 2], and the proof technique bears a remarkable resemblence to the classical shadows protocol [HKP20], which provides a learning algorithm for k-local spin operators.

Proposition 2.20. *For any k-local Hamiltonian H on n qubits, there is a product state $|\psi\rangle$ achieving energy*

$$\langle\psi|H|\psi\rangle \geq \frac{1}{3^k}\lambda_{\max}(H). \tag{2.90}$$

Proof. Let $|\phi\rangle$ be the true (possibly entangled) maximum-energy state achieving

$$\langle\phi|H|\phi\rangle = \lambda_{\max}. \tag{2.91}$$

For each qubit, pick a random basis out of $\{\sigma_X, \sigma_Y, \sigma_Z\}$, and measure in this basis. This gives a product state of single-qubit stabilizer states. Let ρ be the resulting

ensemble of pure product states. We will analyze $\mathrm{Tr}\,(H\rho)$. Measuring each qubit in a random $\{\sigma_X, \sigma_Y, \sigma_Z\}$ basis implements the depolarizing channel with $p = 1/3$:

$$\rho = \mathcal{E}_{1/3}^{\otimes n}(|\phi\rangle\langle\phi|) \quad , \quad \mathcal{E}_{1/3}(\tau) = \frac{1}{3}\tau + \frac{2}{3}\mathbb{1}. \tag{2.92}$$

Using that the depolarizing channel is self-adjoint, we get

$$\mathrm{Tr}\,(H\rho) = \mathrm{Tr}\left(H\mathcal{E}_{1/3}^{\otimes n}(|\phi\rangle\langle\phi|)\right) = \mathrm{Tr}\left(\mathcal{E}_{1/3}^{\otimes n}(H)|\phi\rangle\langle\phi|\right). \tag{2.93}$$

By assumption, the Hamiltonian H is a sum of k-local Paulis. If P is a Pauli string of weight k, then the depolarizing channel acts as

$$\mathcal{E}_{1/3}^{\otimes n}(P) = \frac{1}{3^k}P. \tag{2.94}$$

Thus we get

$$\mathrm{Tr}\left(H\mathcal{E}_{1/3}^{\otimes n}(|\phi\rangle\langle\phi|)\right) = \mathrm{Tr}\left(\mathcal{E}_{1/3}^{\otimes n}(H)|\phi\rangle\langle\phi|\right) = \langle\phi|\left(\frac{1}{3^k}H\right)|\phi\rangle = \frac{1}{3^k}\lambda_{\max}(H). \tag{2.95}$$

□

2.3 Annealed approximation

We here prove Theorem 2.5, as well as the variety of concentration results stated in Table 2.3. Recall that we are interested in models of the form:

$$H = \frac{1}{\sqrt{m}}\sum_{i=1}^{m} g_i A_i, \tag{2.96}$$

where $g_i \sim_{i.i.d.} \mathcal{N}(0,1)$ are standard independent Gaussians and A_i are deterministic matrices. The *Gibbs state* ρ_β at inverse temperature β is defined by

$$\rho_\beta = \frac{e^{-\beta\sqrt{n}H}}{Z_\beta} \quad , \quad Z_\beta = \mathrm{Tr}\left(e^{-\beta\sqrt{n}H}\right), \tag{2.97}$$

where Z_β is called the *partition function* at inverse temperature β. The factor $\sqrt{n}$ ensures that the free energy is extensive and scales proportionally to n.

In this section we show that the commutation index of the terms A_i has an important effect on the concentration properties of the random model H. Denote the commutation index by

$$\Delta := \Delta(\{A_i\}_{i=1}^m). \tag{2.98}$$

Recall that this quantity characterizes the variance of the energy with respect to a fixed state ρ:

$$\sup_{\rho} \mathbb{E}_H |\operatorname{Tr}(H\rho)|^2 = \sup_{\rho} \frac{1}{m} \sum_{i=1}^{m} (\operatorname{Tr}(A_i\rho))^2 = \Delta. \tag{2.99}$$

The value of Δ has implications for relations between the normalized quenched and annealed free energies.

$$\underbrace{\frac{1}{n}\mathbb{E}\log Z_\beta}_{\text{quenched}} \quad vs. \quad \underbrace{\frac{1}{n}\log \mathbb{E}Z_\beta}_{\text{annealed}}. \tag{2.100}$$

The first result is that this variance quantity controls the difference between the two.

Theorem 2.21. (Quenched and annealed free energy)

$$n^{-1}\mathbb{E}\log Z_\beta \leq n^{-1}\log \mathbb{E}Z_\beta \leq n^{-1}\mathbb{E}\log Z_\beta + 4\beta^2\Delta. \tag{2.101}$$

The first inequality always holds by Jensen's inequality; the non-trivial part is the second inequality, which is proved in Section 2.3. This theorem states that a small variance $\Delta \ll \beta^{-2}$ implies that that the annealed free energy well-approximates the quenched free energy, which indicates the absence of spin glass order [BS20]. The next three results give concentration of expectation values, energy, and two-point correlators of the thermal state of H. Two-point correlators are of special interest in the study of the SYK model [BS20; GMV18; KS18; MS16]. Concentration results for Lipschitz bounded functions of the spectrum of the SYK model have also been established in [FTW20]. In here and what follows, $\|\cdot\|$ denotes the operator norm.

Theorem 2.22. (Concentration of expectation values) *For any fixed bounded Hermitian operator X,*

$$\mathbb{P}\left(\left|\operatorname{Tr}\left(X\rho_\beta\right) - \mathbb{E}\operatorname{Tr}\left(X\rho_\beta\right)\right| \geq t\right) \leq 2e^{-t^2/(18\beta^2||X||^2\Delta)}. \tag{2.102}$$

Theorem 2.23. (Concentration of energy)

$$\mathbb{P}\left(\left|\operatorname{Tr}\left(H\rho_\beta\right) - \mathbb{E}\operatorname{Tr}\left(H\rho_\beta\right)\right| \geq t\right) \leq 4\exp\left(-\frac{1}{2\Delta}\left(\sqrt{\frac{t^2}{12\beta^2 n} + \alpha^2} - \alpha\right)\right), \tag{2.103}$$

where $\alpha = \frac{1}{2}\left(1/(4\beta^2 n) + \mathbb{E}[\lambda_{\max}(H)]^2\right)$.

Theorem 2.24 (Concentration of two-point correlators)**.** *For any fixed bounded operators X and Y, denoting*

$$Y(\tau) = \exp\left(i\sqrt{n}H\tau\right) Y \exp\left(-i\sqrt{n}H\tau\right) \tag{2.104}$$

for any $\tau \in \mathbb{R}$, we have

$$\mathbb{P}\left(\frac{1}{2}\left|\operatorname{Tr}\left(XY(\tau)\rho_\beta\right) - \mathbb{E}\operatorname{Tr}\left(XY(\tau)\rho_\beta\right) \pm h.c.\right| \geq t\right) \leq 2e^{-t^2/(6n(5\beta^2+16\tau^2)||X||^2||Y||^2\Delta)}. \tag{2.105}$$

Recall the upper bound on the commutation index of the SYK model given in Theorem 2.10, as well as the $\Theta_q\left(\sqrt{n}\right)$-scaling of the expected maximal energy of the SYK model in this normalization [HO21]. Instantiating Theorem 2.21, Theorem 2.22, Theorem 2.23, and Theorem 2.24 with these parameters thus yields the following results for the SYK model.

Corollary 2.25. (SYK model is annealed.) *For the SYK model H_q^{SYK} where q is even,*

$$\frac{1}{n}\mathbb{E}\log Z_\beta \leq \frac{1}{n}\log \mathbb{E}Z_\beta \leq \frac{1}{n}\mathbb{E}\log Z_\beta + O_q(\beta^2 n^{-q/2}), \tag{2.106}$$

$$\mathbb{P}\left(\left|\operatorname{Tr}\left(X\rho_\beta\right) - \mathbb{E}\operatorname{Tr}\left(X\rho_\beta\right)\right| \geq t\right) \leq 2e^{-\Omega_q(\beta^{-2}n^{q/2-1}t^2)}, \tag{2.107}$$

$$\begin{aligned}\mathbb{P}\left(\left|\operatorname{Tr}\left(H_q^{SYK}\rho_\beta\right) - \mathbb{E}\operatorname{Tr}\left(H_q^{SYK}\rho_\beta\right)\right| \geq t\right) \\ \leq \begin{cases} 4e^{-\Omega_q\left(\beta^{-1}n^{q/2-1/2}t\right)} & t = \Omega(1+\beta n) \\ 4e^{-\Omega_q\left(\min\left(1,\beta^{-2}n^{-2}\right)n^{q/2}t^2\right)} & otherwise \end{cases},\end{aligned} \tag{2.108}$$

$$\mathbb{P}\left(\frac{1}{2}\left|\operatorname{Tr}\left(XY(\tau)\rho_\beta\right) - \mathbb{E}\operatorname{Tr}\left(XY(\tau)\rho_\beta\right) \pm h.c.\right| \geq t\right) \leq 2e^{-\Omega_q\left(\min\left(\beta^{-2},\tau^{-2}\right)n^{q/2-1}t^2\right)}. \tag{2.109}$$

X and Y are any fixed bounded operators.

Importantly, the above result shows that for the standard SYK model with $q = 4$, the quenched free energy in the limit of $n \to \infty$ always equal its annealed approximation for physical temperatures where $\beta = \Theta(\sqrt{n})$. This stands in stark contrast with spin glasses where a transition occurs for some critical temperature β_p into a clustered or 'glassy' phase.

The remainder of this appendix is concerned with establishing Theorems 2.21, 2.22, 2.23, 2.24 for concentration of various observables and free energies. Our general strategy takes advantage of the fact that *Lipschitz continuous* functions of Gaussian random variables exponentially concentrate. We begin by reviewing this concentration property in more detail.

Preliminaries

A function $f : \mathbb{R}^m \to \mathbb{R}$ is said to be *Lipschitz (continuous)* with constant $L \geq 0$ (referred to as the Lipschitz constant) if for all $x, y \in \mathbb{R}^m$, the following inequality holds:

$$|f(x) - f(y)| \leq L\|x - y\|_2, \tag{2.110}$$

where $\|x - y\|_2$ is the Euclidean distance between x and y. Intuitively, this means that the function f does not change too rapidly: the change in f's value is bounded by a linear multiple of the distance between x and y. It can also be seen from the mean-value theorem that the Lipschitz constant is bounded by the maximal gradient:

$$|f(x) - f(y)| \leq \sup_{z \in \mathbb{R}^m} \|\nabla f(z)\| \, \|x - y\|_2 , \tag{2.111}$$

a fact that we will later use.

It is known that a Lipschitz bound for a function f implies concentration when inputs to the function f are Gaussian.

Fact 2.2 (Gaussian concentration of Lipschitz functions, Theorem 2.26 of [Wai19b])**.** *Let $\vec{g} = (g_1, \ldots, g_m)$ be i.i.d. standard Gaussian variables, and $f : \mathbb{R}^m \to \mathbb{R}$ L-Lipschitz. Then for any $t \geq 0$:*

$$\mathbb{P}\big(|f(\vec{g}) - \mathbb{E}f(\vec{g})| \geq t\big) \leq 2e^{-t^2/(2L^2)}. \tag{2.112}$$

We also state a useful fact on the concentration of the operator norm of random matrices of the form of H.

Fact 2.3. (Concentration of the maximal eigenvalue [BBH21, Corollary 4.14]) *Let $\lambda_{max}(H)$ be the maximal eigenvalue of $H = m^{-1/2} \sum_{i=1}^{m} g_i A_i$, where $g_i \sim \mathcal{N}(0, 1)$. We have:*

$$\mathbb{P}\left(\lambda_{max}(H) - \mathbb{E}\lambda_{max}(H) \geq t\right) \leq \exp\left(-\frac{t^2}{2\Delta}\right). \tag{2.113}$$

In the course of proving our results we will also use an equivalent formulation of Fact 2.2 that follows from its sub-Gaussianity [RH23; Ver18; Wai19a].

Lemma 2.26 (Sub-Gaussian MGF bound, Lemma 1.5 of Ref. [RH23])**.** *Given a random variable X with sub-Gaussian concentration bound*

$$\mathbb{P}\big(|X - \mathbb{E}X| \geq t\big) \leq 2e^{-t^2/(2\sigma^2)}, \tag{2.114}$$

it holds that

$$\mathbb{E}[\exp\left(t\left(X - \mathbb{E}X\right)\right)] \leq \exp\left(4\sigma^2 t^2\right). \tag{2.115}$$

Finally, in what follows we use $\|\cdot\|$ to denote the operator norm:

$$\|X\| = \max_{v:\|v\|_2=1} v^\dagger X v, \tag{2.116}$$

where the maximum is taken over vectors v, and $\|\cdot\|_1$ to denote the trace norm:

$$\|X\|_1 = \mathrm{Tr}\left(\sqrt{X^\dagger X}\right). \tag{2.117}$$

Proof of Theorem 2.21

We directly compute the derivatives of $\log Z_\beta$ with respect to each Gaussian g_i

$$\begin{aligned}
\partial_{g_i} \log Z_\beta &= \frac{1}{Z_\beta} \mathrm{Tr}[\partial_{g_i} e^{\beta\sqrt{n}H}] && (2.118)\\
&= \frac{1}{Z_\beta} \mathrm{Tr}\left[\beta\sqrt{\frac{n}{m}} \int_0^1 e^{\beta\sqrt{n}H(1-s)} A_i e^{\beta\sqrt{n}H} ds\right] \\
&\qquad\qquad \text{(Derivative of matrix exponential [Wil67])} \\
&= \beta\sqrt{\frac{n}{m}} \mathrm{Tr}[A_i \rho_\beta]. \qquad \text{(Cyclic property of trace)}
\end{aligned}$$

Therefore, the Lipschitz constant L of $\log Z_\beta$ with respect to the disorder has the gradient bound

$$L^2 \leq \frac{\beta^2 n}{m} \sum_{i=1}^m \mathrm{Tr}[A_i \rho_\beta]^2 \leq \beta^2 n \Delta. \tag{2.119}$$

Now we can bound

$$\frac{\mathbb{E}\left[Z_\beta\right]}{\exp(\mathbb{E}[\log Z_\beta])} = \mathbb{E}\left[\exp(\log Z_\beta - \mathbb{E}[\log Z_\beta])\right] \leq \exp(4\beta^2 n\Delta). \tag{2.120}$$

The inequality uses Fact 2.2 and Lemma 2.26 with $t = 1$. Taking logarithms and rearrange to obtain

$$\frac{1}{n} \log \mathbb{E}[Z_\beta] \leq \frac{1}{n} \mathbb{E}[\log Z_\beta] + 4\beta^2 \Delta, \tag{2.121}$$

as stated.

Proof of Theorem 2.22

The result once again follows from a Lipschitz bound. We use the well-known expression for the derivative of a matrix exponential [Wil67]:

$$\partial_{g_i} \exp(H) = \int_0^1 \exp(tH) \left(\partial_{g_i} H\right) \exp((1-t)H)\, dt. \tag{2.122}$$

From the chain rule we then have

$$\partial_{g_i} \operatorname{Tr}\left(\rho_\beta X\right)$$

$$= \frac{\beta\sqrt{n}}{\sqrt{m}}\left(Z_\beta^{-1} \operatorname{Tr}\left(X \int_0^1 \exp\left(t\beta\sqrt{n}H\right) A_j \exp\left(\left(1-t\right)\beta\sqrt{n}H\right)\, dt\right)\right. \tag{2.123}$$

$$\left. - \operatorname{Tr}\left(X\rho_\beta\right) \operatorname{Tr}\left(A_j\rho_\beta\right)\right). \tag{2.124}$$

Consider now the operator:

$$\sigma_\beta \equiv Z_\beta^{-1} \int_0^1 \exp\left(t\beta\sqrt{n}H\right) X \exp\left(\left(1-t\right)\beta\sqrt{n}H\right)\, dt. \tag{2.125}$$

We can check that σ_β has trace norm bounded by $||X||$. Denoting by $\Sigma_i\left(\cdot\right)$ the ith singular value of $\cdot$ in nonincreasing order, we have by the majorization inequality [Bha97]:

$$\sum_i \Sigma_i\left(AB\right) \leq \sum_i \Sigma_i\left(A\right)\Sigma_i\left(B\right) \tag{2.126}$$

nd the product inequality (for Hermitian B) [HJ91]:

$$\Sigma_i\left(AB\right) \leq \|B\|\,\Sigma_i\left(A\right) \tag{2.127}$$

that

$$\sum_i \Sigma_i\left(\exp\left(t\beta\sqrt{n}H\right) X \exp\left(\left(1-t\right)\beta\sqrt{n}H\right)\right) \tag{2.128}$$

$$\leq \sum_i \Sigma_i\left(\exp\left(t\beta\sqrt{n}H\right) X\right)\Sigma_i\left(\exp\left(\left(1-t\right)\beta\sqrt{n}H\right)\right) \tag{2.129}$$

$$\leq \|X\| \sum_i \Sigma_i\left(\exp\left(t\beta\sqrt{n}H\right)\right)\Sigma_i\left(\exp\left(\left(1-t\right)\beta\sqrt{n}H\right)\right). \tag{2.130}$$

Finally, as $\exp\left(t\beta\sqrt{n}H\right)$ and $\exp\left(\left(1-t\right)\beta\sqrt{n}H\right)$ are Hermitian and positive semidefinite, their singular values are just their eigenvalues. As they are mutually diagonalizable,

$$\|X\| \sum_i \Sigma_i\left(\exp\left(t\beta\sqrt{n}H\right)\right)\Sigma_i\left(\exp\left(\left(1-t\right)\beta\sqrt{n}H\right)\right) \leq \|X\| \operatorname{Tr}\left(\exp\left(t\beta\sqrt{n}H\right)\exp\left(\left(1-t\right)\beta\sqrt{n}H\right)\right)$$

$$= \|X\|\, Z_\beta. \tag{2.131}$$

This implies that σ_β has trace norm bounded by $||X||$. However, while σ_β is Hermitian, it is not necessarily positive semidefinite. We proceed by writing the eigenvalue decomposition:

$$\sigma_\beta = \sum_{\lambda>0} \lambda \left|\lambda\right\rangle \left\langle\lambda\right| - \sum_{\lambda<0} |\lambda| \left|\lambda\right\rangle \left\langle\lambda\right| =: \sigma_\beta^+ - \sigma_\beta^-, \tag{2.132}$$

where $\sigma_\beta^\pm$ are Hermitian and positive semidefinite by construction, and each has trace norm bounded by $\|X\|$ as $\sum_\lambda |\lambda| = \left\|\sigma_\beta\right\|_1 \leq \|X\|$. By the cyclic property of the trace we can then write:

$$\partial_{g_i} \operatorname{Tr}\left(\rho_\beta X\right) = \frac{\beta\sqrt{n}}{\sqrt{m}} \left(\operatorname{Tr}\left(\sigma_\beta^+ A_j\right) - \operatorname{Tr}\left(\sigma_\beta^- A_j\right) - \operatorname{Tr}\left(X\rho_\beta\right) \operatorname{Tr}\left(A_j\rho_\beta\right)\right). \tag{2.133}$$

We thus have

$$\left\|\nabla_{\vec{g}} \operatorname{Tr}\left(\rho_\beta X\right)\right\|_2^2 = \frac{\beta^2 n}{m} \sum_j \left(\operatorname{Tr}\left(\sigma_\beta^+ A_j\right) - \operatorname{Tr}\left(\sigma_\beta^- A_j\right) - \operatorname{Tr}\left(X\rho_\beta\right) \operatorname{Tr}\left(A_j\rho_\beta\right)\right)^2 \tag{2.134}$$

$$\leq \frac{3\beta^2 n}{m} \sum_j \left(\left(\operatorname{Tr}\sigma_\beta^+ A_j\right)^2 + \left(\operatorname{Tr}\sigma_\beta^- A_j\right)^2 + \left(\operatorname{Tr} X\rho_\beta\right)^2 \left(\operatorname{Tr} A_j\rho_\beta\right)^2\right) \tag{2.135}$$

$$\leq 9\beta^2 n ||X||^2 \Delta. \tag{2.136}$$

The result then follows from Fact 2.2.

Proof of Theorem 2.23

We would like an analog of Equation (2.133) where the observable is H. Notice this is g_i-dependent and commutes with ρ_β. We get

$$\partial_{g_i} \operatorname{Tr}\left(\rho_\beta H\right) = \frac{1}{\sqrt{m}} \operatorname{Tr}\left(\rho_\beta A_j\right) + \frac{\beta\sqrt{n}}{\sqrt{m}} \left(\operatorname{Tr}\left(\rho_\beta H A_j\right) - \operatorname{Tr}\left(H\rho_\beta\right) \operatorname{Tr}\left(A_j\rho_\beta\right)\right). \tag{2.137}$$

Let $\lambda_{\max}(H)$ denote the maximal eigenenergy of H, and let $\mathcal{G}_s$ be the set of coefficients $\vec{g}$ where $\lambda_{\max}(H) \leq s + \mathbb{E}\lambda_{\max}(H)$. For $\vec{g} \in \mathcal{G}_s$

$$\begin{aligned}
&\left\|\nabla_{\vec{g}} \operatorname{Tr}\left(\rho_\beta H\right)\right\|_2^2 \\
&\quad= \frac{1}{m}\sum_j \left(\operatorname{Tr}\left(\rho_\beta A_j\right) + \beta\sqrt{n}\operatorname{Tr}\left(\rho_\beta H A_j\right) - \beta\sqrt{n}\operatorname{Tr}\left(H\rho_\beta\right)\operatorname{Tr}\left(A_j\rho_\beta\right)\right)^2 && (2.138)\\
&\quad\leq \frac{3}{m}\sum_j \left(\left(\operatorname{Tr}\rho_\beta A_i\right)^2 + \beta^2 n\left(\operatorname{Tr}\rho_\beta H A_i\right)^2 + \beta^2 n\left(\operatorname{Tr}\rho_\beta H\right)^2\left(\operatorname{Tr}\rho_\beta A_i\right)^2\right) && (2.139)\\
&\quad\leq 3\Delta\left(1 + 2\beta^2 n||H||^2\right) && (2.140)\\
&\quad\leq 3\Delta\left(1 + 2\beta^2 n\left(s + \mathbb{E}\lambda_{\max}(H)\right)^2\right). && (2.141)
\end{aligned}$$

That is, the function $\tau(g) := \operatorname{Tr}\left(\rho_\beta H\right)$ defined on the set $\mathcal{G}_s$ has a Lipschitz constant bounded by

$$L_{\mathcal{G}_s} := 3\Delta\left(1 + 2\beta^2 n\left(s + \mathbb{E}\lambda_{\max}(H)\right)^2\right). \tag{2.142}$$

We now use the Kirszbraun theorem.

Theorem 2.27 (Kirszbraun theorem, Ref. [Val45])**.** *Let $U \subset \mathbb{R}^{d_1}$, and assume $f : U \to \mathbb{R}^{d_2}$ is Lipschitz with Lipschitz constant L. Then, there exists $\hat{f} : \mathbb{R}^{d_1} \to \mathbb{R}^{d_2}$ with Lipschitz constant L such that $\hat{f}(x) = f(x)$ for all $x \in U$.*

In particular, there exists $\hat{\tau}(g)$ with Lipschitz constant given by Equation (2.142) such that $\hat{\tau}(g)$ agrees with $\operatorname{Tr}\left(\rho_\beta H\right)$ on $\mathcal{G}_s$. Furthermore, by Fact 2.3, $\mathbb{P}\left[g \notin \mathcal{G}_s\right] \leq$

$2\exp\left(-\frac{s^2}{2\Delta}\right)$. We use these two properties to calculate:

$$\mathbb{P}\left[\left\|\mathrm{Tr}\left(H\rho_\beta\right) - \mathbb{E}\left[\mathrm{Tr}\left(H\rho_\beta\right)\right]\right\| \geq t\right] \tag{2.143}$$

$$\leq \inf_s \mathbb{P}\left[\left\|\mathrm{Tr}\left(H\rho_\beta\right) - \mathbb{E}\left[\mathrm{Tr}\left(H\rho_\beta\right)\right]\right\| \geq t \wedge g \in \mathcal{G}_s\right] + 2\exp\left(-\frac{s^2}{2\Delta}\right) \tag{2.144}$$

$$= \inf_s \mathbb{P}\left[\left\|\hat{\tau}\left(g\right) - \mathbb{E}\left[\hat{\tau}\left(g\right)\right]\right\| \geq t \wedge g \in \mathcal{G}_s\right] + 2\exp\left(-\frac{s^2}{2\Delta}\right) \tag{2.145}$$

$$\leq \inf_s \mathbb{P}\left[\left\|\hat{\tau}\left(g\right) - \mathbb{E}\left[\hat{\tau}\left(g\right)\right]\right\| \geq t\right] + 2\exp\left(-\frac{s^2}{2\Delta}\right) \tag{2.146}$$

$$\leq \inf_s 2\exp\left(-\frac{t^2}{6\Delta\left(1 + 2\beta^2 n\left(s + \mathbb{E}\left[\lambda_{\max}\left(H\right)\right]\right)^2\right)}\right) + 2\exp\left(-\frac{s^2}{2\Delta}\right) \tag{2.147}$$

$$\leq \inf_s 2\exp\left(-\frac{t^2}{6\Delta\left(1 + 4\beta^2 n\left(s^2 + \mathbb{E}\left[\lambda_{\max}\left(H\right)\right]^2\right)\right)}\right) + 2\exp\left(-\frac{s^2}{2\Delta}\right) \tag{2.148}$$

$$\leq \inf_s 2\exp\left(-\frac{t^2}{24\beta^2\Delta n\left(s^2 + 1/(4\beta^2 n) + \mathbb{E}[\lambda_{\max}(H)]^2\right)}\right) + 2\exp\left(-\frac{s^2}{2\Delta}\right). \tag{2.149}$$

Setting $s^2 = \sqrt{t^2/(12\beta^2 n) + \alpha^2} - \alpha$ where $\alpha = \frac{1}{2}\left(1/(4\beta^2 n) + \mathbb{E}[\lambda_{\max}(H)]^2\right)$ gives the desired result.

Proof of Theorem 2.24

Completely analogously to the proof of Theorem 2.22 we have

$$\begin{aligned}&\partial_{g_i}\mathrm{Tr}\left(\rho_\beta XY\left(\tau\right)\right)\\ &\quad= \frac{\beta\sqrt{n}}{\sqrt{m}}\left(\mathrm{Tr}\left(\sigma_\beta A_j\right) - \mathrm{Tr}\left(XY\left(\tau\right)\rho_\beta\right)\mathrm{Tr}\left(A_j\rho_\beta\right)\right) + \mathrm{Tr}\left(\rho_\beta X\partial_{g_i}Y\left(\tau\right)\right),\end{aligned} \tag{2.150}$$

where

$$\sigma_\beta = Z_\beta^{-1}\int_0^1 \exp\left(t\beta\sqrt{n}H\right)XY\left(\tau\right)\exp\left(\left(1-t\right)\beta\sqrt{n}H\right)\,dt. \tag{2.151}$$

We now focus on the final term of Equation (2.150). We have

$$\partial_{g_i} Y(\tau)$$

$$= \frac{i\tau\sqrt{n}}{\sqrt{m}} \left(\int_0^1 \exp\left(i\tau\sqrt{n}tH\right) A_j \exp\left(i\tau\sqrt{n}(1-t)H\right) \, dt \right) Y \exp\left(-i\tau\sqrt{n}H\right) \tag{2.152}$$

$$+ \text{h.c.} \tag{2.153}$$

$$= \frac{i\tau\sqrt{n}}{\sqrt{m}} \left[\tilde{A}_{j|\tau}, Y(\tau) \right], \tag{2.154}$$

where $\tilde{A}_{j|\tau}$ is the Hermitian, time-averaged operator:

$$\tilde{A}_{j|\tau} = \int_0^1 \exp\left(i\tau\sqrt{n}tH\right) A_j \exp\left(-i\tau\sqrt{n}tH\right) \, dt \tag{2.155}$$

$$= \frac{1}{\tau} \int_0^\tau \exp\left(it\sqrt{n}H\right) A_j \exp\left(-it\sqrt{n}H\right) \, dt. \tag{2.156}$$

We have

$$\frac{1}{4} \left\| \nabla_{\vec{g}} \operatorname{Tr}\left(\rho_\beta XY(\tau)\right) \pm \text{h.c.} \right\|_2^2 \tag{2.157}$$

$$= \frac{n}{4m} \sum_j || \beta \operatorname{Tr}\left(\sigma_\beta A_j\right) - \beta \operatorname{Tr}\left(XY(\tau)\rho_\beta\right) \operatorname{Tr}\left(A_j \rho_\beta\right)$$

$$+ i\tau \operatorname{Tr} \rho_\beta X \left[\tilde{A}_{j|\tau}, Y(\tau) \right] \pm \text{h.c.} ||^2 \tag{2.158}$$

$$\leq \frac{3n}{4m} \sum_j \Big(||\beta \operatorname{Tr}\left(\sigma_\beta A_j\right) \pm \text{h.c.}||^2 + ||\beta \operatorname{Tr}\left(XY(\tau)\rho_\beta\right) \operatorname{Tr}\left(A_j\rho_\beta\right) \pm \text{h.c.}||^2$$

$$+ ||\tau \operatorname{Tr}\left(\rho_\beta X \left[\tilde{A}_{j|\tau}, Y(\tau)\right]\right) \pm \text{h.c.}||^2 \Big). \tag{2.159}$$

The first two terms are conceptually identical to Equation (2.134). To bound the final term, we define

$$\mu^{\pm}_{\beta,\tau} := Y(\tau)\rho_\beta X \pm \text{h.c.}, \tag{2.160}$$

$$\nu^{\pm}_{\beta,\tau} := \rho_\beta XY(\tau) \pm \text{h.c.}, \tag{2.161}$$

such that

$$\operatorname{Tr}\left(\rho_\beta X \left[\tilde{A}_{j|\tau}, Y(\tau)\right]\right) \pm \text{h.c.} = \operatorname{Tr}\left(\mu^{\pm}_{\beta,\tau} \tilde{A}_{j|\tau}\right) - \operatorname{Tr}\left(\nu^{\pm}_{\beta,\tau} \tilde{A}_{j|\tau}\right). \tag{2.162}$$

Note that $\mu^+_{\beta,\tau}, \nu^+_{\beta,\tau}, i\mu^-_{\beta,\tau}, i\nu^-_{\beta,\tau}$ are all Hermitian by construction, and each has trace norm bounded by $2\|X\|\|Y\|$ by Hölder's inequality. Just as in Equation (2.132),

each can be considered as the difference of two positive semidefinite matrices with trace norm bounded by $2\|X\|\|Y\|$. Putting everything together yields

$$\frac{1}{4}\left\|\nabla_{\vec{g}} \operatorname{Tr}\left(\rho_\beta XY(\tau)\right) \pm \text{h.c.}\right\|_2^2 \tag{2.163}$$

$$\leq \frac{3n}{4m}\sum_j \left(||\beta \operatorname{Tr}\left(\sigma_\beta A_j\right) \pm \text{h.c.}||^2 + ||\beta \operatorname{Tr}\left(XY(\tau)\rho_\beta\right)\operatorname{Tr}\left(A_j\rho_\beta\right) \pm \text{h.c.}||^2\right.$$

$$\left. + ||\tau \operatorname{Tr}\left(\rho_\beta X\left[\tilde{A}_{j|\tau}, Y(\tau)\right]\right) \pm \text{h.c.}||^2\right) \tag{2.164}$$

$$\leq \frac{3n}{4}\left(16\beta^2||X||^2||Y||^2\Delta + 4\beta^2||X||^2||Y||^2\Delta + 64\tau^2||X||^2||Y||^2\Delta\right) \tag{2.165}$$

$$= 3n\left(5\beta^2 + 16\tau^2\right)||X||^2||Y||^2\Delta. \tag{2.166}$$

Chapter 3

SUM OF SQUARES SPECTRAL AMPLIFICATION

Simulating quantum many-body systems is one of the most heralded and valuable applications of quantum computing. Although efficient quantum algorithms exist for numerous quantum simulation problems [Kit95; KOS07; Llo96; Wie+10; Ber+15; LC17b], further improvements are essential to fully realize robust large-scale quantum simulation when accounting for the considerable constant factor overhead associated with fault-tolerant quantum computation [Ber+24b; Rub+23; Bab+21]. Quantum algorithm improvements are expected to arise from exploiting specific structures of the simulation instance, thus mitigating worst-case computational costs. Examples of such are exploiting symmetries in chemistry problems or locality in lattice models [LI23; Haa+23]. Here, we introduce a strategy that exploits the low-energy properties of quantum states to improve quantum simulation algorithms. Our framework relies on two distinct algorithmic ideas: sum-of-squares (SOS) representations of Hamiltonians and spectral amplification (SA).

The low-energy setting is one of the most promising areas in quantum simulation. This setting is relevant in the study of matter at low temperatures, including quantum phase transitions and the computation of ground-state properties. In this setting, some of the most important simulation tasks are: i) Estimating the energy of a quantum system with respect to certain quantum state [KOS07], ii) Estimating the ground state energy of a quantum system by phase estimation [Kit95], and iii) Simulating the time evolution of a quantum state under a Hamiltonian [Llo96]. We will show how SOSSA significantly improves the gate complexity of generic methods for all these problems. To this end, SOSSA combines two key ideas: first, it produces a suitable SOS representation of the given Hamiltonian H plus an energy shift β to make the Hamiltonian positive semidefinite and with a small ground state energy, and then it uses SA to amplify the low eigenvalues and the corresponding energy gaps of $H + \beta$ [SB13; LC17a; ZS24].

During the SOS step, $H + \beta$ is processed classically and represented as a sum of positive terms. This can modify properties of H that have an impact on the complexity of the simulation algorithm, such as an ℓ_1-norm that depends on its presentation and the ground state energy. Ideally, the SOS representation is such

	LCU $\lambda = \lambda_{\mathrm{LCU}}$	**Termwise SA** $\lambda = \lambda_{\mathrm{LCU}}$ $\Delta = \Delta_{\mathrm{LCU}}$	**SOSSA** $\lambda = \lambda_{\mathrm{SOS}}$ $\Delta = \Delta_{\mathrm{SOS}}$
Energy estimation	$\Theta(\lambda/\epsilon)$ [KOS07]	$\Theta\big(\sqrt{\Delta\lambda}/\epsilon\big)$ Theorem 3.12	
Phase estimation	$O\big((\lambda/(\sqrt{p}\epsilon))\log\frac{1}{p}\big)$ [Ber+24b]	$O\big(\big(\sqrt{\Delta\lambda}/(\sqrt{p}\epsilon)\big)\log\frac{1}{p}\big)$ Theorem 3.15	
Time evolution	$\Theta(\lambda t + \log\frac{1}{\epsilon})$ [LC17b; LC19]	$\Theta\big(\sqrt{\Delta\lambda}t + \sqrt{\lambda/\Delta}\log\frac{1}{\epsilon}\big)\big)$ [ZS24]	
SYK model	$\lambda_{\mathrm{LCU}} \sim N^2$	$\sqrt{\Delta_{\mathrm{LCU}}\lambda_{\mathrm{LCU}}} \sim N^2$	$\sqrt{\Delta_{\mathrm{SOS}}\lambda_{\mathrm{SOS}}} \sim N^{\frac{3}{2}}$

Table 3.1: (a) Simulation tasks on Hamiltonian H to precision ϵ using LCU, termwise SA, and SOSSA. Presented are the query complexities assuming access to the Hamiltonian only via the block-encodings H/λ_{LCU}, $H_{\mathrm{SA}}/\sqrt{\lambda_{\mathrm{LCU}}}$, and $H_{\mathrm{SOSSA}}/\sqrt{\lambda_{\mathrm{SOS}}}$, respectively. The gate complexities can be determined from the gate complexities of each block-encoding, which can be different, and the additional arbitrary gates in the algorithm. We show: i) estimation of the energy $E = \langle\psi|H|\psi\rangle$, ii) phase estimation of the ground state energy E with initial state $|\psi\rangle$ and ground-state $|\psi_0\rangle$ satisfying $p = |\langle\psi\,|\psi_0\rangle\,|^2 > 0$, and iii) time-evolution to implement e^{-itH} on a state of energy at most E. We assume $-\lambda_{\mathrm{LCU}} \leq E \leq -\lambda_{\mathrm{LCU}} + \Delta_{\mathrm{LCU}}$ for termwise SA and $-\beta \leq E \leq -\beta + \Delta_{\mathrm{SOS}}$ for SOSSA. The lower bound on the ground state energy $-\beta$ is obtained in the SOS step and implies $\Delta_{\mathrm{SOS}} \ll \Delta_{\mathrm{LCU}}$. We also provide adaptive algorithms with improved complexities that do not require knowledge of Δ_{LCU} or Δ_{SOS}. (b) Normalization factors for the SYK model demonstrate an asymptotic speedup in system size N using an appropriate SOS. The gate complexities for all block-encodings is similar in this example.

that $-\beta$ is a tight lower bound on the ground state energy and that the square roots of the positive terms are not too difficult to simulate. This step relates to the well-known optimization task of finding the best SOS lower bound on the ground state energy, which can be solved efficiently in classical preprocessing using semi-definite programming. During the SA step we produce a different Hamiltonian whose eigenvalues are the square roots of those of $H + \beta\mathbb{1}$. Since the low-lying eigenvalues like the ground state energy are now 'amplified', this can alleviate some resources due to, for example, having less stringent requirements in the precision of an estimate. While the SOS step might introduce additional overheads from the complexity of the terms, the SA step has the potential to reduce it, and the method is useful if the overall combination still provides an improvement in gate complexities. Notably, we show this occurs in interesting systems.

In this article we provide the theory of SOSSA and then use it to construct quantum

algorithms that improve over prior art. Our approach is to first improve the query complexities of simulation methods based on SA, which will ultimately improve the gate complexities when the SOS representation is considered. Specifically, for: i) energy estimation, where in contrast with Ref. [Sim+24] our algorithm does not require an upper bound on the expectation to be estimated and works for arbitrary (block-encoded) operators, and ii) ground-state phase estimation, where we generalize results in Ref. [Low+25] in that we do not need an upper bound on the energy and our method works even if the overlap between the trial state and ground state is $p < 1$. The results are summarized in Table 3.1, where we also include those on time-evolution from Ref. [ZS24] for reference. These SA-based algorithms improve the linear query complexities of generic methods on λ to $\sqrt{\Delta\lambda}$, where λ is a parameter related to the norm of the Hamiltonian—the largest possible energy—and Δ is a parameter related to the low energy of the initial state. The relevant low-energy instances arise then when $\Delta \ll \lambda$.

While attaining improved query complexities is key to our approach, we are ultimately concerned with gate complexities. These can be determined from the costs of implementing the corresponding queries, which depend on the SOS representation and the difficulty of simulating its terms. Then, to demonstrate the power of SOSSA, we analyze its performance in applications. In the context of ground state energy estimation of quantum chemistry, recent findings demonstrate that SOSSA provides the best gate complexities currently known [Low+25]. We corroborate these findings further by applying SOSSA to ground state energy estimation of the Sachdev-Ye-Kitaev (SYK) model, where it provides an asymptotic speedup over generic methods by a factor of square root of system size (see Table 3.1). Given that the SYK model exemplifies a strongly correlated condensed matter system, these results suggest the general applicability of SOSSA to many quantum simulation problems.

We conclude that SOSSA provides a useful framework for several quantum simulation tasks, and expect it to be applied to other systems. Furthermore, we complement these findings by providing tight lower bounds for energy and phase estimation that show our quantum algorithms are query optimal in the low-energy setting, and provide a low-depth version for expectation estimation, which scales with the standard quantum limit, that might be of independent interest for near-term applications.

Last, we remark that while SOSSA concerns the combination of SOS representations and SA, each has been extensively studied in prior work. SOS is used in the context of

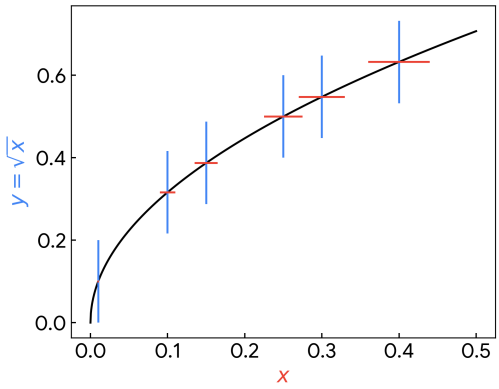

Figure 3.1: Uncertainty propagation through the square root function. By constructing the square root of a positive semidefinite Hamiltonian, we are able to amplify the small eigenvalues due to the divergent behavior of the derivative of $\sqrt{x}$ near zero. Here $x \geq 0$ denotes a rescaled eigenvalue and the relevant low-energy regime occurs for $x \ll 1$. The error bars illustrate that an estimate of an eigenvalue E to precision ϵ can be obtained through an estimate of $\sqrt{E}$ to precision $O(\epsilon/\sqrt{E})$, which becomes coarser as E decreases, thereby reducing the resources needed for phase estimation. The approach can be generalized to other functions $f(x)$ whose derivatives are large or divergent as $x \to 0$ and that could arise from other related constructions, such as the quantum-walk in Ref. [Low+25], where $f(x) = \arccos(x - 1)$ resulted in a similar quadratic amplification.

approximation algorithms [GW95; PNA10], lower bounds on ground-state energies in quantum chemistry [Maz06; Nak+01], and characterizing quantum correlations [NPA08]. These methods often appear in the pseudomoment picture, which is dual to the SOS optimization that we consider here. SA was used to achieve a quadratic quantum speedup within the context of adiabatic quantum computation [SB13] and was more recently used to obtain improved quantum algorithms for simulating time evolution and phase estimation on low-energy states [LC17a; ZS24; Low+25].

SOSSA The goal of SOSSA is to improve the gate complexity of simulation tasks in the low-energy sector. To this end, SOSSA uses SA for reducing query complexities first. Quantum signal processing and the related quantum singular value transform provide the modern machinery for quantum simulation tasks, relying on access to a block-encoding of a Hamiltonian H [LC19; Gil+19; Mar+21]. This is a unitary acting on an enlarged space that contains the matrix H/λ in one of its blocks, where $\lambda \geq \|H\|$ is needed for normalization. In quantum simulation algorithms we are often interested in both, the query and gate complexities. We

define the query complexity to be the number of times the block-encodings are used (including their inverses), and the gate complexity to be the total number of two-qubit gates to implement the algorithm, which includes the number of gates to implement the block-encodings in applications.

A standard approach to construct the block-encoding uses an efficient presentation of H as a linear combination of unitaries (LCU) [CW12; Ber+15]. Suppose a Hamiltonian H acts on a system of N qubits and is presented as a sum of terms as

$$H = \sum_{j=0}^{R-1} g_j \sigma_j \ , \tag{3.1}$$

where $g_j \in \mathbb{R}$ are coefficients and σ_j are (tensor) products of N Pauli matrices (i.e. Pauli strings). The query complexity of various quantum simulation tasks through this LCU scales linearly in the so-called ℓ_1-norm of the Hamiltonian:

$$\lambda_{\mathrm{LCU}} := \sum_{j=0}^{R-1} |g_j| \ . \tag{3.2}$$

The LCU column of Table 3.1 presents some known results.

SA allows one to get around the linear cost in λ_{LCU} in simulation tasks involving states of low energy. SA can only be applied to Hamiltonians that are positive semidefinite. One way to achieve SA is to produce a square root of the Hamiltonian, so that the low-energy spectrum is amplified. Intuitively, this amplification arises because the square root function is steep near zero, i.e. $\sqrt{x} \gg x$ for $x \ll 1$, where x denotes a rescaled eigenvalue. See Figure 3.1 for an illustration. For phase estimation, a precise estimate of a small eigenvalue of the original Hamiltonian can then be obtained by squaring a coarser estimate of an eigenvalue of its square root. For time evolution, the amplified low-energy spectrum allows for a more efficient approximation to the evolution operator using lower degree polynomials in quantum signal processing. In both cases, SA can lead to a lower query complexity than in the worst case. See Appendix 3.2 for results on SA.

The easiest way to apply SA to a Hamiltonian presented as in Eq. (3.1) is to add a shift to each Pauli string σ_j so that it becomes positive semidefinite; for example, $\mathbb{1} \pm \sigma_j \succeq 0$, where $\mathbb{1}$ is the identity matrix. We call this approach termwise SA, but also anticipate that this approach is not generally effective. During this preprocessing, the shift produces

$$H + \lambda_{\mathrm{LCU}} \mathbb{1} = 2 \sum_{j=0}^{R-1} |g_j| \Pi_j \ , \tag{3.3}$$

where $\Pi_j := \frac{1}{2}(\mathbb{1} + \mathrm{sign}(g_j)\sigma_j)$ are simple orthogonal projectors since the Pauli strings σ_j have eigenvalues ± 1. We can define the spectral amplified operator for this specific example to be

$$H_{\mathrm{SA}} = \sqrt{2}\sum_{j=0}^{R-1} |g_j|^{\frac{1}{2}} |j\rangle \otimes \Pi_j = \sqrt{2}\begin{pmatrix} |g_1|^{\frac{1}{2}}\Pi_1 \\ \vdots \\ |g_R|^{\frac{1}{2}}\Pi_R \end{pmatrix}. \tag{3.4}$$

This is a rectangular matrix that acts as a square root of $H + \lambda_{\mathrm{LCU}}\mathbb{1}$ since

$$H_{\mathrm{SA}}^{\dagger} H_{\mathrm{SA}} = H + \lambda_{\mathrm{LCU}}\mathbb{1}. \tag{3.5}$$

Consider now a simulation problem where the corresponding quantum states are supported in the low-energy subspace. We can introduce a parameter Δ_{LCU} that quantifies the low-energy assumption. In particular, for energy estimation we will assume our states of interest $|\psi\rangle$ to satisfy $E = \langle\psi| H |\psi\rangle \leq -\lambda_{\mathrm{LCU}}+\Delta_{\mathrm{LCU}}$. For phase estimation we will assume $\langle\xi|\psi\rangle = 0$ for every eigenstate $|\xi\rangle$ of H of eigenvalue $\langle\xi| H |\xi\rangle > -\lambda_{\mathrm{LCU}} + \Delta_{\mathrm{LCU}}$. For these problems, by using the block-encoding of $H_{\mathrm{SA}}/\sqrt{\lambda_{\mathrm{LCU}}}$ instead of that of H/λ_{LCU}, SA allows us to achieve an improvement in the query complexity. See the termwise SA results in Table 3.1, proven in Thms. 3.12 and 3.15. The net effect is an improved query complexity linear in $\sqrt{\Delta_{\mathrm{LCU}}\lambda_{\mathrm{LCU}}}$ rather than λ_{LCU}. (We can think of $\sqrt{\Delta_{\mathrm{LCU}}\lambda_{\mathrm{LCU}}}$ as the new effective ℓ_1-norm when considering the low-energy subspace.) The improvement occurs when $\Delta_{\mathrm{LCU}} \ll \lambda_{\mathrm{LCU}}$, which is possible in specific instances where the Hamiltonian H is close to 'frustration-free' [SB13].

For this example, an improvement in query complexity ultimately gives an improvement in gate complexity, because the corresponding block-encodings can be implemented with similar gate costs (Lemma 3.6). This would readily achieve the goal of SOSSA, however, for general and frustrated H, the termwise preprocessing outlined might result in a Δ_{LCU} that is comparable to λ_{LCU}. This severely limits the general applicability of termwise SA, and a different kind of preprocessing is desirable to have a significant reduction in gate costs.

Our strategy is then to use a different representation of the Hamiltonian plus a shift, as an SOS of more general operators B_j:

$$H + \beta\mathbb{1} = \sum_{j=0}^{R-1} B_j^{\dagger} B_j \,. \tag{3.6}$$

Using this representation, which implies $H+\beta \succeq 0$, gives also a path to applying SA to general Hamiltonians. Indeed, the previous termwise SA is an example of SOS where $\beta = \lambda_{\mathrm{LCU}}$ and $B_j = |g_j|^{\frac{1}{2}}\Pi_j$. However, we can now consider the B_j's to be linear combinations of Pauli strings σ_j or more general operators. It is important, however, to constrain the B_j's so that they can be efficiently block-encoded; see Chapter 3.

SOSSA then reduces the overall gate complexity by first reducing the query complexity via SA as much as possible. We do this by finding 'good' SOS representations that have the following properties. Let

$$H_{\mathrm{SOSSA}} := \sum_{j=0}^{R-1} |j\rangle \otimes B_j = \begin{pmatrix} B_1 \\ \vdots \\ B_R \end{pmatrix} \tag{3.7}$$

be such that $H_{\mathrm{SOSSA}}^{\dagger} H_{\mathrm{SOSSA}} = H + \beta\mathbb{1}$. Then, we wish to block-encode of $H_{\mathrm{SOSSA}}/\sqrt{\lambda_{\mathrm{SOS}}}$ efficiently, for some normalization factor $\lambda_{\mathrm{SOS}} \geq \|H + \beta\mathbb{1}\|$, which is often very different from λ_{LCU}. Additionally, we wish for the lower bound $-\beta$ to be as close to the ground state energy of H as possible. For energy estimation, we will assume $\langle\psi| H |\psi\rangle \leq -\beta + \Delta_{\mathrm{SOS}}$ and for phase estimation the low-energy state is supported on the subspace of energies in $[-\beta, -\beta + \Delta_{\mathrm{SOS}}]$. These parameters give $\sqrt{\Delta_{\mathrm{SOS}}\lambda_{\mathrm{SOS}}}$, which determines the query complexity when using SA with the Hamiltonian $H + \beta\mathbb{1}$; see the SOSSA results in Table 3.1. The goal of the good SOS representation is to satisfy $\Delta_{\mathrm{SOS}}\lambda_{\mathrm{SOS}} \ll \Delta_{\mathrm{LCU}}\lambda_{\mathrm{LCU}}$ and hence improve upon the query complexity of termwise SA.

After obtaining improved query complexities from the SOS representation, we wish to determine the resulting gate complexities, and we need to account for the gate cost of implementing the block-encoding of $H_{\mathrm{SOSSA}}/\sqrt{\lambda_{\mathrm{SOS}}}$. Unfortunately, this gate cost might be higher than that of implementing the block-encoding of $H_{\mathrm{SA}}/\sqrt{\lambda_{\mathrm{SA}}}$, since the operators B_j are more general and possibly more difficult to simulate. Nevertheless, we find that these two competing effects –the improvement in query complexity versus the increase in the gate complexity of each query– can still give a significant improvement overall, as demonstrated by the examples we studied.

Hence, to obtain improved gate complexities, SOSSA uses SA in combination with good SOS representations to reduce the query complexity as much as possible, even when the gate cost per query can increase.

SOS optimization for SOSSA To achieve greater SA, we would like Δ_{SOS} to be as small as possible, so we would like $-\beta$ to be a tight lower bound on the ground state energy. In this section, we show that optimizing the lower bound $-\beta$ can be achieved efficiently in classical preprocessing using semi-definite programming. Optimization of $-\beta$ alone does not directly optimize the total gate cost but can instead improve the query complexity. However, it can be a useful starting point as demonstrated in Ref. [Low+25]. We will see in Chapter 3 that this optimization is sufficient to achieve asymptotic speedups in the total gate cost for the SYK model.

We can start by selecting an ansatz for the B_j's as low-degree polynomials of some natural operator set $\mathcal{B}$, the 'SOS algebra', of size L. For example, given the Hamiltonian on N qubits, we can choose a constant $k \geq 1$ and let $\mathcal{B}$ be the set of $L = O(N^k)$ monomials, each being a product of k Pauli operators. Then, $B_j(\vec{b}_j) \in \mathcal{B}$ denotes a polynomial of Pauli strings of weight or degree k, where $\vec{b}_j \in \mathbb{C}^L$ are the coefficients of the polynomial. For a fermionic Hamiltonian, we could let $\mathcal{B}$ be the set of degree-k products of fermionic creation and annihilation operators instead, or equivalently in the Majorana operators. Often, we may not want to include all degree-k terms in $\mathcal{B}$, but only a subset like nearest-neighbor products when considering a system in a lattice.

Having selected the set $\mathcal{B}$, we can optimize the lower bound on the ground state energy while constraining the coefficients $\vec{b}_j$ so that Eq. (3.6) is satisfied. This defines the following program:

$$\min_{\vec{b}_j \in \mathbb{C}^L} \beta \quad \text{s.t.} \quad H + \beta \mathbb{1} = \sum_{j=0}^{R-1} B_j(\vec{b}_j)^\dagger B_j(\vec{b}_j).$$

This optimization problem can be directly translated into an SDP whose dimension is polynomial in the operator basis size L and whose linear constraints are derived from the algebra of the polynomials; see Section 3.3. Mathematical problems of this form can be solved in polynomial time with respect to the number of variables and constraints [BV04; BM03; PRW06]. The resulting number of terms R is related to the rank of the primal variable of the SDP and is also polynomial in L.

The dual problem of this SDP is the pseudomoment problem [PNA10; Erd78] where it is well understood that increasing the complexity of B_j by for example allowing a larger set of monomials in $\mathcal{B}$, and thus the size of the SDP, $-\beta$ can be made arbitrarily close to the true ground state energy, implying a smaller energy gap Δ_{SOS}.

Nevertheless, for the quantum algorithm, we note that the savings in complexity from a smaller Δ_{SOS} might be obscured by the higher complexity of simulating the terms B_j. Additionally, the cost of preprocessing from solving the SDP also increases with L, and this classical complexity can also be a factor of consideration in the overall optimization for a specific application.

Application to SYK We can now display our primary example application of the SOSSA framework to the SYK model. We observe an asymptotic speedup in phase estimation of the ground state energy over the standard approach that uses the block-encoding of the full Hamiltonian, which is constructed from an LCU presentation in terms of Pauli strings. The speedup reduces the gate complexity by a factor $\sqrt{N}$, where N is the number of modes. Let $\{\gamma_1, \ldots, \gamma_N\}$ be Majorana operators satisfying anticommutation relations:

$$\gamma_a \gamma_b + \gamma_b \gamma_a = 2\delta_{ab} \mathbb{1} \ . \tag{3.8}$$

The SYK model is described by a fermionic Hamiltonian containing all degree-4 Majorana terms whose coefficients are random Gaussians; that is,

$$H_{\text{SYK}} = \frac{1}{\sqrt{\binom{N}{4}}} \sum_{a,b,c,d} g_{abcd} \gamma_a \gamma_b \gamma_c \gamma_d \ , \tag{3.9}$$

$$g_{abcd} \sim \mathcal{N}(0, 1) \text{ i.i.d. .} \tag{3.10}$$

We will apply SOSSA to the SYK model using degree-2 Majorana operators for the SOS representation, which is described in detail in Section 3.4. That is, the set $\mathcal{B}$ used for the SOS representation are quadratic in the γ_a's Using random matrix theory, in Lemma 3.18 we show that this readily achieves a lower bound where $\beta = O(N)$. Furthermore, the spectrum of H_{SYK} is known to be contained in the interval $[-c\sqrt{N}, c\sqrt{N}]$ for some constant $c > 0$, and hence the energy gap obeys $\Delta_{\text{SOS}} = O(N)$.

To compare the performance of SOSSA with standard approaches or termwise SA, we first note H_{SYK} has number of terms scales like $\sim N^4$. Hence, the asymptotic gate complexities to construct the necessary block-encodings for the LCU, termwise SA, and SOSSA approaches is $\Theta(N^4)$ in all cases.[1] We can thus focus on comparing the query complexities only, which are determined by λ_{LCU}, $\sqrt{\Delta_{\text{LCU}} \lambda_{\text{LCU}}}$ and

[1] We can achieve block encodings with gate complexities $O(N^4)$, and $\Omega(N^4)$ is a lower bound since this is number of parameters in the Hamiltonian.

$\sqrt{\Delta_{\mathrm{SOS}}\lambda_{\mathrm{SOS}}}$. We have

$$\lambda_{\mathrm{LCU}} = \frac{1}{\sqrt{\binom{N}{4}}} \sum_{a,b,c,d} |g_{abcd}| \tag{3.11}$$

for Eq. (3.9), and hence $\lambda_{\mathrm{LCU}} = \Theta(N^2)$ with high probability. Termwise SA does not provide any asymptotic improvement, since the scaling of the ground state energy E is dominated by λ_{LCU} and so Δ_{LCU} will asymptotically scale linearly with λ_{LCU}. Nevertheless, in Section 3.4 we show that we can implement a block encoding of H_{SOSSA} with squared normalization factor $\lambda_{\mathrm{SOS}} = O(N^2)$. Our approach involves the double factorization technique described in Section 3.4 [Bur+21b]. Since the energy gap is $\Delta_{\mathrm{SOS}} = O(N)$, we have $\sqrt{\Delta_{\mathrm{SOS}}\lambda_{\mathrm{SOS}}} = \Theta(N^{\frac{3}{2}})$ for SOSSA, giving an asymptotic speedup by a factor of $\sim \sqrt{N}$ over LCU and termwise SA.

Higher-degree SOS algebras. In Ref. [HO22] it was demonstrated that degree-2 Majorana SOS is unable to recover this scaling, and is limited to a lower bound scaling like $\beta = O(N)$ with high probability. Reference [HO22] introduces a degree-3 SOS, where the B_j's in the SOS representation are cubic in the γ_a's, and which is able to achieve a tighter lower bound to the ground state energy where $\beta = \Omega(\sqrt{N})$. Their SOS algebra $\mathcal{B}$ does not contain all degree-3 Majorana monomials, but rather uses only a particular fragment of the degree-3 terms of size $L = O(N)$.

We can consider applying SOSSA with the degree-3 SOS from Ref. [HO22]. The energy gap will improve to $\Delta_{\mathrm{SOS}} = O(\sqrt{N})$. However, the block-encoding normalization factor scales now as $\lambda_{\mathrm{SOS}} \sim N^{\frac{7}{2}}$; see the bound on λ_{SOS} in Section 3.3. We obtain $\sqrt{\Delta_{\mathrm{SOS}}\lambda_{\mathrm{SOS}}} = \Theta(N^2)$, recovering the query complexity scaling of LCU and termwise SA. Not considering gate cost of block-encoding, which is likely higher than the termwise LCU method, we recover the termwise SA query cost suggesting higher gate complexities.

This further illustrates that optimization of $-\beta$ alone does not guarantee that we will ultimately obtain improved gate complexities. The example analyzed here, which regarded a more involved optimization, readily provided a factor $\sqrt{\Delta_{\mathrm{SOS}}\lambda_{\mathrm{SOS}}}$ that is asymptotically comparable to $\sqrt{\Delta_{\mathrm{SOS}}\lambda_{\mathrm{SOS}}}$. This is a consequence of considering more complex generators for the SOS representations. In addition to the complexity scalings due to these factors, we often expect the block-encodings of the SOS representations to be less efficient to implement. These effects are important when comparing the overall gate complexity of SOSSA to that of termwise SA.

Discussion We described a framework for fast quantum simulation of low-energy states based on SA and SOS representations. To this end, we developed quantum algorithms for energy and phase estimation that improve over prior art and showed that SOSSA gives asymptotic improvements in gate costs with respect to traditional methods when applied to the SYK Hamiltonian. With the addition of Ref. [Low+25], where SOSSA already provided the state of the art for ground state energy estimation in chemical systems, we expect this framework to be generally useful and applicable to other quantum systems.

3.1 Quantum simulation using block-encodings and linear combination of unitaries

In this section we discuss the LCU results in Table 3.1 for energy estimation, phase estimation, and time evolution. To this end, we introduce the notion of a block-encoding of an operator $O \in \mathbb{C}^{M\times M}$. This block-encoding is a unitary $\mathrm{BE}[O/\lambda] \in \mathbb{C}^{M'\times M'}$ acting on an enlarged space ($M' \geq M$) that has O/λ in the first block, for some normalization constant $\lambda \geq \|O\|$,

$$\mathrm{BE}[O/\lambda] := \begin{pmatrix} O/\lambda & \cdot \\ \cdot & \cdot \end{pmatrix} . \tag{3.12}$$

(We use $\|.\|$ for the spectral norm.) When acting on systems of qubits, the first block is specified by the all-zero state of some ancilla register. More explicitly, we can write $\langle 0|_{\mathrm{a}} \mathrm{BE}[O/\lambda] |0\rangle_{\mathrm{a}} = O/\lambda$, where 'a' denotes an ancillary or 'clock' register. The definition can be naturally extended to rectangular operators $O \in \mathbb{C}^{M\times N}$. In this case we need to invoke one projector for the M-dimensional space and one projector for the N-dimensional space. Again, associating these projectors with qubit states, we might write $\langle 0|_{\mathrm{a}} \mathrm{BE}[O/\lambda] |0\rangle_{\mathrm{a}'} = O/\lambda$, where a and a′ are distinct for $M \neq N$.

Block-encodings provide a natural access model for several quantum simulation algorithms. In this case we often assume access to $\mathrm{BE}[H/\lambda]$, the block-encoding of a Hamiltonian H that models certain quantum system. The query complexity of such algorithms is determined by the number of uses to this block-encoding and, to make this complexity optimal, we would like λ to be as small as possible (e.g., as close to $\|H\|$ as possible). This is because various quantum simulation tasks that assume access to $\mathrm{BE}[H/\lambda]$ have (optimal) cost depending linearly on λ, as shown by the following known results.

Theorem 3.1 (Energy estimation from block-encoding [KOS07])**.** *Let $H \in \mathbb{C}^{M\times M}$ be a Hermitian operator, $\lambda \geq \|H\|$ be a normalization factor, and U and unitary*

preparing state $|\psi\rangle = U|0\rangle$*. With* $O(\lambda/\epsilon)$ *calls to* U*,* $BE[H/\lambda]$*, and their inverses, we can measure the energy (expectation)* $\langle\psi| H |\psi\rangle$ *to additive precision* $\epsilon > 0$*.*

Theorem 3.2 (Phase estimation from block-encoding [Ber+18])**.** *Let* $H \in \mathbb{C}^{M\times M}$ *be a Hermitian operator and* $\lambda \geq \|H\|$ *be a normalization factor. With* $O(\lambda/\epsilon)$ *calls to* $BE[H/\lambda]$ *and its inverse, we can perform phase (eigenvalue) estimation on* H *to additive precision* $\epsilon > 0$*.*

Theorem 3.3 (Time evolution from block-encoding [LC19; LC17b])**.** *Let* $H \in \mathbb{C}^{M\times M}$ *be a Hermitian operator and* $\lambda \geq \|H\|$ *be a normalization factor. With* $O(\lambda t + \log\frac{1}{\epsilon})$ *calls to* $BE[H/\lambda]$ *and its inverse, we can implement the time evolution operator* e^{-iHt} *to additive precision* $\epsilon > 0$*.*

We briefly comment on these results. Theorem 3.1 results for performing amplitude estimation, a problem that also reduces to quantum phase estimation as shown in Ref. [KOS07]. The Hamiltonian H is not unitary but the amplitude estimation is done with the block-encoding, to estimate an expectation of $\mathrm{BE}[H/\lambda]$. For additive precision ϵ' this can be done with $O(1/\epsilon')$ uses of $\mathrm{BE}[H/\lambda]$ and the inverse. The result follows from choosing $\epsilon' = \epsilon/\lambda$. Theorem 3.2 results from the standard use of the quantum phase estimation algorithm [Kit95] but, instead of running phase estimation on the unitary $e^{-iH/\lambda}$, it is ran using the walk operator appearing in qubitization [LC19]. We describe this operator below in Lemma 3.4. The benefit of doing this is that the approach does not necessitate of another routine that approximates $e^{-iH/\lambda}$ and the encoding can be done exactly and with less overhead. Theorem 3.3 follows from approximating the action of the evolution e^{-iHt} with a finite series that uses $\mathrm{BE}[H/\lambda]$ and the inverse. The series can be implemented using quantum signal processing and the cost $O(\lambda t + \log\frac{1}{\epsilon})$ is essentially the largest degree appearing in the series.

Lemma 3.4 (Qubitization [LC19])**.** *Let* $H \in \mathbb{C}^{M\times M}$ *be a Hermitian operator and* $\lambda \geq \|H\|$ *be a normalization factor. Let the quantum walk operator* $W := R_{EF_a} \cdot B_E[H/\lambda]$*, where the reflection* $R_{EF_a} = 2|0\rangle\langle 0|_a - \mathbb{1}$ *and assume the block-encoding* $B_E[H/\lambda]$ *is self-inverse. If* $B_E[H/\lambda]$ *is not self-inverse, one can always construct a self-inverse version using one query to controlled-*$B_E[H/\lambda]$ *and its inverse, and two Hadamard gates. Then for any eigenstate* $|\psi_j\rangle$ *of* H *with eigenvalue* E_j*,* W *has eigenstates* $|\psi_{j\pm}\rangle$ *with eigenvalues* $e^{\pm i \arccos(E_j/\lambda)}$*, and* $|\psi_j\rangle|0\rangle_a = \frac{1}{\sqrt{2}}(|\psi_{j+}\rangle + |\psi_{j-}\rangle)$*.*

We also comment on the optimality of these approaches. When given access to the block-encoding, a general method for phase estimation is known to necessitate $\Omega(\lambda/\epsilon)$ uses of the block-encoding and the inverse. Reference [MW23] proves a similar lower bound. The lower bound naturally extends to Theorem 3.1 or otherwise we would be able to perform faster quantum phase estimation via energy estimation, contradicting the previous result. For time evolution, the lower bound $\tilde{\Omega}(t\lambda + \log \frac{1}{\epsilon})$ is given in Ref. [Ber+14]. Since these results are tight, we present them using $\Theta(.)$ notation in Table 3.1. We note, however, that these refer to the query complexities. For specific instances where more structure is known and can be used, the upper bounds might be improved.

Block-encodings from linear combination of unitaries

We discussed quantum simulation when having access to $\mathrm{BE}[H/\lambda]$ but in applications the block-encoding must be constructed from some presentation of the Hamiltonian. A standard approach is based on the linear combination of unitaries (LCU) method [CW12; Ber+15]. Suppose the Hamiltonian is presented as

$$H = \sum_{j=0}^{R-1} g_j \sigma_j, \tag{3.13}$$

where the σ_j's are unitaries, for example Pauli strings, and $g_j \in \mathbb{C}$ are coefficients. When applying LCU, λ is equal to the ℓ_1-norm of the linear combination.

Definition 3.1 (ℓ_1-norm in LCU)**.** *Define λ_{LCU} to be the ℓ_1-norm*

$$\lambda_{LCU} = \sum_{j=0}^{R-1} |g_j| \ . \tag{3.14}$$

Lemma 3.5. (Compilation of LCU [Chi+18]) *Let H be as in Equation* (3.13) *and $\lambda_{LCU} \geq \|H\|$ be as in Equation* (3.14)*. Then, it is possible to construct a block-encoding of H, $BE[H/\lambda_{LCU}]$, using $O(RC_\sigma)$ quantum gates, where C_σ is the gate complexity of the σ_j's. The construction generalizes to arbitrary operators σ_j, which are not necessarily unitary, as long as $\|\sigma\| \leq 1$ and quantum circuits for the $BE[\sigma_j]$'s are given.*

Proof. The proof can be found in Ref.[Chi+18] and here we present a version that uses the notation of this work for completeness. First, we attach a 'clock register' of dimension R. Define the conditional unitary operator

$$\mathrm{SELECT} = \sum_{j=0}^{R-1} |j_\mathrm{a}\rangle\langle j_\mathrm{a}| \otimes \sigma_j \ , \tag{3.15}$$

where the $|j\rangle_{\mathrm{a}} \in \mathbb{C}^R$ are basis states of $\log_2(R)$ qubits. Also, define the state preparation unitary PREPARE that performs

$$\mathrm{PREPARE}\,|0\rangle_{\mathrm{a}} \mapsto |\alpha\rangle_{\mathrm{a}} := \frac{1}{\sqrt{\lambda_{\mathrm{LCU}}}} \sum_{j=0}^{R-1} \sqrt{g_j}\,|j\rangle_{\mathrm{a}} \ . \tag{3.16}$$

Then, $\mathrm{PREPARE}^{-1} \cdot \mathrm{SELECT} \cdot \mathrm{PREPARE}$ gives the block-encoding:

$$\langle 0|_{\mathrm{a}}\,\mathrm{PREPARE}^{-1}\cdot\mathrm{SELECT}\cdot\mathrm{PREPARE}\,|0\rangle_{\mathrm{a}} = \langle a|_{\mathrm{a}}\,\mathrm{SELECT}\,|\alpha\rangle_{\mathrm{a}} = \frac{1}{\lambda_{\mathrm{LCU}}} \sum_{j=0}^{R-1} g_j \sigma_j = \langle 0|_{\mathrm{a}}\,\mathrm{BE}[H/\lambda_{\mathrm{LCU}}]\,|0\rangle \tag{3.17}$$

Note that $|0\rangle_{\mathrm{a}}$ specifies the first block of the matrix. Implementing PREPARE requires accessing the coefficients g_j and needs $O(R)$ quantum gates in the worst case. Implementing SELECT can be done with $O(RC_\sigma)$ gates, and this step dominates the cost. □

Often, we will be interested in instances where the σ_j refer to Pauli strings acting on N qubits involving a constant number of local Pauli operators, in which case $C_\sigma = O(1)$ is constant, or where the σ_j refer to a product of a constant number of fermionic operators acting on N sites, in which case, $C_\sigma = O(\log N)$ for an optimal fermion-to-qubit ternary-tree mapping [Jia+20b]. Alternatively, more specialized constructions [Bab+18a] for SELECT in the Jordan-Wigner representation have gate complexity $O(N)$, and when $R = \Omega(N)$, the cost of block-encoding is dominated by PREPARE with cost $O(R\log(N))$, i.e. $C_\sigma = O(\log(N))$.

3.2 Quantum simulation by spectral amplification

In this section we introduce the basic idea of spectral amplification (SA) and provide the results that give Termwise SA and SOSSA in Table 3.1. The first version of SA was put forward in Ref. [SB13] under the name spectral gap amplification, in which the goal was to amplify the spectral gap for faster adiabatic quantum computing of frustration-free Hamiltonians rather than amplifying the whole low-energy spectrum for arbitrary Hamiltonians as we consider here. More recently, Ref. [ZS24] used SA to speedup time evolution for the low-energy subspace, in the block-encoding framework, which gives the third row in Table 3.1. We refer to the approach as SA because it amplifies all the small eigenvalues of eigenvectors in the low part of the spectrum of a positive semidefinite Hamiltonian.

We give the basic results of SA using block-encodings since this is a natural framework for this approach. Later, we will discuss how to construct these block-

encodings. In SA we consider first a Hamiltonian of the form

$$H' = \sum_{j=0}^{R-1} h_j \in \mathbb{C}^{M\times M} \ , \tag{3.18}$$

where $h_j \succeq 0$ are positive semidefinite Hermitian operators satisfying $\|h_j\| \leq g_j$. We consider a factorization of the h_j's so that $h_j = A_j^\dagger A_j$ for some operators A_j. In the following we assume $A_j \in \mathbb{C}^{M\times M}$ to simplify the exposition. Next we define the corresponding 'spectral amplified' operator H_{SA} by

$$H_{\mathrm{SA}} := \sum_{j=0}^{R-1} |j\rangle_{\mathrm{a}} \otimes A_j \in \mathbb{C}^{MR\times M} \ . \tag{3.19}$$

Note that we have enlarged the space and attached an R-dimensional clock register 'a'. Also note that $|j\rangle_{\mathrm{a}}$ can be replaced by any set of R mutually orthogonal quantum states and does not necessarily have to be a computational basis state. As a block matrix, H_{SA} is given by

$$H_{\mathrm{SA}} = \begin{pmatrix} A_0 \\ \vdots \\ A_{R-1} \end{pmatrix} . \tag{3.20}$$

We will see that the main property that makes SA useful is

$$H' = \sum_{j=0}^{R-1} A_j^\dagger A_j = H_{\mathrm{SA}}^\dagger H_{\mathrm{SA}} \ . \tag{3.21}$$

That is, H_{SA} acts as a square root of H' and now we can use it to produce other operators where the eigenvalues are changed. SA will then use access to H_{SA} and, to this end, we will assume access to unitaries $\mathrm{BE}[A_j/a_j]$, $a_j \geq \|A_j\|$, which are the block-encodings of the individual A_j/a_j's. To this end, it is useful to introduce the following block-encoding normalization factor.

