

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Algorithms and Spectral Gaps for Quantum Problems

Permalink

<https://escholarship.org/uc/item/9q51g4n7>

ISBN

9798189271823

Author

Stier, Zachary

Publication Date

2026-05-01

Peer reviewed|Thesis/dissertation

Algorithms and Spectral Gaps for Quantum Problems

by

Zachary Stier

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Mathematics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Nikhil Srivastava, Chair

Professor Shirshendu Ganguly

Professor Lin Lin

Spring 2026

Algorithms and Spectral Gaps for Quantum Problems

Copyright 2026
by
Zachary Stier

Abstract

Algorithms and Spectral Gaps for Quantum Problems

by

Zachary Stier

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor Nikhil Srivastava, Chair

This dissertation is concerned with a variety of problems in quantum computing and the spectral theory of quantum-mechanical channels and Hamiltonians.

We begin with the problem of quantum gate synthesis. The challenge is to find a fixed, finite set S of quantum gates such that, given a quantum circuit component U acting on n qubits and target precision ε , there is a short, efficiently computable composition of gates from S that approximates U to precision ε . This problem has already led to a rich literature in the case $n = 1$ with deep connections to modern number theory. In this dissertation we present the state-of-the-art algorithm for a single qubit, achieving length $\frac{7}{3} \log \frac{1}{\varepsilon^3}$. The techniques used include the lattice structure of certain dense subgroups of $\text{PU}(2)$ and Lenstra's algorithm for integer programming in convex domains.

The next topic considered is scientific computing algorithms intended to run on a quantum computer. Specifically, we study a framework for computing functions of normal matrices using a quantum computer, with the objective of controlling the ancilla overhead, probability of success, and end precision. This arises as a special case of the *multivariate quantum eigenvalue transform*, for computing function of multiple commuting Hermitian matrices. The techniques applied include approximation of regular functions by low-degree polynomials and modern quantum primitives like quantum signal processing and linear combinations of unitaries.

The third component is a spectral gap comparison result for Davies generators, a class of Lindbladians used for Gibbs state preparation and which are a well-studied example of quantum Markov dynamics, and which contain a canonical classical Markov chain. The main result shows that if the spectrum of the Hamiltonian driving the dynamics is free of long arithmetic progressions, then the spectral gaps of the quantum and classical Markov chains are comparable. Thus, classical techniques may be leveraged to

control the behavior of the quantum dynamics. The proof is based on transferring any “off-diagonal” eigenvector of the Davies generator to an observable which commutes with the Hamiltonian and has a comparable Lindbladian Rayleigh quotient.

Finally, we study the infimum of the spectrum, or *ground state energy*, of a discrete Schrödinger operator on $\theta\mathbb{Z}^n$ parameterized by a potential $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ and a small frequency parameter θ . We relate this ground state energy to that of a corresponding continuous semiclassical Schrödinger operator on $\mathbb{R}^n$ with parameter θ , arising from the same choice of potential, and show that these quantities are equal up to vanishing factors depending on the regularity of V . The proofs are elementary and rely on sampling and interpolation to map low-energy functions for the discrete operator on $\theta\mathbb{Z}^n$ to low-energy functions for the continuous operator on $\mathbb{R}^n$, and vice versa.

Coauthors on these works were J. Basso, T. R. Blackman, Y. Borns-Weil, I. Detherage, S. Ganguly, N. Srivastava, T. Saffat, A. Sinclair, and T.-D. Vuong.

To Howard Ballen

Contents

Contents	ii
List of Figures	iv
1 Introduction	1
1.1 Background and context	1
1.2 Themes	17
1.3 Publication history	19
1.4 Selected preliminaries	19
2 Fast navigation with icosahedral golden gates	22
2.1 Super golden gates	24
2.2 Icosahedral super golden gates	28
2.3 Nearby elements in $\text{PU}(2)$, and factoring in Γ	29
2.4 Algorithm for short paths to diagonal elements	31
2.5 Algorithm for short paths	33
2.6 Examples	35
2.7 Upper-bounding fault-tolerant resources	37
3 A quantum algorithm for functions of multiple commuting Hermitian matrices	38
3.1 Preliminaries	42
3.2 QET for normal matrices	45
3.3 On nonlinear transformation of complex amplitudes	49
3.4 MQET for several commuting Hermitian matrices	50
4 On quantum to classical comparison for Davies generators	54
4.1 Preliminaries	65
4.2 Proof of main comparison theorems	72
4.3 Relating quantum to classical spectral gap	81
4.4 Absence of long proper arithmetic progressions	85

5	On the ground state energies of discrete and semiclassical Schrödinger operators	92
5.1	Preliminaries	97
5.2	Constructing a discrete p.g.s. via sampling	98
5.3	Constructing a continuous p.g.s. via multilinear interpolation	102
5.4	A bound for any potential	109
A	Appendix to Chapter 2 (elements in the number ring)	114
B	Appendix to Chapter 3	119
B.1	Selecting polynomial approximations	119
B.2	Exponentiation of normal matrices	122
B.3	Optimizing the number of block-encoding instances	122
B.4	Chebyshev expansion of Chebyshev interpolants	124
B.5	Nonlinear transformation of complex amplitudes primitives	125
C	Appendix to Chapter 4 (deferred proofs)	127
D	Appendix to Chapter 5 (deferred proofs)	132
	Bibliography	135

List of Figures

1.1	A Deutsch–Jozsa circuit.	5
1.2	Composition of quantum gates corresponds to multiplication	6
1.3	Figure 1.1, rewritten more compactly	21
1.4	A SWAP gate	21
3.1	0- and 1-controlled U -gates.	42
3.2	Numerically-controlled gates	42
3.3	Gates that skip registers	43
3.4	The block-encoding multiplication circuit	44
3.5	Normal QET circuit	46
3.6	MQET circuit	51
5.1	A demonstration of $I_{k,i}^\eta$ and $J_{k',i}^\eta$ in two dimensions	99
5.2	A demonstration of unions of $J_{k',1}^\eta$ in two dimensions	99
B.1	The circuits for Guo, Mitarai, and Fujii’s W and W'	125
B.2	The circuits for Guo, Mitarai, and Fujii’s G and G'	126
B.3	The circuits for Guo, Mitarai, and Fujii’s $\tilde{G}$ and $\tilde{G}'$	126

Acknowledgments

First, I would like to thank my advisor, Nikhil Srivastava. When I arrived at Berkeley I had no idea what I wanted to study, but whenever I'd leave his Math 224A office hours I was always deeply excited about whatever we had just discussed. That feeling hasn't gone away, and over these six years I am grateful to Nikhil for his time, patience, encouragement, and insights. I feel fortunate to have entered the vibrant and collaborative environment that he created and fostered at Berkeley and Simons, and to have learned so much about mathematics itself and the act and joy of doing research mathematics.

Thanks also to Shirshendu Ganguly and Lin Lin. I am grateful that they were willing to serve on my committee; it was a pleasure to collaborate with Shirshendu on what would become Chapter 4, and Lin's guidance and encouragement in his Math 275 led to Chapter 3. I am indebted to Peter Sarnak, for his continued engagement with my mathematical growth even past my time as his undergraduate advisee, and for connecting me with Terrence, leading to Chapter 2.

Next, I am grateful to the coauthors whose joint work forms the basis of this thesis: João, Terrence, Yonah, Izzy, Shirshendu, Nikhil, Tahsin, Alistair, and June. I learned an enormous amount from all of them. I owe similar thanks to my other coauthors, as well as collaborators with whom our work did not lead to a complete publication, whose time, patience, insights, and company were deeply rewarding. For the preparation of this thesis in particular, many thanks to Kris, Izzy, Lin, and Nikhil for many helpful comments, and referential spectral thanks to Vilas and Thomas.

I am grateful to the NSF GRFP for funding my research, as well as the Simons Institute, MSRI, PCMI, UCSB, the AMS, Université Paul Sabatier, Oberwolfach, the Heidelberg Laureate Forum, IPAM, and Kyoto University for hosting pleasant and fulfilling programs.

I am very lucky to have been a part of so many supportive groups in my time in Berkeley: the Math Department; Nikhil's other advisees João, Izzy, Jorge, Theo, and Rikhav; the Discrete Analysis Seminar; Mount Tamalpais College; Good In Theory; and Aunties & Uncs, intramural champions. Throughout graduate school I was also beyond fortunate to have so many wonderful friends in the Bay Area to share in the journey: my cohort; my officemates Carlos, Michael, Catherine, and Abi; and people from Berkeley at large, Stanford, and San Francisco — you know who you are and I am grateful to all of you. Then there are so many great friends from further afield like BCA, Mathcamp, Princeton, and more who I've not been able to see in person nearly as much as I'd like. Extra special thanks to my roommates at the Anthill and Tape Haus — Carlos, Connor, and Rose — for helping make Berkeley really feel like home.

To Katherine I owe too much to fit in this space, but every day with you is full of beauty, curiosity, laughter, and love, and I cannot wait for what the future has in store.

And last, but far from least, is my family: Nana, my parents Stacy and Mitchell, my brother Ken, and my dogs Tashi and Iko; to their love and support I owe everything.

Chapter 1

Introduction

This thesis is primarily based on the content of the following four papers [BS'23; BSS'23; B⁺'25; DSS'24]. Collectively, these works focus on various facets of the theory of quantum computing: classical algorithms to inform hardware design choices, quantum algorithms for scientific computing, and spectral analysis of quantum mechanical channels and Hamiltonians. In this introductory chapter, I proceed through the chapters in an order moving up the computing hierarchy, and lay the groundwork for the later results aimed at a general mathematical audience. Then, I will provide some specialized technical background.

1.1 Background and context

Quantum computing is an exciting, fast-evolving paradigm of computation with the potential to both extend the capabilities of classical computers and more faithfully capture and manipulate the nature of the universe in which we reside. There are continuing developments in quantum computing, both at the theoretical and conceptual levels, and at the stage of physical implementation to begin to realize the potential gains to solving computational problems by harnessing the intrinsic advantages of the model and technology. In particular, over the last four decades, the theoretical study of quantum computers has blossomed substantially. Quantum computing was initially conceived in the early 1980s. Yuri Manin [M'80] and Paul Benioff [B'80] made important early proposals, but none was more influential than that of Richard Feynman [F'82; ACM'24], who in a 1981 lecture propounded quantum computing as a sensible objective for simulating physical processes. Feynman's argument centered on the insufficiency of classical computing, since computing with quantum objects requires exponential complexity, and since by Bell's inequality [B'64], there cannot be "local hidden variables."¹ However, at the more visceral level, Feynman argues that the problem must be solved

¹Feynman actually also provides a new derivation of Bell's inequality.

because nature isn't classical, dammit, and if you want to make a simulation of nature, you'd better make it quantum mechanical, and by golly it's a wonderful problem, because it doesn't look so easy.

Feynman did not conceive of quantum computers for strictly algorithmic purposes. This insight can be attributed to Benioff [B'80] and David Deutsch [D'85], who defined the *quantum Turing machine*. Turing machines are the foundational model of (classical) computers and envision computing as a tape head with the ability to write or erase bits onto an infinite tape, printing and moving depending on the machine's state at the moment, one of a preconfigured finite set. Deutsch's insights include setting up the quantum tape as a sequence of qudits¹ and recognizing the challenge of determining whether the quantum computation has even finished — unlike in classical computing, where one can simply store a separate variable `done?` which is true if the calculation has finished and `false` otherwise, one of the quirks and challenges of quantum computing is its nondeterminism, and reading part of the quantum tape prematurely could cause unintended changes to its state. Accordingly, Deutsch sets aside a “signal” qubit which can safely be measured.

Let us now trace a very brief history of the development of quantum algorithms, touching on aspects that are relevant to this dissertation but certainly missing vast swaths of beautiful modern science and mathematics. The first quantum algorithms primarily solved toy problems. The very first was in fact from the same paper of Deutsch [D'85], which decides whether a function $f : \{0, 1\} \rightarrow \{0, 1\}$ is constant or not. Classically, this requires two queries to f to disambiguate: knowing that (say) $f(0) = 0$ says nothing about whether $f(1)$ is also 0 or not. However, Deutsch, with an appropriate implementation of f for quantum circuitry, is able to decide it in a single query to f .² Not long after, the seminal quantum algorithms of Lov Grover [G'96] and Peter Shor [S'94] provided examples of striking speedups to problems of broad interest: searching unstructured data and integer factorization. Grover's algorithm provides a quadratic speedup: what would require as many as $2^n - 1$ search queries on a list of size 2^n would require merely on the order of $\sqrt{2^n}$ quantum search queries. Shor's algorithm would go further, providing an exponential speedup over the present state of the art for integer factorization.³

¹The more familiar term *qubit* refers specifically to quantum states allowed to realize to only the quantum bits $|0\rangle$ or $|1\rangle$, whereas qudits can realize to a larger set of digits, typically expressed as $|0\rangle, \dots, |d-1\rangle$ for some $d \in \mathbb{N}$.

²Deutsch and Richard Jozsa [DJ92] generalized Deutsch's algorithm to boolean functions $f : \{0, 1\}^n \rightarrow \{0, 1\}$ which take in n bits and are guaranteed to either be constant, or to return 0 exactly half the time and 1 half the time. Classically, one would need as many as $2^{n-1} + 1$ function queries, an exponential demand, in the event of querying the 2^{n-1} inputs which happen to return (say) 0; on a quantum computer, still a single query suffices.

³The problem of the most efficient classical factorization algorithm is still open. Presently, the best algorithm runs in quasipolynomial time, but there are experts [C'16; SR'18] who believe that it will eventually fall into the polynomial regime.

In the 2010s, a new variety of algorithm was developed for near-term use: *variational quantum algorithms* (VQAs) leverage the respective capabilities of present-state classical and quantum computers (cf. [C⁺21], [D⁺25, §20]). Specifically, since the current state-of-the-art quantum computers lack enough logical qubits to solve very hard problems, but can efficiently evaluate quantum circuits, optimization problems are offloaded to more reliable classical solvers. Notable examples of VQAs are *variational quantum eigensolvers* [P⁺14], useful in quantum chemistry,¹ and *quantum approximate optimization algorithms* [FGG⁺14], used for solving combinatorial problems by encoding the problem as another continuous minimization problem.

However, despite the limited state of physical quantum computers at present, there has been enormous progress in recent years. The bellwether result was the 2019 announcement from Google [A⁺19] that their flagship quantum computer could simulate problems (such as sampling from the output distribution of a randomly-generated quantum circuit) that were nigh-impossible classically. While this was met with swift counterarguments from competitors and the community at large (cf. [P⁺19]), it still marked a first instance of quantum computers doing something fundamentally *quantum* at any moderate scale, and developments towards this end continue through the present [K⁺25].

Before proceeding, I shall comment on some somewhat technical core ideas, properties, and terminology of the quantum circuit model that is considered in this dissertation. We can write a quantum state on a single qubit as $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, obeying $|\alpha|^2 + |\beta|^2 = 1$. This normalization automatically tells us about the valid linear transformations (*gates*) that we may apply to a qubit: since they must preserve the length of the vector $(\alpha, \beta) \in \mathbb{C}^2$, they must be *unitary* operators. One example² of such a gate is the *Hadamard matrix*

$$\frac{i}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

This unitarity property is true of systems on any number of qubits, and we can move to multi-qubit systems by taking tensor products. This is in accordance with basis sizes: an n -qubit system should concern length- n bitstrings, each a separate basis element, so the dimension would be 2^n , the same as the dimension of $(\mathbb{C}^2)^{\otimes n}$. For instance, when $n = 2$, the standard basis is $|0\rangle \otimes |0\rangle$, $|0\rangle \otimes |1\rangle$, $|1\rangle \otimes |0\rangle$, and $|1\rangle \otimes |1\rangle$, which are typically abbreviated as $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$. Tensor products of gates similarly carry to each tensor factor; $U \otimes V$ acts on $|\xi\rangle \otimes |\phi\rangle$ as $(U|\xi\rangle) \otimes (V|\phi\rangle)$. This unitarity is also essential to the *no-cloning theorem*, which says that there is no quantum gate U on $2n$ qubits such that, for any $|\phi\rangle$, it is the case that $U(|\phi\rangle \otimes |0\rangle^{\otimes n}) = |\phi\rangle \otimes |\phi\rangle$ [WZ⁺82; D⁺82]. Indeed, in general, quantum gates will act *reversibly*.

If we wish to extract data about a state from the quantum computer, then we may *measure* a qubit or qubits. Instead of simply reporting the pair (α, β) , $|\psi\rangle$ probabilisti-

¹In particular, in estimating the *ground state energy* of central concern in Chapter 5.

²For details on a number of other standard gates, see Section 1.4.3.

cally *collapses*, proportionately with the implicit L^2 probability measure on the basis states: it will return 0 and be in the $|0\rangle$ state with probability $|\alpha|^2$, and return 1 and be in the $|1\rangle$ state with probability $|\beta|^2$. Moreover, note that the qubit is *global phase invariant*: if we multiply the whole qubit by $\omega \in \mathbb{C}$ where $|\omega| = 1$, then we get the same probabilities: $|\omega\alpha|^2 = |\alpha|^2$ and $|\omega\beta|^2 = |\beta|^2$. (It is then apparent that the space of gates on n qubits is in effect the quotient of the full unitary group by the center (scalars ω), namely $U(2^n)/U(1) \cong PU(2^n)$.) The behavior of measuring with probability the square of the coefficient is known as the *Born rule* and is exactly analogous for any number of qubits, and further is true of any measurement basis: if we instead choose to think of $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ in the *plus/minus basis* $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) = H|0\rangle$ and $|-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) = H|1\rangle$, then we have $|\psi\rangle = \frac{\alpha+\beta}{\sqrt{2}}|+\rangle + \frac{\alpha-\beta}{\sqrt{2}}|-\rangle$; the probability of measuring the plus state is $\left|\frac{\alpha+\beta}{\sqrt{2}}\right|^2$ while the probability of measuring the minus state is $\left|\frac{\alpha-\beta}{\sqrt{2}}\right|^2$.¹ Moreover, measurement can be done only partially: if we have the 2-qubit *entangled* system $|\psi\rangle = \frac{2\sqrt{2}}{3}|00\rangle + \frac{1}{3}|11\rangle$ and measure on the second qubit, with probability $\frac{8}{9}$ we return 0 and are left in the state $|0\rangle$, while with probability $\frac{1}{9}$ we return 1 and are left in the state $|1\rangle$ (this determinacy in the outcome of the remaining qubit is why we think of the two qubits as “entangled”).

What if we just “forget about” or “discard”² the second qubit? We do *not* end in $|\psi'\rangle = \frac{2\sqrt{2}}{3}|0\rangle + \frac{1}{3}|1\rangle$. Rather, we need a classical probabilistic mixture of the two outcomes, which we capture using outer products: $\frac{8}{9}|0\rangle\langle 0| + \frac{1}{9}|1\rangle\langle 1|$. This is a *density matrix*, written in the standard (*computational*) basis as $\begin{pmatrix} \frac{8}{9} & 0 \\ 0 & \frac{1}{9} \end{pmatrix}$. Note that the purported result $|\psi'\rangle$ has density matrix $|\psi'\rangle\langle\psi'| = \begin{pmatrix} \frac{8}{9} & \frac{2\sqrt{2}}{9} \\ \frac{2\sqrt{2}}{9} & \frac{1}{9} \end{pmatrix}$ — a decidedly different matrix! This “forgetting” operation is more commonly known as *tracing out* the qubit,³ and gives rise to quantum states with density matrices that *cannot* be factored as $|\psi'\rangle\langle\psi'|$; such unfactorizable states are known as *mixed states*, in contrast to *pure states* like $|0\rangle$, $|1\rangle$, $|\psi\rangle$, and $|\psi'\rangle$.

When working with multiqubit systems, we sometimes classify and group qubits. A *register* is, for our purposes, just a designated subset of our qubits (or “wires”) that we give a collective name. An *ancilla* qubit or register arrives pre-initialized and can be used for “scratch work” or for some other auxiliary reason in the quantum computation.

¹One may readily check that these probabilities sum to 1, as $|\alpha|^2 + |\beta|^2 = 1$. The rewriting of $|\psi\rangle$ into $|\pm\rangle$ can be seen from the relations $\sqrt{2}|0\rangle = |+\rangle + |-\rangle$ and $\sqrt{2}|1\rangle = |+\rangle - |-\rangle$.

²One conceptualization of discarding is to completely lose access to the qubit by throwing it into a black hole.

³For more details about tracing out, see Section 1.4.4.




Figure 1.1: A Deutsch–Jozsa circuit.

As an example, consider Figure 1.1, an instance of the Deutsch–Jozsa algorithm for a specific function; the top three wires are one register, while the fourth wire is the ancilla register; the boxes farthest right indicate that at the end of the algorithm we take a measurement.

Theoretical and engineering challenges for quantum speedups remain even in the wake of Grover’s and Shor’s algorithms and subsequent developments in the 21st century. There is of course the absence at present of sufficiently large quantum computers: the largest have on the order of thousands of qubits in total, but due to the tendency of external noise¹ to corrupt quantum state fidelity, *quantum error correction* is needed. This in turn can reduce the effective number of qubits available for a proper computational task (known as the number of *logical qubits*). Further, there is a need for increasingly specialized circuit components. In some instances, such as Grover’s algorithm, there is no circumventing the need for an oracle specific to the problem instance, by which I mean a specialized circuit component depending on the details of the algorithmic challenge; for instance, if the unstructured search is looking for the fifth element in the list, the Grover oracle will be different than if the target is actually the twelfth element. In others, there are certain fixed gates which must be implemented. Indeed, consider Shor’s algorithm for integer factorization. It relies on a *period-finding* method; that is: given a sequence $a_1, a_2, \dots$ guaranteed to repeat $a_1, a_2, \dots, a_m, a_1, a_2, \dots$, find m . Shor’s solution is quite spiritually similar to the Fourier transform from real and harmonic analysis and the discrete Fourier transform used in applications of Fourier analysis to digital data. However, Shor’s circuit crucially uses quantum gates called *controlled rotations*, which in effect perform certain rotations of increasingly minuscule angles about certain axes, and these gates must be somehow manufactured.

This manufacturing (or *synthesis*) problem is at the center of some very exciting recent work in quantum hardware design and quantum error-correcting codes. The

¹Many modern architectures tend to be held below 1 Kelvin, so noise can take the form of external sources of heat.




Figure 1.2: The gates U and then V applied to $|\psi\rangle$ correspond to just applying VU .

situation is as follows. When assembling quantum circuits to solve specific problems, we will need specific gates, which are parameterized by the problem we are trying to solve (a good example of this is the angles of controlled rotations for Shor’s quantum Fourier transform, which get exponentially small in the number of qubits). As such, it is not necessarily a sensible use of resources to mass-fabricate those specific gates if they will only have limited applicability. Instead, a very close approximation to the right quantum gate will suffice, and it is possible to assemble this approximation from mass-fabricated gates. Composition of gates is multiplicative: if a certain register is subject to the action of the gate U and then the gate V , this is the same as the action of the single gate VU (see Figure 1.2). Thus, gates form a *multiplicative group* modeled as a matrix group, and the unitary nature of state transformations as well as phase invariance imply that this group can be modeled as a *projective unitary group* on n qubits. Recall that the underlying vector space representing the possibilities for the n qubits has dimension 2^n (each length- n bitstring corresponds to a computational basis vector), so the group in particular is $\text{PU}(2^n)$. This group is endowed with a unique metric d which is *bi-invariant*: distance does not change when both inputs are moved in the same way. In symbols, for any $x, y, g, h \in \text{PU}(2^n)$, we have that

$$d(gxh, gyh) = d(x, y).$$

Additionally, there are two reasons to look to use few gates. Not only do we wish to use as few of our mass-fabricated gates as possible simply for cost reasons, but the deeper the circuit, the greater the opportunity for noise to creep in. Thus, the problem we seek to solve is:

Problem (gate synthesis): *Let n be a constant. Find a subset $\mathcal{S} \subset \text{PU}(2^n)$ such that for any $g \in \text{PU}(2^n)$ and $\varepsilon > 0$, there exists a sequence $s_1, \dots, s_m \in \mathcal{S}$ such that m is at most a fixed-degree polynomial in $\log \frac{1}{\varepsilon}$ and $d(g, s_1 \cdots s_m) < \varepsilon$. Then, given g and ε , efficiently find such a sequence $s_1, \dots, s_m$.*

By “efficient” we mean here that the algorithm should run on a classical computer and use a number of elementary operations which is polynomial in $\log \frac{1}{\varepsilon}$. Note that we have the $\text{poly} \log \frac{1}{\varepsilon}$ requirement both on the runtime of the approximate-factorization $g \approx s_1 \cdots s_m$ as well as for m itself. The runtime constraint is mainly the natural formalization of a “polynomial-time algorithm” in this setting, since $\log \frac{1}{\varepsilon}$ is the number of input bits required for the precision parameter; it also represents the logarithm of the cost of a brute-force strategy.

Why should we expect the factorization-length bound to be satisfiable at all? First, there is a heuristic argument. Suppose the words of length ℓ had the idealistic property that the union of ε -radius balls about each word exactly covered the group. This would mean that every element of the group is within ε of some word of length ℓ . Then, we would have

$$\varepsilon^{2^{2n}-1}(\#S)^\ell \approx 1,$$

or $\ell \approx 2^{2n-1} \log_{\#S} \frac{1}{\varepsilon}$,¹ where we have used the dimension of the manifold to approximate the volume of the ε -ball, and we say that the manifold's total volume is 1. More compelling than a heuristic is the *Solovay–Kitaev theorem* [K'97; DN'06], originally from the 1990s, which addresses the problem of polylogarithmic-length approximations, and has remarkably few assumptions for a result of its strength. The theorem says that for *any* set S which generates a dense subgroup,² given $g \in \text{PU}(2^n)$ and $\varepsilon > 0$, there is an approximation using elements of S which has length at most $C \log^3 \frac{1}{\varepsilon}$ for some constant C [KSV'02, Corollary 8.5.1].³ (Not only is this true of any such gate set, but any pair of gates gives rise to a dense subgroup, as the alternative forms a measure-0 subset of the manifold.)

However, despite these results working in great generality, there are improvements to be had, and much can be said with more amenable choices of S . There is a line of work for the case $n = 1$ (a single qubit) which begins with Vadym Kliuchnikov, Dmitri Maslov, and Michele Mosca [KMM'12] that considers number-theoretic gate sets, namely the standard *Clifford+T* gates. In this setting, the exponent of $\log \frac{1}{\varepsilon}$ is actually 1 and the cost model treats the gates generating a finite group (the Clifford gates) as being free and the involution T giving access to the full group as having cost 1. Subsequently, Ori Parzanchevski and Peter Sarnak [PS'18] showed that an algorithm with the same cost structure which produces approximations of length any smaller than $\log_{\#S} \frac{1}{\varepsilon}$ is computationally infeasible.⁴

Moreover, [PS'18] generalizes work of Neil Ross and Peter Selinger [RS'16], showing that for Clifford+T the target length is typically⁵ at most $3 \log_2 \frac{1}{\varepsilon^3}$, to any so-called *golden gate set*. A golden gate set is one with optimal covering properties, introduced in Sarnak's 2015 letter to Scott Aaronson and Andy Pollington [S'15b], and this letter forms the basis for the extension to more general sets, in particular ones with access to larger logarithm bases. The roots of the ideas come from the papers of Alex Lubotzky,

¹Here it is of note that an exponential-in- n term appears in the logarithm's prefactor, and thus it is important that n be a constant, which in turn is reasonable because of locality constraints on quantum hardware.

²By *dense subgroup* here we simply mean the stated condition: that the words $s_1 \cdots s_m$ can be used to approximate *any* g to precision ε .

³Christopher M. Dawson and Michael A. Nielsen [DN'06] give an explicit algorithm with length having exponent 3.97, which is worse than the existence result found in [KSV'02]. In 2023, Greg Kuperberg substantially sharpened the exponent to $\log_\varphi 2 \approx 1.44$ [K'23].

⁴In that this problem is NP-complete.

⁵In particular, on all but a set of negligible volume.

Ralph Phillips, and Sarnak [LPS’86a; LPS’86b; LPS’87; LPS’88] where the first infinite families of graphs were constructed which have maximal spectral gaps of their *graph Laplacians*, in the sense that they match the Alon–Boppana bound [N’91]. The random walk in these graphs spreads out optimally evenly and quickly, analogous to how the golden gates spread out in $\text{PU}(2)$, and in Sarnak’s letter [S’15b], he subsequently connects the coverage property in $\text{PU}(2)$ to the spectral gaps of so-called *Hecke operators* from number theory.

The *icosahedral golden gates* are an example of the *super golden gates* introduced in [PS’18] which have optimal covering properties of the unitary group and have minimal redundancy; the icosahedral gates have access to the largest finite group component of any super golden gate set, and so gain the most from the logarithm base. Compare the bound $3 \log_2 \frac{1}{\varepsilon^3}$ from Clifford+ T to $3 \log_{59} \frac{1}{\varepsilon^3}$; they differ by a factor of $\log_2 59 \approx 5.9$. However, it is possible to do even better. Leveraging ideas appearing in [S’21], Chapter 2 is dedicated to proving the following:

Theorem 1. *Let $S \cup \{\tau\}$ be the set of icosahedral golden gates, where S is a group and τ is the identified involution. For $\varepsilon > 0$ and a typical element $g \in \text{PU}(2)$, there exist $s_1, \dots, s_m \in S \cup \{\tau\}$ such that $d(g, s_1 \cdots s_m) < \varepsilon$ and at most $\frac{7}{3} \log_{59} \frac{1}{\varepsilon^3}$ of the icosahedral gates equal τ . These gates can be computed (classically) in time polynomial in $\log \frac{1}{\varepsilon}$.*

This result is restated in Theorem 2.2.1 and represents an additional improvement of $\frac{9}{7}$ over the general result given in [PS’18], and is the state of the art for this problem.

In order to prove this result, we exploit deep structural properties of the group generated by $S \cup \{\tau\}$ and recognize that factorization reduces to a lattice search in $\mathbb{Z}[i, \varphi] = \{a + ib + \varphi c + i\varphi d : a, b, c, d \in \mathbb{Z}\}$.¹ The condition of being ε -close to the target g is reframed as a specific factorization which initially appeared in a discrete setting in work by Eduardo Carvalho Pinto and Christophe Petit [CP’18] and was made effective in the continuous setting by earlier work [S’21]. Solving the lattice search problems is achieved by A. Paz’s sharpening [P’83] of Hendrik Lenstra’s algorithm [L’83].

There is also exciting work developing subsequent methods for multiple qubits. Shai Evra and Parzanchevski [EP’22] established the golden gate equivalent for *qutrits* (where the group is $\text{PU}(3)$), and Rahul Dalal, Evra, and Parzanchevski [DEP’25] have done the same for 2 and 3 qubits.

Above, we considered a model free of noise or other corruption; we comment here briefly on the need for quantum error-correction [S’95; S’96]. While we wish to act by (say) the icosahedral gates on the logical qubits, we still need to encode this action onto the full set of physical qubits via some encoding scheme in order to be *fault-tolerant* against corruption (faults) from external error sources. There has been recent work on the so-called “exotic” setting which includes the icosahedral golden gates [KT’23].

We now move from classical algorithms for hardware problems to scientific computing tasks to be done on a quantum computer. A core question to address before consid-

¹Here, i is the imaginary unit obeying $i^2 = -1$ and φ is the golden ratio $\frac{1+\sqrt{5}}{2}$.

ering any nontrivial scientific computing is how to even store and transform data with a quantum computer. At face value, the setting is extremely restrictive, since all gates must be unitary and can only multiply, and states must be normalized and nondeterministically lose a good deal of structure as soon as they are measured. Nonetheless, for matricial data there is a technique for embedding into quantum circuitry which is amenable to desired manipulations, known as *block-encoding*. A block-encoding of a matrix A consists of finding a constant $\alpha > 0$ and a larger unitary matrix U for $m > n$ such that the action of U on a state, with the extra qubits measuring to $|0\rangle$, is precisely that of $\frac{A}{\alpha}$. Equivalently, U is a unitary matrix whose top-left block is $\frac{A}{\alpha}$:

$$U = \begin{pmatrix} \frac{A}{\alpha} & \star \\ \star & \star \end{pmatrix}.$$

The parameter α is known as the *subnormalization factor* and accounts for the fact that while $\|U\| = 1$ we might still wish to work with matrices A having $\|A\| > 1$. This adaptability is not free, however. First, the extra qubits in U are provided by ancilla qubits, which requires additional logical and hence physical qubits. Further, large subnormalization factors compound and can decrease the likelihood of correctly measuring on the ancilla qubits. While the ancilla cost is not directly connected to α , lower values of α are possible with more qubits, so there is a meaningful resource tradeoff due to possibly needing to rerun the circuit due to the failure probability.

Since obtaining the correct measurement on the block-encoding ancilla register is not guaranteed, we typically assume when working with block-encodings that we might have to compile the circuit many times, and that we have many copies of the block-encoding as well as of the *prepare oracle* (i.e., a blackbox circuit) for the input vector or state that we wish to evaluate. We assume also that these are our primary costs, so we are interested in using as few instances of the block-encoding oracle as possible, while maximizing the success probability.

Within this model there is a very general goal.

Problem: *Devise algorithms to evaluate matrix functions of square input matrices provided as block-encodings.*

The nature of “matrix functions” here can be broadly construed. One early instance is matrix inversion: $f(A) = A^{-1}$; the algorithm for executing this on a quantum computer is due to Aram Harrow, Avinatan Hassidim, and Seth Lloyd [HHL’08]. Another sensible starting point, given the multiplicative nature of the quantum circuit model, is the matrix product, i.e., to implement $f(A, B) = AB$. Unfortunately, if U_A and U_B are block-encodings of A and B using the same number of ancillae, then simply evaluating $U_A U_B$ is *not* a block-encoding of AB . However, there is a simple combination due to András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe, which entails using different ancilla registers, so that the number of ancillae adds [G⁺’19] (see Proposition 3.1.1 and Figure 3.4). With some more effort, (weighted) sums $f(A, B) = aA + bB$

for $a, b \in \mathbb{C}$ are accessible through the technique known as *linear combinations of unitaries* (LCU), first appearing in work of Andrew M. Childs and Wiebe [CW12]. In principle, this gives immediate access to arbitrary polynomials, say of a single matrix: $f(A) = a_m A^m + \dots + a_1 A + a_0$, simply by composing [G⁺19]’s multiplication circuit and [CW12]’s addition circuit, and an idea for an optimized approach could follow *Horner’s method*: $f(A) = a_0 + A(a_1 + \dots + A(a_{m-1} + A a_m))$. However, there are better approaches that use less ancilla overhead. The main ones are known as *quantum signal processing* (QSP), due to Low and Isaac Chuang [LC16], and the *quantum singular value transform* (QSVT), also from Gilyén, Su, Low, and Wiebe’s work [G⁺19]. Both construct very carefully-chosen sequences of phases interleaved with block-encoding instances, and solving the thorny systems which give rise to the correct phases is a question of active interest [H19]. QSP and QSVT fit into a broader classification of quantum algorithms, the *grand unification*, due to John M. Martyn, Zane M. Rossi, Andrew K. Tan, and Chuang [M⁺21].¹

The above discussion mainly concerns matrix functions which carry over directly to the singular values or eigenvalues. Consider the matrix inverse problem for invertible Hermitian matrices, so $f(A) = A^{-1}$. $A \in \mathbb{R}^{N \times N}$ can be written as $V \Lambda V^\dagger$ where V is unitary and $\Lambda \in \mathbb{R}^N$ is diagonal, so $f(A) = V \Lambda^{-1} V^\dagger = V f(\Lambda) V^\dagger$; that is, we have directly applied f to A ’s eigenvalues. We call this an *eigenvalue transform*, and the same can be done for non-Hermitian matrices via the singular value decomposition: if $A = U \Sigma V^\dagger$ and $f : \mathbb{R} \rightarrow \mathbb{R}$ then the singular value transform of A by f can be conceived as $U f(\Sigma) V^\dagger$. But what about more complicated functions than the inverse? It is clear that the same phenomenon will occur with polynomials, but how would one matricially implement, say, a filter which preserves the eigenvalues falling in the interval $[0, 1]$ and zeroes out all other eigenvalues? The general problem is:

Problem (Quantum Eigenvalue Transformation (QET)): *Let A be a diagonalizable complex $N \times N$ matrix with eigenvalues $\{\lambda_k\} \subset \mathbb{D}$ and corresponding eigenvectors $\{v_k\}$. Given a function $f : \mathbb{C} \rightarrow \mathbb{C}$, a state $|v\rangle = \sum c_k |v_k\rangle$, and $\varepsilon > 0$, output a normalization of the state $f(A)|v\rangle := \sum c_k f(\lambda_j) |v_k\rangle$ to precision ε .*

Approaches mainly consider constructing very close approximations of f using a branch of complex analysis known as *approximation theory* and implementing the polynomial

¹When $\|A\| \leq 1$, block-encoding is known as *dilation*, and there is a rich theory of results about polynomial transformations in this setting. The first result is known as *Székelyfalvi-Nagy’s theorem* and states that there exists a dilation U such that U^n is a dilation of A^n for any $n \in \mathbb{N}$ [P03]. Székelyfalvi-Nagy’s theorem proves *von Neumann’s theorem*: for any polynomial f , $\|f(T)\| \leq \sup_{z \in \mathbb{D}} |f(z)|$. Ando’s theorem is a generalization to two dilations: for A_1 and A_2 which commute, there exist dilations U_1 and U_2 for which $U_1^{n_1} U_2^{n_2}$ is a dilation of $A_1^{n_1} A_2^{n_2}$, and there is a corresponding von Neumann-type inequality. However, this structure breaks down for three dilations, and it is an open research direction [H25] as to how badly the corresponding von Neumann-type inequality fails; i.e., how large can the ratio between $\|f(A_1, A_2, A_3)\|$ and $\sup_{z_1, z_2, z_3 \in \mathbb{D}} |f(z_1, z_2, z_3)|$ get?

using QSP. There is also a generalization, introduced in [BSS'23]:

Problem (Multivariate Quantum Eigenvalue Transformation (MQET)):

Let $A = \{A_\ell\}_{0 \leq \ell \leq r}$ be a family of $r + 1$ pairwise commuting diagonalizable complex $N \times N$ matrices. Let $\{v_k\}$ be a mutual eigenbasis of A and let $\lambda_k = (\lambda_{\ell,k})_{0 \leq \ell \leq r} \in \mathbb{D}^{r+1}$ denote the vector of eigenvalues for the mutual eigenvector v_k . Given a function $f : \mathbb{C}^{r+1} \rightarrow \mathbb{C}$, a state $|v\rangle = \sum c_k |v_k\rangle$, and $\varepsilon > 0$, output a normalization of the state $f(A)|v\rangle := \sum c_k f(\lambda_k) |v_k\rangle$ to precision ε .

A *normal matrix* is one which commutes with its adjoint: M is normal if $MM^\dagger = M^\dagger M$. Normal matrices obey the decomposition $M = A + iB$, where A and B are commuting Hermitian matrices. While self-adjoint operators are by far the most commonly-encountered in quantum mechanics and quantum computing, there are systems that may be driven by normal, non-self-adjoint operators in settings such as decaying subsystems [B'07] and non-Hermitian optics [FEG'17]. Notice that if f is a (monovariate) function of a normal matrix M , then there is g which is a bivariate function of A and B . One key application is in a setting such as [B'07; FEG'17] to consider $e^{-i\beta M t}$, known as *time-evolution* and motivated more for the final chapter. The main result of Chapter 3 is an algorithm for MQET, with an application to normal matrices and with exponentiation being of particular interest.