Definition 3.2. *Define λ to be*

$$\lambda := \sum_{j=0}^{R-1} |a_j|^2 \ . \tag{3.22}$$

We can efficiently implement an appropriate block-encoding of H_{SA} as follows.

Lemma 3.6 (Block-encoding of H_{SA})**.** *Let $H' \succeq 0$ be a Hamiltonian of the form given in Equation* (3.18)*. Assume access to the individual block-encodings $BE[A_j/a_j]$*

and their inverses. Then, we can implement a block-encoding $BE[H_{SA}/\sqrt{\lambda}]$ with $O(R)$ calls to the (controlled) $BE[A_j/a_j]$'s and additional $O(R)$ arbitrary two-qubit gates.

Proof. Define the unitary

$$\text{SELECT} := \sum_{j=0}^{R-1} |j_{\text{b}}\rangle\langle j_{\text{b}}| \otimes \text{BE}\left[\frac{A_j}{a_j}\right] , \quad \langle 0|_{\text{a}} \text{BE}\left[\frac{A_j}{a_j}\right] |0\rangle_{\text{a}} = \frac{A_j}{a_j} , \tag{3.23}$$

and also the unitary PREPARE that prepares the state of the clock register

$$\text{PREPARE}\,|0\rangle_{\text{b}} \mapsto |\alpha\rangle_{\text{b}} := \frac{1}{\sqrt{\lambda}} \sum_{j=0}^{R-1} a_j\,|j\rangle_{\text{b}} \ . \tag{3.24}$$

Then,

$$\begin{aligned}
&\langle 0|_{\text{a}}\,\text{SELECT}\cdot\text{PREPARE}\,|0\rangle_{\text{b}}\,|0\rangle_{\text{a}} \\
&\quad = \langle 0|_{\text{a}}\,\text{SELECT}\,|\alpha\rangle_{\text{b}}\,|0\rangle_{\text{a}} = \sum_{j=0}^{R-1} \frac{a_j}{\sqrt{\lambda}}\,|j\rangle_{\text{b}} \otimes \langle 0|_{\text{a}}\,\text{BE}\left[\frac{A_j}{a_j}\right]|0\rangle_{\text{a}} \qquad (3.25)\\
&\quad = \sum_{j=0}^{R-1} |j\rangle_{\text{b}} \otimes \frac{A_j}{\sqrt{\lambda}} = \langle 0|_{\text{a}}\,\text{BE}\left[\frac{H_{\text{SA}}}{\sqrt{\lambda}}\right] \underbrace{|0\rangle_{\text{b}}\,|0\rangle_{\text{a}}}_{|0\rangle_{\text{a}'}} . \qquad (3.26)
\end{aligned}$$

This shows that SELECT·PREPARE gives the desired block encoding. For SELECT we used the individual controlled $\text{BE}[A_j/a_j]$ once and for PREPARE we used $O(R)$ arbitrary two-qubit gates. □

Equipped with access to $\text{BE}[H_{\text{SA}}/\sqrt{\lambda}]$, we will see in Section 3.2 and Section 3.2 that various simulation tasks can be performed with a cost depending on $\sqrt{\Delta\lambda}$, if the state is supported on the subspace of energy at most $\Delta > 0$ of H. In order to construct the walk operator Lemma 3.4 underlying these tasks, it suffices to consider a block-encoding of the Hamiltonian

$$\mathbf{H}_{\text{SA}} := \begin{pmatrix} \mathbf{0} & H_{\text{SA}}^{\dagger} \\ H_{\text{SA}} & \mathbf{0} \end{pmatrix}, \tag{3.27}$$

which has eigenvalues that coincide with the square roots of those of H. (This $\mathbf{H}_{\text{SA}}$ was used in Ref. [ZS24] for time evolution in the low-energy subspace of H.) Since $\sqrt{x} \gg x$ as $x \to 0$ this readily provides the desired amplification. However, for obtaining improved query and gate complexities, we can instead use $\text{BE}[H_{\text{SA}}/\sqrt{\lambda}]$

to block-encode a shifted version of H'. More precisely, in Lemma 3.7 we construct the following, self-inverse block-encoding:

$$\mathrm{BE}\left[\frac{H'}{\frac{1}{2}\lambda} - \mathbb{1}\right]. \tag{3.28}$$

By construction, any low-energy state of $H' \succeq 0$ with energy $0 \leq E \ll \lambda$ will correspond to an eigenvalue close to -1 of the block-encoded operator $H'/(\lambda/2) - \mathbb{1}$. Since the quantum walk operator has now eigenphases $\arccos(E_j/(\lambda/2) - 1)$, similar to the square root function, the non-linearity of arccos is what allows us to achieve SA.

Lemma 3.7 (Block-encoding of $\frac{1}{2}H_{\mathrm{SA}}^{\dagger}H_{\mathrm{SA}} - \lambda\mathbb{1}$)**.** *Let $H' \succeq 0$ be a Hamiltonian and consider the factorization $H' = H_{SA}^{\dagger}H_{SA}$, where $H_{SA} \in \mathbb{C}^{M\times N}$. Let $BE[H_{SA}/\sqrt{\lambda}] \in \mathbb{C}^{D\times D}$ be a block-encoding of H_{SA}. Then we may block-encode either $BE[\frac{H_{SA}^{\dagger}H_{SA}}{\frac{1}{2}\lambda} - \mathbb{1}]$ using one query to $BE[H_{SA}/\sqrt{\lambda}]$ and its inverse, $O(Q\log(D/M))$ arbitrary two-qubit gates, and one ancillary qubit, or its controlled version using two ancillary qubits.*

Proof. By the definition of block-encodings, given the all-zero state $|0\rangle_{\mathrm{a}'} \in \mathbb{C}^{D/N}$, $|0\rangle_{\mathrm{a}} \in \mathbb{C}^{D/M}$ and any state $|\psi\rangle_{\mathrm{S}} \in \mathbb{C}^{N}$,

$$\mathrm{BE}\left[\frac{H_{\mathrm{SA}}}{\sqrt{\lambda}}\right]|0\rangle_{\mathrm{a}'}|\psi\rangle_{\mathrm{S}} = |0\rangle_{\mathrm{a}}\frac{H_{\mathrm{SA}}}{\sqrt{\lambda}}|\psi\rangle_{\mathrm{S}} + \cdots|0^{\perp}\rangle, \tag{3.29}$$

where $|(\langle 0|_{\mathrm{a}}\otimes\mathbb{1}_{\mathrm{S}})|0^{\perp}\rangle| = 0$. Using quantum singular value transformations [Gil+19] with the polynomial $2x^2 - 1$, we can block-encode

$$\mathrm{BE}\left[\frac{H_{\mathrm{SA}}^{\dagger}H_{\mathrm{SA}}}{\frac{1}{2}\lambda} - \mathbb{1}\right] = \mathrm{BE}\left[\frac{H_{\mathrm{SA}}}{\sqrt{\lambda}}\right]^{\dagger}(\mathrm{Ref}_{\mathrm{a}}\otimes\mathbb{1})\mathrm{BE}\left[\frac{H_{\mathrm{SA}}}{\sqrt{\lambda}}\right], \tag{3.30}$$

where the reflection $\mathrm{Ref}_{\mathrm{a}} := 2|0\rangle\langle 0|_{\mathrm{a}} - \mathbb{1}_{\mathrm{a}}$. This reflection can be implemented using a multi-controlled-CNOT gate which costs one ancillary qubit and $O(\log(D/M))$ two-qubit gates [He+17]. The controlled reflection can be implemented using a multi-controlled-CNOT gate with one additional control, and so uses two additional qubits in total over that of $\mathrm{BE}[H_{\mathrm{SA}}/\sqrt{\lambda}]$. □

Expectation estimation by spectral amplification

In this section, we present new quantum algorithms summarized in Table 3.2 that exploit SA to improve expectation estimation. Given a block-encoding $\mathrm{Be}[H/\lambda]$ of

an arbitrary Hamiltonian H and a unitary preparing the state $|\psi\rangle$, the expectation estimation problem is to estimate $E = \langle\psi| H |\psi\rangle \in [-\lambda, \lambda]$ to additive error $\epsilon \geq 0$ and confidence at least $1 - q \leq 1$. To simplify the presentation, we assume in the following that $E \leq 0$, though we emphasize that our results are symmetric about $E = 0$, meaning that they depend on $-|E|$. In contrast to the traditional approaches that give $O(\lambda/\epsilon)$ query complexity, we show that SA leads to explicit scaling with the improved factor $\sqrt{\lambda(\lambda + E)} \leq \lambda$, which offer greatest advantage in the limit $E/\lambda \to -1$, the 'low-energy sector', and our generalized routines naturally recover previously known results without SA, which concern a different limit in which $|E/\lambda| \ll 1$. In the special case of amplitude estimation [Bra+02] corresponding to the case where $H/\lambda = \Pi$ is a projection, prior work on 'amplified amplitude estimation' [Sim+24] already gives the improved query complexity in this setting. Nevertheless, our algorithms improve on this previous work in these key areas (Row 2 of Table 3.2):

- Given an *a priori* known upper bound $\Delta \geq \lambda + E$, the query complexity of our non-adaptive algorithm in Theorem 3.9 readily scales as $O(\sqrt{\Delta\lambda}/\epsilon)$, improving prior work by a factor $\frac{\Delta}{\Delta-(\lambda+E)}$, which means that our results do not need the known upper bound Δ to be at least a constant multiplicative factor worse than the *a priori* unknown actual value of $\lambda + E$.

- Even without any prior knowledge of E, such as through the upper bound Δ, the query complexity of our adaptive algorithm in Theorem 3.12 scales like $O(\sqrt{(\max\{\epsilon, \lambda + E\}\lambda}/\epsilon)$. The dependence on $\lambda + E$ rather than Δ is a significant improvement as $\lambda + E$ could be arbitrarily smaller than Δ. Moreover, this result also naturally recovers the 'super-Heisenberg' scaling of $O(1/\sqrt{\epsilon})$ when $\epsilon = \Theta(\lambda+E)$ *without prior knowledge* on E, which in previous work [Sim+24] required the specific condition that $\epsilon = \Theta(\lambda + E) = \Theta(\Delta)$.

- Our results are general and apply to arbitrary H that can be block-encoded, instead of only reflections or sums of reflections as in [Sim+24]. Moreover, when we instantiate with $H = 2H_{\mathrm{SA}}^{\dagger} H_{\mathrm{SA}} - \lambda$ in Lemma 3.7, this generalizes the case where H is a reflection and H_{SA} is a projector to arbitrary block-encoded rectangular operators H_{SA}. The low-energy sector is equivalent to assuming small $\frac{1}{\lambda} \langle\psi| H_{\mathrm{SA}}^{\dagger} H_{\mathrm{SA}} |\psi\rangle$, and we give Corollary 3.13 as an example of how our general results expressed in the block-encoding framework easily recover these special cases.

Year	Reference	Query complexity $O(\cdot)$	Needs Δ	Comments
2002	[Bra+02]	$\frac{\lambda}{\epsilon}\log\left(\frac{1}{q}\right)$	No	Equivalent to $H_{\mathrm{SA}} \propto \Pi$,
2024	[Sim+24]	$\frac{\sqrt{\Delta\lambda}}{\epsilon}\log\left(\frac{1}{q}\right)\left(\frac{\Delta}{\Delta-(\lambda+E)}\right)$	Yes	where Π is a projector.
2025	Theorem 3.9	$\frac{\sqrt{\Delta\lambda}}{\epsilon}\log\left(\frac{1}{q}\right)$	Yes	Non-adaptive algorithm
2025	Theorem 3.12	$\frac{\sqrt{\max\{\epsilon,\lambda-\lvert E\rvert\}\lambda}}{\epsilon}\log\left(\frac{1}{q}\right)$	No	Adaptive algorithm

Table 3.2: Cost of estimating an expectation $E = \langle\psi| H |\psi\rangle \in [-\lambda, \lambda]$, to additive error ϵ and failure probability q, given query access to the block-encoding of $\textsc{Be}[H/\lambda]$ and a state preparation unitary for the state $|\psi\rangle$. In some cases, a known upper bound $\Delta \geq \lambda - |E|$ is needed *a priori*.

In the previous section we considered $H' = H_{\mathrm{SA}}^{\dagger} H_{\mathrm{SA}}$ in Equation (3.18) and also $\|H'\| \leq \lambda$. Using Lemma 3.7, we can shift this to $H = 2H' - \lambda$ so that the relevant eigenvalues of H are in $[-\lambda, \lambda]$. To ease the exposition and the comparison with other methods, in the following we will assume access to $\textsc{Be}[H/\lambda]$ where H is arbitrary like Equation (3.13) with eigenvalues are in $[-\lambda, \lambda]$, and the low-energy sector, is that of energies at or near $-\lambda$. At a high level, our non-adaptive algorithm Theorem 3.9 works by performing quantum phase estimation on the walk operator Lemma 3.4. In expectation estimation, this is a walk-operator essentially on the block-encoding of a trivial 1×1 operator $E = \langle\psi| H |\psi\rangle$. Due to the arccos non-linearity, any error ϵ' in the estimate of phase becomes a smaller square-root error ϵ in the estimate of E when it is near $-\lambda$. As the upper bound Δ is known beforehand, we know how to choose ϵ' in phase estimation to achieve the desired $\sqrt{E\Delta}$, scaling.

Achieving the optimal scaling of our Theorem 3.12 without prior knowledge of Δ is significantly more challenging. Without the upper bound Δ, we cannot make a naive choice of ϵ' beforehand. For instance, the conservative choice of $\epsilon' = \Theta(\epsilon/\lambda)$ is guaranteed to achieve the desired accuracy, but this is basically the worst-case scaling. The next idea is to perform a binary search for E using multiple $i = 1, \cdots, i_{\max}$ iterations of phase estimation to accuracy ϵ'_i that decreases geometrically and confidence $1 - q_i$. By learning an estimate $\hat{E}$ of E on the fly and computing ϵ based on $\arccos(\hat{E})$, we can terminate the algorithm when the desired ϵ is achieved. Early termination leads to almost the desired $\tilde{O}(\sqrt{(E+\lambda)\lambda}/\epsilon)$ scaling, but this is still be suboptimal by a logarithmic factor as the failure probability q_i accumulates over multiple steps, and without knowing E beforehand, one has to make a worst-case choice of q_i that assumes the number of iterations $i_{\max} = \Theta(\log(\lambda/\epsilon))$, which

leads to $O(\log \frac{\log(\lambda/\epsilon)}{q})$ scaling. Our solution is to perform the binary search in two steps. First, given ϵ, we perform phase estimation to error $\epsilon' = \Theta(\sqrt{\epsilon/\lambda})$. We prove that this is guaranteed to give us either an estimate $\lambda + \hat{E}$ of $\lambda + E$ to constant multiplicative error, or that $0 \leq \lambda + E \leq \epsilon$, and so we terminate the algorithm and return $-\lambda$ as the estimate of E that is correct to additive error ϵ. Second, using this estimate $\hat{E}$, we now know how many iterations of phase estimation $i_{\max}$ are required, up to a constant additive factor. This information is sufficient for us to make a judicious choice of q_i that by a union bound, achieves the final desired confidence with the desired $O(\log(\frac{1}{q}))$ scaling.

To prove our first result, we require the following known result on phase estimation.

Lemma 3.8. (Phase estimation with confidence intervals [KOS07; Ber+24b; Gil+19]) *Let U be a unitary operator with eigenstates $|\psi_j\rangle$ and corresponding eigenvalues $e^{i\theta_j}$. Let $|\psi\rangle = \sum_j \sqrt{p_j}\,|\psi_j\rangle$ be an arbitrary superposition. Then with $O(\frac{1}{\epsilon}\log\frac{1}{q})$ calls to controlled-U, its inverse, and one copy of $|\psi\rangle$, with probability p_j we estimate $\hat{\theta}_j$ such that $|(\hat{\theta}_j - \theta_j) \mod 2\pi| \leq \epsilon$ with confidence $1 - q$.*

Next, we use the high-confidence quantum phase estimation to obtain our first version of improved expectation or energy estimation.

Theorem 3.9 (Energy estimation by spectral amplification)**.** *Let $H \in \mathbb{C}^{N\times N}$ be a Hamiltonian and assume access to the block-encoding $BE[H/\lambda] \in \mathbb{C}^{D\times D}$. Let $|\psi\rangle \in \mathbb{C}^N$ be prepared by the state preparation unitary P, such that $P\,|0\rangle_S = |\psi\rangle$ and assume it satisfies $\langle\psi|\,H\,|\psi\rangle \leq -\lambda + \Delta$, for some known $\Delta > 0$. Then $E = \langle\psi|\,H\,|\psi\rangle$ can be estimated to additive error ϵ and confidence $1 - q$ using $Q = O(\frac{\sqrt{\lambda\Delta}}{\epsilon}\log\frac{1}{q})$ queries to the block-encoding, state preparation unitary, and their inverses, two ancillary qubits, and $O(Q\log(D/N))$ arbitrary two-qubit gates.*

Proof. Observe that the following quantum circuit is a block-encoding of the 1×1 matrix (E/λ):

$$U := \mathrm{BE}[E/\lambda] = (\mathbb{1}_a \otimes P^\dagger)\mathrm{BE}\left[\frac{H}{\lambda}\right](\mathbb{1}_a \otimes P). \tag{3.31}$$

By Qubitization in Lemma 3.4, the quantum walk operator $W = \mathrm{REF}_{RS}\cdot \mathrm{BE}[E/\lambda]$ has eigenvectors $|\pm\rangle := \frac{1}{\sqrt{2}}(|0\rangle_{S,a} \pm |0^\perp\rangle)$ with eigenphases $\pm\arccos(E/\lambda)$.

We now apply high-confidence quantum phase estimation in Lemma 3.8 to this quantum walk operator with input state $|0\rangle_{S,a}$. For additive error ϵ_{PEA} and confidence

at least $1-q$, this requires $O(\frac{1}{\epsilon_{\text{PEA}}}\log\frac{1}{q})$ queries to W. Note that we will obtain a phase with the + or − sign with probability $1/2$ as the input state $|0\rangle_{\text{S,a}}$ is a uniform superposition of $|\pm\rangle$. However, this sign does not impact our estimate and can be ignored. An error ϵ_{PEA} in the phase translates to a smaller error ϵ in the expectation due to the arccos function. Indeed, using the derivative to propagate errors, it is possible to show

$$\epsilon \leq \sqrt{2\Delta\lambda}\ \epsilon_{\text{PEA}}, \tag{3.32}$$

and hence choosing $\epsilon_{\text{PEA}} = \epsilon/(\sqrt{2\Delta\lambda})$ suffices. This gives the desired complexity. □

We now show that the same scaling is achieved even without prior knowledge on the upper bound for $\langle\psi|\, H|\psi\rangle$. This requires a decision version of phase estimation called gapped phase estimation. Given a unitary U and and an eigenstate $|\psi_j\rangle$ with eigenvalue $e^{i\theta_j}$, the task is to decide with high probability whether θ_j is either in the interval $\mathcal{I}_1$ or $\mathcal{I}_2$, where these intervals are disjoint and separated by some angular gap 2ϵ. Searching for θ_j like in Lemma 3.8 can be reduced to a binary search using gapped phase estimation with optimal query complexity. The following version of gapped phase estimation in particular uses no ancillary qubits beyond the one needed for controlled-U.

Lemma 3.10 (Gapped phase estimation; Appendix D of [LS24])**.** *Let U be a unitary operator and $|\psi\rangle$ be an eigenstate satisfying $U|\psi\rangle = e^{i\theta}|\psi\rangle$. For any eigenstate $U|\psi\rangle = e^{i\theta}|\psi\rangle$ and any θ_0 and ϵ satisfying $0 < \epsilon \leq \theta_0 \leq \epsilon + \theta_0 \leq \frac{\pi}{2}$, $q > 0$, there is a unitary $G_{PE_{\epsilon,\theta_0,q}}$ that prepares the state*

$$G_{PE_{\epsilon,\theta_0,q}}|1\rangle|\psi\rangle = (\alpha(\theta)|0\rangle + \beta(\theta)|1\rangle)|\psi\rangle, \quad |\alpha(\theta)|^2 + |\beta(\theta)|^2 = 1, \tag{3.33}$$

where

$$\forall|\theta| \in [0, \theta_0 - \epsilon], \qquad |\alpha(\theta)|^2 \leq 2q, \qquad |\beta(\theta) - 1| \leq q, \tag{3.34}$$

$$\forall|\theta - \pi| \in [0, \theta_0 - \epsilon], \qquad |\alpha(\theta)|^2 \leq 2q, \qquad |\beta(\theta) + 1| \leq q, \tag{3.35}$$

$$\forall|\theta| \in [\theta_0 + \epsilon, \pi - \theta_0 - \epsilon], \qquad |\beta(\theta)|^2 \leq 2q, \qquad |\alpha(\theta) - 1| \leq q, \tag{3.36}$$

using $Q = O(\frac{1}{\epsilon}\log\frac{1}{q})$ queries to controlled-U and its inverse, $O(Q)$ arbitrary two-qubit gates, and no ancillary qubits.

This version of gapped phase estimation produces measurement probabilities that are symmetric about $\theta = \pi/2$. In situations where it is necessary to distinguish between

the θ or $\pi - \theta$ branches, one may always perform another round of phase estimation on, say, $e^{i\pi/2}U$, using Equation (3.34). However in the application to quantum walk operations, a subtlety is that every eigenstate $|\psi_j\rangle$ of the block-encoded operator maps to a equal superposition of two eigenstate $|\psi_{j,\pm}\rangle$ of the walk operator with eigenvalues $e^{\pm i \arccos\theta_j}$. Hence distinguishing between the $\pm$ branches requires the following approach of controlled gapped phase estimation.

Corollary 3.11 (Controlled gapped phase estimation)**.** *Let U be a unitary operator. For any eigenstate $U|\psi\rangle = e^{i\theta}|\psi\rangle$ and any θ_0, ϵ satisfying $0 < \epsilon \leq \theta_0 \leq \epsilon + \theta_0 \leq \frac{\pi}{2}$, $q > 0$, there is a unitary $CGPE_{\epsilon,\theta_0,q}$ that prepares the state*

$$\begin{aligned} &CGPE_{\epsilon,\theta_0,q}\,|+\rangle\,|1\rangle\,|\psi\rangle = (\alpha_0(\theta)\,|01\rangle + \alpha_1(\theta)\,|11\rangle + \gamma(\theta)\,|-0\rangle)\,|\psi\rangle\,, \\ &|\alpha_0(\theta)|^2 + |\alpha_1(\theta)|^2 + |\gamma(\theta)|^2 = 1, \end{aligned} \tag{3.37}$$

where

$$\forall|\theta| \in [0, \theta_0 - \epsilon], \qquad 1 - |\alpha_0(\theta)|^2 \leq q, \tag{3.38}$$

$$\forall|\theta - \pi| \in [0, \theta_0 - \epsilon], \qquad 1 - |\alpha_1(\theta)|^2 \leq q. \tag{3.39}$$

using $Q = O(\frac{1}{\epsilon}\log\frac{1}{q})$ queries to controlled-U and its inverse, $O(Q)$ arbitrary two-qubit gates, and no ancillary qubits.

Proof. Let $\mathrm{CGPE}_{\epsilon,\theta_0,q'} := |0\rangle\langle 0|\otimes\mathbb{1} + |1\rangle\langle 1|\otimes\mathrm{GPE}_{\epsilon,\theta_0,q'}$ where $\mathrm{GPE}_{\epsilon,\theta_0,q}$ is from Lemma 3.10. Then

$$\mathrm{CGPE}_{\epsilon,\theta_0,q'}\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)\,|1\rangle\,|\psi\rangle = \left[\frac{|0\rangle + \beta(\theta)\,|1\rangle}{\sqrt{2}}\right]|1\rangle\,|\psi\rangle + \frac{1}{\sqrt{2}}\alpha(\theta)\,|1\rangle\,|0\rangle\,|\psi\rangle\,.$$

Now apply the Hadamard gate to obtain the state

$$|\chi(\theta)\rangle = \left[\frac{1+\beta(\theta)}{2}\,|0\rangle + \frac{1-\beta(\theta)}{2}\,|1\rangle\right]|1\rangle\,|\psi\rangle + \frac{1}{\sqrt{2}}\alpha(\theta)\,|-\rangle\,|0\rangle\,|\psi\rangle\,. \tag{3.40}$$

If $|\theta| \in [0, \theta_0 - \epsilon]$, then from Equation (3.34), the probability that we do not measure $|01\rangle$ is at most

$$1 - |\alpha_0(\theta)|^2 = \frac{|1-\beta(\theta)|^2}{4} + \frac{|\alpha(\theta)|^2}{2} \leq \frac{q'^2}{4} + q' \leq \frac{5}{4}q' \leq q, \tag{3.41}$$

where we choose $q' = \frac{4}{5}q$. Similarly, using Equation (3.36) if $|\theta - \pi| \in [0, \theta_0 - \epsilon]$, the probability that we do not measure $|11\rangle$ is also at most q. □

We are now ready to prove the general result of this section, giving the first row of Table 3.1.

Theorem 3.12 (Energy estimation by spectral amplification with no prior)**.** *Let $H \in \mathbb{C}^{N\times N}$ be a Hamiltonian and assume access to the block-encoding $BE[H/\lambda] \in \mathbb{C}^{D\times D}$. Let $|\psi\rangle \in \mathbb{C}^N$ be prepared by the state preparation unitary P, such that $P|0\rangle = |\psi\rangle$. Then $E = \langle\psi| H |\psi\rangle$ can be estimated to any additive error ϵ and confidence $1-q$ using $Q = O(\frac{\sqrt{\max\{\epsilon,\lambda-|E|\}\lambda}}{\epsilon}\log\frac{1}{q})$ queries to the block-encoding, state preparation unitary, and their inverses, two ancillary qubits and $O(Q\log(D/N))$ arbitrary two-qubit gates.*

Proof. In Theorem 3.9, we constructed a unitary walk operator W. Our discussion is made clearer by instead considering the walk operator $-W$, which has eigenvectors $|\pm\rangle$ and eigenphases $\theta_\pm := \pi \mp \arccos(\alpha)$, where $\alpha := \frac{E}{\lambda} \in [-1,1]$. We find it convenient to define $o := \alpha + 1 = E/\lambda + 1 \in [0,2]$. We can choose the principle range of $\theta_\pm$ to be symmetric about 0. Hence $\theta_\pm = \pm\theta \in \pm[0,\pi]$, that is, $\theta := |\theta_+| = |\theta_-|$, and $\theta - \pi/2$ is symmetric around $\alpha = 0$. We search for θ using gapped phase estimation Lemma 3.10 in two steps. In the first step, we assume that $\alpha \in [-1,0] \implies \theta \in [0,\pi/2]$ and estimate θ to additive error ϵ'. As the bounds on the probabilities $|\alpha(\theta)|^2, |\beta(\theta)|^2$ of gapped phase estimation are also symmetric about $\pi/2$, gapped phase estimation on $-W$ cannot distinguish between cases θ or $\pi-\theta$. However, the sign of $\beta(\theta) \approx -\beta(\pi-\theta)$ is sensitive to these cases. In the second step, we therefore distinguish between cases using controlled-gapped phase estimation Corollary 3.11.

Let the iterations of the search be indexed by $i = 0, \cdots, i_{\max}-1$, for some $i_{\max} > 0$. At each iteration, θ is known to be contained in an interval $\mathcal{I}_i := [\mathcal{I}_{i,\mathrm{l}}, \mathcal{I}_{i,\mathrm{r}}]$ with high probability p_i. Hence, $\theta \in \mathcal{I}_0 = [0,\pi/2]$ with probability $p_0 = 1$. At each iteration $i \geq 0$, we split $\mathcal{I}_i$ into thirds

$$\mathcal{I}_{i,\uparrow} := \left[\mathcal{I}_{i,\mathrm{l}} + \frac{1}{3}|\mathcal{I}_i|, \mathcal{I}_{i,\mathrm{r}}\right], \quad \mathcal{I}_{i,\downarrow} := \left[\mathcal{I}_{i,\mathrm{l}}, \mathcal{I}_{i,\mathrm{r}} - \frac{1}{3}|\mathcal{I}_i|\right], \tag{3.42}$$

and we will assign $\mathcal{I}_{i+1}$ to be either $\mathcal{I}_{i,\uparrow}$ or $\mathcal{I}_{i,\downarrow}$. Note that the width $|\mathcal{I}_i| = \frac{\pi}{2}r^i$, where $r = 2/3$.

We determine this assignment using gapped phase estimation $\textsc{Gpe}_{\varphi_i,\theta_i,q_i}$, with $\theta_i = \frac{\mathcal{I}_{i,\mathrm{l}}+\mathcal{I}_{i,\mathrm{r}}}{2}$ at the midpoint of $\mathcal{I}_i$ and $\varphi_i = \frac{1}{6}|\mathcal{I}_i|$. From Lemma 3.10, this prepares a state

$$\textsc{Gpe}_{\varphi_i,\theta_i,q_i}|0\rangle \frac{|+\rangle + |-\rangle}{2} = \frac{1}{\sqrt{2}}\sum_{x\in\{+,-\}}(\alpha(\pm\theta)|0\rangle + \beta(\pm\theta)|1\rangle)|x\rangle\,. \tag{3.43}$$

We measure the $\{|0\rangle, |1\rangle\}$ register to obtain outcome $|m\rangle$. Then we set

$$\mathcal{I}_{i+1} = \begin{cases} \mathcal{I}_{i,\downarrow}, & m = 0, \\ \mathcal{I}_{i,\uparrow}, & m = 1. \end{cases} \tag{3.44}$$

The estimate of θ in $\mathcal{I}_i$ converts to an estimate of

$$\alpha \in -\mathcal{O}s(\mathcal{I}_i) = [-\mathcal{O}s(\mathcal{I}_{i,\mathrm{l}}), -\mathcal{O}s(\mathcal{I}_{i,\mathrm{r}})] = [-\mathcal{O}s(\theta_i - |\mathcal{I}_i|/2), -\mathcal{O}s(\theta_i + |\mathcal{I}_i|/2)], \tag{3.45}$$

$$\epsilon_i = |\mathcal{O}s(\mathcal{I}_i)| = 2\sin(\theta_i)\sin(|\mathcal{I}_i|/2), \tag{3.46}$$

$$E \in \mathcal{H}_i = [\mathcal{H}_{i,\mathrm{l}}, \mathcal{H}_{i,\mathrm{r}}], \quad \mathcal{H}_i := -\lambda \mathcal{O}s\mathcal{I}_i, \tag{3.47}$$

$$\epsilon_i' := |\mathcal{H}_i| = \lambda\epsilon_i, \tag{3.48}$$

where ϵ_i and ϵ_i' is the additive error to which α and E is known respectively.

The probability that this assignment is incorrect, that is $\theta \notin \mathcal{I}_{i+1}$ is given by the maximum of the probabilities

$$\Pr[m = 1 | \theta \in \mathcal{I}_{i,\uparrow} \backslash \mathcal{I}_{i,\downarrow}] = \frac{|\alpha(\theta)|^2 + |\alpha(-\theta)|^2}{2} \leq q_i, \tag{3.49}$$

$$\Pr[m = 0 | \theta \in \mathcal{I}_{i,\downarrow} \backslash \mathcal{I}_{i,\uparrow}] = \frac{|\beta(\theta)|^2 + |\beta(-\theta)|^2}{2} \leq q_i. \tag{3.50}$$

We choose $q_i = \frac{6}{\pi^2}\frac{q}{(i_{\max}-i+1)^2}$. Then by a union bound, the failure probability of amplitude estimation after all $i_{\max}$ steps is at most

$$1 - p_{i_{\max}} \leq \sum_{i=1}^{i_{\max}} q_i \leq \sum_{l=1}^{\infty} \frac{6}{\pi^2}\frac{q}{i^2} = q. \tag{3.51}$$

The query complexity of all the GPE steps is then

$$Q = \mathcal{O}\left(\sum_{i=0}^{i_{\max}-1} \frac{1}{r^i} \log\frac{1}{q_i}\right) = \mathcal{O}\left(\sum_{i=0}^{i_{\max}-1} \left(\frac{3}{2}\right)^i \left((i_{\max} - i + 1) + \log\frac{1}{q}\right)\right) = \mathcal{O}\left(\left(\frac{3}{2}\right)^{i_{\max}} \log\frac{1}{q}\right). \tag{3.52}$$

When the search is complete, there are two cases of interest.

1. $\mathcal{I}_{i_{\max},\mathrm{l}} = 0$: This implies that $\alpha \in [-1, -\mathcal{O}s|\mathcal{I}_{i_{\max}}|]$, and that o is small, that is $o \leq \epsilon_{i_{\max}} = 2\sin^2(|\mathcal{I}_{i_{\max}}|/2) \leq \frac{\pi^2}{8}r^{2i}$. Note further that $o \leq 2\sin^2(|\mathcal{I}_{i_{\max}}|/2)$ if and only if $\mathcal{I}_{i_{\max},\mathrm{l}} = 0$.

2. $\mathcal{I}_{i_{\max},1} > 0$: This implies that $\alpha \geq -Os\mathcal{I}_{i_{\max},1}$ and we have an estimate of $o \geq 1 - Os\mathcal{I}_{i_{\max},1}$ that is bounded away from zero. Let i_f be the first iteration where $\mathcal{I}_{i_f,l} > 0$. Then $\mathcal{I}_{i_f,1} = \frac{1}{2}|\mathcal{I}_{i_f}|$ and

$$\theta \in \left[\frac{1}{2}|\mathcal{I}_{i_f}|, \frac{3}{2}|\mathcal{I}_{i_f}|\right] = \left[\frac{1}{2}, \frac{3}{2}\right]\frac{\pi}{2}r^{i_f}, \tag{3.53}$$

$$o \in \left[1 - Os\left(\frac{1}{2}|\mathcal{I}_{i_f}|\right), 1 - Os\left(\frac{3}{2}|\mathcal{I}_{i_f}|\right)\right], \tag{3.54}$$

$$\epsilon_{i_f} = 2\sin\left(|\mathcal{I}_{i_f}|\right)\sin\frac{|\mathcal{I}_{i_f}|}{2} \leq 4\sin^2\frac{|\mathcal{I}_{i_f}|}{2} = 4\left(1 - Os^2\frac{|\mathcal{I}_{i_f}|}{2}\right) \leq 8\left(1 - Os\frac{|\mathcal{I}_{i_f}|}{2}\right) \leq 8o. \tag{3.55}$$

As ϵ_i decreases monotonically with i, this implies that we obtain a estimate of o to constant multiplicative error. For any later iteration $i > i_f \geq 0$, observe that $|\mathcal{I}_{i_f}| \leq \frac{\pi}{2}\frac{2}{3}$. Hence $\sin\frac{|\mathcal{I}_{i_f}|r^{i-i_f}}{2} \leq \frac{\pi}{3}r^{i-i_f}\sin\frac{|\mathcal{I}_{i_f}|}{2}$, $\sin\frac{3|\mathcal{I}_{i_f}|}{2} \leq \frac{3}{2}\sin|\mathcal{I}_{i_f}|$, and

$$\begin{aligned}\epsilon_i &= 2\sin(\theta_i)\sin\frac{|\mathcal{I}_i|}{2} = 2\sin(\theta_i)\sin\frac{|\mathcal{I}_{i_f}|r^{i-i_f}}{2} \leq 2\sin\frac{3|\mathcal{I}_{i_f}|}{2}\sin\frac{|\mathcal{I}_{i_f}|r^{i-i_f}}{2} \\ &\leq \pi r^{i-i_f}\sin|\mathcal{I}_{i_f}|\sin\frac{|\mathcal{I}_{i_f}|}{2} = \pi r^{i-i_f}\epsilon_{i_f} \leq 8\pi r^{i-i_f}o.\end{aligned} \tag{3.56}$$

We now evaluate the final error $\epsilon'_{i_{\max}}$ of the estimate of E for some choice of $i_{\max}$.
We now evaluate the final error $\epsilon'_{i_{\max}}$ of the estimate of E for some choice of $i_{\max}$.

1. Let $i_{\max,1} := \lceil\frac{1}{2}\log_{1/r}(\frac{\pi^2\lambda}{8\epsilon'})\rceil$. Choose $i_{\max} = i_{\max,1}+d$, where $d = \lceil\log_{1/r}(16\pi)\rceil$, and $q = \Delta_1$. The query complexity is

$$Q_1 = O\left(\frac{\sqrt{\lambda}}{\sqrt{\epsilon'}}\log\frac{1}{\Delta_1}\right), \tag{3.57}$$

and we show that we either obtain an estimate of E to error at most ϵ', or obtain an estimate of $(\lambda - |E|) \in [0, \lambda]$ to at most a constant multiplicative error of $\frac{1}{2}$ as follows.

   a) Case $\mathcal{I}_{i_{\max,1},l} = 0$:

$$\epsilon'_{i_{\max,1}} \leq \lambda\frac{\pi^2}{8}r^{2i_{\max,1}} \leq \epsilon'. \tag{3.58}$$

   b) Case $\mathcal{I}_{i_{\max,1},l} > 0$:

$$\epsilon'_{i_{\max,1}+d} \leq \lambda 8\pi r^{d+i_{\max,1}-i_f}o = 8\pi r^{d+i_{\max,1}-i_f}h \leq 8\pi r^d(\lambda - |E|) \leq \frac{1}{2}(\lambda - |E|). \tag{3.59}$$

Hence

$$\lambda - |E| \leq \lambda - |\mathcal{H}_{i_{\max},\mathrm{r}}| \leq \lambda - |E| + \epsilon'_{i_{\max}} \leq \frac{3}{2}(\lambda - |E|) \leq \frac{3}{2}(\lambda - |\mathcal{H}_{i_{\max},\mathrm{r}}|). \tag{3.60}$$

2. Let $i_{\max,2} := \lceil \log_{1/r}(\frac{\sqrt{(\lambda - |\mathcal{H}_{i_{\max},\mathrm{r}}|)\lambda}}{c\epsilon'}) \rceil$, where $c = \frac{2}{\sqrt{3}\pi}$. Choose $i_{\max} = i_{\max,2}$, and $q = \Delta_2$. The query complexity is

$$Q_2 = O\left(\frac{\sqrt{(\lambda - |\mathcal{H}_{i_{\max},\mathrm{r}}|)\lambda}}{\epsilon'} \log \frac{1}{\Delta_2}\right), \tag{3.61}$$

and we obtain an estimate of E to error at most ϵ' as follows.

$$\begin{aligned}
\epsilon_{i_{\max}} &= 2\sin(\theta_i)\sin(|\mathcal{I}_i|/2) \leq \sin(\theta_i)\frac{\pi}{2}r^{i_{\max}} \leq \sin(\arccos(-(o + \epsilon_{i_{\max}} - 1)))\frac{\pi}{2}r^{i_{\max}} \\
&= \sqrt{(2 - o - \epsilon_{i_{\max}})(o + \epsilon_{i_{\max}})}\frac{\pi}{2}r^{i_{\max}} \leq \sqrt{2\,(o + \epsilon_{\max})}\frac{\pi}{2}r^{i_{\max}}. \qquad (3.62) \\
\epsilon'_{i_{\max}} &= \lambda\epsilon_{i_{\max}} \leq \lambda\sqrt{2\left(o + \epsilon_{i_{\max}}\right)}\frac{\pi}{2}r^{i_{\max}} \leq \sqrt{2\lambda}\sqrt{(\lambda - |E|) + \epsilon'_{i_{\max}}}\frac{\pi}{2}r^{i_{\max}} \\
&\leq c\frac{\pi}{\sqrt{2}}\sqrt{\frac{(\lambda - |E|) + \epsilon'_{i_{\max}}}{\lambda - |\mathcal{H}_{i_{\max},r}|}}\epsilon' \leq c\frac{\sqrt{3}\pi}{2}\epsilon' \leq \epsilon'. \qquad (3.63)
\end{aligned}$$

If $\mathcal{I}_{i_{\max,2},1} \geq 1 - \epsilon'/2$, the estimate for E in the full range $[-\lambda, \lambda]$ is already correct to error ϵ' and we terminate the algorithm. Otherwise, we now have to determine which of $h \in [-\lambda, -\epsilon'/2)$ or $h \in (\epsilon'/2, \lambda]$ is true. This is accomplished using controlled gapped phase estimation from Corollary 3.11 using $\mathrm{CGPE}_{\pi/2-\mathcal{I}_{i_{\max},\mathrm{r}},\pi/2,\Delta_3}$. With failure probability at most Δ_3, we measure and obtain $|01\rangle$ if $\theta \in [0, \mathcal{I}_{i_{\max},\mathrm{r}}]$, and obtain $|11\rangle$ if $\theta \in [\pi - \mathcal{I}_{i_{\max},\mathrm{r}}, \pi]$. Using the inequality

$$\frac{\pi}{2} - \mathcal{I}_{i_{\max},\mathrm{r}} = \frac{\pi}{2} - \arccos(|\mathcal{H}_{i_{\max},\mathrm{r}}|/\lambda) = \arcsin(|\mathcal{H}_{i_{\max},\mathrm{r}}|/\lambda) \geq \frac{|\mathcal{H}_{i_{\max},\mathrm{r}}|}{\lambda} \geq \frac{|E|}{\lambda}\,. \tag{3.64}$$

This has query complexity

$$Q_3 = O\left(\frac{\lambda}{\max(\epsilon', |E|)}\log\frac{1}{\Delta_3}\right) = \begin{cases} O(\log\frac{1}{\Delta_3}), & |E| \geq \lambda/2, \\ O(\frac{\sqrt{\lambda}\sqrt{\lambda - |E|}}{\epsilon'}\log\frac{1}{\Delta_3}), & |E| < \lambda/2. \end{cases} \tag{3.65}$$

We then set $\Delta_1 = \Delta_2 = \Delta_3 = q'/3$. Then the overall query complexity to estimate E to additive error ϵ' and failure probability at most q' is

$$Q = Q_1 + Q_2 + Q_3 = O\left(\frac{\sqrt{\lambda}}{\epsilon'}\left(\sqrt{\epsilon'} + \sqrt{\lambda - |E|}\right)\log\frac{1}{q'}\right). \tag{3.66}$$

The gate complexities and qubit overhead follow from the construction of the walk operator and gapped phase estimation procedure. Finally, observe that $\sqrt{\lambda}\sqrt{\lambda - |E|} \leq \sqrt{\lambda + E}\sqrt{\lambda - E}$, and also that $x + y = O(\max\{x, y\})$.

□

As a bonus, we now give a proof of amplified amplitude amplification, which is a special case of Theorem 3.12 instantiated with the spectral amplified block-encoding Lemma 3.7. This improves on recent work, e.g. [Sim+24, Lemma 2], which required a known upper $\Delta \geq |\langle\psi|\Pi|\psi\rangle|$, and had poorer query complexity $O(\frac{\sqrt{\Delta}}{\epsilon}\log(\frac{1}{q})\frac{\Delta}{\Delta - |\langle\psi|\Pi|\psi\rangle|})$.

Corollary 3.13 (Amplified Amplitude Estimation)**.** *Let the projector $\Pi^2 = \Pi$ be block-encoded by BE$[\Pi]$. Let the state preparation orcale $P|0\rangle = |\psi\rangle$. Then $\langle\psi|\Pi|\psi\rangle$ can be estimated to additive error ϵ with confidence $1-q$ using $O\left(\frac{\sqrt{\max\{\epsilon, \langle\psi|\Pi|\psi\rangle(1-\langle\psi|\Pi|\psi\rangle)\}}}{\epsilon}\log\frac{1}{q}\right)$ queries to BE$[\Pi]$, P and their inverse.*

Proof. Let $H = 2\Pi/\lambda - \mathbb{1}$, where $\lambda = 1$. Using Lemma 3.7, we block-encode BE$[H]$ using one query to BE$[\Pi]$ and its inverse. As $E = \langle\psi|H|\psi\rangle = 2\langle\psi|\Pi|\psi\rangle - 1$, Theorem 3.12 states that the query complexity to estimate $\langle\psi|H|\psi\rangle$ to additive error ϵ and confidence $1 - q$ is

$$O\left(\frac{\sqrt{\lambda\max\{\epsilon, \lambda - |E|\}}}{\epsilon}\log\frac{1}{q}\right) = O\left(\frac{(\sqrt{\max\{\epsilon, \langle\psi|\Pi|\psi\rangle(1 - \langle\psi|\Pi|\psi\rangle)\}}}{\epsilon}\log\frac{1}{q}\right). \tag{3.67}$$

□

Phase estimation by spectral amplification

In this section, we present new quantum algorithms summarized in Table 3.3 that exploit SA to improve phase estimation of the energy of quantum ground states. Given a block-encoding BE$[H/\lambda]$ and a unitary P preparing the state $|\psi\rangle$, such that the overlap $p > 0$ with the ground state $|\psi_0\rangle$ is known *a priori*, the goal is to estimate the corresponding ground state energy E, where $H|\psi_0\rangle = E|\psi_0\rangle$, to additive error ϵ and confidence $1-q$. In contrast to the typical approach that scale like $O(\lambda/\epsilon)$ queries, we show that SA leads to explicit scaling with the improved factor $\sqrt{\lambda(\lambda + E)} \leq \lambda$. Our generalized routines naturally recover previously known results without SA, which is the $|E| \ll \lambda$ limit, in addition to achieving better qubit or query complexities. Our results in terms of queries to BE$[H/\lambda]$ is most general: If there

is an operator $H' = H_{\mathrm{SA}}^\dagger H_{\mathrm{SA}}$ and we are provided the block-encoding $\mathrm{BE}[H_{\mathrm{SA}}/\sqrt{\lambda}]$, then Lemma 3.7 informs us that we may identify $H/\lambda \leftarrow (2H_{\mathrm{SA}}^\dagger H_{\mathrm{SA}} - \lambda)/\lambda$. If the ground state energy of $H_{\mathrm{SA}}^\dagger H_{\mathrm{SA}}$ is small, we automatically realize SA as the factor $E_0 = \langle\psi_0| H_{\mathrm{SA}}^\dagger H_{\mathrm{SA}} |\psi_0\rangle - \lambda$ is close to $-\lambda$. We note that the expectation estimation results of Section 3.2 could be reduced to the specific $p = 1$ case of this section. Out algorithms improve on previous work (Row 3 of Table 3.3) in a few key areas:

- In the case where $|\psi\rangle$ is supported on the eigenstates with energy less than an *a priori* known upper bound $\Delta \geq \lambda + E$, our non-adaptive algorithm Theorem 3.14 returns an estimate of the energy of one of these eigenstates to additive error ϵ using $O(\sqrt{\lambda\Delta}/\epsilon)$ queries.

- In the case where we specifically want the ground state energy, for which, like previous work, we have no *a priori* known upper bound, our adaptive algorithm Theorem 3.15 matches the scaling of previous results in the parameters p, q, and scale with the improved factor $\sqrt{\lambda \max\{\epsilon, \lambda + E\}}/\epsilon$. This factor is upper bounded by previous results, e.g. the $|E| \ll \lambda$ case. Moreover, it exhibits novel super-Heisenberg scaling like $O(\sqrt{\lambda/\epsilon})$ when $\epsilon = \Theta(\lambda + E)$, a result which was previously unknown.

- Through the use of improved gapped phase estimation techniques Lemma 3.10, we reduce qubit overhead from a logarithimic factor to just a constant 2. This could be relevant in practical implementations of the algorithm.

At a high-level, the proof our non-adaptive algorithm Theorem 3.14 is very similar to that of Theorem 3.9 – perform phase estimation to accuracy $\epsilon' = \lambda/\epsilon$, and propagate the arccos nonlinearity with the *a priori* known upper bound Δ. The proof of the adaptive algorithm Theorem 3.15 also mirrors that of Theorem 3.12 in that it uses binary search by multiple iterations of gapped phase estimation, and also performs gapped phase estimation in two steps: First to a phase estimation accuracy of $\epsilon' = \sqrt{\epsilon/\lambda}$, which is guaranteed to either tell us that E is $-\lambda$ to error ϵ and so we terminate the algorithm, or give us an estimate $\lambda + \hat{E}$ of $\lambda + E$ to constant multiplicative error. Second, use the estimate $\hat{E}$ to choose the number of additional phase estimation and their accuracy and confidence parameters.

Year	Reference	Query complexity $O(\cdot)$		Extra qubits	Comments
		Block-encoding	State preparation		
2017	[GTC18]	$\frac{\lambda^{3/2}}{\sqrt{p}\epsilon^{3/2}} \log^{\Theta(1)}\left(\frac{\lambda}{p\epsilon q}\right)$	$\frac{1}{\sqrt{p}}\sqrt{\frac{\lambda}{\epsilon}} \log^{\Theta(1)}\left(\frac{\lambda}{p\epsilon q}\right)$	$O\left(\log\left(\frac{1}{\epsilon}\right)\right)$	
2020	[LT20]	$\frac{\lambda}{\sqrt{p}\epsilon} \log\left(\frac{\lambda}{\epsilon}\right) \log\left(\frac{1}{p}\right) \log\left(\frac{\log(\lambda/\epsilon)}{q}\right)$	$\frac{1}{\sqrt{p}} \log\left(\frac{\lambda}{\epsilon}\right) \log\left(\frac{\log(\lambda/\epsilon)}{q}\right)$	$O\left(\log\left(\frac{1}{p}\right)\right)$	
2024	[Ber+24b]	$\frac{\lambda}{\sqrt{p}\epsilon} \log\left(\frac{1}{p}\right) \log\left(\frac{1}{q}\right)$	$\frac{1}{\sqrt{p}} \log\left(\frac{\lambda}{\epsilon}\right) \log\left(\frac{\log(\lambda/\epsilon)}{q}\right)$	$O\left(\log\left(\frac{\lambda \log(1/q)}{\epsilon p}\right)\right)$	
2025	[Low+25], Theorem 3.14	$\frac{\sqrt{\Delta\lambda}}{\epsilon} \log\left(\frac{1}{q}\right)$	1	$O(\log(1/\epsilon))$	$((\lambda + E)$ of $\|\psi\rangle) \leq \Delta$, $p = 1$.
2025	Theorem 3.15	$\frac{\sqrt{\max\{\epsilon, \lambda+E\}\lambda}}{\sqrt{p}\epsilon} \log\left(\frac{1}{p}\right) \log\left(\frac{1}{q}\right)$	$\frac{1}{\sqrt{p}} \log\left(\frac{\lambda}{\epsilon}\right) \log\left(\frac{\log(\lambda/\epsilon)}{q}\right)$	2	

Table 3.3: Cost of estimating an *a* priori unknown ground state energy E of H to additive error ϵ and failure probability q, given query access to the block-encoding of $\mathrm{BE}[H/\lambda]$ and a state preparation unitary for a trial state with overlap $\sqrt{p}$ with the ground state.

First we consider phase estimation, the result in the first line of Table 3.1. This is the key result which enabled the improved quantum chemistry compilations in Ref. [Low+25]. The following proof is based on $\mathbf{H}_{\mathrm{SA}}$ for simplicity; a similar proof follows using H_{SA} combined with Lemma 3.7.

Theorem 3.14 (Phase estimation with SA)**.** *Let $H \succeq 0$ be a Hamiltonian of the form given in Equation* (3.18)*. Let $|\psi\rangle$ be a low-energy eigenstate supported on the subspace of energy of H at most $\Delta > 0$. We can perform phase estimation of H on the state $|\psi\rangle$ to additive precision ϵ with $O(\sqrt{\Delta\lambda}/\epsilon)$ calls to $BE[H_{SA}/\sqrt{\lambda}]$, the block-encoding of H_{SA}, and its inverse.*

Proof. Consider the related problem of performing quantum phase estimation on the Hamiltonian $\mathbf{H}_{\mathrm{SA}}$ of Eq. (3.27) within additive precision $\epsilon > 0$. This can be done by using the block-encoding $\mathrm{BE}[\mathbf{H}_{\mathrm{SA}}/\sqrt{\lambda}]$, and this block encoding can be easily implemented with one call to (controlled) $\mathrm{BE}[H_{\mathrm{SA}}/\sqrt{\lambda}]$ and one call to its inverse. (An explicit construction of this block-encoding is in Ref. [ZS24].) Quantum phase estimation necessitates $O(\sqrt{\lambda}/\epsilon)$ uses of these block-encodings for this precision, which is known as the 'Heisenberg limit' [Ber+18; KOS07]. Next we note that

$$(\mathbf{H}_{\mathrm{SA}})^2 = \begin{pmatrix} H & \mathbf{0} \\ \mathbf{0} & H_{\mathrm{SA}} H_{\mathrm{SA}}^\dagger \end{pmatrix}, \tag{3.68}$$

implying that, if $E \geq 0$ is the desired eigenvalue of H (i.e., $H|\psi\rangle = E|\psi\rangle$), $\pm\sqrt{E}$ are also eigenvalues of $\mathbf{H}_{\mathrm{SA}}$. Indeed, the (at most) two dimensional subspace spanned

by $\{|\psi\rangle|0\rangle, \mathbf{H}_{\mathrm{SA}}|\psi\rangle|0\rangle\}$ is invariant under the action of $\mathbf{H}_{\mathrm{SA}}$ and is the one that gives rise to the $\pm\sqrt{E}$. Hence, phase estimation using $\mathbf{H}_{\mathrm{SA}}$ with initial state $|\psi\rangle|0\rangle$ will produce an estimate of $\sqrt{E}$ or $-\sqrt{E}$.

Let ζ be such an estimate; for example consider $\sqrt{E}$ in which case we have $|\zeta-\sqrt{E}| \leq \epsilon$. (The analysis can be applied to the case of $-\sqrt{E}$.) If we take the square we obtain

$$E - 3\sqrt{E}\epsilon \leq (\sqrt{E}-\epsilon)^2 \leq \zeta^2 \leq (\sqrt{E}+\epsilon)^2 \leq E + 3\sqrt{E}\epsilon\,, \tag{3.69}$$

where we assumed the precision to satisfy $\epsilon \leq \sqrt{E}$. Hence, if we set $\epsilon \leq \epsilon/(3\sqrt{E})$ we can obtain E within additive precision ϵ, which is the desired goal.

To this end, we run quantum phase estimation with $\mathbf{H}_{\mathrm{SA}}$, initial state $|\psi\rangle|0\rangle$, and additive precision $\epsilon \leq \epsilon/(3\sqrt{E})$. The number of calls to (controlled) $\mathrm{BE}[H_{\mathrm{SA}}/\sqrt{\lambda}]$ and its inverse will be $O(\sqrt{E\lambda}/\epsilon)$. The result follows from the assumption $E \leq \Delta$. □

Previous work [Ber+24b] summarized in Table 3.3 solves this problem with a query complexity of

$$\text{Queries to BE}\left[\frac{H}{\lambda}\right] = O\left(\frac{\lambda}{\sqrt{p}\epsilon}\log\left(\frac{1}{p}\right)\log\left(\frac{1}{q}\right)\right), \tag{3.70}$$

$$\text{Queries to } P = O\left(\frac{1}{\sqrt{p}}\ln\left(\frac{\lambda}{\epsilon}\right)\ln\left(\frac{\ln(\lambda/\epsilon)}{\sqrt{q}}\right)\right). \tag{3.71}$$

The key idea is reducing the estimation of E to a fuzzy binary search problem. Roughly speaking, at each iteration $i = 0, \cdots, i_{\max}-1$, an interval $\mathcal{I}_i$ containing E with high confidence $1-q_i$ is found with query complexity $O(\frac{1}{|\mathcal{I}_i|}\log(\frac{1}{q_i}))$. As i increases, the width $|\mathcal{I}_i|$ is chosen to decrease geometrically. Then by a union bound over, the final iteration estimates E_{SOS} with additive error $|\mathcal{I}_i|/2$, failure probability $q = \sum_{i=0}^{i_{\max}-1} q_i$. Now suppose that E is close to $-\lambda$, such as by a appropriate H_{SA} construction. Then a modification of exploiting the arccos walk nonlinearity in the original proof of Equation (3.70) leads to an improved spectral amplified query complexity.

Theorem 3.15 (Fast ground state energy estimation by spectral amplification with no prior)**.** *Let any Hermitian $H \in \mathbb{C}^{N\times N}$ have ground state $|\psi_0\rangle$ with unknown energy E. Let $|\psi\rangle$ be prepared by the state preparation unitary $P|0\rangle_{\mathrm{S}} = |\psi\rangle$ such that the overlap $|\langle\psi|\psi_0\rangle| \geq \sqrt{p}$. Let $BE[H/\lambda] \in \mathbb{C}^{D\times D}$ be a block-encoding of H_{SA}. Then if $E \in [-\lambda, 0]$, it can be estimated to any additive error ϵ and confidence $1-q$*

using

$$\textit{Queries to BE} \left[\frac{H}{\lambda}\right] = O\left(\frac{\sqrt{\max\{\epsilon, \lambda - |E|\}\lambda}}{\sqrt{p}\epsilon} \log\left(\frac{1}{p}\right) \log\left(\frac{1}{q}\right)\right), \tag{3.72}$$

$$\textit{Queries to } P = O\left(\frac{1}{\sqrt{p}} \log\left(\frac{\lambda}{\epsilon}\right) \log\left(\frac{\log(\lambda/\epsilon)}{q}\right)\right), \tag{3.73}$$

and their inverses, and two ancillary qubits and $O(Q \log(D/N))$ *arbitrary two-qubit gates.*

Proof of Equation (3.72). We find it convenient to work in units of phase. Then using qubitization Lemma 3.4 on $\textsc{Be}[H/\lambda]$, we form a quantum walk W. For each eigenstate $|\psi_j\rangle$ of $H|\psi_j\rangle = E_j|\psi_j\rangle$, let $|\psi_j\rangle = \frac{1}{\sqrt{2}}(|\psi_j^+\rangle + |\psi_j^-\rangle)$, where $|\psi_j^\pm\rangle$ are all mutually orthogonal. Then the eigenphases of the quantum walk are

$$-W|\psi_j^\pm\rangle = \exp(i\Theta_\pm), \quad \Theta_{j,\pm} = \pi \mp i \arccos(\alpha_j), \quad \alpha_j = o_j - 1, \quad o_j = \frac{E_j + \lambda}{\lambda}, \tag{3.74}$$

where the minimum energy is $E := E_0$. Let $\alpha := \alpha_0$, $o := o_0 \in [0, 2]$, and $\theta := |\Theta_{0,\pm}| \in [0, \pi]$, where we choose the principal range of $\Theta_{j,\pm} \in [-\pi/2, \pi/2]$. Let us further assume that $o \in [0, 1]$, hence $\theta \in \mathcal{I}_0 := [0, \pi/2]$.

Let us now perform a search for θ. Let the iterations of the search be indexed by $i = 0, \cdots, i_{\max} - 1$, for some $i_{\max} > 0$ At each iteration, θ is known to be contained in an interval $\mathcal{I}_i := [\mathcal{I}_{i,\mathrm{l}}, \mathcal{I}_{i,\mathrm{r}}]$ with high probability p_i. At each iteration $i \geq 0$, we split $\mathcal{I}_i$ into thirds

$$\mathcal{I}_{i,\uparrow} := \left[\mathcal{I}_{i,\mathrm{l}} + \frac{1}{3}|\mathcal{I}_i|, \mathcal{I}_{i,\mathrm{r}}\right], \quad \mathcal{I}_{i,\downarrow} := \left[\mathcal{I}_{i,\mathrm{l}}, \mathcal{I}_{i,\mathrm{r}} - \frac{1}{3}|\mathcal{I}_i|\right], \tag{3.75}$$

and we will assign $\mathcal{I}_{i+1}$ to be either $\mathcal{I}_{i,\uparrow}$ or $\mathcal{I}_{i,\downarrow}$. Note that the width $|\mathcal{I}_i| = \frac{\pi}{2} r^i$, where $r = 2/3$.

We determine this assignment using gapped phase estimation $\textsc{Gpe}_{\varphi_i,\theta_i,\delta_i}$ Lemma 3.10, with $\theta_i = \frac{\mathcal{I}_{i,\mathrm{l}} + \mathcal{I}_{i,\mathrm{r}}}{2}$ at the midpoint of $\mathcal{I}_i$ and $\varphi_i = \frac{1}{6}|\mathcal{I}_i|$. This prepares a state

$$\begin{aligned}
&\textsc{Gpe}_{\varphi_i,\theta_i,\delta_i}|0\rangle|0\rangle|\psi_t\rangle \\
&= \textsc{Gpe}_{\varphi_i,\theta_i,\delta_i}|0\rangle\left(\sqrt{\frac{p}{2}}(|\psi_{0,+}\rangle + |\psi_{0,-}\rangle) + \sum_{j>0}\cdots\sum_{\pm}|\psi_{j,\pm}\rangle\right) \\
&= \sqrt{\frac{p}{2}}\sum_{\pm}(\alpha(\pm\theta)|0\rangle + \beta(\pm\theta)|1\rangle)|\psi_{0,\pm}\rangle) + \sum_{j>0}\cdots\sum_{\pm}(\alpha(\Theta_{j,\pm})|0\rangle + \beta(\Theta_{j,\pm})|1\rangle)|\psi_{j,\pm}\rangle\,.
\end{aligned} \tag{3.76}$$

Let us measure the $\{|0\rangle, |1\rangle\}$ register to obtain outcome $|m\rangle$. Let $\kappa = |(|0\rangle\langle 0| \otimes I)\textsc{Gpe}_{\varphi_i,\theta_i,\delta_i}|0\rangle|0\rangle|\psi_t\rangle|^2$ be the probability that we obtain outcome $m = 1$. Assuming that θ is promised to be one of the two following cases:

1. $\theta > \theta_i + \varphi_i$: The probability that we obtain outcome $m = 1$ is

$$\begin{aligned}\kappa &= \sum_j |\langle\psi_j|\psi_t\rangle|^2 \frac{1}{2}\sum_{\pm} |\beta(\Theta_{j,\pm})|^2 \\ &\leq \sum_j |\langle\psi_j|\psi_t\rangle|^2 \max_{j,\pm} |\beta(\Theta_{j,\pm})|^2 = \max_{j,\pm} |\beta(\Theta_{j,\pm})|^2 \leq \delta_i. \end{aligned} \tag{3.77}$$

2. $\theta \leq \theta_i - \varphi_i$: The probability that we obtain outcome $m = 1$ is

$$\begin{aligned}\kappa &= \sum_j |\langle\psi_j|\psi_t\rangle|^2 \frac{1}{2}\sum_{\pm} |\alpha(\Theta_{j,\pm})|^2 \\ &= \frac{p}{2}\sum_{\pm}(1 - |\alpha(\pm\theta)|^2) + \sum_{j>0} |\langle\psi_j|\psi_t\rangle|^2 \frac{1}{2}\sum_{\pm} |\alpha(\Theta_{j,\pm})|^2 \\ &\geq p(1-\delta_i). \end{aligned} \tag{3.78}$$

Let us choose $\delta_i = p/4$. Hence, we can determine which of the two cases hold by deciding whether $\kappa \leq \delta_i = p/4$ or $\kappa \geq p(1-\delta_i) \geq p(1-p/4) \geq 3p/4$. Now define the operator O and state $|\psi\rangle$ to be

$$O := (|0\rangle\langle 0| \otimes I)\textsc{Gpe}_{\varphi_i,\theta_i,\delta_i}, \quad |\psi\rangle = |0\rangle|0\rangle|\psi_t\rangle, \quad \Rightarrow \quad \kappa = \langle\psi| O^\dagger O |\psi\rangle. \tag{3.79}$$

As we can block-encode $\mathrm{BE}[O/1]$ using a single controlled-Not gate and 1 query to $\textsc{Gpe}_{\varphi_i,\theta_i,\delta_i}$, we can solve this decision problem by applying either Theorem 3.9 or Theorem 3.12 to find an estimate $\hat{\kappa}$ of $\kappa \leq 3p/4$ to additive error $p/2$ and failure probability q_i using $Q_{P,i} = O\left(\frac{1}{\sqrt{p}}\log\frac{1}{q_i}\right)$ queries to P and its inverse, and $Q_{H,i} = O\left(\frac{1}{\sqrt{p}r^i}\log\frac{1}{p}\log\frac{1}{q_i}\right)$ queries to $\mathrm{BE}[H/\lambda]$.

Hence, we make the assignment

$$\mathcal{I}_{i+1} = \begin{cases} \mathcal{I}_{i,\downarrow}, & \hat{\kappa} > p/4, \\ \mathcal{I}_{i,\uparrow}, & \hat{\kappa} \leq p/4. \end{cases} \tag{3.80}$$

The estimate of θ in $\mathcal{I}_i$ converts to an estimate of

$$\alpha \in -\mathcal{O}s(\mathcal{I}_i) = [-\mathcal{O}s(\mathcal{I}_{i,\mathrm{l}}), -\mathcal{O}s(\mathcal{I}_{i,\mathrm{r}})] = [-\mathcal{O}s(\theta_i - |\mathcal{I}_i|/2), -\mathcal{O}s(\theta_i + |\mathcal{I}_i|/2)], \tag{3.81}$$

$$\epsilon_i = |\mathcal{O}s(\mathcal{I}_i)| = 2\sin(\theta_i)\sin(|\mathcal{I}_i|/2),$$
$$E + \lambda \in \mathcal{H}_i = [\mathcal{H}_{i,\mathrm{l}}, \mathcal{H}_{i,\mathrm{r}}], \quad \mathcal{H}_i := \lambda(1 - \mathcal{O}s\mathcal{I}_i),$$
$$\epsilon_i' := |\mathcal{H}_i| = \lambda\epsilon_i,$$

where ϵ_i and ϵ_i' is the additive error to which α and E is known respectively.

The remainder of the proof mirrors that of Theorem 3.12 starting from Equation (3.82). We choose $q_i = \frac{6}{\pi^2}\frac{q/2}{(i_{\max}-i+1)^2}$. Then by a union bound, the failure probability after all $i_{\max}$ steps is at most $q/2$ and we have identified with at least probability $1 - q/2$ that $\theta \in \mathcal{I}_{i_{\max}}$ with additive error at most $|\mathcal{I}_{i_{\max}}| = \frac{\pi}{2}r^{i_{\max}}$. The total query complexity is then

$$Q_H = \sum_{i=0}^{i_{\max}-1} \mathcal{O}\left(\frac{1}{\sqrt{p}r^i}\log\frac{1}{p}\log\frac{1}{q_i}\right) = \mathcal{O}\left(\left(\frac{2}{3}\right)^{i_{\max}}\frac{1}{\sqrt{p}}\log\frac{1}{p}\log\frac{1}{q}\right), \tag{3.82}$$

$$Q_P = \sum_{i=0}^{i_{\max}-1} \mathcal{O}\left(\frac{1}{\sqrt{p}}\log\frac{1}{q_i}\right) = \mathcal{O}\left(\frac{i_{\max}}{\sqrt{p}}\log\frac{i_{\max}}{q}\right). \tag{3.83}$$

We now make two choices of $i_{\max}$:

1. Let $i_{\max} = i_{\max,1} + \Theta(1)$, where $i_{\max,1} = \frac{1}{2}\log_{1/r}\frac{\lambda}{\epsilon'}$. Then we either determine that $E + \lambda \le \epsilon'$ and terminate the algorithm, or obtain an estimate $\hat{E} = \Theta(E)$ to constant multiplicative error.

2. Let $i_{\max} = i_{\max,2} + \Theta(1)$, where $i_{\max,2} = \log_{1/r}\frac{\sqrt{\hat{E}\lambda}}{\epsilon'}$. Then we estimate E to additive error at most ϵ'.

The overall query complexity for these two loops over different $i_{\max}$ is

$$Q_H = \mathcal{O}\left(\frac{\sqrt{\lambda}}{\sqrt{p}\epsilon'}(\sqrt{\epsilon'} + \sqrt{E+\lambda})\log\frac{1}{p}\log\frac{1}{q}\right), \tag{3.84}$$

$$Q_P = \mathcal{O}\left(\frac{\log(\lambda/\epsilon')}{\sqrt{p}}\log\frac{\log(\lambda/\epsilon')}{q}\right). \tag{3.85}$$

□

Spectral amplification and linear combination of unitaries

Thus far we provided the query complexities for various simulation tasks using spectral amplification. These results required access to $\mathrm{BE}[H_{\mathrm{SA}}/\sqrt{\lambda}]$, which can be constructed from the $\mathrm{BE}[A_j/a_j]$ as explained in Lemma 3.6. In applications we will need to construct these block-encodings from some presentation of the Hamiltonian terms and here we discuss the LCU presentation within the context of SA. To this end, we will assume that the terms $h_j \succeq 0$ in Eq. (3.18) can be expressed as $h_j = A_j^\dagger A_j$, for some operators A_j that are expressed as

$$A_j = \sum_{l=0}^{L-1} a_{jl}\sigma_{jl} \ , \tag{3.86}$$

where the coefficients $a_{jl} \geq 0$ without loss of generality and the σ_{jl}'s are unitary, e.g., Pauli strings. We seek to construct a block-encoding for H_{SA} from accessing these unitaries. We can obtain the following result.

Lemma 3.16. *Let $H_{SA} = \sum_{j=0}^{R-1} |j\rangle_a \otimes A_j$ be as in Eq.* (3.19)*, where A_j is presented in Eq.* (3.86)*. Then, we can implement a block-encoding $BE[H_{SA}/\sqrt{\lambda}]$, where*

$$\lambda := \sum_{j=0}^{R-1} (\|\vec{a}_j\|_1)^2 \ , \ \|\vec{a}_j\|_1 := \sum_{l=0}^{L-1} a_{jl} \ . \tag{3.87}$$

This requires $O(R \times L \times C)$ gates, where C is the gate cost of the σ_{jl}'s.

Proof. The proof follows the steps of Lemma 3.6. To implement SELECT, we need to replace $\mathrm{BE}[A_j/a_j]$ in that proof by

$$\mathrm{BE}\left[\frac{A_j}{\sum_{l=0}^{L-1} a_{jl}}\right] = \mathrm{BE}\left[\frac{\sum_{l=0}^{L-1} a_{jl}\sigma_{jl}}{\sum_{l=0}^{L-1} a_{jl}}\right] \ . \tag{3.88}$$

Note that we have effectively replaced a_j in Lemma 3.6 by $\sum_{l=0}^{L-1} a_{jl} \equiv \|\vec{a}_j\|_1$ and A_j by the LCU. In that proof, $\lambda = \sum_j a_j$ and this would give the normalization factor $\lambda = \sum_j (\|\vec{a}_j\|_1)^2$ in this case.

The block-encoding can be constructed using standard techniques since it involves an LCU. We can define

$$\mathrm{SELECT}_j := \sum_{l=0}^{L-1} |l_{\mathrm{c}}\rangle\langle l_{\mathrm{c}}| \otimes \sigma_{jl} \tag{3.89}$$

and

$$\text{PREPARE}_j \left|0\right\rangle_{\text{c}} \mapsto \frac{1}{\sqrt{\sum_{l=0}^{L-1} a_{jl}}} \sum_{l=0}^{L-1} \sqrt{a_{jl}} \left|l\right\rangle_{\text{c}} \ , \tag{3.90}$$

so that

$$\left\langle 0\right|_{\text{c}} (\text{PREPARE}_j)^\dagger \cdot \text{SELECT} \cdot \text{PREPARE}_j \left|0\right\rangle_{\text{c}} = \frac{\sum_{l=0}^{L-1} a_{jl}\sigma_{jl}}{\sum_{l=0}^{L-1} a_{jl}} \ . \tag{3.91}$$

That is, $(\text{PREPARE}_j)^\dagger \cdot \text{SELECT} \cdot \text{PREPARE}_j$ gives the desired $\text{BE}[A_j/\sum_{l=0}^{L-1} a_{jl}] = \text{BE}[A_j/\|\vec{a}_j\|_1]$. The total gate cost is then dominated by the query cost of Lemma 3.6, which is $O(R)$, times the gate cost of this block-encoding $\text{BE}[A_j/\|\vec{a}_j\|_1]$, which is $O(L \times C)$, where C is the gate cost of the σ_{jl}. □

3.3 SOS optimization and example SOSSA block-encoding

In this section we provide necessary details for constructing the mathematical program that determines an SOS representation of a Hamiltonian given an operator basis for the SOS generators such that the lower bound energy is maximized. Recall that maximizing the lower bound is an important algorithmic component in the SOSSA protocol that can lower the query complexity and end-to-end gate complexity if balanced with the cost of the block-encoding implementation. We also provide an example block-encoding for H_{SOSSA} expressed as a polynomial of Pauli operators.

SOS optimization

Section 3 advocates for an SOS representation with increased complexity beyond termwise SA, which resulted in a loose lower bound that potentially negates an advantage through SA. In the following, we will restrict our exposition to Hamiltonians composed of Pauli strings but the construction can be applied more generally. We will also start with the formulation of the SOS optimization as a mathematical program and return later to the motivation of this form. Termwise SA uses a restricted algebra to form an SOS operator (in that case a projector) made from the linear combination of an identity operator and the Pauli operator of the Hamiltonian (Eq. (3.3)). The lower bound $-\beta$ on the ground state energy can be improved by providing more variational freedom in the SOS generators, the B_j operators of Eq. (3.6). The variational protocol can be formulated as a semidefinite program (SDP). As an illustrative example, consider the set of operators constituting the span

of one and two-qubit Pauli operators over N-qubits organized into a column vector

$$\vec{X}^{(2)} = \left(\mathbb{1}, X_1, \ldots, Z_N, X_1X_2, \ldots, Z_{N-1}Z_N\right)^T. \tag{3.92}$$

If our Hamiltonian is 2-local, then this algebra is sufficient to represent the Hamiltonian. In the following, we consider $\vec{X}^{(k)}$ which involves operators involving degree-k products of Pauli operators. A particular SOS generator is defined as $B_j := \vec{b}_j\vec{X}^{(k)}$ for a row vector, $\vec{b}_j \in \mathbb{C}^{\binom{N}{k}}$, of complex coefficients to be optimized in order to satisfy Eq. (3.6). To show the optimization protocol is an SDP, we form the SOS representation of the Hamiltonian from a set of operators $\{B_j\}$:

$$\sum_{j=0}^{R-1} B_j^\dagger B_j = \sum_{j=0}^{R-1} \left(\vec{b}_j\vec{X}^{(k)}\right)^\dagger \left(\vec{b}_j\vec{X}^{(k)}\right) \tag{3.93}$$

$$= \left(\vec{X}^{(k)\dagger}\right)^T \left(\sum_j \vec{b}_j^\dagger\vec{b}_j\right)\vec{X}^{(k)} \tag{3.94}$$

$$= \left(\vec{X}^{(k)\dagger}\right)^T G\vec{X}^{(k)}\ , \tag{3.95}$$

which can be used to construct the program

$$\begin{aligned} &\min_G\ \beta \\ &\text{s.t.}\ \ H + \beta\mathbb{1} = \left(\vec{X}^{(k)\dagger}\right)^T G\vec{X}^{(k)} \\ &G \succeq 0. \end{aligned} \tag{3.96}$$

The above SDP equality constraints are short-hand for relating the coefficients of the operators obtained by expanding the right-hand side of the equality to the left-hand side (the Hamiltonian and the shift). This computationally the equality constraint is represented by letting the Hamiltonian coefficients in matrix form, $\mathbf{H} \in \mathbb{C}^{L\times L}$ where $L = |\vec{X}^{(k)}|$ provide the constants resulting from the Hilbert-Schmidt inner products, $\langle\mathbf{A}|G\rangle$, between constraint matrix $\mathbf{A}$ and G that encodes the coefficient relationship of the equality constraint. Similarly, the cost function corresponds to minimizing the coefficient associated with the identity operators, which in this case corresponds to the diagonal elements of the Gram matrix G. The SDP of Eq. (3.96) can be solved in polynomial time with respect to the linear dimension of G. Once the SDP is solved an SOS representation can be recovered by expressing G in its eigenbasis $G = \sum_j \mu_j\vec{a}_j\vec{a}_j^\dagger$, where $\{\vec{a}_j\}$ which yields a definition for each B_j

$$\vec{X}^{(k)\dagger}G\vec{X}^{(k)} = \sum_{j=0}^{R-1}(\sqrt{\mu_j}\vec{a}_j^\dagger\vec{X}^{(k)})^\dagger(\sqrt{\mu_j}\vec{a}_j^\dagger\vec{X}^{(k)}) = \sum_{j=0}^{R-1} B_j^\dagger B_j. \tag{3.97}$$

The number of B_j's in the generating polynomials is equal to the rank R of the Gram matrix G.

The intuition behind providing more variational freedom through the form of the SOS generator is that for any $-\beta$ that is less than the ground state energy E of some Hamiltonian H we can define their sum $\tilde{H} = H + \beta\mathbb{1}$ to be a positive semidefinite operator whose square root $\tilde{H} = \sqrt{\tilde{H}}\sqrt{\tilde{H}}$ which is potentially a different, and potentially more complicated, many-body body operator. Thus if the equality of Eq. (3.6) is satisfied then the SOS Hamiltonian serves as a certificate that the energy is lower bounded by $-\beta$. The certificate perspective is the primary method for constructing constraints on pseudomoment matrices in 'outer' approximation methods that seek to minimize representations of marginals and approximately constrain them through the sum-of-square construction [Hal07; Wit15; Erd78; Nak+01]. In fact, it can be shown that the SOS Hamiltonians can be determined from the primal problem directly [Low+25].

SOSSA block-encoding for Pauli operators

Utilizing the solution of the SDP to construct the SOS generators B_j, we can now replace these in the spectral amplified Hamiltonian; that is, we can replace $A_j \to B_j$ in Eqs. (3.19), (3.21). Expanding Eq. (3.97), each SOS generator is a linear combination of Pauli operators

$$B_j = \sum_{l=0}^{L-1} b_{jl}\sigma_l \ , \tag{3.98}$$

where $\{\sigma_l\}$ is the monomial basis and L is the number of monomials in the SOS optimization. To block encode H_{SOSSA} of Eq. (3.7), which is essentially H_{SA} in Eq. (3.19) after replacing $A_j \to B_j$, gives us the following block-encoding normalization factor.

Definition 3.3 (SOS λ)**.** *Define λ_{SOS} to be*

$$\lambda_{SOS} := \sum_{j=0}^{R-1} \|\vec{b}_j\|_1^2 \ , \ \|\vec{b}_j\|_1 = \sum_{l=0}^{L-1} b_{jl} \ . \tag{3.99}$$

This is the normalization factor in Lemma 3.16 for this choice of operators B_j, which readily implies the following.

Corollary 3.17 (Compilation of SOSSA with Pauli strings)**.** *Let $H_{SOSSA} = \sum_{j=0}^{R-1} |j\rangle\otimes B_j$, where the B_j's are in Eq.* (3.98)*. Then, we can construct the block-encoding $BE[H_{SOSSA}/\sqrt{\lambda_{SOS}}]$ with $\tilde{O}(R \times L)$ gates.*

Since the Pauli strings are of constant weight, C is a constant in Lemma 3.16.

In Section 3.4 we rely on upper and lower bounds on λ_{SOS}. To derive these, consider a general Hamiltonian $H \in \mathbb{C}^{M\times M}$ such that after the shift it produces $H + \beta\mathbb{1} = \sum_{j=0}^{R-1} B_j^\dagger B_j$ as in Eq. (3.6). Pauli strings are orthogonal and their trace is zero, implying

$$\sum_{j=0}^{R-1} \|\vec{b}_j\|_2^2 = \frac{1}{M}\mathrm{Tr}H + \beta \ . \tag{3.100}$$

This gives us bounds of λ_{SOS} via Cauchy-Schwarz:

$$\frac{1}{M}\mathrm{Tr}H + \beta \ \leq \ \lambda_{\mathrm{SOS}} \ \leq \ L\left(\frac{1}{M}\mathrm{Tr}H + \beta\right). \tag{3.101}$$

3.4 Application to SYK model

In this section we describe the application of the SOSSA framework to the SYK model. We will see a factor of $\sqrt{N}$ speedup in both query and gate complexities when compared to the standard LCU compilation. This is accomplished by showing that Δ_{SOS} scales linearly in system size with high probability and λ_{SOS} scales quadratically.

Let $\{\gamma_1, \ldots, \gamma_N\}$ be Majorana operators satisfying anticommutation relations:

$$\gamma_a\gamma_b + \gamma_b\gamma_a = 2\delta_{ab}\mathbb{1} \ . \tag{3.102}$$

The SYK model is described by a fermionic Hamiltonian containing all degree-4 Majorana terms whose coefficients are random Gaussians where

$$H_{\mathrm{SYK}} = \frac{1}{\sqrt{\binom{N}{4}}} \sum_{a,b,c,d} g_{abcd}\gamma_a\gamma_b\gamma_c\gamma_d \ , \tag{3.103}$$

$$g_{abcd} \sim \mathcal{N}(0,1) \text{ i.i.d. } . \tag{3.104}$$

To implement SOSSA, we will use the degree-2 Majorana SOS, which we describe in Section 3.4. We will then compile the resulting SOSSA Hamiltonian using the double factorization technique, which we recap in Section 3.4. In Section 3.4, we use random matrix theory to analyze the energy gap Δ_{SOS} of the degree-2 Majorana SOS on the SYK model, and we also analyze the normalization factor λ_{SOS} resulting from the double factorization block encoding. Putting these together demonstrate our asymptotic improvements.

Degree-2 Majorana SOS

For fermionic systems, a natural SOS ansatz is the degree-2 Majorana SOS, which we will use for the SYK model. In this section we describe how to implement degree-2 Majorana SOS for a general fermionic Hamiltonians of the form

$$H = i \sum_{a,b=1}^{N} K_{ab}\gamma_a\gamma_b - \sum_{a,b,c,d=1}^{N} J_{abcd}\gamma_a\gamma_b\gamma_c\gamma_d. \tag{3.105}$$

Our SOS is defined by the basis $\{\mathbb{1}, \gamma_a, \gamma_a\gamma_b\}$. Using Section 3.3, we can write the SOS relaxation as ($\beta \in \mathbb{R}$)

$$\begin{aligned} \mathrm{SOS}(H) \quad = \quad & \min \beta \\ & \text{s.t.} \quad H + \beta\mathbb{1} = \vec{X}^\dagger G \vec{X}, \\ & \qquad G \succeq 0, \end{aligned} \tag{3.106}$$

where

$$\vec{X} = \left(\mathbb{1}, \gamma_1, \ldots, \gamma_N, i\gamma_1\gamma_2, \ldots, i\gamma_{N-1}\gamma_N\right) \tag{3.107}$$

and G is the Gram matrix with dimension $1 + N + \binom{N}{2}$. Comparing the coefficients of $\mathbb{1}$ in the polynomial constraint equation reveals that $\beta = \mathrm{Tr}\, G$, since all other operators in the generating set are traceless. Comparing coefficients of the other monomials gives us the linear constraints on the SDP. Thus we can explicitly write the SDP

$$\min \mathrm{Tr}\, G \tag{3.108}$$

$$\begin{aligned} \text{s.t.} \quad & 0 = G(\mathbb{1}, \gamma_a) + G(\gamma_a, \mathbb{1}) + i\sum_c G(i\gamma_a\gamma_c, \gamma_c) - i\sum_c G(\gamma_c, i\gamma_a\gamma_c) && \forall a \\ & K_{ab} = G(\mathbb{1}, i\gamma_a\gamma_b) + G(i\gamma_a\gamma_b, \mathbb{1}) - iG(\gamma_a, \gamma_b) + iG(\gamma_b, \gamma_a) \\ & \qquad - i\sum_c G(i\gamma_a\gamma_c, i\gamma_b\gamma_c) + i\sum_c G(i\gamma_b\gamma_c, i\gamma_a\gamma_c) && \forall a, b \\ & 0 = G(\gamma_a, i\gamma_b\gamma_c) + G(\gamma_b, i\gamma_c\gamma_a) + G(\gamma_c, i\gamma_a\gamma_b) \\ & \qquad + G(i\gamma_b\gamma_c, \gamma_a) + G(i\gamma_c\gamma_a, \gamma_b) + G(i\gamma_a\gamma_b, \gamma_c) && \forall a, b, c \\ & J_{abcd} = G(i\gamma_a\gamma_b, i\gamma_c\gamma_d) + G(i\gamma_c\gamma_d, i\gamma_a\gamma_b) \\ & \qquad - G(i\gamma_a\gamma_c, i\gamma_b\gamma_d) - G(i\gamma_b\gamma_d, i\gamma_a\gamma_c) \\ & \qquad + G(i\gamma_a\gamma_d, i\gamma_b\gamma_c) + G(i\gamma_b\gamma_c, i\gamma_a\gamma_d) && \forall a, b, c, d \\ & G \succeq 0. \end{aligned}$$

Double factorization

The degree-2 Majorana SOS leads to a representation of the Hamiltonian that takes the form

$$H + \beta \mathbb{1} = \sum_j B_j^\dagger B_j \ , \tag{3.109}$$

$$B_j = e_j \mathbb{1} + \sum_a f_{j,a} \gamma_a + \sum_{ab} g_{j,ab} \gamma_a \gamma_b . \tag{3.110}$$

A similar factorization occurs in the 2-body term of quantum chemistry Hamiltonians. Using the direct block encoding strategy, Definition 3.3 gives for this SOS representation a SOS λ of

$$\lambda_{\mathrm{SOS}} = \sum_{j=0}^{R-1} \left(|e_j| + \|f_j\|_1 + \sum_{ab} |g_{j,ab}| \right)^2 , \tag{3.111}$$

where R is the rank of the Gram matrix G.

We can achieve a much better λ_{SOS} using the concept of double factorization from quantum chemistry [Bur+21b], which we now describe. See [Lee+21b] for a comprehensive discussion of quantum algorithms techniques for quantum chemistry. Given a quadratic polynomial in the Majorana operators $\sum_{ab} g_{ab} \gamma_a \gamma_b$, let's first decompose g_{ab} into its real and imaginary parts $g_{ab} = g_{ab}^{(R)} + i g_{ab}^{(I)}$. Without loss of generality, $g_{ab}^{(R)}$ and $g_{ab}^{(I)}$ are real antisymmetric, since γ_a and γ_b always anticommute. There are Gaussian unitaries $U^{(R)}, U^{(I)}$ which rotate the quadratic polynomials into the block-diagonal forms

$$\left(\sum_{ab} g_{ab}^{(R)} \gamma_a \gamma_b \right) = U^{(R)\dagger} \sum_a \tilde{g}_a^{(R)} \gamma_{2a-1} \gamma_{2a} U^{(R)} \tag{3.112}$$

$$\left(\sum_{ab} g_{ab}^{(I)} \gamma_a \gamma_b \right) = U^{(I)\dagger} \sum_a \tilde{g}_a^{(I)} \gamma_{2a-1} \gamma_{2a} U^{(I)} . \tag{3.113}$$

This is because a Gaussian unitary U acts on the N-vector of Majorana operators $(\gamma_1, \ldots, \gamma_N)$ via a real orthogonal matrix O, and we can always design U so that O block diagonalizes a given antisymmetric matrix g_{ab}:

$$\begin{pmatrix} g_{11} & g_{12} & \cdots & g_{1,N} \\ g_{21} & g_{22} & & \vdots \\ \vdots & & \ddots & \\ g_{N,1} & \cdots & & g_{N,N} \end{pmatrix} = O^T \begin{pmatrix} 0 & g_1 & & & \\ -g_1 & 0 & & & \\ & & \ddots & & \\ & & & 0 & g_N \\ & & & -g_N & 0 \end{pmatrix} O . \tag{3.114}$$

Let's find such Gaussian unitaries $U_j^{(R)}, U_j^{(I)}$ for each $g_{j,ab}$. This lets us write B_j as

$$B_j = e_j \mathbb{1} + \sum_a f_{j,a}\gamma_a + U^{(R)\dagger} \sum_a \tilde{g}^{(R)}_{(j,a)} \gamma_{2a-1}\gamma_{2a} U^{(R)} + U^{(I)\dagger} \sum_a i\tilde{g}^{(I)}_{(j,a)} \gamma_{2a-1}\gamma_{2a} U^{(I)}. \tag{3.115}$$

We can efficiently implement the Gaussian unitaries U_j on the quantum computer. Thus we can follow Definition 3.3 to derive λ_{SOS} for the double factorized SOS representation, and the proof of Corollary 3.17 will go through. We get

$$\lambda_{\text{SOS}} = \sum_{j=0}^{R-1} \left(|e_j| + \|f_j\|_1 + \|\tilde{g}_j^{(R)}\|_1 + \|\tilde{g}_j^{(I)}\|_1\right)^2. \tag{3.116}$$

This could be significantly cheaper than directly block encoding B_j, which would have λ_{SOS} as give in Equation (3.111).

Application to SYK model

We can now bring all of these ingredients together to analyze the performance of degree-2 Majorana SOSSA with double factorization on the SYK model, when compared to the standard LCU approach.

The asymptotic gate complexities to construct the necessary block-encodings for LCU and SOSSA are both $\Theta(N^4)$. First, notice that the number of terms in H_{SYK} is $\sim N^4$, with each coefficient independently random. This creates a $\Omega(N^4)$ information-theoretic lower bound on the gate cost of any faithful block encoding. LCU can be implemented with gate cost proportional to the number of terms. Double factorization too allows for a $O(N^4)$ gate cost: the cost of the B_j are dominated by the costs of the Gaussian unitaries, which are $O(N^2)$, and the number of B_j is at most $O(N^2)$. The asymptotic equivalence of the block encoding gate costs means that we can focus on comparing the query complexities λ_{LCU} and $\sqrt{\Delta_{\text{SOS}}\lambda_{\text{SOS}}}$.

For the SYK model,

$$\lambda_{\text{LCU}} = \frac{1}{\sqrt{\binom{N}{4}}} \sum_{a,b,c,d} |g_{abcd}| \tag{3.117}$$

which scales like $\lambda_{\text{LCU}} \sim N^2$. We will see that $\sqrt{\Delta_{\text{SOS}}\lambda_{\text{SOS}}} \sim N^{\frac{3}{2}}$, giving a factor of $\sqrt{N}$ improvement in the query complexity.

We begin by analyzing the scaling of Δ_{SOS}. This will depend on the quality of the SOS lower bound $-\beta$ arising from the degree-2 Majorana SOS described in Section 3.4. The SYK Hamiltonian is traceless, $\text{Tr}(H_{\text{SYK}}) = 0$, implying that

the ground state energy is negative, and it is known that with high probability the minimum and maximum eigenvalues are proportional to $\sqrt{N}$ [HO22]. The following lemma shows that applying degree-2 Majorana SOS to the SYK models gives a lower bound $-\beta$, where $\beta = O(N)$, giving an energy gap of $\Delta_{\text{SOS}} = O(N)$.

Lemma 3.18. *Degree-2 Majorana SOS achieves a lower bound of $-\beta$ on the ground energy of H_{SYK}, for $\beta = O(N)$, with high probability.*

Proof. Our proof strategy is to use the dual problem to the SOS optimization. The SOS optimization obeys Slater's condition and thus exhibits strong duality—see Section 5.2.3 of [BV04]. Thus, in order to show the lower bound where $\beta = O(N)$, it is sufficient to show that the lowest achievable energy in the dual problem is at least $-cN$ for some constant $c > 0$. That is, no solution to the dual problem can take a value lower than $-cN$.

Appendix B of [PNA10] allows us to express the dual optimization problem to the SOS. Define a degree-2 pseudoexpectation to be a matrix $\tilde{\rho}$ whose rows and columns are indexed by quadratic Majorana operators, and where we impose that $\tilde{\rho}$ is positive semi-definite and obeys algebraic constraints

$$\tilde{\rho}(i\gamma_a\gamma_b, i\gamma_c\gamma_d) = -\tilde{\rho}(i\gamma_a\gamma_c, i\gamma_b\gamma_d) = \tilde{\rho}(i\gamma_a\gamma_d, i\gamma_b\gamma_c) \tag{3.118}$$

$$=\tilde{\rho}(i\gamma_b\gamma_c, i\gamma_a\gamma_d) = -\tilde{\rho}(i\gamma_b\gamma_d, i\gamma_a\gamma_c) = \tilde{\rho}(i\gamma_c\gamma_d, i\gamma_a\gamma_b). \tag{3.119}$$

Let J be the matrix with rows/columns indexed by subsets of $\{1, \ldots, N\}$ of size 2, with entries

$$J(ab, cd) = \frac{1}{\sqrt{\binom{N}{4}}} g_{abcd}, \tag{3.120}$$

where g_{abcd} are the SYK coefficients. The dual problem is to minimize $\text{Tr}(J\tilde{\rho})$ over degree-2 pseudoexpectations $\tilde{\rho}$.

To complete the proof, we show that any degree-2 pseudoexpectation $\tilde{\rho}$ obeys $\text{Tr}(J\tilde{\rho}) \geq -cN$ with high probability. We have

$$\text{Tr}(J\tilde{\rho}) \geq -\|J\| \cdot \|\tilde{\rho}\|_1 = -\|J\| \cdot \text{Tr}(\tilde{\rho}) = -\binom{N}{2}\|J\| , \tag{3.121}$$

where $\|.\|$ is the operator norm. The first inequality was an application of matrix Holder's inequality [Bau11]. Random matrix theory tells us that $\|J\| = O(N^{-1})$ with high probability. For example, see Theorem 4.4.5 in Ref. [Ver18]. Thus $\text{Tr}(J\tilde{\rho}) \geq -cN$ for any degree-2 pseudoexpectation $\tilde{\rho}$ with high probability over the SYK disorder. □

Next we analyze λ_{SOS}, which will involve the double factorization block encoding technique described in Section 3.4. The degree-2 Majorana SOS with double factorization writes the SYK Hamiltonian as

$$H_{\mathrm{SYK}} + \beta\mathbb{1} = \sum_j B_j^\dagger B_j \tag{3.122}$$

$$B_j = e_j\mathbb{1} + \sum_a f_{j,a}\gamma_a + U^{(R)^\dagger}\sum_a \tilde{g}^{(R)}_{(j,a)}\gamma_{2a-1}\gamma_{2a}U^{(R)} + U^{(I)^\dagger}\sum_a i\tilde{g}^{(I)}_{(j,a)}\gamma_{2a-1}\gamma_{2a}U^{(I)}, \tag{3.123}$$

where the U_j are Gaussian unitaries. From Equation (3.116), we can calculate

$$\begin{aligned}
\lambda_{\mathrm{SOS}} &= \sum_j \left(|e_j| + \|f_j\|_1 + \|\tilde{g}_j^{(R)}\|_1 + \|\tilde{g}_j^{(I)}\|_1\right)^2 && (3.124)\\
&\le 4N\Big(\sum_j |e_j|^2 + \sum_j \|\tilde{f}_j\|_2^2 + \sum_j \|\tilde{g}_j^{(R)}\|_2^2 + \sum_j \|\tilde{g}_j^{(I)}\|_2^2\Big) \quad \text{(Cauchy-Schwarz)} && (3.125)\\
&\le 4N\cdot\overline{\mathrm{Tr}}\Big(\sum_j B_j^\dagger B_j\Big) && (3.126)\\
&= 4N\cdot\overline{\mathrm{Tr}}\big(H_{\mathrm{SYK}} + \beta\mathbb{1}\big) && (3.127)\\
&= 4N\beta && (3.128)\\
&= O(N^2), \quad \text{(Lemma 3.18)} && (3.129)
\end{aligned}$$

where $\overline{\mathrm{Tr}}$ denotes the normalized trace so that $\overline{\mathrm{Tr}}(\mathbb{1}) = 1$. Putting this together with the above analysis showing $\Delta_{\mathrm{SOS}} = O(N)$ gives $\sqrt{\Delta_{\mathrm{SOS}}\lambda_{\mathrm{SOS}}} = O(N^{\frac{3}{2}})$.

Chapter 4

LEARNING FERMIONIC OBSERVABLES

In many quantum experiments, the key objective is to prepare a target quantum state, which can be a large many-body state and may exhibit significant entanglement. As quantum computers become more powerful, such experiments can increasingly be carried out through quantum simulations, where a quantum algorithm prepares the desired state. The next challenge is to measure and extract physically relevant properties from this prepared state, which typically involves preparing many copies of the same unknown quantum state and performing measurements on those copies. Here, sample complexity—the number of state copies needed—can be a major bottleneck, especially since generating each copy might be expensive. In particular, full state tomography quickly becomes intractable for large many-body systems, as its sample complexity scales polynomially with the (exponentially large) dimension of the Hilbert space.

Physical properties are often given by expectation values of observables. An important example is the k-RDM (reduced density matrix) in quantum chemistry, which consists of the expectation values of local fermionic operators. One can efficiently compute most physically relevant local observables of a fermionic system (e.g., dipole moment, charge density, and importantly—energy) from just the 1- and 2-RDMs as a consequence of fermions being identical particles that interact pairwise. A number of important methods for post-processing the output of quantum simulations require k-RDMs with larger k. For example, subspace expansion techniques for approximating excited states from ground states via linear response typically require the 4-RDM [McC+17; Yos+22]. Perturbation theory [Guo+16; Sha+17] and multi-reference configuration interaction methods [Tak+20] for relaxing ground state calculations in small basis sets towards their continuum (large basis) limit often require the 4-RDM but converge even faster given access to higher order RDMs. There are popular impurity model schemes for extrapolating finite simulations of condensed phase fermionic systems towards their thermodynamic limits (e.g., density matrix embedding theory [Wou+16]) and hybrid quantum-classical schemes for quantum Monte Carlo [Hug+22b], which require the full 1-RDM.

A large body of literature exists on methods for computing the fermionic 2-RDM

matrix elements as a means of estimating the energy of chemical systems during the course of a quantum variational algorithm (see e.g., [VYI20; PZC23; BBO20; ZRM21]). The sample complexity of these schemes scale like $\sim n^k$ for the k-RDM. This could be a large polynomial in the system size which becomes highly expensive for large systems, and poses a large bottleneck which limits the utility of VQE algorithms. In this work we will see that we can *exponentially improve* this sample complexity to $\sim k \log n$ through use of entangled measurements.

The task of extracting many expectation values from copies of an unknown quantum state has been studied in quantum computing under the name *shadow tomography*. In this chapter we discuss protocols for shadow tomography, as introduced by Aaronson [Aar18], specialized to the case of local fermionic observables. Let $\mathcal{F}_k^{(n)}$ denote the set of k-body fermionic operators on n fermionic modes, where $k = O(1)$, and let ρ be an unknown n-qubit quantum state. Given copies of ρ, we would like to learn the expectation values $\mathrm{Tr}(\Gamma\rho)$ to precision ε for every $\Gamma \in \mathcal{F}_k^{(n)}$.

Definition 4.1 (Fermionic shadow tomography)**.** *The shadow tomography task for the k-body fermionic operators $\mathcal{F}_k^{(n)}$ is as follows. We are given copies of an unknown n-qubit state ρ, and our goal is to output estimates y_Γ such that with high probability*[1] *we have $|y_\Gamma - \mathrm{Tr}(\Gamma\rho)| \leq \epsilon$ for all $\Gamma \in \mathcal{F}_k^{(n)}$.*

In applications of shadow tomography, samples are often expensive and are the main criterion for efficiency. For example, suppose one applies shadow tomography in the readout stage of a quantum simulation algorithm. The sample complexity directly translates to the number of shots one must run of the quantum computation. Alternatively, if one is using shadow tomography to learn from a physical experiment, the sample complexity represents the number of repeats of the experiment.

One can use a very naive tomography protocol to perform this task. Suppose we are performing shadow tomography on a set $\mathcal{S}$ of operators. For a given operator $\Gamma \in \mathcal{S}$, if we measure its value $O((\log|\mathcal{S}|)/\epsilon^2)$ times (using one copy of ρ for each measurement), then we can ensure that the sample mean is within ϵ of $\mathrm{Tr}(\rho\Gamma)$ with probability at least $1 - 0.01/|\mathcal{S}|$. If we follow this procedure for each of the operators in $\mathcal{S}$, the union bound guarantees that with high probability they will all be ϵ-close to their true values. This algorithm uses $O((|\mathcal{S}|\log|\mathcal{S}|)/\epsilon^2)$ copies of the unknown state ρ and is computationally efficient.

[1] Throughout this chapter, we use "with high probability" to mean with probability at least 99%, say.

Remarkably, Aaronson described a protocol for shadow tomography of any set of bounded observables (such as Pauli operators) that uses exponentially fewer copies of the unknown state than the naive algorithm [Aar18; AR19; Bra+17]. The best known scaling of the number of samples for learning m general observables is $O(n(\log^2 m)/\varepsilon^4)$ [BO21]. However, the general shadow tomography schemes suffer from two major caveats: they are explicitly exponential in computational runtime (even when the number of observables m scales polynomially with n), and they require entangled measurements on many copies of the unknown state ρ at a time.

These caveats can be avoided for certain restricted sets of observables such as low-weight Pauli operators. For k-local Paulis with $k = O(1)$, there are simple and computationally efficient protocols to learn m observables with $O((\log m)/\epsilon^2)$ single-copy measurements [CW20; EHF19; BBO20; Jia+20a; HKP20]. The classical shadows framework [HKP20] provides a broader family of learning protocols that can also handle other sets of non-Pauli observables with single-copy measurements, notably including rank-1 observables which are relevant to fidelity estimation. However, classical shadows and other single-copy learning strategies become inefficient for higher weight Pauli operators.

To go beyond low-weight Paulis one can use a shadow tomography protocol developed in [HKP21] which learns *any* subset $\mathcal{S}$ of Pauli operators using $O((\log |\mathcal{S}|)/\epsilon^4)$ copies of ρ, and $\mathrm{poly}(|\mathcal{S}|, n, 1/\epsilon)$ runtime. The protocol proceeds in two stages. In the first stage—*learning magnitudes*—one computes estimates of the magnitudes $|\mathrm{Tr}(\rho P)|$ to within $\epsilon/4$ error (say), for all Paulis $P \in \mathcal{S}$. Remarkably, this can be achieved efficiently using only $O((\log |\mathcal{S}|)/\epsilon^4)$ two-copy measurements using the well-known Bell sampling procedure [Mon17]. It is based on measuring copies of $\rho \otimes \rho$ in the basis which simultaneously diagonalizes the operators $P \otimes P$ for all Paulis P. In the second stage—*learning signs*—one computes the signs of all Paulis $P \in \mathcal{S}$ that were estimated to have nonnegligible magnitude in the first stage. The learning signs protocol from Ref. [HKP21] proceeds by a sequence of gentle measurements which requires entangling $O((\log |\mathcal{S}|)/\epsilon^2)$ copies of ρ.

The requirement to perform joint entangled measurements on many copies of the unknown state ρ prohibits the use of shadow tomography techniques in a practical setting. Suppose we are interested in a 100-qubit quantum state ρ; for example, perhaps it is the groundstate of a chemical molecule. Using a quantum computer with 200 logical qubits, we can perform entangling measurements on two copies

$\rho \otimes \rho$. However, suppose we are interested in precision $\epsilon = 0.001$, and $\mathcal{S}$ is the set of 2-body fermionic operators (so $|\mathcal{S}| \simeq \binom{200}{4}$). In order to perform an entangling measurement on $O((\log |\mathcal{S}|)/\epsilon^2)$ copies of ρ, we would need over a *billion* logical qubits in our quantum computer.

This motivates us to define a notion of *triply efficient* shadow tomography.

Definition 4.2 (Triply efficient shadow tomography)**.** *A shadow tomography protocol on set $\mathcal{S}$ of operators is* triply efficient *if:*

Sample efficiency: The number of samples scales as poly$(\log |\mathcal{S}|, 1/\epsilon)$.

Computational efficiency: The classical and quantum computation is poly$(|\mathcal{S}|, n, 1/\epsilon)$.

Few-copy measurements: The algorithm uses joint measurements on a constant number of copies of ρ (ideally 1 or 2).

Is there a triply efficient shadow tomography protocol? While this question is well-posed for arbitrary subsets of observables $\mathcal{S}$, we restrict our attention to three subsets that are practically motivated and representative of the complexity of the problem:

$\mathcal{P}_k^{(n)}$: The set of k-local Pauli operators on n qubits, where $k = O(1)$.

$\mathcal{F}_k^{(n)}$: The set of k-body fermionic operators on n fermionic modes, where $k = O(1)$.

$\mathcal{P}^{(n)}$: The set of all Pauli operators on n qubits.

As we describe below, k-local Pauli operators and k-body fermionic operators arise in a variety of applications in many-body physics and quantum chemistry and are one of the most common algorithmic applications of shadow tomography. The set of all Pauli operators is of exponential size in n, and perhaps not as practically relevant, but seems important to study nonetheless since it represents a general Pauli learning task. Surprising tomography algorithms are possible in this case, including an algorithm we describe that, for any constant error ϵ, compresses the output (of size 4^n) into a polynomial-sized description from which an ϵ-estimate of the expected value of any Pauli observable can be extracted efficiently.

This brings us to the main question that guided this chapter: *Do there exist triply efficient shadow tomography protocols for these observables?*

In fact, the schemes based on random single-copy measurements described in Refs. [HKP20; CW20; BBO20; Jia+20a; EHF19] already achieve triply efficient shadow tomography for the set $\mathcal{P}_k^{(n)}$ of k-local Pauli operators with $k = O(1)$. In this chapter we present triply efficient shadow tomography algorithms for the set $\mathcal{F}_k^{(n)}$ of k-body fermionic operators for $k = O(1)$. Furthermore, our algorithms only use Clifford measurements on two copies of ρ at a time.

We will see that it is impossible to perform sample-efficient shadow tomography using only single-copy measurements for the set of k-body fermionic operators (Theorem 4.2). Taken together, our protocols and the single-copy lower bounds demonstrate that two-copy measurements are necessary and sufficient for Pauli and fermionic shadow tomography. Further, we see a striking difference between local Paulis, for which single-copy measurements suffice, and local fermionic operators where entangled measurements are necessary.

Local observables

In order to describe our results, let us now define the sets of local observables that are relevant to qubit and fermionic systems. Let $|P|$ denote the Pauli-weight of an operator $P \in \mathcal{P}^{(n)}$, i.e., the number of qubits on which it acts nontrivially. For example, $|X \otimes \mathbb{1} \otimes Y \otimes \mathbb{1} \otimes Z| = 3$.

A broad class of quantum many-body systems that arise in condensed matter physics are described by systems of spins with $k = O(1)$ particle interactions. Such systems are described by a Hamiltonian operator which can be expressed as a sum of operators from the set

$$\mathcal{P}_k^{(n)} = \{P \in \mathcal{P}^{(n)} : |P| = k\} \qquad \textbf{(k-local Pauli operators)} \tag{4.1}$$

of all weight-k Pauli observables. Note that $|\mathcal{P}_k^{(n)}| = 3^k \binom{n}{k}$ and $\log |\mathcal{P}_k^{(n)}| = O(k \log n)$. Learning all Paulis in the set $\mathcal{P}_k^{(n)}$ is quite useful—it allows one to reconstruct all the k-qubit reduced density matrices of the state ρ and compute for example the expected value of any k-local Hamiltonian operator. So shadow tomography with the set $\mathcal{P}_k^{(n)}$ is particularly relevant for characterizing ground states of quantum spin systems with few-body interactions.

A different subset of Pauli operators describes few-body interactions between fermionic particles, such as the electronic structure of molecules. Just as before, there is a fermionic locality parameter k but it is fundamentally different from the one defined by Pauli weight. To describe it, one fixes any subset of $2n$ anticommuting

n-qubit Pauli operators:

$$\gamma_1, \gamma_2, \ldots, \gamma_{2n} \in \mathcal{P}^{(n)} \qquad \forall a, b : \gamma_a \gamma_b + \gamma_b \gamma_a = 2\delta_{ab}\mathbb{1}, \tag{4.2}$$

where δ_{ab} is 1 if $a = b$ and 0 otherwise. Note that since $\gamma_a \in \mathcal{P}^{(n)}$ we also have $\gamma_a^\dagger = \gamma_a$ for each $1 \leq i \leq 2n$. These are known as the Majorana fermion operators associated with a fermionic system with n modes (a fermionic mode is a state that can either be occupied or unoccupied by a fermionic particle). Note that there is a freedom here—a particular choice of operators in Equation (4.2) is a fermion-to-qubit mapping that describes how we associate the degrees of freedom of the n-mode fermionic system with those of the n-qubit Hilbert space. Such mappings have a long history and there are several choices that are used in practice to design algorithms for fermionic systems on a quantum computer (see, e.g., Refs [JW28; BK02; SRL12; Jia+20a; Der+21], and Section 4.3). However, our discussion and the results described below apply to any choice of fermion-to-qubit mapping.

With our Majoranas (Equation (4.2)) in hand, let us now define the Majorana monomials as

$$\Gamma(x) = i^{|x|\cdot(|x|-1)/2} \gamma_1^{x_1} \gamma_2^{x_2} \ldots \gamma_{2n}^{x_{2n}} \qquad \forall x \in \{0,1\}^{2n}. \tag{4.3}$$

The overall phase factor $i^{|x|\cdot(|x|-1)/2}$ ensures that $\Gamma(x)$ is Hermitian for all $x \in \{0,1\}^{2n}$. In fact, the 4^n operators

$$\{\Gamma(x) : x \in \{0,1\}^{2n}\} \tag{4.4}$$

coincide with the 4^n n-qubit Pauli operators in $\mathcal{P}^{(n)}$, up to (efficiently computable) signs. Now let us define the k-body fermionic operators

$$\mathcal{F}_k^{(n)} = \{\Gamma(x) : |x| = 2k\}, \qquad \textbf{(k-body fermionic operators)} \tag{4.5}$$

which should be compared with Equation (4.1). Note that $|\mathcal{F}_k^{(n)}| = \binom{2n}{2k}$ and $\log |\mathcal{F}_k^{(n)}| = O(k \log n)$. A system of n fermionic modes with k-particle interactions is described by a Hamiltonian operator that is a sum of terms from $\mathcal{F}_k^{(n)}$. Note that $\mathcal{F}_k^{(n)}$ consists of the Majorana monomials in Equation (4.4) of degree $2k$; typically only these even-degree monomials are relevant to physics and chemistry due to conservation of fermionic parity.

The notion of k-locality for fermions is fundamentally more expressive than that of of k-locality for spin systems; the former subsumes the latter in the sense that

compact qubit-to-fermion mappings exist which embed the k-local n-qubit Pauli operators within a subset of the k-body fermionic operators on $O(n)$ fermionic modes (see for example Ref. [Bra+19]), whereas fermion-to-qubit mappings necessarily represent the Majorana fermion operators Equation (4.2) using Pauli operators of average weight at least $\Omega(\log(n))$ [Jia+20a]. The expectation values of the fermionic operators $\mathcal{F}_k^{(n)}$ comprise the matrix elements of what physicists and chemists refer to as the k-body reduced density matrix, or k-RDM.

Single-copy measurements.

It would be very practical if we could achieve triply efficient Pauli shadow tomography using only measurements on one copy of ρ at a time—and for k-local Paulis, it can be done. The shadow tomography task for $\mathcal{P}_k^{(n)}$ can be performed, using a time-efficient algorithm, using only *single-copy* measurements on $O(3^k(\log|\mathcal{P}_k^{(n)}|)/\epsilon^2) = O(3^k(k\log n)/\epsilon^2)$ copies of ρ [HKP20; CW20; BBO20; Jia+20a; EHF19].

One might hope that we can similarly achieve triply efficient Pauli shadow tomography for $\mathcal{F}_k^{(n)}$ and $\mathcal{P}^{(n)}$ as well. Unfortunately, this is not possible in either case, even if we only care about sample efficiency and allow unbounded computation time. It was shown in Ref [Che+22] that sample-efficient shadow tomography with single-copy measurements is impossible for the set of all Paulis.

Theorem 4.1 ([Che+22])**.** *There is no sample-efficient shadow tomography protocol with single-copy measurements for the set $\mathcal{S} = \mathcal{P}^{(n)}$ of all Paulis. In particular, any protocol based on single-copy measurements must consume $\Omega(2^n/\varepsilon^2)$ copies of ρ.*

For k-body fermionic operators, one can also establish a lower bound, see Section 4.1.[2]

Theorem 4.2. *There is no sample-efficient single-copy shadow tomography protocol for the set $\mathcal{F}_k^{(n)}$ of k-body fermionic operators. In particular, for $k = O(1)$, any protocol based on single-copy measurements must consume $\Omega(n^k/\varepsilon^2)$ copies of ρ.*

Several efficient shadow tomography algorithms based on single-copy measurements are known which achieve the $\Omega(n^k/\epsilon^2)$ lower bound, up to a log factor [BBO20; ZRM21; Wan+22; Hug+22a] (see also Ref. [Low22] which describes how this

[2] Theorem 1 of [BBO20] contains a lower bound for single-copy non-adaptive Clifford measurements; on the other hand, Theorem 4.2 applies to arbitrary and even adaptive single-copy measurements.

scaling can be improved if the particle number is asymptotically smaller than the number of fermionic modes).

The lower bound in Theorem 4.2 demonstrates a significant difference between learning local fermionic observables and local Pauli observables; for local Paulis there are sample-efficient single-copy protocols, whereas this is impossible for local fermionic observables.

In order to achieve triply efficient shadow tomography for $\mathcal{F}_k^{(n)}$ and $\mathcal{P}^{(n)}$, we will need to measure two or more copies of ρ at a time. Before jumping into two-copy measurements, we review the single-copy algorithm for k-local Pauli operators and offer a new interpretation in terms of fractional graph colorings that will be used in our algorithms.

The classical shadows single-copy protocol for k-local Paulis is very simple: it is based on measuring each qubit of ρ (in each copy of ρ) in a random single-qubit Pauli basis X, Y, or Z uniformly at random [HKP20]. The postprocessing of the measurement data to compute expected values is equally simple.

The high-level format of this protocol is as follows: one selects a Clifford basis at random according to some probability distribution p, and then measures in that basis. The distribution has the property that each of the Paulis P in the set $\mathcal{S}$ of observables of interest (in the above, $\mathcal{S} = \mathcal{P}_k^{(n)}$) has a high chance (at least 3^{-k}) of being diagonal in a basis sampled from p. Every time we pick a basis in which a Pauli is diagonal, we learn some information about its expected value and hence the sample complexity of the protocol is inversely related to the probability of being diagonal in a randomly sampled basis.

We reinterpret this measurement strategy, and other protocols based on random single-copy Clifford measurements, as arising from fractional colorings of the *commutation graph* of the observables $\mathcal{S}$, defined as follows:

Definition 4.3. *The commutation graph $G(\mathcal{S})$ of a set $\mathcal{S} \subseteq \mathcal{P}^{(n)}$ of Pauli operators is the graph with vertex set $\mathcal{S}$ and an edge between every pair of anticommuting operators.*

An independent set in $G(\mathcal{S})$ corresponds to a set of commuting observables that can be measured simultaneously via a Clifford measurement. Similarly, a coloring of this graph with χ colors describes a learning strategy with deterministic single-copy Clifford measurements, based on measuring χ disjoint sets of commuting

Pauli observables. Such deterministic graph coloring strategies for learning Pauli observables have been explored previously, see for example Ref. [JGM19; VYI20]. But the protocol for local Pauli observables described above is not based on a coloring of the commutation graph $G(\mathcal{S})$: the measurement bases are associated with *overlapping* sets of commuting Pauli observables, and are chosen randomly rather than deterministically. As we will see, a probabilistic Clifford measurement strategy can be viewed as defining a *fractional coloring* of $G(\mathcal{S})$. A fractional coloring is a well-studied relaxation of the notion of graph coloring; see Section 4.1 for details. We show that the sample complexity of single-copy learning with Clifford measurements is upper bounded by the fractional chromatic number of $G(\mathcal{S})$, which is the size of the smallest fractional coloring of $G(\mathcal{S})$. In Section 4.1, we prove the following theorem.

Theorem 4.3. *Let $\mathcal{S} \subseteq \mathcal{P}^{(n)}$. Suppose the commutation graph $G(\mathcal{S})$ admits a fractional coloring of size χ that can be sampled by a classical randomized algorithm with runtime T. Then there is an algorithm using only single-copy Clifford measurements of ρ which can estimate $\mathrm{Tr}(P\rho)$ within error ϵ for all $P \in \mathcal{S}$ with high probability using*

$$O(\chi(\log|\mathcal{S}|)/\epsilon^2) \tag{4.6}$$

copies of ρ. The runtime of the algorithm is $O((T + n^3) \cdot \chi(\log|\mathcal{S}|)/\epsilon^2)$.

The single-copy measurement strategies based on fractional colorings from Theorem 4.3 have the special feature that they only use Clifford measurements, which have efficient classical descriptions. Such protocols learn a *compressed classical representation* of ρ that can be used to compute Pauli observables; see Section 4.1 for details.

It is worth noting that the fractional chromatic number of any graph with N vertices is at most a factor of $1+\log(N)$ smaller than its chromatic number [Lov75], so we could in principle work exclusively with standard colorings and only lose this log factor. While true, this neglects the fact that protocols are often more naturally phrased using fractional coloring, just as classical algorithms are often more naturally phrased using randomness, even if they can ultimately be derandomized.

The power of two copies.

In light of Theorem 4.1 and Theorem 4.2 we see that there exist sets of Pauli observables for which sample-efficient shadow tomography cannot be achieved with one-copy measurements. Are two-copy measurements enough?

Our starting point here is a new algorithm that shows that sample-efficient shadow tomography is indeed possible in the general case with two-copy measurements. The protocol has three steps.

The first step is the learning magnitudes subroutine from Ref. [HKP21], which we now review. In this step we perform Bell sampling to measure copies of $\rho \otimes \rho$ in the Clifford basis that diagonalizes the commuting Pauli observables $P \otimes P$ for all $P \in \mathcal{P}^{(n)}$. Since these operators commute, we can learn the all observables $\mathrm{Tr}((P \otimes P)(\rho \otimes \rho)) = \mathrm{Tr}(\rho P)^2$ for $P \in \mathcal{S}$ to error δ using $O((\log |\mathcal{S}|)/\delta^2)$ measurements. By choosing $\delta = \Theta(\epsilon^2)$, we see that $O((\log |\mathcal{S}|)/\epsilon^4)$ two-copy measurements suffices to compute estimates $\{u_P\}_{P \in \mathcal{S}}$ such that

$$|u_P - |\mathrm{Tr}(\rho P)|| \leq \epsilon/4 \qquad \text{for all} \quad P \in \mathcal{S}, \tag{4.7}$$

with high probability. After we have learned the magnitudes in this way, we may find that some of our estimates are negligible; if our estimate u_P from the first stage is less than $3\epsilon/4$ then 0 is an ϵ-approximation to the expected value $\mathrm{Tr}(\rho P)$ and we can forget about this Pauli P going forward. So in the second stage of the algorithm we are only concerned with observables in the set

$$\mathcal{S}_\epsilon = \{P \in \mathcal{S} : |u_P| \geq 3\epsilon/4\}. \tag{4.8}$$

This set $\mathcal{S}_\epsilon$ is a random variable determined by the output of Bell sampling, but the condition in Equation (4.7) implies that with high probability we have

$$|\mathrm{Tr}(\rho P)| \geq \epsilon/2 \quad \text{for all} \quad P \in \mathcal{S}_\epsilon. \tag{4.9}$$

To complete the learning task it suffices to then compute the sign of $\mathrm{Tr}(\rho P)$ for all Paulis $P \in \mathcal{S}_\epsilon$.

Improved algorithms using graph theory.

In order to resolve the signs, we employ a general framework for two-copy shadow tomography which is based on fractional graph coloring. Our framework can be viewed as an extension of a heuristic learning algorithm proposed in Appendix E.2.d of Ref. [HKP21]; in this work we use it to obtain algorithms with rigorous performance guarantees.

We have already seen that single-copy tomography for any set of observables $\mathcal{S}$ reduces to fractional graph coloring for the commutation graph $G(\mathcal{S})$. Likewise, via Bell sampling, two-copy tomography reduces to fractional graph coloring for

the commutation graph $G(\mathcal{S}_\epsilon)$. That is, we propose to use the single-copy algorithm to learn all observables in $\mathcal{S}_\epsilon$, once we have already determined $\mathcal{S}_\epsilon$ using an initial stage of Bell sampling. But how do the two-copy measurements help us?

A key insight is that the Paulis in $\mathcal{S}_\epsilon$ cannot be very anticommuting. Intuition from the Heisenberg uncertainty principle tells us that anticommuting (traceless) observables cannot simultaneously be large on a quantum state, since the quantum state cannot simultaneously be an eigenvector of anticommuting observables. But the high probability event in Equation (4.9) implies that the Paulis in $\mathcal{S}_\epsilon$ are simultaneously large on the state ρ, and thus cannot anticommute with each other too often. This can be formalized as an upper bound on the clique number of their commutation graph. In Section 4.1 we show the following:

Lemma 4.4. *The largest clique in the commutation graph $G(\mathcal{S}_\epsilon)$ has size at most $4/\epsilon^2$ with high probability.*

Going forward, our aim is to exploit this upper bound on the clique number to compute good (fractional) colorings.

Unfortunately, it is well known that the chromatic (or fractional chromatic) number is not upper bounded by any function of the clique number in general [3]. However, such upper bounds can be established for certain families of graphs, a research direction pioneered by Gyárfás [Gyá87]. A family of graphs for which this is possible is called *chi-bounded* and the upper bound on chromatic number is said to be expressed in terms of a chi-binding function (see Refs. [SR19; SS20] for recent surveys). Our shadow tomography learning task for a set of observables $\mathcal{S}$ thus reduces to establishing a suitable chi-binding function for the family of induced subgraphs of the commutation graph $G(\mathcal{S})$; see Section 4.1 for details.

We show that the family of induced subgraphs of the commutation graph of k-body fermionic observables admits a polynomial chi-binding function (that does not depend on n).

Lemma 4.5. *Let $k \geq 1$, and let G' be any induced subgraph of the commutation graph $G(\mathcal{F}_k^{(n)})$ of k-body fermionic observables, and let ω be the size of the largest clique in G'. Then the fractional chromatic number of G' satisfies*

$$\chi_f(G') \leq p_k(\omega), \tag{4.10}$$

[3] For example, there exists a family of triangle-free graphs with chromatic number $\Omega(\sqrt{m/\log m})$ where m is the number of vertices [Kim95].

where p_k is a polynomial. Moreover, for any $k = O(1)$ we can sample from a fractional coloring of G' with size $p_k(\omega)$ using a classical algorithm with runtime $\mathrm{poly}(n)$. *The polynomials for $k = 1, 2$ are $p_1(\omega) = \omega + 1$ and $p_2(\omega) = O(\omega^8)$.*

The proof of Lemma 4.5 is provided in Section 4.3. As discussed above, the commutation graph $G(\mathcal{S}_\epsilon)$ is an induced subgraph of $G(\mathcal{S})$ with clique number at most $O(1/\epsilon^2)$. For k-body fermionic observables $\mathcal{S} = \mathcal{F}_k^{(n)}$, Lemma 4.5 tells us there is an efficiently computable fractional coloring of $G(\mathcal{S}_\epsilon)$ with at most $\mathrm{poly}(1/\epsilon^2)$ colors. We can then use the single-copy learning protocol from Theorem 4.3 to learn all observables in $\mathcal{S}_\epsilon$. This reduction, which describes how to convert Lemma 4.5 into a two-copy learning protocol, is formalized in Lemma 4.13. Putting it all together gives the following theorem.

Theorem 4.6. *Let $k = O(1)$. There exists a triply efficient shadow tomography protocol for the set $\mathcal{S} = \mathcal{F}_k^{(n)}$ of k-body fermionic observables that uses only two-copy Clifford measurements.*

This triply efficient protocol has sample complexity

$$O\left(\frac{\log|\mathcal{F}_k^{(n)}|}{\epsilon^4} + \frac{p_k(4/\epsilon^2)\log|\mathcal{F}_k^{(n)}|}{\epsilon^2}\right) = O((k\log n)p_k(4/\epsilon^2)/\epsilon^2), \tag{4.11}$$

where p_k is the polynomial from Lemma 4.5 that depends on the locality k, and we also used the fact that $p_k(\omega) = \Omega(\omega)$ for all $k \geq 1$. For each $k \geq 1$ we obtain an exponential improvement over single-copy learning protocols in terms of the sample complexity as a function of system size n. For $k = 1$ our learning algorithm has sample complexity $O((\log n)/\epsilon^4)$, and the measurements and postprocesssing are simple to implement. We anticipate that this learning algorithm could find applications in quantum simulations of chemistry and fermionic physics. With our current analysis, the degree of the polynomial p_k increases very rapidly as a function of k rendering the scheme less practical for $k \geq 2$. We hope this could be improved in future work. An upper bound on the ϵ-dependence of the sample complexity is $\sim \epsilon^{-O((2k)^{k+1})}$. For $k = 2, 3$ the sample complexity is $\sim \epsilon^{-18}$ and $\sim \epsilon^{-110}$ respectively.

It is natural to ask how far we can push this two-copy framework based on Bell sampling and fractional coloring. Below, we show that it provides a nontrivial shadow tomography protocol for *any* subset of Pauli observables $\mathcal{S} \subseteq \mathcal{P}^{(n)}$. This

gives hope that our framework could lead to triply efficient shadow tomography in the general case.

We shall exploit the fact that the longest induced path in the commutation graph $G(\mathcal{P}^{(n)})$ contains at most $2n + 1$ vertices; see Section 4.2 for details. The following upper bound on chromatic number then follows from a seminal result in chi-boundedness due to Gyárfás [Gyá87]; see Section 4.2.

Lemma 4.7. *Let G' be any induced subgraph of the commutation graph $G(\mathcal{P}^{(n)})$, and let ω be the size of the largest clique in G'. The chromatic number of G' is upper bounded as*

$$\chi(G') \leq (2n+1)^{\omega-1}. \tag{4.12}$$

Moreover, a coloring with this many colors can be computed by a classical algorithm with runtime poly$(|G'|, n^{\omega})$.

To get a shadow tomography algorithm for any set of Pauli observables $\mathcal{S} \subseteq P^{(n)}$, we follow the strategy outlined above and formalized in Lemma 4.13. That is, we apply Lemma 4.7 to the subgraph $G' = G(\mathcal{S}_\epsilon)$ induced by the set $\mathcal{S}_\epsilon$ computed via Bell sampling. From Lemma 4.4 we have that with high probability the largest clique in G' has size $\omega = O(1/\epsilon^2)$. So we get an coloring of $G(\mathcal{S}_\epsilon)$ with at most $n^{O(1/\epsilon^2)}$ colors, that can be computed with runtime poly$(|\mathcal{S}|, n^{1/\epsilon^2})$. When $\epsilon = \Omega(1)$ is a small constant, this protocol is time-efficient, has sample complexity poly(n), and only uses two-copy measurements, for any subset of Pauli observables $\mathcal{S}$. This gives a sample-efficient protocol only when $|\mathcal{S}|$ is exponentially large as a function of n. At a technical level this is a consequence of the factor of n appearing in Equation (4.12), and we do not know if this can be avoided.

This leaves open the question of triply efficient shadow tomography for arbitrary subsets of Pauli observables. However, we will see that it provides insight into a related question concerning compressed classical representations of quantum states.

Rapid-retrieval Pauli compression.

Can we compress an n-qubit quantum state into a small amount of classical information, so that the compressed classical description is sufficient to extract the expectation values of any bounded observable to within a small constant error? This question has been studied using tools from communication complexity and it is known that an exponential classical description size is necessary if one wishes to

recover bounded observables in the general case; a representation size $\tilde{\Theta}(\sqrt{2^n})$ is necessary and sufficient for n-qubit pure states [Raz99; Gav+07; GS18].

On the other hand, if we restrict our attention to the set of n-qubit Pauli observables (or other sets of observables with only singly exponential size), a classical description of size $\mathrm{poly}(n)$ exists and can be computed using the matrix multiplicative weights algorithm [Aar04; Aar+18]. However, a significant drawback of known methods for this task is that they require exponential classical runtime to extract the expected value of a given Pauli observable from the compressed classical representation.

A consequence of Lemma 4.7 is that this exponential cost can be avoided, at least for any small constant precision $\epsilon = \Omega(1)$. That is, one can compress an n-qubit state ρ into $\mathrm{poly}(n)$ classical bits. Given this classical data and an n-qubit Pauli P, there is an *efficient* classical algorithm to estimate $\mathrm{Tr}(\rho P)$ to within ϵ-error. Moreover, such a representation can be learned from $\mathrm{poly}(n)$ samples of ρ.

Corollary 4.8 (Rapid-retrieval Pauli compression)**.** *Let ρ be an n-qubit quantum state. Let $\epsilon \in (0, 1)$ be a constant independent of n. Using two-copy Clifford measurements on* $\mathrm{poly}(n)$ *copies of ρ, along with $2^{O(n)}$ runtime, we can (with high probability) learn a compressed classical representation of ρ, call it $D(\rho, \epsilon)$, that consists of* $\mathrm{poly}(n)$ *bits. An ϵ-approximation to the expected value* $\mathrm{Tr}(\rho P)$ *of any Pauli observable $P \in \mathcal{P}^{(n)}$ can be extracted from $D(\rho, \epsilon)$ using a classical algorithm with* $\mathrm{poly}(n)$ *runtime.*

The classical description $D(\rho, \epsilon)$ consists of a list of all the Clifford measurement bases and measurement outcomes used in the two-copy learning algorithm discussed above; see Section 4.1 for details. In particular, Corollary 4.8 is obtained by combining Lemma 4.7 and Lemma 4.14.

Discussion and open questions

In this chapter we have provided the first triply efficient shadow tomography protocols for the set of k-body fermionic observables. We have also provided a route to strengthening and generalizing our results via a connection between two-copy tomography and graph theory techniques related to chi-boundedness.

There are many questions left open by our work. Is it possible to improve the upper bounds from Equation (4.10) and Equation (4.12)—e.g., can we establish better chi-binding functions for the (families of) commutation graphs of interest? Is rapid-retrieval compression possible for smaller error parameters, e.g., $\epsilon = 1/\mathrm{poly}(n)$?

Can we devise triply efficient learning algorithms for any subset of Pauli observables?

One route towards resolving these questions would be via improved algorithms for coloring the commutation graph $G(\mathcal{S}_\epsilon)$. The following conjecture asserts an efficient fractional coloring of the commutation graph of any subset of Paulis that has simultaneously large expected values in a quantum state.

Conjecture 4.9. *Let ρ be an n-qubit state, $\delta \in (0,1)$, and let $B \subseteq \mathcal{P}^{(n)}$ be the set of all Paulis P such that $|\operatorname{Tr}(\rho P)| \geq \delta$. There is a fractional coloring of the commutation graph $G(B)$ of size $O(1/\delta^2)$.*

If this conjecture holds, and in addition the fractional coloring is suitably efficient,[4] then we would obtain a triply efficient Pauli shadow tomography algorithm for *any* subset $\mathcal{S}$ of Pauli observables. Moreover, the learning algorithm would also output a rapid-retrieval Pauli compression of all observables in $\mathcal{S}$, of size $O(n^2(\log|\mathcal{S}|)/\epsilon^4)$; see Section 4.1.

To address Conjecture 4.9, it is natural to ask if Lemma 4.4 can be strengthened by showing that $O(1/\epsilon^2)$ is in fact an upper bound on the *fractional clique number*—a well-known linear programming relaxation of the clique number [SU11]. The existence of a suitable fractional coloring stated in Conjecture 4.9 would then follow from linear programming duality, which asserts that the fractional clique number of any graph equals its fractional chromatic number.

4.1 Commutation, learning, and coloring

In this section we describe properties of the commutation graph $G(\mathcal{S})$ of a set of Pauli observables $\mathcal{S} \subseteq \mathcal{P}^{(n)}$, and the connection between these properties and shadow tomography algorithms.

Lower bounds from commutation index

We begin by discussing lower bounds on learning arising from anticommutativity of a set of observables. This is quantified by the *commutation index* introduced in Section 2.1. Ref. [Che+22] shows that the inverse of the commutation index is a lower bound on the sample complexity of single-copy shadow tomography for $\mathcal{S}$[5]. It

[4] In particular, we require that a fractional coloring of size $O(1/\delta^2)$ for any subset $R \subseteq B$ can be sampled in time $\operatorname{poly}(|R|, n, 1/\delta)$.

[5] Maximization over states as stated in [Che+22, Equation 79] can be replaced by maximization over density matrices via convexity.

can be interpreted as describing a tension between anticommutation and learnability.

Theorem 4.10 (Theorem 5.5, [Che+22])**.** *Shadow tomography to precision ϵ for a set $\mathcal{S}$ of Pauli observables with single-copy measurements requires at least*

$$\Omega\left(\frac{1}{\epsilon^2 \Delta(\mathcal{S})}\right) \tag{4.13}$$

copies of ρ. This holds even for adaptive measurement strategies.[6]

Using results on the commutation index from Section 2.1 we get the following lower bounds on the sample complexity of learning local operators with single-copy measurements.

Theorem 4.11. *Let $1 \leq k \leq \log_3(2n+1)$. Any (possibly adaptive) single-copy protocol which learns $\mathrm{Tr}(P\rho)$ to precision ϵ for all k-local n-qubit Paulis $P \in \mathcal{P}_k^{(n)}$ with constant probability requires $\Omega(3^k/\epsilon^2)$ copies of ρ.*

Proof. Follows from Theorem 2.9 and Theorem 4.10. □

The lower bound in Theorem 4.11 matches the sample complexity of known single-copy protocols such as classical shadows, up to a factor of $k \log n$ [HKP20].

For local fermionic observables, one can obtain a much stronger single-copy sample complexity lower bound which scales polynomially in the system size n. This is telling us that local fermionic observables are much harder to learn than local Pauli observables. The following theorem implies that there is no sample efficient single-copy protocol for k-body fermionic observables, as stated in Theorem 4.2.

Theorem 4.12. *Any (possibly adaptive) single-copy protocol which learns $\mathrm{Tr}(\Gamma\rho)$ to precision ϵ for all k-body Majorana operators Γ on n fermionic modes with constant probability requires number of copies scaling as $\Omega(n^k/\epsilon^2)$, for any fixed $k \geq 1$.*

Proof. Follows from Theorem 2.10 and Theorem 4.10. □

The $\Omega(n^k/\epsilon^2)$ single-copy sample complexity lower bound in Theorem 4.12 matches what is achieved by the single-copy protocols in [BBO20; Jia+20a; Wan+22; Hug+22a], up to a factor of $k \log(n)$.

[6] The lower bound holds even if the shadow tomography scheme is only able to output the absolute values of the expectation values to precision ϵ.

Fractional coloring and single-copy Clifford learning

Here we describe the connection between shadow tomography algorithms that learn Pauli observables $\mathcal{S} \subseteq \mathcal{P}^{(n)}$ using probabilistic Clifford measurements, and fractional colorings of the commutation graph $G(\mathcal{S})$. A fractional coloring is a relaxation of the usual notion of graph coloring [SU11].

Definition 4.4. *Let $G = (V, E)$ be a graph. A fractional coloring of G of size χ is a probability distribution q over independent sets $I \subseteq V$ with the property that*

$$\forall v \in V : \quad \Pr_{I\sim q}(v \in I) \geq 1/\chi. \tag{4.14}$$

The fractional chromatic number $\chi_f(G)$ of G is the size of the smallest fractional coloring of G.

Note that the size χ of a fractional coloring need not be an integer. Also note that a (standard, non-fractional) coloring of G with χ colors can be regarded as a fractional coloring of size χ, corresponding to a uniform distribution over color classes, so the fractional chromatic number of a graph is upper bounded by its chromatic number.

Theorem 4.3, restated below, asserts that if we have a fractional coloring of the commutation graph $G(\mathcal{S})$ of small size, then we can learn the expectation values of all observables in $\mathcal{S}$ with few single-copy Clifford measurements. In particular, the sample complexity of the algorithm scales linearly with the size of the fractional coloring.

In the following, samples from the fractional coloring are represented as binary vectors of length $|\mathcal{S}|$ whose support is an independent set in $G(\mathcal{S})$. In this setting, the runtime to produce a single sample from a fractional coloring always satisfies $T \geq |\mathcal{S}|$.

Theorem 4.3. *Let $\mathcal{S} \subseteq \mathcal{P}^{(n)}$. Suppose the commutation graph $G(\mathcal{S})$ admits a fractional coloring of size χ that can be sampled by a classical randomized algorithm with runtime T. Then there is an algorithm using only single-copy Clifford measurements of ρ which can estimate $\mathrm{Tr}(P\rho)$ within error ϵ for all $P \in \mathcal{S}$ with high probability using*

$$O(\chi(\log|\mathcal{S}|)/\epsilon^2) \tag{4.6}$$

copies of ρ. The runtime of the algorithm is $O((T + n^3) \cdot \chi(\log|\mathcal{S}|)/\epsilon^2)$.

Proof. An independent set I in the commutation graph $G(\mathcal{S})$ consists of a set of mutually commuting Pauli operators, which can be simultaneously measured

by applying a Clifford circuit and measuring in the computational basis. Such a Clifford circuit can be computed using a classical algorithm with $O(n^3)$ runtime, via standard techniques in the stabilizer formalism [AG04].

If we draw an independent set I from a fractional coloring q with size χ and then measure ρ in the corresponding Clifford basis, the result gives us a measurement of Pauli $P \in \mathcal{S}$ whenever $P \in I$. Note that it is also possible that we get more useful measurements than this—the Clifford unitary may diagonalize some Paulis P that are not in I.

Suppose we repeat this process independently N times, sampling independent sets $I_1, I_2, \ldots, I_N$ and measuring N independent identical copies of ρ in the corresponding Clifford bases $C_1, C_2, \ldots, C_N$. For each $P \in \mathcal{S}$, let $x_P^j \in \{-1, 0, 1\}$ be the random variable that is equal to the measured outcome of P if it is diagonalized by C_j, and zero otherwise.

We have

$$\Pr[x_P^j \in \{-1, 1\}] \geq \Pr[P \in I_j] \geq \frac{1}{\chi} \qquad j \in [N] \qquad P \in \mathcal{S}, \tag{4.15}$$

where we used the fact that q is a fractional coloring of size χ.

For each Pauli $P \in \mathcal{S}$, let

$$R_P = \{j : x_P^j \in \{-1, 1\}\} \quad \text{and} \quad N_P = |R_P|. \tag{4.16}$$

Consider the sample mean

$$\tilde{P} = \frac{1}{N_P} \sum_{j \in R_P} x_P^j. \tag{4.17}$$

Conditioned on a fixed value $N_P \geq 1$, this sample mean $\tilde{P}$ is an average of N_P independent ± 1-valued random variables. It satisfies

$$\mathbb{E}(\tilde{P}) = \mathrm{Tr}(\rho P) \qquad \text{and} \qquad \mathrm{Var}(\tilde{P}) \leq \frac{1}{N_P}. \tag{4.18}$$

By Chebyshev's inequality we have

$$\Pr\left[|\tilde{P} - \mathrm{Tr}(\rho P)| \geq \epsilon \;\middle|\; N_P \geq 100/\epsilon^2\right] \leq 0.01. \tag{4.19}$$

From Equation (4.15) we see that by taking $N = O(\chi/\epsilon^2)$ we can ensure that, for a given Pauli P, we have $N_P \geq 100/\epsilon^2$ with probability at least 0.99 (say). Therefore,

$$\Pr\left[|\tilde{P} - \mathrm{Tr}(\rho P)| \leq \epsilon\right] \geq \Pr\left[N_P \geq 100/\epsilon^2\right] \cdot \Pr\left[|\tilde{P} - \mathrm{Tr}(\rho P)| \leq \epsilon \;\middle|\; N_P \geq 100/\epsilon^2\right] \tag{4.20}$$

$$\geq 0.99^2 \tag{4.21}$$

for each Pauli $P \in \mathcal{S}$.

Now repeat the above process L times, generating sample means $\tilde{P}_1, \ldots, \tilde{P}_L$ for each $P \in \mathcal{S}$, and consider the median-of-means estimator

$$\lambda_P = \text{median}(\tilde{P}_1, \tilde{P}_2, \ldots, \tilde{P}_L). \tag{4.22}$$

By choosing $L = O(\log |\mathcal{S}|)$ we can ensure that, for each $P \in \mathcal{S}$ we have $|\lambda_P - \text{Tr}(\rho P)| \leq \epsilon$ with probability at least $0.01/|\mathcal{S}|$. By a union bound we get all the expected values in $\mathcal{S}$ to within ϵ with probability at least 0.99. The total number of samples of ρ and the total number of samples from the fractional coloring q used in the algorithm are both at most $LN = O(\chi(\log |\mathcal{S}|)/\epsilon^2)$.

Now consider the runtime of the protocol. The independent sets I in the fractional coloring are specified explicitly as subsets of $\mathcal{S}$, so the median-of-means estimator λ_P can be computed for all $P \in \mathcal{S}$ with a runtime $O(LN|\mathcal{S}|)$ once we have already obtained all the measurement data $\{x_P^j\}$. The total runtime is therefore upper bounded as $O((T + n^3 + |\mathcal{S}|)LN)$, where the first term is the cost of sampling the fractional colorings, the second term is the cost of computing a Clifford circuit for each sample (and applying this circuit to measure the state), and the third term is the cost of postprocessing. Since $T \geq |\mathcal{S}|$ the runtime simplifies to $O((T+n^3)LN)$. □

Finally, let us show that single-copy measurement strategies from Theorem 4.3 learn a *compressed classical representation* of ρ that encodes the expected values of all Pauli observables from $\mathcal{S}$ (to within error ϵ). Indeed, each Clifford measurement basis has an efficient classical description consisting of $O(n^2)$ bits. We can imagine a version of the learning algorithm described in Theorem 4.3 where, after the measurements are performed using Equation (4.6) copies of ρ, the resulting measurement outcomes and measurement bases are packaged up into a classical description $D(\rho, \mathcal{S}, \epsilon)$ of size

$$O(n^2 \chi(\log |\mathcal{S}|)/\epsilon^2). \tag{4.23}$$

Here there is a factor of n^2 for each measurement basis (each measurement outcome only requires n bits to describe and so describing the outcomes requires asymptotically fewer bits than describing the bases). The efficient protocol for extracting expected values $\text{Tr}(\rho P)$ with $P \in \mathcal{S}$ (up to ϵ error) can be performed using only the compressed classical description $D(\rho, \mathcal{S}, \epsilon)$.

Cliques and two-copy Clifford learning

As discussed in the introduction, we propose a framework for two-copy learning that uses an initial stage of Bell sampling (as in Ref. [HKP21]) to determine a set $\mathcal{S}_\epsilon \subseteq \mathcal{S}$ that with high probability satisfies Equation (4.9), which we restate:

$$|\operatorname{Tr}(\rho P)| \geq \epsilon/2 \quad \text{for all} \quad P \in \mathcal{S}_\epsilon. \tag{4.24}$$

This step uses $O((\log|\mathcal{S}|)/\epsilon^4)$ Clifford measurements on copies of $\rho \otimes \rho$. Then we aim to learn all observables in $\mathcal{S}_\epsilon$ using the fractional coloring approach described in the previous section. In particular, we aim to find an (efficiently sampleable) fractional coloring of the commutation graph $G(\mathcal{S}_\epsilon)$.

Lemma 4.13 (Template for two-copy Clifford shadow tomography)**.** *Suppose that $G(\mathcal{S}_\epsilon)$ admits a fractional coloring of size χ that can be sampled by a randomized algorithm with runtime T. Then there is an algorithm which performs shadow tomography for $\mathcal{S}$ using*

$$O((\log|\mathcal{S}|)/\epsilon^4 + \chi(\log|\mathcal{S}|)/\epsilon^2) \tag{4.25}$$

two-copy Clifford measurements and runtime $O(|\mathcal{S}|(\log|\mathcal{S}|)/\epsilon^4 + (T+n^3)\chi(\log|\mathcal{S}|)/\epsilon^2)$.

Proof. The first step uses $O((\log|\mathcal{S}|)/\epsilon^4)$ two-copy Bell measurements and classical runtime $O(|\mathcal{S}|(\log|\mathcal{S}|)/\epsilon^4)$ to compute the set $\mathcal{S}_\epsilon$. We output zero as our estimate for the expected value of any Pauli in $\mathcal{S} \setminus \mathcal{S}_\epsilon$. Then we use the single-copy learning protocol from Theorem 4.3 to compute estimates of $\operatorname{Tr}(\rho P)$ for all $P \in \mathcal{S}_\epsilon$. This second step uses runtime $O((T+n^3)\chi(\log|\mathcal{S}|)/\epsilon^2)$. □

The protocol from Lemma 4.13 is based on Clifford measurements, which have an efficient classical description. Because of this, and the fact that the classical post-processing is simple, we obtain the following classical compressed representation of ρ.