Theorem 2. *Consider r pairwise commuting Hermitian matrices with their block-encodings each requiring m ancillae and a function f with a r -variate polynomial approximation g agreeing to within ε on all of $[-1, 1]^r$. Then, there is an algorithm requiring on the order of $r(\deg g)^r$ block-encoding instances, $r(m + \log \deg g)$ additional ancillae, and $(\deg g + 2)^r$ subnormalization factor,¹ which solves the MQET problem with precision ε . When $r = 2$, this solves QET with a normal input matrix.*

To make these algorithms work, we will exploit that polynomials admit relatively cheap decompositions, from the perspective of block-encoding parameters and phase factors. We rely heavily on decomposition into the *Chebyshev polynomials*, a common and elegant choice of “base” (orthogonal) polynomial, which are trivial to implement with QSP. Other polynomials used in the approximation are not necessarily Chebyshev, and their QSP phase factors must be classically computed. Then, that the polynomials are in *commuting* variables is essential for plugging in the matrix outputs. When devising the initial polynomial approximation g , we also exploit the Chebyshev basis for $r - 1$ of the input variables to write the dependence in the final variable not as an integral as would be required in the general setting, but rather as an exact sum.

¹We anticipate that instantiations of this algorithm would perform with better parameters than the stated guarantees, as they derive from lossy inequalities in the analysis. The framework of Rossi, Jack L. Ceroni, and Chuang [RCC'25] achieves some better parameter guarantees (often linear) by avoiding LCU entirely, but requires far more phase factor computations.

Similar work by Rossi and Chuang [RC’22] and extended with Ceroni [RCC’25], known as *multivariable QSP* or *M-QSP*, achieves many of the deliverables, although based on a separate approximation-theoretic conjecture and with a different cost model.

We now return to Feynman’s original vision for quantum computing: greater access to the laws of nature by leveraging quantum mechanics on a quantum machine. The problem we focus on here is known as *Gibbs state preparation*, based on the *Gibbs measure*. The measure is over possible molecular configurations σ (say, for particle spins) subject to a given temperature T and an energy function H (called the *Hamiltonian*) of σ . Each configuration should be weighted as¹

$$\rho_\beta(\sigma) \propto e^{-\beta H(\sigma)},$$

where it is customary to work in terms of the inverse temperature $\beta = \frac{1}{T}$, which arises much more commonly in this context. Explicitly, the Gibbs measure on states is therefore

$$\rho_\beta(\sigma) = \frac{e^{-\beta H(\sigma)}}{\sum_{\sigma'} e^{-\beta H(\sigma')}},$$

where the denominator is the *partition function*.

While the Gibbs measure is a classical probability measure, there is a natural corresponding quantum object. The Hamiltonian can more generally be realized as an operator (in classical mechanics, an operator acting on position and momentum in each spatial dimension), and quantum Hamiltonian mechanics are similarly driven by an operator H , still known as the Hamiltonian, which is self-adjoint. We work instead with the *Gibbs state*

$$\rho_\beta := \frac{e^{-\beta H}}{\text{tr } e^{-\beta H}}.$$

Notice that, if we view H ’s eigenvalues as the energies $H(\sigma)$, then the Gibbs state’s eigenvalues are precisely the quantum eigenvalue transform (with the same inverse-temperature parameter β) which yields the Gibbs measure, and the partition function $\text{tr } e^{-\beta H}$ plays the same role.

¹Let’s consider a brief motivational sketch [W’24; W’25; G’02; E’25]. Recall the *Gibbs free energy* from thermodynamics: at fixed temperature and pressure, the Gibbs free energy is a linear combination of the enthalpy and the temperature times the entropy. In statistical mechanics we are focused more on macroscopic statistics and phenomena that arise from specific arrangements of molecules in space, which themselves are generally not predictable to any fine-grained detail; thus, we must express some uncertainty on which particular arrangement is present at any given moment. We can therefore imagine working with some probability distribution over possible particle arrangements and replace the enthalpy by H ; when we move to the expectation over the states given by our measure of choice, we then have the enthalpy term replaced by the expected energy, and the entropy term replaced by the entropy of the probability measure on states. In this way the Gibbs free energy can now be viewed as a function F from probability measures on a set of discrete states to $\mathbb{R}$. In chemistry, nature automatically maximizes the Gibbs free energy, and thus in turn we should seek the probability distribution which maximizes F . After a brief calculation [W’24], we recover the stated proportionality relation.

In classical probability theory and Markov chain theory, the problem of *Gibbs sampling* is of special interest. It asks: given a Hamiltonian and an inverse temperature β , can one sample from the Gibbs measure ρ_β ? i.e., is there an efficient algorithm that returns σ with probability $\rho_\beta(\sigma)$? Approximate algorithms are acceptable, i.e., returning σ with probability $\mathbb{P}[\sigma]$ where $|\mathbb{P}[\sigma] - \rho_\beta(\sigma)| < \varepsilon$, so long as the algorithm does not depend badly on the inverse approximation parameter $\frac{1}{\varepsilon}$. Arguably the most influential algorithm for this problem is known as the *Metropolis–Hastings algorithm*, based on work by Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, Edward Teller [M⁺53], and W. K. Hastings [H70]. The approach proceeds by walking in the graph of states and at each stage considering making a move, but rejecting the move with a certain probability, with each part balanced to respect the target Gibbs measure. Naturally, there is a similar quantum problem:

Problem (Gibbs state preparation): *Given a Hamiltonian and an inverse temperature β , can one prepare the Gibbs state on a quantum register with reasonable resource costs?*

The Gibbs state in particular is typically not a pure state, so this preparation algorithm will require measuring or tracing out some ancillae along the way. This has been an area of great interest in the quantum computing community in recent years. As noted above, since the Gibbs state is a quantum eigenvalue transformation of the Hamiltonian, some of the scientific computing approaches from above could be considered, as was done in works by Joran van Apeldoorn, Anirban Narayan Chowdhury, Chuang, Gilyén, Sander Gribling, Low, Rolando D. Somma, Su, Wiebe, Ronald de Wolf, and Theodore J. Yoder [CS’17; V⁺’17; VG’18; G⁺’19; YLC’14], but the resource costs have proved prohibitive [D⁺’25, §15].

Instead, a motivating picture is that the starting state can be arbitrary, but we wish to transform it to one at temperature T , so the natural thing to do is to “leave the system in a refrigerator (heat bath) of temperature T .” How can this be done? A refrigerator consists of air kept at a fixed temperature which barrages the food left inside, until the temperature reaches equilibrium. So-called *Monte-Carlo* or *Lindbladian* approaches simulate this effect by iteratively applying certain carefully-chosen gates with ancillae which get traced out before the next application [L’76]; after many rounds the state converges to the Gibbs state. We call the Lindbladian operator $\mathcal{L}$ and consider the sequence of transitions

$$\sigma \mapsto \mathcal{L}(\sigma) \mapsto \mathcal{L}^2(\sigma) \mapsto \dots \xrightarrow{\text{limit?}} \rho_\beta.$$

We consider a specific style of Lindbladian known as a *Davies generator* [D’76] which is subject to the same structural constraints as the general Lindbladian (as established by Vittorio Gorini, Andrzej Kossakowski, and E. C. G. Sudarshan [GKS’76]) and obeys an additional self-adjointness property, guaranteeing that $\mathcal{L}$ ’s eigenvalues are all real [RFZ’10]. In particular, we are concerned with the *least nontrivial eigenvalue* of $\mathcal{L}$,

which controls the *mixing time*, i.e., the iteration at which the above sequence of applications of $\mathcal{L}$ approaches ρ_β .

The role of this eigenvalue can be cleanly seen for *random walks on graphs*. Consider a connected network on n nodes where you are positioned at a fixed node but would like to end up at a uniformly random node in the network in as few steps as possible, with some imprecision ε allowed in the end distribution.¹ For simplicity, suppose every node has the same number of neighbors. The network has an associated matrix known as the *graph Laplacian* L which depends on the vertex structure; L has least eigenvalue 0 and least nontrivial eigenvalue γ . Then, by simply choosing to walk to a neighbor chosen uniformly at random at each step, the task will be completed after $t(\varepsilon)$ steps, obeying

$$\left(\frac{1}{\gamma} - 1\right) \log \frac{1}{2\varepsilon} \leq t(\varepsilon) \leq \frac{1}{\gamma} \log \frac{n}{\varepsilon}.$$

(This result can be found in the book of David A. Levin, Yuval Peres, and Elizabeth L. Wilmer [LPW'17, Chapter 12].) So, the behavior of this iterative random walk algorithm is well-controlled by $\frac{1}{\gamma}$. The same is true of many other algorithms that can be represented similarly by a matrix.

This quantity was influentially studied for $\mathcal{L}$ by Kristan Temme [T'13]. He made an interesting connection between the quantum realm (of the Davies generator itself as a quantum Markov chain) to the classical realm (relating the least nontrivial eigenvalue of a certain, smaller, classical Markov chain, which is in principle accessible through entirely classical means) with a pessimistic outlook revealed by some worst-case behavior. The main result of Chapter 4 is to show that generically, the situation is not nearly so dire:

Theorem 3. *Suppose the Hamiltonian H has no 3-term arithmetic progression amongst its eigenvalues: for all $\lambda < \mu < \nu$ eigenvalues of H , $\nu - \mu \neq \mu - \lambda$. Then, letting $\lambda_{\mathcal{L}}$ be the least nontrivial eigenvalue of $\mathcal{L}$ and λ be the least nontrivial eigenvalue of a certain corresponding classical chain, we have that $\lambda_{\mathcal{L}}$ lies between $\frac{1}{2}\lambda$ and λ . If instead the longest arithmetic progression is length D , then $\lambda_{\mathcal{L}}$ lies between $\frac{1}{2D}\lambda$ and λ .*

(This result is restated in Theorem 4.0.11.) Moreover, the condition that H is void of 3-term arithmetic progressions is verified for small perturbations of many commonly-encountered quantum Hamiltonian families. The proof of this result works with the *Bohr spectrum*, consisting of the differences of the Hamiltonian's eigenvalues. The first step is to reduce to considering only Hermitian inputs. Then, for a fixed Bohr difference ω , the eigenvalues of the Gibbs state are related to an average of consecutive eigenvalues at difference ω ; the number of terms being averaged is at most the longest arithmetic progression among H 's eigenvalues.

¹This imprecision is known as the *total variation distance*; its definition is not directly relevant to the discussion at hand.

We now turn to a problem lying just outside of quantum computing and back to Feynman's vision of simulation of quantum systems. It concerns more connections between the eigenvalues of quantum-mechanical operators and classical analogues. Specifically, we begin by considering the *Schrödinger equation*. The Schrödinger equation simply says that the evolution in time of a quantum state as determined by a Hamiltonian is governed by

$$i\hbar\partial_t|\psi(t, x)\rangle = H|\psi(t, x)\rangle,$$

where $\psi(t)$ is the state, changing as t increases from 0, the initial time; i is the imaginary unit and $\hbar \approx 1.05 \times 10^{-34} Js$ is the (*reduced*) *Planck constant*; and H is the Hamiltonian (the *Schrödinger operator*), corresponding to the energy.¹ While in principle H could be any operator, physically we are interested in the H capturing the energies that are most commonly encountered: *kinetic* and *potential* energy of one particle (a *single-body* system). The kinetic energy is captured in this setting by a rescaling of the second derivative² or *Laplacian*, which in one-dimensional space for a particle having mass m looks like³

$$-\frac{\hbar^2}{2m}\partial_x^2.$$

The potential energy is more abstractly captured by some *potential function* V from $\mathbb{R}^n$ to $\mathbb{R}$. (In Newtonian mechanics, the most familiar potential function is mgh , where m is the mass, g is local gravitational acceleration, and h is the input parameter of height over some baseline.) We will focus on *periodic potentials*: for each of the n standard basis directions v_i , for all $x \in \mathbb{R}^n$, $V(x) = V(x + v_i)$. (In one dimension, this is just to say that $V(x) = V(x + 1)$.) All told, it is common to work with the self-adjoint Hamiltonian

$$Hf(x) := -\frac{\hbar^2}{2m}\Delta f(x) + V(x)f(x).$$

It is important to note that even though the prescribed particle is situated in only n dimensions, H itself is an infinite-dimensional operator.

We may also move past the highly specific value taken on by $\hbar$; we replace it with a positive parameter θ which we drive to 0. The study of such operators is known as *semiclassical analysis* [Z'12], and the regime of interest is precisely the limit $\theta \downarrow 0$. Thus, we are interested in

$$H_\theta f(x) = -\frac{\theta^2}{2m}\Delta f(x) + V(x)f(x).$$

¹It is important to note that $\psi(t, x)$ is also a function from all of $\mathbb{R}^n \ni x$ to $\mathbb{C}$ when the particle is in n -dimensional space, which will become more relevant momentarily.

²Kinetic energy comes from momentum or velocity, which are themselves merely first derivatives. The presence of the second derivative can be seen in one dimension from taking inner products of $\partial_x|\psi(t, x)\rangle$ against itself: since ∂_x is self-adjoint, this inner product is $\langle\psi(t, x)|\partial_x^*\partial_x|\psi(t, x)\rangle = \langle\psi(t, x)|\partial_x^2|\psi(t, x)\rangle$.

³In higher dimensions, ∂_x^2 is often written as Δ .

In analogy with the Gibbs state preparation Lindbladian, the behavior is driven by something like the least nontrivial eigenvalue. However, it is important to note that in the setting of infinite-dimensional operators, it is common to have no “true” eigenvalues in the sense of solving an eigenvalue equation $Hx = \lambda x$. Instead, we work with the generalized idea of the *spectrum*: $\sigma(H)$ consists of all $\lambda \in \mathbb{C}$ for which $H - \lambda$ is not invertible, and it is always the case that for $\lambda \in \sigma(H)$ and $\varepsilon > 0$, there exists a “pseudoeigenfunction” x such that $\|(H - \lambda)x\| < \varepsilon$. In finite dimensions, these definitions coincide, and it remains the case in finite or infinite dimensions that it can be difficult to access precise information about the spectrum of large and complicated operators or systems. Therefore, one may try to work with or simulate on a computer a simplified system which is discretized. Consider a mesh located only at the integer points, scaled by the *semiclassical parameter* θ , i.e. $\theta\mathbb{Z}^n \subset \mathbb{R}^n$. We can simulate the action of H . V is trivial to imagine extending to the lattice mesh — just take the same value. For the derivatives, the path is by *finite differences*. For a function f of a single real variable x and a small parameter δ , we can estimate $f'(x) = \frac{f(x+\delta)-f(x)}{\delta}$, and subsequently $f''(x) = \frac{f(x+2\delta)-2f(x+\delta)+f(x)}{\delta^2}$. As such, after recentering so that $f''(x) = \frac{f(x+\delta)-2f(x)+f(x-\delta)}{\delta^2}$, the discretization of H_θ is precisely

$$\tilde{H}_\theta g(\theta k) := \left(2n g(\theta k) - \sum_{j \sim k} g(\theta j) \right) + V(\theta k) g(\theta k)$$

where $j \sim k$ refers to those lattice points in $\mathbb{Z}^n$ differing from k by exactly 1 in a single entry. Here, notice that θ^2 has cancelled out, and that for simplicity we treat $m = \frac{1}{2}$.

Problem: Are the spectral infima of H_θ and $\tilde{H}_\theta$ the same in the limit $\theta \downarrow 0$?

Since $\tilde{H}_\theta$ is intended to replicate the behavior to some order of H_θ , one would expect the answer to be yes. The main result of Chapter 5 is to verify this in explicit terms for generic choices of periodic potential V and semiclassical parameter θ . Let λ_θ be the spectral infimum of H_θ , and $\tilde{\lambda}_\theta$ be the spectral infimum of $\tilde{H}_\theta$.

Theorem 4. For any periodic potential V and any θ irrational,

$$\tilde{\lambda}_\theta \leq \lambda_\theta.$$

Additionally, for any sufficiently regular periodic potential V and any θ sufficiently small, there are functions ε_1 and ε_2 obeying $\lim_{\theta \downarrow 0} \varepsilon_1(\theta) = \lim_{\theta \downarrow 0} \varepsilon_2(\theta) = 0$ such that

$$(1 + \varepsilon_1(\theta))\tilde{\lambda}_\theta \geq \lambda_\theta - \varepsilon_2(\theta).$$

Further, for a wider class of sufficiently regular potentials V (not necessarily periodic), there exists a constant c_n depending only on n and V 's regularity for which

$$c_n \tilde{\lambda}_\theta \geq \lambda_\theta.$$

The exact forms of ε_1 and ε_2 are stated in Theorem 5.3.1, and the exact form of c_n is stated in Theorem 5.0.11. Using techniques from semiclassical analysis, similar results with sharper asymptotic ε_1 and ε_2 but no control over the precise constants in these functions were obtained by Simon Becker, Jens Wittsten, and Maciej Zworski [BWZ'26]. Previously, this problem had been studied for very specific choices of V , such as the *almost-Mathieu operator* $V(x) \propto 1 - \cos(2\pi x)$ [AJ'09] in one dimension; the best results are due to Florin Boca and Alexandru Zaharescu [BZ'05] but require very specialized and delicate trigonometric calculations.

Our results follow by far more generalizable considerations of *test functions*: functions which exhibit the stretching behavior characteristic of eigenvalues at least as extreme as the elements of the spectrum. The first stated result above follows by considering a continuous pseudoeigenfunction of the (continuous) semiclassical operator, and describes a discrete function that can be extracted which gives rise to the behavior captured by the inequality. One difficulty which must be overcome is that continuous functions are only defined up to “sets of measure zero,” which in practice means that they may be changed at isolated points without changing the overall behavior or structure of the function. In particular, this means that simply sampling the test function on the discrete lattice is insufficient. The other two results correspond to the other direction: starting with a discrete test function, we construct a test function on each hypercube in $\mathbb{R}^n$ with vertices in $\theta\mathbb{Z}^n$, defined via a natural expected value. The result is a function that agrees with the given values on the lattice and is multiaffine in each input variable within each hypercube; moreover, the result is in fact continuous, because the “attachments” at the hypercube faces depend only on the lattice points forming that face. Subsequently, we heavily use techniques from an area known as *Fourier analysis of Boolean functions* — the analogue to traditional Fourier analysis, but for affine transformations of functions of the form $\{0, 1\}^n \rightarrow \{0, 1\}$ [O'14].

It is worth noting that the motivation for the problem presented here followed the reasoning from *semiclassical* to *discrete*. In the original paper of Nikhil Srivastava, Isabel Detherage, and myself [DSS'24], we actually present the *reverse* perspective: beginning with the discrete operators, and their applications in geometric group theory, how can one learn more given information about the corresponding semiclassical operator? This perspective is preserved in Chapter 5.

1.2 Themes

There are some themes that play out across several components of this dissertation.

Spectral gaps. Prominently, *spectral gaps* are at the core of Chapter 4 and Chapter 5 — both concern connecting the least nontrivial eigenvalue of a quantum operator with the least nontrivial eigenvalue of a naturally-arising classical analogue. In particular, Theorem 4.0.11 relates the spectral gap of a quantum operator and a corre-

sponding classical operator to within a factor of 2 when the underlying Hamiltonian’s spectrum has no arithmetic progressions. Theorem 5.0.6, Theorem 5.0.9, and Theorem 5.0.11 also connect the spectral gaps of related quantum and classical operators, giving equality in the limit of the parameter $\theta \downarrow 0$ with effective bounds on the θ -dependence. There is also the spectral gap connection for Chapter 2, where the idea to construct the (super) golden gate sets is deeply related to spectral problems and spectral graph theory.