Lemma 4.14 (Rapid-retrieval compression)**.** *The shadow tomography protocol described in Lemma 4.13 learns a compressed classical representation of ρ consisting of B bits, where*

$$B = O(n(\log|\mathcal{S}|)/\epsilon^4 + n^2\chi(\log|\mathcal{S}|)/\epsilon^2). \tag{4.26}$$

With high probability, this compressed representation has the following rapid-retrieval property. There is a classical algorithm which, given this classical data

and any Pauli $P \in \mathcal{S}$, outputs an estimate of $\mathrm{Tr}(\rho P)$ *to within* ϵ *error. The runtime of the algorithm is* $O(B)$.

Proof. The compressed representation consists of the $O((\log|\mathcal{S}|)/\epsilon^4)$ Bell samples (each one is $2n$ bits) as well as a classical description of each Clifford measurement basis and measurement outcome used in the second stage of the learning protocol. That is, $N = O(\chi(\log|\mathcal{S}|)/\epsilon^2)$ Clifford measurement bases $C_1, C_2, \ldots, C_N$ (each described by a circuit with $O(n^2)$ one- and two-qubit Clifford gates) and corresponding measurement outcomes $z^1, z^2, \ldots, z^N \in \{0,1\}^n$. Given a Pauli $P \in \mathcal{S}$ and this classical data, we can compute an ϵ-error estimate of $\mathrm{Tr}(\rho P)$ in the following way. First, using the Bell samples, we determine if $P \in \mathcal{S}_\epsilon$. This step requires us to compute a sample mean over the Bell samples, using runtime $O(n(\log|\mathcal{S}|)/\epsilon^4)$. If $P \notin \mathcal{S}_\epsilon$, we output 0 as our estimate. If $P \in \mathcal{S}_\epsilon$ then we compute the median-of-means estimator from Equation (4.22). To do this we have to compute indicator functions $x_P^j \in \{1, 0, -1\}$ that describe the measured outcome of Pauli P for each Clifford measurement basis j, which is given by

$$x_P^j = \langle z^j | C_j^\dagger P C_j | z^j \rangle. \tag{4.27}$$

The RHS is computed using the stabilizer formalism: we update the Pauli P by conjugating each gate in the circuit C_j one-by-one, and this process takes a total runtime $O(n^2)$ since there are $O(n^2)$ one- and two-qubit Clifford gates in the circuit. The total runtime to extract the estimate of $\mathrm{Tr}(\rho P)$ is therefore

$$O(n(\log|\mathcal{S}|)/\epsilon^4 + n^2\chi(\log|\mathcal{S}|)/\epsilon^2). \tag{4.28}$$

□

Lemma 4.13 forms the basis of several of the two-copy protocols that we present in this work. To use this framework one needs to find an efficiently sampleable fractional coloring of $G(\mathcal{S}_\epsilon)$.

A challenge here is that the set $\mathcal{S}_\epsilon$ and its commutation graph depend in a potentially complicated way on the unknown state ρ. Ideally, we would like to understand any structural properties of this graph that can be leveraged to compute good fractional colorings. In this paper we we will only exploit two simple properties: (A) $G(\mathcal{S}_\epsilon)$ is an induced subgraph of $G(\mathcal{S})$ and (B) with high probability, $G(\mathcal{S}_\epsilon)$ does not have large cliques, as described in Lemma 4.4, which we restate and prove below.

Lemma 4.4. *The largest clique in the commutation graph* $G(\mathcal{S}_\epsilon)$ *has size at most* $4/\epsilon^2$ *with high probability.*

Proof. Recall that $\mathcal{S}_\epsilon$ satisfies Equation (4.24) with high probability. We show that in this case the largest clique in $G(\mathcal{S}_\epsilon)$ has size at most $4/\epsilon^2$.

Suppose there is a clique in $G(\mathcal{S}_\epsilon)$ of size ω. The vertices of the clique are a set of pairwise anticommuting Pauli operators $P_1, P_2, \ldots, P_\omega$. Applying Lemma 2.6 with these operators and the state ρ gives

$$\sum_{j=1}^{\omega} \mathrm{Tr}(P_j \rho)^2 \leq 1. \tag{4.29}$$

On the other hand from Equation (4.24) we have $\mathrm{Tr}(P_j\rho)^2 \geq \epsilon^2/4$ for each $1 \leq j \leq \omega$. Plugging into the above gives $\omega\epsilon^2/4 \leq 1$, and therefore the size of the maximal clique is upper bounded as $\omega \leq 4/\epsilon^2$. □

To use our framework to learn Pauli observables $\mathcal{S} \subseteq \mathcal{P}^{(n)}$, it suffices to establish a so-called *chi-binding function* for the family of induced subgraphs of $G(\mathcal{S})$. That is, we seek a function $g(\omega)$ such that for any induced subgraph G' of $G(\mathcal{S})$ with largest clique of size ω, there is a fractional coloring of G' with size $\chi \leq g(\omega)$. A statement of this form implies—via Lemma 4.13 and Lemma 4.4—a shadow tomography algorithm that uses two-copy Clifford measurements and has sample complexity

$$O((\log|\mathcal{S}|)/\epsilon^4 + g(4/\epsilon^2)(\log|\mathcal{S}|)/\epsilon^2). \tag{4.30}$$

Lemma 4.13 also gives an upper bound on the runtime of the protocol in terms of the time required to sample from the fractional coloring.

4.2 Coloring commutation graphs with bounded clique number

In this section we prove Lemma 4.7, restated below.

Lemma 4.7. *Let G' be any induced subgraph of the commutation graph $G(\mathcal{P}^{(n)})$, and let ω be the size of the largest clique in G'. The chromatic number of G' is upper bounded as*

$$\chi(G') \leq (2n+1)^{\omega-1}. \tag{4.12}$$

Moreover, a coloring with this many colors can be computed by a classical algorithm with runtime poly$(|G'|, n^\omega)$.

We shall use the following algorithmic version of a result of Gyárfás, which we prove below.

Lemma 4.15 (Algorithmic version of Thm. 2.4 of [Gyá87])**.** *Suppose G is a graph on m vertices whose longest induced path has ℓ vertices, with $\ell \geq 1$, and clique number ω. Then there is a classical algorithm which colors G using $\ell^{\omega-1}$ colors and runtime $O(m^2\omega)$.*

Proof of Lemma 4.7. Let G' be an induced subgraph of $G(\mathcal{P}^{(n)})$. Below we show that $G(\mathcal{P}^{(n)})$, and therefore also G', does not contain any induced paths with more than $2n+1$ vertices. The claim then follows by applying Lemma 4.15.

Suppose $P_1, P_2, \ldots, P_s$ is an induced path in $G(\mathcal{P}^{(n)})$, i.e.,

$$P_iP_{i+1} = -P_{i+1}P_i \text{ for } \quad 1 \leq i \leq s-1 \quad \text{and} \quad [P_i, P_j] = 0 \quad |i-j| \geq 2. \tag{4.31}$$

Define Pauli operators

$$Q_r = P_1P_2 \ldots P_r \qquad 1 \leq r \leq s. \tag{4.32}$$

We now use Equation (4.31) to show that these operators are pairwise anticommuting. To see this note that

$$Q_{r+a} = (P_1P_2 \ldots P_r)\,(P_{r+1}P_{r+2} \ldots P_{r+a}) = -\,(P_{r+1}P_{r+2} \ldots P_{r+a})\,(P_1P_2 \ldots P_r)\,, \tag{4.33}$$

where we used $P_rP_{r+1} = -P_{r+1}P_r$ and the fact that $[P_i, P_j] = 0$ whenever $i \leq r-1$ and $j \geq r+1$. Therefore

$$Q_rQ_{r+a} = -Q_r\,(P_{r+1}P_{r+2} \ldots P_{r+a})\,(P_1P_2 \ldots P_r) = -Q_{r+a}Q_r, \tag{4.34}$$

which shows that $\{Q_j\}_{j\in[s]}$ are pairwise anticommuting Pauli operators.

It is a well known fact that the n qubit Hilbert space does not contain any set of pairwise anticommuting Pauli operators with size greater than $2n+1$ (see for example Appendix G of Ref. [BBO20]). Thus $s \leq 2n+1$. □

The algorithm of Lemma 4.15 relies on a non-standard graph traversal algorithm, which can be interpreted as a combination of depth-first-search and breadth-first-search. The graph search algorithm begins with an arbitrary seed vertex v in G, and generates a spanning tree of G with root v. We call it *neighbour-first search*.

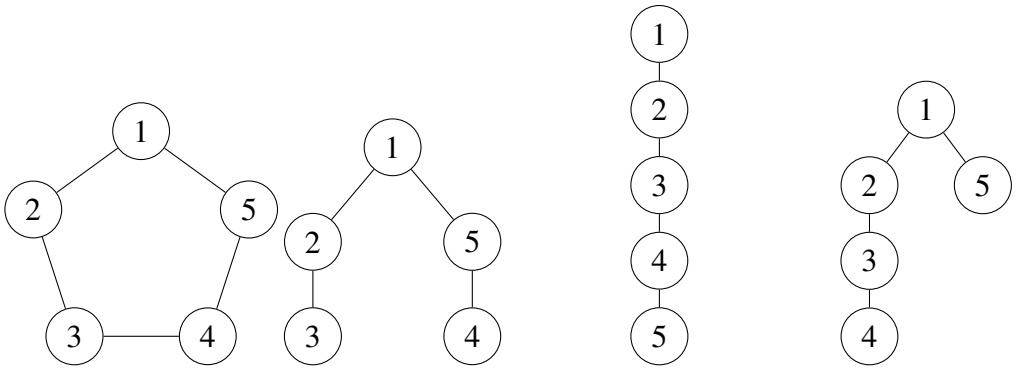

(a) Original graph (b) BFS spanning tree (c) DFS spanning tree (d) NFS spanning tree

Figure 4.1: An example showing the spanning trees generated by breadth-first search (BFS), depth-first search (DFS), and our algorithm neighbour-first search (NFS) for the cycle on 5 vertices.

Algorithm 1 Neighbour-first search (NFS)

Input: Connected graph G, seed vertex v.
Output: Spanning tree T of G with root v.

NFS(G, v):

1. If T is empty, initialize $T = \{v\}$.
2. For each neighbour w of v which is not yet in T, add w to T as a child of v.
3. For each child w of v:
    a) Do NFS(G, w).

The spanning tree T output by Algorithm 1 is associated with a partition of the vertex set of G into levels, which are the vertices at a fixed distance from the root of T. (The number of levels is the depth of the T plus one.)

Lemma 4.16. *Let graph $G = (V, E)$ have m vertices, clique number ω, and longest induced path with ℓ vertices. Algorithm 1 has runtime $O(|V| + |E|) = O(m^2)$ and outputs a spanning tree T of G with the following properties:*

- *The depth of T is no larger than $\ell - 1$.*
- *The vertices in any level of T induce a subgraph of G with clique number at most $\omega - 1$.*

Proof. No vertex can share an edge with any ancestors in T other than its parent, since if it shared an edge with an ancestor higher than its parent then it would have appeared at a higher level as a neighbour of the ancestor. This means a path from the root down T forms an induced path, and the depth of T cannot be longer than the longest induced path in G.

Consider two vertices w_1 and w_2 in the same level t of the spanning tree T. Then the children of w_1 cannot share any edges in G with the children of w_2. This is because the children of w_1 constitute a connected component of the subgraph of G induced by all vertices that are not in the first t levels of T. Armed with this observation, consider a clique of size ω within a level of T. When combined with the common parent, this would form a clique of size $\omega + 1$ in G, a contradiction. Thus the clique number of any level of T is at most $\omega - 1$.

Finally, similar to breadth-first search or depth-first search, since each edge is examined at most twice, the time complexity is $O(|V| + |E|) = O(m^2)$. □

Proof of Lemma 4.15. We can prove the theorem by induction on ω. Let $\mathcal{A}_\omega$ denote the coloring algorithm which applies to graphs of clique number ω. When $\omega = 1$, there are no edges and there is an algorithm $\mathcal{A}_1$ which can color the graph using a single color in $O(m^2)$ time. For the inductive step, assume there is a coloring algorithm $\mathcal{A}_{\omega-1}$ using $\ell^{\omega-2}$ colors and runtime $O(m^2\omega)$ for any graph of clique number $\omega - 1$ and longest induced path ℓ.

The coloring algorithm $\mathcal{A}_\omega$ for graphs of clique number ω is as follows. First apply the neighbour-first search algorithm to find spanning tree T of G. Then for each level of T, apply $\mathcal{A}_{\omega-1}$. For each level, we use a disjoint set of colors. Since there are at most ℓ levels in the T, the number of colors used by $\mathcal{A}_\omega$ is at most $\ell^{\omega-1}$ by the induction hypothesis.

It remains to analyze the runtime of $\mathcal{A}_\omega$. Say the neighbour-first search step has runtime at most Cm^2 in the worst case for some constant C. We will show that the runtime of $\mathcal{A}_\omega$ is at most $Cm^2\omega$. From the induction hypothesis, the applications of $\mathcal{A}_{\omega-1}$ have total runtime $\sum_i Cm_i^2(\omega - 1) \leq Cm^2(\omega - 1)$, where m_i is the number of vertices in the i^{th} layer. Here we used $\sum_i m_i^2 \leq \left(\sum_i m_i\right)^2 = m^2$. Thus the runtime of $\mathcal{A}_\omega$ is $Cm^2 + Cm^2(\omega - 1) \leq Cm^2\omega$. □

4.3 Learning local fermionic operators

In this section we consider shadow tomography for local fermionic observables.

A system of n fermionic modes is associated with a set of $2n$ Majorana fermion operators, which are mutually anticommuting Hermitian observables $\{\gamma_a\}_{a\in[2n]}$ that act on a Hilbert space of dimension 2^n. We can represent them by a set of Pauli operators satisfying

$$\gamma_1, \gamma_2, \ldots, \gamma_{2n} \in \mathcal{P}^{(n)} \qquad \gamma_a\gamma_b + \gamma_b\gamma_a = 2\delta_{ab}\mathbb{1}. \tag{4.35}$$

There are a variety of specific choices (fermion-to-qubit mappings) that satisfy the above, including the Jordan-Wigner mapping [JW28], the Bravyi-Kitaev mapping [BK02], and the ternary tree mapping [Vla19; Jia+20a]. The latter two mappings have the desirable property that each Majorana fermion operator is represented by a Pauli operator of low weight $O(\log n)$.

We are interested in k-body fermionic observables as defined in Chapter 4. We write

$$\Gamma(x) = i^{|x|\cdot(|x|-1)/2}\gamma_1^{x_1}\gamma_2^{x_2}\ldots\gamma_{2n}^{x_{2n}} \qquad x \in \{0,1\}^{2n} \tag{4.36}$$

for the Majorana monomials, and

$$\mathcal{F}_k^{(n)} = \{\Gamma(x) : |x| = 2k\} \tag{4.37}$$

for the set of k-body Majorana operators on n fermionic modes.

In the case of 1-body observables there is a simple and practical algorithm for coloring induced subgraphs of $G(\mathcal{F}_1^{(n)})$ with bounded clique number.

Lemma 4.17. *Let G' be any induced subgraph of the commutation graph $G(\mathcal{F}_1^{(n)})$ of 1-body fermionic observables, and let ω be the size of the largest clique in G'. There is a classical algorithm with runtime $O(n^2\omega)$ that computes a coloring of G' with at most $\omega + 1$ colors.*

Proof. Let $G' = G(\mathcal{S})$ be the subgraph of $G(\mathcal{F}_1^{(n)})$ induced by some subset $\mathcal{S} \subseteq \mathcal{F}_1^{(n)}$ of 1-body fermionic observables. Let ω be the maximum size of a clique in G'.

Consider an auxiliary graph $H(\mathcal{S})$ defined as follows. This graph $H(\mathcal{S})$ has $2n$ vertices labeled by the Majorana fermion operators $\{\gamma_1, \gamma_2, \ldots, \gamma_{2n}\}$. For each observable $i\gamma_a\gamma_b \in \mathcal{S}$, we include an edge $\{\gamma_a, \gamma_b\}$ in $H(\mathcal{S})$. Two elements of $\mathcal{S}$ commute if and only if they do not share any Majorana fermion operators. For example, $i\gamma_1\gamma_2$ anticommutes with $i\gamma_2\gamma_3$ but commutes with $i\gamma_3\gamma_4$. Thus a commuting set of 1-body fermionic observables corresponds to a matching in

$H(\mathcal{S})$, and partitioning $\mathcal{S}$ into commuting sets corresponds to an edge coloring of $H(\mathcal{S})$.

Now observe that our graph of interest G' is the *line graph* of $H(\mathcal{S})$. An edge coloring of $H(\mathcal{S})$ gives a vertex coloring of G'. The edge coloring algorithm of Misra and Gries [MG92] computes an edge-coloring of a graph H using no more than $\deg(H) + 1$ colors, where $\deg(H)$ is the maximum degree of any vertex in H. But the edges connecting to a single vertex in our graph $H(\mathcal{S})$ form a clique in G', so the degree of $H(\mathcal{S})$ is at most ω. The runtime of the edge coloring algorithm is asymptotically upper bounded by the number of vertices times the number of edges, which in our case is $O(n \cdot n\omega)$. □

It is also possible to directly vertex color the given graph G' using Brooks' theorem, which states that the chromatic number of a graph is at most its maximum degree +1, since a high degree vertex also yields a large clique. This argument yields a slightly looser bound of 2ω.

We now prove Lemma 4.5, restated below.

Lemma 4.5. *Let $k \geq 1$, and let G' be any induced subgraph of the commutation graph $G(\mathcal{F}_k^{(n)})$ of k-body fermionic observables, and let ω be the size of the largest clique in G'. Then the fractional chromatic number of G' satisfies*

$$\chi_f(G') \leq p_k(\omega), \tag{4.10}$$

where p_k is a polynomial. Moreover, for any $k = O(1)$ we can sample from a fractional coloring of G' with size $p_k(\omega)$ using a classical algorithm with runtime $\mathrm{poly}(n)$. *The polynomials for $k = 1, 2$ are $p_1(\omega) = \omega + 1$ and $p_2(\omega) = O(\omega^8)$.*

Recall the definition of Majorana monomials from Equation (4.36). In the following we shall use the commutation relations of these operators which we now derive. Using Equation (4.2) we get

$$\gamma_j \Gamma(y) = (-1)^{|y|+y_j} \Gamma(y) \gamma_j \qquad y \in \{0,1\}^{2n} \quad j \in [2n]. \tag{4.38}$$

Applying the above for all indices j in the support of $x \in \{0,1\}^{2n}$ gives

$$\Gamma(x)\Gamma(y) = (-1)^{|x||y|+y\cdot x} \Gamma(y)\Gamma(x) \qquad x, y \in \{0,1\}^{2n}. \tag{4.39}$$

Claim 4.1. *Suppose $x, y \in \{0,1\}^{2n}$ are such that $|x|, |y|$ are either both even, or both odd. If $x_j = y_j = 0$ for some $j \in [2n]$ then*

$$[\gamma_j \Gamma(x), \gamma_j \Gamma(y)] = 0 \quad \textit{if and only if} \quad [\Gamma(x), \Gamma(y)] = 0. \tag{4.40}$$

Proof. Follows directly from Equation (4.39). □

In our proof, it will be helpful to consider Majorana monomials of both odd and even degree. Write

$$\mathcal{M}_r^{(n)} = \{\Gamma(x) : |x| = r,\ x \in \{0,1\}^{2n}\} \tag{4.41}$$

for the set of degree-r Majorana monomials, so that $\mathcal{M}_{2k}^{(n)} = \mathcal{F}_k^{(n)}$ are the k-body fermionic observables of interest.

Proof. The proof is by induction in r. Our inductive hypothesis is that, for any induced subgraph H of the commutation graph $G(\mathcal{M}_r^{(n)})$ of degree-r Majorana monomials with largest clique of size at most ω, we can sample from a fractional coloring of H with $f_r(\omega)$ colors using a classical algorithm with runtime $t_r(n)$ such that $t_r(n) = \text{poly}(n)$ for any constant $r = O(1)$. Here $f_r(\omega)$ is a polynomial that we determine below. Ultimately we are interested in the even values of r and we have $p_k(\omega) = f_{2k}(\omega)$ where p_k is the polynomial in the statement of Lemma 4.5.

The base case is $r = 2$. We saw in Lemma 4.17 that if $\mathcal{S} \subseteq \mathcal{M}_2^{(n)}$ and its commutation graph $G(\mathcal{S})$ has no cliques larger than ω, then $G(\mathcal{S})$ can be colored with $\omega+1$ colors using a classical algorithm with runtime $O(n^2\omega) = O(\text{poly}(n))$ since $\omega \leq |\mathcal{M}_2^{(n)}| = O(n^2)$. Thus $f_2(\omega) = \omega + 1$ and we can efficiently sample from the coloring by selecting a color uniformly at random.

In the following two claims we handle the induction step separately for the odd and even values of r.

Claim 4.2. *Suppose $r \geq 3$ is odd. Then*

$$f_r(\omega) = r\omega f_{r-1}(\omega). \tag{4.42}$$

Proof. Let $r \geq 3$ be odd, let G' be an induced subgraph of $G(\mathcal{M}_r^{(n)})$, and suppose the largest clique in G' has size at most ω. Let $V \subseteq \mathcal{M}_r^{(n)}$ be the vertex set of G'. Let

$\Gamma(x^1), \Gamma(x^2), \ldots, \Gamma(x^L) \in V$ be a maximal set of pairwise anticommuting operators in V. We can construct such a set by starting at any vertex of G' and greedily adding vertices until this is no longer possible. By definition, this set is a clique in G' and therefore $L \leq \omega$.

Let $I \subseteq [2n]$ be the set of all indices of Majoranas that appear in these operators. Since each has weight r, we have

$$|I| \leq r\omega. \tag{4.43}$$

For convenience let us relabel the Majorana fermion operators so that $I = \{1, 2, \ldots, T\}$ where $T \leq r\omega$. Then define

$$\mathcal{S}_i = \{\Gamma(z) \in V : z_i = 1, \text{ and } z_j = 0 \text{ for all } 1 \leq j \leq i-1\}. \tag{4.44}$$

We now show that V can be partitioned as

$$V = \mathcal{S}_1 \sqcup \mathcal{S}_2 \sqcup \ldots \sqcup \mathcal{S}_T. \tag{4.45}$$

By definition, the sets on the RHS are disjoint and each contained in V so all we need to show is that for any $\Gamma(y) \in V$ there is some $i \in T$ such that $\Gamma(y) \in \mathcal{S}_i$. So let $\Gamma(y) \in V$ be given. Since the set $\Gamma(x^1), \Gamma(x^2), \ldots, \Gamma(x^L) \in V$ is a maximal set of pairwise anticommuting operators, we must have

$$[\Gamma(y), \Gamma(x^j)] = 0 \quad \text{for some } j \in [L]. \tag{4.46}$$

Since $|y| = |x| = r$ are both odd we see from Equation (4.39) that this implies $y \cdot x^j \neq 0$. Therefore $y_i = 1$ for some index $i \in \{1, 2, \ldots, T\}$. Let $\ell \in [T]$ be the smallest index such that $y_\ell = 1$. Then $\Gamma(y) \in \mathcal{S}_\ell$ and we have shown V can be partitioned as in Equation (4.45).

Now for each $1 \leq i \leq T$ consider the commutation graph $G(\mathcal{S}_i)$. Each operator $\Gamma(z) \in \mathcal{S}_i$ has $z_i = 1$. From Claim 4.1, the commutation graph of $\mathcal{S}_i$ is therefore unchanged if we flip $z_i \leftarrow 0$ for all $\Gamma(z) \in \mathcal{S}_i$. Define

$$\mathcal{S}'_i = \{\Gamma(z \oplus \hat{e}_i) : \Gamma(z) \in \mathcal{S}_i\}. \tag{4.47}$$

We have shown that the commutation graph $G(\mathcal{S}_i)$ coincides with the commutation graph $G(\mathcal{S}'_i)$, where the set $\mathcal{S}'_i \subseteq \mathcal{M}^{(n)}_{r-1}$ only contains degree-$(r-1)$ Majorana monomials. Moreover, $G(V)$ does not contain any clique larger than ω, so neither

does its induced subgraph $G(\mathcal{S}_i)$. Therefore $G(\mathcal{S}'_i)$ does not contain any clique of size greater than ω.

By our inductive hypothesis, for each $1 \le i \le T$, we can sample efficiently from a fractional coloring of $G(\mathcal{S}'_i) = G(\mathcal{S}_i)$ with size at most $f_{r-1}(\omega)$. Now let us define a fractional coloring of V in which we choose an index $i \in [T]$ uniformly at random and then sample an independent set in $\mathcal{S}_i$ according to the fractional coloring of $G(\mathcal{S}_i)$. Note that any independent set in $G(\mathcal{S}_i)$ is also an independent set in $G(V)$, so this defines a valid fractional coloring. Moreover, the probability of any vertex $u \in G(V)$ being sampled is equal to the probability that we choose i such that $u \in \mathcal{S}_i$ (this probability is $1/T$) times the probability that the sampled independent set of $\mathcal{S}_i$ contains u (this is at least $1/f_{r-1}(\omega)$ by our inductive hypothesis). This procedure samples a fractional coloring of size

$$T \cdot f_{r-1}(\omega) \le r\omega \cdot f_{r-1}(\omega) \tag{4.48}$$

as claimed. To sample from the fractional coloring, we need to first construct a maximal set of pairwise anticommuting operators $\Gamma(x^1), \ldots, \Gamma(x^L)$, from which we can define the set I and the partition Equation (4.45). As noted above, this step can be performed by starting at an arbitrary vertex $\Gamma(x^1)$ of G' and then growing the set one operator at a time until this is no longer possible. This step has $\mathrm{poly}(n)$ runtime because the graph has at most $|\mathcal{M}_r^{(n)}| = \binom{2n}{r} = \mathrm{poly}(n)$ vertices. The next step is to choose an index $1 \le i \le T$ at random and sample an independent set of $\mathcal{S}_i$ uniformly at random using a fractional coloring of $G(\mathcal{S}_i)$ which by our inductive hypothesis can be done in $\mathrm{poly}(n)$ time. □

Claim 4.3. *Suppose $r \ge 4$ is even. Then*

$$f_r(\omega) = (f_{r-1}(\omega))^r \,. \tag{4.49}$$

Proof. Let $r \ge 4$ be even, let G' be an induced subgraph of $G(\mathcal{M}_r^{(n)})$, and let ω be the size of the largest clique in G'. Let $V \subseteq \mathcal{M}_r^{(n)}$ be the vertex set of G'. For each $1 \le i \le 2n$, define

$$W_i = \{\Gamma(x) \in V : x_i = 1\}. \tag{4.50}$$

Note any clique in $G(W_i)$ has size at most ω. From Claim 4.1, the commutation graph of W_i is unchanged if we flip $z_i \leftarrow 0$ for all $\Gamma(z) \in W_i$. Define

$$W'_i = \{\Gamma(z \oplus \hat{e}_i) : \Gamma(z) \in W_i\}. \tag{4.51}$$

Then $G(W_i) = G(W_i')$ and any clique in $G(W_i')$ has size at most ω. Moreover, W_i' is a set of degree-$(r-1)$ Majorana monomials, and by our inductive hypothesis we can efficiently sample a coloring of $G(W_i')$ with size at most $f_{r-1}(\omega)$. Let q_i be the corresponding fractional coloring of W_i, for each $1 \leq i \leq 2n$.

Now let us randomly sample a set $\Omega \subseteq V$ as follows. First, select independent sets $I_1 \sim q_1, I_2 \sim q_2, \ldots, I_{2n} \sim q_{2n}$ according to the fractional colorings described above. Then let

$$\Omega = \{\Gamma(x) \in V : \Gamma(x) \in I_j \text{ for all } j \in [2n] \text{ such that } x_j = 1\}. \tag{4.52}$$

Since $|x| = r$ for all $\Gamma(x) \in V$, we have

$$\Pr(\Gamma(x) \in \Omega) \geq \left(\frac{1}{f_{r-1}(\omega)}\right)^r \qquad \Gamma(x) \in V. \tag{4.53}$$

Now let us show that Ω is an independent set in V; this implies that the above procedure samples from a fractional coloring of V with $(f_{r-1}(\omega))^r$ colors. So suppose $\Gamma(x), \Gamma(y) \in \Omega$. We will show that $\Gamma(x), \Gamma(y)$ commute; equivalently, there is no edge between the corresponding vertices in G'. First suppose $x \cap y = \emptyset$. In this case, since $|x| = |y| = r$ are both even, it follows directly that $[\Gamma(x), \Gamma(y)] = 0$. If on the other hand $x_j = y_j = 1$ for some $j \in [2n]$, Then $\Gamma(x), \Gamma(y) \in W_j \cap \Omega$. But $W_j \cap \Omega \subseteq I_j$ is an independent set in the commutation graph of W_j, and therefore $[\Gamma(x), \Gamma(y)] = 0$.

The algorithm we have described above only involves identifying the subsets of vertices W_i for $1 \leq i \leq 2n$ (which can be done in linear time in the number of vertices of G', which is upper bounded polynomially in n), and then using $O(n)$ calls to the subroutine for sampling fractional colorings of commutation graphs of degree-$(r-1)$ Majorana monomials with clique number at most ω. Since this subroutine has poly(n) runtime by our inductive hypothesis, so does the algorithm described above. □

Putting together Claim 4.2 and Claim 4.3 and Lemma 4.17 we see that the sizes $f_r(\omega)$ of the fractional colorings are polynomial functions of ω with degree that depends only on r. For even values of r the polynomials $p_k(\omega) = f_{2k}(\omega)$ satisfy the recurrence

$$p_1(\omega) = \omega + 1 \quad \text{and} \quad p_k(\omega) = ((2k-1)\omega p_{k-1}(\omega))^{2k} \qquad k \geq 2. \tag{4.54}$$

For the 2-body and 3-body fermionic observables we get

$$p_2(\omega) = O(\omega^8) \quad , \quad p_3(\omega) = O(\omega^{54}). \tag{4.55}$$

In general, we have the upper bound

$$p_k(\omega) \leq (2k\omega)^{(2k)^{k+1}}. \tag{4.56}$$

□

Chapter 5

LEARNING BOSONIC OBSERVABLES

Efficiently extracting information from quantum states is a central task in quantum information science. It is crucial in physical experiments and will be critical in simulations run on future quantum computers. Often learning the details of an entire quantum state is not required, but rather we would like to extract the expectation values of a set of interesting observables.

Naively, measurements of many properties of quantum states are constrained by the uncertainty principle for non-commuting operators. Additionally, tomographic techniques for exactly learning a quantum state up to a stringent standard like worst case observable error (or trace distance) are known to scale exponentially in the number of qubits or polynomially in the size of the Hilbert space [OW16; AA23]. Surprisingly however, the development of shadow tomography techniques demonstrated that one can learn a set of expectation values of even non-commuting observables with high probability with a shockingly small number of samples, scaling only polylogarithmically in the number of observables [Aar18; Aar+18; AR19; BO21]. While powerful, the general schemes suffer from two large caveats: they are computationally inefficient, and they require immense quantum memories, sometimes millions of times the size of the original state, to enable huge entangled measurements. Classical shadows [HKP20] were developed to circumvent both of these limitations—they are computationally efficient, and require only single-copy measurements for a wide class of useful observables. However, classical shadows place limitations on the sets of observables that are available, for example some schemes are only able to learn observables which are either local or low rank.

It is now known that many of the limitations of classical shadows performed only on single copies at a time are fundamental. For general quantum states, certain collections of observables can only be learned with a logarithmic number of samples by exploiting entangled measurements across multiple copies of a state [Che+22; ACQ22; Hua+22]. This was shown to be true even for some of the simplest large sets of observables, namely Pauli operators on n qubits. Phrased a different way, the ability to make entangled measurements on copies of a quantum state can grant exponentially more power in learning tasks. Since such schemes require a

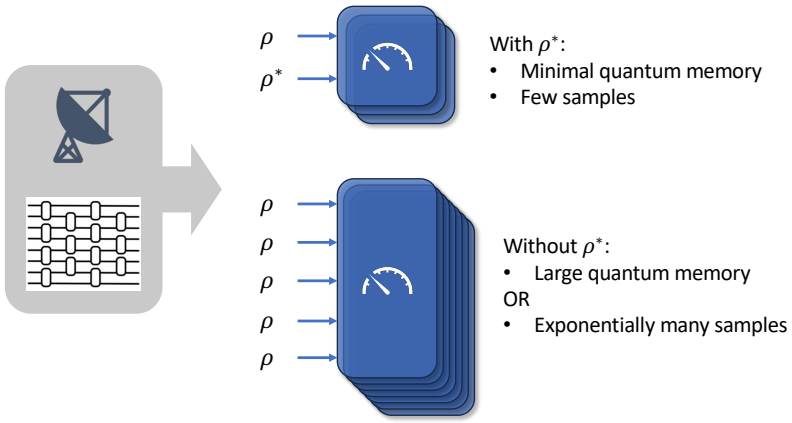

Figure 5.1: A cartoon of the techniques and novel resources in this work. Here we quantify the advantage endowed by minimal quantum memories containing $\rho \otimes \rho^*$ in learning about natural properties of quantum states coming from quantum sensors or digital quantum simulations. (Left) The resources ρ and ρ^* are available from both computational sources such as a known quantum circuit or natural sources like certain quantum sensor setups. (Right) This resource provides an exponential advantage in queries and computation for some learning tasks when only a minimal quantum memory (or constant number of copies) is available, as in most devices in the foreseeable future. Without ρ^*, one needs either a large quantum memory or exponentially many samples. Several applications of this technique are introduced and we believe this will motivate the development of further applications of minimal quantum memories.

quantum memory to store simultaneous copies of an unknown state and this type of advantage cannot be overcome even by an arbitrary amount of classical computation when samples are limited, this constitutes a promising future application of quantum computers. It was demonstrated experimentally in Ref. [Hua+22] that the advantage persists even for small numbers of qubits and in the presence of noise. Importantly, Ref. [Hua+22] showed that exponential advantages are available using only 2 copies at a time of the state in quantum memory, or an example of a minimal quantum memory for which the number of copies required is independent of the learning task. This demonstrated the existence of learning tasks for which extremely limited quantum resources could provide huge advantages, in contrast to the need of general shadow tomography to have memories that could be millions of times larger than the system of interest even when tasked with estimating observables to a precision of only 10^{-3}.

In this chapter, we explore a novel resource for learning which is able to grant exponential advantages using only a minimal quantum memory (space for only 2

copies)—the ability to make joint measurements on an unknown quantum state with its complex conjugate, denoted $\rho \otimes \rho^*$. We give a learning task that can be achieved with low sample complexity using measurements on $\rho \otimes \rho^*$. In contrast, without access to ρ^*, the same learning task requires exponentially more measurements, unless the size of the quantum memory is allowed to expand to practically unrealistic sizes as a function of the learning task. It is known that quantum memory access to ρ and ρ^* can provide large advantages for learning tasks [Mon17; GNW21; Gre+23; WCL23; HBK23; HLK23; MM22; MA23]. However, to the knowledge of the authors this is the first time such an advantage has been demonstrated for a shadow tomography task. We prove a lower bound showing that that copies of ρ^* and ρ without quantum memory are also insufficient for the learning task.

Our exponential separation holds for a natural and physically motivated set of observables. The setting is a d-dimensional Hilbert space that discretizes position and momentum space with a natural limit of a continuous bosonic mode, and the operators we learn if we take the infinite dimensional limit are the bosonic displacement operators [Bra+23]. These operators more naturally correspond to real-space arrays of quantum sensors. Recent developments in reconfigurable atom arrays [Blu+24] may provide a fruitful test bed for applications in a sensing context for example, especially given their wide bandwidth and sensitivity in other applications [OM99; Sed+12; Hol+17; Wad+17; Sim+21]. In addition to the learning algorithm using $\rho \otimes \rho^*$, we develop a version of classical shadows tailored to the d-dimensional bosonic setting. It uses a uniform distribution over the generalized d-dimensional Clifford group to make single-copy measurements with good predictive power. Despite the more limited power of single copies, we identify a wide class of quantities that are efficiently learnable.

Let's begin with some background on the operators and states that we aim to learn, and then state our main theorems showing the exponential power of access to the complex conjugate resource in a minimal quantum memory. Shifts in discrete position and momentum space are given by the d-dimensional clock and shift operators, Z and X, which are generalizations of the qubit Pauli operators to d

dimensions.

$$X = \begin{pmatrix} 0 & 0 & \dots & 0 & 1 \\ 1 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & 0 \end{pmatrix}, \tag{5.1}$$

$$Z = \begin{pmatrix} 1 & 0 & \dots & 0 \\ 0 & \omega & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \omega^{d-1} \end{pmatrix}, \tag{5.2}$$

where $\omega := e^{i2\pi/d}$. The operators map naturally to a 1D discrete line in real space, and can be interchanged via the d-dimensional quantum Fourier transform. Combined position and momentum shifts may be lumped together into displacement operators $D_{q,p}$, defined briefly as

$$D_{q,p} = e^{i\pi qp/d} X^q Z^p \tag{5.3}$$

and in more detail in Section 5.1.

Definition 5.1. *The* displacement amplitudes *of a d-dimensional state ρ are*

$$y_{q,p} = \mathrm{Tr}\left(D_{q,p}\rho\right). \tag{5.4}$$

The displacement operators may be used to form a basis for quantum states, and there are d^2 displacement amplitudes.

The central task we consider in this chapter is to estimate all $y_{q,p}$ to precision ε given copies of an unknown quantum state ρ.

Task 5.2. (Informal) *Given access to a quantum state ρ, estimate all the displacement amplitudes $\{y_{q,p}\}$ to precision ε with high probability.*

Our first result is a sample complexity lower bound showing the minimal size of a conventional quantum memory required to efficiently perform this task without access to the resource ρ^*. We assume we can measure K copies $\rho^{\otimes K}$ at a time, possibly in entangled bases, but allow no access to ρ^*.

Theorem 5.1. *Let d be the dimension of the Hilbert space. Assume d is prime. Any protocol which learns the magnitudes of all displacement amplitudes to precision ε with probability $2/3$ by measuring copies of $\rho^{\otimes K}$ for $K \leq 1/(12\varepsilon)$ requires $\Omega(\sqrt{d}/(K^2\varepsilon^2))$ measurements.*

The full proof of this is given in Section 5.3. Using a quantum memory with $\rho^{\otimes K}$, performing the learning task whilst consuming a number of copies scaling only as polylog(d) is impossible, even if K grows as large as $1/(12\varepsilon)$. Hence it is impossible to efficiently perform the task with a minimal quantum memory. Our lower bound techniques do not strictly cover the case when d is not prime. However, one would hope that a successful learning protocol would function equally well for any d; Theorem 5.1 is sufficient to rule out protocols which work for all d.

This negative result may lead one to conclude that quantum memories are not as powerful as one might hope for physical learning tasks, but there is a resolution to this challenge which reveals an interesting subtlety in the power of quantum computing in analyzing quantum data. While a quantum memory containing K states $\rho^{\otimes K}$ is insufficient, measurements on the state $\rho \otimes \rho^*$ are able to learn the displacement amplitudes of ρ up to a sign. Not only does the learning algorithm using $\rho \otimes \rho^*$ have logarithmic sample complexity, but the algorithm is very simple and computationally efficient. Note that we use the term "up to a sign" to convey that because these are unitary but not always Hermitian operators, they are complex valued and we are able to learn some, but not all, phase information about that value with this procedure. With the detailed proof and algorithm given in Section 5.2, we show

Theorem 5.2. *There is an algorithm which can learn all displacement amplitudes up to a possible minus sign, with precision ε, using $O(\log d/\varepsilon^4)$ samples. The algorithm makes measurements only on copies of $\rho \otimes \rho^*$ contained in a minimal quantum memory. Moreover, the algorithm is computationally efficient.*

When viewed together, these theorems highlight the exponential advantage of using ρ^* as a resource in learning tasks. This naturally leads one to wonder whether the state ρ^* has this power in the single copy setting, but indeed lower bounds rule this out and show that entangled measurements are also a necessary component of the learning algorithm in Theorem 5.2. Our following theorem articulates this more precisely.

Theorem 5.3. *Let d be the dimension of the Hilbert space. Any single-copy protocol which learns the magnitudes of all displacement amplitudes to precision ε with probability $2/3$ requires a number of copies scaling as $\Omega(d/\varepsilon^2)$. This holds even if the protocol has access to single-copy measurements of both ρ and ρ^*.*

With the detailed proof given in Section 5.3, this result reiterates the conclusion that entangled measurements using quantum memories have dramatically more power than those that can process only a single copy at a time, and indeed that ρ^* inside a minimal quantum memory is a powerful and novel resource.

Our algorithm to learn the displacement amplitudes up to a sign using $\rho \otimes \rho^*$ is inspired by the Pauli shadow tomography algorithm in [HKP21]. The difficulty of measuring many Paulis with single-copy measurements arises because they are highly non-commuting, and so any particular measurement basis will only give information about a small subset of Paulis. The insight to get around this using two-copy measurements is as follows. When two Paulis P and Q do not commute, they anticommute $PQ = -QP$. Thus $P \otimes P$ always commutes with $Q \otimes Q$, and the set of operators of the form $P \otimes P$ on two registers mutually commute for all Paulis P. This means there is a simultaneous eigenbasis of all $P \otimes P$ which is entangled across the two registers; it turns out that this eigenbasis is the Bell basis on each pair of qubits. Measuring $\rho \otimes \rho$ in the Bell basis then allows simultaneous estimation of all $\operatorname{Tr}(P \otimes P)(\rho \otimes \rho) = (\operatorname{Tr} P\rho)^2$, which gives estimates of all magnitudes $|\operatorname{Tr} P\rho|$.

Applying this trick to the displacement operators requires some care, and reveals why the complex conjugate state ρ^* is important. Displacement operators $D_{q,p}$ obey the commutation relation

$$D_{q',p'} D_{q,p} = e^{i2\pi(qp'-q'p)/d} D_{q,p} D_{q',p'}. \tag{5.5}$$

Unlike the Pauli case, it is not true that $D_{q,p} \otimes D_{q,p}$ mutually commute for all q, p, since the complex phase acquired from braiding the operators does not square to 1 in general. However, the operators $D_{q,p} \otimes D_{-q,p}$ *do* mutually commute, since the complex phase in the commutation relation will cancel with its complex conjugate. As in the Pauli case, this implies there is an entangled simultaneous eigenbasis of all $D_{q,p} \otimes D_{-q,p}$. This turns out to be the generalized Bell basis, which can be accessed via the mod-d quantum fourier transform and a controlled shift. If we measure $\rho \otimes \rho$ in the generalized Bell basis, we can derive estimates of all quantities of the form $\operatorname{Tr}(D_{q,p} \otimes D_{-q,p})(\rho \otimes \rho)$. However, this does not give us information about $\operatorname{Tr} D_{q,p}\rho$. Rather, if we measure $\rho \otimes \rho^*$ in the generalized Bell basis, we get estimates of the quantities

$$\begin{aligned} \operatorname{Tr}\left((D_{q,p} \otimes D_{-q,p})(\rho \otimes \rho^*)\right) &= \operatorname{Tr}\left(D_{q,p}\rho\right) \operatorname{Tr}\left(D_{q,p}^T \rho^*\right) \\ &= \operatorname{Tr}\left(D_{q,p}\rho\right)^2 = y_{q,p}^2 \end{aligned} \tag{5.6}$$

from which we can learn all displacement amplitudes $y_{q,p}$ up to a $\pm$ sign.

Finally, we highlight two potential applications of this chapter. As a first application, we consider cases where an explicit quantum circuit is known for a state we wish to study. When conducting a physical experiment, we are often preparing some natural quantum state and subsequently performing measurements. In quantum simulation, we aim to design quantum algorithms that simulate Nature, so that the quantum algorithm prepares the physical quantum state of interest. One may wonder, is there any advantage to having a quantum algorithm which prepares state ρ, rather than accessing copies of ρ through an experimental setup? In particular, are there natural learning tasks where performing quantum simulation gives a big benefit?

Task 5.2 answers this question in the affirmative. While other polynomial advantages to having access to the source code are known [KO23], the learning task here demonstrates an exponential advantage over black box access. In general experimental setups, it can sometimes be unclear how to access the complex conjugate ρ^* of the state of interest ρ. However, if we have a quantum algorithm which prepares ρ we can easily access ρ^* on our quantum computer—we simply complex conjugate the quantum algorithm itself.

More concretely, suppose unitary U prepares state ρ via $\rho = \mathrm{Tr}_{\bar{S}}\left(U|0...0\rangle\langle 0...0|U^\dagger\right)$, and we have an efficient quantum circuit for U. Then by complex conjugating every gate in the circuit we can implement unitary U^*, which will prepare state ρ^* via $\rho^* = \mathrm{Tr}_{\bar{S}}\left(U^*|0...0\rangle\langle 0...0|U^T\right)$. Theorem 5.1 and Theorem 5.2 then exhibit an exponential cost saving for Task 5.2 from having a description of a quantum circuit which prepares a quantum state of interest ρ.

As a second category of applications, we consider unknown quantum states collected from nature. For example, these states could be gathered via quantum sensors or transduced from other quantum systems. The ability to learn about unknown states ρ with exponentially fewer samples using a minimal quantum memory with only $K = 2$ prompted experimental demonstrations of this idea showing they were robust even with noisy operations today [Hua+22]. Despite these promising results, connecting these advantages to existing quantum sensor states today has been challenging, as many quantum sensors today are single qubit, ensembles of single qubits, stretched single qubits like GHZ states, or cavity modes [DRC17].

In contrast to the collection of qubit case, the operators considered in d-dimensions here connect naturally with quantum sensor arrays in regular spatial arrangements

with connections to applications like very long baseline interferometry enabled by quantum communication [GJC12]. The displacement operators are a natural description of discrete position and momentum for real-space arrays, especially in a quantum regime where few excitations are expected and background thermal noise is high. For these scenarios, we argue that one way to view the results here is as a specialized form of mixedness testing, where for a natural class of signals, exponentially fewer samples are required to detect the presence of the signal when combined with a sea of background noise. For applications like detection of radio signals as in NMR, it is common for an infinite temperature background to be quite strong, and this may find applications in that area. These results provide a compelling setting for which exponential advantage in signal detection is possible with quantum memory.

5.1 Background on bosonic systems

Displacement operators

Given a d-dimensional quantum system $\mathbb{C}^d$ with basis $\{|0\rangle, \ldots, |d-1\rangle\}$, we define the operators X and Z by

$$X : |j\rangle \to |j+1\rangle \tag{5.7}$$

$$Z : |j\rangle \to \omega^j |j\rangle. \tag{5.8}$$

Here and throughout, the addition inside the ket is modulo d, and

$$\omega := e^{i2\pi/d}. \tag{5.9}$$

The matrix representations of X and Z are

$$X = \begin{pmatrix} 0 & 0 & 0 & \ldots & 0 & 1 \\ 1 & 0 & 0 & \ldots & 0 & 0 \\ 0 & 1 & 0 & \ldots & 0 & 0 \\ 0 & 0 & 1 & \ldots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \ldots & 1 & 0 \end{pmatrix}, \tag{5.10}$$

$$Z = \begin{pmatrix} 1 & 0 & 0 & \ldots & 0 \\ 0 & \omega & 0 & \ldots & 0 \\ 0 & 0 & \omega^2 & \ldots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \ldots & \omega^{d-1} \end{pmatrix}. \tag{5.11}$$

X, Z are traceless and unitary but not Hermitian in general. X is a *shift* rotation, generated by the discrete momentum operator, and Z is a *phase* rotation, generated by the discrete position operator. For $d = 2$, X and Z coincide with the usual Pauli matrices, and in this case they are in fact Hermitian. X, Z obey the commutation relation

$$\begin{aligned} ZX &= \omega XZ \\ \implies \quad Z^p X^q &= \omega^{qp} X^q Z^p. \end{aligned} \tag{5.12}$$

Define the *displacement operator* $D_{q,p}$ by

$$D_{q,p} = e^{i\pi qp/d} X^q Z^p. \tag{5.13}$$

$D_{q,p}$ acts on basis vectors by

$$D_{q,p} : |j\rangle \to e^{i\pi(q+2j)p/d} |j+q\rangle. \tag{5.14}$$

$D_{q,p}$ is a unitary which can be interpreted as shifting by the vector (q, p) in discrete position-momentum phase space.

Proposition 5.4. *The displacement operators have properties*

$$D_{q,p}^{-1} = D_{q,p}^\dagger = D_{-q,-p} \tag{5.15}$$

$$D_{q,p}^* = D_{q,-p} \tag{5.16}$$

$$D_{q,p}^T = D_{-q,p} \tag{5.17}$$

$$D_{q',p'} D_{q,p} = e^{i2\pi(qp'-q'p)/d} D_{q,p} D_{q',p'} \tag{5.18}$$

$$D_{q,p}^k = D_{kq,kp} \tag{5.19}$$

Proposition 5.5. *$\{D_{q,p}\}$ form a basis of the $(d \times d)$-dimensional space of operators on $\mathbb{C}^d$. Moreover, they are orthogonal in the Hilbert-Schmidt inner product.*

$$\mathrm{Tr}\left(D_{q,p}^\dagger D_{q',p'}\right) = d \cdot \delta_{q,q'} \delta_{p,p'}. \tag{5.20}$$

Proposition 5.5 lets us decompose any d-dimensional density matrix ρ as

$$\rho = \frac{1}{d} \sum_{q,p} \mathrm{Tr}\left(D_{q,p}\rho\right) D_{q,p}^\dagger. \tag{5.21}$$

This can be viewed as the d-dimensional analog of the Bloch vector representation

$$\rho = \frac{1}{d}\Big(\mathbb{1} + \sum_{(q,p)\neq(0,0)} y_{q,p} D_{q,p}^\dagger\Big) \quad , \quad y_{q,p} = \mathrm{Tr}\left(D_{q,p}\rho\right). \tag{5.22}$$

Bosonic limit $d \to \infty$

For an overview of continuous bosonic modes, see [Bra+23].

As d increases, the d-dimensional system described above forms an increasingly good discrete approximation to an infinite-dimensional continuous bosonic mode. To formalize this, let's set

$$x = \sqrt{\frac{\pi}{d}}(q, p). \tag{5.23}$$

x lives on a lattice with spacing $\sim 1/\sqrt{d}$ and size $\sim \sqrt{d}$ with periodic boundary conditions. As d goes to infinity, x becomes a continuous phase space variable in the plane.

For the quantum harmonic oscillator with frequency ω and mass m, the characteristic length and momentum scales are $\sqrt{\hbar/m\omega}$ and $\sqrt{\hbar m\omega}$ respectively. Thus if we can measure the position to precision δ, this is effectively measuring the particle on a discrete lattice with

$$d \sim \frac{\hbar}{m\omega\delta^2}. \tag{5.24}$$

The analog of Proposition 5.5 in the continuous limit is

$$\operatorname{Tr} D_x D_{x'} = \pi\delta^2(x - x') \tag{5.25}$$

which leads to the infinite-dimensional Bloch representation

$$\rho = \frac{1}{\pi}\int [d^2x]\, y_x D_x^\dagger \quad , \quad y_x = \operatorname{Tr}(D_x\rho)\,. \tag{5.26}$$

D_x represents a shift in phase plane by x. The commutation relation Equation (5.18) becomes

$$D_{x'}D_x = e^{i2x^T J x'} D_x D_{x'}, \tag{5.27}$$

where J is the *symplectic form*, defined by

$$J = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}. \tag{5.28}$$

If we have n bosonic modes, the phase space variable becomes

$$\vec{x} = (\vec{q}, \vec{p}) = (q_1, \ldots, q_n, p_1, \ldots, p_n). \tag{5.29}$$

Gaussian operations

Gaussian unitaries can be defined in three equivalent ways:

- They are the *normalizer* of the Heisenberg-Weyl group. That is, they conjugate displacement operators to themselves.
- They are symplectic transformations of phase space. (See Equation (5.31).)
- They are generated by Hamiltonians quadratic in position and momentum operators.

A matrix $S \in \mathbb{R}^{2n\times 2n}$ is *symplectic* if

$$SJS^T = J \quad , \quad J = \begin{pmatrix} 0 & I_n \\ -I_n & 0 \end{pmatrix}. \tag{5.30}$$

J is the symplectic form. Denote the set of symplectic matrices $\mathrm{Sp}(2n)$.

For every symplectic $S \in \mathrm{Sp}(2n)$, there is a corresponding Gaussian unitary U_S on n bosonic modes. This acts on the position and momentum operators via

$$U_S^\dagger \vec{x} U_S = S\vec{x}. \tag{5.31}$$

The adjoint action of a Gaussian unitary on the displacement operators is

$$U_S^\dagger D(\vec{x}) U_S = D(S^{-1}\vec{x}). \tag{5.32}$$

Note that for a single mode, the symplectic condition $SJS^T = J$ simply becomes $\det S = 1$. Gaussian unitaries are the infinite-dimensional version of Clifford operations.

5.2 Learning displacement amplitudes

Definition 5.3. *The* displacement amplitudes *of a d-dimensional state ρ are*

$$y_{q,p} = \mathrm{Tr}\left(D_{q,p}\rho\right). \tag{5.33}$$

Note however that the coefficients $y_{q,p}$ are now *complex* in general, since $D_{q,p}$ are not always Hermitian. The Hermiticity of ρ is reflected in the relation

$$y_{q,p}^* = y_{-q,-p}. \tag{5.34}$$

Given copies of a d-dimensional state and its conjugate $\rho \otimes \rho^*$, the first phase of our learning algorithm will aim to learn the displacement amplitudes $\{\pm y_{q,p}\}$ up to a sign $\pm$. The second phase will use entangled measurements across a few more copies to resolve the signs.

Estimating displacement amplitudes up to a sign

Using commutation relation Equation (5.18), it can be checked that the following operators mutually commute when acting on two systems $\mathbb{C}^d \otimes \mathbb{C}^d$:

$$\{D_{q,p} \otimes D_{-q,p}\}. \tag{5.35}$$

This suggests that we can measure two systems in the joint eigenbasis of these operators. If we placed $\rho \otimes \rho^*$ in the two systems for some d-dimensional state ρ, we would have

$$\begin{aligned} \operatorname{Tr}\left((D_{q,p} \otimes D_{-q,p})(\rho \otimes \rho^*)\right) &= \operatorname{Tr}\left(D_{q,p}\rho\right) \operatorname{Tr}\left(D_{q,p}^T \rho^*\right) \\ &= \operatorname{Tr}\left(D_{q,p}\rho\right)^2 = y_{q,p}^2. \end{aligned} \tag{5.36}$$

Here we used Proposition 5.4.

We now construct the desired eigenbasis. It will form a d-dimensional generalization of the Bell basis on 2 qubits. Define

$$|\Phi_{0,0}\rangle = \frac{1}{\sqrt{d}} \sum_j |j\rangle|-j\rangle \tag{5.37}$$

$$\begin{aligned} |\Phi_{a,b}\rangle &= (X^a Z^b \otimes \mathbb{1})|\Phi_{0,0}\rangle \\ &= \frac{1}{\sqrt{d}} \sum_j e^{i2\pi bj/d} |j+a\rangle|-j\rangle. \end{aligned} \tag{5.38}$$

It can be checked that $\{|\Phi_{a,b}\rangle\}$ forms an orthonormal basis of $\mathbb{C}^d \otimes \mathbb{C}^d$. Furthermore, we can calculate

$$(D_{q,p} \otimes D_{-q,p})|\Phi_{a,b}\rangle = e^{i2\pi(ap-bq)/d}|\Phi_{a,b}\rangle. \tag{5.39}$$

Algorithm 2 Learning displacement amplitudes up to a sign.

Input:

- A list of M displacement indices $\mathcal{S} \in \mathbb{Z}_d^2$, $|\mathcal{S}| = M$.
- Precision ε.
- N copies of $\rho \otimes \rho^*$, where ρ is an unknown quantum state in d dimensions.

Output: Estimates $\hat{u}_{q,p}$ for each $(q,p) \in \mathcal{S}$ satisfying $|\pm \hat{u}_{q,p} - y_{q,p}| \leq \varepsilon$ for one of the choices of sign $\pm$.

Algorithm:

1. Measure each copy of $\rho \otimes \rho^*$ in the basis $\{|\Phi_{a,b}\rangle\}$, receiving outcomes $\{(a^{(k)}, b^{(k)})\}_{k=1}^N$.
2. Now given $(q,p) \in \mathcal{S}$, compute

$$\hat{v}_{q,p} = \frac{1}{N} \sum_{k=1}^{N} \exp\left(i2\pi(a^{(k)}p - b^{(k)}q)/d\right). \tag{5.40}$$

3. If $|\hat{v}_{q,p}| \leq \frac{2}{3}\varepsilon^2$, output $\hat{u}_{q,p} = 0$.
4. Else output $\hat{u}_{q,p} = \sqrt{\hat{v}_{q,p}}$, choosing the primary square root without loss of generality.

Theorem 5.6. *Algorithm 2 succeeds with high probability using $N = O(\log M/\varepsilon^4)$ copies of $\rho \otimes \rho^*$.*

Proof. Suppose we measure $\rho \otimes \rho^*$ in the basis $\{|\Phi_{a,b}\rangle\}$, and get the distribution $(a,b) \sim \mathcal{P}$. From Equation (5.36) and Equation (5.39), we can construct an estimator for $y_{q,p}^2$.

$$\begin{aligned} y_{q,p}^2 &= \mathrm{Tr}\left((D_{q,p} \otimes D_{-q,p})(\rho \otimes \rho^*)\right) \\ &= \mathbb{E}_{(a,b)\sim\mathcal{P}} \ \mathrm{Tr}\left((D_{q,p} \otimes D_{-q,p})|\Phi_{a,b}\rangle\langle\Phi_{a,b}|\right) \\ &= \mathbb{E}_{(a,b)\sim\mathcal{P}} \ \exp\left(i2\pi(ap - bq)/d\right). \end{aligned} \tag{5.41}$$

Applying Hoeffding's inequality in the complex plane tells us that with $N = O(\log M/\varepsilon^4)$ copies,

$$\left|\hat{v}_{q,p} - y_{q,p}^2\right| \leq \frac{1}{3}\varepsilon^2 \tag{5.42}$$

for any M displacement operators $D_{q,p}$ with high probability.

Case 1: $|\hat{v}_{q,p}| \leq \frac{2}{3}\varepsilon^2$. Then with high probability

$$|y_{q,p}^2| \leq |\hat{v}_{q,p}| + |\hat{v}_{q,p} - y_{q,p}^2| \leq \frac{2}{3}\varepsilon^2 + \frac{1}{3}\varepsilon^2 = \varepsilon^2 \implies |y_{q,p}| \leq \varepsilon \tag{5.43}$$

and the estimate $\hat{u}_{q,p} = 0$ satisfies $|\hat{u}_{q,p} - y_{q,p}| \leq \varepsilon$.

Case 2: $|\hat{v}_{q,p}| > \frac{2}{3}\varepsilon^2$. Then $|\hat{u}_{q,p}| > \frac{\sqrt{2}}{\sqrt{3}}\varepsilon$. Negating $\hat{u}_{q,p}$ if necessary, suppose we choose the correct hemisphere for $\hat{u}_{q,p}$. This guarantees that

$$|\hat{u}_{q,p} + y_{q,p}| \geq |\hat{u}_{q,p}| > \frac{\sqrt{2}}{\sqrt{3}}\varepsilon. \tag{5.44}$$

Now

$$\frac{1}{3}\varepsilon^2 \geq |\hat{v}_{q,p} - y_{q,p}^2| \tag{5.45}$$

$$= |\hat{u}_{q,p}^2 - y_{q,p}^2| = |\hat{u}_{q,p} + y_{q,p}| \cdot |\hat{u}_{q,p} - y_{q,p}| \tag{5.46}$$

$$\geq \frac{\sqrt{2}}{\sqrt{3}}\,\varepsilon \cdot |\hat{u}_{q,p} - y_{q,p}| \tag{5.47}$$

$$\implies |\hat{u}_{q,p} - y_{q,p}| \leq \frac{\sqrt{2}}{2\sqrt{3}}\varepsilon \tag{5.48}$$

finishing the proof. □

Implementing the measurements

Proposition 5.7. *If we encode the d-dimensional system in $O(\log d)$ qubits on a quantum computer, then we can measure in the basis $\{|\Phi_{a,b}\rangle\}$ in time poly$\log d$.*

Proof. The *quantum Fourier transform* W in d dimensions is the map

$$W : |b\rangle \to \frac{1}{\sqrt{d}} \sum_j e^{-i2\pi bj/d} |j\rangle \tag{5.49}$$

with matrix representation

$$W = \frac{1}{\sqrt{d}} \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & \omega^{-1} & \omega^{-2} & \dots & \omega^{-(d-1)} \\ 1 & \omega^{-2} & \omega^{-4} & \dots & \omega^{-2(d-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{-(d-1)} & \omega^{-2(d-1)} & \dots & \omega^{-(d-1)^2} \end{pmatrix} \tag{5.50}$$

W is a unitary. In relation to X, Z, it has the properties

$$WXW^\dagger = Z \tag{5.51}$$
$$WZW^\dagger = X^\dagger. \tag{5.52}$$

If we encode the d-dimensional system in $O(\log d)$ qubits, then W can be implemented in time polylog d.

Define the *controlled shift* operator CX on two d-dimensional systems to act by

$$CX : |j\rangle|l\rangle \to |j + l\rangle|l\rangle. \tag{5.53}$$

W and CX together can be used to transform the standard product basis on two d-dimensional systems to the entangled $\{|\Phi_{a,b}\rangle\}$ basis.

$$(CX)^{-1} \cdot (\mathbb{1} \otimes W) \cdot |a\rangle|b\rangle = |\Phi_{a,b}\rangle \tag{5.54}$$

This lets us implement a measurement in the basis $\{|\Phi_{a,b}\rangle\}$ by first inverting this transformation and then measuring in the standard product basis. □

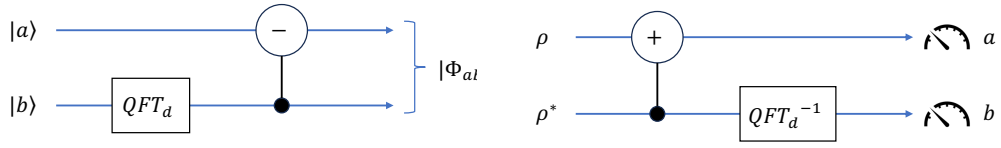

(a) Transforming between bases

(b) Algorithm 2

Figure 5.2

n subsystems

Suppose now we have n d-dimensional subsystems, with Hilbert space $(\mathbb{C}^d)^{\otimes n}$. We can promote q, p, a, b to vectors $\vec{q}, \vec{p}, \vec{a}, \vec{b} \in \mathbb{Z}_d^n$ and define

$$D_{\vec{q},\vec{p}} = D_{q_1,p_1} \otimes \cdots \otimes D_{q_n,p_n} \tag{5.55}$$
$$|\Phi_{\vec{a},\vec{b}}\rangle = |\Phi_{a_1,b_1}\rangle \otimes \cdots \otimes |\Phi_{a_n,b_n}\rangle. \tag{5.56}$$

All statements in Proposition 5.4 hold with q, p replaced with the vectors $\vec{q}, \vec{p}$. In the commutation relation Equation (5.18), the products $qp', q'p$ are to be replaced with dot products $\vec{q} \cdot \vec{p'}$, $\vec{q'} \cdot \vec{p}$.

Following Proposition 5.5, $\{D_{\vec{q},\vec{p}}\}$ form a basis of the $(d^n \times d^n)$-dimensional space of operators on $(\mathbb{C}^d)^{\otimes n}$, orthogonal in the Hilbert-Schmidt inner product.

$$\mathrm{Tr}\left(D_{\vec{q},\vec{p}} D_{\vec{q}',\vec{p}'}\right) = d \cdot \delta_{q_1,q_1'} \delta_{p_1,p_1'} \ldots \delta_{q_n,q_n'} \delta_{p_n,p_n'}. \tag{5.57}$$

We have the Bloch vector decomposition

$$\rho = \frac{1}{d^n}\left(\mathbb{1} + \sum_{\vec{q},\vec{p}} y_{\vec{q},\vec{p}} D^{\dagger}_{\vec{q},\vec{p}}\right) \quad , \quad y_{\vec{q},\vec{p}} = \mathrm{Tr}\left(D(\vec{q},\vec{p})\rho\right). \tag{5.58}$$

The analog of Equation (5.41) is

$$y^2_{\vec{q},\vec{p}} = \mathbb{E}_{(\vec{a},\vec{b})} \exp\left(i2\pi(\vec{a}\cdot\vec{p} - \vec{b}\cdot\vec{q})/d\right). \tag{5.59}$$

Theorem 5.8. *The natural generalization of Algorithm 2 to n qudits has the same guarantee as in Theorem 5.6, and likewise the measurements can be implemented in time poly*log *d.*

Learning algorithm in infinite dimensions

Using Gaussian unitaries, we can phrase Algorithm 2 in the infinite-dimensional setting. Step 1 of Algorithm 2 is to measure $\rho \otimes \rho^*$ in the generalized Bell basis $\{|\Phi_{a,b}\rangle\}$. This is achieved by applying a controlled shift Gaussian operation, followed by *homodyne measurements* along certain quadratures. Let the position and momentum variables of the first and second register be q_1, p_1, q_2, p_2.

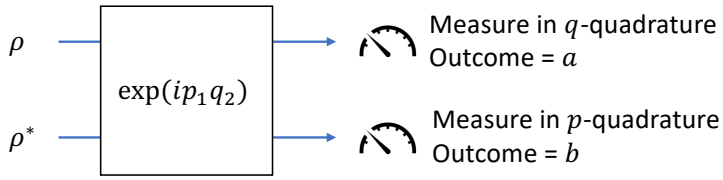

Figure 5.3: Algorithm 2 in infinite dimensions

On two modes, the controlled shift corresponds to the Gaussian unitary

$$CX = \exp ip_1 q_2. \tag{5.60}$$

It has symplectic matrix

$$CX = U_S \quad , \quad S = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \tag{5.61}$$

In the Heisenberg picture, the procedure shown in Figure 5.3 transforms the variables q_1, p_2 to the observables

$$a = q_1 + q_2 \quad , \quad b = -p_1 + p_2. \tag{5.62}$$

We finish the procedure by measuring the first register in the q quadrature, and the second register in the p quadrature. Note measuring in the p quadrature is equivalent to doing a quantum Fourier transform, or phase shift, before measuring in the q quadrature.

Define

$$\alpha = (a, b). \tag{5.63}$$

The estimator for the learning algorithm has a nice expression in terms of the symplectic product. The analogue of Equation (5.41) is

$$y_x^2 = \mathbb{E}_\alpha \exp\left(i2\alpha^T J x\right). \tag{5.64}$$

An example of a class of states where the displacement amplitudes are non-trivial are the GKP codestates. These are the states stabilized by two displacement operators D_x and $D_{x'}$ for some x and x' satisfying $x^T J x' = \pi$.

5.3 Sample complexity lower bounds

The aim of this section is to show that the use of entangled measurements across $\rho \otimes \rho^*$ are essential to our learning task. Any strategy which uses only single-copy measurements on ρ and ρ^*, even an adaptive one, necessarily requires exponentially many copies. On the other hand, any strategy which uses only ρ, and does not have access to ρ^*, must necessarily make entangled measurements across $\Omega(1/\varepsilon)$ copies at a time.

The proofs will follow the techniques introduced in [Che+22; Che+21]. The proof idea is as follows. First we identify an ensemble of states $\{\rho_x\}$ indexed by x such that performing the learning task on ρ_x allows one to identify the label x, and in particular distinguish ρ_x from the maximally mixed state. We then aim to show that single-copy measurements, or alternatively measurements without ρ^*, cannot sufficiently distinguish ρ_x from the maximally mixed state. More specifically, the expected distribution of measurement outcomes on ρ_x for a uniformly random x differs only by an exponentially small amount from the distribution of outcomes on the maximally mixed state. Via Le Cam's two-point method [YAL97], this implies

that such measurements require exponential sample complexity in order to perform the distinguishing task. This in turn gives a sample complexity lower bound on the original learning task.

Here we state two lemmas adapted from [Che+22] which will be useful to us.

Lemma 5.9. ([Che+22] Lemma 4.8) *When we only consider the classical outcome of the POVM measurement and neglect the post-measurement quantum state, then any POVM can be simulated by a rank-1 POVM with some postprocessing. A rank-1 POVM is defined by* $\{w_s, |\psi_s\rangle\}_s$ *where* $w_s > 0$ *and*

$$\sum_s w_s |\psi_s\rangle\langle\psi_s| = \mathbb{1}. \tag{5.65}$$

The outcomes probabilities are

$$\mathbb{P}_\rho(s) = w_s \langle\psi_s|\rho|\psi_s\rangle. \tag{5.66}$$

This lemma lets us consider only rank-1 POVMs without loss of generality.

Lemma 5.10. ([Che+22] Lemma 5.4) *Suppose we have a many vs one distinguishing task*

- (YES) $\rho = \rho_x$ *for some random* x *and some family of states* $\{\rho_x\}$.
- (NO) $\rho = \frac{1}{d}\mathbb{1}$ *maximally mixed.*

If all outcomes l *of a protocol satisfy*

$$\frac{\mathbb{E}_x \mathbb{P}_{\rho_x}(l)}{\mathbb{P}_{\mathbb{1}/d}(l)} \geq 1 - \delta \tag{5.67}$$

then the probability of success is at most $(1+\delta)/2$. *Note that the outcome* l *refers to the entire history of many measurements, which are possibly adaptive.*

We will also use a set of operators defined in [Asa+16]. These are given by

$$E_{q,p} = \chi D_{q,p} + \chi^* D_{-q,-p} \quad , \quad \chi = \frac{1+i}{2} \tag{5.68}$$

and are known as *displacement observables*.

Proposition 5.11. [Asa+16] *The displacement observables are Hermitian and have properties*

$$E^*_{q,p} = E^T_{q,p} = E_{-q,p} \tag{5.69}$$

$$\|E_{q,p}\|_{op} \leq \sqrt{2} \tag{5.70}$$

$$\mathrm{Tr}\left(E_{q,p} E_{q',p'}\right) = d \cdot \delta_{q,q'} \delta_{p,p'}. \tag{5.71}$$

Single copies of ρ and ρ^*

Theorem 5.12. *Let d be the dimension of the Hilbert space. Any single-copy protocol which learns $|\operatorname{Tr}\left(D_{q,p}\rho\right)|$ to precision ε for all $(q,p) \in \mathbb{Z}_d^2$ with probability $2/3$ requires $\Omega(d/\varepsilon^2)$ copies. This holds even if the protocol has access to single-copy measurements of both ρ and ρ^*.*

Proof. Suppose protocol $\mathcal{A}$ is able to learn $|\operatorname{Tr}\left(D_{q,p}\rho\right)|$ to precision ε for all $(q,p) \in \mathbb{Z}_d^2 \setminus \{(0,0)\}$ with probability $2/3$. Consider the task in Proposition 5.13 below with ε replaced by $3\sqrt{2}\varepsilon$. We will argue that $\mathcal{A}$ is able to succeed at this task with probability $2/3$.

$$\operatorname{Tr}\left(D_{q,p}\rho_{q,p,r}\right) = \frac{3\sqrt{2}r\varepsilon}{d}\operatorname{Tr}\left(D_{q,p}E_{q,p}\right) \tag{5.72}$$

$$= \begin{cases} 3\sqrt{2}r\varepsilon & d \text{ even},\ q = p = d/2 \\ 3\sqrt{2}\chi^* r\varepsilon & \text{otherwise} \end{cases} \tag{5.73}$$

$$\implies \ |\operatorname{Tr}\left(D_{q,p}\rho_{q,p,r}\right)| \geq 3\varepsilon \tag{5.74}$$

using Proposition 5.5. Thus $\mathcal{A}$ can distinguish any $\rho_{q,p,r}$ from the maximally mixed state, which has

$$\operatorname{Tr}\left(D_{q,p}\frac{\mathbb{1}}{d}\right) = 0 \ \forall q, p. \tag{5.75}$$

In this case, Proposition 5.13 gives us a sample complexity lower bound of $\Omega(d/\varepsilon^2)$. □

Proposition 5.13. *Let d be the dimension of the Hilbert space. Consider the task of distinguishing between the following two scenarios:*

- (YES) $\rho = \rho_{q,p,r} = \frac{1}{d}\left(\mathbb{1} + r\varepsilon E_{q,p}\right)$ *for some uniformly random $(q,p) \in \mathbb{Z}_d^2 \setminus \{(0,0)\}$ and some uniformly random sign $r \in \{\pm 1\}$. (We assume $0 < \varepsilon < 1$.)*
- (NO) $\rho = \frac{1}{d}\mathbb{1}$ *maximally mixed.*

Any single-copy protocol requires $\Omega(d/\varepsilon^2)$ copies in order to succeed with probability $2/3$. This holds even if the protocol has access to single-copy measurements of both ρ and ρ^*.

Proof of Theorem 5.12 using Proposition 5.13.

$$\operatorname{Tr}\left(D_{q,p}\rho_{q,p,r}\right) = \frac{3\sqrt{2}r\varepsilon}{d}\operatorname{Tr}\left(D_{q,p}E_{q,p}\right) \tag{5.76}$$

$$= \begin{cases} 3\sqrt{2}r\varepsilon & d \text{ even},\ q = p = d/2 \\ 3\sqrt{2}\chi^* r\varepsilon & \text{otherwise} \end{cases} \tag{5.77}$$

$$\implies \ |\operatorname{Tr}\left(D_{q,p}\rho_{q,p,r}\right)| \geq 3\varepsilon \tag{5.78}$$

$$\operatorname{Tr}\left(D_{q,p}\frac{\mathbb{1}}{d}\right) = 0 \ \forall q, p. \tag{5.79}$$

□

Proof. Suppose single-copy protocol $\mathcal{A}$ uses T copies of ρ or ρ^*, and at step t applies the rank-1 POVM $\{w_s^t, |\psi_s^t\rangle\}_s$. By Lemma 5.9, this is without loss of generality. Note the slight abuse of notation, since the POVM of later steps are allowed to depend on the outcomes of earlier measurements. Suppose the outcome of the measurements are $l = (s_1, \ldots, s_T)$. We have

$$\mathbb{P}_\rho(l) = \prod_{t=1}^{T} w_{s_t}^t \langle\psi_{s_t}^t|\rho^t|\psi_{s_t}^t\rangle, \tag{5.80}$$

where ρ^t is either ρ or ρ^* for each t. Note that

$$\rho_{q,p,r}^* = \frac{1}{d}\left(\mathbb{1} + r\varepsilon E_{-q,p}\right) \tag{5.81}$$

using Equation (5.69). Let

$$\gamma_t = \begin{cases} +1 & \rho^t = \rho \\ -1 & \rho^t = \rho^* \end{cases}. \tag{5.82}$$

The aim is to establish an inequality like Equation (5.67), so that we can invoke

Lemma 5.10

$$\frac{\mathbb{E}_{(q,p)\neq(0,0)}\mathbb{E}_r\mathbb{P}_{\rho_{q,p,r}}(l)}{\mathbb{P}_{\mathbb{1}/d}(l)} = \mathbb{E}_{(q,p)\neq(0,0)}\mathbb{E}_r \prod_{t=1}^{T} \frac{w^t_{s_t} + r\varepsilon w^t_{s_t}\langle\psi^t_{s_t}|E_{\gamma_t q,p}|\psi^t_{s_t}\rangle}{w^t_{s_t}} \tag{5.83}$$

$$= \mathbb{E}_{(q,p)\neq(0,0)}\mathbb{E}_r \exp\Big(\sum_{t=1}^{T}\log\left(1 + r\varepsilon\langle\psi^t_{s_t}|E_{\gamma_t q,p}|\psi^t_{s_t}\rangle\right)\Big) \tag{5.84}$$

$$\geq \exp\Big(\sum_{t=1}^{T}\mathbb{E}_{(q,p)\neq(0,0)}\mathbb{E}_r \log\left(1 + r\varepsilon\langle\psi^t_{s_t}|E_{\gamma_t q,p}|\psi^t_{s_t}\rangle\right)\Big) \tag{5.85}$$

$$= \exp\Big(\sum_{t=1}^{T}\frac{1}{2}\mathbb{E}_{(q,p)\neq(0,0)}\log\left(1 - \varepsilon^2\langle\psi^t_{s_t}|E_{\gamma_t q,p}|\psi^t_{s_t}\rangle^2\right)\Big) \tag{5.86}$$

$$\geq \exp\Big(-\sum_{t=1}^{T}\varepsilon^2\mathbb{E}_{(q,p)\neq(0,0)}\langle\psi^t_{s_t}|E_{\gamma_t q,p}|\psi^t_{s_t}\rangle^2\Big) \tag{5.87}$$

$$\geq \exp(-T\varepsilon^2\Gamma) \tag{5.88}$$

$$\geq 1 - T\varepsilon^2\Gamma, \tag{5.89}$$

where

$$\Gamma = \sup_{|\psi\rangle}\ \mathbb{E}_{(q,p)\neq(0,0)}\langle\psi|E_{\pm q,p}|\psi\rangle^2. \tag{5.90}$$

In Equation (5.85) we used Jensen's inequality. In Equation (5.87) we used $\log 1 - x \geq -2x\ \forall x \in [0, 0.79]$, which is valid as long as $\varepsilon \leq 0.62$ by Equation (5.70).

It remains to upper bound Γ. At this point, it is clear that we can drop the $\pm$ coming from the use of the conjugate state, since it is averaged over all q, p. We can express Γ as

$$\Gamma = \sup_{|\psi\rangle}\langle\psi|\langle\psi|\mathbb{E}_{(q,p)\neq(0,0)}(E_{q,p}\otimes E_{q,p})|\psi\rangle|\psi\rangle. \tag{5.91}$$

It can be checked using Equation (5.68) that

$$E_{q,p}\otimes E_{q,p} + E_{-q,-p}\otimes E_{-q,-p} = D_{q,p}\otimes D_{-q,-p} + D_{-q,-p}\otimes D_{q,p} \tag{5.92}$$

and thus we can write

$$\mathbb{E}_{(q,p)\neq(0,0)}(E_{q,p}\otimes E_{q,p}) = \mathbb{E}_{(q,p)\neq(0,0)}(D_{q,p}\otimes D_{-q,-p}). \tag{5.93}$$

It can be checked that

$$\sum_{q,p}(D_{q,p}\otimes D_{-q,-p}) = d\cdot \mathrm{SWAP} \tag{5.94}$$

and so

$$\mathbb{E}_{(q,p)\neq(0,0)}(D_{q,p}\otimes D_{-q,-p}) = \frac{d}{d^2-1}\cdot \mathrm{SWAP} - \frac{1}{d^2-1}\cdot \mathbb{1}. \tag{5.95}$$

Putting this all together, we get

$$\Gamma = \frac{1}{d^2-1}\sup_{|\psi\rangle}\langle\psi|\langle\psi|(d\cdot \mathrm{SWAP}-\mathbb{1})|\psi\rangle|\psi\rangle = \frac{d-1}{d^2-1} = \frac{1}{d+1}. \tag{5.96}$$

Returning to Equation (5.89), we have

$$\frac{\mathbb{E}_{(q,p)\neq(0,0)}\mathbb{E}_r\mathbb{P}_{\rho_{q,p,r}}(l)}{\mathbb{P}_{\mathbb{1}/d}(l)} \geq 1 - \frac{T\varepsilon^2}{d+1}. \tag{5.97}$$

By Lemma 5.10, our single-copy protocol $\mathcal{A}$ succeeds with probability at most $(1+T\varepsilon^2/(d+1))/2$. This completes the proof of Proposition 5.13: to succeed with probability $2/3$, $\mathcal{A}$ requires $T=\Omega(d/\varepsilon^2)$. □

Entangled measurements without ρ^*

Theorem 5.14. *Let d be the dimension of the Hilbert space. Assume d is prime. Any protocol which learns $|\operatorname{Tr}\left(D_{q,p}\rho\right)|$ to precision ε for all $(q,p)\in\mathbb{Z}_d^2$ with probability $2/3$ by measuring copies of $\rho^{\otimes K}$ for $K\leq 1/(12\varepsilon)$ requires $\Omega(\sqrt{d}/(K^2\varepsilon^2))$ measurements.*

Proof. Suppose protocol $\mathcal{A}$ is able to learn $|\operatorname{Tr}\left(D_{q,p}\rho\right)|$ to precision ε for all $(q,p)\in\mathbb{Z}_d^2\setminus\{(0,0)\}$ with probability $2/3$. Consider the task in Proposition 5.15 below with ε replaced by $3\sqrt{2}\varepsilon$. We will argue that $\mathcal{A}$ is able to succeed at this task with probability $2/3$.

$$\operatorname{Tr}\left(D_{q,p}\rho_{q,p,r}\right) = \frac{3\sqrt{2}r\varepsilon}{d}\operatorname{Tr}\left(D_{q,p}E_{q,p}\right) \tag{5.98}$$

$$= \begin{cases} 3\sqrt{2}r\varepsilon & d \text{ even},\ q=p=d/2 \\ 3\sqrt{2}\chi^* r\varepsilon & \text{otherwise}\end{cases} \tag{5.99}$$

$$\implies |\operatorname{Tr}\left(D_{q,p}\rho_{q,p,r}\right)| \geq 3\varepsilon \tag{5.100}$$

using Proposition 5.5. Thus $\mathcal{A}$ can distinguish any $\rho_{q,p,r}$ from the maximally mixed state, which has

$$\operatorname{Tr}\left(D_{q,p}\frac{\mathbb{1}}{d}\right) = 0\ \forall q,p. \tag{5.101}$$

In this case, Proposition 5.15 says that $\mathcal{A}$ requires either entangled measurements across at least $1/2\sqrt{2} \cdot (3\sqrt{2}\varepsilon) = 1/(12\varepsilon)$ copies of ρ at a time, or $\Omega(\sqrt{d}/(K^2\varepsilon^2))$ total copies. □

Proposition 5.15. *Let d be the dimension of the Hilbert space. Assume d is prime. Consider the task of distinguishing between the following two scenarios:*

- (YES) $\rho = \rho_{q,p,r} = \frac{1}{d}\left(\mathbb{1} + r\varepsilon E_{q,p}\right)$ *for some uniformly random* $(q, p) \in \mathbb{Z}_d^2 \setminus \{(0,0)\}$ *and some uniformly random sign* $r \in \{\pm 1\}$. *(We assume* $0 < \varepsilon < 1$.*)*
- (NO) $\rho = \frac{1}{d}\mathbb{1}$ *maximally mixed.*

Any protocol succeeding with probability $2/3$ *which measures copies of* $\rho^{\otimes K}$ *for* $K \leq 1/(2\sqrt{2}\varepsilon)$ *requires* $\Omega(\sqrt{d}/(K^2\varepsilon^2))$ *measurements.*

Proof of Theorem 5.14 using Proposition 5.15.

$$\operatorname{Tr}\left(D_{q,p}\rho_{q,p,r}\right) = \frac{3\sqrt{2}r\varepsilon}{d}\operatorname{Tr}\left(D_{q,p}E_{q,p}\right) \tag{5.102}$$

$$= \begin{cases} 3\sqrt{2}r\varepsilon & d \text{ even, } q = p = d/2 \\ 3\sqrt{2}\chi^* r\varepsilon & \text{otherwise} \end{cases} \tag{5.103}$$

$$\implies \quad |\operatorname{Tr}\left(D_{q,p}\rho_{q,p,r}\right)| \geq 3\varepsilon \tag{5.104}$$

$$\operatorname{Tr}\left(D_{q,p}\frac{\mathbb{1}}{d}\right) = 0 \;\forall q, p \tag{5.105}$$

□

Proof. Suppose protocol $\mathcal{A}$ measures K copies of ρ at a time, where $K \leq 1/(2\sqrt{2}\varepsilon)$. At step t, $\mathcal{A}$ applies the rank-1 POVM $\{w_s^t, |\psi_s^t\rangle\}_s$, where t goes from 1 up to T. By Lemma 5.9, this is without loss of generality. Note the slight abuse of notation, since the POVM of later steps are allowed to depend on the outcomes of earlier measurements. Suppose the outcome of the measurements are $l = (s_1, \ldots, s_T)$. We have

$$\mathbb{P}_\rho(l) = \prod_{t=1}^{T} w_{s_t}^t \langle\psi_{s_t}^t|\rho^{\otimes K}|\psi_{s_t}^t\rangle. \tag{5.106}$$

The aim is to establish an inequality like Equation (5.67), so that we can invoke Lemma 5.10

$$\frac{\mathbb{E}_{(q,p)\neq(0,0)}\mathbb{E}_r\mathbb{P}_{\rho_{q,p,r}}(l)}{\mathbb{P}_{\mathbb{1}/d}(l)} = \mathbb{E}_{(q,p)\neq(0,0)}\mathbb{E}_r \prod_{t=1}^{T} \langle\psi^t_{s_t}|\big(\mathbb{1} + r\varepsilon E_{q,p}\big)^{\otimes K}|\psi^t_{s_t}\rangle \tag{5.107}$$

$$= \mathbb{E}_{(q,p)\neq(0,0)}\mathbb{E}_r \prod_{t=1}^{T} F^t_{q,p,r} \tag{5.108}$$

$$= \mathbb{E}_{(q,p)\neq(0,0)}\mathbb{E}_r \exp\Big(\sum_{t=1}^{T} \log F^t_{q,p,r}\Big) \tag{5.109}$$

$$\geq \exp\Big(\sum_{t=1}^{T} \mathbb{E}_{(q,p)\neq(0,0)}\mathbb{E}_r \log F^t_{q,p,r}\Big) \tag{5.110}$$

$$= \exp\Big(\sum_{t=1}^{T} \frac{1}{2}\mathbb{E}_{(q,p)\neq(0,0)} \log\big(F^t_{q,p,+1}\cdot F^t_{q,p,-1}\big)\Big) \tag{5.111}$$

$$= \exp\Big(\sum_{t=1}^{T} \frac{1}{2}\mathbb{E}_{(q,p)\neq(0,0)} \log\big(1 - G^t_{q,p}\big)\Big) \tag{5.112}$$

$$\geq \exp\Big(\sum_{t=1}^{T} \frac{1}{2}\mathbb{E}_{(q,p)\neq(0,0)} \log\big(1 - \max(0, G^t_{q,p})\big)\Big) \tag{5.113}$$

$$\geq \exp\Big(-\sum_{t=1}^{T} \mathbb{E}_{(q,p)\neq(0,0)} \max(0, G^t_{q,p})\Big) \tag{5.114}$$

$$\geq 1 - \sum_{t=1}^{T} \mathbb{E}_{(q,p)\neq(0,0)} \max(0, G^t_{q,p}), \tag{5.115}$$

where

$$F^t_{q,p,r} := \langle\psi^t_{s_t}|\big(\mathbb{1} + r\varepsilon E_{q,p}\big)^{\otimes K}|\psi^t_{s_t}\rangle \tag{5.116}$$

$$G^t_{q,p} := 1 - F^t_{q,p,+1}\cdot F^t_{q,p,-1}. \tag{5.117}$$

In Equation (5.110) we used Jensen's inequality. In Equation (5.114) we used $\log 1 - x \geq -2x\ \forall x \in [0, 0.79]$, which is valid as long as $G^t_{q,p} \leq 0.79$. This is guaranteed by the assumption $K \leq 1/(2\sqrt{2}\varepsilon)$, since

$$F^t_{q,p,r} = \langle\psi|(I + r\varepsilon E_{q,p})^{\otimes K}|\psi\rangle \tag{5.118}$$

$$\geq (1 - \varepsilon\sqrt{2})^K \tag{5.119}$$

$$\geq 1 - \sqrt{2}K\varepsilon \qquad \text{(provided } \varepsilon \leq 1/\sqrt{2}) \tag{5.120}$$

$$\geq 1/2 \qquad \text{for the choice of } K, \tag{5.121}$$

so

$$G^t_{q,p} = 1 - F^t_{q,p,+1} \cdot F^t_{q,p,-1} \leq 3/4. \tag{5.122}$$

We would like to upper bound $\mathbb{E}_{(q,p)\neq(0,0)} \max(0, G^t_{q,p})$. Let's calculate

$$\begin{aligned}
G^t_{q,p} &= 1 - \langle\psi^t_{s_t}|\big(\mathbb{1} + \varepsilon E_{q,p}\big)^{\otimes K}|\psi^t_{s_t}\rangle\langle\psi^t_{s_t}|\big(\mathbb{1} - \varepsilon E_{q,p}\big)^{\otimes K}|\psi^t_{s_t}\rangle && (5.123)\\
&= 1 - \sum_{S,S'\subset[K]} (-1)^{|S'|}\varepsilon^{|S|+|S'|}\langle\psi^t_{s_t}|E^{\otimes S}_{q,p}|\psi^t_{s_t}\rangle\langle\psi^t_{s_t}|E^{\otimes S'}_{q,p}|\psi^t_{s_t}\rangle && (5.124)\\
&= 1 - \langle\psi^t_{s_t}|\big(\mathbb{1} + H^0_{q,p}\big)|\psi^t_{s_t}\rangle^2 + \langle\psi^t_{s_t}|H^1_{q,p}|\psi^t_{s_t}\rangle^2 && (5.125)\\
&= -2\langle\psi^t_{s_t}|H^0_{q,p}|\psi^t_{s_t}\rangle - \langle\psi^t_{s_t}|H^0_{q,p}|\psi^t_{s_t}\rangle^2 + \langle\psi^t_{s_t}|H^1_{q,p}|\psi^t_{s_t}\rangle^2 && (5.126)\\
&\leq -2\langle\psi^t_{s_t}|H^0_{q,p}|\psi^t_{s_t}\rangle + \langle\psi^t_{s_t}|H^1_{q,p}|\psi^t_{s_t}\rangle^2 && (5.127)\\
\implies &\mathbb{E}_{(q,p)\neq(0,0)} \max(0, G^t_{q,p}) && (5.128)\\
&\leq 2\max_{|\psi\rangle} \mathbb{E}_{(q,p)\neq(0,0)}|\langle\psi|H^0_{q,p}|\psi\rangle| + \max_{|\psi\rangle} \mathbb{E}_{(q,p)\neq(0,0)}\langle\psi|H^1_{q,p}|\psi\rangle^2 && (5.129)\\
&\leq 2\sqrt{\Gamma^0} + \Gamma^1, && (5.130)
\end{aligned}$$

where

$$\begin{aligned}
H^0_{q,p} &= \sum_{S\subset[K],|S|\text{even},S\neq\emptyset} \varepsilon^{|S|}E^{\otimes S}_{q,p} && (5.131)\\
H^1_{q,p} &= \sum_{S\subset[K],|S|\text{odd}} \varepsilon^{|S|}E^{\otimes S}_{q,p} && (5.132)\\
\Gamma^0 &= \max_{|\psi\rangle} \mathbb{E}_{(q,p)\neq(0,0)}\langle\psi|H^0_{q,p}|\psi\rangle^2 && (5.133)\\
\Gamma^1 &= \max_{|\psi\rangle} \mathbb{E}_{(q,p)\neq(0,0)}\langle\psi|H^1_{q,p}|\psi\rangle^2 && (5.134)
\end{aligned}$$

and we used Cauchy-Schwarz in the final step. Returning to Equation (5.115), we have

$$\frac{\mathbb{E}_{(q,p)\neq(0,0)}\mathbb{E}_r\mathbb{P}_{\rho_{q,p,r}}(l)}{\mathbb{P}_{\mathbb{1}/d}(l)} \geq 1 - T\big(2\sqrt{\Gamma^0} + \Gamma^1\big). \tag{5.135}$$

It remains to upper bound Γ^0 and Γ^1. Let's first deal with Γ^0.

$$\begin{aligned}
\Gamma^0 &= \max_{|\psi\rangle} \mathbb{E}_{(q,p)\neq(0,0)}\langle\psi|\langle\psi|H^0_{q,p}\otimes H^0_{q,p}|\psi\rangle|\psi\rangle && (5.136)\\
&\leq \left\|\mathbb{E}_{(q,p)\neq(0,0)} H^0_{q,p}\otimes H^0_{q,p}\right\|_{\text{op}} && (5.137)\\
&\leq \sum_{S,S'\subset[K],|S|,|S'|\text{even},S,S'\neq\emptyset} \varepsilon^{|S|+|S'|}\left\|\mathbb{E}_{(q,p)\neq(0,0)} E^{\otimes S\cup S'}_{q,p}\right\|_{\text{op}} && (5.138)\\
&\leq \sum_{4\leq k\leq 2K, k \text{ even}} (2K)^k\varepsilon^k\left\|\mathbb{E}_{(q,p)\neq(0,0)} E^{\otimes k}_{q,p}\right\|_{\text{op}}. && (5.139)
\end{aligned}$$

Lemma 5.17, which is stated and proved at the end of the section, tells us that

$$\left\| \sum_{q,p} E_{q,p}^{\otimes k} \right\|_{\text{op}} \leq 2^{k/2} d \tag{5.140}$$

$$\implies \left\| \mathbb{E}_{(q,p)\neq(0,0)} E_{q,p}^{\otimes k} \right\|_{\text{op}} \leq \frac{2^{k/2} d}{d^2-1} + \frac{1}{d^2-1} \leq \frac{2^{k/2}}{d-1}. \tag{5.141}$$

Plugging this into Equation (5.139) gives

$$\Gamma^0 \leq \frac{1}{d-1} \sum_{4\leq k\leq 2K, k \text{ even}} \left(2\sqrt{2}K\varepsilon\right)^k. \tag{5.142}$$

This is a geometric series, which is dominated by its first term since by assumption $K \leq 1/(2\sqrt{2}\varepsilon)$. We get

$$\Gamma^0 = O(K^4\varepsilon^4/d). \tag{5.143}$$

The calculation for Γ^1 is similar.

$$\Gamma^1 \leq \sum_{2\leq k\leq 2K, k \text{ even}} (2K)^k \varepsilon^k \left\| \mathbb{E}_{(q,p)\neq(0,0)} E_{q,p}^{\otimes k} \right\|_{\text{op}} \tag{5.144}$$

$$= O(K^2\varepsilon^2/d). \tag{5.145}$$

Returning to Equation (5.135), we have

$$\frac{\mathbb{E}_{(q,p)\neq(0,0)} \mathbb{E}_r \mathbb{P}_{\rho_{q,p,r}}(l)}{\mathbb{P}_{\mathbb{1}/d}(l)} \geq 1 - O\left(TK^2\varepsilon^2/\sqrt{d}\right). \tag{5.146}$$

By Lemma 5.10, our protocol $\mathcal{A}$ succeeds with probability at most $(1+O(TK^2\varepsilon^2/\sqrt{d}))/2$. To succeed with probability $2/3$, $\mathcal{A}$ requires $T = \Omega(\sqrt{d}/(K^2\varepsilon^2))$. This completes the proof of Proposition 5.15, modulo Lemma 5.17. □

Lemma 5.16. *Let the Hilbert space dimension d be prime. For any $1 \leq m \leq k$ with k even,*

$$\left\| \sum_{q,p} D_{q,p}^{\otimes m} \otimes D_{-q,-p}^{\otimes(k-m)} \right\|_{op} = d \tag{5.147}$$

Proof. Denote

$$\mathcal{D}(m,k) = \sum_{q,p} D_{q,p}^{\otimes m} \otimes D_{-q,-p}^{\otimes(k-m)} \tag{5.148}$$

Consider the action on a basis state $|a_1\rangle \dots |a_k\rangle$.

$$\mathcal{D}(m,k)|a_1\rangle \dots |a_k\rangle \tag{5.149}$$

$$= \sum_{q,p} e^{i\pi kqp/d} \left((X^q)^{\otimes m} \otimes (X^{-q})^{\otimes(k-m)}\right)\left((Z^p)^{\otimes m} \otimes (Z^{-p})^{\otimes(k-m)}\right)|a_1\rangle \dots |a_k\rangle \tag{5.150}$$

$$= \sum_{q,p} e^{i(2\pi/d)p\left(\frac{k}{2}q + a_1 + \dots + a_m - a_{m+1} - \dots - a_k\right)}|a_1+q\rangle \dots |a_m+q\rangle|a_{m+1}-q\rangle \dots |a_k+q\rangle \tag{5.151}$$

$$= d|a_1+\hat{q}\rangle \dots |a_m+\hat{q}\rangle|a_{m+1}-\hat{q}\rangle \dots |a_k-\hat{q}\rangle, \tag{5.152}$$

where $\hat{q}$ is the unique solution to

$$\frac{k}{2}q + a_1 + \dots + a_m - a_{m+1} - \dots - a_k = 0 \mod d. \tag{5.153}$$

Here we used that d is prime, so that $\mathbb{Z}_d$ is a field.

Let g be the multiplicative inverse of $k/2$ mod d. The operator $\mathcal{D}(m,k)/d$ implements a linear map on basis vectors over $\mathbb{Z}_d^k$ given by the matrix

$$I_k + \begin{pmatrix} g & \dots & g & -g & \dots & -g \\ \vdots & & \vdots & \vdots & & \vdots \\ g & \dots & g & -g & \dots & -g \end{pmatrix}. \tag{5.154}$$

This matrix is invertible over $\mathbb{Z}_d^k$ with inverse

$$I_k + \begin{pmatrix} h & \dots & h & -h & \dots & -h \\ \vdots & & \vdots & \vdots & & \vdots \\ h & \dots & h & -h & \dots & -h \end{pmatrix}, \tag{5.155}$$

where h solves

$$g + h + (2m-k)gh = 0. \tag{5.156}$$

We have shown that $\mathcal{D}(m,k)/d$ is in fact a permutation matrix. In particular,

$$||\mathcal{D}(m,k)||_{\text{op}} = d. \tag{5.157}$$

□

Lemma 5.17. *Let the Hilbert space dimension d be prime. For any even k,*

$$\left\| \sum_{q,p} E_{q,p}^{\otimes k} \right\|_{op} \leq 2^{k/2} d \tag{5.158}$$

Proof. First expand

$$E_{q,p}^{\otimes k} = \left(\chi D_{q,p} + \chi^* D_{-q,-p}\right)^{\otimes k} \tag{5.159}$$

$$= \sum_{S \subset [k]} \chi^{|S|} (\chi^*)^{k-|S|} D_{q,p}^{\otimes S} \otimes D_{-q,-p}^{\otimes [k]\setminus S} \tag{5.160}$$

Now sum over q, p and take the operator norm.

$$\left\| \sum_{q,p} E_{q,p}^{\otimes k} \right\|_{\text{op}} \leq \frac{1}{2^{k/2}} \sum_{S \subset [k]} \left\| \sum_{q,p} D_{q,p}^{\otimes S} \otimes D_{-q,-p}^{\otimes [k]\setminus S} \right\|_{\text{op}} \tag{5.161}$$

$$= 2^{k/2} d \tag{5.162}$$

using Lemma 5.16. □

We suspect Lemma 5.17 is *not* tight, and the correct bound is $2d$ independent of k, but it is sufficient for our purposes.

Chapter 6

TOPOLOGICAL DATA ANALYSIS

In quantum complexity theory, the goal is to characterize the capacity of quantum computers for solving computational problems. There is no requirement that the problems considered be quantum mechanical in nature, but most of the key results in the field do consider such problems. See [Osb12; Gha+15] for reviews. The field has been important both for understanding the potential of quantum computers and for giving insight into fundamental physics. However, the capacity of quantum computers to solve problems that do not appear inherently quantum mechanical is less well understood.

In this chapter, we examine a classic problem in computational topology and find that it can be characterized by quantum complexity classes. Topology studies properties of spaces that only depend on the continuity and connectivity between points and do not depend on distances between points. In particular, we are concerned with *homology*, a branch of topology that describes k-dimensional holes in a topological space. For example, a circle constitutes a 1-dimensional hole, a hollow sphere a 2-dimensional hole, and so on. The types of spaces we consider are motivated by the practical application of topological data analysis (TDA). Our result thus has implications for quantum advantage in TDA and the quantum TDA algorithm introduced in [LGZ16]. In TDA, one applies techniques from topology to extract global information from data in a way that is resistant to local noise. For background on TDA and its applications, see [Was18; Pet+14; GGB16; Rei+17].

Results

The computational problem we consider is perhaps the simplest question in homology—given some space, does it have a k-dimensional hole or not? Our main result is that, for certain types of input spaces and a suitable promise on the gap, this problem is QMA_1-hard and contained in QMA. The complexity class QMA is the quantum analogue of the class MA, and QMA_1 is a one-sided error version of QMA. (See Chapter 6 for full definitions.)

A problem of interest in TDA is to compute the number of holes in the *clique complex* of a graph—the clique complex of a graph is the simplicial complex

formed by mapping every $k+1$-clique in the graph to a k-simplex. (See Chapter 6 for formal definition.) We consider the problem of deciding whether a weighted graph's clique complex has a k-dimensional hole or not. The decision version of the clique homology problem was shown to be QMA_1-hard in [CK22]. However, the problem is only believed to be inside QMA when the combinatorial Laplacian has a promise gap. Our key contribution is showing that the clique homology problem remains QMA_1-hard when the gap is imposed. Thus we can provide upper and lower bounds on the complexity of this classical problem in terms of quantum complexity classes.

Problem 6.1. (Gapped clique homology) *Fix functions $k : \mathbb{N} \to \mathbb{N}$ and $g : \mathbb{N} \to [0, \infty)$, with $g(n) \geq 1/\mathrm{polyn}$, $k(n) \leq n$. The input to the problem is a vertex-weighted graph $\mathcal{G}$ on n vertices. The task is to decide whether:*

- ***YES*** *The $k(n)^{th}$ homology group of $Cl(\mathcal{G})$ is non-trivial $H_k(\mathcal{G}) \neq 0$.*
- ***NO*** *The $k(n)^{th}$ homology group of $Cl(\mathcal{G})$ is trivial $H_k(\mathcal{G}) = 0$ and the weighted combinatorial Laplacian Δ_k has minimum eigenvalue $\lambda_{\min}(\Delta_k) \geq g(n)$.*

The condition on the minimum eigenvalue $\lambda_{\min}(\Delta_k)$ can be interpreted as a promise that the graph is far from having a hole in the NO case; see below.

Theorem 6.1. *Problem 6.1 is QMA_1-hard and contained in QMA.*

At this point, we would like to make two clarifying remarks on our main result. First of all, one may wonder why we cannot get tight bounds on the complexity of Problem 6.1; for example, could it be QMA-hard, or contained in QMA_1? We suspect that the true complexity of Problem 6.1 is QMA_1. However, this is a somewhat fragile argument to make since containment in QMA_1 depends on the specific choice of universal gate set for the verifier quantum circuit. Secondly, we would like to clarify the comparison to the previous results of [CK22]. The best known upper bound on the complexity of the decision version of the clique homology problem, as considered in [CK22], is PSPACE; we can draw an analogy to the local Hamiltonian problem without a promise gap, which is PSPACE-complete [FL16a]. Since PSPACE contains QMA and QMA_1, the statement that the decision version is QMA_1-hard does not provide compelling evidence that the clique homology problem has a quantum structure. On the other hand, our results study a related

problem for which the complexity can be upper and lower bounded by quantum complexity classes.

Techniques

The main technical contribution is the development of tools which allow us to lower bound the eigenvalues of the Laplacian operator from Hodge theory. With this purpose, we introduce a new technique to the field—a powerful tool from algebraic topology known as *spectral sequences*. In particular, we exploit a connection between spectral sequences and Hodge theory presented in [For94]. In homology, spectral sequences can play a powerful role analogous to perturbation theory in the analysis of perturbative gadgets [KKR06]. The use of spectral sequences to perform a kind of perturbation theory on the Laplacians of simplicial complexes is novel to quantum information theory. Further, the ability to lower bound the Laplacian eigenvalues may be of independent interest. In addition, our work extends the simplicial surgery technique used in prior quantum complexity work [CK22].

Our hardness proof will proceed by reducing from a particular local Hamiltonian problem. We would like to encode the Hamiltonian problem into the gapped clique homology problem. For this purpose, our main focus will be to establish the following theorem.

Theorem 6.2. (Main theorem) *Given a local Hamiltonian* H on n qubits, we can construct a vertex-weighted graph $\mathcal{G}$ on polyn vertices and a k such that the combinatorial Laplacian $\Delta_k(Cl(\mathcal{G}))$ satisfies*

$$\lambda_{\min}(H) = 0 \implies \lambda_{\min}(\Delta_k(Cl(\mathcal{G}))) = 0 \tag{6.1}$$

$$\lambda_{\min}(H) \geq \frac{1}{polyn} \implies \lambda_{\min}(\Delta_k(Cl(\mathcal{G}))) \geq \frac{1}{polyn} \tag{6.2}$$

** There are some conditions on the form of H, but the class is sufficiently expressive to be QMA_1-hard.*

In the language of the theorem, the bulk of the work is to establish $\lambda_{\min}(\Delta_k(\mathrm{Cl}(\mathcal{G}))) \geq 1/\mathrm{poly}n$ in the case $\lambda_{\min}(H) \geq 1/\mathrm{poly}n$. Here, spectral sequences will be essential.

Spectral gap of Laplacian

How can we interpret the gap at the bottom of the spectrum of the combinatorial Laplacian $\lambda_{\min}(\Delta_k) \geq 1/\mathrm{poly}n$? At $k = 0$, the combinatorial Laplacian Δ^0 is equal

to the usual graph Laplacian L plus the projector onto the constant vector:

$$\Delta^0 = L + \begin{pmatrix} 1 & \dots & 1 \\ \vdots & & \vdots \\ 1 & \dots & 1 \end{pmatrix} \tag{6.3}$$

Thus the smallest eigenvalue of Δ^0 corresponds to the first non-zero eigenvalue of the graph Laplacian L. This is precisely the eigenvalue which controls graph expansion; it appears in the well-known Cheeger inequality which relates the Laplacian spectrum to geometric connectivity of the graph [LP17]. This provides a geometric interpretation of the minimum eigenvalue of Δ^0—it measures how far the graph is from being disconnected. A similar geometric interpretation holds for higher dimensional combinatorial Laplacians Δ_k. Indeed, higher-dimensional Cheeger inequalities have been studied [GS14; SKM14; PRT16], with connections to the field of *high-dimensional expanders* [Lub18]. A large minimum eigenvalue $\lambda_{\min}(\Delta_k)$ means that the graph is far from having a k-dimensional hole.

This leads us to an exciting future direction: Can we use graph operations and gadgets to perform gap amplification on the combinatorial Laplacian? In light of our QMA_1-hardness result, this may have connections to the yet elusive *quantum PCP conjecture* [AAV13].

Implications

Related to deciding the existence of a hole is the problem of computing the normalized number of holes. For this problem there is an efficient quantum algorithm, known as the quantum TDA algorithm. (For a discussion of how this algorithm works, and what precisely we mean by 'normalized number of holes' we refer readers to Chapter 6.) A significant motivation for our work was understanding the complexity of the problem solved by this quantum TDA algorithm and its speedup over classical algorithms [LGZ16; GK19; GCD22; Uba+21; Hay22; MGB22; Ber+24a; Akh+22; SL22; ASS22]. Unlike many other quantum computing applications in machine learning [Tan19; GLT18; Chi+22], the quantum TDA algorithm has resisted *dequantization*, and researchers still debate the presence of a speedup over the best possible classical algorithm.

Our result can be seen as providing some suggestion that the quantum TDA *cannot* be dequantized. We have shown that deciding whether clique complexes have holes is just as hard as deciding if a generic local Hamiltonian is frustration-free. We can likewise translate the problem solved by the quantum TDA algorithm into a problem

phrased in local Hamiltonians. It is known that for generic local Hamiltonians this problem is very unlikely to be tractable on a classical computer. More precisely, it was shown in [GCD22] that this problem is DQC1-hard (see Chapter 6). However, the DQC1-hardness for generic Hamiltonians is inconclusive since some classical algorithms may exist that could exploit some unique structure in clique complexes to outperform algorithms for generic quantum Hamiltonians. Our work suggests that problems involving clique complexes do *not* possess exploitable structure, and are just as hard as the corresponding problem on a general Hamiltonian.

The results in this chapter do not answer these questions conclusively – our reduction does not immediately imply that *all* problems regarding clique complexes are just as hard as the corresponding problem translated into the language of quantum Hamiltonians. Nevertheless, we have developed techniques which are able to lower bound the eigenvalues of the Laplacian operator when reducing from a quantum Hamiltonian. We anticipate this could open up new possibilities in searching for quantum advantage in topological data analysis.

Discussion

The problem of deciding whether or not a space contains a hole makes no explicit reference to quantum mechanics, so it is surprising that its complexity turns out to be characterised by quantum complexity classes. Other examples of this kind are rare. One story to compare to is that of the *Jones polynomial*. Estimating the Jones polynomial, an invariant from knot theory, was shown to be BQP-complete in [AJL06]. That result is an aspect of a deep connection between topological quantum field theory and knot theory. Similarly, our result is an aspect of a deep connection between *supersymmetry* and homology, which was previously explored in [CC21; Cri20; CK22]. Incidentally, both of these connections were explored in the 1980s by Witten [Wit82; Wit89].

These results in quantum complexity suggest that a fruitful avenue for studying the possibility of quantum advantage in seemingly classical problems is to look for 'hidden quantumness'—mathematical problems that, at first glance, do not appear quantum, but can be mapped to specific families of quantum systems. One of the critical areas where quantum computers will offer practical advantanges over their classical counterparts is in studying quantum systems. By looking for such examples of 'hidden quantumness', we may be able to extend the utility of quantum computation into more fields.

Prior work

Understanding the possiblity of quantum advantage in TDA has inspired a number of works in this area in recent years.

The first key result in the field was the quantum TDA algorithm of [LGZ16]. It gives an approximation to the *normalized* k^{th}-Betti number of a simplicial complex. The k^{th}-Betti number β_k is the number of k dimensional holes in the complex, and the normalized Betti number is given by $\beta_k/|S_k|$ where S_k is the set of k-simplices. We can understand the quantum TDA algorithm as running phase estimation on the Laplacian Δ_k. The input state for the phase estimation is the maximally mixed state over k-simplices:

$$\frac{1}{|S_k|}\sum_{x\in S_k}|x\rangle\langle x|. \tag{6.4}$$

The algorithm effectively samples the eigenvalues of the Laplacian Δ_k, to any $1/\text{poly}(n)$ precision. By counting the fraction of times a zero eigenvalue is observed, we get an estimate of $\beta_k/|S_k|$ to any $1/\text{poly}(n)$ additive error. In order for the input state to be prepared efficiently for a clique complex, the graph must be *clique dense* (see [GCD22]).

In [GCD22], the authors initiated the investigation into the complexity of the problem solved by the quantum TDA algorithm. They showed that if one applies the quantum TDA algorithm with a generic local Hamiltonian in the place of Δ_k, then it is able to solve a DQC1-hard problem. [1]

Inspired by the connection between homology and supersymmetry, in [CC21] it was shown that the problem of deciding whether a general chain complex (a generalisation of simplicial complexes) has an k-dimensional hole is QMA_1-hard and contained in QMA (given a suitable promise gap). Moreover, it was shown in the same paper that estimating the normalized Betti numbers of a general chain complex is DQC1-hard and contained in BQP.

Other papers have considered the *decision* version of Problem 6.1:

Problem 6.2. (Decision clique homology) *Let $\mathcal{G}$ be a graph on n vertices, given by its adjacency matrix. We are also given an integer k. The task is to decide*

- ***YES*** *The k^{th} homology group of $Cl(\mathcal{G})$ is non-trivial $H_k(\mathcal{G}) \neq 0$.*

[1] DQC1 is the 'one clean qubit' model of quantum computation where the initial state is limited to a single qubit in the state $|0\rangle$, along with a supply of maximally mixed qubits. (See [Bra08, Section 6.3] for a formal definition.) It does not capture the full power of quantum computation, but is thought to be impossible to simulate efficiently with classical computation.

- ***NO*** *The k^{th} homology group of $Cl(\mathcal{G})$ is trivial $H_k(\mathcal{G}) = 0$.*

In [AS16] it was shown that Problem 6.2 is NP-hard, and in [SL22] it was shown that Problem 6.2 remains NP-hard when restricted to clique dense graphs. These results culminated in [CK22] where it was shown that Problem 6.2 is QMA_1-hard, including when restricted to clique-dense graphs. However, Problem 6.2 is not believed to be inside QMA. Moreover, the constructions used for the reductions in previous works do not satisfy the necessary gap to guarantee containment in QMA. It should be noted the previous result holds for both weighted and unweighted graphs, whereas our results hold for weighted graphs only.

On the more applied side of the field, a number of recent papers have looked more closely at the quantum TDA algorithm. In [Uba+21; Hay22; MGB22; Ber+24a; Akh+22] improvements were made to the algorithm in [LGZ16] which make it more practical to run.

Simplicial homology

The building blocks of simplicial complexes are *simplices*. Simplices can be thought of as the generalisation of triangles and tetrahedron to arbitrary dimensions. A 0-simplex is a point, a 1-simplex is a line, 2- and 3-simplices are triangles and tetrahedra respectively. In higher dimensions a k-simplex is defined as a k-dimensional polytope which is the convex hull of $k+1$-vertices. We can denote a k-simplex by its vertices:

$$\sigma = [x_0 \dots x_k], \tag{6.5}$$

where the x_i are the vertices of the simplex. Simplices are oriented, with the orientation induced by the ordering of the vertices. So permuting the vertices in a simplex leads to an equivalent simplex, possibly up to an overall sign:

$$[x_{\pi(0)} \dots x_{\pi(k)}] = \text{sgn}(\pi)[x_0 \dots x_k]. \tag{6.6}$$

A *simplicial complex* $\mathcal{K}$ is a collection of simplices satisfying two requirements: (A) if a simplex is in $\mathcal{K}$ then all its faces are also in $\mathcal{K}$ and (B) the intersection of any two simplices in $\mathcal{K}$ is a face of both the simplices. Intuitively we can think of constructing a simplicial complex by gluing simplices together along faces.

In general, we can consider linear combinations of k-simplices, known as *k-chains*. The space of k-chains forms a vector space. To define the notion of a hole in a

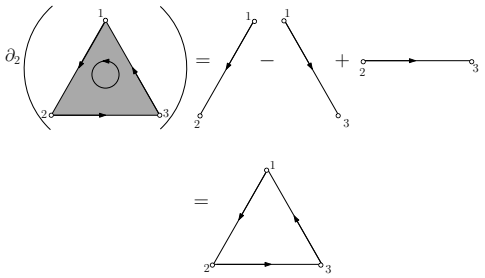

Figure 6.1: The boundary operators action on a 2-simplex (i.e. a triangle).

simplicial complex we need to introduce the boundary operator ∂_k. The boundary operator acts on k-simplices as

$$\partial_k [x_0 \dots x_k] = \sum_{j=0}^{k} (-1)^j [x_0 \dots \hat{x}_j \dots x_k], \tag{6.7}$$

where the notation $[x_0 \cdots \hat{x}_j \cdots x_k]$ means that the j^{th} vertex is deleted. In Figure 6.1 we demonstrate the action of the boundary map on a 2-simplex. As the name suggests, the boundary map acting on a simplex gives the boundary of that simplex. The action of ∂_k can be extended by linearity to collections of simplices.

We define any object that does not have a boundary as a cycle. So a cycle c satisfies $c \in \text{Ker}\, \partial_k$. All boundaries are cycles, because boundaries don't themselves have a boundary. In other words, the boundary operator is nilpotent:

$$\partial_k \circ \partial^{k+1} = 0. \tag{6.8}$$

How should we define what is a hole in a simplicial complex? Intuitively a hole is a cycle which is not the boundary of anything. So a hole h satisfies $h \in \text{Ker}\, \partial_k$, but there does not exist any v such that $h = \partial^{k+1} v$. Formally this means that holes are elements of the *homology group*:

$$H_k = \frac{\text{Ker}\, \partial_k}{\text{Im} \partial^{k+1}}. \tag{6.9}$$

Note that the homology group is a quotient group, meaning that its elements are equivalence classes. We can think of these equivalence classes as being sets of cycles that can be continuously deformed into each other. Cycles which are boundaries can be continuously deformed to a single point, so these are trivial elements in homology. If two non-trivial cycles cannot be continuously deformed into one another then they are the boundaries of different holes, and so are different elements of homology.

It is possible to define a coboundary operator:

$$d_k = (\partial^{k+1})^\dagger \tag{6.10}$$

which can in turn be used to define the *Laplacian*:

$$\Delta_k = d^{k-1}\partial_k + \partial^{k+1}d_k \tag{6.11}$$

which is a positive semi-definite operator.

There is a close relationship between the Laplacian and the homology, which is described by *Hodge theory*. The most basic proposition of Hodge theory is the following.

Proposition 6.3. $\mathrm{Ker}\,\Delta_k$ *is canonically isomorphic to* H_k.

Proof. Suppose $\Delta_k|\psi\rangle = 0$. This means that

$$0 = \langle\psi|\Delta_k|\psi\rangle \tag{6.12}$$

$$= ||\partial_k|\psi\rangle||^2 + ||d_k|\psi\rangle||^2, \tag{6.13}$$

where we used that $d^{k-1} = (\partial^k)^\dagger$ and $\partial^{k+1} = (d^k)^\dagger$. Thus

$$||\partial_k|\psi\rangle||^2 = ||d_k|\psi\rangle||^2 = 0 \tag{6.14}$$

$$\implies \partial_k|\psi\rangle = d_k|\psi\rangle = 0. \tag{6.15}$$

Thus

$$|\psi\rangle \in \mathrm{Ker}\,\partial_k \tag{6.16}$$

and

$$|\psi\rangle \in \mathrm{Ker}\,d_k = (\mathrm{Im}(d_k)^\dagger)^\perp = (\mathrm{Im}\partial^{k+1})^\perp. \tag{6.17}$$

Each homology class $[|\psi\rangle] \in H_k = \mathrm{Ker}\,\partial_k/\mathrm{Im}\partial^{k+1}$ will have a unique representative orthogonal to $\mathrm{Im}\partial^{k+1}$. □

The proposition tells us that each homology class has a unique *harmonic* representative, where harmonic means that it is in the kernel of the Laplacian. The equation $\Delta_k|\psi\rangle = 0$ is a high-dimensional generalization of Laplace's equation, and $\mathrm{Ker}\,\Delta_k$ is sometimes referred to as the *harmonic subspace*.

Clique complexes

The computational complexity of determining whether or not a simplicial complex has a trivial homology group depends, of course, on how the simplicial complex is provided as input. If we are given the simplicial complex as a list of simplices, the problem of deciding the homology is in ¶. This is because we are doing linear algebra over a space whose dimension is equal to the number of k-simplices [DC91]. To make the question more interesting, we would like a succinct description of the simplicial complex. This is the purpose of this section. Motivated from the practical task of topological data analysis, we will study clique complexes—a class of simplicial complexes that can be represented by a graph $\mathcal{G}$.[2]

Definition 6.3. *The* clique complex *of a graph $\mathcal{G}$, denoted $Cl(\mathcal{G})$, is the simplicial complex consisting of the* cliques *of $\mathcal{G}$. A $k+1$-clique becomes a k-simplex.*

Now the input size is the $n \times n$ adjacency matrix of $\mathcal{G}$, where n is the number of vertices. Yet there could be up to $\binom{n}{k+1}$ k-simplices. If k is growing with n, the number of k-simplices and hence the dimension of $C_k(\mathcal{G})$ could be exponential in n. Combined with the relevance of the self-adjoint operator Δ_k, we can start to see the emergence of quantum mechanical concepts in these homological objects. Hiding in this succinct graph is an exponential-dimensional Hilbert space $C_k(\mathcal{G})$ with a Hamiltonian Δ_k!

It should be noted that not all simplicial complexes arise as clique complexes. For example, the 'hollow triangle' $\{[x_0x_1], [x_1x_2], [x_2x_0]\}$ cannot be the clique complex of any graph.

Quantum complexity theory

There are two complexity classes we will be interested in throughout this chapter. The first, QMA is often referred to as the quantum analogue of the classical complexity class, NP. It is the set of problems where a proof (in the form of a quantum state) can be checked efficiently by a quantum computer. The second complexity class we deal with is a slight modification of QMA known as QMA_1. This is the 'perfect completeness' version of QMA. This means that in YES cases we require the verifier to accept on valid witnesses with probability 1, while we still allow some probability of error in NO cases.

[2] We note there do exist succinct descriptions of general simplicial complexes, e.g. by providing a list of vertices and maximal faces, or a list of vertices and minimal non-faces. We do not consider these input representations in this work.

The question of whether QMA_1 is strictly contained within QMA is open. The complexity classes are known to be distinct relative to a particular quantum oracle [Aar08]. However, the versions of QMA and QMA_1 where the proofs are restricted to be classical bit strings *are* equal, for certain choices of universal gate sets in the perfect completeness case [Jor+12]. Equivalence of the classes for certain choices of universal gate set is also known to be true in the setting where there is exponentially small completeness-soundness gap [FL16b].

Containment in QMA

We stated Problem 6.1 without any reference to quantum mechanics, and Theorem 6.1 may appear surprising. However, this classical problem does exhibit some characteristic properties of quantum mechanics. Note first that the basis of the vector space $C_k(\mathcal{G})$ are the k-simplices of $\mathrm{Cl}(\mathcal{G})$. There could be as many as $\binom{n}{k+1}$ of these. So if k is growing linearly with n, the chainspace could have dimension exponential in n. The emergence of an exponential-dimensional Hilbert space from a small object is a characteristic property of quantum mechanics. Moreover, we can see the combinatorial Laplacian as playing the role of a quantum Hamiltonian, strengthening the link to quantum mechanics.

We can demonstrate containment of Problem 6.1 in QMA as follows [CC21]. The Laplacian Δ_k of $\mathrm{Cl}(\mathcal{G})$ is a sparse Hermitian operator. If $\mathrm{Cl}(\mathcal{G})$ has a hole then Δ_k has a zero eigenvalue, and if it has no hole then via the promise every eigenvalue of Δ_k is bounded away from zero by g. Our QMA verification protocol is simply to run quantum phase estimation [Kit95] on the witness state. We accept if the measured energy is smaller than $g/2$, and reject otherwise. In YES cases a valid witness will be an eigenstate of Δ_k with eigenvalue zero, and in NO cases all possible witnesses will fail with high probability.

Hardness construction

The challenging aspect of our work is to demonstrate QMA_1-hardness of Problem 6.1. Containment tells us we can frame the topological problem in quantum mechanical terms. Hardness gives us a converse: we can convert a quantum Hamiltonian to a topological object whose topology reflects the minimum eigenvalue of the Hamiltonian.

We will establish hardness by reducing from Quantum-4-SAT. In [Bra11] the authors show QMA_1-hardness of Quantum-4-SAT by constructing a family of local Hamiltonians H which encode the computational histories of a QMA_1-verification

circuits, such that

a. If there is an accepting witness to the QMA_1-verification circuit then H has a zero energy eigenstate.

b. If there is no accepting witness then the minimum eigenvalue of H is bounded away from zero.

We will encode H into some graph $\mathcal{G}$ such that

i. If H has a zero energy groundstate there is a hole in $\mathrm{Cl}(\mathcal{G})$.

ii. If the minimum eigenvalue of H is bounded away from zero then the minimum eigenvalue of the Laplacian is bounded away from zero.

This is the content of Theorem 6.2, which completes the hardness argument.

First we would like to develop a notion of *tensor products* for clique complexes. This will be provided by the join operation, which we now describe.

Definition 6.4. *Given two simplicial complexes $\mathcal{K}, \mathcal{L}$, we define their* join *to be the simplicial complex consisting of simplices $\sigma \otimes \tau := \sigma \cup \tau$ for all $\sigma \in \mathcal{K}, \tau \in \mathcal{L}$.*

The chain spaces and homology of the join are the tensor products of the constituents, which is often referred to as the *Kunneth formula*.

Fact 6.1. *There are canonical isomorphisms*

$$C_k(\mathcal{K} * \mathcal{L}) \cong \bigoplus_{i+j=k-1} C_i(\mathcal{K}) \otimes C_j(\mathcal{L}) \tag{6.18}$$

$$H_k(\mathcal{K} * \mathcal{L}) \cong \bigoplus_{i+j=k-1} H_i(\mathcal{K}) \otimes H_j(\mathcal{L}). \tag{6.19}$$

Moreover, if $|\psi\rangle \in C^i(\mathcal{K})$ and $|\varphi\rangle \in C^j(\mathcal{L})$ where $i + j = k - 1$, then the Laplacian acts as

$$\Delta_k(|\psi\rangle \otimes |\varphi\rangle) = (\Delta_i|\psi\rangle) \otimes |\varphi\rangle + |\psi\rangle \otimes (\Delta_j|\varphi\rangle). \tag{6.20}$$

Since we are building our simplicial complexes as the clique complexes of graphs, we must be able to implement the join at the level of the graphs. This is achieved by taking the two constituent graphs $\mathcal{G}$ and $\mathcal{G}'$ and including all edges between $\mathcal{G}$ and $\mathcal{G}'$.

Definition 6.5. *The* join *of two graphs $\mathcal{G} = (V, E)$ and $\mathcal{G}' = (V', E')$ is the graph $\mathcal{G} * \mathcal{G}'$ with vertices $V \cup V'$ and edges $E \cup E' \cup \{(u, v) : u \in V, v \in V'\}$.*

Crucially, the clique complex of the join of two graphs is the join of the clique complexes of the graphs.

The first step is to construct a graph $\mathcal{G}_1$ such that $\mathrm{Cl}(\mathcal{G}_1)$ has two holes.[3] This means that the homology of $\mathrm{Cl}(\mathcal{G}_1)$ can encode the Hilbert space of one qubit. We then form our base qubit graph $\mathcal{G}_n$ by taking the n-fold join of $\mathcal{G}_1$. That is, we take n copies of $\mathcal{G}_1$ where vertices are connected all-to-all between the different copies. Constructed in this way, $\mathcal{G}_n$ has 2^n $2n-1$-dimensional holes, one for each computational basis state. Moreover, $\mathcal{G}_n$ has a tensor-product-like structure—each copy of $\mathcal{G}_1$ in the join can be identified with a qubit. At this stage, the kernel of the Laplacian is isomorphic to the entire encoded Hilbert space of n qubits—$\mathcal{G}_n$ by itself corresponds to the empty zero Hamiltonian.

The next step is to design *gadgets* which implement terms in the H. First we decompose H into a sum of local rank-1 projectors. Then for each m-local term $|\phi\rangle\langle\phi|$ we take the graph $\mathcal{G}_m$ (the m-fold join of the single qubit graph) and design a gadget which 'fills in the hole' in $\mathrm{Cl}(\mathcal{G}_m)$ corresponding to the state $|\phi\rangle$. Constructively, this involves adding extra *gadget vertices* to the graph, and adding edges between these new vertices and the original vertices from $\mathcal{G}_m$. This serves to lift the cycle corresponding to $|\phi\rangle$ out of the homology by rendering it a boundary. The clique complex of the resulting m-qubit graph has $2^m - 1$ holes of dimension $2m - 1$, encoding a m-local projector acting on a system of m qubits.

To construct a graph $\hat{\mathcal{G}}_n$ which implements the Hamiltonian $H = \sum_i |\phi_i\rangle\langle\phi_i|$ the procedure is as follows:

1. Start with the graph $\mathcal{G}_n$.
2. For each term $|\phi_i\rangle\langle\phi_i|$ in H, insert the gadget implementing that term onto the copies of $\mathcal{G}_1$ corresponding to qubits in the support of $|\phi_i\rangle\langle\phi_i|$.
3. For each term $|\phi_i\rangle\langle\phi_i|$, connect its gadget vertices *all to all* with the vertices of $\mathcal{G}_n$ corresponding to qubits outside the support of $|\phi_i\rangle\langle\phi_i|$.

[3] This was also the first step in [CK22]; however, the graph we choose here is a different one. In order to show hardness of the gapped problem, we need that the natural inner product on simplices respects that encoded computational basis states should be orthogonal. This feature was not present in the encoding of [CK22].

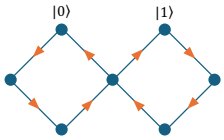

(a) Our one-qubit graph is the 'bowtie' graph pictured here. It has two 1D holes, which encode the qubit in the 1-homology.

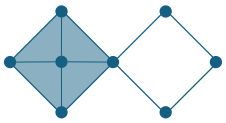

(b) The gadget to implement the projector onto $|0\rangle\langle 0|$ is simply to add an additional vertex and connect it to the hole representing the $|0\rangle$ state. This fills in the $|0\rangle$ hole and removes it from the homology, providing it some energy on the Laplacian.

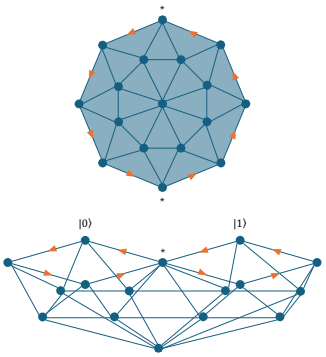

(c) The gadget to implement the projector onto $|-\rangle\langle -|$ is more involved. We must fill in the cycle given by $|-\rangle \propto |0\rangle - |1\rangle$.

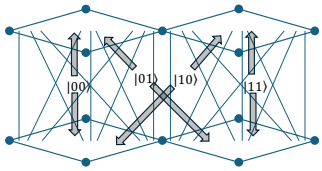

(d) This is a depiction of the two-qubit graph, which consists of the join of two bowtie graphs. The vertices in the distinct bowties are connected all-to-all by edges.

Figure 6.2: Some illustrations of the hardness construction.

4. Do *not* connect any gadget vertices coming from different Hamiltonian terms.

With this candidate reduction in hand we then need to show that it satisfies the necessary properties. Demonstrating that the resulting graph satisfies Item i is straightforward—our construction of the gadgets fills in the holes in $\mathrm{Cl}(\hat{\mathcal{G}}_n)$ corresponding to states that are lifted in H. If the Hamiltonian H has a zero energy groundstate then there is a state $|\psi\rangle$ in the Hilbert space of n-qubits that satisfies every projector in H. This state corresponds to a hole in $\mathrm{Cl}(\mathcal{G}_n)$ that has not been filled in by any gadget, and thus the hole remains in $\mathrm{Cl}(\hat{\mathcal{G}}_n)$ has non-trivial homology. If, on the other hand, H is not satisfiable, there is no state in the Hilbert space of n qubits that satisfies every projector in H. Therefore, the process of 'filling in holes' via gadgets has removed 2^n holes from the homology of $\mathrm{Cl}(\mathcal{G}_n)$ to construct $\mathrm{Cl}(\hat{\mathcal{G}}_n)$. We demonstrate when constructing the gadgets that the method of constructing gadgets does not introduce any new 'spurious' homology classes into the complex. Therefore, since all 2^n holes have been removed, and no holes have been

introduced, the resulting complex $\mathrm{Cl}(\hat{\mathcal{G}}_n)$ has trivial homology.[4]

Demonstrating that the graph satisfies Item ii is more challenging, and constitutes the main technical contribution of this work. At first glance there is no reason why Item ii should hold. We have encoded the ground space of H into the homology of $\mathrm{Cl}(\hat{\mathcal{G}}_n)$. But the Laplacian of $\mathrm{Cl}(\hat{\mathcal{G}}_n)$ has many more excited states than the spectrum of H, and it is plausible that the excited spectrum of the Laplacian includes very low energy states.

Proof of gap

Our construction consists of *weighted* simplicial complexes. We should clarify that the weighting will *not* affect the homology of the simplicial complex; rather, it will only affect the Laplacian operator. In order to define the Laplacian, we implicitly chose an inner product on the chain spaces C_k, which consisted of the most basic choice of declaring the simplices themselves to form an orthonormal basis. We will now relax this so that the simplices are orthogonal with weights. The more general inner product is

$$\langle \sigma | \tau \rangle = \begin{cases} w(\sigma)^2 & \sigma = \tau \\ 0 & \text{otherwise} \end{cases} \tag{6.21}$$

for $\sigma, \tau \in \mathcal{K}^k$.

This involves assigning a weight $w(\sigma) \geq 0$ to each simplex σ in the simplicial complex $\mathcal{K}$. There are two issues associated with the generality of this definition. Firstly, recall that we introduced the clique complex to provide a more succinct description of a simplicial complex. We likewise need the inner product to be succinctly describable, so listing the weights of all simplices is not possible. The second is that we would like the inner product to respect the join operation. That is, after taking the join of two simplicial complexes $\mathcal{K} * \mathcal{L}$, we would like the inner product on $\bigoplus_{i+j=k-1} C_i(\mathcal{K}) \otimes C_j(\mathcal{L})$ to be induced from those on $C_i(\mathcal{K})$ and $C_j(\mathcal{L})$.

To solve these issues, we add more structure to the definition. Each *vertex* v in the simplicial complex is assigned a weight $w(v)$, and the weights of the higher simplices are induced via

$$w(\sigma) = \prod_{v \in \sigma} w(v). \tag{6.22}$$

[4]This same argument was used in [CK22] with a different graph construction to show that the decision version of the clique homology problem is QMA_1-hard—see Chapter 6 for details.

This allows a graph $\mathcal{G}$ with weighted *vertices* to induce a weighted clique complex. Note that the edges of $\mathcal{G}$ are still binary (present or not present). This also ensures that, in the join construction, the weight of the tensor product of two simplices σ, τ is the product of the individual weights $w(\sigma \otimes \tau) = w(\sigma)w(\tau)$.

In order to get a handle on the excited spectrum of a single gadget, we need to ensure that the states we lift out of homology do not mix with the rest of the spectrum. In order to do this we weight the gadget vertices by some polynomially small parameter $\lambda \ll 1$. This ensures that the gadgets can be viewed as perturbations of the original qubit complex. We then want to analyze the spectrum and eigenspaces, which sounds similar to the domain of perturbation theory and perturbative gadgets [KKR06]. However, perturbation theory is not able to provide the level of generality and control we require for our purposes.[5]

Here, the main innovation of our work enters. We use an advanced tool from algebraic topology known as *spectral sequences* [For94; Cho06; McC01] to analyze the spectrum of each gadget. We will now elaborate on this key technique.

Our complex has some vertices weighted by λ, which is a small perturbative parameter $\lambda \ll 1$. In order to understand the low-energy spectrum of the Laplacian, we would like to expand the kernel of the Laplacian perturbatively in λ. This gives us a sequence of vector spaces $E_0^k, E_1^k, E_2^k, \ldots$ which provide increasingly close approximations to $\operatorname{Ker} \Delta_k$.

$$E_j^k \to \operatorname{Ker} \Delta_k \text{ as } j \to \infty. \tag{6.23}$$

Taking all orders of λ into account gives the true kernel of the Laplacian, which is isomorphic to the homology.

$$\operatorname{Ker} \Delta_k \cong H_k. \tag{6.24}$$

From the weighting of the complex we can obtain a *filtration* on the chain complex, and a filtered chain complex has an associated *spectral sequence*, which consists of vector spaces $e_{j,l}^k$. At $j = 0$, our spectral sequence $e_{0,l}^k$ consists of the k-simplices that are weighted by λ^l. For each 'page' j there are coboundary maps

$$d_{j,l}^k : e_{j,l}^k \to e_{j,l+j}^{k+1} \tag{6.25}$$

[5]The difficulty is that we would like to go to arbitrarily high orders of perturbation theory. Using generic perturbation theory tools, this would be close to impossible.

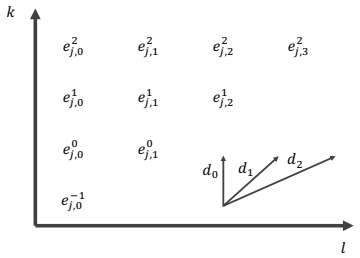

Figure 6.3: One can visualize the pages of the spectral sequence as a 2D array of vector spaces for each j. To move to the next page, we take the cohomology everywhere with respect to the coboundary map $d^{\cdot}_{j,\cdot}$.

and to obtain the page $e^k_{j+1,l}$ we take the coboundary of the previous page $e^k_{j,l}$. Now consider the direct sums

$$e^k_j := \bigoplus_l e^k_{j,l}. \tag{6.26}$$

The vector spaces e^k_j provides a sequence of approximations to the true homology group H_k as j increases.

$$e^k_j \to H_k \text{ as } j \to \infty. \tag{6.27}$$

Intuitively, subsequent pages of the spectral sequence take into account more terms in the filtration, which in our case corresponds to including terms of higher order in λ.

A result from algebraic topology tells us that these two sequences are in fact isomorphic [For94].

$$E^k_j \cong e^k_j \tag{6.28}$$

That is, the perturbative expansion of the Laplacian groundspace is isomorphic to the spectral sequence of the filtration. This provides a connection between Hodge theory and spectral sequences. Our strategy from here is clear: We compute the spectral sequence e^k_j algebraically and use the isomorphism of Ref. [For94] to deduce the perturbative expansion E^k_j of the Laplacian kernel and hence the low-energy spectrum of the Laplacian. Our use of this isomorphism is reminiscent of many methods in algebraic topology, where one converts difficult questions in analysis and topology into the easier language of algebra.

Once we have completed the analysis of the spectrum and eigenspace of the Laplacian corresponding to a single gadget using spectral sequences, we must still analyze what happens when we put all the gadgets together. The argument that the lowest eigenvalue of the Laplacian is bounded away from zero in NO cases proceeds in two steps:

1. In a NO case, any state must have large overlap with the excited subspace of at least one of the gadgets.

2. States with low energy must have small overlap with the excited subspace of all the gadgets. Hence, by Step 1, such low energy states do not exist.

The first point is straightforward—in NO cases the Hamiltonian H is not satisfiable, so it is not possible to construct a global state which is in the ground state of every projector. Therefore the overlap of any global state with the zero energy groundstate of each gadget must be bounded away from one for at least one of the gadgets. The second point is technically challenging—it involves detailed understanding of the structure of the combined Laplacian and its eigenspaces.

6.1 Preliminaries

$\mathbf{QMA}_1$ and Quantum m-SAT

Let's begin by formally defining QMA and QMA_1.

Definition 6.6 (QMA [Kit+02])**.** *A problem $A = (A_{yes}, A_{no})$ is in QMA if there is a P-uniform family of polynomial-time quantum circuits (the "verifier") V_n, one for each input size n, such that*

- *If $x \in A_{yes}$, there exists a poly(n)-qubit witness state $|w\rangle$ such that $\mathbb{P}[V_n(x, |w\rangle) = 1] \geq \frac{2}{3}$,*
- *If $x \in A_{no}$, then for any poly(n)-qubit witness state $|w\rangle$, $\mathbb{P}[V_n(x, |w\rangle) = 1] \leq \frac{1}{3}$.*

The constants $\frac{1}{3}$, $\frac{2}{3}$ in the definition of QMA are conventional. The definition of QMA is equivalent as long as the acceptance and failure probabilities are separated by some inverse polynomial in the problem size [Kit+02].

Definition 6.7 (QMA_1 [Bra11; GN16])**.** *A problem $A = (A_{yes}, A_{no})$ is in QMA_1 if there is a P-uniform family of polynomial-time quantum circuits (the "verifier") V_n, one for each input size n, such that*

- *If $x \in A_{yes}$, there exists a poly(n)-qubit witness state $|w\rangle$ such that $\mathbb{P}[V_n(x, |w\rangle) = 1] = 1$,*
- *If $x \in A_{no}$, then for any poly(n)-qubit witness state $|w\rangle$, $\mathbb{P}[V_n(x, |w\rangle) = 1] \leq \frac{1}{3}$.*

The canonical QMA_1-complete problem is Quantum m-SAT, in which you are asked to decide whether a local Hamiltonian composed of positive semi-definite local terms has an exactly zero eigenvalue, given a suitable promise gap. In this work we will reduce from Quantum 4-SAT to the problem of deciding whether or not a particular homology group of a clique complex is non-trivial.

Problem 6.8. (Quantum m-SAT[Bra11]) *Fix function $g : \mathbb{N} \to [0, \infty)$ with $g(n) \geq 1/\mathit{polyn}$. The input to the problem is a a list of m-local projectors $\Pi_j \in \mathcal{P}$. $\mathcal{P}$ are a set of projectors obeying certain constraints.*[6] *Let*

$$H = \sum_j \Pi_j. \tag{6.29}$$

The task is to decide whether:

- ***YES*** *H is satisfiable i.e. there exists a state $|\psi\rangle$ with $H|\psi\rangle = 0$.*
- ***NO*** *The minimum eigenvalue of H is at least g.*

Quantum m-SAT is known to be QMA_1-complete for $m \geq 3$ [Bra11; GN16]. Quantum 2-SAT is known to be in P [Bra11].

In [Bra11] a history state construction was used to reduce from a general problem in QMA_1 to Quantum 4-SAT. Computational history states are of the form

$$|\Phi\rangle_{CQ} = \sum_{t=0}^{T} |t\rangle |\psi_t\rangle, \tag{6.30}$$

where $\{|t\rangle\}$ is an orthonormal basis for $\mathcal{H}_C$, the clock register, and the $|\psi_t\rangle = \Pi_{i=0}^{t} U_i |\psi_0\rangle$ for some initial state $|\psi_0\rangle$ and some set of unitaries $\{U_i\}$. The first register of $|\Phi\rangle_{QC}$ encodes the time, while the second register (the 'computational' register) encodes the state of the quantum circuit at time t.

The idea in [Bra11] is to construct a local Hamiltonian (composed of projectors), whose zero energy ground states are history states, this Hamiltonian is given by:

$$H_{\mathrm{hs}} = H_{\mathrm{in}} + H_{\mathrm{clock}} + H_{\mathrm{prop}}, \tag{6.31}$$

[6] We have been deliberately vague in defining the constraints that the set of projectors must satisfy. That is because the constraints depend on the gate set, and there is not a standard definition. However, this is only an issue for showing *containment* in QMA_1. For showing QMA_1-hardness there is no need to include any constraints on the form of the projectors in Problem 6.8. Since we are interested in showing QMA_1-hardness and containment in QMA (where we only need to constrain the locality of the projectors), we will not impose any constraints on the allowable projectors (beyond their locality) throughout this work.

where H_{in} constrains the starting state $|\psi_0\rangle$ of the circuit, H_{clock} penalises any states in the clock register which don't encode valid times, and H_{prop} penalises any states where $|\psi_t\rangle \neq U_t |\psi_{t-1}\rangle$. H_{hs} has a degenerate zero energy ground state, where all computations which start in a valid state $|\psi_0\rangle$ satisfy every constraint. To break this degeneracy, and encode QMA_1-verification circuits it is necessary to add one extra term to the Hamiltonian:

$$H_{\text{Bravyi}} = H_{\text{hs}} + H_{\text{out}}, \tag{6.32}$$

where H_{out} penalises any computation which outputs NO, and gives zero energy to any computation which outputs YES.

In Table 6.1 we give an overview of the rank-1 projectors that we need to be able to implement in order to reduce from the construction in [Bra11]. In Section 6.2 we will construct gadgets to implement each of these states.

Note on gateset for QMA_1: It is important to note that the definition of QMA_1 (see Definition 6.7) implicitly depends on a choice of universal gate set. In standard QMA (see Definition 6.6) the definition is independent of gate set, since all universal gate sets can approximate any unitary evolution. However, the requirement of perfect completeness in QMA_1 means that it may be necessary to implement a given unitary evolution exactly.

When choosing a universal gate set for our construction we require that every gate in the set should have only rational coefficients. This ensures that the only states we need to lift are integer states—see Section 6.2. Based on this requirement, we choose the universal gate set: $\mathcal{G} = \{\text{CNOT}, U\}$ where U is the 'Pythagorean gate' [CK22]:

$$U = \frac{1}{5}\begin{pmatrix} 3 & 4 \\ -4 & 3 \end{pmatrix}. \tag{6.33}$$

This choice of gate set is shown to be universal in [ADH97, Theorem 3.3] and [Shi03, Theorem 1.2].

Simplicial homology

We will re-introduce the subject of homology more formally and from a 'cohomology-first' perspective, since this will play nicely with the weighting introduced in Section 6.1, and later spectral sequences in Section 6.3.