Rayleigh quotients. Relatedly, the ability to even access these spectral gaps in Chapter 4 and Chapter 5 is due to the *Rayleigh quotient*. For a self-adjoint finite-dimensional operator A acting on a space $\mathcal{H}$ and nonzero test function $x \in \mathcal{H}$, the Rayleigh quotient is defined as $\frac{x^\top Ax}{x^\top x}$. For x a λ -eigenvector, this will precisely equal λ , but the Courant–Fischer minmax theorem says, in part, that if A ’s least eigenvalues are $\lambda_0 < \lambda_1$, with eigenvectors x_0 and x_1 , then¹

$$\lambda_0 = \inf_{\substack{x \in \mathcal{H} \\ x \neq 0}} \frac{x^\top Ax}{x^\top x},$$

$$\lambda_1 = \inf_{\substack{x \in \mathcal{H} \\ x \neq 0 \\ x \perp x_0}} \frac{x^\top Ax}{x^\top x}.$$

Therefore, exhibiting, for instance, x for which the Rayleigh quotient is at most λ implies that $\lambda_0 \leq \lambda$. This idea is exploited heavily and fruitfully in both of those chapters, where the numerator of the Rayleigh quotient is split as in the proof of Proposition 4.1.10 and in Claim 5.2.6 and Claim 5.2.7; the proof of Proposition 4.1.10 also splits the Rayleigh quotient’s denominator.

Low-degree polynomial approximation is crucial in Chapter 3 and Chapter 5. In the former, access to low-degree polynomial approximations gives access to useful scientific computing primitives like quantum signal processing, which is well-studied and optimized for the situation of polynomial eigenvalue transforms. In particular, the degree of the approximation is the main parameter driving the cost of the algorithm provided in Theorem 3.0.1. In the latter, the multiaffine approximation to local hypercube behavior is essential to granting detailed information about the Rayleigh quotients of such test functions — more complicated structure would have made for possibly impassable analysis, but the multiaffineness is the “sweet spot” of conceptual simplicity and meaningful first-order approximation to actual test functions. Concretely, it enables the boolean Fourier decompositions which are used in the lower bound portion through results like Proposition 5.3.5.

¹The expression for λ_0 is also true when A is infinite-dimensional.

Efficient classical computation enabling improved quantum techniques is seen throughout. In Chapter 2, the entire algorithm happens on a classical computer, and the purpose is to enable shorter quantum circuit depth. In Chapter 3, the QSP phase factors must be computed classically and allow for the visible improvement in ancilla overhead, as well as the integrals-turned-sums due to the involvement of the Chebyshev basis. Finally, accessing improved analysis of quantum systems via sharp analysis of classical operators is one way to motivate Chapter 5.

1.3 Publication history

Chapter 2 has previously appeared in *Quantum Information and Computation* as [BS'23] and was coauthored with Terrence Richard Blackman. Chapter 3 has previously appeared as [BSS'23] and was coauthored with Yonah Borns-Weil and Tahsin Saffat. Chapter 4 has previously appeared in *Quantum Information and Processing* as [B⁺'25] and was coauthored with João Basso, Shirshendu Ganguly, Alistair Sinclair, Nikhil Srivastava, and June Vuong. Chapter 5 has previously appeared in *Pure and Applied Analysis* as [DSS'24] and was coauthored with Izzy Detherage and Nikhil Srivastava. I am grateful to each of my coauthors for their collaboration throughout my PhD experience!

1.4 Selected preliminaries

Here I will introduce some of the somewhat specialized background content that appears in multiple chapters of this dissertation. Concepts and notation appearing in only a single chapter have been left *in situ*.

1.4.1 Notation

Sets. Take $n \in \mathbb{N}$. We use $[n]$ to refer to the set $\{1, \dots, n\}$. We use $[[n]]$ to refer to $\{0, \dots, n-1\}$.

Probability. We use $\mathbb{E}$ to denote the expectation of a random variable, and $\mathbb{E}_x[f(x, y)]$ to denote the partial expectation of an expression with respect to the random variable x only, treating the other variables as constants.

The Iverson bracket. For a proposition P , we use the *Iverson bracket* $[P]$ to resolve to 1 if P is true and 0 otherwise. For instance, $[\pi \in \mathbb{R}] = 1$ while $[i \in \mathbb{R}] = 0$.

1.4.2 Linear algebra

Rayleigh quotient. For a linear operator T , we let its domain be $D(T)$. We introduce the *Rayleigh quotient* for any operator T and function $f \in D(T)$:

$$\mathcal{R}_T(f) = \frac{\langle f, Tf \rangle}{\langle f, f \rangle}.$$

Norms. We use $\|\cdot\|$ to denote the ℓ_2 and L^2 norms of functions, and the operator norms of operators. We use $\|\cdot\|_p$ to denote the ℓ_p and L^p norms more generally. We use $\|\cdot\|_F$ to denote the Frobenius norm of an operator.

Gibbs states. We reiterate the *Gibbs state* at inverse temperature β for a Hamiltonian H :

$$\rho_\beta := \rho_\beta(H) := \frac{\exp(-\beta H)}{\text{tr} \exp(-\beta H)}.$$

1.4.3 Single-qubit gates

Pauli gates. We introduce the *Pauli gate set*, each of which is a 2×2 unitary matrix:

$$X := \begin{pmatrix} & 1 \\ 1 & \end{pmatrix}, \quad Y := \begin{pmatrix} & -i \\ i & \end{pmatrix}, \quad Z := \begin{pmatrix} 1 & \\ & -1 \end{pmatrix}, \quad I := \begin{pmatrix} 1 & \\ & 1 \end{pmatrix}.$$

Additionally, these matrices form a real basis for the 2×2 Hermitian matrices.

The Clifford group. Let

$$H := \frac{i}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \text{ and } T := \begin{pmatrix} e^{i\pi/8} & \\ & e^{-i\pi/8} \end{pmatrix}.$$

Define the *Clifford group* to be the group $C_{24} = \langle H, T^2 \rangle \cong S_4$, the symmetric group on four symbols which is also the group of symmetries of the cube and octahedron.

Single-site gates. Let $P \in \mathbb{C}^{2 \times 2}$ be any single-qubit operator (such as one of the Pauli matrices). Then the operator P_i , acting on the i th qubit of an n -qubit system, is defined as:

$$P_i := I^{\otimes(i-1)} \otimes P \otimes I^{\otimes(n-i)}.$$

1.4.4 Quantum circuits

Registers. We can sometimes rewrite circuits with shorthand to indicate acting on entire registers at once. In Figure 1.3, note the shorthand indicating that the top wire is actually a register of three qubits.




Figure 1.3: Figure 1.1, rewritten more compactly, with U the balanced function oracle.




Figure 1.4: A SWAP gate

The SWAP gate. We define the two-qubit gate

$$\text{SWAP} : |\psi\rangle \otimes |\psi'\rangle \mapsto |\psi'\rangle \otimes |\psi\rangle,$$

denoted in circuitry as shown in Figure 1.4. SWAP has the matricial form

$$\begin{pmatrix} & 1 & \\ 1 & & \\ & 1 & \end{pmatrix}$$

in the computational basis.

Tracing out. When we trace out a register, we denote this with the circuit component . In a bipartite system on registers $\mathcal{A}$ and $\mathcal{B}$, if we have (mixed) state σ and trace out $\mathcal{A}$ then we write the result as the state

$$\text{tr}_{\mathcal{B}} \sigma = \sum_{|\psi\rangle \in \mathcal{A}} \left(\langle \psi |_{\mathcal{A}} \otimes \text{id}_{\mathcal{B}} \right) \sigma \left(| \psi \rangle_{\mathcal{A}} \otimes \text{id}_{\mathcal{B}} \right)$$

where $\mathcal{A}$ is any basis for the qubits in register $\mathcal{A}$ (often taken to just be the computational basis). This represents a classical probability distribution on the possible outcomes on register $\mathcal{B}$ after measuring $\mathcal{A}$, and is invariant under the choice of $\mathcal{A}$.

Chapter 2

Fast navigation with icosahedral golden gates

Lubotzky, Phillips, and Sarnak, in a series of papers [LPS’86a; LPS’86b; LPS’87; LPS’88], explicitly constructed topological generators with optimal covering properties for the compact Lie group $\text{PU}(2)$. Such generators for projective unitary groups find an interesting application in quantum computing where a fundamental design challenge is to determine an optimal, fault-tolerant decomposition of a quantum gate. For classical computing a single bit state is an element of $\{0, 1\}$. A classical gate implements functions on binary inputs. The only nontrivial single bit logic operation is NOT, which takes 0 to 1 and 1 to 0 (though it is also possible for the codomain to contain more than one bit). In the quantum setting, we recall that qubits are points $u = (u_1, u_2) \in \mathbb{C}^2$ up to a mutual phase $e^{2\pi i \theta}$ in each component, such that $|u|^2 = |u_1|^2 + |u_2|^2 = 1$. A gate here *cannot* output more than one qubit, and thus must be a 2×2 projective unitary.

A *universal gate set* is a finite set of gates, $S := \{s_1, s_2, \dots, s_k : s_\ell \in \text{PU}(2)\}$, that can approximate, in the bi-invariant metric on the compact Lie group, any matrix arbitrarily well. That is, the group generated by S must be topologically dense in $\text{PU}(2)$. The Solovay–Kitaev theorem [NC’10] guarantees that universal gate sets can efficiently approximate quantum operations for unitaries on a constant number of qubits. Universal gate sets typically consists of a finite group $C \leq \text{PU}(2)$ together with an extra element τ , which we will take to be an involution, so that the subgroup generated by C and τ covers $\text{PU}(2)$ with minimal τ -count, and simultaneously navigates $\text{PU}(2)$ efficiently. That is, given some gate in $\text{PU}(2)$ and desired precision ε , there is an efficient algorithm (polynomial-time in the input size) that with high probability finds a short word in S to that precision, typically of length $O(\log(1/\varepsilon))$. The deep insight of Sarnak [S’15b] is that the construction and optimality of universal single-qubit quantum gate sets can be understood in terms of the arithmetic of quaternion algebras. Specifically, identifying C with a subgroup of the group of units in a definite quaternion algebra over a totally real number field provides a coherent framework within which one can systematically address the question of optimality of topological generators for $\text{PU}(2)$.

The question we address within this context is that of finding the “best” topological generators of $\text{PU}(2)$ among those universal gate sets, the *super golden gates* of Parzanchevski and Sarnak [PS’18], which are known to possess optimal covering properties and efficient navigation. In particular, there are only finitely many super golden gates [PS’18, pp. 870–871] and the respective finite subgroups C can be realized as the group of symmetries of the Platonic solids: for the tetrahedron, A_4 ; for the cube and the octahedron, S_4 ; and for the dodecahedron and icosahedron, A_5 .

We demonstrate that the *icosahedral* super golden gates admit a factorization directly analogous to the one obtained by Stier [S’21], and that this gives the best-known preconstant to the first-power logarithm in the approximation length. The exact factor of the improvement is $\log_2 59 \approx 5.9$. This improvement is due to the fact that the icosahedral super golden gates have a growth rate that is on the order of 59^k , while gate set studied in [S’21], the Clifford+ T gates, have growth rate of order 2^k . We note also that the *icosahedral* super-golden-gates represent the greatest number of distinct gates with bounded τ -count.

The commonly used Clifford+ T gate set provides a set of elementary gates that is universal and consists only of a small number of gates, all of which are very well compatible with many established error correction schemes and can be physically implemented in all quantum technologies that seem promising for large-scale quantum computations [NWD’20]. In this case the finite group C is the Clifford group C_{24} of order 24 in $\text{PU}(2)$. At least one non-Clifford gate must be added to the basic gate set in order to achieve universality. A common choice for this additional gate is the T -gate (or $\frac{\pi}{8}$ -gate). The T -gate is not the only possible extension of the Clifford group but it is considered to be the most practical one. This is due to the availability of fault-tolerant implementations of the T -gate. For this reason, the Clifford+ T gate set is considered the most promising candidate for practical quantum computing. We recall, for the convenience of the reader, some of the salient details of the single qubit Clifford+ T gate set. We take $S := \{H, T\}$ [S’15b; S’21; RS’16]. The Clifford+ T universal gate set is an example of a golden gate set [S’15b; PS’18]. The remarkable observation of [S’15b] is that the H and T gates of the Clifford+ T gate set come from the definite quaternion algebra

$$\mathcal{A} := \left(\frac{-1, -1}{\mathbb{Q}(\sqrt{2})} \right).$$

Kliuchnikov, Maslov, and Mosca [KMM’12] characterized the group $\langle S \rangle$ and demonstrated an efficient algorithm to factor exactly, by carefully studying the powers of $\sqrt{2}$ in the denominators of matrix entries. Ross and Selinger [RS’16] then focused on diagonal matrices near to $\langle S \rangle$, with the goal of finding ε -close factorizations of length $c \log_2(1/\varepsilon^3)$; the base of 2 is intrinsic to the structure of the Clifford+ T gates. By studying *upright sets* in the plane, which function analogously to the simplices of Lenstra’s algorithm [L’83; P’83], they achieved the leading coefficient of $1+o(1)$ (in that restricted case). Parzanchevski and Sarnak [PS’18] generalized Ross and Selinger’s approach to

any golden gate set (with the base of the logarithm correspondingly changing per the gate set), and by considering Euler angles reached the coefficient of $3 + o(1)$ for approximations to generic elements. Working with the (finite) LPS Ramanujan graphs (see Section 2.1.1 and [LPS'88]), Carvalho Pinto and Petit [CP'18] factorize in the equivalent of $7/3 + o(1)$. This approach is adapted by Stier [S'21] for the same coefficient with Clifford+ T gates. We combine aspects of the techniques of [CP'18; S'21], in the proposed icosahedral setting, to reduce this bound to $\frac{7}{3} \log_{59}(1/\varepsilon^3)$.

2.1 Super golden gates

We recall here essential ideas related to the arithmetic of quaternion algebras [V'80; V'21], S -arithmetic groups [M'01, §C], and Ramanujan graphs [LP'20] which lead to the icosahedral super golden gates [PS'18].

2.1.1 Quaternion algebras

A quaternion algebra $\mathcal{A}$ over a field $\mathbb{F}$ is a central simple algebra of dimension four over $\mathbb{F}$ (cf. [M'06, §5.2] or [J'97, p. 15]). It follows from Wedderburn's structure theorem on simple algebras [D'47] that every quaternion algebra over any field $\mathbb{F}$ is either isomorphic to $M_2(\mathbb{F})$ or a division algebra with center $\mathbb{F}$ [L'06]. If the characteristic of $\mathbb{F}$ is not 2 then it is always possible to find a basis $\{1, i, j, k\}$ for $\mathcal{A}$ over $\mathbb{F}$ such that

$$i^2 = \alpha, j^2 = \beta, k = ij = -ji$$

where $\alpha, \beta \in \mathbb{F}^\times$. We designate such an algebra by

$$\mathcal{A} = \left(\frac{\alpha, \beta}{\mathbb{F}} \right).$$

Evidently $q \in \mathcal{A}$ is of the form

$$q = x_0 + x_1 i + x_2 j + x_3 k, \tag{2.1.1}$$

where $x_\ell \in \mathbb{F}$. In this notation, Hamilton's quaternions arise as

$$\mathbb{H} = \left(\frac{-1, -1}{\mathbb{R}} \right).$$

The conjugate of q as in (2.1.1), denoted by $\bar{q}$, is equal to $x_0 - x_1 i - x_2 j - x_3 k$. For each $q \in \mathcal{A}$ we define the (reduced) norm map $N : \mathcal{A} \rightarrow \mathbb{F}$ by

$$N(q) := q\bar{q} = x_0^2 - \alpha x_1^2 - \beta x_2^2 + \alpha\beta x_3^2$$

and the (reduced) trace map $\text{tr} : \mathcal{A} \rightarrow \mathbb{F}$ by

$$\text{tr}(q) := q + \bar{q} = 2x_0.$$

We note that every element $q \in \mathcal{A}$ satisfies the quadratic equation

$$q^2 - \text{tr}(q)q + N(q) = 0.$$

It is possible to make everything completely explicit by embedding $\mathcal{A}$ in $M_2(F(\sqrt{\alpha}))$ by, for example,

$$i \mapsto \begin{pmatrix} \sqrt{\alpha} & \\ & -\sqrt{\alpha} \end{pmatrix} \text{ and } j \mapsto \begin{pmatrix} & \beta \\ 1 & \end{pmatrix}.$$

Evidently, if α is a square in $\mathbb{F}$, then $\mathcal{A} \cong M_2(\mathbb{F})$. A necessary and sufficient condition for $\mathcal{A} \cong M_2(\mathbb{F})$ is that α is the norm of an element in $\mathbb{F}(\sqrt{\beta})$ with respect to $\mathbb{F}$ [V21]. Of course, one may interchange α and β in this remark.

Specializing to quaternion algebras over the rational field $\mathbb{Q}$, or over one of the completions $\mathbb{Q}_p$ or $\mathbb{R}$ (the completion $\mathbb{Q}_\infty$), let $\mathcal{A}$ be a quaternion algebra over $\mathbb{Q}$ and let p be a prime (or ∞). We define

$$\mathcal{A}_p := \mathcal{A} \otimes_{\mathbb{Q}} \mathbb{Q}_p.$$

Either $\mathcal{A}_p \cong M_2(\mathbb{Q}_p)$ or $\mathcal{A}_p$ is a division algebra over $\mathbb{Q}_p$. In this case we will have that $\mathcal{A}_p \cong \mathbb{H}_p$ and say that $\mathcal{A}$ is *ramified* at p . When $\mathcal{A}_p \cong M_2(\mathbb{Q}_p)$ we will say that $\mathcal{A}$ is *unramified* or *split* at p . If $\mathcal{A}$ is ramified at ∞ it is called a *definite rational quaternion algebra*—that is, if $\mathcal{A}$ is definite then

$$\mathcal{A}_\infty = \mathcal{A} \otimes_{\mathbb{Q}} \mathbb{R} \cong \mathbb{H}.$$

If $\mathcal{A}$ is split at ∞ , it is called an *indefinite rational quaternion algebra*—that is, if $\mathcal{A}$ is indefinite then

$$\mathcal{A}_\infty = \mathcal{A} \otimes_{\mathbb{Q}} \mathbb{R} \cong M_2(\mathbb{R}).$$

Definite quaternion algebras over totally real number fields

We now turn to the quaternion algebras that are the objects of interest in this chapter, definite quaternion algebras over totally real number fields. Let $\mathcal{A}$ be a quaternion algebra over a number field $\mathbb{F}$, v be a place of $\mathbb{F}$, and $\mathbb{F}_v$ be the completion of $\mathbb{F}$ at v . Recall that a *totally real number field* is a finite algebraic extension of $\mathbb{Q}$ all of whose complex embeddings lie entirely in $\mathbb{R}$. For example, the field $\mathbb{Q}(\sqrt{d})$ is totally real for positive, integral d . If every infinite place of F is ramified in $\mathcal{A}$, we say that $\mathcal{A}$ is a *totally definite quaternion algebra*. Consequently, if $\mathcal{A}$ is a totally definite quaternion algebra over a number field $\mathbb{F}$, then $\mathbb{F}$ is necessarily totally real. Moreover, if $\mathbb{F}$ is a quadratic field, the number of finite places which are ramified in $\mathcal{A}$ is even.