Definition 6.9. *A* simplicial complex $\mathcal{K}$ *is a collection of subsets* $\mathcal{K} = \mathcal{K}^0 \cup \mathcal{K}^1 \cup \ldots$ *such that*

Term in H_{Bravyi}	*Penalizes state* $\|\psi_S\rangle$
$H_{\text{prop}}t'$	$\frac{1}{\sqrt{2}}(\|1011\rangle - \|1000\rangle)$
$H_{\text{prop}}t(\text{CNOT})$	$\frac{1}{\sqrt{2}}(\|0110\rangle - \|0101\rangle)$
$H_{\text{prop}}t(\text{CNOT})$	$\frac{1}{\sqrt{2}}(\|0010\rangle - \|0001\rangle)$
$H_{\text{prop}}t(U_{Pyth.})$	$\frac{1}{5\sqrt{2}}(-5\|011\rangle + 4\|100\rangle + 3\|101\rangle)$
$H_{\text{prop}}t(U_{Pyth.})$	$\frac{1}{5\sqrt{2}}(-5\|010\rangle + 3\|100\rangle - 4\|101\rangle)$
$H_{\text{prop}}t(\text{CNOT})$	$\frac{1}{\sqrt{2}}(\|1101\rangle - \|1010\rangle)$
$H_{\text{prop}}t(\text{CNOT})$	$\frac{1}{\sqrt{2}}(\|1011\rangle - \|1100\rangle)$
$H^{(1)}_{\text{clock}}$	$\|00\rangle$
$H^{(2)}_{\text{clock}}$	$\|11\rangle$
$H_{\text{in}}, H_{\text{out}}$	$\|011\rangle$
$H^{(6)}_{\text{clock}}, H^{(4)}_{\text{clock}}, H^{(5)}_{\text{clock}}, H^{(3)}_{\text{clock}}$	$\|1100\rangle$
$H^{(4)}_{\text{clock}}$	$\|0111\rangle$
$H^{(5)}_{\text{clock}}$	$\|0001\rangle$

Table 6.1: Projectors needed for quantum 4-SAT with universal gate set $\mathcal{G}$. Note we collated projectors which are the same up to re-ordering the qubits involved.

- $\sigma \in \mathcal{K}^k$ *have* $|\sigma| = k + 1$.
- *If* $\sigma \in \mathcal{K}$*, then* $\tau \in \mathcal{K}$ *for all* $\tau \subset \sigma$.

Intuitively, a simplicial complex is a higher dimensional generalization of a graph. It has vertices $\mathcal{K}^0$, edges $\mathcal{K}^1$, triangles $\mathcal{K}^2$, tetrahedra $\mathcal{K}^3$, et cetera.

From a simplicial complex we can derive a *chain complex*. Let $C^k(\mathcal{K})$ be the complex vector space formally spanned by $\mathcal{K}^k$. This involves picking a conventional ordering for each simplex $\sigma = [v_0, \ldots, v_k]$ and identifying $|[\pi(v_0), \ldots, \pi(v_k)]\rangle = \text{sgn}(\pi)|\sigma\rangle$ for any permutation $\pi \in S_{k+1}$. Here $\text{sgn}(\pi)$ denotes the sign of the permutation, and it is known as the *orientation* of the simplex. $C^k(\mathcal{K})$ is known as a *chain space*.

For a k-simplex $\sigma \in \mathcal{K}^k$, let $\text{up}(\sigma) \subset \mathcal{K}^0$ be the subset of vertices v such that $\sigma \cup \{v\} \in \mathcal{K}^{k+1}$ is a $k+1$-simplex.

Define the *coboundary map* $d^k : C^k(\mathcal{K}) \to C^{k+1}(\mathcal{K})$ to act as

$$d^k|\sigma\rangle = \sum_{v \in \text{up}(\sigma)} |\sigma \cup \{v\}\rangle \tag{6.34}$$

for a k-simplex $\sigma \in \mathcal{K}^k$—see Figure 6.4.

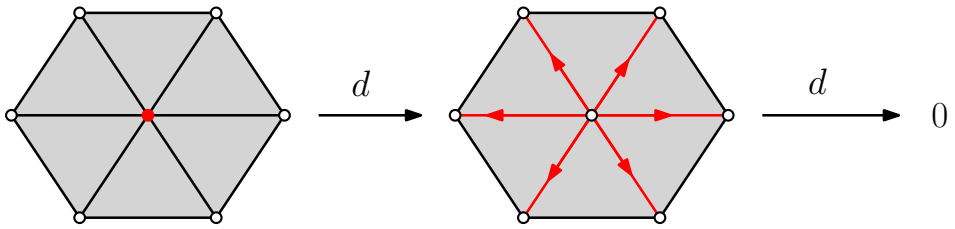

Figure 6.4: The action of the coboundary map. In each step the coboundary map is acting on the parts of the complex shown in red. The third figure is empty because the lines that are acted on by the coboundary map in the second figure all have coboundaries composed of two triangles, and these coboundaries cancel out with those of the lines on either side as they have opposite orientation.

This gives a chain of vector spaces with linear maps between them.

$$C^{-1}(\mathcal{K}) \xrightarrow{d^{-1}} C^0(\mathcal{K}) \xrightarrow{d^0} C^1(\mathcal{K}) \xrightarrow{d^1} C^2(\mathcal{K}) \xrightarrow{d^2} \dots \tag{6.35}$$

Here we define $C^{-1}(\mathcal{K}) = \text{span}\{\emptyset\} \cong \mathbb{C}$ the 1-dimensional vector space, and d^0 maps the empty set $|\emptyset\rangle$ to the uniform superposition $\sum_{v\in\mathcal{K}^0} |v\rangle \in C^0(\mathcal{K})$. With this convention, we are working with the *reduced* cohomology.

It can be checked that $d^k \circ d^{k-1} = 0$ for each k. This gives a *chain complex*.

Definition 6.10. *A* chain complex *is a chain of complex vector spaces C^k with linear maps $d^k : C^k \to C^{k+1}$ which satisfy $d^k \circ d^{k-1} = 0$ for all k.*

$$C^{-1} \xrightarrow{d^{-1}} C^0 \xrightarrow{d^0} C^1 \xrightarrow{d^1} C^2 \xrightarrow{d^2} \dots \tag{6.36}$$

$d^k \circ d^{k-1} = 0$ means that $\text{Im}\, d^{k-1} \subseteq \text{Ker}\, d^k$ for each k. This allows us to define the *cohomology groups* as

$$H^k = \frac{\text{Ker}\, d^k}{\text{Im}\, d^{k-1}}. \tag{6.37}$$

Let $(C^k)^*$ be the *dual space* of C^k. Formally, this is the space of linear functionals $f : C^k \to \mathbb{C}$. Let $\partial^k = (d^{k-1})^* : (C^k)^* \to (C^{k-1})^*$ the dual map of d^{k-1}. ∂^k are known as *boundary maps*, as introduced in Chapter 6. As noted before, these act as

$$\partial^k |\sigma\rangle = \sum_{v\in\sigma} |\sigma \setminus \{v\}\rangle \tag{6.38}$$

for a k-simplex $\sigma \in \mathcal{K}^k$. (Here, by $|\sigma\rangle$ and $|\sigma \setminus \{v\}\rangle$, we technically mean the indicator functions of these simplices, which are members of the dual spaces $(C^k)^*, (C^{k-1})^*$.) See Figure 6.1 for a diagrammatic representation.

We get the chain complex

$$(C^{-1})^* \xleftarrow[\partial^0]{} (C^0)^* \xleftarrow[\partial^1]{} (C^1)^* \xleftarrow[\partial^2]{} (C^2)^* \xleftarrow[\partial^3]{} \dots \tag{6.39}$$

From $d^k \circ d^{k-1} = 0 \; \forall k$ we get that $\partial^k \circ \partial^{k+1} = 0 \; \forall k$. This allows us to define the *homology groups* as

$$H^{k*} = \frac{\operatorname{Ker} \partial^k}{\operatorname{Im} \partial^{k+1}}. \tag{6.40}$$

The homology groups H^{k*} are the dual spaces of the cohomology groups H^k.

Hodge theory

Our next move will be to choose an inner product on C^k, thus rendering it a *Hilbert space*. This is equivalent to choosing an isomorphism between the space and its dual $C^k \leftrightarrow C^{k*}$. The most basic choice is to declare the simplices themselves to form an orthonormal basis. That is, for $\sigma, \tau \in \mathcal{K}^k$

$$\langle \sigma | \tau \rangle = \begin{cases} 1 & \sigma = \tau \\ 0 & \text{otherwise.} \end{cases} \tag{6.41}$$

We will modify this choice later, but for now let's consider this case. We can now drop the asterisks in the notation and identify $C^{k*} = C^k$. The k-boundary map is now the adjoint of the $k-1$-coboundary map

$$\partial^k = (d^{k-1})^\dagger. \tag{6.42}$$

The inner product allows us to define the *Laplacian* as

$$\Delta^k : C^k \to C^k \tag{6.43}$$

$$\Delta^k = d^{k-1}\partial^k + \partial^{k+1} d^k. \tag{6.44}$$

This is a positive semi-definite self-adjoint operator on the Hilbert space C^k. In fact, we can split up the definition and write

$$\Delta^{\downarrow k} = d^{k-1}\partial^k \tag{6.45}$$

$$\Delta^{\uparrow k} = \partial^{k+1} d^k \tag{6.46}$$

$$\Delta^k = \Delta^{\downarrow k} + \Delta^{\uparrow k} \tag{6.47}$$

and now both $\Delta^{\downarrow k}$ and $\Delta^{\uparrow k}$ are individually positive semi-definite.

Fact 6.2.

$$\langle\psi|\Delta^{\downarrow k}|\psi\rangle = ||\partial^k|\psi\rangle||^2 \tag{6.48}$$
$$\langle\psi|\Delta^{\uparrow k}|\psi\rangle = ||d^k|\psi\rangle||^2 \tag{6.49}$$
$$\langle\psi|\Delta^k|\psi\rangle = ||\partial^k|\psi\rangle||^2 + ||d^k|\psi\rangle||^2 \tag{6.50}$$
$$\tag{6.51}$$

Recall the basic theorem of Hodge theory Proposition 6.3, restated here:

Proposition 6.3. $\mathrm{Ker}\,\Delta_k$ *is canonically isomorphic to* H_k.

The proposition tells us that each homology class has a unique *harmonic* representative, where harmonic means that it is in the kernel of the Laplacian. The equation $\Delta^k|\psi\rangle = 0$ is a high-dimensional generalization of Laplace's equation, and $\mathrm{Ker}\,\Delta^k$ is sometimes referred to as the *harmonic subspace*.

Pairing of Laplacian eigenstates

From the chain complex

$$C^{-1} \xleftarrow[\partial^0]{} C^0 \xleftarrow[\partial^1]{} C^1 \xleftarrow[\partial^2]{} C^2 \xleftarrow[\partial^3]{} \dots \tag{6.52}$$

we can build the *graded* vector space

$$C = C^{-1} \oplus C^0 \oplus C^1 \oplus \dots \tag{6.53}$$

We will refer to this as the *Fock space*, where this terminology comes from a connection to supersymmetric quantum systems. ∂^k and d^k give maps

$$\partial, d : C \to C \tag{6.54}$$

by acting blockwise. The Laplacian becomes

$$\Delta = d\partial + \partial d : C \to C \tag{6.55}$$
$$\Delta^{\downarrow} = d\partial : C \to C \tag{6.56}$$
$$\Delta^{\uparrow} = \partial d : C \to C \tag{6.57}$$

$\{\partial, d\}$ generate a C^*-algebra representation on Fock space. Moreover, the Laplacian Δ commutes with ∂ and d, and hence commutes with this representation. Thus we can write each Δ-eigenspace as a sum of irreducible $\{\partial, d\}$-subrepresentations.

Proposition 6.4. *The kernel of the Laplacian* $\operatorname{Ker}\Delta$ *consists of states that are annihilated by both* ∂ *and* d*. These are one-dimensional* $\{\partial, d\}$*-irreps, or* singlets*. Let* $E > 0$ *be an eigenvalue of the Laplacian* Δ*. The* E*-eigenspace is a direct sum of 2-dimensional subspaces* $\{|\psi^{\uparrow}\rangle, |\psi^{\downarrow}\rangle\}$ *with* $|\psi^{\uparrow}\rangle \in C^k$ *for some* k *and* $|\psi^{\downarrow}\rangle \in C^{k+1}$ *such that*

$$\partial|\psi^{\uparrow}\rangle = 0 \quad , \quad d|\psi^{\uparrow}\rangle \propto |\psi^{\downarrow}\rangle \tag{6.58}$$

$$\partial|\psi^{\downarrow}\rangle \propto |\psi^{\uparrow}\rangle \quad , \quad d|\psi^{\downarrow}\rangle = 0 \tag{6.59}$$

which implies

$$\Delta^{\downarrow}|\psi^{\uparrow}\rangle = 0 \quad , \quad \Delta^{\uparrow}|\psi^{\uparrow}\rangle = E|\psi^{\uparrow}\rangle \tag{6.60}$$

$$\Delta^{\downarrow}|\psi^{\downarrow}\rangle = E|\psi^{\downarrow}\rangle \quad , \quad \Delta^{\uparrow}|\psi^{\downarrow}\rangle = 0. \tag{6.61}$$

These 2-dimensional subspaces are likewise $\{\partial, d\}$*-irreps, or* doublets*. We refer to the states* $|\psi^{\uparrow}\rangle$ *as 'paired up', and the states* $|\psi^{\downarrow}\rangle$ *as 'paired down'.*

Proof. (Sketch.) $\langle\psi|\Delta|\psi\rangle = 0$ if and only if $\partial|\psi\rangle = d|\psi\rangle = 0$ by Fact 6.2, and the states in the kernel of the Laplacian are precisely the *singlets* in this representation.

Now recall that $\partial : C^{k+1} \to C^k$, $d : C^k \to C^{k+1}$ and $\partial^2 = d^2 = 0$. These properties imply that all non-singlet irreducible $\{\partial, d\}$-subrepresentations are necessarily *doublets* supported on neighboring blocks $C^k \oplus C^{k+1}$ for some k. □

Joins

The *join* will be an important operation for us on simplicial complexes.

Definition 6.11. *Given two simplicial complexes* $\mathcal{K}$ *and* $\mathcal{L}$*, define their* join $\mathcal{K} * \mathcal{L}$ *to be the simplicial complex consisting of simplices* $\sigma \otimes \tau := \sigma \cup \tau$ *for all* $\sigma \in \mathcal{K}$*,* $\tau \in \mathcal{L}$*.*

The chain spaces and homology of the join are given by the Kunneth formula.

Fact 6.3. (Kunneth formula) *There are canonical isomorphisms*

$$C^k(\mathcal{K} * \mathcal{L}) \cong \bigoplus_{i+j=k-1} C^i(\mathcal{K}) \otimes C^j(\mathcal{L}) \tag{6.62}$$

$$H^k(\mathcal{K} * \mathcal{L}) \cong \bigoplus_{i+j=k-1} H^i(\mathcal{K}) \otimes H^j(\mathcal{L}) \tag{6.63}$$

We would also like to relate the Laplacian of $\mathcal{K} * \mathcal{L}$ to the Laplacians of $\mathcal{K}$ and $\mathcal{L}$.

Lemma 6.5. *If $|\psi\rangle \in C^i(\mathcal{K})$ and $|\varphi\rangle \in C^j(\mathcal{L})$ where $i + j = k - 1$, then*

$$\Delta^k(|\psi\rangle \otimes |\varphi\rangle) = (\Delta^i|\psi\rangle) \otimes |\varphi\rangle + |\psi\rangle \otimes (\Delta^j|\varphi\rangle). \tag{6.64}$$

Proof. If σ is an i-simplex of $\mathcal{K}$ and τ is a j-simplex of $\mathcal{L}$, and $i + j = k - 1$, it can be checked that

$$\partial^k(|\sigma\rangle \otimes |\tau\rangle) = (\partial^i|\sigma\rangle) \otimes |\tau\rangle + (-1)^{|\sigma|}|\sigma\rangle \otimes (\partial^j|\tau\rangle) \tag{6.65}$$

and

$$d^k(|\sigma\rangle \otimes |\tau\rangle) = (d^i|\sigma\rangle) \otimes |\tau\rangle + (-1)^{|\sigma|}|\sigma\rangle \otimes (d^j|\tau\rangle). \tag{6.66}$$

Now

$$\Delta^k = d^{k-1}\partial^k + \partial^{k+1}d^k \tag{6.67}$$

so

$$\Delta^k(|\sigma\rangle \otimes |\tau\rangle) = (\Delta^i|\sigma\rangle) \otimes |\tau\rangle + |\sigma\rangle \otimes (\Delta^j|\tau\rangle). \tag{6.68}$$

By linearity, this extends to chains $|\psi\rangle \in C^i(\mathcal{K})$, $|\varphi\rangle \in C^j(\mathcal{L})$. □

Since we are building our simplicial complexes as the clique complexes of graphs, we must be able to implement the join at the level of the graphs. This is achieved by taking the two constituent graphs $\mathcal{G}$ and $\mathcal{G}'$ and including all edges between $\mathcal{G}$ and $\mathcal{G}'$.

Definition 6.12. *The* join *of two graphs $\mathcal{G} = (V, E)$ and $\mathcal{G}' = (V', E')$ is the graph $\mathcal{G} * \mathcal{G}'$ with vertices $V \cup V'$ and edges $E \cup E' \cup \{(u, v) : u \in V, v \in V'\}$.*

Fact 6.4. *The clique complex of the join of two graphs is the join of the clique complexes of the graphs.*

Generalized octrahedra

Let $\mathfrak{g}_1$ denote the graph consisting of two disjoint points (with no edges), and let $\mathfrak{g}_n$ be the n-fold *join* of $\mathfrak{g}_1$.[7]

$$\mathfrak{g}_n = \mathfrak{g}_1 * \cdots * \mathfrak{g}_1 \quad (n \text{ times}) \tag{6.69}$$

[7] By Fact 6.4, it is equivalent to think of the join as acting at the level of graphs or at the level of simplicial complexes.

It will be useful to develop an intuitive interpretation for the complex $\mathfrak{g}_n$. The Kunneth formula tells us that it should have

$$\dim H^k = \begin{cases} 1 & k = n-1 \\ 0 & \text{otherwise.} \end{cases} \tag{6.70}$$

In fact, $\mathfrak{g}_n$ is topologically homeomorphic to the $(n-1)$-sphere S^{n-1}. But which triangulation of S^{n-1} does it form in particular?

We can interpret $\mathfrak{g}_n$ as a "generalized octahedron". $\mathfrak{g}_2$ is the square loop, and $\mathfrak{g}_3$ is the standard octahedron. The $(n-1)$-simplices of $\mathfrak{g}_n$ are the n-cliques of the 1-skeleton. There are 2^n of these, corresponding to choosing one vertex from each copy of $\mathfrak{g}_1$. Notice how the number of simplices which make up the higher dimensional octahedron is exponential in the number of vertices. These generalized octahedron are sometimes referred to as cross-polytopes.

Weighting

In our construction, we would like to consider *weighted* simplicial complexes. The purpose of this section is to define a natural notion of weighting. We should clarify that the weighting will *not* affect the homology of the simplicial complex; rather, it will only affect the Laplacian operator.

Recall that, in order to define the Laplacian, we had to choose an inner product on the chain spaces C^k. We went with the most basic choice of declaring the simplices themselves to form an orthonormal basis. We will now relax this so that the simplices are orthogonal with weights. The more general inner product is

$$\langle \sigma | \tau \rangle = \begin{cases} w(\sigma)^2 & \sigma = \tau \\ 0 & \text{otherwise} \end{cases} \tag{6.71}$$

for $\sigma, \tau \in \mathcal{K}^k$.

This involves assigning a weight $w(\sigma) \geq 0$ to each simplex σ in the simplicial complex $\mathcal{K}$. There are two issues associated with the generality of this definition. Firstly, recall that we introduced the clique complex to provide a more succinct description of a simplicial complex. We likewise need the inner product to be succinctly describable, so listing the weights of all simplices is not possible. The second is that we would like the inner product to respect the join operation. That is, after taking the join of two simplicial complexes $\mathcal{K} * \mathcal{L}$, we would like the

inner product on $\bigoplus_{i+j=k-1} C^i(\mathcal{K}) \otimes C^j(\mathcal{L})$ to be induced from those on $C^i(\mathcal{K})$ and $C^j(\mathcal{L})$.

To solve these issues, we add more structure to the definition. Each *vertex* v in the simplicial complex is assigned a weight $w(v)$, and the weights of the higher simplices are induced via

$$w(\sigma) = \prod_{v \in \sigma} w(v). \tag{6.72}$$

This allows a graph $\mathcal{G}$ with weighted *vertices* to induce a weighted clique complex. Note that the edges of $\mathcal{G}$ are still binary (present or not present). This also ensures that, in the join construction, the weight of the tensor product of two simplices σ, τ is the product of the individual weights $w(\sigma \otimes \tau) = w(\sigma)w(\tau)$.

How do the coboundary and boundary operators now act on the weighted complex? Let's first consider the coboundary operator. We would like to transform to bases which are orthonormal in the new inner products. Our new basis for $C^k(\mathcal{K})$ will be $\{|\sigma'\rangle : \sigma \in \mathcal{K}^k\}$, where

$$|\sigma'\rangle = \frac{1}{w(\sigma)}|\sigma\rangle \tag{6.73}$$

is the unit vector of $|\sigma\rangle$. d^k originally acted as

$$d^k|\sigma\rangle = \sum_{v \in \mathrm{up}(\sigma)} |\sigma \cup \{v\}\rangle. \tag{6.74}$$

Written in the new orthonormal basis, this becomes

$$\begin{aligned} d^k|\sigma'\rangle &= \sum_{v \in \mathrm{up}(\sigma)} \frac{w(\sigma \cup \{v\})}{w(\sigma)} |(\sigma \cup \{v\})'\rangle \\ &= \sum_{v \in \mathrm{up}(\sigma)} w(v)|(\sigma \cup \{v\})'\rangle \end{aligned} \tag{6.75}$$

∂^k is defined by

$$\partial^k = (d^{k-1})^\dagger \tag{6.76}$$

which thus acts by

$$\partial^k|\sigma'\rangle = \sum_{v \in \sigma} w(v)|(\sigma \setminus \{v\})'\rangle. \tag{6.77}$$

From here onwards, we drop the primes on the standard orthonormal basis.

We anticipate it will be useful to the reader to explicitly describe the action of the Laplacian Δ^k of a weighted clique complex $\mathcal{G}$. Let σ, τ be two k-simplices. Let's say that σ and τ have a *similar common lower simplex* if we can remove a vertex

v_σ from σ and a vertex v_τ from τ such that we get the same $k-1$-simplex, with the same orientation. Let's say that they have a *dissimilar common lower simplex* if the same holds but with opposite orientation. Let's say that σ and τ are *upper adjacent* if they are lower adjacent and their union forms a $k+1$-simplex. σ and τ being upper adjacent means that they are the faces of a common $k+1$-simplex.

Fact 6.5. *(Similar to [Gol02, Theorem 3.3.4])*

$$\langle\sigma|\Delta^k|\tau\rangle = \begin{cases} \left(\sum_{u\in up(\sigma)} w(u)^2\right) + \\ \left(\sum_{v\in\sigma} w(v)^2\right) + 1 & \text{If } \sigma = \tau. \\ \\ w(v_\sigma)w(v_\tau) & \text{If } \sigma \text{ and } \tau \text{ have a similar common lower simplex and are not upper adjacent. } v_\sigma \text{ and } v_\tau \text{ are the vertices removed from } \sigma \text{ and } \tau \text{ respectively to get the common lower simplex.} \\ \\ -w(v_\sigma)w(v_\tau) & \text{If } \sigma \text{ and } \tau \text{ have a dissimilar common lower simplex and are not upper adjacent. } v_\sigma \text{ and } v_\tau \text{ are the vertices removed from } \sigma \text{ and } \tau \text{ respectively.} \\ \\ 0 & \text{Otherwise. This includes the case that } \sigma \text{ and } \tau \text{ have no common lower simplex, and the case that they are upper adjacent.} \end{cases} \tag{6.78}$$

As a corollary, we can see that the Laplacian Δ^k is a poly(n)-sparse matrix.

Thickening

Given a graph whose clique complex, $\mathcal{K}$, is a triangulation of S^n. We would like to find a graph whose clique complex is topologically $\mathcal{K}\times I$ where $I=[0,1]$. Here, we describe a construction which we call *thickening* which achieves this.

First, the following lemma tells us how to *triangulate* $\mathcal{K}\times I$.

Lemma 6.6. *Let $\mathcal{K}$ be a simplicial complex. Order the vertices $\mathcal{K}^0$. Let $\mathcal{L}$ be the simplicial complex with vertices $\mathcal{L}^0=\mathcal{K}^0\times\{0,1\}$ and simplices*

$$[(u_1,0)(u_2,0)\dots(u_a,0)] \tag{6.79}$$

whenever $[u_1u_2\dots u_a]\in\mathcal{K}$,

$$[(u_1,1)(u_2,1)\dots(u_a,1)] \tag{6.80}$$

whenever $[u_1 u_2 \dots u_a] \in \mathcal{K}$ *and finally,*

$$[(u_1, 0)(u_2, 0) \dots (u_a, 0)(v_1, 1) \dots (v_b, 1)] \tag{6.81}$$

whenever

- $u_1 < \dots < u_a \leq v_1 < \dots < v_b$
- $[u_1 \dots u_a] \in \mathcal{K}$
- $[v_1 \dots v_b] \in \mathcal{K}$
- *if* $u_a = v_1$ *then* $[u_1 \dots u_a v_2 \dots v_b] \in \mathcal{K}$
- *if* $u_a \neq v_1$ *then* $[u_1 \dots u_a v_1 \dots v_b] \in \mathcal{K}$.

Then $\mathcal{L}$ *is a triangulation of* $\mathcal{K} \times I$.

Proof. To demonstrate that $\mathcal{L}$ is a triangulation of $\mathcal{K} \times I$ we must show that $\mathcal{L}$ subdivides $\mathcal{K} \times I$, i.e. that $\mathcal{K} \times I$ is filled by simplices, and no two simplices overlap.

Consider the first point. It is immediate from the definition of $\mathcal{L}$ that if the maximal simplices in $\mathcal{K}$ are n-simplices than the maximal simplices in $\mathcal{L}$ are $n+1$-simplices. The maximal simplices in $\mathcal{L}$ are of the form

$$[(u_1, 0)(u_2, 0) \dots (u_a, 0)(u_a, 1)(v_2, 1) \dots (v_b, 1)], \tag{6.82}$$

where $a + b = n$ Note that every facet of a maximal simplex in $\mathcal{L}$ is either:

- of the form $[(u_1, 0)(u_2, 0) \dots (u_a, 0)]$ – i.e. lies on the boundary $\mathcal{K} \times \{0\}$ of the triangulation
- of the form $[(u_1, 1)(u_2, 1) \dots (u_a, 1)]$ – i.e. lies on the boundary $\mathcal{K} \times \{1\}$ of the triangulation
- shared between at least two maximal simplices

The first two claims are trivial. To see that the third claim is true note that $\mathcal{K}$ is a closed manifold. Therefore consider a maximal simplex $s \in \mathcal{L}$. By the definition of $\mathcal{L}$, there exists a corresponding maximal simplex $s' \in \mathcal{K}$, such that for any vertex v we remove from s to construct a facet there is a corresponding vertex v' we can remove from s'. Since $\mathcal{K}$ is closed we can always add a *different* vertex w' to $s' \setminus v'$

to give a different maximal simplex $t' \in \mathcal{K}$. There will be a corresponding vertex w we can add to the facet of s to give a different maximal facet $t \in \mathcal{L}$, such that t and s share the facet obtained from s by removing v. Therefore there are no 'gaps' in $\mathcal{L}$, it completely fills $\mathcal{K} \times I$.

Consider now the second point—the simplices of $\mathcal{L}$ must not overlap. Assume for contradiction that they do overlap. Then two simplices $s, t \in \mathcal{L}$ must share an intersection which is not a simplex. But it is straightforward to check that any two maximal simplices of $\mathcal{L}$ either do not intersect, or intersect on a simplex. Therefore the simplices of $\mathcal{L}$ do not overlap. □

Finally, we can implement this triangulation with a clique complex using the following lemma.

Lemma 6.7. *Let $\mathcal{K}$ be a clique complex. Order the vertices $\mathcal{K}^0$. Let $\mathcal{G}$ be the graph with vertices $\mathcal{K}^0 \times \{0, 1\}$ and edges*

$$\{((u, 0), (v, 0)) : (u, v) \in \mathcal{K}^1\} \tag{6.83}$$
$$\cup \{((u, 1), (v, 1)) : (u, v) \in \mathcal{K}^1\} \tag{6.84}$$
$$\cup \{((v, 0), (v, 1)) : v \in \mathcal{K}^0\} \tag{6.85}$$
$$\cup \{((u, 0), (v, 1)) : (u, v) \in \mathcal{K}^1, u < v\} \tag{6.86}$$

The clique complex of $\mathcal{G}$ is $\mathcal{L}$ from Lemma 6.6, and in particular has clique complex triangulating $\mathcal{K} \times I$.

Proof. It is straightforward to check that the cliques in $\mathcal{G}$ are precisely those listed in Lemma 6.6. □

Perturbation of subspaces

In this section, we introduce a notion of a perturbation of a subspace. This will be useful in stating a central lemma in the argument, Lemma 6.9, and investigating its consequences.

Definition 6.13. *Consider a subspace $\mathcal{U} \subseteq \mathcal{V}$ of a complex vector space $\mathcal{V}$. Let $\mathcal{U}_\lambda \subseteq \mathcal{V}$ be a family of subspaces indexed by the continuous parameter $\lambda \in [0, 1]$. We say that $\mathcal{U}_\lambda$ is a $O(\lambda)$-perturbation of $\mathcal{U}$ if there exists orthonormal bases $\{|u\rangle\}_u$ for $\mathcal{U}$ and $\{|u, \lambda\rangle\}_u$ for each $\mathcal{U}_\lambda$ such that*

$$|||u, \lambda\rangle - |u\rangle|| = O(\lambda) \ \forall u. \tag{6.87}$$

Using this definition, we prove a two-part lemma which will be useful later.

Lemma 6.8. *Suppose $\mathcal{U}_\lambda$ is a $O(\lambda)$-perturbation of subspace $\mathcal{U} \subseteq \mathcal{V}$. Let Π be the orthogonal projection onto $\mathcal{U}$ and Π_λ orthogonal projection onto $\mathcal{U}_\lambda$ for each λ. Then*

1. *$||\Pi_\lambda - \Pi||_{\text{op}} \leq O(\lambda)$ in operator norm.*
2. *If $|\psi_\lambda\rangle$ is a parametrized family of states such that $\langle\psi_\lambda|\Pi|\psi_\lambda\rangle = O(\lambda^2)$, then $\langle\psi_\lambda|\Pi_\lambda|\psi_\lambda\rangle = O(\lambda^2)$.*

Proof. Using the condition Equation (6.87), we can write

$$|u,\lambda\rangle = |u\rangle + O(\lambda)|\tilde{u}_\lambda\rangle \tag{6.88}$$

for some normalized vector $|\tilde{u}_\lambda\rangle$. Then

$$\begin{aligned}
\Pi_\lambda &= \sum_u |u,\lambda\rangle\langle u,\lambda| && (6.89)\\
&= \sum_u \big(|u\rangle + O(\lambda)|\tilde{u}_\lambda\rangle\big)\big(\langle u| + O(\lambda)\langle\tilde{u}_\lambda|\big) && (6.90)\\
&= \Big(\sum_u |u\rangle\langle u|\Big) + O(\lambda)\sum_u |u\rangle\langle\tilde{u}_\lambda| + O(\lambda)\sum_u |\tilde{u}_\lambda\rangle\langle u| + O(\lambda^2) && (6.91)\\
&= \Pi + O(\lambda)\sum_u |u\rangle\langle\tilde{u}_\lambda| + O(\lambda)\sum_u |\tilde{u}_\lambda\rangle\langle u| + O(\lambda^2). && (6.92)
\end{aligned}$$

We can immediately read off Part 1, that $||\Pi_\lambda - \Pi|| \leq O(\lambda)$ in operator norm. As for Part 2,

$$\langle\psi_\lambda|\Pi_\lambda|\psi_\lambda\rangle = \langle\psi_\lambda|\Pi|\psi_\lambda\rangle + O(\lambda)\sum_u\langle\psi_\lambda|u\rangle\langle\tilde{u}_\lambda|\psi_\lambda\rangle + O(\lambda)\sum_u\langle\psi_\lambda|\tilde{u}_\lambda\rangle\langle u|\psi_\lambda\rangle + O(\lambda^2) \tag{6.93}$$

$$= O(\lambda^2) + O(\lambda)\sum_u\langle\psi_\lambda|u\rangle\langle\tilde{u}_\lambda|\psi_\lambda\rangle + O(\lambda)\sum_u\langle\psi_\lambda|\tilde{u}_\lambda\rangle\langle u|\psi_\lambda\rangle \tag{6.94}$$

$$\leq O(\lambda^2) + O(\lambda)\Big(\sum_u |\langle u|\psi_\lambda\rangle|^2\Big)^{\frac{1}{2}}\Big(\sum_u |\langle\tilde{u}_\lambda|\psi_\lambda\rangle|^2\Big)^{\frac{1}{2}} \tag{6.95}$$

$$\leq O(\lambda^2) + O(\lambda)\Big(\langle\psi_\lambda|\Pi|\psi_\lambda\rangle\Big)^{\frac{1}{2}} \tag{6.96}$$

$$\leq O(\lambda^2). \tag{6.97}$$

We used the assumption $\langle\psi_\lambda|\Pi|\psi_\lambda\rangle = O(\lambda^2)$, and then Cauchy-Schwarz, and finally the assumption again. □

6.2 Hamiltonian to homology gadgets

Given an instance of quantum m-SAT H on n qubits, we are aiming to find a weighted graph $\mathcal{G}$ and some k such that the ground energy of the k^{th}-order Laplacian Δ^k is related to the ground energy of H. We will do this by first constructing a *qubit graph* $\mathcal{G}_n$ whose harmonic states can be identified with qubit states $|\psi\rangle \in (\mathbb{C}^2)^{\otimes n}$.[8] This means the k-homology of $\mathcal{G}_n$ will be 2^n-dimensional. The clique complex of this graph will have no $k+1$-simplices, so $\Delta^{\uparrow k} = 0$.

The m-local Hamiltonian H is a sum of rank-1 projectors

$$H = \sum_{i=1}^{t} \phi_i, \tag{6.98}$$

where each $\phi_i = |\phi_i\rangle\langle\phi_i|$ for a state $|\phi_i\rangle$ on at most m qubits.

Definition 6.14. *$|\phi\rangle$ is an* integer state *if it can be written as*

$$|\phi\rangle = \frac{1}{\mathcal{Z}} \sum_{z\in\{0,1\}^m} a_z |z\rangle, \tag{6.99}$$

where $a_z \in \mathbb{Z}$ are integers, and $\mathcal{Z} = (\sum_z |a_z|^2)^{\frac{1}{2}}$ is a normalization factor.

We will assume $\{|\phi_i\rangle\}$ are integer states. For each term ϕ_i we will add a *gadget* $\mathcal{T}_i$ to the graph which aims to implement the effect of this term on the groundspace of harmonic states. $\mathcal{T}_i$ will consist of additional vertices and edges which add $k+1$-simplices to the clique complex so that $\Delta^{\uparrow k}$ implements ϕ_i. Topologically, this is achieved by designing $\mathcal{T}_i$ to be a triangulation of a $k+1$-manifold whose *boundary* is the cycle corresponding to $|\phi_i\rangle$.

In our construction, the clique complex will be *weighted* as described in Section 6.1. Recall that this is done by assigning weights to the vertices of the underlying graph. The vertices of the original qubit graph $\mathcal{G}_n$ will all have weight 1. The vertices of the gadgets $\mathcal{T}_i$ will have weight $\frac{1}{\text{poly}n} \leq \lambda \ll 1$. The weights are much smaller than 1, although only polynomially so.

Qubit graph

Let $\mathcal{G}_1$ be the *bowtie graph* (see Figure 6.5). The clique complex has no 2-simplices, and the 1-homology is isomorphic to $\mathbb{C}^2$, spanned by equivalence classes of the

[8] Here we do not use the same graph as in [CK22] because in order to maintaing a link between the spectrum of the Laplacian and the spectrum of the Quantum 4-SAT Hamiltonian we are reducing from we need our basis states to correspond to orthogonal holes in the clique complex, not just distinct homology classes.

left and right loops. Further, the loops themselves form an orthonormal basis of the harmonic states $\mathrm{Ker}\,\Delta^1 \subseteq C^1$. Identify $|0\rangle$ with the left loop, (i.e. the cycle $[xa_3] + [a_3a_2] + [a_2a_4] + [a_4x]$) and $|1\rangle$ with the right loop (i.e. the cycle $[xb_3] + [b_3b_2] + [b_2b_4] + [b_4x]$). Note that $C^1 \cong \mathbb{C}^8$ since there are 8 1-simplices (edges). The other 6 states have some higher energies on Δ^1.

To construct $\mathcal{G}_n$, we employ the *join* operation.

$$\mathcal{G}_n = \mathcal{G}_1 * \cdots * \mathcal{G}_1 \quad (n \text{ times}). \tag{6.100}$$

Using the Kunneth formula Fact 6.3, we see that $H^{2n-1} \cong (\mathbb{C}^2)^{\otimes n}$. Moreover, the computational qubit states $|z_1\rangle \otimes \cdots \otimes |z_n\rangle$ for $z_i \in \{0, 1\}$ form a natural orthonormal basis for the groundspace $\mathrm{Ker}\,\Delta^{2n-1} \subseteq C^{2n-1}$. Note this means that we take $k = 2n - 1$. Denote

$$\mathcal{H}_n := \mathrm{Ker}\,\Delta^{2n-1} \cong (\mathbb{C}^2)^{\otimes n}. \tag{6.101}$$

What is the cycle corresponding to a computational basis state $|z\rangle, z \in \{0, 1\}^n$? $|z\rangle$ is the join of n copies of the square loop. The square loop is a two-fold join $\mathfrak{g}_2 = \mathfrak{g}_1 * \mathfrak{g}_1$ from Section 6.1. Thus $|z\rangle$ is a copy of the $(2n - 1)$-dimensional octahedron $\mathfrak{g}_{2n}$. $|z\rangle$ is topologically homeomoprhic to S^{2n-1}.

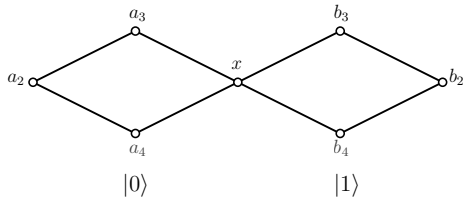

Figure 6.5: The clique-complex that encodes a single-qubit.

Single gadget

We now describe the gadget for a single term ϕ where $\phi = |\phi\rangle\langle\phi|$. For this section we drop the subscript i labelling the terms.

Let's first work on implementing the Hamiltonian term ϕ on the harmonic subspace $\mathcal{H}_m \subseteq C^{2m-1}(\mathcal{G}_m)$, $\mathcal{H}_m \cong (\mathbb{C}^2)^{\otimes m}$ of the m-qubit graph. This will involve adding vertices and edges to $\mathcal{G}_m$, and these additional vertices and edges make up the gadget. Let the final graph be denoted $\hat{\mathcal{G}}_m$. Eventually we can then add the remaining $n - m$ qubits and find the relevant gadget for $\mathcal{G}_n$ by joining $\mathcal{G}_{n-m}$

$$\hat{\mathcal{G}}_n = \hat{\mathcal{G}}_m * \mathcal{G}_{n-m}. \tag{6.102}$$

Motivated by the equation $\langle\psi|\Delta^{\uparrow k}|\psi\rangle = ||d^k|\psi\rangle||^2$, we would like the gadget to be a *triangulation* of a *manifold* whose *boundary* is precisely the cycle $|\phi\rangle$. This can be viewed topologically as *filling in* the cycle $|\phi\rangle$. Throughout this section $\mathcal{J}$ refers to the cycle we want to fill in with the gadget (i.e. if we are constructing a gadget for the projector $|\phi\rangle\langle\phi|$ then $\mathcal{J}$ is the cycle that corresponds to the state $|\phi\rangle$), and $\mathcal{J}^0$ refers to the vertices in $\mathcal{J}$. The procedure is as follows.

1. Construct a clique complex $\mathcal{K}$ which is a triangulation of S^{2m-1} and a relation:

$$R \subseteq \{(z_i, z_j)|z_i, z_j \in \mathcal{K}^0\} \tag{6.103}$$

   such that the map

$$\begin{aligned} &f : \mathcal{K}^0 \to \mathcal{J}^0 \\ &f(z_i) = z_j \ \text{ where } \ (z_i, z_j) \in R \end{aligned} \tag{6.104}$$

   is a surjective function, and the simplicial complex given by

$$\{f(\sigma) \cap \mathcal{J}^0 | \sigma \in \mathcal{K}\} \tag{6.105}$$

   is a copy $\mathcal{J}$ of the cycle $|\phi\rangle$.[9]

2. Let $\mathcal{L}$ be the thickening of $\mathcal{K}$ described by Lemma 6.6. The vertices of $\mathcal{L}$ are $\mathcal{L}^0 = \mathcal{K}^0 \times \{0, 1\}$. By Lemma 6.7, $\mathcal{L}$ is the clique complex of some graph. This creates a thickened S^{2m-1} shell with outside layer $\mathcal{K}^0 \times \{0\}$ and inside layer $\mathcal{K}^0 \times \{1\}$.

3. Add a central vertex v_0 which connects to all vertices of the inside layer $\mathcal{K}^0 \times \{1\}$. By the definition of clique complexes (Chapter 6), this automatically introduces into the complex all the simplices which contain these new edges. We will refer to this process as *coning off the cycle*.[10] Denote the simplicial complex at this step by $\hat{\mathcal{K}}$.

4. Let $\mathcal{V} = \{\mathcal{J}^0 \cup \mathcal{K}^0 \times \{1\} \cup v_0\}$. Apply $f(\cdot)$ to $\mathcal{K}^0 \times \{0\}$ and construct the cycle:

$$\tilde{\mathcal{K}} = \{f(\sigma) \cap \mathcal{V} | \sigma \in \hat{\mathcal{K}}\} \tag{6.106}$$

   as set out in eq. (6.105) to get $\mathcal{J}$, a copy of the cycle $|\phi\rangle$.

[9] See the end of this section for a an example of how the function f is applied and conditions on R for f to be a surjective function.

[10] The terminology arises from the concept of constructing a *mapping cone*—the cycle $\mathcal{L}$ can be represented as a map of a sphere onto the complex the process of adding this central vertex and the associated simplices can be seen as constructing the mapping cone.

5. The weights of the vertices in the outer layer $\mathcal{J}^0$, which will be identified with vertices in the qubit graph $\mathcal{G}_m$, remain set to 1.

6. The weights of the inside layer $\mathcal{K}^0 \times \{1\}$ and the central vertex v_0 are set to λ where $\frac{1}{\text{poly}n} \leq \lambda \ll 1$.

7. The above forms the gadget. It remains to simply glue $\mathcal{J}$ onto the cycle $|\phi\rangle \in \text{Cl}(\mathcal{G}_m)$ by identifying the vertices in $\mathcal{J}^0$ with the equivalent vertices in $\text{Cl}(\mathcal{G}_m)$. The resulting complex is the clique complex of some new graph. Let the new graph with the gadget glued in be denoted $\hat{\mathcal{G}}_m$.

The above procedure provides a method to fill-in any cycle, provided it is possible to complete Item 1. For basis states Item 1 is trivial—the cycles themselves are already triangulations of S^{2m-1} and clique complexes. For arbitrary states we do not have a general method for carrying out the procedure. However, for integer states (see Definition 6.14) we can construct a general method for Item 1. The method relies on the simplicial surgery techniques, first introduced in [CK22].

Note on the function $f(\cdot)$: The function $f(\cdot)$ acts on the *vertices* of a simplicial complex. We then construct a new simplicial complex out of the new set of vertices, according to the method set out in eq. (6.105). To see how this works explicitly we consider a simple example. Take a simplicial complex, specified by its maximal faces:

$$K = \{[x_0x_1x_2x_3], [x'_0x'_1x'_2x_3], [x_2x_3x_4x_5], [x_0x'_1x'_2, x_3], [x'_0x'_1x_1x'_2x_2x_3x_4x_5]\} \tag{6.107}$$

Define a relation:

$$R = \{(x'_i, x_i)|i \in [0, 2]\} \cup \{(x_i, x_i)|i \in [0, 5]\} \tag{6.108}$$

Acting with $f(\cdot)$ on K then gives a simplicial complex defined by maximal faces:

$$f(K) = \{[x_0x_1x_2x_3], [x_2x_3x_4x_5], [x_0x_1x_2, x_3], [x_0x_1x_2x_3x_4x_5]\} \tag{6.109}$$

Note that any simplices $\sigma \in K$ which do not contain any pairs from the relation R map to simplices of the same dimension under $f(\cdot)$ (the vertices may have changed, but the number of vertices in the simplex is unchanged). Simplices $\sigma \in K$ which contain one or more pairs of vertices from R will map under $f(\cdot)$ to lower dimensional simplices, because simplices cannot contain two of the same vertex, so the number of vertices remaining in the simplex has decreased.

Note on the relation R: In order for $f(\cdot)$ to be a surjective function $f : \mathcal{K}^0 \to J^0$ as required the relation R must satisfy:

- R must be functional, i.e. for all $x, y, z \in \mathcal{K}^0$, $(x, y) \in R$ and $(x, z) \in R$ implies $y = z$
- for all $y \in \mathcal{J}^0$ there must exist $x \in \mathcal{K}^0$ such that $(x, y) \in R$

6.3 Spectral sequences

The purpose of this section is to prove the following lemma, which concerns the spectrum of the Laplacian of a single gadget $\hat{\mathcal{G}}_m$ from Section 6.2.

Lemma 6.9. (Single gadget lemma) *Let $\hat{\mathcal{G}}_m$ be the weighted graph described in Section 6.2, implementing the projector onto the integer state $|\phi\rangle$ on m qubits. Let $\hat{\Delta}^k$ be the Laplacian of this graph. Recall the definition of a subspace perturbation from Section 6.1.*

- *$\hat{\Delta}^{2m-1}$ has a $(2^m - 1)$-dimensional kernel, which is a $O(\lambda)$-perturbation of the subspace $\{|\psi\rangle \in \mathcal{H}_m : \langle\phi|\psi\rangle = 0\}$. Note $\mathcal{H}_m$ is embedded as $\mathcal{H}_m \subseteq C^{2m-1}(\mathcal{G}_m) \subseteq C^{2m-1}(\hat{\mathcal{G}}_m)$.*
- *The first excited state $|\hat{\phi}\rangle$ of $\hat{\Delta}^{2m-1}$ above the kernel is a $O(\lambda)$-perturbation of $|\phi\rangle \in \mathcal{H}_m$, and it has energy $\Theta(\lambda^{4m+2})$.*
- *The next lowest eigenvectors have eigenvalues $\Theta(\lambda^2)$, and they are $O(\lambda)$-perturbations of sums of $(2m-1)$-simplices touching the central vertex v_0.*
- *The rest of the eigenvalues are $\Theta(1)$.*

Spectral sequences are a tool from algebraic topology which (among other things) analyze the homology of *filtered* chain complexes. It turns out that weighting a subset of vertices by $\lambda \ll 1$, as we do in our construction, naturally gives rise to a certain filtration. In this setting, Ref. [For94] showed a beautiful relationship between the spectral sequence and the perturbative eigenspaces of the Hodge theoretic Laplacian. It is this relationship which we exploit in this section to prove Lemma 6.9. For a light introduction to spectral sequences, see [Cho06]; for a comprehensive textbook, see [McC01].

Spectral sequence of a filtration

Spectral sequences will work best for us in the *cohomology* picture. First we define a *filtration*.

Definition 6.15. *Suppose we have a cohomological chain complex*

$$C^{-1} \xrightarrow{d^{-1}} C^0 \xrightarrow{d^0} C^1 \xrightarrow{d^1} C^2 \xrightarrow{d^2} \dots \tag{6.110}$$

A filtration *on this chain complex is a nested sequence of subspaces*

$$C^k = \mathcal{U}_0^k \supseteq \mathcal{U}_1^k \supseteq \mathcal{U}_2^k \supseteq \dots \tag{6.111}$$

for each n such that

$$d^k(\mathcal{U}_l^k) \subseteq \mathcal{U}_l^{k+1} \; \forall \; k, l. \tag{6.112}$$

Our filtrations will be bounded, in the sense that $U_j^k = \{0\}$ for sufficiently large j, for each k.

We can now develop the spectral sequence of such a filtration. The spectral sequences will consist of *pages* indexed by j. Each page is an array of vector spaces $e_{j,l}^k$, one for each dimension k and filtration level l.

The zeroth page is simply

$$e_{0,l}^k = \mathcal{U}_l^k / \mathcal{U}_{l+1}^k. \tag{6.113}$$

The chain complex coboundary map d^k induces coboundary maps

$$d_{0,l}^k : e_{0,l}^k \to e_{0,l}^{k+1} \tag{6.114}$$

since if two chains differ by an element of $\mathcal{U}_{l+1}^k$, then their coboundaries will differ by an element of $\mathcal{U}_{l+1}^{k+1}$.

Define the first page of the spectral sequence to be the cohomology of the zeroth page with respect to $d_{0,l}^k$, entrywise for each k, l.

$$e_{1,l}^k = \operatorname{Ker} d_{0,l}^k / \operatorname{Im} d_{0,l}^{k-1} \tag{6.115}$$

Now the coboundary map d^k induces coboundary maps

$$d_{1,l}^k : e_{1,l}^k \to e_{1,l+1}^{k+1} \tag{6.116}$$

This is because (a) the coboundary of any representative of an element in $\operatorname{Ker} d_{0,l}^k$ is a cocycle in $\mathcal{U}_{l+1}^{k+1}$, and thus is the representative of some element of $\operatorname{Ker} d_{0,l+1}^{k+1}$;

(b) if we chose a different representative of the element in $\operatorname{Ker} d^k_{0,l}$, the resulting element of $\operatorname{Ker} d^{k+1}_{0,l+1}$ would differ by an element of $\operatorname{Im} d^k_{0,l+1}$, so we end up with the same element of $e^{k+1}_{1,l+1}$; and (c) if our element of $\operatorname{Ker} d^{k+1}_{0,l+1}$ differed by an element of $\operatorname{Im} d^{k-1}_{0,l}$, we get the exact same element of $\operatorname{Ker} d^{k+1}_{0,l+1}$.

In general, at page j there are induced coboundary maps

$$d^k_{j,l} : e^k_{j,l} \to e^{k+1}_{j,l+j} \tag{6.117}$$

and page $j+1$ is defined to be the cohomology of page j entrywise with respect to these coboundary maps

$$e^k_{j+1,l} = \operatorname{Ker} d^k_{j,l} / \operatorname{Im} d^{k-1}_{j,l-j}. \tag{6.118}$$

One should have in mind the entries of a single page j laid out in an array as follows. On this array, $d^k_{j,l}$ will map from a space to the one which is one step upwards and j steps to the right.

k	:			
1	$e^1_{j,0}$	$e^1_{j,1}$	$e^1_{j,2}$	
0	$e^0_{j,0}$	$e^0_{j,1}$	$e^0_{j,2}$	
-1	$e^{-1}_{j,0}$	$e^{-1}_{j,1}$	$e^{-1}_{j,2}$	$\dots$
	0	1	2	l

We can express the spaces $e^k_{j,l}$ more explicitly.

Definition 6.16. *Let $C^k = \mathcal{U}^k_0 \supseteq \mathcal{U}^k_1 \supseteq \mathcal{U}^k_2 \supseteq \dots$ be a filtered chain complex with coboundary d. Define*

$$Z^k_{j,l} = \mathcal{U}^k_l \cap (d^k)^{-1}(\mathcal{U}^{k+1}_{l+j}) \tag{6.119}$$

$$B^k_{j,l} = \mathcal{U}^k_l \cap d^{k-1}(\mathcal{U}^{k-1}_{l-j}). \tag{6.120}$$

With these definitions in place, it turns out that the terms $e^k_{j,l}$ of the spectral sequence are equal to

$$e^k_{j,l} = Z^k_{j,l} / (B^k_{j-1,l} + Z^k_{j-1,l+1}). \tag{6.121}$$

Let $e_j^k = \bigoplus_l e_{j,l}^k$. The point of spectral sequences is that, for sufficiently large j, e_j^k is isomorphic to the cohomology of the complex

$$e_\infty^k \cong H^k = \text{Ker}\, d^k / \text{Im}\, d^{k-1}. \tag{6.122}$$

The e_j^k spaces form an algebraic sequence of approximations to the true cohomology.

Relationship to Hodge theory

From our construction we have a clique complex $\mathcal{G}$ where a subset of the vertices are weighted by $\lambda \ll 1$, and the rest by 1. The weight of a simplex is defined to be the product of the weights of the vertices involved in the simplex. In this context, there is a natural filtration on the chain complex C. Let

$$\mathcal{U}_l^k = \text{span}\{\sigma \in \mathcal{G}^k : w(\sigma) \in \{\lambda^l, \lambda^{l+1}, \dots\}\}. \tag{6.123}$$

This is the span of the k-simplices which involve at least l 'gadget' vertices. These spaces are nested

$$C^k = \mathcal{U}_0^k \supseteq \mathcal{U}_1^k \supseteq \mathcal{U}_2^k \supseteq \dots \tag{6.124}$$

and it can also be checked that

$$d^k(\mathcal{U}_l^k) \subseteq \mathcal{U}_l^{k+1}. \tag{6.125}$$

Thus we have a filtration.

We are interested in the low energy eigenstates of the Laplacian Δ^k of $\mathcal{G}$. In particular, we would like to examine the eigenvalues which are zero to first order in λ, and then second order, and so on. This is reminiscent of perturbation theory from quantum mechanics. Recall that

$$\langle\psi|\Delta^k|\psi\rangle = ||\partial^k|\psi\rangle||^2 + ||d^k|\psi\rangle||^2. \tag{6.126}$$

Motivated by this, define the isomorphism

$$\rho_\lambda^k : C^k \to C^k \tag{6.127}$$

$$|\sigma\rangle \to w(\sigma)|\sigma\rangle, \tag{6.128}$$

where $\sigma \in \mathcal{G}^k$ is a k-simplex which has weight $w(\sigma)$. Morally, ρ_λ^k maps from the weighted chainspace to the unweighted chainspace. From this, construct the maps

$$\partial_\lambda^k = \rho_\lambda^k \circ \partial^k \circ (\rho_\lambda^k)^{-1} \tag{6.129}$$

$$d_\lambda^k = \rho_\lambda^k \circ d^k \circ (\rho_\lambda^k)^{-1}. \tag{6.130}$$

Define the spaces

$$E_j^k = \{|\psi\rangle \in C^k \ : \ \exists\, |\psi_\lambda\rangle = |\psi\rangle + \lambda|\psi_1\rangle + \lambda^2|\psi_2\rangle + \cdots + \lambda^j|\psi_j\rangle \in C^k[\lambda] \tag{6.131}$$

$$\text{s.t. } \partial_\lambda^k|\psi_\lambda\rangle \in \lambda^j C^{k-1}[\lambda],\ d_\lambda^k|\psi_\lambda\rangle \in \lambda^j C^{k+1}[\lambda]\}, \tag{6.132}$$

where $C^k[\lambda]$ is the space of polynomials in λ with coefficients in C^k. Further define

$$E_{j,l}^k = E_j^k \cap \mathcal{U}_l^k \cap (\mathcal{U}_{l+1}^k)^\perp \tag{6.133}$$

so that

$$E_j^k = \oplus_l E_{j,l}^k. \tag{6.134}$$

E_j^k is the space of vectors which have perturbations in λ which give energies of size $O(\lambda^{2j})$ on the Laplacian Δ^k. Here and throughout this section, $O(\lambda^l)$ is used as shorthand for polynomials in λ which contain no terms of degree less than l. Taking $j \to \infty$ should give $E_\infty^k = H^k(C)$, and these spaces E_j^k form a Hodge-theoretic sequence of approximations to the true homology.

Proposition 6.10. *The space*

$$\{|\psi\rangle \in C^k \ : \ |\psi\rangle \text{ is an eigenvector of } \Delta^k \text{ with eigenvalue } O(\lambda^{2j})\} \tag{6.135}$$

is a $O(\lambda)$-perturbation of E_j^k, in the sense of Section 6.1.

Proof. This follows from Theorem 2 of [For94], combined with Rellich's theorem stated at the end of the introduction of [For94]. □

Theorem 7 from [For94] tells us remarkably that these E_j^k spaces are *isomorphic* to the e_j^k spaces of the filtration. The E_j^k spaces are our real objects of interest, but the e_j^k spaces are tractable to calculate. This is reminiscent of a standard methodology in algebraic topology where we prove difficult topological and analytic statements by turning them into algebraic statements. This will become our strategy to prove Lemma 6.9. The rest of this section is devoted to describing the isomorphism.

Recall Definition 6.16. $\lambda^{-l}\rho_\lambda^k$ creates isomorphisms

$$Z_{j,l}^k \to \widetilde{Z}_{j,l}^k \tag{6.136}$$

$$:= \{|\psi_\lambda\rangle = |\psi_0\rangle + \lambda|\psi_1\rangle + \lambda^2|\psi_2\rangle + \dots \ : \ |\psi_i\rangle \in \mathcal{U}_{l+i}^k,\ d_\lambda|\psi_\lambda\rangle \in \lambda^j C[\lambda]\} \tag{6.137}$$

$$Z_{j-1,l+1}^k \to \lambda\widetilde{Z}_{j-1,l+1}^k \tag{6.138}$$

$$:= \{|\psi_\lambda\rangle = \lambda|\psi_1\rangle + \lambda^2|\psi_2\rangle + \dots \ : \ |\psi_i\rangle \in \mathcal{U}_{l+i}^k,\ d_\lambda|\psi_\lambda\rangle \in \lambda^j C[\lambda]\} \tag{6.139}$$

$$B_{j-1,l}^k \to \widetilde{B}_{j-1,l}^k \tag{6.140}$$

$$:= \lambda^{-j+1} d_\lambda\{|\psi_\lambda\rangle = |\psi_0\rangle + \lambda|\psi_1\rangle + \lambda^2|\psi_2\rangle + \dots \ : \ |\psi_i\rangle \tag{6.141}$$

$$\in \mathcal{U}_{l-j+1+i}^{k-1},\ d_\lambda|\psi_\lambda\rangle \in \lambda^{j-1}C[\lambda]\} \tag{6.142}$$

$$= \lambda^{-j+1} d_\lambda \widetilde{Z}_{j-1,l-j+1}^k. \tag{6.143}$$

Let $|\psi\rangle \in E_{j,l}^k$. Let

$$|\psi_\lambda\rangle = |\psi\rangle + \lambda|\psi_1\rangle + \lambda^2|\psi_2\rangle + \dots \tag{6.144}$$

be the polynomial from the definition of E_j^k. Let $|\psi_{i,l}\rangle$ denote the projection of $|\psi_i\rangle$ onto $\mathcal{U}_l^k$, and write

$$|\psi_\lambda\rangle = \left[|\psi\rangle + \sum_{i>0} \lambda^i |\psi_{i,l+i}\rangle\right] + \left[\sum_{i>0}\sum_{c\neq i} \lambda^i |\psi_{i,l+c}\rangle\right] \tag{6.145}$$

$$= |\psi_\lambda^{(0)}\rangle + |\psi_\lambda^{(1)}\rangle. \tag{6.146}$$

Proposition 6.11. (Theorem 7 from [For94])

$|\psi_\lambda^{(0)}\rangle \in \widetilde{Z}_{j,l}^k$ and the map

$$E_{j,l}^k \to \widetilde{Z}_{j,l}^k / (\lambda\widetilde{Z}_{j-1,l+1}^k + \widetilde{B}_{j-1,l}^k) \cong e_{j,l}^k \tag{6.147}$$

$$|\psi\rangle \to [|\psi_\lambda^{(0)}\rangle] \tag{6.148}$$

is an isomorphism. The corresponding map from $[|\psi\rangle] \in e_{j,l}^k$ to $E_{j,l}^k$ is to project the representative $|\psi\rangle$ onto $\mathcal{U}_l^k$.

Example spectral sequence

In this section, we get some practice with spectral sequences by calculating the spectral sequence of the weighted complex shown below. The vertices on the perimeter have weight 1, and the vertices in the interior have weight $\lambda \ll 1$. This is a filling in of a hexagon, and we are concerned with the $k = 1$ homology. This

complex does not correspond to any Hamiltonian on any number of qubits, but rather it is like a rank-1 projector on a 1-dimensional Hilbert space. Regardless, let's refer to the weight 1 vertices on the boundary as *qubit* vertices, and the weight λ vertices in the bulk as *gadget* vertices. It will serve as a simple example which highlights some key features which will be present in the general gadget construction.

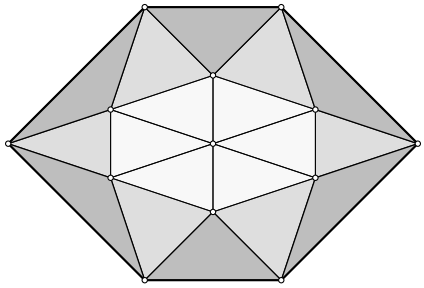

Figure 6.6: The complex for the example spectral sequence. 2-simplices are shown darker and bolder if they are more heavily weighted.

The relevant chainspaces are

$$C^2 = \mathcal{U}_0^2 \supseteq \mathcal{U}_1^2 \supseteq \mathcal{U}_2^2 \supseteq \mathcal{U}_3^2 \tag{6.149}$$

$$C^1 = \mathcal{U}_0^1 \supseteq \mathcal{U}_1^1 \supseteq \mathcal{U}_2^1 \tag{6.150}$$

$$C^0 = \mathcal{U}_0^0 \supseteq \mathcal{U}_1^0. \tag{6.151}$$

The chain C^2 is truncated at $\mathcal{U}_3^2$ since there are no triangles with more than 3-gadget vertices, and similarly for edges C^1 and vertices C^0.

The zeroth page of the spectral sequence is $e_{0,l}^k$ for $k = -1, 0, 1, 2$, $l = 0, \ldots, k+1$. $e_{0,l}^k$ can be thought of as the space spanned by the k-simplices of weight λ^l. There are no triangles consisting only of qubit vertices, so $e_{0,0}^2 = \{0\}$. Below we see pictorial representations of Page 0.

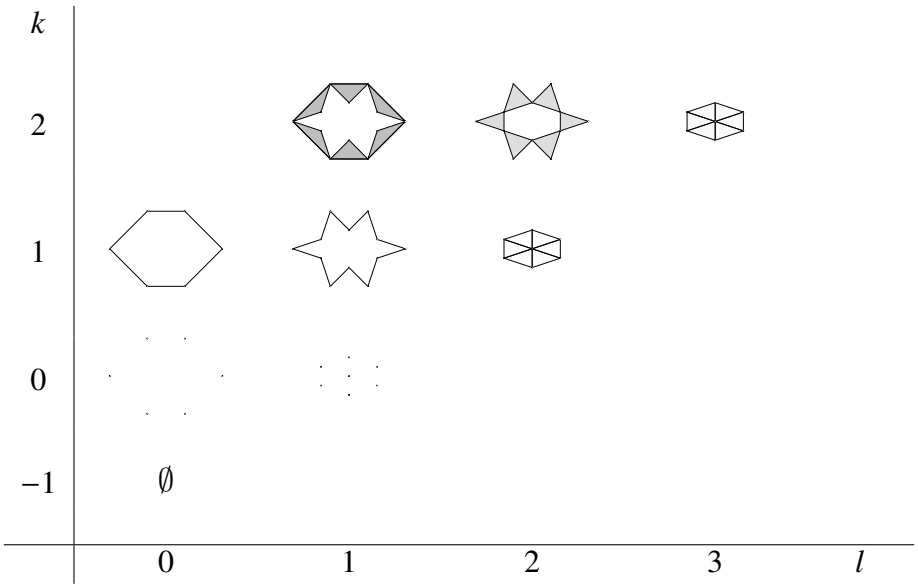

The coboundary map of the zeroth page maps

$$d^k_{0,l} : e^k_{0,l} \to e^{k+1}_{0,l}. \tag{6.152}$$

It maps upwards one step in the diagram, and it can be thought of as a qubit vertex coboundary map, which adds a qubit vertex to the simplex. If no qubit vertex can be added to the k-simplex σ, then $d^k_{0,l}|\sigma\rangle = 0$.

We are now in a position to compute the first page $e^k_{1,l}$ of the spectral sequence, which is defined as

$$e^k_{1,l} = \mathrm{Ker}\, d^k_{0,l} / \mathrm{Im}\, d^{k-1}_{0,l}. \tag{6.153}$$

It may be more intuitive to bear in mind that this is the same as

$$e^k_{1,l} = \mathrm{Ker}\, \partial^k_{0,l} / \mathrm{Im}\, \partial^{k+1}_{0,l}, \tag{6.154}$$

where

$$\partial^k_{0,l} : e^k_{0,l} \to e^{k-1}_{0,l} \tag{6.155}$$

is the qubit boundary map, which removes qubit vertices. (Formally, we can define $\partial^k_{0,l} = (d^{k-1}_{0,l})^\dagger$.)

The first column $l = 0$, we simply get the homology of the qubit complex, which is topologically a single loop S^1. Thus $e^0_{1,0} = e^2_{1,0} = \{0\}$, but $e^1_{1,0}$ is the 1-dimensional space spanned by the loop of qubit vertices. For the purposes of this section, denote

this state by $|\text{loop}\rangle$. The $e^1_{0,1}$ looks like it also has a loop—that is, a 1-chain without boundary. However, this 1-chain is in fact the boundary of the uniform superposition of triangles in $e^2_{0,1}$, so the homology at this position is zero $e^1_{1,1} = \{0\}$. The triangle chainspaces $e^2_{0,1}$, $e^2_{0,2}$ have no cycles, so their homologies are zero $e^2_{1,1} = e^2_{1,2} = \{0\}$. In $e^0_{0,1}$, the only gadget vertex which is *not* the boundary of an edge in $e^1_{0,1}$ is the central vertex, so $e^0_{1,1}$ is the 1-dimensional space spanned by this vertex. Similarly, the edges on the 'outside' of $e^1_{0,2}$ are the boundaries of triangles in $e^2_{0,2}$, so the ones which survive in the homology $e^2_{1,2}$ are the spokes of the 'star' touching the central vertex. For $e^2_{1,3}$, both boundary maps involved are zero, so $e^2_{1,3} = e^2_{0,3}$ i.e. all the central triangles. See the below pictorial representations of the first page.

Page 1

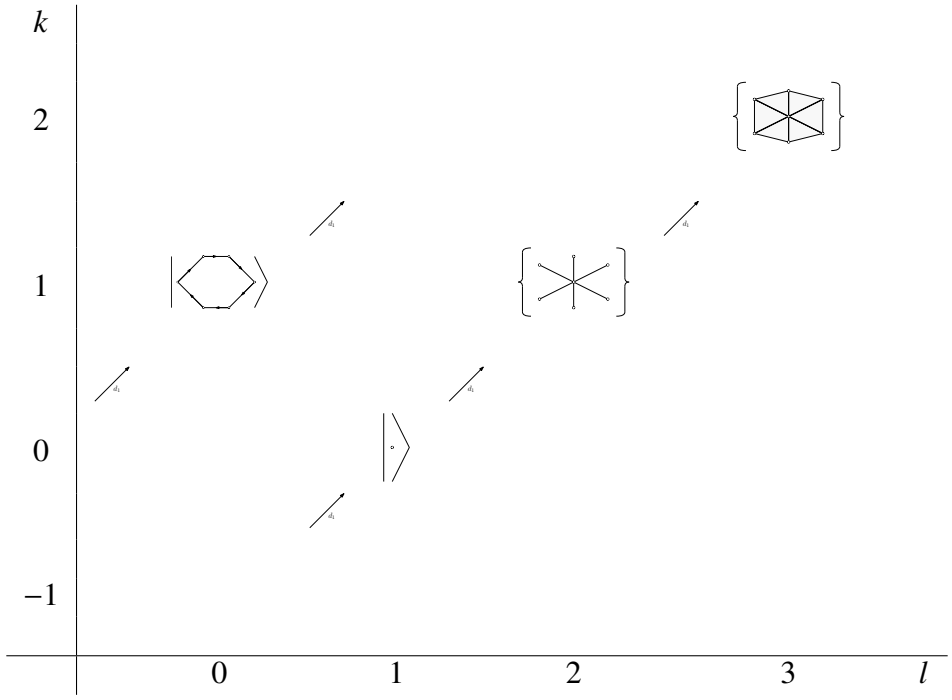

The coboundary map of the first page maps

$$d^k_{1,l} : e^k_{1,l} \to e^{k+1}_{1,l+1}. \tag{6.156}$$

It maps 'diagonally' one step up and one step to the right in the diagram. It can be thought of as a gadget vertex coboundary map, which adds a gadget vertex to the simplex. If no gadget vertex can be added to the k-simplex σ, then $d^k_{1,l}|\sigma\rangle = 0$. As before, it may be more intuitive to visualize the gadget vertex boundary map $\partial^k_{1,l} = (d^{k-1}_{1,l-1})^\dagger$.

Using this, we aim to compute the second page $e^k_{2,l}$ of the spectral sequence, which is defined as

$$e^k_{2,l} = \operatorname{Ker} d^k_{1,l}/\operatorname{Im} d^{k-1}_{1,l-1} = \operatorname{Ker} \partial^k_{1,l}/\operatorname{Im} \partial^{k+1}_{1,l+1}. \tag{6.157}$$

Both boundary maps acting on $e^1_{1,0}$ are zero, so $e^1_{2,0} = e^1_{1,0}$, which recall is the span of the state $|\text{loop}\rangle$. Apart from $e^1_{1,0}$, we have a 3-chain

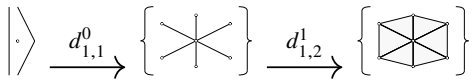

(6.158)

$e^0_{1,1}$ contains only the central vertex v_0, which is *not* in the kernel of $d^0_{1,1}$, so $e^0_{2,1} = \{0\}$. $e^1_{1,2}$ consists of the star of edges touching v_0. The coboundaries of these edges has support on the triangles wedged in between them. The only way for a superposition of these edges to have coboundaries cancelling on each of these triangles is to be proportional to the uniform superposition. But this state is precisely the coboundary of the central vertex $d^0_{1,1}|v_0\rangle$. Thus the middle term has no homology and $e^1_{2,2} = \{0\}$. It remains to calculate the space $e^2_{2,3}$. By counting dimensions, there should be a single state in the homology of $e^2_{1,3}$. This is because there are the same number of edges spanning $e^1_{1,2}$ as triangles spanning $e^2_{1,3}$, and precisely one state in $e^1_{1,2}$ was killed by $d^1_{1,2}$, namely the uniform superposition $d^0_{1,1}|v_0\rangle$. Thus there is a unique state in $e^2_{2,3}$, which is the state in $e^2_{1,3}$ orthogonal to the image of $d^1_{1,2}$. But $(\operatorname{Im} d^1_{1,2})^\perp = \operatorname{Ker} \partial^2_{1,3}$, so we are looking for the unique state in the kernel of the gadget vertex boundary map $\partial^2_{1,3}$. We can now see that this state is the uniform superposition over the triangles, with matching orientations. Denote this state by $|\text{core}\rangle$.

Page 2

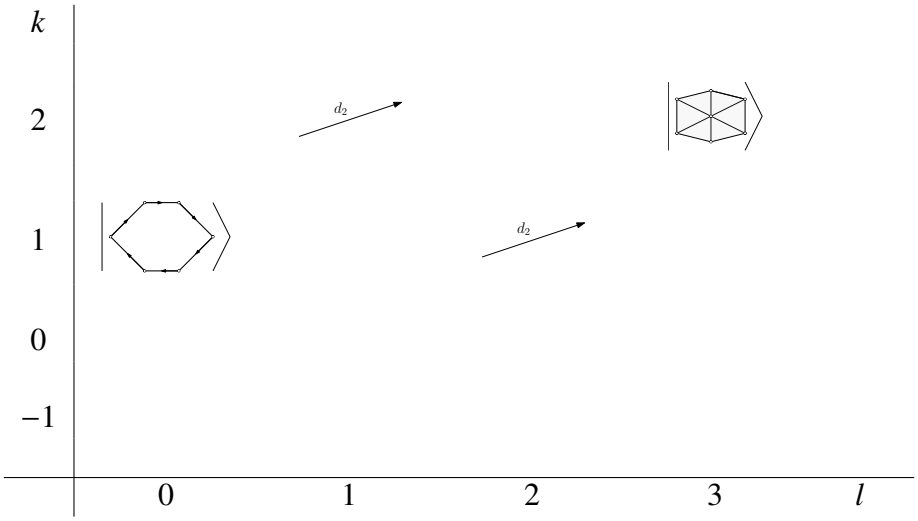

We have seen that page 2 contains only two states: $|\text{loop}\rangle \in e^1_{2,0}$ and $|\text{core}\rangle \in e^2_{2,3}$. It turns out that page 3 of the spectral sequence must be identical to page 2. To see why this is the case, consider the coboundary map of page 2 $d^k_{2,l}$. This maps via the 'knight move' one step up and two steps to the right; see Figure 6.7.

$$d^k_{2,l} : e^k_{2,l} \to e^{k+1}_{2,l+2} \tag{6.159}$$

Thus $|\text{loop}\rangle$ is mapped into $e^2_{2,2} = \{0\}$, and there is nothing in $e^1_{2,1} = \{0\}$ to map to $|\text{core}\rangle$. Since all relevant coboundary maps are zero, taking the homology leaves the page unchanged.

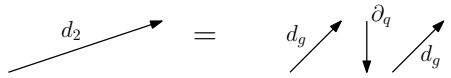

Figure 6.7: The coboundary map of page 2, $d_2 = d_g \cdot \partial_q \cdot d_g$. This acts as a 'knight move' as shown above.

Page 3

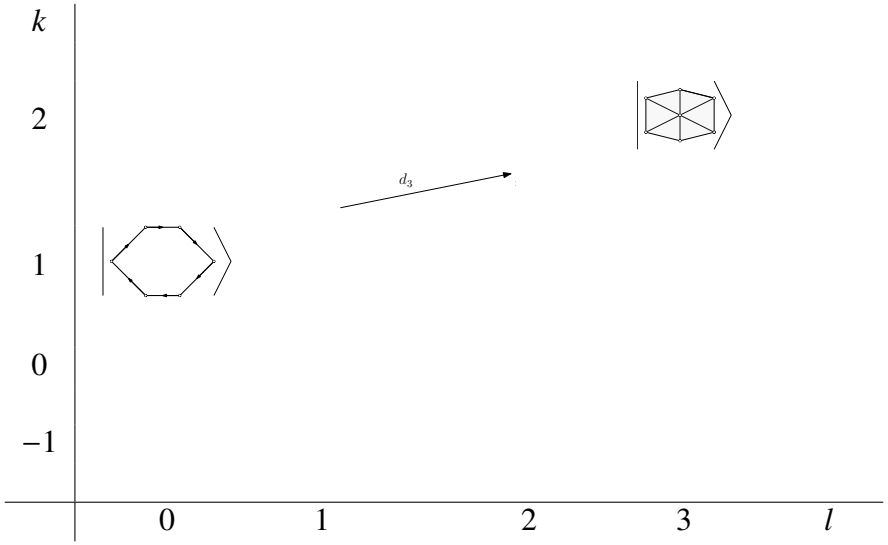

This argument no longer applies when we look at page 4. Now, the coboundary map of page 3 $d^k_{3,l}$ maps one step up and *three* steps to the right. In particular, $d^1_{3,0}$ maps

$$d^1_{3,0} : e^1_{3,0} \to e^2_{3,3} \tag{6.160}$$

We will argue that in fact (ignoring normalizations)

$$d^1_{3,0}\,|\text{loop}\rangle \;=\; |\text{core}\rangle \tag{6.161}$$

and these two states cancel each other out when we take the homology, leaving page 4 completely empty. To see this, we need to examine the coboundary map $d^1_{3,0}$ more closely. Returning to the zeroth page $e^k_{0,l}$, denote by $\partial^k_{\text{qubit},l} = \partial^k_{0,l}$ the 'qubit boundary map'

$$\partial^k_{\text{qubit},l} : e^k_{0,l} \to e^{k-1}_{0,l} \tag{6.162}$$

which acts by removing qubit vertices. Further, denote by $d^k_{\text{gadget},l}$ the 'gadget coboundary map'

$$d^k_{\text{gadget},l} : e^k_{0,l} \to e^{k+1}_{0,l+1} \tag{6.163}$$

which acts by adding gadget vertices. Now we can think of the map $d^1_{3,0}$ as acting by

$$d^1_{3,0} = d^1_{\text{gadget},2} \circ \partial^2_{\text{qubit},2} \circ d^1_{\text{gadget},1} \circ \partial^2_{\text{qubit},1} \circ d^1_{\text{gadget},0} \tag{6.164}$$

as shown in Figure 6.8 With this new understanding, let's examine $d^1_{3,0}|\text{loop}\rangle$. From Figure 6.9, we can see that indeed $d^1_{3,0}|\text{loop}\rangle = |\text{core}\rangle$.

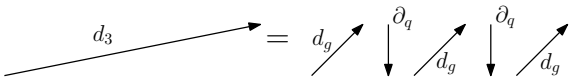

Figure 6.8: The coboundary map of page 2, $d_3 = d_g \cdot \partial_q \cdot d_g \cdot \partial_q \cdot d_g$. This acts by moving one step up and three to the right as shown.

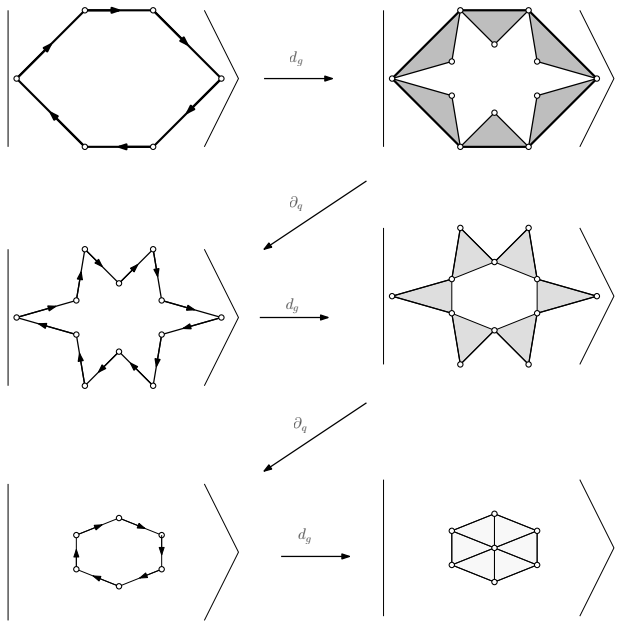

Figure 6.9: Applying the d_3 map to the $|\text{loop}\rangle$ gives the $|\text{core}\rangle$ state as claimed.

Page 4

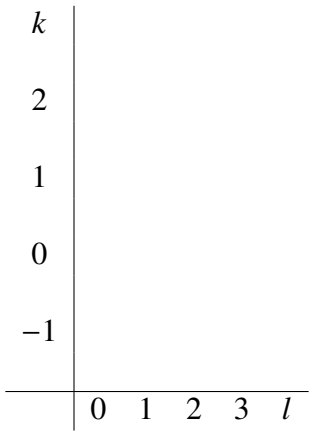

Proof of Lemma 6.9

We will apply the machinery of spectral sequences to the construction $\hat{\mathcal{G}}_m$ described in Section 6.2. Our real objects of interest for Lemma 6.9 are the spaces E_j^k. Our plan is to calculate the spaces $e_{j,l}^k$ and use the isomorphism in Proposition 6.11 to analyze E_j^k. The isomorphism in Proposition 6.11 allows us to learn about the analytical spaces E_j^k using easier algebraic techniques.

In Section 6.2, we were concerned with the $(2m-1)$-homology, and we were filling in a $(2m-1)$-cycle $|\phi\rangle \in H^{2m-1}(\mathcal{G}_m)$ by adding a gadget of $2m$-simplices. Recall $\mathcal{G}_m^k \subseteq \hat{\mathcal{G}}_m^k$ is the original qubit complex before adding the gadget. Equivalently, $\mathcal{G}_m^k$ is the simplicial subcomplex consisting of simplices of weight 1. For this section, we will drop the m subscripts $\mathcal{G}^k \subseteq \hat{\mathcal{G}}^k$. The original vertices $\mathcal{G}^0$ have weight 1, and will be referred to as *qubit vertices*. The added vertices $\hat{\mathcal{G}}^0 \setminus \mathcal{G}^0$ have weight λ, and will be referred to as *gadget vertices*. The weight of a simplex is defined to be the product of the weights of the vertices involved in the simplex. Recalling the construction of the gadgets from Section 6.2, introduce the following notations:

- Let [bulk] denote the set of simplices which involve the central vertex v_0, and let [non-bulk] be the complement of this set. It can be checked that [non-bulk] $\subseteq \hat{\mathcal{G}}$ is in fact a simplicial subcomplex, but [bulk] is *not*.
- Let $\Omega_j^k \subseteq \hat{\mathcal{G}}^k$ be the k-simplices of weights $\{1, \lambda, \ldots, \lambda^j\}$. It can be checked that $\Omega_j \subseteq \hat{\mathcal{G}}$ is a simplicial subcomplex for each j, and $C^k(\Omega_j) = (\mathcal{U}_{j+1}^k)^\perp$. Tautologically, $\Omega_0^k = \mathcal{G}^k$ is the qubit complex.

Spectral sequence of general gadget

We now calculate the spectral sequence of the general gadget construction from Section 6.2. The following lemma, stated in generality, will prove useful.

Lemma 6.12. *Let $\mathcal{P}$ be a simplicial complex and $\mathcal{Q} \subseteq \mathcal{P}$ a simplicial subcomplex. If $\mathcal{Q}$ has no $(k-1)$-cohomology, then*

$$C^k(\mathcal{Q})^\perp \cap d^{k-1}(C^{k-1}(\mathcal{Q})^\perp) = C^k(\mathcal{Q})^\perp \cap \operatorname{Im} d^{k-1}. \tag{6.165}$$

Proof. It is clear that

$$C^k(\mathcal{Q})^\perp \cap d^{k-1}(C^{k-1}(\mathcal{Q})^\perp) \subseteq C^k(\mathcal{Q})^\perp \cap \operatorname{Im} d^{k-1}. \tag{6.166}$$

It remains to show the opposite inclusion.

Let $|\alpha\rangle \in C^k(\mathcal{Q})^\perp \cap \operatorname{Im} d^{k-1}$. We must show that $|\alpha\rangle \in C^k(\mathcal{Q})^\perp \cap d^{k-1}(C^{k-1}(\mathcal{Q})^\perp)$. We know there is a $|\beta\rangle \in C^{k-1}(\mathcal{P})$ with $|\alpha\rangle = d^{k-1}|\beta\rangle$. Now

$$d^{k-1}(|\beta\rangle_{C^{k-1}(\mathcal{Q})})_{C^k(\mathcal{Q})} = 0, \tag{6.167}$$

where the subscript denotes *restriction*. Here comes the key step: $\mathcal{Q}$ has no $(k-1)$-cohomology, which implies that there exists a $|\gamma\rangle \in C^{k-2}(\mathcal{Q})$ with

$$|\beta\rangle_{C^{k-1}(\mathcal{Q})} = (d^{k-2}|\gamma\rangle)_{C^{k-1}(\mathcal{Q})}. \tag{6.168}$$

Now $(d^{k-1} \circ d^{k-2})|\gamma\rangle = 0$, so

$$d^{k-1}(d^{k-2}|\gamma\rangle)_{C^{k-1}(\mathcal{Q})} + d^{k-1}(d^{k-2}|\gamma\rangle)_{C^{k-1}(\mathcal{Q})^\perp} = 0. \tag{6.169}$$

Consider

$$|\beta\rangle_{C^{k-1}(\mathcal{Q})^\perp} - (d^{k-2}|\gamma\rangle)_{C^{k-1}(\mathcal{Q})^\perp} \in C^{k-1}(\mathcal{Q})^\perp. \tag{6.170}$$

This has

$$d^{k-1}\Big(|\beta\rangle_{C^{k-1}(\mathcal{Q})^\perp} - (d^{k-2}|\gamma\rangle)_{C^{k-1}(\mathcal{Q})^\perp}\Big) = |\alpha\rangle. \tag{6.171}$$

□

Lemma 6.12 is telling us that, if $\mathcal{Q} \subseteq \mathcal{P}$ is a simplicial subcomplex with no $(k-1)$-homology, then all k-coboundaries in $C^k(\mathcal{Q})^\perp$ are in fact the coboundaries of chains in $C^{k-1}(\mathcal{Q})^\perp$.

The zeroth page of the spectral sequence is simply $e^k_{0,l} = \mathcal{U}^k_l / \mathcal{U}^k_{l+1}$. By choosing the representative orthogonal to $\mathcal{U}^k_{l+1}$ in each equivalence class, we can think of $e^k_{0,l}$

as $\mathcal{U}_l^k \cap (\mathcal{U}_{l+1}^k)^\perp$, which are the simplices of weight λ^l. This will be a general trick in this section, to replace a quotient by an intersection with the orthogonal space. There are no simplices of dimension higher than $2m$, so there is nothing above the row $k = 2m$; that is, $e_{0,l}^k = \{0\}\ \forall\ k > 2m$. There are no k-simplices with more than $k+1$ gadget vertices, so there is nothing below the diagonal $l = k+1$; that is, $e_{0,l}^k = \{0\}\ \forall\ l > k+1$. There are no $2m$-simplices consisting only of qubit vertices, so $e_{0,0}^{2m} = \{0\}$. The first column $l = 0$ is simply the chainspaces of the qubit complex $C^k(\mathcal{G})$. The diagonal $l = k+1$ consists simply of $e_{0,k+1}^k = \mathcal{U}_{k+1}^k$, which are the simplices involving no qubit vertices. Note these are *not* the same as the bulk chainspaces $C^k([\text{bulk}])$. The spaces $e_{0,l}^k$ for $0 < l < k+1$ consist of simplices which are a mixture of qubit and gadget vertices, which make up the *thickening* of the gadget construction.

Page 0

k							
$2m$	$\{0\}$	$\dots$	$\dots$	$e_{0,2m-1}^{2m}$	$e_{0,2m}^{2m}$	$\mathcal{U}_{2m+1}^{2m}$	
$2m-1$	$C^{2m-1}(\mathcal{G})$	$\dots$	$\dots$	$e_{0,2m-1}^{2m-1}$	$\mathcal{U}_{2m}^{2m-1}$		
$2m-2$	$C^{2m-2}(\mathcal{G})$	$\dots$	$\dots$	$\mathcal{U}_{2m-1}^{2m-2}$			
$\vdots$	$\vdots$		$\iddots$				
$\vdots$	$\vdots$	$\iddots$					
	0	$\dots$	$\dots$	$2m-1$	$2m$	$2m+1$	l

The coboundary map $d_{0,l}^k$ of the zeroth page maps upwards one step from $e_{0,l}^k$ to $e_{0,l}^{k+1}$. When acting on the representative in $\mathcal{U}_l^k \cap (\mathcal{U}_{l+1}^k)^\perp$, it can be thought of as a qubit vertex coboundary map, which adds a qubit vertex to the simplex.

We are now in a position to compute the first page $e_{1,l}^k$ of the spectral sequence, which is defined as $e_{1,l}^k = \text{Ker}\, d_{0,l}^k / \text{Im}\, d_{0,l}^{k-1}$. For the first column $l = 0$, this simply gives us the cohomology of the qubit complex $e_{1,0}^k = H^k(\mathcal{G})$. These cohomology groups are zero except for $H^{2m-1}(\mathcal{G})$. Next let's look at the diagonal $l = k+1$. The claim is that these spaces are isomorphic to the bulk chainspaces $e_{1,k+1}^k = C^k([\text{bulk}])$.

The spaces $e^{k-1}_{1,k+1}$ are zero, so $\operatorname{Im} d^{k-1}_{0,k+1} = \{0\}$, and $e^k_{1,k+1}$ are simply the cocycles $e^k_{1,k+1} = \operatorname{Ker} d^k_{0,k+1}$. The simplices in $\mathcal{U}^k_{k+1}$ which are *not* in $C^k([\text{bulk}])$ do not vanish under $d^k_{0,k+1}$, and thus are not cocycles. It remains to look at $e^k_{1,l}$ for $1 < l < k+1$. It turns out that these spaces are all zero.

Claim 6.6. *$e^k_{1,l} = \{0\}$ for $1 \le l \le k$, $k = 1, \ldots, 2m$.*

Proof. We will split this up into two cases: $k \neq 2m-1$ and $k = 2m-1$. The first case $k \neq 2m-1$ will be easier, and the case $k = 2m-1$ will require an extra idea. Note the first case $k \neq 2m-1$ includes $k < 2m-1$ and $k = 2m$. The strategy in this proof is to apply Lemma 6.12 to show that all states in the kernel of the relevant outgoing coboundary map are also in the image of the relevant incoming coboundary map.

Let's begin with the rows $k \neq 2m-1$. Recalling the interpretation $e^k_{0,l} = \mathcal{U}^k_l \cap (\mathcal{U}^k_{l+1})^\perp$, let $|\alpha\rangle \in \mathcal{U}^k_l \cap (\mathcal{U}^k_{l+1})^\perp$ with $|\alpha\rangle \in \operatorname{Ker} d^k_{0,l}$. This tells us that $d^k|\alpha\rangle \in \mathcal{U}^{k+1}_{l+1}$. Restricted to the subcomplex Ω_l, this says that $(d^k|\alpha\rangle)_{C^{k+1}(\Omega_l)} = 0$. That is, $|\alpha\rangle$ is a *cocycle* in Ω_l. (Here and throughout, the subscript refers to restriction or orthogonal projection onto this subspace.) By considering the gadget construction from Section 6.2, Ω_l has no k-homology for $k \neq 2m-1$. (Ω_l *does* in fact have $(2m-1)$-homology, so for this case we will need an extra trick. But for now, $k \neq 2m-1$ and Ω_l has no k-homology.) This means that $|\alpha\rangle$ is not only a *cocycle* in Ω_l but necessarily also a *coboundary*. Now apply Lemma 6.12 with $\mathcal{P} = \Omega_l$ and $\mathcal{Q} = \Omega_{l-1}$. Note $|\alpha\rangle \in \mathcal{U}^k_l$ so indeed $|\alpha\rangle \perp C^k(\Omega_{l-1})$. We get that $|\alpha\rangle = (d^{k-1}|\beta\rangle)_{C^k(\Omega_l)}$ for some $|\beta\rangle \in C^{k-1}(\Omega_l) \cap C^{k-1}(\Omega_{l-1})^\perp$. But $C^{k-1}(\Omega_l) \cap C^{k-1}(\Omega_{l-1})^\perp = \mathcal{U}^{k-1}_l \cap (\mathcal{U}^{k-1}_{l-1})^\perp$, so $|\beta\rangle \in e^{k-1}_{0,l}$. Noting that $(d^{k-1}|\beta\rangle)_{C^k(\Omega_l)}$ is precisely $(d^{k-1}|\beta\rangle)_{(\mathcal{U}^k_{l+1})^\perp} = d^{k-1}_{0,l}|\beta\rangle$, we get that $|\alpha\rangle \in \operatorname{Im} d^{k-1}_{0,l}$. The conclusion is that $\operatorname{Ker} d^k_{0,l} = \operatorname{Im} d^k_{0,l}$ and $e^k_{0,l}$ has no homology, so $e^k_{1,l} = \{0\}$.

Now we move onto the case $k = 2m-1$. As before, suppose $|\alpha\rangle \in \mathcal{U}^{2m-1}_l \cap (\mathcal{U}^{2m-1}_{l+1})^\perp$ with $|\alpha\rangle \in \operatorname{Ker} d^{2m-1}_{0,l}$. Here, it is *not* true that Ω_l has no $(2m-1)$-homology. However, since $|\alpha\rangle$ has no support on the qubit complex $\mathcal{G}$, the idea is that *it cannot tell* that the complex has homology. We are able to imagine closing all the $(2m-1)$-cycles with fictitious auxiliary vertices, and Lemma 6.12 will tell us that this move is not detected by $|\alpha\rangle$.

Let's formalize this. Recalling Section 6.2, $\mathcal{G} \subseteq \hat{\mathcal{G}}$ consists of 2^m copies of S^{2m-1} corresponding to the m-bit strings. For each copy of S^{2m-1} introduce an extra 'auxiliary' vertex (of weight 1) and connecting it with all vertices in this S^{2m-1}, thus

'filling in the hole'. Denote the new objects after this operation with a star such as $\hat{\mathcal{G}}^*$. The effect of the auxiliary vertices is that $\hat{\mathcal{G}}^*$ has no $(2m-1)$-cohomology. This includes Ω_l^* for each l: these subcomplexes no longer have any $(2m-1)$-cohomology.

Continuing the argument, $|\alpha\rangle \in \mathcal{U}_l^{2m-1} \cap (\mathcal{U}_{l+1}^{2m-1})^\perp$ with $|\alpha\rangle \in \operatorname{Ker} d_{0,l}^{2m-1}$, which tells us that $d^{2m-1}|\alpha\rangle \in \mathcal{U}_{l+1}^{2m}$ and, restricting to the subcomplex Ω_l, $(d^{2m-1}|\alpha\rangle)_{C^{2m}(\Omega_l)} = 0$. That is, $|\alpha\rangle$ is a *cocycle* in Ω_l. Now $|\alpha\rangle \in \mathcal{U}_l^{2m-1}$ so $|\alpha\rangle$ has no support on the qubit complex $\mathcal{G}$, $|\alpha\rangle_{\mathcal{G}} = 0$. Thus $|\alpha\rangle$ is likewise a *cocycle* in Ω_l^*. But Ω_l^* has no $(2m-1)$-homology, so $|\alpha\rangle$ is in fact a *coboundary* in Ω_l^*. Now apply a Lemma 6.12 with $\mathcal{P} = \Omega_l^*$ and $\mathcal{Q} = \Omega_{l-1}^*$. We get that $|\alpha\rangle = (d^{2m-2}|\beta\rangle)_{C^{2m-1}(\Omega_l^*)}$ for some $|\beta\rangle \in C^{2m-2}(\Omega_l^*) \cap C^{2m-2}(\Omega_{l-1}^*)^\perp$. But $C^{2m-2}(\Omega_l^*) \cap C^{2m-2}(\Omega_{l-1}^*)^\perp = \mathcal{U}_l^{2m-2} \cap (\mathcal{U}_{l-1}^{2m-2})^\perp$, so $|\beta\rangle \in e_{0,l}^{2m-2}$. Noting that $(d^{2m-2}|\beta\rangle)_{C^k(\Omega_l^*)} = (d^{2m-2}|\beta\rangle)_{C^k(\Omega_l)}$ is precisely $d_{0,l}^{2m-2}|\beta\rangle$, we get that $|\alpha\rangle \in \operatorname{Im} d_{0,l}^{2m-2}$. The conclusion is that $\operatorname{Ker} d_{0,l}^{2m-1} = \operatorname{Im} d_{0,l}^{2m-1}$ and $e_{0,l}^{2m-1}$ has no homology, so $e_{1,l}^{2m-1} = \{0\}$. □

Page 1

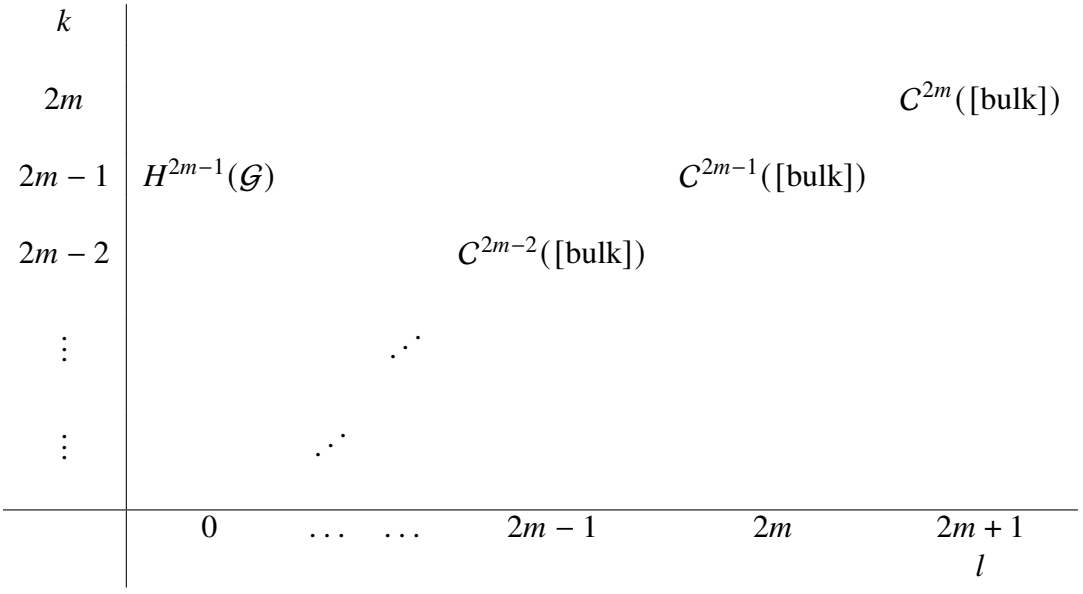

The page 1 coboundary map $d_{1,l}^k$ maps diagonally upwards and to the right one step from $e_{0,l}^k$ to $e_{0,l}^{k+1}$. It can be thought of as a gadget vertex coboundary map, which adds a gadget vertex to the simplex. Acting on $C^k([\text{bulk}])$, this has the same action as the regular coboundary map $\hat{d}^k$ of $\hat{\mathcal{G}}^k$.

Let's now calculate the second page of the spectral sequence, defined as $e_{2,l}^k = \operatorname{Ker} d_{1,l}^k / \operatorname{Im} d_{1,l-1}^{k-1}$. $H^{2m-1}(\mathcal{G})$ remains unchanged, since both the relevant cobound-

ary maps are zero. Next we show that all terms on the diagonal $l = k + 1$ vanish, except for $e^{2m}_{2,2m+1}$.

Claim 6.7. $e^k_{2,k+1} = \{0\}\ \forall\ k = 0, \ldots, 2m-1$

Proof. This proof will be very similar to that of Claim 6.6, except that we will apply Lemma 6.12 with $\mathcal{P}$ being the entire complex and $\mathcal{Q}$ being [non-bulk]. We will again split up into the two cases $k < 2m-1$ and $k = 2m-1$, with the case $k = 2m-1$ requiring the same extra idea. The relevant coboundary maps are now those of the first page $d^k_{1,k+1}$, and recall $e^k_{2,k+1}$ is defined to be the homology of the chain:

$$C^{k-1}([\text{bulk}]) \xrightarrow{d^{k-1}_{1,k}} C^k([\text{bulk}]) \xrightarrow{d^k_{1,k+1}} C^{k+1}([\text{bulk}]) \tag{6.172}$$

Let's again begin with $k < 2m-1$. Let $|\alpha\rangle \in C^k([\text{bulk}])$ with $|\alpha\rangle \in \operatorname{Ker} d^k_{1,k+1}$. Acting on $C^k([\text{bulk}])$, $d^k_{1,k+1}$ simply acts as the original coboundary map d^k, so in fact $|\alpha\rangle \in \operatorname{Ker} d^k$ is a *cocycle*. But the complex has no k-homology for $k \neq 2m-1$, so $|\alpha\rangle$ must also be a *coboundary*. Now apply Lemma 6.12 with $\mathcal{P} = \hat{\mathcal{G}}$ and $\mathcal{Q} =$ [non-bulk]. We get that $|\alpha\rangle \in d^{k-1}(C^{k-1}([\text{non-bulk}])^\perp) = d^{k-1}_{1,k}(C^{k-1}([\text{bulk}]))$. We conclude that $\operatorname{Ker} d^k_{1,k+1} = \operatorname{Im} d^{k-1}_{1,k}$ and $e^k_{2,k+1} = \{0\}$.

Now we move onto the case $k = 2m-1$. Recalling Section 6.2, $\mathcal{G} \subseteq \hat{\mathcal{G}}$ consists of 2^m copies of S^{2m-1} corresponding to the m-bit strings. For each copy of S^{2m-1} introduce an extra auxiliary vertex (of weight 1) and connecting it with all vertices in this S^{2m-1}, thus filling in the hole. Denote the new objects after this operation with a star such as $\hat{\mathcal{G}}^*$. The effect of the auxiliary vertices is that $\hat{\mathcal{G}}^*$ has no $(2m-1)$-homology.

As before, suppose $|\alpha\rangle \in C^{2m-1}([\text{bulk}])$ with $|\alpha\rangle \in \operatorname{Ker} d^{2m-1}_{1,2m} = \operatorname{Ker} d^{2m-1}$. Now $|\alpha\rangle$ has no support on the qubit complex $\mathcal{G}$, thus $|\alpha\rangle$ is likewise a cocycle in $\hat{\mathcal{G}}^*$. But $\hat{\mathcal{G}}^*$ has no $(2m-1)$-homology, so $|\alpha\rangle$ is in fact a *coboundary* in $\hat{\mathcal{G}}^*$. Now apply Lemma 6.12 with $\mathcal{P} = \hat{\mathcal{G}}^*$ and $\mathcal{Q} = [\text{non-bulk}]^*$. We get that $|\alpha\rangle \in d^{2m-2}(C^{2m-2}([\text{non-bulk}]^*)^\perp)$. But $C^{2m-2}([\text{non-bulk}]^*)^\perp$ is simply $C^{2m-2}([\text{bulk}])$, so $|\alpha\rangle \in d^{2m-2}(C^{2m-2}([\text{bulk}])) = d^{2m-2}_{1,2m-1}(C^{2m-2}([\text{bulk}]))$. We conclude that $\operatorname{Ker} d^{2m-1}_{1,2m} = \operatorname{Im} d^{2m-2}_{1,2m-1}$ and $e^{2m-1}_{2,2m} = \{0\}$.

The interpretation of this move is the same as in the proof of Claim 6.6. Since $|\alpha\rangle$ has no support on the qubit complex $\mathcal{G}$, *it cannot tell* that the complex has some $(2m-1)$-homology. We are able to imagine closing all the $(2m-1)$-cycles with

fictitious auxiliary vertices, and Lemma 6.12 will tell us that this is not detected by $|\alpha\rangle$. □

Finally let's investigate $e_{2,2m+1}^{2m}$. The entire space $C^{2m}([\text{bulk}])$ is in $\operatorname{Ker} d_{1,2m+1}^{2m}$ since $e_{1,2m+2}^{2m+1} = \{0\}$, so we have $e_{2,2m+1}^{2m} = C^{2m}([\text{bulk}])/\operatorname{Im} d_{1,2m}^{2m-1}$. As usual, we can pick the representative orthogonal to the space we are quotienting and write $e_{2,2m+1}^{2m} = C^{2m}([\text{bulk}]) \cap (\operatorname{Im} d_{1,2m}^{2m-1})^{\perp}$. But $(\operatorname{Im} d_{1,2m}^{2m-1})^{\perp} = \operatorname{Ker} \partial_{1,2m+1}^{2m}$ where $\partial_{1,l}^{k} = (d_{1,l-1}^{k-1})^{\dagger}$ can be interpreted as a gadget vertex boundary map which removes a gadget vertex from the simplex. So $e_{2,2m+1}^{2m} = C^{2m}([\text{bulk}]) \cap \operatorname{Ker} \partial_{1,2m+1}^{2m}$. Since [bulk] is isomorphic to the interior of a $2m$-ball, the only state in $C^{2m}([\text{bulk}])$ which has no $\partial_{1,2m+1}^{2m}$-boundary is the uniform superposition over all the $2m$ simplices in $[\text{bulk}]^{2m}$, with appropriate orientations. Let's denote this state by $|\text{core}\rangle \in C^{2m}([\text{bulk}])$.

Page 2

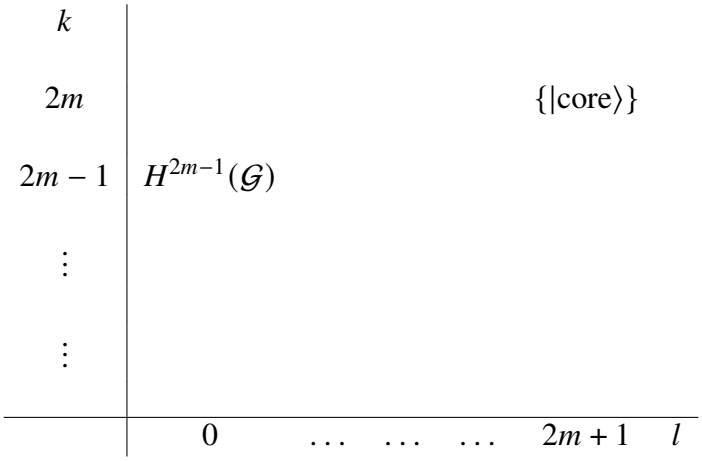

Pages 3 to $2m+1$ will remain unchanged. This can be seen purely from the direction in which the coboundaries $d_{j,l}^{k}$ map. $d_{j,l}^{k}$ moves one step up and j steps to the right, so these must necessarily be zero maps.

$$d_{j,l}^{k} : e_{j,l}^{k} \to e_{j,l+j}^{k+1} \tag{6.173}$$

Since all relevant coboundary maps are zero, taking the homology leaves the page unchanged.

This argument no longer applies when we look at page $2m+2$. Now, the coboundary map of page $2m+1$ $d_{2m+1,l}^{k}$ maps one step up and $2m+1$ steps to the right. In particular, $d_{2m+1,0}^{1}$ maps

$$d_{2m+1,0}^{1} : H^{2m-1}(\mathcal{G}) \to \{|\text{core}\rangle\}. \tag{6.174}$$

We will argue that in fact (ignoring normalizations)

$$d^1_{2m+1,0}\,|\phi\rangle \;=\; |\text{core}\rangle, \tag{6.175}$$

where $|\phi\rangle \in H^{2m-1}(\mathcal{G})$ is the cycle being filled by the gadget $\hat{\mathcal{G}}$. Thus $|\phi\rangle$ is lifted out of the homology at this page. To see this, we need to examine the coboundary map $d^1_{2m+1,0}$ more closely. Returning to the zeroth page $e^k_{0,l}$, denote by $\partial^k_{\text{qubit},l} = \partial^k_{0,l}$ the 'qubit boundary map'

$$\partial^k_{\text{qubit},l} : e^k_{0,l} \to e^{k-1}_{0,l} \tag{6.176}$$

which acts by removing qubit vertices. Further, denote by $d^k_{\text{gadget},l}$ the 'gadget coboundary map'

$$d^k_{\text{gadget},l} : e^k_{0,l} \to e^{k+1}_{0,l+1} \tag{6.177}$$

which acts by adding gadget vertices. Now we can think of the map $d^1_{2m+1,0}$ as acting by

$$d^1_{2m+1,0} = d^{2m-1}_{\text{gadget},2m} \circ \partial^{2m}_{\text{qubit},2m} \circ \cdots \circ d^1_{\text{gadget},1} \circ \partial^2_{\text{qubit},1} \circ d^1_{\text{gadget},0}. \tag{6.178}$$

With this new understanding, we can see that $d^1_{2m+1,0}|\phi\rangle = |\text{core}\rangle$. Recalling Section 6.2, the cycle $|\phi\rangle$ on the outside layer gets transported through the thickening of the gadget to the inside layer, at which point it gets sent to $|\text{core}\rangle$ by the final gadget vertex coboundary map $d^{2m-1}_{\text{gadget},2m}$.

Page 2

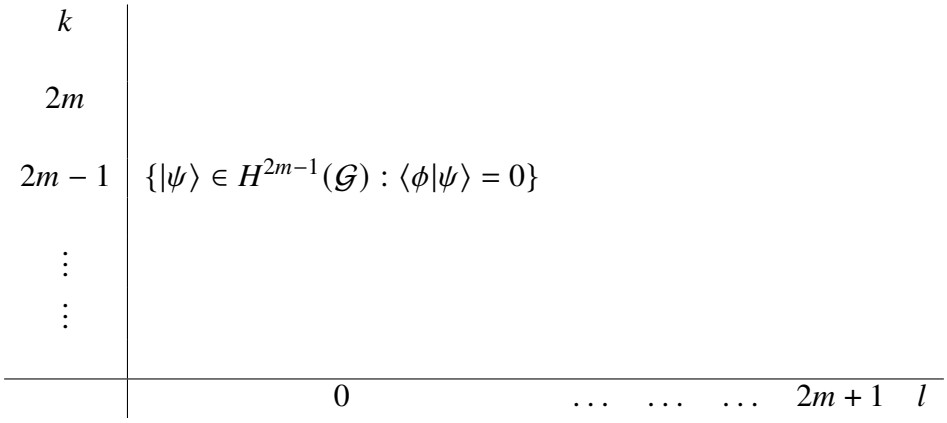

Finishing the proof

By Theorem 4 of [For94], the eigenvalues of the Laplacian decay like $\Theta(\lambda^{2j})$ for some j. Thus if we take the eigenspace of eigenvalues $O(\lambda^{2j})$ and intersect it with

the orthogonal complement of the eigenspace of eigenvalues $O(\lambda^{2j+2})$, we get the eigenspace of eigenvalues that are exactly $\Theta(\lambda^{2j})$.

Proposition 6.10 tells us that the Laplacian eigenspace of eigenvalues $O(\lambda^{2j})$ is a $O(\lambda)$-perturbation of E_j^{2m-1}, in the sense of Section 6.1. Thus the Laplacian eigenspace of eigenvalues $\Theta(\lambda^{2j})$ is a $O(\lambda)$-perturbation of $E_j^{2m-1} \cap (E_{j+1}^{2m-1})^\perp$.

Thus to complete the proof of Lemma 6.9, we use the isomorphism in Proposition 6.11 to derive the spaces E_j^k. Recall

$$E_j^k = \bigoplus_l E_{j,l}^k \tag{6.179}$$

and by Proposition 6.11 we can get $E_{j,l}^k$ by taking $e_{j,l}^k$ and projecting a representative from each equivalence class onto $\mathcal{U}_l^k$. This gives:

- $E_0^{2m-1} = C^{2m-1}(\hat{\mathcal{G}})$, with $E_{0,l}^{2m-1} = \mathcal{U}_l^{2m-1} \cap (\mathcal{U}_l^{2m-1})^\perp$ for each l. (These are tautological.)
- $E_1^{2m-1} = H^{2m-1}(\mathcal{G}) \oplus C^{2m-1}([\text{bulk}])$, with $E_{1,0}^{2m-1} = H^{2m-1}(\mathcal{G})$ and $E_{1,2m}^{2m-1} = C^{2m-1}([\text{bulk}])$.
- $E_j^{2m-1} = E_{j,0}^{2m-1} = H^{2m-1}(\mathcal{G})$ for all $j = 2, \ldots, 2m+1$.
- $E_j^{2m-1} = E_{j,0}^{2m-1} = \{|\psi\rangle \in H^{2m-1}(\mathcal{G}) : \langle\phi|\psi\rangle = 0\}$ for all $j \geq 2m+2$.

(Recall that $H^{2m-1}(\mathcal{G}) \cong \operatorname{Ker} \Delta^{2m-1}$ where Δ^k is the Laplacian of the qubit complex $\mathcal{G}^k$. When the space $H^{2m-1}(\mathcal{G})$ appears above, it is the harmonic representative from $\operatorname{Ker} \Delta^k \subseteq C^k(\mathcal{G}) \subseteq C^k(\hat{\mathcal{G}})$ which is present.)

6.4 Combining gadgets

In this section, we describe how to combine many gadgets together to simulate a local Hamiltonian. We prove our main theorem, Theorem 6.14.

Recall from Section 6.2 $\hat{\mathcal{G}}_m$ denotes the complex of a single gadget. We must be careful to keep track of what we think of as the original qubit complex and the additional gadget complex. The *vertices* $\hat{\mathcal{G}}_m^0$ can be partitioned into the vertices of the original qubit graph $\mathcal{G}_m^0$, which have weight 1, and the added gadget vertices, which have weight λ. Let these added gadget vertices be denoted $\mathcal{T}^0$. The *k-simplices* $\hat{\mathcal{G}}_m^k$ can be partitioned into the k-simplices of the original qubit complex $\mathcal{G}_m^k$, and the k-simplices which contain at least one vertex from $\mathcal{T}^0$. Let these extra

k-simplices be denoted $\mathcal{T}^k$. Note, however, that $\mathcal{T}$ is *not* a simplicial complex in its own right, since $\mathcal{T}^k$ involve vertices outside of $\mathcal{T}^0$. We can decompose the chain space of $\hat{\mathcal{G}}_m$ as a direct sum

$$C^k(\hat{\mathcal{G}}_m) = C^k(\mathcal{G}_m) \oplus C^k(\mathcal{T}). \tag{6.180}$$

In Section 6.3 we used the powerful tool of spectral sequences to understand the spectrum of the Laplacian of a single gadget $\hat{\mathcal{G}}_m$, captured by Lemma 6.9. This lemma will be an essential ingredient later when we come to analyze the spectrum of many gadgets combined.

Padding with identity

In order to add the remaining $n - m$ qubits, we join the graph $\mathcal{G}_{n-m}$ to get $\hat{\mathcal{G}}_n = \hat{\mathcal{G}}_m * \mathcal{G}_{n-m}$. This is analogous to tensoring a Hamiltonian term with identity on all the qubits outside of its support.

There is another way to look at the final complex $\hat{\mathcal{G}}_n$. To get $\hat{\mathcal{G}}_n$, we implement the gadget $\mathcal{T}$ described in Section 6.2 on the copies of $\mathcal{G}_1$ corresponding to the qubits on which ϕ is supported, and then connecting the vertices of the gadget $\mathcal{T}^0$ *all to all* with the qubit vertices in the copies of $\mathcal{G}_1$ corresponding to qubits outside the support of ϕ.

Lemma 6.13. *Let $\hat{\Delta}'^k$ be the Laplacian of $\hat{G}_n$.*

- *$\hat{\Delta}'^{2n-1}$ has a $(2^m - 1) \cdot 2^{n-m}$-dimensional kernel, which is a $O(\lambda)$-perturbation of the subspace $\{|\psi\rangle \in \mathcal{H}_m : \langle\phi|\psi\rangle = 0\}$ tensored with $\mathcal{H}_{n-m}$.*
- *The first excited eigenspace of $\hat{\Delta}'^{2n-1}$ above the kernel is the 2^{n-m}-dimensional space $|\hat{\phi}\rangle \otimes \mathcal{H}_{n-m}$, where $|\hat{\phi}\rangle$ is a $O(\lambda)$-perturbation of $|\phi\rangle \in \mathcal{H}_m$, and it has energy $\Theta(\lambda^{4m+2})$.*
- *The next lowest eigenvectors have eigenvalues $\Theta(\lambda^2)$, and they are $O(\lambda)$-perturbations of sums of $(2m-1)$-simplices touching the central vertex v_0, tensored with $\mathcal{H}_{n-m}$.*
- *The rest of the eigenvalues are $\Theta(1)$.*

Proof. By Fact 6.3, the new chainspace is

$$C^{2n-1}(\hat{\mathcal{G}}_n) = C^{2m-1}(\hat{\mathcal{G}}_m) \otimes C^{2(n-m)+1}(\mathcal{G}_{n-m}). \tag{6.181}$$

By Lemma 6.5, the new Laplacian on $\hat{\mathcal{G}}_n$ is

$$\hat{\Delta}'^{2n-1} = \hat{\Delta}^{2m-1} \otimes \mathbb{1} + \mathbb{1} \otimes \Delta^{2(n-m)+1}, \tag{6.182}$$

where $\Delta^{2(n-m)+1}$ here is the Laplacian on $C^{2(n-m)+1}(\mathcal{G}_{n-m})$. Notice that the two terms on the right-hand-side commute. $\Delta^{2(n-m)+1}$ has kernel $\mathcal{H}_{n-m}$ and its first excited eigenvalue is $\Theta(1)$. The conclusions follow from Lemma 6.9. □

Define the new gadget simplices of $\hat{\mathcal{G}}_n$ to be $\mathcal{T}' = \mathcal{T} * \mathcal{G}_{n-m}$, and we have

$$C^{2n-1}(\mathcal{T}') = C^{2m-1}(\mathcal{T}) \otimes C^{2(n-m)-1}(\mathcal{G}_{n-m}). \tag{6.183}$$

Some facts about a single gadget

We now take the opportunity to prove some useful facts about the complex with a single gadget added. Suppose we implement the gadget corresponding to integer state $|\phi\rangle$, and let $|\hat{\phi}\rangle$ be as in Lemma 6.13. First we show that states in the subspace $|\hat{\phi}\rangle \otimes \mathcal{H}_{n-m}$ are *cycles*. In other words, these states are *paired up* in the language of Section 6.1.

Claim 6.8. $\hat{\partial}^{2n-1}(|\hat{\phi}\rangle \otimes \mathcal{H}_{n-m}) = 0$

Proof. $\mathcal{H}_{n-m}$ consists of cycles, so it is sufficient to show $\hat{\partial}^{2m-1}|\hat{\phi}\rangle = 0$. $|\hat{\phi}\rangle$ is an eigenstate of $\hat{\Delta}^{2m-1}$, thus by Proposition 6.4 it must be paired up $|\hat{\phi}\rangle \in \operatorname{Im} \hat{\partial}^{2m}$ or paired down $|\hat{\phi}\rangle \in \operatorname{Im} \hat{d}^{2m-2}$. By Lemma 6.9, $|\hat{\phi}\rangle$ is a $O(\lambda)$ perturbation of $|\phi\rangle \in \mathcal{H}_m$. $|\phi\rangle$ is a cycle $\hat{\partial}^{2m-1}|\phi\rangle = 0$, so $|\phi\rangle$ is orthogonal to $\operatorname{Im} \hat{d}^{2m-2}$. This guarantees that the $O(\lambda)$-perturbation $|\hat{\phi}\rangle = |\phi\rangle + O(\lambda)$ is *not* contained in $\operatorname{Im} \hat{d}^{2m-2}$. Thus $|\hat{\phi}\rangle$ is paired up and $\hat{\partial}^{2m-1}|\hat{\phi}\rangle = 0$. □

Next we will state some facts about the *bulk* of the gadget. Let [bulk] denote the simplices touching the central vertex v_0 of gadget $\mathcal{T}$ (see Section 6.2), with chainspaces $C^k([\text{bulk}])$, and let $\Pi^{[k]}$ be the projection onto $C^k([\text{bulk}])$.

Claim 6.9. *All states* $|\psi\rangle$ *have* $||\Pi^{[2n-2]}\hat{\partial}^{2n-1}|\psi\rangle|| = O(\lambda)|||\psi\rangle||$ *and* $||\Pi^{[2n]}\hat{d}^{2n-1}|\psi\rangle|| = O(\lambda)|||\psi\rangle||$.

Proof. All vertices touching [bulk] have weight λ. Equations 6.75, 6.77 then give the conclusion. □

The proof of the following claim relies on a fact from Section 6.3.

Claim 6.10. *A normalized state* $|\psi\rangle \in C^{2n-1}([\text{bulk}])$ *has* $||\Pi^{[2n-2]}\hat{\partial}^{2n-1}|\psi\rangle|| = \Omega(\lambda)$ *or* $||\Pi^{[2n]}\hat{d}^{2n-1}|\psi\rangle|| = \Omega(\lambda)$.

Proof. Since all vertices relevant to these claims are weighted by λ, we can consider the *unweighted* complex. Then it is sufficient to show that either $||\Pi^{[2n-2]}\hat{\partial}^{2n-1}|\psi\rangle|| \neq 0$ or $||\Pi^{[2n]}\hat{d}^{2n-1}|\psi\rangle|| \neq 0$. Returning to the weighted complex simply introduces a factor of λ.

Suppose $||\Pi^{[2n]}\hat{d}^{2n-1}|\psi\rangle|| = 0$. That is, $\Pi^{[2n]}\hat{d}^{2n-1}|\psi\rangle = 0$. This in fact tells us that $\hat{d}^{2n-1}|\psi\rangle = 0$, since $\hat{d}^{2n-1}$ maps $C^{2n-1}([\text{bulk}])$ into $C^{2n}([\text{bulk}])$.

Claim 6.7 from Section 6.3 is telling us that $C^{2n-1}([\text{bulk}]) \cap \text{Ker}\,\hat{d}^{2n-1} = \hat{d}^{2n-2}(C^{2n-2}([\text{bulk}]))$. Note the proof of Claim 6.7 does *not* rely on any spectral sequence machinery, but only Lemma 6.12 and the argument in the proof of Claim 6.7 where we close the homology with auxiliary vertices.

The result is that, from $||\Pi^{[2n]}\hat{d}^{2n-1}|\psi\rangle|| = 0$, we can deduce that $|\psi\rangle \in \hat{d}^{2n-2}(C^{2n-2}([\text{bulk}]))$. Remembering $\hat{\partial}^{2n-1} = (\hat{d}^{2n-2})^\dagger$, this guarantees that $\hat{\partial}^{2n-1}|\psi\rangle$ has some component in $C^{2n-2}([\text{bulk}])$ and $||\Pi^{[2n-2]}\hat{\partial}^{2n-1}|\psi\rangle|| \neq 0$. □

Combining gadgets

We have seen how to construct a gadget $\mathcal{T}$ to implement a single local rank-1 projector $\phi = |\phi\rangle\langle\phi|$ where $|\phi\rangle$ is an integer state. We would now like to implement a Hamiltonian

$$H = \sum_{i=1}^{t} \phi_i \tag{6.184}$$

which is a sum of such terms. This will involve adding a gadget $\mathcal{T}_i$ for each term ϕ_i. We would like to add these gadgets in an independent way. From Section 6.4, we have a way of gluing in a gadget $\mathcal{T}_i'$ onto $\mathcal{G}_n$ which corresponds to ϕ_i. To implement multiple terms ϕ_i, we simply glue in the gadgets separately. One may wonder if we should include any edges between the gadget vertices if different gadgets; we do *not* include any of these edges.

The full procedure is described as follows, which includes the padding-with-identity step from Section 6.4:

1. For each i, add gadget $\mathcal{T}_i$ to the copies of $\mathcal{G}_1$ corresponding to qubits in the support of $|\phi_i\rangle$.

2. For each i, connect gadget vertices $\mathcal{T}_i^0$ *all to all* with the qubit vertices in the copies of $\mathcal{G}_1$ corresponding to qubits outside the support of $|\phi_i\rangle$ to get $\mathcal{T}_i'$.

3. Do *not* connect any gadget vertices from different gadgets $\mathcal{T}_i^0 \leftrightarrow \mathcal{T}_j^0$.

From now on, we drop the primes on $\mathcal{T}_i'$, $\hat{\Delta}_i'^k$. Let the final weighted graph after adding many gadgets in this way be denoted $\hat{\mathcal{G}}_n$, with Laplacian $\hat{\Delta}^k$.

Fact 6.11. *After adding many gadgets with this procedure, the total chain space can be decomposed as*

$$C^{2n-1}(\hat{\mathcal{G}}_n) = C^{2n-1}(\mathcal{G}_n) \oplus C^{2n-1}(\mathcal{T}_1) \oplus \cdots \oplus C^{2n-1}(\mathcal{T}_t). \tag{6.185}$$

The following claim says that the *up* Laplacian respects this decomposition.

Claim 6.12. *When acting on the entire chainspace $C^{2n-1}(\hat{\mathcal{G}}_n)$, we can write the* up *Laplacian $\hat{\Delta}^{\uparrow 2n-1}$ of the entire complex $\hat{\mathcal{G}}_n$ as the sum of the up Laplacians $\hat{\Delta}_i^{\uparrow 2n-1}$ of the individual gadgets $C^{2n-1}(\mathcal{G}_n) \oplus C^{2n-1}(\mathcal{T}_i)$.*

$$\hat{\Delta}^{\uparrow 2n-1} = \hat{\Delta}_1^{\uparrow 2n-1} + \cdots + \hat{\Delta}_t^{\uparrow 2n-1} \tag{6.186}$$

Proof. It is sufficient to check

$$\langle\psi|\hat{\Delta}^{\uparrow 2n-1}|\psi\rangle = \langle\psi_1|\hat{\Delta}_1^{\uparrow 2n-1}|\psi_1\rangle + \cdots + \langle\psi_t|\hat{\Delta}_t^{\uparrow 2n-1}|\psi_t\rangle \tag{6.187}$$

for all states $|\psi\rangle \in C^{2n-1}(\hat{\mathcal{G}}_n)$, where $|\psi_i\rangle$ is the component in $C^{2n-1}(\mathcal{G}_n) \oplus C^{2n-1}(\mathcal{T}_i)$. But

$$\begin{aligned}
\langle\psi|\hat{\Delta}^{\uparrow 2n-1}|\psi\rangle &= ||\hat{d}^{2n-1}|\psi\rangle||^2 && (6.188)\\
&= ||\hat{d}_1^{2n-1}|\psi_1\rangle||^2 + \cdots + ||\hat{d}_t^{2n-1}|\psi_t\rangle||^2 && (6.189)\\
&= \langle\psi_1|\hat{\Delta}_1^{\uparrow 2n-1}|\psi_1\rangle + \cdots + \langle\psi_t|\hat{\Delta}_t^{\uparrow 2n-1}|\psi_t\rangle, && (6.190)
\end{aligned}$$

where in the second line, we used that the different gadgets do not share any $2n$-simplices. That is, there are no $2n$-simplices which contain vertices from more than one gadget. □

The hope is that, after this procedure to combine the gadgets, they will implement a version of the Hamiltonian H on the simulated qubit subspace $\mathcal{H}_n$. This will be reflected in our main theorem Theorem 6.14, which states that a $1/\text{poly}(n)$ lower

bound on the spectrum of H results in a $1/\text{poly}(n)$ lower bound on the spectrum of the Laplacian $\hat{\Delta}^{2n-1}$.

The high level overview of the proof of Theorem 6.14 is as follows. Assume that $\lambda_{\min}(H) \geq 1/\text{poly}(n)$. We will argue by contradiction. Suppose there is a state $|\varphi\rangle$ of extremely low energy on $\hat{\Delta}^{2n-1}$. Consider a single gadget $\mathcal{T}_i$. $|\varphi\rangle$ is forced to have small overlap with the excited eigenspaces of the single-gadget Laplacian $\hat{\Delta}_i^{2n-1}$. Thus the restriction of $|\varphi\rangle$ to $C^{2n-1}(\hat{\mathcal{G}}_n) \oplus C^{2n-1}(\mathcal{T}_i)$ must lie close to $\text{Ker}\,\hat{\Delta}_i^{2n-1}$. Lemma 6.13 tells us that $\text{Ker}\,\hat{\Delta}_i^{2n-1}$ is close to the space in the simulated qubit subspace $\mathcal{H}_n$ which is *orthogonal to* the state $|\phi_i\rangle$ which is being filled in by gadget $\mathcal{T}_i$. Thus the restriction of $|\varphi\rangle$ to $C^{2n-1}(\hat{\mathcal{G}}_n)$ is close to the subspace of $\mathcal{H}_n$ orthogonal to $|\phi_i\rangle$. This must hold for each gadget $\mathcal{T}_i$, so the restriction of $|\varphi\rangle$ to $C^{2n-1}(\hat{\mathcal{G}}_n)$ must be in $\mathcal{H}_n$ and simultaneously orthogonal to all states $\{|\phi_i\rangle\}$. But the $1/\text{poly}(n)$ spectrum lower bound on $H = \sum_i \phi_i$ forbids this.

Theorem 6.14. (Main theorem, formal) *Suppose H is a m-local Hamiltonian on n qubits with t terms, where each term is a rank-1 projector onto an integer state. Starting from H, let $\hat{\mathcal{G}}_n$ be the weighted graph described above, with Laplacian $\hat{\Delta}^k$. For any $g > 0$, there is a constant $c > 0$ sufficiently small such that setting*

$$\lambda = ct^{-1}g \tag{6.191}$$

$$E = c\lambda^{4m+2}t^{-1}g \tag{6.192}$$

gives

$$\lambda_{\min}(H) = 0 \implies \lambda_{\min}(\hat{\Delta}^{2n-1}) = 0 \tag{6.193}$$

$$\lambda_{\min}(H) \geq g \implies \lambda_{\min}(\hat{\Delta}^{2n-1}) \geq E \tag{6.194}$$

Proof. First we show that $\lambda_{\min}(H) = 0 \implies \lambda_{\min}(\hat{\Delta}^{2n-1}) = 0$. We first argue that the complex $\hat{\mathcal{G}}_n$ has some $(2n-1)$-homology, and then we invoke Proposition 6.3.

Recall that the clique complex of the graph $\mathcal{G}_n$ (which is the n-fold join of the initial qubit graph $\mathcal{G}_1$) has a homology group with rank 2^n. This graph corresponds to the zero Hamiltonian—every state in the Hilbert space of n qubits is a zero-energy ground state. When we add the gadget for the term $|\phi_i\rangle\langle\phi_i|$ to the graph $\mathcal{G}_n$, we fill in the cycles $|\phi_i\rangle \otimes \mathbb{C}^{\otimes(n-m)}$ by rendering them the boundaries of some $2n$-dimensional objects. Crucially, we do not fill in any cycles other than $|\phi_i\rangle \otimes \mathbb{C}^{\otimes(n-m)}$ by adding this gadget. Thus any state $|\psi\rangle$ which is orthogonal to $\text{span}\{|\phi_i\rangle\}$ will still give an

element of homology. The state $|\psi\rangle$ satisfying $H|\psi\rangle = 0$ corresponds to a cycle $|\psi\rangle \in \mathcal{H}_n$ which is orthogonal to $\text{span}\{|\phi_i\rangle\}$. Thus $\hat{\mathcal{G}}_n$ has non-trivial $(2n-1)$-homology $H^{2n-1}(\hat{\mathcal{G}}_n) \neq \{0\}$. Invoking Proposition 6.3 tells us that $\text{Ker}\,\hat{\Delta}^{2n-1} \neq \{0\}$. That is, $\lambda_{\min}(\hat{\Delta}^{2n-1}) = 0$.

It remains to tackle the case $\lambda_{\min}(H) \geq g$. We begin with some notation.

Definition 6.17.

- *Let $\hat{\partial}^k$, $\hat{d}^k$, $\hat{\Delta}^k$ be the boundary, coboundary maps and Laplacian on the entire complex with all gadgets $C^k(\hat{\mathcal{G}}_n)$, with ∂^k, d^k, Δ^k still reserved for the original qubit complex $C^k(\mathcal{G}_n)$.*

- *Let $\hat{\partial}_i^k$, $\hat{d}_i^k$, $\hat{\Delta}_i^k$ be the boundary, coboundary maps and Laplacian on the single gadget complex $C^k(\mathcal{G}_n) \oplus C^k(\mathcal{T}_i)$.*

- *We can decompose the chainspaces of the entire complex as*

$$C^k(\hat{\mathcal{G}}_n) = C^k(\mathcal{G}_n) \oplus C^k(\mathcal{T}_1) \oplus \cdots \oplus C^k(\mathcal{T}_t). \tag{6.195}$$

  Let Π_0^k be projection onto $C^k(\mathcal{G}_n)$ and Π_i^k projection onto $C^k(\mathcal{T}_i)$. These are a complete set of projectors on $C^k(\hat{\mathcal{G}})$:

$$\Pi_0^k + \sum_i \Pi_i^k = id. \tag{6.196}$$

- *Let $[\text{bulk}]_i$ denote the simplices touching the central vertex v_0 of gadget $\mathcal{T}_i$ (see Section 6.2), with chainspaces $C^k([\text{bulk}])$. Let $[\text{bulk}] = \sqcup_i [\text{bulk}]_i$, so $C^k([\text{bulk}]) = \bigoplus_i C^k([\text{bulk}]_i)$. Let $\Pi_i^{[k]}$ be the projection onto $C^k([\text{bulk}]_i)$, and $\Pi^{[k]} = \bigoplus_i \Pi_i^{[k]}$ projection onto $C^k([\text{bulk}])$.*

- *Lemma 6.13 gives us an orthogonal decomposition of $C^{2n-1}(\mathcal{G}_n) \oplus C^{2n-1}(\mathcal{T}_i)$ into four spaces, for each gadget $\mathcal{T}_i$.*

  1. *The space of eigenvectors of eigenvalues $\Theta(1)$ - call it $\mathcal{A}_i$. Let $\Pi_i^{(\mathcal{A})}$ be the projection onto $\mathcal{A}_i$.*
  2. *The space of eigenvectors of eigenvalues $\Theta(\lambda^2)$ - call it $\mathcal{B}_i$. Let $\Pi_i^{(\mathcal{B})}$ be the projection onto $\mathcal{B}_i$. $\mathcal{B}_i$ is a $O(\lambda)$-perturbation of $C^{2n-1}([\text{bulk}]_i)$.*
  3. *The eigenspace with eigenvalue $\Theta(\lambda^{4m_i+2})$. Let $\hat{\Phi}_i$ project onto this space. $\text{Im}\,\hat{\Phi}_i$ is a $O(\lambda)$-perturbation of $|\phi_i\rangle \otimes \mathcal{H}_{n-m} \subseteq \mathcal{H}_n$. Let Φ_i project onto $|\phi_i\rangle \otimes \mathcal{H}_{n-m}$.*

4. *The kernel of* $\hat{\Delta}_i^{2n-1}$. *Let* $\hat{\Phi}_i^{\perp}$ *project onto* $\operatorname{Ker} \hat{\Delta}_i^{2n-1}$. $\operatorname{Im} \hat{\Phi}_i^{\perp}$ *is a* $O(\lambda)$*-perturbation of* $|\phi_i\rangle^{\perp} \otimes \mathcal{H}_{n-m} \subseteq \mathcal{H}_n$. *Let* $\Phi_i^{\perp}$ *project onto* $|\phi_i\rangle^{\perp} \otimes \mathcal{H}_{n-m}$.

We have a complete set of projectors

$$\Pi_i^{(\mathcal{A})} + \Pi_i^{(\mathcal{B})} + \hat{\Phi}_i + \hat{\Phi}_i^{\perp} = \Pi_0^{2n-1} + \Pi_i^{2n-1}. \tag{6.197}$$

- *Finally, let* $\Pi_n^{(\mathcal{H})}$ *be the projection onto* $\mathcal{H}_n$*, with image in* $C^{2n-1}(\mathcal{G}_n)$.

Now for the proof. Assume $\lambda_{\min}(H) \geq g$, and assume for contradiction that the *normalized* state $|\varphi\rangle \in C^{2n-1}(\hat{\mathcal{G}}_n)$ has

$$\langle\varphi|\hat{\Delta}^{2n-1}|\varphi\rangle < E, \tag{6.198}$$

where E is defined in the theorem statement. We will show a contradiction by deriving $\langle\varphi|\varphi\rangle < 1$.

Consider the following calculation.

$$\langle\varphi|\varphi\rangle = \langle\varphi|\big(\Pi_0^{2n-1} + \sum_i \Pi_i^{2n-1}\big)|\varphi\rangle \tag{6.199}$$

$$= \sum_i \langle\varphi|\big(\Pi_0^{2n-1} + \Pi_i^{2n-1}\big)|\varphi\rangle - (t-1)\langle\varphi|\Pi_0^{2n-1}|\varphi\rangle \tag{6.200}$$

$$= \sum_i \langle\varphi|\big(\hat{\Phi}_i^{\perp} + \hat{\Phi}_i + \Pi_i^{(\mathcal{A})} + \Pi_i^{(\mathcal{B})}\big)|\varphi\rangle - (t-1)\langle\varphi|\Pi_0^{2n-1}|\varphi\rangle \tag{6.201}$$

$$= \left(\langle\varphi| \sum_i \hat{\Phi}_i^{\perp}|\varphi\rangle - (t-1)\langle\varphi|\Pi_0^{2n-1}|\varphi\rangle\right) + \langle\varphi| \sum_i \hat{\Phi}_i|\varphi\rangle \tag{6.202}$$

$$+ \langle\varphi| \sum_i \Pi_i^{(\mathcal{A})}|\varphi\rangle + \langle\varphi| \sum_i \Pi_i^{(\mathcal{B})}|\varphi\rangle. \tag{6.203}$$

Here we used Equations 6.197 and 6.196.

Examining the term in brackets, we have

$$\langle\varphi| \sum_i \hat{\Phi}_i^{\perp}|\varphi\rangle - (t-1)\langle\varphi|\Pi_0^{2n-1}|\varphi\rangle \tag{6.204}$$

$$= \langle\varphi| \sum_i \big(\Phi_i^{\perp} + O(\lambda)\big)|\varphi\rangle - (t-1)\langle\varphi|\Pi_0^{2n-1}|\varphi\rangle \tag{6.205}$$

$$= \langle\varphi| \sum_i \Phi_i^{\perp}|\varphi\rangle - (t-1)\langle\varphi|\Pi_0^{2n-1}|\varphi\rangle + O(\lambda t) \tag{6.206}$$

$$= \langle\varphi| \sum_i \big(\Pi_n^{(\mathcal{H})} - \Phi_i\big)|\varphi\rangle - (t-1)\langle\varphi|\Pi_0^{2n-1}|\varphi\rangle + O(\lambda t) \tag{6.207}$$

$$= \langle\varphi|\Pi_n^{(\mathcal{H})}|\varphi\rangle - \langle\varphi| \sum_i \Phi_i|\varphi\rangle - (t-1)\langle\varphi|\big(\Pi_0^{2n-1} - \Pi_n^{(\mathcal{H})}\big)|\varphi\rangle + O(\lambda t). \tag{6.208}$$

In the second line we used a consequence of Lemma 6.13, combined with Part 1 of Lemma 6.8: $\hat{\Phi}_i^{\perp} = \Phi_i^{\perp} + O(\lambda)$.

Consider the term $\langle\varphi| \sum_i \Phi_i |\varphi\rangle$. $\sum_i \Phi_i$ is precisely the implementation of the Hamiltonian H on the simulated n-qubit subspace $\mathcal{H}_n$. This is where we use the assumption on the minimum eigenvalue of H.

$$\langle\varphi| \sum_i \Phi_i |\varphi\rangle = \big(\langle\varphi|\Pi_n^{(\mathcal{H})}\big) \sum_i \Phi_i \big(\Pi_n^{(\mathcal{H})}|\varphi\rangle\big) \tag{6.209}$$

$$\geq g \; ||\Pi_n^{(\mathcal{H})}|\varphi\rangle||^2 \tag{6.210}$$

$$= g \; \langle\varphi|\Pi_n^{(\mathcal{H})}|\varphi\rangle. \tag{6.211}$$

Putting this all together, we have

$$\langle\varphi|\varphi\rangle \leq \langle\varphi|\Pi_n^{(\mathcal{H})}|\varphi\rangle \cdot (1-g) - (t-1)\langle\varphi|\big(\Pi_0^{2n-1} - \Pi_n^{(\mathcal{H})}\big)|\varphi\rangle + O(\lambda t) \tag{6.212}$$

$$+ \langle\varphi| \sum_i \hat{\Phi}_i |\varphi\rangle + \langle\varphi| \sum_i \Pi_i^{(\mathcal{A})}|\varphi\rangle + \langle\varphi| \sum_i \Pi_i^{(\mathcal{B})}|\varphi\rangle. \tag{6.213}$$

The first term is strictly less than 1, and the second term is non-positive since $\Pi_0^{2n-1} - \Pi_n^{(\mathcal{H})} \succeq 0$. If we can somehow show that the final three terms are small, then this will give the contradiction $\langle\varphi|\varphi\rangle < 1$. And indeed, it is true that the low-energy assumption $\langle\varphi|\hat{\Delta}^{2n-1}|\varphi\rangle < E$ forces the terms involving $\{\hat{\Phi}^i\}$, $\{\Pi_i^{(\mathcal{A})}\}$, $\{\Pi_i^{(\mathcal{B})}\}$ to be small. This forms the content of Lemmas 6.15, 6.16, 6.17, whose proofs are postponed to Section 6.5.

Lemma 6.15. *For m-local Hamiltonian H with t terms, recall $\hat{\Delta}^k$ is the Laplacian of the corresponding graph $\hat{\mathcal{G}}_n$. Let $\{\hat{\Phi}_i\}$ be as defined in Definition 6.17. If $|\varphi\rangle \in C^{2n-1}(\hat{\mathcal{G}}_n)$ is a state with $\langle\varphi|\hat{\Delta}^{2n-1}|\varphi\rangle < E$, then*

$$\langle\varphi| \sum_i \hat{\Phi}_i |\varphi\rangle = O(\lambda^{-(4m+2)} E t) \tag{6.214}$$

Lemma 6.16. *For m-local Hamiltonian H with t terms, recall $\hat{\Delta}^k$ is the Laplacian of the corresponding graph $\hat{\mathcal{G}}_n$. Let $\{\Pi_i^{(\mathcal{A})}\}$ be as defined in Definition 6.17. If $|\varphi\rangle \in C^{2n-1}(\hat{\mathcal{G}}_n)$ is a state with $\langle\varphi|\hat{\Delta}^{2n-1}|\varphi\rangle < O(\lambda^2)$, then*

$$\langle\varphi| \sum_i \Pi_i^{(\mathcal{A})}|\varphi\rangle = O(\lambda^2 t) \tag{6.215}$$

Lemma 6.17. *For m-local Hamiltonian H with t terms, recall $\hat{\Delta}^k$ is the Laplacian of the corresponding graph $\hat{\mathcal{G}}_n$. Let $\{\Pi_i^{(\mathcal{B})}\}$ be as defined in Definition 6.17. If $|\varphi\rangle \in C^{2n-1}(\hat{\mathcal{G}}_n)$ is a state with $\langle\varphi|\hat{\Delta}^{2n-1}|\varphi\rangle < O(\lambda^4)$, then*

$$\langle\varphi| \sum_i \Pi_i^{(\mathcal{B})}|\varphi\rangle = O(\lambda^2 t) \tag{6.216}$$

Note that we set $E = o(\lambda^{4m+2})$ in the statement of Theorem 6.14, so $\langle\varphi|\hat{\Delta}^{2n-1}|\varphi\rangle < E$ implies the conditions of Lemmas 6.16, 6.17. Returning to the calculation armed with Lemmas 6.15, 6.16, 6.17, we get

$$\langle\varphi|\varphi\rangle = \langle\varphi|\Pi_n^{(\mathcal{H})}|\varphi\rangle \cdot (1-g) - (t-1)\langle\varphi|\big(\Pi_0^{2n-1} - \Pi_n^{(\mathcal{H})}\big)|\varphi\rangle + O(\lambda t) \tag{6.217}$$

$$+ O(\lambda^{-(4m+2)}Et) + O(\lambda^2 t) + O(\lambda^2 t) \tag{6.218}$$

Choosing

$$\lambda = ct^{-1}g \tag{6.219}$$

$$E = c\lambda^{4m+2}t^{-1}g \tag{6.220}$$

for a constant c sufficiently small, this becomes

$$1 = \langle\varphi|\varphi\rangle \tag{6.221}$$

$$\leq \langle\varphi|\Pi_n^{(\mathcal{H})}|\varphi\rangle \cdot (1-g) - (t-1)\langle\varphi|\big(\Pi_0^{2n-1} - \Pi_n^{(\mathcal{H})}\big)|\varphi\rangle + \frac{1}{10}g \tag{6.222}$$

$$\leq 1 - g + \frac{1}{10}g \tag{6.223}$$

$$< 1. \tag{6.224}$$

a contradiction. This concludes the proof of Theorem 6.14. □

6.5 Postponed proofs

In this section, we prove Lemmas 6.15, 6.16, 6.17 one at a time. The proof of Lemma 6.15 is comparatively simple, since we need only use that the energy of $|\varphi\rangle$ on the *up* Laplacian is small $\langle\varphi|\hat{\Delta}^{\uparrow 2n-1}|\varphi\rangle < E$. Claim 6.12 tells us that the up Laplacians play nicely with the decomposition into separate gadgets, and there is no interference between gadgets. Unfortunately, we start running into more trouble when we must consider also down Laplacians, which *can* interfere between different gadgets. For this reason, the proofs of Lemmas 6.16, 6.17 are more involved.