Unit groups in orders in definite quaternion algebras over totally real number fields

Let $\mathcal{A}$ be quaternion algebra over a field $\mathbb{F}$ and let $O_{\mathbb{F}}$ be the ring of integers in $\mathbb{F}$. An element $q \in \mathcal{A}$ is *integral* if $N(q), \text{tr}(q) \in O_{\mathbb{F}}$. An *order* $O \subseteq \mathcal{A}$ is a $O_{\mathbb{F}}$ -algebra of integral elements such that [S'00, p. 2]

$$O \otimes_{O_{\mathbb{F}}} \mathbb{F} \cong \mathcal{A}.$$

Observe that as an example of an order we always have

$$O := O_{\mathbb{F}} \oplus O_{\mathbb{F}}i \oplus O_{\mathbb{F}}j \oplus O_{\mathbb{F}}k. \quad (2.1.2)$$

If O is an $O_{\mathbb{F}}$ -order in a definite quaternion algebra $\mathcal{A}$ over the totally real field $\mathbb{F}$ then the group of units of reduced norm 1, i.e.

$$O^1 := \{\alpha \in O : N(\alpha) = 1\} \quad (2.1.3)$$

is a finite group [V'21, p. 289]. These finite groups will correspond to the groups of rotational symmetries of the platonic solids [V'21, pp. 172–173].

Ramanujan graphs

Two key ideas are needed for the construction of golden gate and super golden gate sets: a S -arithmetic unit quaternion group and a Ramanujan graph. We outline the essential ideas of Lubotzky, Phillips, and Sarnak's [LPS'88] “LPS” construction of these objects.

Ramanujan graphs are graphs whose spectrum is bounded optimally. Let X be a finite connected k -regular graph and A its adjacency matrix.

Definition 2.1.4. The graph X is called a Ramanujan graph if every eigenvalue λ of A satisfies either $|\lambda| = k$ or $|\lambda| \leq 2\sqrt{k-1}$.

LPS Ramanujan graphs [DSV'03] arise as Cayley graphs of $\text{PSL}(2, \mathbb{F}_q)$ (for $\mathbb{F}_q$ the finite field on q elements). When considering these graphs we interchangeably refer to their elements by their group-theoretic properties as matrices and their graph-theoretic relations as vertices. [LPS'88] establishes that for any prime $p \equiv 1 \pmod{4}$ there are infinitely many $(p+1)$ -regular Ramanujan graphs. We use the notation $X^{p,q}$ where p and q are distinct primes congruent to 1 modulo 4 to represent such graphs. The construction comes from number theory by way of the generalized Ramanujan conjecture [PS'18, p. 873] [H'+21]. The symmetric space $\text{PGL}(2, \mathbb{Q}_p)/\text{PGL}(2, \mathbb{Z}_p)$ can be identified with a $(p+1)$ -regular infinite tree. $\text{PGL}(2, \mathbb{Z}[1/p])$ acts from the left on $\text{PGL}(2, \mathbb{Q}_p)/\text{PGL}(2, \mathbb{Z}_p)$. The generalized Ramanujan conjecture, a theorem in this case, implies that the quotient of $\text{PGL}(2, \mathbb{Q}_p)/\text{PGL}(2, \mathbb{Z}_p)$ by any congruence subgroup

of $\mathrm{PGL}(2, \mathbb{Z}[1/p])$, a $(p+1)$ -regular graph, is a Ramanujan graph. By considering an appropriate congruence subgroup of $\mathrm{PGL}(2, \mathbb{Z}[1/p])$ we can identify the quotient of this symmetric space with a Cayley graph associated to $\mathrm{PSL}(2, \mathbb{F}_q)$ or $\mathrm{PGL}(2, \mathbb{F}_q)$, depending on the value of the Legendre symbol $\left(\frac{p}{q}\right)$ [S'15a; BCC'19].

p -arithmetic unit quaternion groups

Golden gate and super golden gate sets for $\mathrm{PU}(2)$ require the construction of a p -arithmetic group (a special case of S -arithmetic groups, where S is a collection of places of $\mathbb{Z}$). Let $G \leq \mathrm{GL}(n)$ be an algebraic group defined over $\mathbb{Z}[1/p]$ with $G(\mathbb{R})$ compact. A p -arithmetic group Λ is a subgroup of $G(\mathbb{Z}[1/p]) \leq G(\mathbb{R}) \times G(\mathbb{Q}_p)$, and has *congruence subgroup* $\Lambda(N) := \{g \in \Lambda : g \equiv I \pmod{N}\}$. (For a broader perspective, see e.g. [H'26; L'26].) That is, in our compact Lie group $\mathrm{PU}(2)$ we take only rational numbers whose denominators are powers of a fixed prime p as coefficients in the matrices.

One can also make the same construction over the ring of integers $O_{\mathbb{F}}$ of a totally real number field $\mathbb{F}$ strictly containing $\mathbb{Q}$, at which point the role of p is played by $\mathfrak{p}$ a prime ideal of $O_{\mathbb{F}}$, so that the inverted elements (those allowed in denominators) are now all of $\mathfrak{p} \cap \mathbb{F}^{\times}$.

2.1.2 Golden gates

Golden gates are special unit groups in quaternion algebras over totally real number fields derived from the p -arithmetic groups [S'15b]. They give variants of optimal generators for $\mathrm{PU}(2)$ and connect the arithmetic of quaternion algebras to quantum computation on a single qubit. The “golden” characterization is to be understood by way of an interesting link between p -arithmetic unit groups coming from unit quaternion groups and the Ramanujan graphs $X^{p,q}$, which we explicate below.

Recall once more that in classical computation, one decomposes any function into basic logical gates such as XOR, AND, and NOR, and that in quantum computation, the classical bits are replaced by qubits, which are vectors in projective Hilbert space $\mathbb{C}P^n$, and that the logical gates are *all* the elements of the projective unitary group $\mathrm{PU}(2)$. Let S be a subgroup of $\mathrm{PU}(2)$ and denote by $S^{(\ell)}$ the set of ℓ -fold products of elements in S . If $\langle S \rangle = \bigcup_{\ell \geq 0} S^{(\ell)}$ is dense in $\mathrm{PU}(2)$ (with respect to the standard bi-invariant metric

$d^2(A, B) := 1 - \frac{|\mathrm{tr}(A^*B)|}{2}$) then S is universal. That is, any gate can be approximated with arbitrary precision as a product of elements of S .

The notion of a golden gate set is much stronger, requiring [LP'20]:

1. *Optimal covering of $\mathrm{PU}(2)$ by $\langle S \rangle$* : for every ℓ the set $S^{(\ell)}$ distributes in $\mathrm{PU}(2)$ as a perfect sphere packing (or randomly placed points) would, up to a negligible factor.

2. *Approximation*: given $A \in \text{PU}(2)$ and $\varepsilon > 0$, there is an efficient algorithm to find some $A' \in B_\varepsilon(A)$ (the ε -ball around A) such that $A' \in S^{(\ell)}$ with ℓ (almost) minimal.
3. *Compiling*: given $A \in \langle S \rangle$ as a matrix, there is an efficient algorithm to write A as a word in S of the smallest possible length.

These requirements ensure that any gate can be approximated and compiled as an efficient circuit using the gates in S .

2.1.3 Super golden gates

Each super golden gate set is composed of a finite group C and an involution τ , which lie in a $\mathfrak{p}$ -arithmetic group for $\mathfrak{p}$ a prime ideal of the integers $\mathcal{O}_{\mathbb{F}}$ of a totally definite quaternion algebra $\mathcal{A}$, over a totally real number field $\mathbb{F}$. We require that:

1. C acts simply transitively on the neighbors of any given vertex in X^p , the $(p+1)$ -regular tree, for p the norm of $\mathfrak{p}$.
2. τ is an involution which takes a vertex to one of its neighbors.

$\mathfrak{p}$ -arithmetic unit quaternion group act transitively on the vertices of the corresponding X^p . The $\mathfrak{p}$ -arithmetic groups which act transitively on the vertices of X^p are called the *golden gates* and the $\mathfrak{p}$ -arithmetic groups which act transitively on both the vertices and edges of X^p are called the *super golden gates*.

2.2 Icosahedral super golden gates

Recall that there are only finitely many such super golden gate sets and in each case the finite group C is identified with the group of rotational symmetries of a platonic solid: for the tetrahedron, A_4 ; for the cube and the octahedron, S_4 ; and for the dodecahedron and icosahedron, A_5 . Each of these finite groups can be precisely identified with a $\mathfrak{p}$ -arithmetic unit quaternion group coming from one of the following quaternion algebras [S'15b]:

- tetrahedral gates: $\mathcal{A} = \left(\frac{-1, -1}{\mathbb{Q}} \right)$;
- octahedral gates: $\mathcal{A} = \left(\frac{-1, -1}{\mathbb{Q}(\sqrt{2})} \right)$; and
- icosahedral gates: $\mathcal{A} = \left(\frac{-1, -1}{\mathbb{Q}(\sqrt{5})} \right)$.

We consider the quaternion algebra $\mathcal{A}$ over the *golden field* $\mathbb{Q}(\sqrt{5})$ [PS'18, p.895]. A maximal order $\mathcal{O}$ in $\mathcal{A} := \left(\frac{-1, -1}{\mathbb{Q}(\sqrt{5})} \right)$ is given by the ring of icosians. The unit group $\mathcal{O}^\times$

is the platonic icosahedral group, generated by $\left\langle \frac{1+i+j+k}{2}, \frac{1+\varphi^{-1}j+\varphi k}{2} \right\rangle$, where $\varphi := \frac{1+\sqrt{5}}{2}$ is the golden ratio, and $O^1 \cong A_5$. In $\text{PU}(2)$, this corresponds to

$$C_{60} = \left\langle \begin{pmatrix} 1 & 1 \\ i & -i \end{pmatrix}, \begin{pmatrix} 1 & \varphi - i/\varphi \\ \varphi + i/\varphi & -1 \end{pmatrix} \right\rangle =: \langle \rho, \sigma \rangle.$$

We take as our involution

$$\tau = \tau_{60} := \begin{pmatrix} 2 + \varphi & 1 - i \\ 1 + i & -2 - \varphi \end{pmatrix}.$$

For the prime ideal $\mathfrak{p} = (7 + 5\varphi)$ one has $\mathbb{F}_{\mathfrak{p}} \cong \mathbb{Q}_{59}$, and the generated group $\Gamma = \langle C_{60}, \tau_{60} \rangle$ is the full $(7 + 5\varphi)$ -arithmetic group of O . As such, we establish:

Theorem 2.2.1. *Subject to standard number-theoretic heuristic conjectures, there exists a factorization of any $g \in \text{PU}(2)$ to precision ε using τ -count at most $(7/3 + o(1)) \log_{59}(1/\varepsilon^3)$.*

In particular, if g is near to a diagonal matrix then we just apply the approach in Section 2.4 (for additive error) to get a path of length at most $(1 + o(1)) \log_{59}(1/\varepsilon^3)$, and otherwise run the approach in Section 2.5.

My “standard number-theoretic heuristic conjectures,” we refer to two sorts of assumptions about the solvability of integer problems. One concerns convex (in fact, linear) constraints: once the domain of consideration is large enough, generically there will be an integer lattice point in the interior. The other concerns sum of square problems in the number ring; Cramér’s conjecture may be interpreted as a claim about representability by certain quadratic forms on $\mathbb{Z}$ [S’19], and Sardari’s conjecture (*) has a similar flavor for somewhat more complicated domains. In this work we insist on such a heuristic for number rings extending $\mathbb{Z}$: presented with many elements of the number ring, *one* of them ought to be representable as a sum of two squares.

Notice that no other choice of golden gates is both super and has a greater logarithm base, so that these are the “best” generators of $\text{PU}(2)$ given the present state of knowledge, as an interesting application of number theory in that geometric setting. However, we make no claim in this work as to the relevance of icosahedral golden gates to fault-tolerant quantum computing, as such relationships are not yet established; given the present stage of research into quantum computers, we offer the results of this chapter merely to advance understanding of a gate set with specific known advantages, and many more unknown qualities.

2.3 Nearby elements in $\text{PU}(2)$, and factoring in Γ

In this section, we establish a key technical lemma regarding approximations in the matrix group, and the method for exact synthesis for elements of the relevant subgroup.

For $\alpha, \beta \in \mathbb{C}$ satisfying $|\alpha|^2 + |\beta|^2 = 1$ we write

$$u(\alpha, \beta) := \begin{pmatrix} \alpha & \beta \\ -\bar{\beta} & \bar{\alpha} \end{pmatrix},$$

and for $\theta \in \mathbb{R}$ we write

$$u(\theta) := u(e^{i\theta}, 0) = \begin{pmatrix} e^{i\theta} & \\ & e^{-i\theta} \end{pmatrix}.$$

We restate below [S'21, Lemma 5] which holds independent of our choice of universal gate set.

Lemma 2.3.1. *Select absolute constants $\delta, \varepsilon_0 > 0$ and put $C = \sqrt{\frac{1}{2} + \frac{1}{2}\left(\frac{2+\delta}{\varepsilon_0}\right)^2}$. Take γ_1 and γ_2 in either $\text{SU}(2)$ or $\text{PU}(2)$ and write them as $\gamma_1 = u(\alpha_1, \beta_1)$ and $\gamma_2 = u(\alpha_2, \beta_2)$. If $\|\alpha_1\| - \|\alpha_2\| < \varepsilon$ for some $\varepsilon < \delta$ and $\min\{|\alpha_1|, |\alpha_2|\} < \sqrt{1 - \varepsilon_0^2}$ then for*

$$\begin{aligned} \theta_1 &= \frac{1}{2}(\arg \alpha_1 - \arg \alpha_2 + \arg \beta_1 - \arg \beta_2), & \delta_1 &= u(\theta_1) \\ \theta_2 &= \frac{1}{2}(\arg \alpha_1 - \arg \alpha_2 - \arg \beta_1 + \arg \beta_2), & \delta_2 &= u(\theta_2) \end{aligned}$$

we have the approximation $\delta_1 \gamma_2 \delta_2$ to γ_1 , satisfying

$$d(\gamma_1, \delta_1 \gamma_2 \delta_2) < C\varepsilon.$$

Loosely, this result can be thought of as a sufficient condition to transform one 2×2 unitary into (approximately) another based only on a weak condition and by “tuning” with diagonal (rotation) matrices on either side. The condition is merely that the “starting” and “target” matrices have top-left entries nearby in absolute value, and that neither is very near to 1.

We now move to factorizing. Put $\eta = 7 + 5\varphi$ for the sequel. Observe that η is a positive real number, and the generator of $\mathfrak{p}$ above. For our purposes, we will encounter elements of $\Gamma = \langle \rho, \sigma, \tau \rangle \leq \text{PU}(2)$ only as $\mathbb{Z}[\varphi]$ -quaternions with norm a power of η , envisioned as elements of the Cayley graph for $\langle \rho, \sigma, \tau \rangle \leq \mathbb{Q}(\varphi) \text{U}_2(\mathbb{Z}[\varphi])$, which [PS'18] shows acts transitively with respect to the distance measure of τ -count, which can be detected by counting the power of η in the quaternion norm, after quotienting out $\mathbb{Z}[\varphi]$ -scalars. This leads to the following factoring algorithm.

Algorithm 2.3.2. Let C be a set of representatives of $\langle \rho, \sigma \rangle \leq \text{PU}(2)$ lifted to $\mathbb{Q}(\varphi) \text{U}_2(\mathbb{Z}[\varphi])$. Given α a lift of some element of Γ with τ -count k , it can be factored by determining which (unique) $c \in C$ gives rise to $\alpha c \tau$ of τ -count $k - 1$, which in turn requires no more than 60 multiplications of three matrices. c is found if and only if $\alpha c \tau$ has all coefficients divisible by η in $\mathbb{Z}[\varphi]$.

This algorithm has the base case of simply comparing α against all 60 elements of C , another trivial computational task. The idea of the general case is that α has the representation

$$c_0 \prod_{i \in [n]} \tau c_i$$

where c_i is not the identity for $i \neq 0, n$. Then c will be projectively equal to c_n^{-1} , giving rise to

$$\alpha c \tau = c_0 \prod_{i \in [n-1]} \tau c_i \tau c_n c \tau = z p c_0 \prod_{i \in [n-1]} \tau c_i$$

for $z = \det(c_n c)$ (coprime to η).

2.4 Algorithm for short paths to diagonal elements

2.4.1 Algorithm

We proceed by, for given diagonal element $\delta = u(\theta)$ and ε , seeking $\gamma \in \Gamma$ with $d(\delta, \gamma) < \varepsilon$. Knowing that γ is projectively equal to an element

$$\begin{pmatrix} x_0 + x_1 i & x_2 + x_3 i \\ -x_2 + x_3 i & x_0 - x_1 i \end{pmatrix}$$

for $x_0, x_1, x_2, x_3 \in \mathbb{Z}[\varphi]$ satisfying

$$x_0^2 + x_1^2 + x_2^2 + x_3^2 = \eta^m \tag{2.4.1}$$

for some $m \in \mathbb{N}$, we also have that it is sufficient to satisfy

$$x_0 \cos \theta + x_1 \sin \theta \geq \eta^{m/2} (1 - 2\varepsilon^2) \tag{2.4.2}$$

(note that this is precisely [RS'16, (13) and Problem 7.4] (see there for the derivation)), following the readily computable observation that $\|\delta - \gamma\| \geq \sqrt{2}d(\delta, \gamma)$ and setting the goal $\|\delta - \gamma\| \leq \sqrt{2}\varepsilon$, since [RS'16] operates with respect to the operator norm. Observe, unfortunately, that (2.4.1) is a quadratic constraint and (2.4.2) is a linear constraint. We now explain how to transform them into a practicable sequence of integer programs.¹

Fix m . We seek $x_\ell \in \mathbb{Z}[\varphi]$ satisfying (2.4.1) and (2.4.2). Define $y_\ell = x_\ell / \eta^{m/2}$. Artificially add the constraint

$$1 - \varepsilon^2 - y_1 \sin \theta \geq 0. \tag{2.4.3}$$

¹Fundamentally, this is the same idea as in [RS'16], as their study of upright ellipses accomplishes the same task as Lenstra's algorithm.

Observe from (2.4.1) that $y_0^2 \leq 1 - y_1^2$. Then, we have the sequence of implications (mainly by algebraic manipulation)

$$\begin{aligned}
 y_0 \cos \theta &\geq 1 - \varepsilon^2 - y_1 \sin \theta \geq 0 \\
 y_0^2 \cos^2 \theta &\geq (1 - \varepsilon^2)^2 + y_1^2 \sin^2 \theta - 2(1 - \varepsilon^2)y_1 \sin \theta \\
 \cos^2 \theta &\geq (y_1 - (1 - \varepsilon^2) \sin \theta)^2 - (1 - \varepsilon^2)^2 \sin^2 \theta + (1 - \varepsilon^2)^2 \\
 \cos^2 \theta (2\varepsilon^2 - \varepsilon^4) &\geq (y_1 - (1 - \varepsilon^2) \sin \theta)^2 \\
 |y_1 - (1 - \varepsilon^2) \sin \theta| &\leq |\cos \theta| \sqrt{2 - \varepsilon^2} \varepsilon
 \end{aligned} \tag{2.4.4}$$

and so we have reduced our consideration to just one of the four variables. As $[\mathbb{Q}(\varphi) : \mathbb{Q}] = 2$, we explicitly work with the Galois group elements

$$\begin{aligned}
 \sigma_+ : 1 &\mapsto 1 \\
 \varphi &\mapsto \varphi, \\
 \sigma_- : 1 &\mapsto 1 \\
 \varphi &\mapsto \varphi^\bullet,
 \end{aligned}$$

where $\varphi^\bullet$ is φ 's Galois conjugate $\frac{1-\sqrt{5}}{2} = 1 - \varphi$, both of which are real embeddings, yielding the additional constraints

$$|\sigma_\pm y_1| \leq 1. \tag{2.4.5}$$

Now, putting $x_1 = c + d\varphi$ for some $c, d \in \mathbb{Z}$, (2.4.3), (2.4.4), and (2.4.5) transform into

Problem 2.4.6. *Given m , find $(c, d) \in \mathbb{Z}^2$ satisfying the five linear constraints*

$$\begin{aligned}
 (c + d\varphi) \sin \theta &\leq \eta^{m/2} (1 - \varepsilon^2), \\
 |c + d\sigma_\pm \varphi| &\leq (\sigma_\pm \eta)^{m/2}, \\
 |c + d\varphi - \eta^{m/2} (1 - \varepsilon^2) \sin \theta| &\leq \eta^{m/2} |\cos \theta| \sqrt{2 - \varepsilon^2} \varepsilon.
 \end{aligned}$$

The existence of such a pair is determinable in $O(\text{poly } m)$ time, by Lenstra's algorithm [L83; P83].

Putting $x_0 = a + b\varphi$ for some $a, b \in \mathbb{Z}$, for a particular solution to Problem 2.4.6, (2.4.1) and (2.4.2) transform into

Problem 2.4.7. *Given m and $x_1 = c + d\varphi$, find $(a, b) \in \mathbb{Z}^2$ satisfying the four linear constraints*

$$\begin{aligned}
 |a + b\sigma_\pm \varphi| &\leq (\sigma_\pm \eta)^{m/2} \sqrt{1 - (\sigma_\pm x_1)^2}, \\
 (a + b\varphi) \cos \theta &\leq \eta^{m/2} (1 - x_1 \sin \theta), \\
 (a + b\varphi) \cos \theta &\geq \eta^{m/2} (1 - \varepsilon^2 - x_1 \sin \theta).
 \end{aligned}$$

Again, the existence of such a pair is determinable in $O(\text{poly } m)$ time, by Lenstra's algorithm.

Finally, if we have solutions to Problem 2.4.7 and Problem 2.4.6, we seek to solve

Problem 2.4.8. *Given m , $x_0 = a + b\varphi$, and $x_1 = c + d\varphi$, find $(x_2, x_3) \in \mathbb{Z}[\varphi]^2$ satisfying (2.4.1).*

Here we change techniques. The objective is to write $\eta^m - x_0^2 - x_1^2$ as a sum of squares in $\mathbb{Z}[\varphi]$. Assuming efficient factorization in $\mathbb{Z}$ (or $\mathbb{Z}[\varphi]$, also a PID), Problem 2.4.8 is efficiently solvable via Theorem A.5 (the general approach being basically identical to the classical algorithms for $\mathbb{Z}[i]$ or $\mathbb{Z}[e^{i\pi/8}]$, cf. [RS'16, §C] for the latter).

We have succeeded in the core of the algorithm. To handle given δ , begin by fixing $m = 0$. For given m , attempt to solve Problem 2.4.6; for each solution, attempt to solve Problem 2.4.7; for each solution, attempt to solve Problem 2.4.8. If this results in a tuple (x_0, x_1, x_2, x_3) satisfying all three problems, halt and return

$$\frac{1}{\eta^{m/2}} \begin{pmatrix} x_0 + x_1 i & x_2 + x_3 i \\ -x_2 + x_3 i & x_0 - x_1 i \end{pmatrix}.$$

Otherwise, if all possibilities have been exhausted, increment m .

2.4.2 Analysis

We refer the reader to [PS'18, §2.3]'s analysis of timing and correctness for Ross and Selinger's algorithm generalized to any golden gate set, where the conclusion is that for desired precision ε , the factorization length becomes $(1 + o(1)) \log_{59}(1/\varepsilon^3)$ with required computational time remaining $O(\text{poly } \log(1/\varepsilon))$.

2.5 Algorithm for short paths

Here we largely adopt the structure and wording from [S'21, §4 and §5], as the core concepts and heuristics that make the algorithm work are the same in the icosahedral setting.

2.5.1 Algorithm

Select absolute constants $\tilde{\varepsilon}, \varepsilon_0 > 0$. Take any $g = u(\alpha, \beta) \in \text{PU}(2)$ where $|\alpha| < \sqrt{1 - \varepsilon_0^2}$, and pick $\varepsilon < \tilde{\varepsilon}$. We wish to approximate g in $\mathfrak{d}$ using $\gamma \in \Gamma$ of the form

$$\gamma = \frac{1}{\eta^{k/2}} \begin{pmatrix} x_0 + x_1 i & x_2 + x_3 i \\ -x_2 + x_3 i & x_0 - x_1 i \end{pmatrix} \quad (2.5.1)$$

having k , the factorization length, minimized, and so we begin with $k = 0$. (We also have $x_0, x_1, x_2, x_3 \in \mathbb{Z}[\varphi]$.) In particular, the objective is to approximate g as $\gamma_1\gamma\gamma_2$ where $\gamma_1, \gamma_2 \in \Gamma$ approximate well-chosen diagonals computable using Section 2.4, and $\gamma \in \Gamma$ has factorization computable using Algorithm 2.3.2. We will see that γ is designed to have factorization typically shorter than that of γ_1 and γ_2 , giving rise to the desired improvement.

In order to apply Lemma 2.3.1 we need to have $\left| \frac{x_0+x_1i}{\eta^{k/2}} \right| = \sqrt{\frac{x_0^2+x_1^2}{\eta^k}}$ near $|\alpha|$ (that is, within ε). Because $\left| \frac{x_0+x_1i}{\eta^{k/2}} \right| + |\alpha| \geq |\alpha|$ which is fixed, it suffices to find candidate values for $x_0, x_1 \in \mathbb{Z}[\varphi]$ with $\left| \left| \frac{x_0+x_1i}{\eta^{k/2}} \right|^2 - |\alpha|^2 \right| < \varepsilon|\alpha|$, rewritten to

$$|x_0^2 + x_1^2 - |\alpha|^2 \eta^k| < \varepsilon|\alpha|\eta^k. \quad (2.5.2)$$

Viewing γ as an element of $\text{SU}(2)$, we also have $\det \gamma = 1$, i.e. $x_0^2 + x_1^2 + x_2^2 + x_3^2 = \eta^k$. As $x_\ell \in \mathbb{Z}[\varphi] \subset \mathbb{Q}(\varphi) \subset \mathbb{R}$, it follows that $\sigma_\pm(x_0^2 + x_1^2) + \sigma_\pm(x_2^2 + x_3^2) = \sigma_\pm(\eta^k)$, and so

$$\sigma_\pm(x_0^2 + x_1^2) \leq (\sigma_\pm \eta)^k. \quad (2.5.3)$$

Now, let $m = x_0^2 + x_1^2 \in \mathbb{Z}[\varphi]$. Considering $\mathbb{Z}[\varphi]$ as an integer lattice, we adapt (2.5.2) and (2.5.3) and seek to solve

$$|m - |\alpha|^2 \eta^k| < \varepsilon|\alpha|\eta^k \quad (2.5.4)$$

$$m \leq \eta^k \quad (2.5.5)$$

$$m^\bullet \leq (\eta^\bullet)^k \quad (2.5.6)$$

$$m, m^\bullet \geq 0 \quad (2.5.7)$$

which are convex constraints on m 's lattice components. Since this is an integer programming problem in two dimensions, we apply Lenstra's algorithm [L'83; P'83] to efficiently list all such lattice points m . For each m , using Theorem A.5, we attempt to write m as a sum of two squares; if possible, say $m = x_0^2 + x_1^2$, we then attempt to write $\tilde{m} = \eta^k - m$ as a sum of two squares. If possible, say $\tilde{m} = x_2^2 + x_3^2$, so we simply halt and return γ corresponding to (2.5.1). However, if $\tilde{m}$ may not be represented as a sum of two squares, we simply move on to the next value of m and try this process again. If this fails for all m arising from k , we increment k and run Lenstra's algorithm for the new inequalities.

Supposing we have halted and constructed γ , we compute δ_1 and δ_2 guaranteed by Lemma 2.3.1. These are efficiently approximable by Section 2.4 to γ_1 and γ_2 , respectively. Chaining together the three approximations as $\gamma_1\gamma\gamma_2$ gives the final desired approximation to g .

2.5.2 Analysis

We begin the analysis by establishing the τ -count and tightness of the approximation. In particular, $d(\gamma_1, \delta_1) < \varepsilon$ and $d(\gamma_2, \delta_2) < \varepsilon$ with factorization lengths of γ_ℓ each up to $(1 + o(1)) \log_{59}(1/\varepsilon^3)$. By Lemma 2.3.1, $d(\gamma, \delta_1 \gamma \delta_2) < C\varepsilon$. Therefore, $d(g, \gamma_1 \gamma \gamma_2) < (C + 2)\varepsilon$ (by the triangle inequality) and since γ has a factorization of τ -count up to $(\frac{1}{3} + o(1)) \log_{59}(1/\varepsilon^3)$, this constitutes a factorization of an element in g 's neighborhood of τ -count up to $(\frac{7}{3} + o(1)) \log_{59}(1/\varepsilon^3)$.

The efficiency of this algorithm—that is, that it runs in time $O(\text{poly log}(1/\varepsilon))$ —is because we expect to halt when $59^k \varepsilon \in O(1)$ (so only $k \approx \frac{1}{3} \log_{59}(1/\varepsilon^3)$ calls are expected), and only call polynomially-many polynomial-time subroutines. The dominant subroutines are calls to Lenstra's algorithm which as shown in [L'83] runs in time polynomial in the size of the constraints for any fixed dimension n . Indeed, here we have only $n = 2$ dimensions, $m = 6$ linear constraints (two per absolute value), and the largest value a in the constraints is η^k , so the runtime is polynomial in $nm \log a \in \Theta(k)$.

The reason we expect to halt when $59^k \varepsilon \in O(1)$ is that heuristically, we expect to halt when the area enclosed by (2.5.4)–(2.5.7) is $O(\text{poly log}(1/\varepsilon))$. Conveniently, the region is a rectangle since the vectors $\langle 1, \varphi \rangle$ and $\langle 1, \varphi^\bullet \rangle$ are orthogonal, so assuming in the limit that $\varepsilon \ll \min\left\{|\alpha|, \frac{1-|\alpha|^2}{|\alpha|}\right\}$ (so that (2.5.5) and the “bottom” inequality of (2.5.4) are redundant), we compute length-times-width of

$$\frac{2\varepsilon|\alpha|\eta^k}{\sqrt{1+\varphi^2}} \cdot \frac{(\eta^\bullet)^k}{\sqrt{1+(1-\varphi)^2}} = \frac{2|\alpha|}{\sqrt{5}} \cdot 59^k \varepsilon \in O(\text{poly log}(1/\varepsilon))$$

whence we find $k \in (1 + o(1)) \log_{59}(1/\varepsilon)$ as expected.

When attempting to write elements of $\mathcal{O}$ as a sum of two squares, we primarily rest on a belief, in the style of Cramér's conjecture and a conjecture of Sardari [S'19, (*)] that sums of squares are dense in $\mathbb{N}$. Seeking to analogize [S'19, (*)] in particular, we note that the operative aspect is that a dense cluster of lattice points will represent at least one sum of two squares, and that such a point thus will be found quickly through Lenstra's algorithm.

The significance of this result is to accomplish a factorization in $\text{PU}(2)$ in analogue to [S'21], but using a gate set with additional desirable properties beyond those enjoyed by the Clifford+ T gates.

2.6 Examples

The algorithm is implemented in Python and was used to produce the following examples of more “classical” single-qubit quantum gate generators. Recall H and T from the Clifford+ T gate set and pick precision $\varepsilon = 1/10^{10}$. We yield

$$T \approx \xi_0 \tau \xi_1 \tau \xi_2 \tau \xi_3 \tau \xi_4 \tau \xi_5 \tau \xi_6 \tau \xi_7 \tau \xi_8 \tau \xi_9 \tau \xi_{10} \tau \xi_{11} \tau \xi_{12} \tau \xi_{13} \tau \xi_{14} \tau \xi_{15} \tau \xi_{16} \tau \xi_{17} \tau \xi_{18} \tau \xi_{19}$$

$$\begin{aligned}
 &= (\sigma\rho\rho\rho\rho)\tau(\rho\sigma\sigma\rho\sigma\sigma)\tau(\sigma\rho\sigma\rho\sigma)\tau(\rho\sigma\rho\sigma\rho\sigma)\tau(\rho\sigma\sigma\rho \\
 &\quad \sigma\rho\sigma\sigma\rho\sigma\rho)\tau(\sigma\sigma\rho\sigma\rho\sigma\sigma\rho\sigma\rho)\tau(\rho)\tau(\sigma\rho\sigma\sigma\rho\sigma\rho\sigma\sigma\rho) \\
 &\quad \tau(\rho)\tau(\rho\sigma\sigma\rho\sigma\rho\sigma\sigma\rho)\tau(\rho\sigma\rho\sigma\rho\sigma\sigma)\tau(\sigma\sigma\rho\sigma\rho\sigma\sigma\rho) \\
 &\quad \tau(\sigma\rho\sigma\rho\sigma\sigma)\tau(\sigma\rho\sigma\sigma\rho\sigma\rho\sigma\sigma\rho)\tau(\sigma\rho\sigma\rho\sigma\sigma\rho\sigma)\tau(\sigma\rho \\
 &\quad \sigma\rho\sigma)\tau(\rho\sigma\sigma\rho\sigma\rho\sigma\sigma\rho\sigma\rho)\tau(\rho\sigma\sigma\rho\sigma\rho\sigma\sigma\rho\sigma)\tau(\sigma\rho\sigma)
 \end{aligned}$$

where $\xi_i \in \langle \sigma, \rho \rangle$ correspond to the parenthesized terms following; this achieves precision in d of $1.28/10^{10}$. The τ -count is 19; compare this with the predicted $\log_{59}(1/\varepsilon^3) \approx 16.9$. We remark that the discrepancy can be attributed to computational limitations: at key steps in the algorithm, we must run the Tonelli–Shanks algorithm for finding quadratic residues modulo some prime q , which has worst-case behavior $\Theta(q)$. Therefore it is only practical to abandon on instances where some prime factor is at least 10^6 , something that occurred more than 328 times before the above-stated approximation was found.

For the second example of H , the central element is determined to be the following:

$$\begin{aligned}
 \gamma &= \xi_0\tau\xi_1\tau\xi_2\tau\xi_3\tau\xi_4\tau\xi_5\tau\xi_6\tau\xi_7\tau\xi_8\tau\xi_9 \\
 &= (\rho\sigma\sigma\rho\sigma\rho\sigma\sigma\rho\sigma\rho)\tau(\rho\sigma\rho\sigma\rho\sigma\sigma\rho\sigma\sigma)\tau(\sigma\rho\sigma\rho\sigma\sigma)\tau \\
 &\quad (\rho\sigma\sigma\rho\sigma\rho\sigma\sigma\rho)\tau(\rho\sigma\sigma\rho)\tau(\sigma\rho\sigma\sigma\rho\sigma\rho)\tau(\sigma\rho\sigma\rho\sigma\sigma\rho \\
 &\quad \sigma\sigma)\tau(\rho\sigma\rho\sigma\sigma\rho\sigma\rho\sigma\sigma\rho)\tau(\rho\sigma\sigma\rho\sigma\rho\sigma)\tau(\rho\sigma\sigma\rho\sigma\rho)
 \end{aligned}$$

(ξ_i serves the same function as above). It has τ -count 9; compare this with the predicted $\frac{1}{3}\log_{59}(1/\varepsilon^3) \approx 5.6$. As before, the discrepancy can be attributed to non-vanishing of $o(1)$ factors for explicit ε and to having to abandon possibilities with very large prime factors. The outer diagonal elements are determined to be the following:

$$\begin{aligned}
 \gamma_1 &\approx \tilde{\xi}_0\tau\tilde{\xi}_1\tau\tilde{\xi}_2\tau\tilde{\xi}_3\tau\tilde{\xi}_4\tau\tilde{\xi}_5\tau\tilde{\xi}_6\tau\tilde{\xi}_7\tau\tilde{\xi}_8\tau\tilde{\xi}_9\tau\tilde{\xi}_{10}\tau\tilde{\xi}_{11}\tau\tilde{\xi}_{12}\tau\tilde{\xi}_{13}\tau\tilde{\xi}_{14}\tau\tilde{\xi}_{15}\tau\tilde{\xi}_{16}\tau\tilde{\xi}_{17}\tau\tilde{\xi}_{18} \\
 &= (\sigma\rho\sigma\rho)\tau(\rho\sigma\rho\sigma\sigma\rho\sigma\rho\sigma\sigma\rho\sigma)\tau(\rho\sigma\sigma\rho\sigma\sigma)\tau(\rho\sigma\sigma\rho\sigma \\
 &\quad \rho\sigma\sigma\rho\sigma\rho)\tau(\rho\sigma\rho\sigma\rho\sigma)\tau(\sigma\rho\sigma\rho)\tau(\sigma\sigma\rho\sigma)\tau(\sigma\sigma\rho\sigma\rho\sigma) \\
 &\quad \tau(\sigma\rho\sigma\rho\sigma\sigma\rho\sigma)\tau(\sigma\rho\sigma\rho\sigma\sigma\rho\sigma\rho)\tau(\rho\sigma\sigma\rho\sigma\rho\sigma)\tau(\rho\sigma \\
 &\quad \rho\sigma\sigma\rho)\tau(\sigma\sigma\rho\sigma)\tau(\sigma\sigma\rho\sigma\rho\sigma\sigma)\tau(\sigma\sigma)\tau(\sigma\rho\sigma\rho\sigma\sigma\rho\sigma\sigma) \\
 &\quad \tau(\sigma\rho\sigma\rho\sigma\sigma\rho\sigma)\tau(\sigma\sigma\rho\sigma\rho\sigma\sigma\rho\sigma\rho)\tau(\rho\sigma\rho\sigma\sigma\rho\sigma\sigma) \\
 \gamma_2 &\approx \hat{\xi}_0\tau\hat{\xi}_1\tau\hat{\xi}_2\tau\hat{\xi}_3\tau\hat{\xi}_4\tau\hat{\xi}_5\tau\hat{\xi}_6\tau\hat{\xi}_7\tau\hat{\xi}_8\tau\hat{\xi}_9\tau\hat{\xi}_{10}\tau\hat{\xi}_{11}\tau\hat{\xi}_{12}\tau\hat{\xi}_{13}\tau\hat{\xi}_{14}\tau\hat{\xi}_{15}\tau\hat{\xi}_{16}\tau\hat{\xi}_{17}\tau\hat{\xi}_{18} \\
 &= (\sigma\rho\sigma\rho\sigma\sigma\rho)\tau(\sigma\sigma\rho\sigma\rho\sigma\sigma\rho)\tau(\sigma\rho\sigma\rho\sigma\sigma\rho\sigma\sigma)\tau(\sigma\rho\sigma \\
 &\quad \sigma\rho\sigma\rho)\tau(\sigma\sigma\rho\sigma\rho\sigma)\tau(\rho\sigma\rho\sigma\sigma)\tau(\sigma\rho\sigma\sigma\rho\sigma\rho)\tau(\rho\sigma\rho\sigma \\
 &\quad \rho\sigma\sigma\rho\sigma\rho)\tau(\rho\sigma\rho\sigma\sigma\rho\sigma\sigma)\tau(\rho\sigma\sigma\rho\sigma)\tau(\rho\sigma\rho\sigma\rho\sigma\sigma\rho \\
 &\quad \sigma\rho)\tau(\rho\sigma\sigma\rho)\tau(\rho\sigma\sigma\rho\sigma\rho\sigma\sigma\rho)\tau(\rho\sigma\sigma\rho\sigma)\tau(\rho\sigma\rho\sigma\sigma\rho \\
 &\quad \sigma\rho\sigma\sigma)\tau(\rho\sigma\rho\sigma\sigma\rho\sigma\rho\sigma\sigma)\tau(\sigma\sigma\rho\sigma\rho\sigma\sigma\rho\sigma)\tau(\sigma\rho)\tau(\rho)
 \end{aligned}$$

($\tilde{\xi}_i$, $\hat{\xi}_i$ serves the same function as above). They both have τ -counts of 18, with 75 collective abandoned cases; compare this with the predicted $\log_{59}(1/\varepsilon^3) \approx 16.9$. Multiplying out $\gamma_1\gamma\gamma_2$ gives distance in d of $1.28/10^{10}$ (that this difference equals the previous one is a coincidence; they disagree at the third decimal place).