Here is some notation which will be useful in the proof of Lemmas 6.16, 6.17.

Definition 6.18. *Recall the block decomposition*

$$C^{2n-1}(\hat{\mathcal{G}}_n) = C^{2n-1}(\mathcal{G}_n) \oplus C^{2n-1}(\mathcal{T}_1) \oplus \cdots \oplus C^{2n-1}(\mathcal{T}_t) \tag{6.225}$$

Write

$$|\varphi\rangle = |\omega_0\rangle + |\omega_1\rangle + \cdots + |\omega_t\rangle \tag{6.226}$$

where $|\omega_0\rangle \in C^{2n-1}(\mathcal{G}_n)$, $|\omega_i\rangle \in C^{2n-1}(\mathcal{T}_i)$. *Further, introduce the notation*

$$|\varphi_i\rangle = (\Pi_0^{2n-1} + \Pi_i^{2n-1})|\varphi\rangle = |\omega_0\rangle + |\omega_i\rangle \tag{6.227}$$

Proof of Lemma 6.15

Proof. Recall the definition of $|\hat{\phi}_i\rangle$ from Lemma 6.13. Lemma 6.13 tells us that $|\hat{\phi}_i\rangle \otimes \mathcal{H}_{n-m}$ is a $\hat{\Delta}_i^{2n-1}$-eigenspace with eigenvalue λ^{4m_i+2}. Combined with Claim 6.8, we can see that these states are in fact $\hat{\Delta}_i^{\uparrow 2n-1}$-eigenstates with eigenvalue λ^{4m_i+2}. Thus $\hat{\Delta}_i^{\uparrow 2n-1} \succeq \lambda^{4m_i+2} \cdot \hat{\Phi}_i$; states in $\operatorname{Im} \hat{\Phi}_i$ have energy at least λ^{4m_i+2} on $\hat{\Delta}_i^{\uparrow 2n-1}$, and states orthogonal to $\operatorname{Im} \hat{\Phi}_i$ have energy at least zero. Now $\hat{\Delta}^{2n-1} \succeq \hat{\Delta}^{\uparrow 2n-1} = \hat{\Delta}_1^{\uparrow 2n-1} + \cdots + \hat{\Delta}_t^{\uparrow 2n-1} \succeq \hat{\Delta}_i^{\uparrow 2n-1}$ using Claim 6.12, so in fact $\hat{\Delta}^{2n-1} \succeq \lambda^{4m_i+2} \cdot \hat{\Phi}_i$. But our state $|\varphi\rangle$ only has energy $\langle\varphi|\hat{\Delta}^{2n-1}|\varphi\rangle < E$ on $\hat{\Delta}^{2n-1}$. Thus

$$\langle\varphi|\hat{\Phi}_i|\varphi\rangle \leq \lambda^{-(4m_i+2)} \cdot \langle\varphi|\hat{\Delta}^{2n-1}|\varphi\rangle \tag{6.228}$$

$$= O(\lambda^{-(4m+2)} E) \tag{6.229}$$

$$\implies \langle\varphi| \sum_i \hat{\Phi}_i |\varphi\rangle = O(\lambda^{-(4m+2)} Et) \tag{6.230}$$

□

Proof of Lemma 6.16

Proof. Our strategy is to first show that $|\varphi_i\rangle$ is low energy on $\hat{\Delta}_i^{2n-1}$.

$$\langle\varphi|\hat{\Delta}^{2n-1}|\varphi\rangle < O(\lambda^2) \implies \begin{cases} \langle\varphi|\hat{\Delta}^{\downarrow 2n-1}|\varphi\rangle < O(\lambda^2) & \implies ||\hat{\partial}^{2n-1}|\varphi\rangle||^2 < O(\lambda^2) \\ \langle\varphi|\hat{\Delta}^{\uparrow 2n-1}|\varphi\rangle < O(\lambda^2) & \implies ||\hat{d}^{2n-1}|\varphi\rangle||^2 < O(\lambda^2) \end{cases} \tag{6.231}$$

Now

$$||\hat{d}^{2n-1}|\varphi\rangle||^2 < O(\lambda^2) \tag{6.232}$$

$$\implies ||\hat{d}_1^{2n-1}|\varphi_1\rangle||^2 + \cdots + ||\hat{d}_t^{2n-1}|\varphi_1\rangle||^2 < O(\lambda^2) \tag{6.233}$$

$$\implies \left\|\hat{d}_i^{2n-1}|\varphi_i\rangle\right\|^2 < O(\lambda^2)\ \forall i \tag{6.234}$$

using Claim 6.12. It is not so simple for the boundaries. Write

$$\left\|\hat{\partial}_i^{2n-1}|\varphi_i\rangle\right\|^2 = \left\|\Pi_0^{2n-2}\hat{\partial}_i^{2n-1}|\varphi_i\rangle\right\|^2 + \left\|\Pi_i^{2n-2}\hat{\partial}_i^{2n-1}|\varphi_i\rangle\right\|^2 \tag{6.235}$$

For the restriction to the gadget chainspace $C^{2n-2}(\mathcal{T}_i)$, we indeed have

$$\left\|\Pi_i^{2n-2}\hat{\partial}_i^{2n-1}|\varphi_i\rangle\right\|^2 < O(\lambda^2) \tag{6.236}$$

since the $(2n-2)$-simplices in gadget $\mathcal{T}_i$ are not shared with any other gadgets. (This is a similar reasoning to the proof of Claim 6.12.) However, the $(2n-2)$-simplices in the qubit complex $\mathcal{G}_n^{2n-2}$ touches $(2n-1)$-simplices from many gadgets, and thus

the boundaries of multiple $\{|\varphi_j\rangle\}_j$ could contribute to $\Pi_0^{2n-2}\hat{\partial}^{2n-1}|\varphi\rangle$ and possibly cancel out. Here we need a different argument. (To go from the first line to the second line in what follows, note that $\hat{\partial}^{2n-1}$ and $\hat{\partial}_i^{2n-1}$ have identical actions when acting on $C^{2n-1}(\mathcal{G}) \oplus C^{2n-1}(\mathcal{T}_i)$.)

$$\Pi_0^{2n-2}\hat{\partial}^{2n-1}|\varphi\rangle = \Pi_0^{2n-2}\hat{\partial}^{2n-1}\big(|\omega_0\rangle + |\omega_i\rangle\big) + \Pi_0^{2n-2}\hat{\partial}^{2n-1}\sum_{j\neq i}|\omega_j\rangle \tag{6.237}$$

$$= \Pi_0^{2n-2}\hat{\partial}_i^{2n-1}|\varphi_i\rangle + \Pi_0^{2n-2}\hat{\partial}^{2n-1}\sum_{j\neq i}|\omega_j\rangle \tag{6.238}$$

$$\implies \Pi_0^{2n-2}\hat{\partial}_i^{2n-1}|\varphi_i\rangle = \Pi_0^{2n-2}\hat{\partial}^{2n-1}|\varphi\rangle - \Pi_0^{2n-2}\hat{\partial}^{2n-1}\sum_{j\neq i}|\omega_j\rangle \tag{6.239}$$

$$\implies \left\|\Pi_0^{2n-2}\hat{\partial}_i^{2n-1}|\varphi_i\rangle\right\| \leq \left\|\Pi_0^{2n-2}\hat{\partial}^{2n-1}|\varphi\rangle\right\| + \left\|\Pi_0^{2n-2}\hat{\partial}^{2n-1}\sum_{j\neq i}|\omega_j\rangle\right\| \tag{6.240}$$

$$< O(\lambda) + \left\|\Pi_0^{2n-2}\hat{\partial}^{2n-1}\sum_{j\neq i}|\omega_j\rangle\right\| \tag{6.241}$$

Here we used triangle inequality. But,

$$\left\|\Pi_0^{2n-2}\hat{\partial}^{2n-1}\sum_{j\neq i}|\omega_j\rangle\right\| = O(\lambda) \tag{6.242}$$

since the vertices $\mathcal{T}_j^0$ have weight λ, recalling Equation 6.77. Thus

$$\left\|\Pi_0^{2n-2}\hat{\partial}_i^{2n-1}|\varphi_i\rangle\right\| < O(\lambda) + O(\lambda) = O(\lambda) \tag{6.243}$$

$$\implies \langle\varphi_i|\hat{\Delta}_i^{2n-1}|\varphi_i\rangle = \left\|\hat{d}_i^{2n-1}|\varphi_i\rangle\right\|^2 + \left\|\Pi_i^{2n-2}\hat{\partial}_i^{2n-1}|\varphi_i\rangle\right\|^2 + \left\|\Pi_0^{2n-2}\hat{\partial}_i^{2n-1}|\varphi_i\rangle\right\|^2 \tag{6.244}$$

$$= O(\lambda^2) + O(\lambda^2) = O(\lambda^2) \tag{6.245}$$

$\mathcal{A}_i$ is a $\hat{\Delta}_i^{2n-1}$-eigenspace with eigenvalues $\Theta(1)$, so $\hat{\Delta}_i^{2n-1} \succeq \Theta(1)\cdot\Pi_i^{(\mathcal{A})}$. Thus

$$\langle\varphi_i|\Pi_i^{(\mathcal{A})}|\varphi_i\rangle \leq \Theta(1)\cdot\langle\varphi_i|\hat{\Delta}_i^{2n-1}|\varphi_i\rangle = O(\lambda^2) \tag{6.246}$$

$$\implies \langle\varphi|\sum_i \Pi_i^{(\mathcal{A})}|\varphi\rangle = O(\lambda^2 t) \tag{6.247}$$

□

Proof of Lemma 6.17

Proof. We will use what we are told by Lemma 6.13 about the form of $\mathcal{B}_i$. Namely, $\mathcal{B}_i$ is a $O(\lambda)$-perturbation of the space $C^{2n-1}([\text{bulk}]_i)$ of $(2n-1)$-simplices touching the central vertex of gadget $\mathcal{T}_i$.

We must rule out the possibility that, when we combine many gadgets, new states of very low energy emerge with high overlap on the $\{\mathcal{B}_i\}$ subspaces. What is the situation we are fighting against? We know states with high overlap on a single $\mathcal{B}_i$ subspace must have energy at least $\Omega(\lambda^2)$. But this state could be a cocycle, and perhaps the $\Omega(\lambda)$ boundary lives in the qubit complex $\mathcal{C}^{2n-2}(\mathcal{G})$. If this could happen, then perhaps many such states from different gadgets could be superposed in such a way that their boundaries destructively interfere on the qubit complex, leading to a state of very low energy.

However, Claim 6.10 will let us show that not only must the states in an individual $\mathcal{B}_i$ have coboundaries or boundaries of size $\Omega(\lambda)$, but further these $\Omega(\lambda)$ coboundaries and boundaries must be supported on $[\text{bulk}]_i$.

We will show that, if $\langle\varphi|\hat{\Delta}^{2n-1}|\varphi\rangle < O(\lambda^4)$, then $\langle\varphi_i|\Pi_i^{(\mathcal{B})}|\varphi_i\rangle = O(\lambda^2)$ for each i. Lemma 6.13 and Part 2 of Lemma 6.8 tell us that it is enough to show $\langle\varphi_i|\Pi_i^{[2n-1]}|\varphi_i\rangle = O(\lambda^2)$ for each i. We will show the contrapositive. That is, we will assume

$$\langle\varphi_i|\Pi_i^{[2n-1]}|\varphi_i\rangle \notin O(\lambda^2) \tag{6.248}$$

and aim to derive a contradiction with $\langle\varphi|\hat{\Delta}^{2n-1}|\varphi\rangle \leq O(\lambda^4)$. Our strategy to do this is to use Claim 6.10 to first derive that either $||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}|\varphi_i\rangle||$ or $||\Pi_i^{[2n]}\hat{d}_i^{2n-1}|\varphi_i\rangle||$ must be big. That is, the components of the boundary and coboundary on $[\text{bulk}]_i$ cannot both be small. Then since $\{[\text{bulk}_i]\}$ are separated from each other and 'cannot interfere', this will necessitate that $\langle\varphi|\hat{\Delta}^{2n-1}|\varphi\rangle$ is big, providing the contradiction.

We will use the subspace decomposition

$$\begin{aligned}
\Pi_0^{2n-1} + \Pi_i^{2n-1} &= \Pi_i^{[2n-1]} + \Pi_i^{[2n-1]\perp} \\
&= \Pi_i^{[2n-1]} + \Pi_i^{[2n-1]\perp}\big(\Pi_i^{(\mathcal{A})} + \Pi_i^{(\mathcal{B})} + \hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big) \\
&= \Pi_i^{[2n-1]} + \Pi_i^{[2n-1]\perp}\Pi_i^{(\mathcal{A})} + \Pi_i^{[2n-1]\perp}\Pi_i^{(\mathcal{B})} && (6.249) \\
&\quad + \big(\hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big) - \Pi_i^{[2n-1]}\big(\hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big), && (6.250)
\end{aligned}$$

where $\Pi_i^{[2n-1]\perp} := \Pi_0^{2n-1} + \Pi_i^{2n-1} - \Pi_i^{[2n-1]}$.

Let's apply Claim 6.10 to the state $\Pi_i^{[2n-1]}|\varphi_i\rangle \in \mathcal{C}^{2n-1}([\text{bulk}]_i)$. Equation 6.248 gives $||\Pi_i^{[2n-1]}|\varphi_i\rangle|| \notin O(\lambda)$. From this, the two cases from Claim 6.10 are $||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\Pi_i^{[2n-1]}|\varphi_i\rangle|| \notin O(\lambda^2)$ and $||\Pi_i^{[2n]}\hat{d}_i^{2n-1}\Pi_i^{[2n-1]}|\varphi_i\rangle|| \notin O(\lambda^2)$.

Case 1. $||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\Pi_i^{[2n-1]}|\varphi_i\rangle|| \notin O(\lambda^2)$

In this case, we will show $||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}|\varphi_i\rangle||$ cannot be small. Using Equation 6.250,

$$\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}|\varphi_i\rangle = \Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\Pi_i^{[2n-1]}|\varphi_i\rangle \tag{6.251}$$
$$+ \Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\Pi_i^{[2n-1]\perp}\Pi_i^{(\mathcal{A})}|\varphi_i\rangle \tag{6.252}$$
$$+ \Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\Pi_i^{[2n-1]\perp}\Pi_i^{(\mathcal{B})}|\varphi_i\rangle \tag{6.253}$$
$$+ \Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\big(\hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big)|\varphi_i\rangle \tag{6.254}$$
$$- \Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\Pi_i^{[2n-1]}\big(\hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big)|\varphi_i\rangle \tag{6.255}$$
$$\implies ||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}|\varphi_i\rangle|| \geq ||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\Pi_i^{[2n-1]}|\varphi_i\rangle|| \tag{6.256}$$
$$- ||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\Pi_i^{[2n-1]\perp}\Pi_i^{(\mathcal{A})}|\varphi_i\rangle|| \tag{6.257}$$
$$- ||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\Pi_i^{[2n-1]\perp}\Pi_i^{(\mathcal{B})}|\varphi_i\rangle|| \tag{6.258}$$
$$- ||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\big(\hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big)|\varphi_i\rangle|| \tag{6.259}$$
$$- ||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\Pi_i^{[2n-1]}\big(\hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big)|\varphi_i\rangle|| \tag{6.260}$$

by triangle inequality. We have $||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\Pi_i^{[2n-1]}|\varphi_i\rangle|| \notin O(\lambda^2)$ by assumption. To conclude that $||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}|\varphi_i\rangle||$ is big, we will argue that the remaining terms on the right hand side are small.

We know from before that $||\Pi_i^{(\mathcal{A})}|\varphi_i\rangle|| = O(\lambda)$, thus by Claim 6.9 $||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\Pi_i^{[2n-1]\perp}\Pi_i^{(\mathcal{A})}|\varphi_i\rangle|| = O(\lambda^2)$. From Lemma 6.13, $||\Pi_i^{[2n-1]\perp}\Pi_i^{(\mathcal{B})}|| = ||\Pi_i^{[2n-1]\perp}\big(\Pi_i^{[2n-1]} + O(\lambda)\big)|| = O(\lambda)$ and $||\Pi_i^{[2n-1]}\big(\hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big)|| = ||\Pi_i^{[2n-1]}\big(\Phi_i + \Phi_i^{\perp} + O(\lambda)\big)|| = O(\lambda)$. Thus by Claim 6.9 $||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\Pi_i^{[2n-1]\perp}\Pi_i^{(\mathcal{B})}|\varphi_i\rangle|| = O(\lambda^2)$ and $||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\Pi_i^{[2n-1]}\big(\hat{\Phi}_i+\hat{\Phi}_i^{\perp}\big)|\varphi_i\rangle|| = O(\lambda^2)$. Finally, $||\hat{\partial}_i^{2n-1}\big(\hat{\Phi}_i+\hat{\Phi}_i^{\perp}\big)|\varphi_i\rangle|| = O(\lambda^{2m_i+1}) = O(\lambda^2)$ so $||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}\big(\hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big)|\varphi_i\rangle|| = O(\lambda^2)$. Altogether, we get

$$||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}|\varphi_i\rangle|| \notin O(\lambda^2) \tag{6.261}$$

Case 2. $||\Pi_i^{[2n]}\hat{d}_i^{2n-1}\Pi_i^{[2n-1]}|\varphi_i\rangle|| \notin O(\lambda^2)$

This will be similar to Case 1. Again using Equation 6.250,

$$\Pi_i^{[2n]}\hat{d}_i^{2n-1}|\varphi_i\rangle = \Pi_i^{[2n]}\hat{d}_i^{2n-1}\Pi_i^{[2n-1]}|\varphi_i\rangle \tag{6.262}$$
$$+ \Pi_i^{[2n]}\hat{d}_i^{2n-1}\Pi_i^{[2n-1]\perp}\Pi_i^{(\mathcal{A})}|\varphi_i\rangle \tag{6.263}$$
$$+ \Pi_i^{[2n]}\hat{d}_i^{2n-1}\Pi_i^{[2n-1]\perp}\Pi_i^{(\mathcal{B})}|\varphi_i\rangle \tag{6.264}$$
$$+ \Pi_i^{[2n]}\hat{d}_i^{2n-1}\big(\hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big)|\varphi_i\rangle \tag{6.265}$$
$$- \Pi_i^{[2n]}\hat{d}_i^{2n-1}\Pi_i^{[2n-1]}\big(\hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big)|\varphi_i\rangle \tag{6.266}$$
$$\implies ||\Pi_i^{[2n]}\hat{d}_i^{2n-1}|\varphi_i\rangle|| \geq ||\Pi_i^{[2n]}\hat{d}_i^{2n-1}\Pi_i^{[2n-1]}|\varphi_i\rangle|| \tag{6.267}$$
$$- ||\Pi_i^{[2n]}\hat{d}_i^{2n-1}\Pi_i^{[2n-1]\perp}\Pi_i^{(\mathcal{A})}|\varphi_i\rangle|| \tag{6.268}$$
$$- ||\Pi_i^{[2n]}\hat{d}_i^{2n-1}\Pi_i^{[2n-1]\perp}\Pi_i^{(\mathcal{B})}|\varphi_i\rangle|| \tag{6.269}$$
$$- ||\Pi_i^{[2n]}\hat{d}_i^{2n-1}\big(\hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big)|\varphi_i\rangle|| \tag{6.270}$$
$$- ||\Pi_i^{[2n]}\hat{d}_i^{2n-1}\Pi_i^{[2n-1]}\big(\hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big)|\varphi_i\rangle|| \tag{6.271}$$

By assumption, $||\Pi_i^{[2n]}\hat{d}_i^{2n-1}\Pi_i^{[2n-1]}|\varphi_i\rangle|| \notin O(\lambda^2)$. We know from before that $||\Pi_i^{(\mathcal{A})}|\varphi_i\rangle|| = O(\lambda)$, thus by Claim 6.9 $||\Pi_i^{[2n]}\hat{d}_i^{2n-1}\Pi_i^{[2n-1]\perp}\Pi_i^{(\mathcal{A})}|\varphi_i\rangle|| = O(\lambda^2)$. From Lemma 6.13, $||\Pi_i^{[2n-1]\perp}\Pi_i^{(\mathcal{B})}|| = ||\Pi_i^{[2n-1]\perp}\big(\Pi_i^{[2n-1]} + O(\lambda)\big)|| = O(\lambda)$ and $||\Pi_i^{[2n-1]}\big(\hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big)|| = ||\Pi_i^{[2n-1]}\big(\Phi_i + \Phi_i^{\perp} + O(\lambda)\big)|| = O(\lambda)$. Thus by Claim 6.9 $||\Pi_i^{[2n]}\hat{d}_i^{2n-1}\Pi_i^{[2n-1]\perp}\Pi_i^{(\mathcal{B})}|\varphi_i\rangle|| = O(\lambda^2)$ and $||\Pi_i^{[2n]}\hat{d}_i^{2n-1}\Pi_i^{[2n-1]}\big(\hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big)|\varphi_i\rangle|| = O(\lambda^2)$. Finally, $||\hat{d}_i^{2n-1}\big(\hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big)|\varphi_i\rangle|| = O(\lambda^{2m_i+1}) = O(\lambda^2)$ so $||\Pi_i^{[2n]}\hat{d}_i^{2n-1}\big(\hat{\Phi}_i + \hat{\Phi}_i^{\perp}\big)|\varphi_i\rangle|| = O(\lambda^2)$. Altogether, we get

$$||\Pi_i^{[2n]}\hat{d}_i^{2n-1}|\varphi_i\rangle|| \notin O(\lambda^2) \tag{6.272}$$

We have concluded that either $||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}|\varphi_i\rangle|| \notin O(\lambda^2)$ or $||\Pi_i^{[2n]}\hat{d}_i^{2n-1}|\varphi_i\rangle|| \notin O(\lambda^2)$, so

$$||\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}|\varphi_i\rangle||^2 + ||\Pi_i^{[2n]}\hat{d}_i^{2n-1}|\varphi_i\rangle||^2 \notin O(\lambda^4) \tag{6.273}$$
$$\implies \langle\varphi_i|\hat{d}_i^{2n-2}\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}|\varphi_i\rangle + \langle\varphi_i|\hat{\partial}_i^{2n}\Pi_i^{[2n]}\hat{d}_i^{2n-1}|\varphi_i\rangle \notin O(\lambda^4) \tag{6.274}$$

Now consider the operators $\hat{d}_i^{2n-2}\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1}$, $\hat{\partial}_i^{2n}\Pi_i^{[2n]}\hat{d}_i^{2n-1}$. On $[\text{bulk}]_i$, $\hat{\partial}_i^k$ and $\hat{d}_i^k$ are the same as $\hat{\partial}^k$ and $\hat{d}^k$, thus

$$\hat{d}_i^{2n-2}\Pi_i^{[2n-2]}\hat{\partial}_i^{2n-1} = \hat{d}^{2n-2}\Pi_i^{[2n-2]}\hat{\partial}^{2n-1} \tag{6.275}$$
$$\hat{\partial}_i^{2n}\Pi_i^{[2n]}\hat{d}_i^{2n-1} = \hat{\partial}^{2n}\Pi_i^{[2n]}\hat{d}^{2n-1} \tag{6.276}$$

So now we can deduce

$$\langle\varphi|\hat{\Delta}^{2n-1}|\varphi\rangle = \langle\varphi|\hat{d}^{2n-2}\hat{\partial}^{2n-1}|\varphi\rangle + \langle\varphi|\hat{\partial}^{2n}\hat{d}^{2n-1}|\varphi\rangle \tag{6.277}$$

$$\geq \langle\varphi|\hat{d}^{2n-2}\Pi_i^{[2n-2]}\hat{\partial}^{2n-1}|\varphi\rangle + \langle\varphi|\hat{\partial}^{2n}\Pi_i^{[2n]}\hat{d}^{2n-1}|\varphi\rangle \tag{6.278}$$

$$\notin \mathcal{O}(\lambda^4) \tag{6.279}$$

a contradiction to $\langle\varphi|\hat{\Delta}^{2n-1}|\varphi\rangle \leq \mathcal{O}(\lambda^4)$.

We have concluded that $\langle\varphi_i|\Pi_i^{(\mathcal{B})}|\varphi_i\rangle = \mathcal{O}(\lambda^2)$. This tells us

$$\langle\varphi|\sum_i \Pi_i^{(\mathcal{B})}|\varphi\rangle = \mathcal{O}(\lambda^2 t) \tag{6.280}$$

□

BIBLIOGRAPHY

[23] "Suppressing quantum errors by scaling a surface code logical qubit". In: *Nature* 614.7949 (2023), pp. 676–681.

[AA23] Anurag Anshu and Srinivasan Arunachalam. "A survey on the complexity of learning quantum states". In: *Nature Reviews Physics* (2023), pp. 1–11.

[Aar+18] Scott Aaronson et al. "Online learning of quantum states". In: *Advances in neural information processing systems* 31 (2018).

[Aar04] Scott Aaronson. "Limitations of quantum advice and one-way communication". In: *Proceedings. 19th IEEE Annual Conference on Computational Complexity, 2004.* IEEE. 2004, pp. 320–332.

[Aar08] Scott Aaronson. *On Perfect Completeness for QMA*. 2008. arXiv: 0806.0450 [quant-ph].

[Aar18] Scott Aaronson. "Shadow tomography of quantum states". In: *Proceedings of the 50th annual ACM SIGACT symposium on theory of computing*. 2018, pp. 325–338.

[AAV13] Dorit Aharonov, Itai Arad, and Thomas Vidick. "Guest column: the quantum PCP conjecture". In: *Acm sigact news* 44.2 (2013), pp. 47–79.

[AB22] Anurag Anshu and Nikolas P. Breuckmann. "A construction of combinatorial NLTS". In: *Journal of Mathematical Physics* 63.12 (Dec. 2022), p. 122201. ISSN: 0022-2488. DOI: 10.1063/5.0113731. URL: https://doi.org/10.1063/5.0113731.

[ABN23] Anurag Anshu, Nikolas P Breuckmann, and Chinmay Nirkhe. "NLTS Hamiltonians from good quantum codes". In: *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 2023, pp. 1090–1096.

[ACQ22] Dorit Aharonov, Jordan Cotler, and Xiao-Liang Qi. "Quantum algorithmic measurement". In: *Nature communications* 13.1 (2022), p. 887.

[ADH97] Leonard M. Adleman, Jonathan DeMarrais, and Ming-Deh A. Huang. "Quantum Computability". In: *SIAM Journal on Computing* 26.5 (1997), pp. 1524–1540. DOI: 10.1137/S0097539795293639. eprint: https://doi.org/10.1137/S0097539795293639. URL: https://doi.org/10.1137/S0097539795293639.

[AG04] Scott Aaronson and Daniel Gottesman. "Improved simulation of stabilizer circuits". In: *Physical Review A* 70.5 (2004), p. 052328.

[AGK23] Eric R Anschuetz, David Gamarnik, and Bobak Kiani. "Combinatorial NLTS from the overlap gap property". In: *arXiv preprint arXiv:2304.00643* (2023).

[AGK24] Eric R Anschuetz, David Gamarnik, and Bobak T Kiani. "Bounds on the ground state energy of quantum p-spin Hamiltonians". In: *arXiv preprint arXiv:2404.07231* (2024).

[Aha+09] Dorit Aharonov et al. "The power of quantum systems on a line". In: *Communications in mathematical physics* 287.1 (2009), pp. 41–65.

[AJL06] Dorit Aharonov, Vaughan Jones, and Zeph Landau. "A polynomial quantum algorithm for approximating the Jones polynomial". In: *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*. 2006, pp. 427–436.

[Akh+22] Ismail Yunus Akhalwaya et al. "Towards quantum advantage on noisy quantum computers". In: *arXiv preprint arXiv:2209.09371* (2022).

[AR19] Scott Aaronson and Guy N Rothblum. "Gentle measurement of quantum states and differential privacy". In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. 2019, pp. 322–333.

[AS16] Michał Adamaszek and Juraj Stacho. "Complexity of simplicial homology and independence complexes of chordal graphs". In: *Computational Geometry* 57 (2016), pp. 8–18.

[Asa+16] Ali Asadian et al. "Heisenberg-Weyl Observables: Bloch vectors in phase space". In: *Physical Review A* 94.1 (2016), p. 010301.

[ASS22] Simon Apers, Sayantan Sen, and Dániel Szabó. "A (simple) classical algorithm for estimating Betti numbers". In: *arXiv preprint arXiv:2211.09618* (2022).

[Bab+18a] Ryan Babbush et al. "Encoding Electronic Spectra in Quantum Circuits with Linear T Complexity". In: *Phys. Rev. X* 8 (4 Oct. 2018), p. 041015. DOI: 10.1103/PhysRevX.8.041015. URL: https://link.aps.org/doi/10.1103/PhysRevX.8.041015.

[Bab+18b] Ryan Babbush et al. "Low-Depth Quantum Simulation of Materials". In: *Phys. Rev. X* 8 (1 2018), p. 011044. DOI: 10.1103/PhysRevX.8.011044. URL: https://link.aps.org/doi/10.1103/PhysRevX.8.011044.

[Bab+21] Ryan Babbush et al. "Focus beyond quadratic speedups for error-corrected quantum advantage". In: *PRX quantum* 2.1 (2021), p. 010103.

[Bañ23] Mari Carmen Bañuls. "Tensor network algorithms: A route map". In: *Annual Review of Condensed Matter Physics* 14.1 (2023), pp. 173–191.

[Bau11] Bernhard Baumgartner. "An inequality for the trace of matrix products, using absolute values". In: *arXiv preprint arXiv:1106.6189* (2011).

[BBH21] Afonso S. Bandeira, March Tian Boedihardjo, and Ramon van Handel. "Matrix Concentration Inequalities and Free Probability". In: 2021.

[BBO20] Xavier Bonet-Monroig, Ryan Babbush, and Thomas E O'Brien. "Nearly optimal measurement scheduling for partial tomography of quantum states". In: *Physical Review X* 10.3 (2020), p. 031064.

[BE95] Peter Borwein and Tamas Erdelyi. *Polynomials and Polynomial Inequalities*. Vol. 161. Graduate Texts in Mathematics. New York: Springer, 1995. ISBN: 978-0-387-94509-5.

[Ber+14] Dominic W Berry et al. "Exponential improvement in precision for simulating sparse Hamiltonians". In: *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*. 2014, pp. 283–292.

[Ber+15] Dominic W Berry et al. "Simulating Hamiltonian dynamics with a truncated Taylor series". In: *Physical review letters* 114.9 (2015), p. 090502.

[Ber+18] Dominic W Berry et al. "Improved techniques for preparing eigenstates of fermionic Hamiltonians". In: *npj Quantum Information* 4.1 (2018), p. 22.

[Ber+24a] Dominic W Berry et al. "Analyzing prospects for quantum advantage in topological data analysis". In: *PRX Quantum* 5.1 (2024), p. 010319.

[Ber+24b] Dominic W. Berry et al. *Rapid initial state preparation for the quantum simulation of strongly correlated molecules*. 2024. arXiv: 2409.11748 [quant-ph].

[BF71] Oriol Bohigas and Jorge Flores. "Two-body random Hamiltonian and level density". In: *Physics Letters B* 34.4 (1971), pp. 261–263.

[Bha97] Rajendra Bhatia. *Matrix Analysis*. Springer, 1997.

[BK02] Sergey B Bravyi and Alexei Yu Kitaev. "Fermionic quantum computation". In: *Annals of Physics* 298.1 (2002), pp. 210–226.

[Blu+24] Dolev Bluvstein et al. "Logical quantum processor based on reconfigurable atom arrays". In: *Nature* 626.7997 (2024), pp. 58–65.

[BM03] Samuel Burer and Renato DC Monteiro. "A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization". In: *Mathematical programming* 95.2 (2003), pp. 329–357. URL: https://link.springer.com/article/10.1007/s10107-002-0352-8.

[BO21] Costin Bădescu and Ryan O'Donnell. "Improved quantum data analysis". In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 1398–1411.

[Bra+02] Gilles Brassard et al. “Quantum amplitude amplification and estimation”. In: *Contemporary Mathematics* 305 (2002), pp. 53–74. URL: https://arxiv.org/abs/quant-ph/0005055.

[Bra+17] Fernando GSL Brandão et al. “Quantum SDP solvers: Large speed-ups, optimality, and applications to quantum learning”. In: *arXiv preprint arXiv:1710.02581* (2017).

[Bra+19] Sergey Bravyi et al. “Approximation algorithms for quantum many-body problems”. In: *Journal of Mathematical Physics* 60.3 (2019).

[Bra+23] Anthony J Brady et al. “Advances in Bosonic Quantum Error Correction with Gottesman-Kitaev-Preskill Codes: Theory, Engineering and Applications”. In: *arXiv preprint arXiv:2308.02913* (2023).

[Bra08] Fernando GSL Brandão. “Entanglement theory and the quantum simulation of many-body physics”. In: *arXiv preprint arXiv:0810.0026* (2008).

[Bra11] Sergey Bravyi. “Efficient algorithm for a quantum analogue of 2-SAT”. In: *Contemporary Mathematics* 536 (2011), pp. 33–48.

[BS20] CL Baldwin and B Swingle. “Quenched vs Annealed: Glassiness from SK to SYK”. In: *Physical Review X* 10.3 (2020), p. 031026.

[Bur+21a] Vera von Burg et al. “Quantum computing enhanced computational catalysis”. In: *Physical Review Research* 3.3 (July 2021), p. 033055. ISSN: 2643-1564. DOI: 10.1103/physrevresearch.3.033055. URL: http://dx.doi.org/10.1103/PhysRevResearch.3.033055.

[Bur+21b] Vera von Burg et al. “Quantum computing enhanced computational catalysis”. In: *Physical Review Research* 3.3 (2021), p. 033055.

[BV04] Stephen P Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[CC21] Chris Cade and P Marcos Crichigno. “Complexity of supersymmetric systems and the cohomology problem”. In: *arXiv preprint arXiv:2107.00011* (2021).

[Cha+20] Christopher Chamberland et al. *Building a fault-tolerant quantum computer using concatenated cat codes*. 2020. arXiv: 2012.04108 [quant-ph].

[Che+21] Sitan Chen et al. “A hierarchy for replica quantum advantage”. In: *arXiv preprint arXiv:2111.05874* (2021).

[Che+22] Sitan Chen et al. “Exponential separations between learning with and without quantum memory”. In: *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2022, pp. 574–585.

[Che+23] Chi-Fang Chen et al. “Sparse random Hamiltonians are quantumly easy”. In: *arXiv preprint arXiv:2302.03394* (2023).

[Che+24a] Chi-Fang Chen et al. “A new approach to strong convergence”. In: *arXiv preprint arXiv:2405.16026* (2024).

[Che+24b] Chi-Fang Chen et al. “Efficient unitary designs and pseudorandom unitaries from permutations”. In: *arXiv preprint arXiv:2404.16751* (2024).

[Chi+18] Andrew M. Childs et al. “Toward the first quantum simulation with quantum speedup”. In: *Proceedings of the National Academy of Sciences* 115 (38 Sept. 2018), pp. 9456–9461. ISSN: 0027-8424. DOI: 10.1073/pnas.1801723115. URL: https://www.pnas.org/content/115/38/9456%20http://www.pnas.org/lookup/doi/10.1073/pnas.1801723115.

[Chi+22] Nai-Hui Chia et al. “Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning”. In: *Journal of the ACM* 69.5 (2022), pp. 1–72.

[Cho06] Timothy Y Chow. “You could have invented spectral sequences”. In: *Notices of the AMS* 53 (2006), pp. 15–19.

[CK22] Marcos Crichigno and Tamara Kohler. “Clique Homology is QMA1-hard”. In: *arXiv preprint arXiv:2209.11793* (2022).

[Cri20] P. Marcos Crichigno. “Supersymmetry and Quantum Computation”. In: (Nov. 2020). arXiv: 2011.01239 [quant-ph].

[CT17] Giuseppe Carleo and Matthias Troyer. “Solving the quantum many-body problem with artificial neural networks”. In: *Science* 355.6325 (2017), pp. 602–606.

[CW12] Andrew M Childs and Nathan Wiebe. “Hamiltonian simulation using linear combinations of unitary operations”. In: *arXiv preprint arXiv:1202.5822* (2012).

[CW20] Jordan Cotler and Frank Wilczek. “Quantum overlapping tomography”. In: *Physical Review Letters* 124.10 (2020), p. 100401.

[Dal+23a] Alexander M Dalzell et al. “Quantum algorithms: A survey of applications and end-to-end complexities”. In: *arXiv preprint arXiv:2310.03011* (2023).

[Dal+23b] Alexander M Dalzell et al. “Sparse random Hamiltonians are quantumly easy”. In: *arXiv preprint arXiv:2302.03394* (2023).

[DC91] B.R. Donald and D.R. Chang. “On the complexity of computing the homology type of a triangulation”. In: *[1991] Proceedings 32nd Annual Symposium of Foundations of Computer Science*. 1991, pp. 650–661. DOI: 10.1109/SFCS.1991.185432.

[Del73] Philippe Delsarte. "An algebraic approach to the association schemes of coding theory". In: *Philips Res. Rep. Suppl.* 10 (1973), pp. vi+–97.

[Der+21] Charles Derby et al. "Compact fermion to qubit mappings". In: *Physical Review B* 104.3 (2021), p. 035118.

[DRC17] Christian L Degen, Friedemann Reinhard, and Paola Cappellaro. "Quantum sensing". In: *Reviews of modern physics* 89.3 (2017), p. 035002.

[EH17] Lior Eldar and Aram W. Harrow. "Local Hamiltonians Whose Ground States Are Hard to Approximate". In: *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. 2017, pp. 427–438. DOI: 10.1109/FOCS.2017.46.

[EHF19] Tim J. Evans, Robin Harper, and Steven T. Flammia. "Scalable Bayesian Hamiltonian learning". In: *arXiv preprint arXiv:1912.07636* (2019).

[Erd78] Robert M Erdahl. "Representability". In: *International Journal of Quantum Chemistry* 13.6 (1978), pp. 697–718. URL: https://onlinelibrary.wiley.com/doi/10.1002/qua.560130603.

[Fac+19] Davide Facoetti et al. "Classical glasses, black holes, and strange quantum liquids". In: *Physical Review B* 100.20 (2019), p. 205108.

[Fey82] Richard P. Feynman. "Simulating physics with computers". In: *International Journal of Theoretical Physics* 21.6-7 (1982), pp. 467–488. DOI: 10.1007/BF02650179.

[FH13] Michael H Freedman and Matthew B Hastings. "Quantum systems on non-k-hyperfinite complexes: A generalization of classical statistical mechanics on expander graphs". In: *arXiv preprint arXiv:1301.1363* (2013).

[FL16a] Bill Fefferman and Cedric Lin. "Quantum Merlin Arthur with exponentially small gap". In: *arXiv preprint arXiv:1601.01975* (2016).

[FL16b] Bill Fefferman and Cedric Yen-Yu Lin. *A Complete Characterization of Unitary Quantum Space*. 2016. arXiv: 1604.01384 [quant-ph].

[For94] Robin Forman. "Hodge theory and spectral sequences". In: *Topology* 33.3 (1994), pp. 591–611.

[FTW19] Renjie Feng, Gang Tian, and Dongyi Wei. "Spectrum of SYK model". In: *Peking Mathematical Journal* 2 (2019), pp. 41–70.

[FTW20] Renjie Feng, Gang Tian, and Dongyi Wei. "Spectrum of SYK model III: large deviations and concentration of measures". In: *Random Matrices: Theory and Applications* 9.02 (2020), p. 2050001.

[FW70] JB French and SSM Wong. "Validity of random matrix theories for many-particle systems". In: *Physics Letters B* 33.7 (1970), pp. 449–452.

[Gam21] David Gamarnik. “The overlap gap property: A topological barrier to optimizing over random structures”. In: *Proceedings of the National Academy of Sciences* 118.41 (2021), e2108492118.

[Gav+07] Dmitry Gavinsky et al. “Exponential separations for one-way quantum communication complexity, with applications to cryptography”. In: *Proceedings of the 39th annual ACM Symposium on Theory of Computing (STOC 2007)*. 2007, pp. 516–525.

[GCD22] Casper Gyurik, Chris Cade, and Vedran Dunjko. “Towards quantum advantage via topological data analysis”. In: *Quantum* 6 (2022), p. 855.

[GGB16] Chad Giusti, Robert Ghrist, and Danielle S Bassett. “Two’s company, three (or more) is a simplex”. In: *Journal of computational neuroscience* 41.1 (2016), pp. 1–14.

[Gha+15] Sevag Gharibian et al. “Quantum hamiltonian complexity”. In: *Foundations and Trends® in Theoretical Computer Science* 10.3 (2015), pp. 159–282.

[GHG23] Carlos de Gois, Kiara Hansenne, and Otfried Gühne. “Uncertainty relations from graph theory”. In: *Physical Review A* 107.6 (2023), p. 062211.

[GI09] Daniel Gottesman and Sandy Irani. “The quantum and classical complexity of translationally invariant tiling and Hamiltonian problems”. In: *2009 50th Annual IEEE Symposium on Foundations of Computer Science*. IEEE. 2009, pp. 95–104.

[Gil+19] András Gilyén et al. “Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics”. In: *Proceedings of the 51st annual ACM SIGACT symposium on theory of computing*. 2019, pp. 193–204.

[GJC12] Daniel Gottesman, Thomas Jennewein, and Sarah Croke. “Longer-baseline telescopes using quantum repeaters”. In: *Physical review letters* 109.7 (2012), p. 070503. URL: https://arxiv.org/pdf/1107.2939.pdf.

[GJV18] Antonio M García-García, Yiyang Jia, and Jacobus JM Verbaarschot. “Exact moments of the Sachdev-Ye-Kitaev model up to order 1/N2”. In: *Journal of High Energy Physics* 2018.4 (2018), pp. 1–43.

[GK19] Sam Gunn and Niels Kornerup. “Review of a quantum algorithm for Betti numbers”. In: *arXiv preprint arXiv:1906.07673* (2019).

[GLT18] András Gilyén, Seth Lloyd, and Ewin Tang. “Quantum-inspired low-rank stochastic regression with logarithmic dependence on the dimension”. In: *arXiv preprint arXiv:1811.04909* (2018).

[GMV18] Guy Gur-Ari, Raghu Mahajan, and Abolhassan Vaezi. "Does the SYK model have a spin glass phase?" In: *Journal of High Energy Physics* 2018.11 (2018), pp. 1–20.

[GN16] David Gosset and Daniel Nagaj. "Quantum 3-SAT is QMA _1-complete". In: *SIAM Journal on Computing* 45.3 (2016), pp. 1080–1128.

[GNW21] David Gross, Sepehr Nezami, and Michael Walter. "Schur–Weyl duality for the Clifford group with applications: Property testing, a robust Hudson theorem, and de Finetti representations". In: *Communications in Mathematical Physics* 385.3 (2021), pp. 1325–1393.

[Gol02] Timothy E Goldberg. "Combinatorial Laplacians of simplicial complexes". In: *Senior Thesis, Bard College* 6 (2002).

[GPS00] Antoine Georges, Olivier Parcollet, and Subir Sachdev. "Mean field theory of a quantum Heisenberg spin glass". In: *Physical review letters* 85.4 (2000), p. 840.

[Gre+23] Sabee Grewal et al. "Improved Stabilizer Estimation via Bell Difference Sampling". In: *arXiv preprint arXiv:2304.13915* (2023).

[GS14] Anna Gundert and May Szedlák. "Higher dimensional Cheeger inequalities". In: *Proceedings of the thirtieth annual symposium on Computational geometry*. 2014, pp. 181–188.

[GS18] David Gosset and John Smolin. "A compressed classical description of quantum states". In: *arXiv preprint arXiv:1801.05721* (2018).

[GTC18] Yimin Ge, Jordi Tura, and J. Ignacio Cirac. *Faster ground state preparation and high-precision ground energy estimation with fewer qubits*. 2018. arXiv: 1712.03193 [quant-ph]. URL: https://arxiv.org/abs/1712.03193.

[Guo+16] Sheng Guo et al. "Electron Valence State Perturbation Theory Based on a Density Matrix Renormalization Group Reference Function, with Applications to the Chromium Dimer and a Trimer Model of Poly-Phenylenevinylene". In: *Journal of Chemical Theory and Computation* 12.4 (Apr. 2016), pp. 1583–1591. ISSN: 1549-9618. DOI: 10.1021/acs.jctc.5b01225.

[Gur17] Razvan Gurau. "Quenched equals annealed at leading order in the colored SYK model". In: *Europhysics letters* 119.3 (2017), p. 30003.

[GV16] Antonio M García-García and Jacobus JM Verbaarschot. "Spectral and thermodynamic properties of the Sachdev-Ye-Kitaev model". In: *Physical Review D* 94.12 (2016), p. 126010.

[GW95] Michel X Goemans and David P Williamson. "Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming". In: *Journal of the ACM (JACM)* 42.6 (1995), pp. 1115–1145.

[Gyá87] András Gyárfás. "Problems from the world surrounding perfect graphs". In: *Applicationes Mathematicae* 19.3-4 (1987), pp. 413–441. DOI: 10.4064/am-19-3-4-413-441.

[Haa+23] Jeongwan Haah et al. "Quantum Algorithm for Simulating Real Time Evolution of Lattice Hamiltonians". In: *SIAM Journal on Computing* 52 (6 Dec. 2023), FOCS18-250-FOCS18–284. ISSN: 0097-5397. DOI: 10.1137/18M1231511.

[Hal07] William Hall. "Compatibility of subsystem states and convex geometry". In: *Phys. Rev. A* 75 (3 Mar. 2007), p. 032102. DOI: 10.1103/PhysRevA.75.032102. URL: https://link.aps.org/doi/10.1103/PhysRevA.75.032102.

[Hay22] Ryu Hayakawa. "Quantum algorithm for persistent betti numbers and topological data analysis". In: *Quantum* 6 (2022), p. 873.

[HBK23] Tobias Haug, Kishor Bharti, and Dax Enshan Koh. "Pseudorandom unitaries are neither real nor sparse nor noise-robust". In: *arXiv preprint arXiv:2306.11677* (2023).

[He+17] Yong He et al. "Decompositions of n-qubit Toffoli Gates with Linear Circuit Complexity". In: *International Journal of Theoretical Physics* 56.7 (2017), pp. 2350–2361. DOI: 10.1007/s10773-017-3389-4. URL: https://doi.org/10.1007/s10773-017-3389-4.

[Her+23a] Yaroslav Herasymenko et al. "Fermionic Hamiltonians without trivial low-energy states". In: *arXiv preprint arXiv:2307.13730* (2023).

[Her+23b] Yaroslav Herasymenko et al. "Optimizing sparse fermionic hamiltonians". In: *Quantum* 7 (2023), p. 1081.

[HJ91] Roger A. Horn and Charles R. Johnson. "Singular value inequalities". In: *Topics in Matrix Analysis*. Cambridge University Press, 1991, pp. 134–238.

[HKP20] Hsin-Yuan Huang, Richard Kueng, and John Preskill. "Predicting many properties of a quantum system from very few measurements". In: *Nature Physics* 16.10 (2020), pp. 1050–1057.

[HKP21] Hsin-Yuan Huang, Richard Kueng, and John Preskill. "Information-theoretic bounds on quantum advantage in machine learning". In: *Physical Review Letters* 126.19 (2021), p. 190505.

[HLK23] Tobias Haug, Soovin Lee, and MS Kim. "Efficient stabilizer entropies for quantum computers". In: *arXiv preprint arXiv:2305.19152* (2023).

[HO21] Matthew B. Hastings and Ryan O'Donnell. *Optimizing Strongly Interacting Fermionic Hamiltonians*. 2021. DOI: 10.48550/ARXIV.2110.10701. URL: https://arxiv.org/abs/2110.10701.

[HO22] Matthew B Hastings and Ryan O'Donnell. "Optimizing strongly interacting fermionic Hamiltonians". In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. 2022, pp. 776–789.

[Hol+17] Christopher L Holloway et al. "Atom-based RF electric field metrology: from self-calibrated measurements to subwavelength and near-field imaging". In: *IEEE Transactions on Electromagnetic Compatibility* 59.2 (2017), pp. 717–728. URL: https://ieeexplore.ieee.org/abstract/document/7812705.

[HTS21a] Arijit Haldar, Omid Tavakol, and Thomas Scaffidi. "Variational wave functions for Sachdev-Ye-Kitaev models". In: *Physical Review Research* 3.2 (2021), p. 023020.

[HTS21b] Arijit Haldar, Omid Tavakol, and Thomas Scaffidi. "Variational wave functions for Sachdev-Ye-Kitaev models". In: *Phys. Rev. Res.* 3 (2 Apr. 2021), p. 023020. DOI: 10.1103/PhysRevResearch.3.023020. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.3.023020.

[Hua+22] Hsin-Yuan Huang et al. "Quantum advantage in learning from experiments". In: *Science* 376.6598 (2022), pp. 1182–1186.

[Hug+22a] William J Huggins et al. "Nearly optimal quantum algorithm for estimating multiple expectation values". In: *Physical Review Letters* 129.24 (2022), p. 240501.

[Hug+22b] William J. Huggins et al. "Unbiasing fermionic quantum Monte Carlo with a quantum computer". In: *Nature* 603.7901 (Mar. 2022), pp. 416–420. ISSN: 0028-0836. DOI: 10.1038/s41586-021-04351-z.

[JGM19] Andrew Jena, Scott Genin, and Michele Mosca. "Pauli partitioning with respect to gate sets". In: *arXiv preprint arXiv:1907.07859* (2019).

[Jia+20a] Zhang Jiang et al. "Optimal fermion-to-qubit mapping via ternary trees with applications to reduced quantum states learning". In: *Quantum* 4 (2020), p. 276.

[Jia+20b] Zhang Jiang et al. "Optimal fermion-to-qubit mapping via ternary trees with applications to reduced quantum states learning". In: *Quantum* 4 (June 2020), p. 276. ISSN: 2521-327X. DOI: 10.22331/q-2020-06-04-276. URL: https://doi.org/10.22331/q-2020-06-04-276.

[Jor+12] Stephen P. Jordan et al. *Achieving perfect completeness in classical-witness quantum Merlin-Arthur proof systems*. 2012. arXiv: 1111.5306 [quant-ph].

[JW28] P. Jordan and E. Wigner. "Über das Paulische Äquivalenzverbot". In: *Zeitschrift für Physik* 47.9 (Sept. 1928), pp. 631–651. ISSN: 0044-3328. DOI: 10.1007/BF01331938. URL: https://doi.org/10.1007/BF01331938.

[Kim95] Jeong Han Kim. "The Ramsey number $R(3,t)$ has order of magnitude $t^2/\log t$". In: *Random Structures & Algorithms* 7.3 (1995), pp. 173–207.

[Kin+24] Robbie King et al. "Strongly interacting fermions are non-trivial yet non-glassy". In: *arXiv preprint arXiv:2408.15699* (2024). URL: https://arxiv.org/abs/2408.15699.

[Kin+25] Robbie King et al. "Triply efficient shadow tomography". In: *Proceedings of the 2025 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2025. DOI: 10.1137/1.9781611978322.27.

[Kit+02] Alexei Yu Kitaev et al. *Classical and quantum computation*. 47. American Mathematical Soc., 2002.

[Kit15a] Alexei Kitaev. *A Simple Model of Quantum Holography*. Talks at Kavli Institute for Theoretical Physics. Santa Barbara, U.S.A., 2015.

[Kit15b] Alexei Kitaev. *Hidden Correlations in the Hawking Radiation and Thermal Noise*. Talk at Kavli Institute for Theoretical Physics. Santa Barbara, U.S.A., 2015.

[Kit95] A Yu Kitaev. "Quantum measurements and the Abelian stabilizer problem". In: *arXiv preprint quant-ph/9511026* (1995).

[KK24] Robbie King and Tamara Kohler. "Gapped Clique Homology on Weighted Graphs is QMA 1-Hard and Contained in QMA". In: *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2024, pp. 493–504. DOI: 10.1109/FOCS61266.2024.00039.

[KKR06] Julia Kempe, Alexei Kitaev, and Oded Regev. "The complexity of the local Hamiltonian problem". In: *Siam journal on computing* 35.5 (2006), pp. 1070–1097.

[Knu93] Donald E Knuth. "The sandwich theorem". In: *arXiv preprint math/9312214* (1993).

[KO23] Robin Kothari and Ryan O'Donnell. "Mean estimation when you have the source code; or, quantum Monte Carlo methods". In: *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2023, pp. 1186–1215.

[KOS07] Emanuel Knill, Gerardo Ortiz, and Rolando D Somma. "Optimal quantum measurements of expectation values of observables". In: *Physical Review A—Atomic, Molecular, and Optical Physics* 75.1 (2007), p. 012328.

[KS18] Alexei Kitaev and S Josephine Suh. "The soft mode in the Sachdev-Ye-Kitaev model and its gravity dual". In: *Journal of High Energy Physics* 2018.5 (2018), pp. 1–68.

[KWM24] Robbie King, Kianna Wan, and Jarrod R. McClean. "Exponential Learning Advantages with Conjugate States and Minimal Quantum Memory". In: *PRX Quantum* 5 (4 Oct. 2024), p. 040301. DOI: 10.1103/PRXQuantum.5.040301.

[LC17a] Guang Hao Low and Isaac L Chuang. "Hamiltonian simulation by uniform spectral amplification". In: *arXiv preprint arXiv:1707.05391* (2017).

[LC17b] Guang Hao Low and Isaac L. Chuang. "Optimal Hamiltonian Simulation by Quantum Signal Processing". In: *Phys. Rev. Lett.* 118 (1 2017), p. 010501. DOI: 10.1103/PhysRevLett.118.010501. URL: https://link.aps.org/doi/10.1103/PhysRevLett.118.010501.

[LC19] Guang Hao Low and Isaac L Chuang. "Hamiltonian simulation by qubitization". In: *Quantum* 3 (2019), p. 163. DOI: 10.22331/q-2019-07-12-163.

[LCV07] Yi-Kai Liu, Matthias Christandl, and Frank Verstraete. "Quantum computational complexity of the N-representability problem: QMA complete". In: *Physical review letters* 98.11 (2007), p. 110503.

[Lee+21a] Joonho Lee et al. "Even More Efficient Quantum Computations of Chemistry Through Tensor Hypercontraction". In: *PRX Quantum* 2.3 (July 2021), p. 030305. ISSN: 2691-3399. DOI: 10.1103/prxquantum.2.030305. URL: http://dx.doi.org/10.1103/PRXQuantum.2.030305.

[Lee+21b] Joonho Lee et al. "Even more efficient quantum computations of chemistry through tensor hypercontraction". In: *PRX Quantum* 2.3 (2021), p. 030305.

[Lee+22] Seunghoon Lee et al. *Is there evidence for exponential quantum advantage in quantum chemistry?* 2022. DOI: 10.48550/ARXIV.2208.02199. URL: https://arxiv.org/abs/2208.02199.

[LGZ16] Seth Lloyd, Silvano Garnerone, and Paolo Zanardi. "Quantum algorithms for topological and geometric analysis of data". In: *Nature communications* 7.1 (2016), p. 10138.

[LI23] Ignacio Loaiza and Artur F Izmaylov. "Block-invariant symmetry shift: Preprocessing technique for second-quantized hamiltonians to improve their decompositions to linear combination of unitaries". In: *Journal of Chemical Theory and Computation* 19.22 (2023), pp. 8201–8209. URL: https://pubs.acs.org/doi/10.1021/acs.jctc.3c00912.

[Lin24] William Linz. "L-systems and the Lovász number". In: *arXiv preprint arXiv:2402.05818* (2024).

[Llo96] Seth Lloyd. "Universal Quantum Simulators". In: *Science* 273.5278 (1996), pp. 1073–1078. ISSN: 0036-8075. DOI: 10.1126/science.273.5278.1073. URL: https://science.sciencemag.org/content/273/5278/1073.

[Lov75] László Lovász. "On the ratio of optimal integral and fractional covers". In: *Discrete mathematics* 13.4 (1975), pp. 383–390.

[Low+25] Guang Hao Low et al. "Fast quantum simulation of electronic structure by spectrum amplification". In: *arXiv preprint arXiv:2502.15882* (2025).

[Low22] Guang Hao Low. "Classical shadows of fermions with particle number symmetry". In: *arXiv preprint arXiv:2208.08964* (2022).

[LP17] David A Levin and Yuval Peres. *Markov chains and mixing times*. Vol. 107. American Mathematical Soc., 2017.

[LS24] Guang Hao Low and Yuan Su. *Quantum linear system algorithm with optimal queries to initial state preparation*. 2024. arXiv: 2410.18178 [quant-ph]. URL: https://arxiv.org/abs/2410.18178.

[LT20] Lin Lin and Yu Tong. "Near-optimal ground state preparation". In: *Quantum* 4 (Dec. 2020), p. 372. ISSN: 2521-327X. DOI: 10.22331/q-2020-12-14-372. URL: https://doi.org/10.22331/q-2020-12-14-372.

[Lub18] Alexander Lubotzky. "High dimensional expanders". In: *Proceedings of the international congress of mathematicians: Rio de Janeiro 2018*. World Scientific. 2018, pp. 705–730.

[MA23] Jisho Miyazaki and Seiseki Akibue. "Non-locality of conjugation symmetry: characterization and examples in quantum network sensing". In: *arXiv preprint arXiv:2309.12523* (2023).

[Mar+21] John M Martyn et al. "Grand unification of quantum algorithms". In: *PRX quantum* 2.4 (2021), p. 040203.

[Maz06] David A Mazziotti. "Quantum chemistry without wave functions: Two-electron reduced density matrices". In: *Accounts of chemical research* 39.3 (2006), pp. 207–215.

[McA+20] Sam McArdle et al. "Quantum computational chemistry". In: *Rev. Mod. Phys.* 92 (1 2020), p. 015003. DOI: 10.1103/RevModPhys.92.015003. URL: https://link.aps.org/doi/10.1103/RevModPhys.92.015003.

[McC+17] Jarrod R. McClean et al. “Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states”. In: *Phys. Rev. A* 95 (4 Apr. 2017), p. 042308. DOI: 10.1103/PhysRevA.95.042308. URL: https://link.aps.org/doi/10.1103/PhysRevA.95.042308.

[McC01] John McCleary. *A user's guide to spectral sequences*. 58. Cambridge University Press, 2001.

[MG92] Jayadev Misra and David Gries. “A constructive proof of Vizing's theorem”. In: *Information Processing Letters* 41.3 (1992), pp. 131–133.

[MGB22] Sam McArdle, András Gilyén, and Mario Berta. “A streamlined quantum algorithm for topological data analysis with exponentially fewer qubits”. In: *arXiv preprint arXiv:2209.12887* (2022).

[MM22] Jisho Miyazaki and Keiji Matsumoto. “Imaginarity-free quantum multiparameter estimation”. In: *Quantum* 6 (2022), p. 665.

[Mon17] Ashley Montanaro. “Learning stabilizer states by Bell sampling”. In: *arXiv preprint arXiv:1707.04012* (2017).

[MS16] Juan Maldacena and Douglas Stanford. “Remarks on the Sachdev-Ye-Kitaev model”. In: *Physical Review D* 94.10 (2016), p. 106002.

[MW23] Nikhil S Mande and Ronald de Wolf. “Tight bounds for quantum phase estimation and related problems”. In: *arXiv preprint arXiv:2305.04908* (2023).

[Nak+01] Maho Nakata et al. “Variational calculations of fermion second-order reduced density matrices by semidefinite programming algorithm”. In: *The Journal of Chemical Physics* 114.19 (2001), pp. 8282–8292. URL: https://pubs.aip.org/aip/jcp/article-abstract/114/19/8282/460915/Variational-calculations-of-fermion-second-order?redirectedFrom=fulltext.

[NI21] Yusuke Nomura and Masatoshi Imada. “Dirac-type nodal spin liquid revealed by refined quantum many-body solver using neural-network wave function, correlation ratio, and level spectroscopy”. In: *Physical Review X* 11.3 (2021), p. 031034.

[NPA08] Miguel Navascués, Stefano Pironio, and Antonio Acín. “A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations”. In: *New Journal of Physics* 10.7 (2008), p. 073013.

[OM99] Andreas Osterwalder and Frédéric Merkt. “Using high Rydberg states as electric field sensors”. In: *Physical review letters* 82.9 (1999), p. 1831. URL: https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.82.1831.

[Osb12] Tobias J Osborne. "Hamiltonian complexity". In: *Reports on progress in physics* 75.2 (2012), p. 022001.

[OW16] Ryan O'Donnell and John Wright. "Efficient quantum tomography". In: *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. 2016, pp. 899–912.

[Par79] Giorgio Parisi. "Infinite number of order parameters for spin-glasses". In: *Physical Review Letters* 43.23 (1979), p. 1754.

[Pet+14] Giovanni Petri et al. "Homological scaffolds of brain functional networks". In: *Journal of The Royal Society Interface* 11.101 (2014), p. 20140873.

[PNA10] Stefano Pironio, Miguel Navascués, and Antonio Acin. "Convergent relaxations of polynomial optimization problems with noncommuting variables". In: *SIAM Journal on Optimization* 20.5 (2010), pp. 2157–2180.

[PRT16] Ori Parzanchevski, Ron Rosenthal, and Ran J Tessler. "Isoperimetric inequalities in simplicial complexes". In: *Combinatorica* 36.2 (2016), pp. 195–227.

[PRW06] Janez Povh, Franz Rendl, and Angelika Wiegele. "A boundary point method to solve semidefinite programs". In: *Computing* 78 (2006), pp. 277–286. URL: https://link.springer.com/article/10.1007/s00607-006-0182-2.

[PZC23] Linqing Peng, Xing Zhang, and Garnet Kin-Lic Chan. "Fermionic Reduced Density Low-Rank Matrix Completion, Noise Filtering, and Measurement Reduction in Quantum Simulations". In: *Journal of Chemical Theory and Computation* 19.24 (Dec. 2023), pp. 9151–9160. ISSN: 1549-9618. DOI: 10.1021/acs.jctc.3c00851.

[Raz99] Ran Raz. "Exponential separation of quantum and classical communication complexity". In: *Proceedings of the thirty-first annual ACM symposium on Theory of computing*. 1999, pp. 358–367.

[Rei+17] Michael Reimann et al. "Cliques of Neurons Bound into Cavities Provide a Missing Link between Structure and Function". In: *Frontiers in Computational Neuroscience* 11 (June 2017), pp. 1–16. DOI: 10.3389/fncom.2017.00048. URL: https://hal.archives-ouvertes.fr/hal-01706964.

[RH23] Philippe Rigollet and Jan-Christian Hütter. "High-dimensional statistics". In: *arXiv preprint arXiv:2310.19244* (2023).

[Rub+23] Nicholas C Rubin et al. "Fault-tolerant quantum simulation of materials using Bloch orbitals". In: *PRX Quantum* 4.4 (2023), p. 040303.

[SB13] Rolando D Somma and Sergio Boixo. "Spectral gap amplification". In: *SIAM Journal on Computing* 42.2 (2013), pp. 593–610.

[Sch11] Ulrich Schollwöck. “The density-matrix renormalization group in the age of matrix product states”. In: *Annals of physics* 326.1 (2011), pp. 96–192.

[Sed+12] Jonathon A Sedlacek et al. “Microwave electrometry with Rydberg atoms in a vapour cell using bright atomic resonances”. In: *Nature physics* 8.11 (2012), pp. 819–824. URL: https://www.nature.com/articles/nphys2423.

[SH20] Markus Schmitt and Markus Heyl. “Quantum many-body dynamics in two dimensions with artificial neural networks”. In: *Physical Review Letters* 125.10 (2020), p. 100503.

[Sha+17] Sandeep Sharma et al. “Combining Internally Contracted States and Matrix Product States To Perform Multireference Perturbation Theory”. In: *Journal of Chemical Theory and Computation* 13.2 (2017), pp. 488–498. DOI: 10.1021/acs.jctc.6b00898. URL: https://doi.org/10.1021/acs.jctc.6b00898.

[Sha+20] Or Sharir et al. “Deep autoregressive models for the efficient variational simulation of many-body quantum systems”. In: *Physical review letters* 124.2 (2020), p. 020503.

[Shi03] Yaoyun Shi. “Both Toffoli and controlled-NOT need little help to do universal quantum computing”. In: *Quantum Information & Computation* 3.1 (2003), pp. 84–92.

[Sim+21] Matthew T Simons et al. “Rydberg atom-based sensors for radio-frequency electric field metrology, sensing, and communications”. In: *Measurement: Sensors* 18 (2021), p. 100273. URL: https://www.sciencedirect.com/science/article/pii/S2665917421002361.

[Sim+24] Sophia Simon et al. *Amplified Amplitude Estimation: Exploiting Prior Knowledge to Improve Estimates of Expectation Values*. 2024. arXiv: 2402.14791 [quant-ph]. URL: https://arxiv.org/abs/2402.14791.

[SKM14] John Steenbergen, Caroline Klivans, and Sayan Mukherjee. “A Cheeger-type inequality on simplicial complexes”. In: *Advances in Applied Mathematics* 56 (2014), pp. 56–77.

[SL22] Alexander Schmidhuber and Seth Lloyd. “Complexity-Theoretic Limitations on Quantum Algorithms for Topological Data Analysis”. In: *arXiv preprint arXiv:2209.14286* (2022).

[SO12] Attila Szabo and Neil S Ostlund. *Modern quantum chemistry: introduction to advanced electronic structure theory*. Courier Corporation, 2012.

[SR19] Ingo Schiermeyer and Bert Randerath. "Polynomial χ-binding functions and forbidden induced subgraphs: a survey". In: *Graphs and Combinatorics* 35.1 (2019), pp. 1–31.

[SRL12] Jacob T. Seeley, Martin J. Richard, and Peter J. Love. "The Bravyi-Kitaev transformation for quantum computation of electronic structure". In: *The Journal of Chemical Physics* 137.22 (Dec. 2012), p. 224109. ISSN: 0021-9606. DOI: 10.1063/1.4768229.

[SS20] Alex Scott and Paul Seymour. *A survey of χ-boundedness*. 2020. arXiv: 1812.07500 [math.CO].

[SSC22] Or Sharir, Amnon Shashua, and Giuseppe Carleo. "Neural tensor contractions and the expressive power of deep neural quantum states". In: *Physical Review B* 106.20 (2022), p. 205136.

[SU11] E.R. Scheinerman and D.H. Ullman. *Fractional Graph Theory: A Rational Approach to the Theory of Graphs*. Dover books on mathematics. Dover Publications, 2011. ISBN: 9780486485935. URL: https://books.google.com/books?id=zzFxD8kPWigC.

[SW24] Brian Swingle and Mike Winer. "Bosonic model of quantum holography". In: *Physical Review B* 109.9 (2024), p. 094206.

[SY93] Subir Sachdev and Jinwu Ye. "Gapless spin-fluid ground state in a random quantum Heisenberg magnet". In: *Physical review letters* 70.21 (1993), p. 3339.

[Tak+20] Tyler Takeshita et al. "Increasing the Representation Accuracy of Quantum Simulations of Chemistry without Extra Quantum Resources". In: *Phys. Rev. X* 10 (1 Jan. 2020), p. 011004. DOI: 10.1103/PhysRevX.10.011004. URL: https://link.aps.org/doi/10.1103/PhysRevX.10.011004.

[Tal00] Michel Talagrand. "Rigorous low-temperature results for the mean field p-spins interaction model". In: *Probability theory and related fields* 117 (2000), pp. 303–360.

[Tan19] Ewin Tang. "A quantum-inspired classical algorithm for recommendation systems". In: *Proceedings of the 51st annual ACM SIGACT symposium on theory of computing*. 2019, pp. 217–228.

[Uba+21] Shashanka Ubaru et al. "Quantum topological data analysis with linear depth and exponential speedup". In: *arXiv preprint arXiv:2108.02811* (2021).

[Val45] F. A. Valentine. "A Lipschitz Condition Preserving Extension for a Vector Function". In: *American Journal of Mathematics* 67.1 (1945), pp. 83–93. ISSN: 00029327, 10806377. DOI: 10.2307/2371917. (Visited on 01/21/2025).

[Ver18] Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Vol. 47. Cambridge University Press, 2018.

[Vla19] Alexander Yu Vlasov. "Clifford algebras, Spin groups and qubit trees". In: *arXiv preprint arXiv:1904.09912* (2019).

[VYI20] Vladyslav Verteletskyi, Tzu-Ching Yen, and Artur F. Izmaylov. "Measurement optimization in the variational quantum eigensolver using a minimum clique cover". In: *The Journal of Chemical Physics* 152.12 (Mar. 2020). ISSN: 0021-9606. DOI: 10.1063/1.5141458.

[Wad+17] Christopher G Wade et al. "Real-time near-field terahertz imaging with atomic optical fluorescence". In: *Nature Photonics* 11.1 (2017), pp. 40–43. URL: https://www.nature.com/articles/nphoton.2016.214.

[Wai19a] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*. Vol. 48. Cambridge university press, 2019.

[Wai19b] Martin J. Wainwright. "Basic tail and concentration bounds". In: *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019, pp. 21–57. DOI: 10.1017/9781108627771.002.

[Wan+22] Kianna Wan et al. "Matchgate shadows for fermionic quantum simulation". In: *arXiv preprint arXiv:2207.13723* (2022).

[Was18] Larry Wasserman. "Topological data analysis". In: *Annual Review of Statistics and Its Application* 5 (2018), pp. 501–532.

[WCL23] Ya-Dong Wu, Giulio Chiribella, and Nana Liu. "Quantum-enhanced learning of continuous-variable quantum states". In: *arXiv preprint arXiv:2303.05097* (2023).

[Wie+10] Nathan Wiebe et al. "Higher order decompositions of ordered operator exponentials". In: *Journal of Physics A: Mathematical and Theoretical* 43.6 (2010), p. 065203.

[Wil67] R. M. Wilcox. "Exponential Operators and Parameter Differentiation in Quantum Physics". In: *Journal of Mathematical Physics* 8.4 (Apr. 1967), pp. 962–982. ISSN: 0022-2488. DOI: 10.1063/1.1705306.

[Wit15] Peter Wittek. "Algorithm 950: Ncpol2sdpa—sparse semidefinite programming relaxations for polynomial optimization problems of noncommuting variables". In: *ACM Transactions on Mathematical Software (TOMS)* 41.3 (2015), pp. 1–12. URL: https://dl.acm.org/doi/10.1145/2699464.

[Wit82] Edward Witten. "Supersymmetry and Morse theory". In: *Journal of differential geometry* 17.4 (1982), pp. 661–692.

[Wit89] Edward Witten. “Quantum field theory and the Jones polynomial”. In: *Communications in Mathematical Physics* 121.3 (1989), pp. 351–399.

[Wou+16] Sebastian Wouters et al. “A Practical Guide to Density Matrix Embedding Theory in Quantum Chemistry”. In: *Journal of Chemical Theory and Computation* 12.6 (2016), pp. 2706–2719. ISSN: 1549-9618. DOI: 10.1021/acs.jctc.6b00316. URL: http://pubs.acs.org/doi/abs/10.1021/acs.jctc.6b00316.

[XSW23] Zhen-Peng Xu, René Schwonnek, and Andreas Winter. “Bounding the joint numerical range of Pauli strings by graph parameters”. In: *arXiv preprint arXiv:2308.00753* (2023).

[YAL97] Bin Yu, Fano Assouad, and Lucien Le Cam. “Festschrift for lucien le cam”. In: *Assouad, Fano, and Le Cam*. Vol. 423. Springer, 1997, p. 435.

[Yos+22] Nobuyuki Yoshioka et al. “Generalized Quantum Subspace Expansion”. In: *Physical Review Letters* 129.2 (July 2022), p. 020502. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.129.020502.

[ZRM21] Andrew Zhao, Nicholas C Rubin, and Akimasa Miyake. “Fermionic partial tomography via classical shadows”. In: *Physical Review Letters* 127.11 (2021), p. 110504.

[ZS24] Alexander Zlokapa and Rolando D Somma. “Hamiltonian simulation for low-energy states with optimal time dependence”. In: *Quantum* 8 (2024), p. 1449.