2.7 Upper-bounding fault-tolerant resources

We are grateful to one of the anonymous referees from *Quantum Information and Computation* for the remark which comprises this section, reproduced here for the reader's benefit with only light editing.

Note that it is possible to use existing results to implement super golden gates on common fault-tolerant quantum computers, as well as any gate specified by a matrix with entries in some number field. Therefore, it is theoretically possible to upper-bound the cost of icosahedral golden gates.

Consider a gate U with entries in some number field K . This gate can be synthesized in three steps using three well-known results.

First, reduce it to implementing an n -qubit unitary U' with entries in one of the fields $\mathbb{Q}(\zeta_8)$, $\mathbb{Q}(i)$, or $\mathbb{Q}$, by employing the Catalytic Embedding idea presented in [G'22; A⁺23].

Second, synthesize U' using an $(n + 2)$ -qubit unitary U'' with entries in the rings $\mathbb{Z}[\zeta_8, 1/2]$, $\mathbb{Z}[i, 1/2]$, or $\mathbb{Z}[1/2]$, utilizing [BRS'15, Theorem 18].

Third, employ exact synthesis techniques for unitaries [GS'13; AGR'20] with entries in $\mathbb{Z}[\zeta_8, 1/2]$, $\mathbb{Z}[i, 1/2]$, or $\mathbb{Z}[1/2]$ to find a Clifford+ T circuit for the unitary U'' . Additionally, synthesize the catalyst state up to the required accuracy using existing approximate synthesis techniques. The cost of the catalyst state can be neglected since it is reused.

Furthermore, the results presented in [B⁺20] can be utilized to numerically establish lower bounds on the T -cost of synthesizing U'' . This approach allows us to determine if icosahedral golden gates can become practical if multi-qubit synthesis algorithms are improved.

Chapter 3

A quantum algorithm for functions of multiple commuting Hermitian matrices

One fundamental matrix problem is the computation of matrix-valued functions. Specifically, recall the QET:

Problem (Quantum Eigenvalue Transformation (QET)): *Let A be a diagonalizable complex $N \times N$ matrix with eigenvalues $\{\lambda_k\} \subset \mathbb{D}$ and corresponding eigenvectors $\{v_k\}$. Given a function $f : \mathbb{C} \rightarrow \mathbb{C}$, a state $|v\rangle = \sum c_k |v_k\rangle$, and $\varepsilon > 0$, output (a scaling of) the state $f(A)|v\rangle := \sum c_k f(\lambda_j) |v_k\rangle$ to precision ε .*

(Here, $\mathbb{D}$ is the closed unit disk in $\mathbb{C}$.) QET may be implemented using the quantum signal processing algorithm (QSP) for a degree- D polynomial f and a Hermitian matrix A with constant ancilla overhead and $O(D)$ instances of a quantum circuit implementation of A , for instance as a *block-encoding*, which comes with some prescribed additional ancillae and an intrinsic constant known as the *subnormalization factor* (see Section 3.1.2); f can be provided as a classical blackbox,¹ and $|v\rangle$ can be given via some prepare oracle U of which we have only one copy.² As we will need $|v\rangle$'s prepare oracle exactly once and f 's blackbox is only accessed on the classical side, we are more concerned with minimizing the number of instances of A 's block-encoding. The QSP method requires that f is L^∞ -bounded.

QET is not well-understood for other classes of matrices. We propose the following generalization of QET that makes sense for several commuting matrices. One generic setting in which this would be of use is in the case of multiple commuting observ-

¹i.e., a blackbox from which we may query one input at a time. In principle f could take time polynomial in the problem parameters to evaluate, but since this is all done on the classical side we treat it as a constant-time blackbox.

²i.e., we initialize all registers to a tensor of $|0\rangle$'s and on U 's register apply it to obtain the input state $|v\rangle$, while U acts on other states unpredictably.

ables; it may be the case that one wishes to compute some function of several such observables, without having to compute individual measurements.

We point out that commuting matrices have a simultaneous eigenbasis, in particular commuting Hermitians have a simultaneous orthonormal eigenbasis. Recall MQET:

Problem (Multivariate Quantum Eigenvalue Transformation (MQET)):

Let $A = \{A_\ell\}_{0 \leq \ell \leq r}$ be a family of $r + 1$ pairwise commuting diagonalizable complex $N \times N$ matrices. Let $\{v_k\}$ be a mutual eigenbasis of A and let $\lambda_k = (\lambda_{\ell,k})_{0 \leq \ell \leq r} \in \mathbb{D}^{r+1}$ denote the vector of eigenvalues for the mutual eigenvector v_k . Given a function $f : \mathbb{C}^{r+1} \rightarrow \mathbb{C}$, a state $|v\rangle = \sum c_k |v_k\rangle$, and $\varepsilon > 0$, output (a scaling of) the state $f(A)|v\rangle := \sum c_k f(\lambda_k) |v_k\rangle$ to precision ε .

The matrices A_ℓ can be provided as block-encodings, the function $f : \mathbb{C}^{r+1} \rightarrow \mathbb{C}$ can be provided as a blackbox, and the state $|v\rangle$ can be provided by a single prepare oracle, all as in QET.

To motivate our method for MQET, first consider the problem of QET for normal matrices. Suppose we want to compute QET for a normal matrix M and function $f : \mathbb{C} \rightarrow \mathbb{C}$ that is L^∞ -bounded on the unit square. We retain the boundedness assumption from the case where M is Hermitian because it occurs as a special case of the problem we are currently considering. Let us also make the additional assumption that f is a polynomial. This is a mild assumption as any L^∞ -bounded function on the unit square can be approximated by L^∞ -bounded polynomials. Now, make the critical observation that any normal matrix M decomposes as $A + iB$, where A and B are commuting Hermitian matrices. We look for a decomposition

$$f(A + iB) \approx \sum_k P_k(A) Q_k(B)$$

for some polynomials P_k, Q_k . Then, we could implement $P_k(A)$ and $Q_k(B)$ using QSP and combine the products using the linear combination of unitaries primitive. The number of block-encoding instances required grows with the number of terms in the sum, and the subnormalization factor grows with the L^∞ norms of the polynomials P_k, Q_k appearing in the decomposition. A reasonable idea would be to let $P_k(x) = x^k$. Then, the number of terms in the sum is $1 + \deg f$, but the L^∞ norm of Q_k blows up rapidly. However, by letting $P_k(x) = T_k(x)$, the Chebyshev polynomial of degree k , we can obtain a decomposition where the number of terms in the sum is still $1 + \deg f$, but Q_k remains L^∞ -bounded. It is not clear in general how to compute the optimal decomposition of the polynomial f .

The method just outlined for QET of normal matrices inspires our method for MQET of commuting Hermitian matrices. Observe that $f(A + iB) = g(A, B)$ where $g : [-1, 1]^2 \rightarrow \mathbb{C}$ is a bivariate L^∞ -bounded polynomial. In general, we may compute MQET for a family $A = \{A_\ell\}_{0 \leq \ell \leq r}$ of pairwise commuting Hermitian matrices and

$f : [-1, 1]^{r+1} \rightarrow \mathbb{C}$ as follows. Compute a decomposition

$$f(A) \approx \sum_k P_k^{(0)}(A_0) P_k^{(1)}(A_1) \dots P_k^{(r)}(A_r)$$

where k is indexed by tuples $(a_i)_{0 \leq i < r}$ and $P_k^{(i)} = T_{a_i}$ for $0 \leq i < r$. The number of terms scales with the degree vector of f and the the polynomials appearing in the decomposition remain L^∞ -bounded.

We believe that this framework is promising and will have many applications beyond those provided in this work.

Let $1 + \deg f = D = 2^d$. We present the following:

Theorem 3.0.1. *Consider any pairwise commuting Hermitians $A = \{A_\ell\}_{0 \leq \ell \leq r}$. Given block-encodings $U_{A_\ell} \in \text{BE}_{\alpha, m}(A_\ell; \varepsilon)$, target precision $\varepsilon' > 0$, and an $(r + 1)$ -variate polynomial g approximating a given continuous L^∞ -bounded function $f : [-1, 1]^{r+1} \rightarrow \mathbb{C}$ as $\|f - g\|_{L^\infty[-1, 1]^{r+1}} \in O(\varepsilon')$, there is an algorithm (Algorithm 1) for implementing a block-encoding of $f(A)$ using $O(rD^{r+1})$ instances of the unitaries U_{A_ℓ} , $O(r(m + d))$ additional ancillae, and $O((D + 2)^r)$ subnormalization factor. Applying this matrix to $|v\rangle$ solves the given MQET instance with precision ε' .*

As a consequence, we may specialize to the case of bivariate functions to obtain the following result about QET for normal matrices:

Theorem 3.0.2. *Consider any normal matrix M . Given block-encoding $U_M \in \text{BE}_{\alpha, m}(M; \varepsilon)$, target precision $\varepsilon' > 0$, and a bivariate polynomial g approximating a given L^∞ -bounded function $f : \mathbb{C} \rightarrow \mathbb{C}$ as $\|f - g\|_{L^\infty([-1, 1] + i[-1, 1])} \in O(\varepsilon')$, there is an algorithm (Algorithm 2) for implementing a block-encoding of $f(M)$ using $O(D^2)$ instances of the unitary U_M , $O(m + d)$ additional ancillae, and $O(D)$ subnormalization factor. Applying this matrix to $|v\rangle$ solves the given QET instance with precision ε' .*

A further application is to perform a sort of ‘Hamiltonian’ simulation, where the ‘Hamiltonian’ operator is actually just any normal matrix (not necessarily Hermitian).

This framework fits into the recent development of quantum matrix functions, with hallmarks such as Grover’s search algorithm [G’96] and its generalization amplitude amplification, matrix inversion (Harrow, Hassidim, and Lloyd [HHL’08]), the linear combination of unitaries algorithm (Childs and Wiebe [CW’12], which we use as a subroutine), and quantum singular value transformation (Gilyén, Su, Low, and Wiebe [G’19]). Our results even more deeply rely on QSP of Low and Chuang [LC’16], and furnish approximate versions for both the specialized setting of normal operators as well as families of commuting operators, albeit with potentially steep costs and without so general a framework as offered by QSP and QSVT (see the *grand unification* of Martyn, Rossi, Tan, and Chuang [M’21]). This is similar in spirit to the recent work of Rossi and Chuang [RC’22] on *multivariable quantum signal processing* (M-QSP);

many of the deliverables obtained in this chapter also appear, possibly with different associated costs, in [RC'22].

We briefly describe some differences between this work and [RC'22]. M-QSP achieves extremely short block-encodings, using linearly-many instances of the block-encodings of the input matrices, but relying on a separate conjecture about the Fejér–Riesz theorem, on a relatively restricted class of functions, and requiring that the *block-encodings* commute—not just that the underlying, block-encoded matrices themselves commute. In dropping constraints on functions (and their polynomial approximations) and the structure of the block-encodings, we incur greater cost in block-encoding parameters and total instances used.

Our application here to normal matrices recalls the aims of the recent work of Takahira, Ohashi, Sogabe, and Usuda [T⁺'20; T⁺'21] which uses quadrature and Cauchy's integral formula to implement matrix functions, though [T⁺'20; T⁺'21]'s and Section 3.2's approaches are ultimately fundamentally incomparable. The idea in [T⁺'21] is for any holomorphic function and any diagonalizable matrix to approximate the integral via Cauchy's integral formula using an appropriately-chosen contour containing the matrix's eigenvalues. In order to compare resulting block-encoding parameters, one must somehow contrast the parameters chosen for a given quadrature against a polynomial approximation. As the philosophies of the algorithms are simply incompatible, any comparison cannot hold, and indeed there are instances where the quadrature approach is advantageous (such as when the eigenvalues in $\mathbb{D}$ are well-understood and lie in a convenient contour's interior), while there are also instances where the polynomial decomposition approach is advantageous (such as when the function is nearly an indicator function in $\mathbb{D}$, which admits poor L^∞ holomorphic approximation).

Lastly, the work of Guo, Mitarai, and Fujii [GMF'21] concerns a quantum operation distinct from matrix functions, instead known as nonlinear transformation of complex amplitudes (NTCA), in which the individual amplitudes of a state (with respect to a canonical basis) are transformed by a polynomial. Our approach gives rise to a solution to this problem for a less constrained class of functions than those considered in [GMF'21].

The paper is organized as follows. In Section 3.1 we describe present methods for implementing QET. This includes the framework of block-encoding and QSP for Hermitian matrices, as well as some newer literature on QET algorithms for other families of matrices. In Section 3.2 we describe the proposed method for normal matrices, revealed to be a special case of the algorithm for MQET (Section 3.4). Section 3.3 details an application of the algorithm for normal matrices to NTCA.




 Figure 3.1: 0- and 1-controlled U -gates.


Figure 3.2: An example of the new notation for multiply-controlled gates.

3.1 Preliminaries

3.1.1 Mathematics notation

Here, $\mathbb{D}$ is the *closed* unit disk in $\mathbb{C}$. Also, in this chapter, when we take a p -norm $\|\cdot\|_p$, we specifically restrict the domain to $[-1, 1]$, regardless of whether the function may be extended to more of $\mathbb{R}$ or $\mathbb{C}$. We use the shorthand “ $f : S \rightarrow \mathbb{C}$ is L^∞ -bounded” (for any domain S) to say that $\|f\|_{L^\infty S} \in O(1)$, independently of other ambient parameters (such as the number of qubits in the relevant circuit). Finally, T_k is the k th Chebyshev polynomial.

3.1.2 Circuit notation

When we provide examples of quantum circuits, we will predominantly use standard notation, and will introduce two new notations intended for visual ease.

The first is a compact notation for controls on ancilla registers. We say (consistent with the standard notation) that a single-wire control as in Figure 3.1 is *0-controlled* on the left and *1-controlled* on the right. Suppose we have a m -ancilla register A , indexed from the top as $0, 1, \dots, m-1$, and that U acts on the register B with a 0- or 1-control from each wire of A . We will compactly identify this with a number written in binary: the j -indexed wire is the j th place value, and it is 0 if the corresponding control is 0-controlled, and 1 if 1-controlled. If the resulting binary value is b , then this will be indicated by grouping all m wires and putting $|b\rangle$ in a rounded control box. For instance, we have the equivalent circuits depicted in Figure 3.2.

We also introduce notation for multi-wire gates that do not act on some wires that




Figure 3.3: An example of the new notation for gates that don't incorporate a specific register.

they are “on top of.” For instance, if there are three registers A , B , and C , which for broader clarity must be provided in that order, but a specific gate U only acts on A and C , then we will depict a “break” in the wire(s) for B . See for an example Figure 3.3, which also depicts a formulation via SWAP gates which circumvents this notation, assuming that (say) A and B comprise the same number of qubits. Another example is Figure 3.4.

3.1.3 Block-encodings

Computations with general matrices are most easily performed in the block-encoding model, which we summarize here.

Intuitively, a block-encoding of a not necessarily unitary matrix A is a large unitary matrix

$$U_A = \begin{pmatrix} \frac{1}{\alpha} A & \star \\ \star & \star \end{pmatrix}$$

where the $\star$'s represent arbitrary matrices (not necessarily all square). More specifically, given a n -qubit matrix A , we say that a $(n + a)$ -qubit matrix U_A is an (α, a, ε) -block-encoding of A if

$$\|A - \alpha(|0^a\rangle \otimes I_n) U_A (|0^a\rangle \otimes I_n)\| < \varepsilon.$$

This is also written as $U_A \in \text{BE}_{\alpha,a}(A; \varepsilon)$. α is known as the *subnormalization factor* and is related to the number of repetitions needed to obtain the correct state after applying a block-encoding (since one wishes to measure the first register to be $|0^a\rangle$).

To highlight its applicability, we note that if each matrix entry may be efficiently computed then so too can a block-encoding, by the universality of the Toffoli gate in classical circuitry and the Toffoli+ H gates in quantum circuitry [S'02; A'03]. Further, $(S, n+1, 0)$ -block-encodings are implementable using specialized oracles which identify the locations of nonzero entries in A , in the event that A has at most S nonzero entries in each row and in each column, and in the case that (e.g.) A has exactly S bands the same oracles give $(S, \log S + 1, 0)$ -block-encodings [L'22, §6.3]. However, for our purposes, we will fully blackbox the implementation details of our given block-encodings, as matrix-specific considerations quite often inform the choice of block-encoding model.

Various functions are relatively easy (but not free) to compute with block-encoded matrices. We highlight two, which we will use later on.




Figure 3.4: [G'+19]'s multiplication circuit.

Proposition 3.1.1 (block-encoding of product [G'+19, Lemma 53]). *Given n -qubit matrices A, B and $U_A \in \text{BE}_{\alpha,a}(A; \varepsilon_A)$, $U_B \in \text{BE}_{\beta,b}(B; \varepsilon_B)$, then $(I_a \otimes U_A)(I_b \otimes U_B) \in \text{BE}_{\alpha\beta, a+b}(AB; \alpha\varepsilon_B + \beta\varepsilon_A)$, where I_a, I_b act on the ancilla registers of A, B respectively.¹*

See also Figure 3.4 for a depiction of this circuit.

Proposition 3.1.2 (block-encoding linear combinations of unitaries (LCU)). *Given access to a prepare oracle V satisfying*

$$V|0^n\rangle = \frac{1}{\|\mathbf{a}\|_1} \sum_{0 \leq j < K} \sqrt{a_j} |j\rangle$$

for $K = 2^k$, and $U_0, \dots, U_{K-1} \in \text{U}(N)$ arranged in a select oracle

$$U := \sum_{0 \leq j < K} |j\rangle\langle j| \otimes U_j,$$

then we have that

$$(V^\dagger \otimes I_n)U(V \otimes I_n) \in \text{BE}_{\|\mathbf{a}\|_1, k} \left(\sum_{0 \leq j < K} a_j U_j; 0 \right).$$

If U and V are implemented so that actually $(V^\dagger \otimes I_n)U(V \otimes I_n) \in \text{BE}_{\|\mathbf{a}\|_1, k} \left(\sum_{0 \leq j < K} a_j U_j; \varepsilon \right)$

then we will write $\text{LCU}_\varepsilon((U_i)_{0 \leq i < K}, \mathbf{a})$ for the circuit.

Using quantum signal processing (QSP) (for Hermitian matrices) or quantum singular value transform many other matrix functions can be implemented using polynomial approximations and block-encodings, cf. [L'22, Chapters 7–8]. See in particular the algorithm named *QET for Hermitian matrices* [L'22, §7.5].

¹That is, suppose the work register is labeled 3, and we have an a -qubit register 1 and a b -qubit register 2. Then we arrange U_A to be acting on 1 and 3, and U_B to be acting on 2 and 3.

3.2 QET for normal matrices

Recall the notion of block-encoding non-unitary matrices into larger, unitary matrices which may then be implemented as quantum circuit primitives. If $U_A \in \text{BE}_{\alpha, m}(A; \varepsilon)$ is such a block-encoding and $f : \mathbb{C}^{N \times N} \rightarrow \mathbb{C}^{N \times N}$ is a matrix function, one seeks to implement a circuit for $U_{f(A)} \in \text{BE}_{\alpha', m'}(f(A); \varepsilon')$ using only the standard quantum primitives as well as U_A (and $U_A^\dagger$), with as few calls to U_A and $U_A^\dagger$ as possible, m' as small as possible (minimal number of ancillae), and α' as small as possible (maximal success probability). For A Hermitian, this problem is discussed extensively in [L'22, Chapter 7]. This analysis crucially depends on the A -block being Hermitian in order to construct an alternating sequence of $\mathbb{C}$ -spaces, leaving the question of block-encoding matrix functions for *normal* matrices unresolved by these techniques. In this section we put forth an algorithm to implement quantum circuits for a given matrix function f of a given normal matrix M . We assume that it is feasible to produce an arbitrarily good approximation to f by a complex polynomial $g(z) = G(x, y)$ where $z = x + iy$ (see Section B.1), into which we may plug in $z = M$.¹ Classically computing phase factors for QSP is a feasible task, which has been studied in Haah [H'19] and Dong, Meng, Whaley, and Lin [D+'21]. However, at no point do we assume that computing M 's eigenvalues is tractable.

For clarity, we recall that a normal matrix is one which commutes with its Hermitian adjoint, and normal matrices are precisely those which are unitarily diagonalizable.

We defer a treatment of the most general case to Section 3.4.

3.2.1 Algorithm

We describe an algorithm for obtaining a block-encoding of $f(M)$ given a block-encoding $U_M \in \text{BE}_{\alpha, m}(M; \varepsilon)$ of a normal matrix M and a function $f : \mathbb{D} \rightarrow \mathbb{D}$. We let f be (γ, k) -extensible (Property B.1.7) for $\gamma, k \in O(1)$. This is a technical condition, defined in Section B.1, which just ensures that f is well-approximable. Let us assume that f is well-approximated on $\mathbb{D}$ by a polynomial g .

Consider the decomposition

$$M = A + iB \tag{3.2.1}$$

where A and B are commuting Hermitians with the same dimensions as M . It is not hard to obtain block-encodings of A and B . In fact, all three are simultaneously diagonalizable, where if $M = U\mathcal{D}U^\dagger$ then $A = U(\text{re } \mathcal{D})U^\dagger$ and $B = U(\text{im } \mathcal{D})U^\dagger$. (It is this observation that powers Section B.2.) In light of this, we immediately see how to

¹The technical formulation of this condition is that f on $\mathbb{D}$ satisfies Property B.1.7, (γ, k) -extensibility for (say) $\gamma, k \in O(1)$.




Figure 3.5: Circuit for block-encoding of $f(M)$. See Section 3.4.1 for the algorithm, and Algorithm 1 for a pseudocode description. V comes from LCU (Proposition 3.1.2), and the remaining notation on the top wire is introduced in Section 3.1.2. We may let P_k be T_k , as discussed in Section 3.2.3.

implement block-encodings of A and B :

$$\begin{aligned} U_A &= \text{LCU}_\varepsilon\left(\left(U_M, U_M^\dagger\right), (1/2, 1/2)\right), & U_A &\in \text{BE}_{\alpha, m+1}(A; \varepsilon) \\ U_B &= \text{LCU}_\varepsilon\left(\left(U_M, U_M^\dagger\right), (-i/2, i/2)\right), & U_B &\in \text{BE}_{\alpha, m+1}(B; \varepsilon). \end{aligned}$$

As suggested in at the beginning of this chapter, we can expand $g(x + iy) = G(x, y)$ in terms of polynomials in x and y so that

$$g(M) = G(A, B) = \sum_{0 \leq k < D} P_k(A) Q_k(B) \quad (3.2.2)$$

for some $P_k, Q_k \in \mathbb{C}[z]$ of degree less than D (where g is also of degree less than D ; suppose for convenience that D is a power of 2). In the following, we describe how to use this decomposition, along with QSP, to obtain a block-encoding of $g(M)$, and how the parameters of the block-encoding depend on the decomposition we use. The implementation is also summarized in Figure 3.5. We will find that the subnormalization factor of the block-encoding of $g(M)$ is proportional to $\|\beta\|_1$, where $\beta \in \mathbb{R}_{\geq 0}^D$ is defined by $\beta_k := \|P_k\|_\infty \|Q_k\|_\infty$. Our method for optimizing $\|\beta\|_1$ is to let $P_k = T_k$ and

$$Q_k(y) = \frac{2^{1-\delta_{k,0}}}{\pi} \int_{-1}^1 g(x + iy) T_k(x) \frac{dx}{\sqrt{1-x^2}},$$

so that

$$g(x + iy) = \sum_{0 \leq k < D} T_k(x) Q_k(y) \quad (3.2.3)$$

In Proposition 3.2.4, we show that this guarantees $\|\beta\|_1 \leq 2D$. Our entire algorithm is summarized in Algorithm 1.

Elaborating on how to use the decomposition in (3.2.2) to implement a block-encoding of $g(M)$, first, normalize to $\tilde{P}_k = \frac{P_k}{\|P_k\|_\infty}$ and $\tilde{Q}_k = \frac{Q_k}{\|Q_k\|_\infty}$. Figure 3.5 describes how to obtain $g(M)$ as quantum circuitry, we (classically) compute the phase factors Φ_k, Ψ_k for the polynomials $x \mapsto \tilde{P}_k(x), \tilde{Q}_k(x)$ (respectively), e.g. via [D⁺21].¹ Then, we use QSP to compute $U_{\tilde{P}_k(A)} \in \text{BE}_{\alpha, m+2}(\tilde{P}_k(A); \varepsilon)$ and $U_{\tilde{Q}_k(B)} \in \text{BE}_{\alpha, m+2}(\tilde{Q}_k(B); \varepsilon)$, and compute via [G⁺19, Lemma 53] the block-encoding of the product $U_{\tilde{P}_k(A)\tilde{Q}_k(B)} \in \text{BE}_{\alpha^2, 2m+4}(\tilde{P}_k(A)\tilde{Q}_k(B); 2\alpha\varepsilon)$. Finally,

$$U_{g(M)} = \text{LCU}_{2\alpha\varepsilon} \left((\tilde{P}_k(A)\tilde{Q}_k(B))_{0 \leq k < D}, \beta \right), \quad U_{g(M)} \in \text{BE}_{\alpha^2 \|\beta\|_1, 2m+4+d}(g(M); 2\alpha\varepsilon).$$

Algorithm 1 block-encoding of $f(M)$

Input: $U_M \in \text{BE}_{\alpha, m}(M; \varepsilon)$ for $M \in \mathbb{C}^{N \times N}$ normal; $g \in \mathbb{C}[z, \bar{z}]$ of degree less than D such that $\|f - g\|_\infty$ is small for $f : \mathbb{D} \rightarrow \mathbb{D}$

Output: $U_{f(M)} \in \text{BE}_{\alpha', m'}(f(M); \varepsilon')$

- 1: $U_A \leftarrow \text{LCU} \left((U_M, U_M^\dagger), (1/2, 1/2) \right)$ ▷ $U_A \in \text{BE}_{\alpha, m+1}(A; \varepsilon)$
 - 2: $U_B \leftarrow \text{LCU} \left((U_M, U_M^\dagger), (-i/2, i/2) \right)$ ▷ $U_B \in \text{BE}_{\alpha, m+1}(B; \varepsilon)$
 - 3: **for** $0 \leq k < D$ **do**
 - 4: $P_k \leftarrow T_k$ ▷ A Chebyshev polynomial
 - 5: $Q_k(y) \leftarrow \frac{2^{1-\delta_{k,0}}}{\pi} \int_{-1}^1 g(x + iy) T_k(x) \frac{dx}{\sqrt{1-x^2}}$ ▷ These satisfy (3.2.3)
 - 6: $\beta_k \leftarrow \|P_k\|_\infty \|Q_k\|_\infty$
 - 7: $(\tilde{P}_k, \tilde{Q}_k) \leftarrow (P_k / \|P_k\|_\infty, Q_k / \|Q_k\|_\infty)$
 - 8: **Compute** $U_{\tilde{P}_k(A)} \in \text{BE}_{\alpha, m+2}(\tilde{P}_k(A); \varepsilon)$ ▷ Using QSP
 - 9: **Compute** $U_{\tilde{Q}_k(B)} \in \text{BE}_{\alpha, m+2}(\tilde{Q}_k(B); \varepsilon)$ ▷ Using QSP
 - 10: **Compute** $S_k \in \text{BE}_{\alpha^2, 2m+4}(\tilde{P}_k \tilde{Q}_k; 2\alpha\varepsilon)$
 - 11: **end for**
 - 12: **return** $\text{LCU}((S_k)_{0 \leq k < D}, \beta)$
-

3.2.2 Analysis of the algorithm

In the previous section, we kept a running tally of the ancilla count; the final count was $2m+4+d$. We also track precision— $2\alpha\varepsilon$ —and subnormalization factor— $\alpha^2 \|\beta\|_1$ —above.

¹For clarity, we assume that one's implementation of phase factor computation is exact. However, error here is relatively insignificant: if each phase factor of Φ_k is off by ε' additively then we may crudely bound that the block-encoding of $\tilde{P}_k(A)$ is off by $D\varepsilon'$, but since this computation is done classically it is sensible to assert that its error vanishes relative to the error on the quantum computer.

We can track the number of applications of U_M (and its adjoint) as follows. U_A , $U_A^\dagger$, U_B , and $U_B^\dagger$ each requires two such gates; applying QSP for a degree less-than- D polynomial multiplies the number of instances by at most $2D$. Computing products adds this cost. Running LCU adds these costs; the final LCU call is on D unitaries. Thus there are at most $4(D-1)D \in O(D^2)$. Within the framework of decomposing a function by a polynomial of degree less than D as in (3.2.2) for which each monomial of the form $x^k y^{D-1-k}$ has nonzero coefficient (as is generically the case), the decomposition described using integration against the Chebyshevs uses as few U_M s as possible; see Section B.3.

Observe that all nontrivial costs are incurred precisely when we deal with the many polynomials arising from the decomposition (3.2.1), by computing their paired products and then add using LCU (all of these inducing possibly large multiples of the base cost for each polynomial coming from QSP).

3.2.3 Subnormalization factor

Proposition 3.2.4. *If $G(x, y)$ is degree less than D in both x and y and satisfies $\|G\|_{L^\infty[-1,1]^2} \leq 1$, then there are efficiently-computable degree-less-than- D polynomials P_k and Q_k for $0 \leq k < D$ such that $G(x, y) = \sum_{0 \leq k < D} P_k(x)Q_k(y)$ with $\|P_k\|_\infty \leq 1$ and $\|Q_k\|_\infty \leq 2$.*

Proof. Let $P_k := T_k$ and

$$\tilde{Q}_k(y) := \int_{-1}^1 G(x, y) T_k(x) \frac{dx}{\sqrt{1-x^2}}. \quad (3.2.5)$$

Let $Q_0 = \frac{1}{\pi} \tilde{Q}_0$ and $Q_k = \frac{2}{\pi} \tilde{Q}_k$ for $k > 0$. It is well-known that Chebyshev polynomials satisfy $\|T_k\|_\infty = 1$. Also, for $y \in [-1, 1]$,

$$|Q_k(y)| \leq \frac{2}{\pi} \int_{-1}^1 |G(x, y)| |T_k(x)| \frac{dx}{\sqrt{1-x^2}} \leq \frac{2}{\pi} \int_{-1}^1 \frac{dx}{\sqrt{1-x^2}} = 2.$$

Chebyshev polynomials form an orthogonal set with respect to the measure $(1-x^2)^{-1} dx$. Further, T_k has norm π for $k = 0$ and $\pi/2$ otherwise. It follows that by construction $G(x, y) = \sum_{0 \leq k < D} P_k(x)Q_k(y)$.

To compute $\tilde{Q}_k$, we use the techniques of Section B.1 and Section B.4: first write G in the form (B.1.2). and evaluate (3.2.5) as a sum using Proposition B.4.4. ■

The algorithm described in Section 3.2.1 implements a block-encoding of $f(M)$ with subnormalization factor bounded by $2\alpha D$ assuming that $|f(x + iy)| \leq 1$ for $x, y \in [-1, 1]$. Because of the L^∞ constraint on f it is nontrivial to implement QET for several interesting functions. However, the approximation $\tilde{f}$ is not required to be holomorphic. This allows for a certain freedom. For example, if we wish to implement QET

for the function $f(z) = z^n$, assuming that the eigenvalues are contained on the unit disk, then we actually need to choose a polynomial approximation $\tilde{f}$ with the following properties:

- $|\tilde{f}(z)| \leq 1$ for $\operatorname{re} z, \operatorname{im} z \in [-1, 1]$
- $|f(z) - \tilde{f}(z)| < \epsilon$ for $|z| \leq 1$, for some approximation parameter ϵ

Because the function $f(z)$ blows up outside of the unit disk, it seems these two conditions are in contention. Perhaps it is possible then that these two conditions can be simultaneously satisfied by a bivariate polynomial with reasonable degree. There are similar questions, for example, for functions such as $f(z) = \kappa/z$ on the domain $|z| > \kappa$.

3.3 On nonlinear transformation of complex amplitudes

We now demonstrate that the methods of Section 3.2 naturally provide an extension to the *nonlinear transformation of complex amplitudes* as developed in [GMF'21].

Problem (Nonlinear Transformation of Complex Amplitudes (NTCA)):

Let $F : \mathbb{D} \rightarrow \mathbb{C}$ be any function. Given a state $|v\rangle = \sum c_k |k\rangle$ in the computational basis and $\epsilon > 0$, output (a scaling of) the state $\sum F(c_k) |k\rangle$ to precision ϵ .

In the setting of [GMF'21], we have access to a (unitary) state preparation oracle U , as well as its adjoint and controlled versions, such that $U|0\rangle^{\otimes n} = |v\rangle$. [GMF'21] addresses this problem with $\epsilon = 0$ for

$$F(z) = F_1(z) = P(\operatorname{re} z) + Q(\operatorname{im} z) \quad (3.3.1)$$

for arbitrary continuous functions $P, Q : [-1, 1] \rightarrow \mathbb{C}$. Specifically, if P and Q can be approximated up to error $\frac{\epsilon}{4N}$ by degree- D polynomials $\tilde{P}$ and $\tilde{Q}$, respectively, such that $\tilde{F}_1(z) := \tilde{P}(\operatorname{re} z) + \tilde{Q}(\operatorname{im} z)$ approximates F_1 , and $\gamma = \sup_{x \in [-1, 1]} \max\{|P(x)|, |Q(x)|\}$, then¹

$\frac{1}{\sqrt{c}} \sum_{0 \leq k < N} F_1(c_k) |k\rangle$ can be approximated to L^2 error ϵ/N , using in expectation

$$O\left(D\gamma \sqrt{\frac{N}{\sum_{0 \leq k < N} |\tilde{F}_1(c_k)|^2}}\right)$$

¹Suppose the normalization factor is $c := \sum_{0 \leq k < N} |F(c_k)|^2$.

controlled U and $U^\dagger$ gates.

The algorithm in [GMF'21] uses the technique of *block-encoding of amplitudes*. In particular, let $A = (I \otimes W^\dagger) \tilde{G} (I \otimes W)$ and $B = (I \otimes W^\dagger) \tilde{G}' (I \otimes W)$, where W , $\tilde{G}$, and $\tilde{G}'$ are defined as in [GMF'21, Figures 1–3]; for completeness, we also include their definitions and a brief explanation in Section B.5. Then A and B are block-encodings of matrices which map $|k\rangle$ to $(\operatorname{re} c_k)|k\rangle$ and $(\operatorname{im} c_k)|k\rangle$, respectively, that is, they are block-encodings of diagonal matrices containing the data of the original amplitudes.

Applying the algorithm from Section 3.2.1 to A and B , we may obtain a block-encoding of a matrix that maps $|k\rangle$ to $F_2(c_k)|k\rangle$, for all F_2 which are (γ, k) -extensible, which is defined in Property B.1.7—allowing for an additional kind of function as those specified in (3.3.1). Thus, this resolves NTCA with expected

$$O\left(D^2 \gamma \sqrt{\frac{N}{\sum_{0 \leq k < N} |F_2(c_k)|^2}}\right)$$

controlled U and $U^\dagger$ operations. A concrete instance of a newly-attainable function is

$$f(z) = \frac{5}{1024}(z + \bar{z})^6(z - \bar{z})^4 + \frac{17}{512}(z + \bar{z})(z - \bar{z})^8,$$

or $f(x + iy) = 5x^6y^4 + 17xy^8$.

3.4 MQET for several commuting Hermitian matrices

We generalize the ideas in Section 3.2 following the observation that QET of normal matrices is a special case of MQET of two commuting Hermitian matrices.

Consider a continuous function $f : [-1, 1]^{r+1} \rightarrow \mathbb{D}$ for which we select a polynomial approximation g , perhaps via the techniques of Section B.1. Suppose the maximum total degree of any monomial in g is less than D .

Generalizing the method from Section 3.2, we write down an expansion of g analogous to (3.2.2). Suppose we find polynomials $P_s^{(k)}, Q_s \in \mathbb{C}[z]$ for s ranging in $[[D]]^r$ and k ranging from 0 to $r - 1$ satisfying

$$g(x) = \sum_{s \in [[D]]^r} Q_s(x_r) \prod_{0 \leq k < r} P_s^{(k)}(x_k). \quad (3.4.1)$$

Remark 3.4.2. This sum has a factor of D fewer terms than the number of monomials. This is crucial because we avoid applying LCU to each of the monomials, which causes the subnormalization factor to needlessly increase.




Figure 3.6: Circuit for block-encoding of $f(A_0, \dots, A_r)$. Here we let $0 \leq s < D^r$ denote the integer given in base D by s . See Section 3.4.1 for the algorithm, and Algorithm 2 for a pseudocode description. V comes from LCU (Proposition 3.1.2), and the remaining notation on the top wire is introduced in Section 3.1.2. We may let $P_s^{(k)}$ be T_{s_k} , as discussed in Section 3.4.3.

In the following subsection, we describe in detail how to use this decomposition, along with QSP, to obtain a block-encoding of $g(A)$, and how the parameters of the block-encoding depend on the decomposition we use. The implementation is summarized in Figure 3.6. We will find that the subnormalization factor of the block-encoding is proportional to $\|\beta\|_1$, where $\beta \in \mathbb{R}_{\geq 0}^{[D+1]^r}$ is defined by $\beta_s := \|Q_s\|_\infty \prod_{0 \leq k < r} \|P_s^{(k)}\|_\infty$. We

optimize $\|\beta\|_1$ by letting $P_s^{(k)} = T_{s_k}$ (the s_k th Chebyshev polynomial) and defining Q_s by an orthogonal projection formula proposed in Section 3.4.3; we show that this ensures $\|\beta\|_1 \leq 2D^r$. Our entire algorithm is summarized in Algorithm 2.

3.4.1 Algorithm

We describe an algorithm for obtaining a block-encoding of $g(A)$ given block-encodings $U \in \prod_{0 \leq k \leq r} \text{BE}_{\alpha_k, m_k}(A_k; \varepsilon_k)$, as well a decomposition of g as in (3.4.1).

For convenience, we introduce some notation. Let

$$T(s, A) := Q_s(A_r) \prod_{0 \leq k < r} P_s^{(k)}(A_k),$$

so that $g(A) = \sum_{s \in [[D]]^r} T(s, A)$. Let also

$$\alpha := \prod_{0 \leq k \leq r} \alpha_k, \quad m := \sum_{0 \leq k \leq r} m_k, \quad \varepsilon := \sum_{0 \leq k \leq r} \alpha_k \varepsilon_k. \quad (3.4.3)$$

Let $\beta \in \mathbb{R}_{\geq 0}^{[[D+1]]^r}$ be the vector where $\beta_s = \|Q_s\|_\infty \prod_{0 \leq k < r} \|P_s^{(k)}\|_\infty$. Normalize to $\tilde{P}_s^{(k)} := \frac{P_s^{(k)}}{\|P_s^{(k)}\|_\infty}$ and $\tilde{Q}_s := \frac{Q_s}{\|Q_s\|_\infty}$. To obtain $g(A)$ as quantum circuitry, we (classically) compute the phase factors $\Phi_s^{(k)}, \Psi_s$ for the polynomials $x \mapsto \tilde{P}_s^{(k)}(x), \tilde{Q}_s(x)$ (respectively), e.g. via [D⁺21]. Then, we use QSP to compute $U_{\tilde{P}_s^{(k)}(A_k)} \in \text{BE}_{\alpha_k, m+2}(\tilde{P}_s^{(k)}(A_k); \varepsilon_k)$ and $U_{\tilde{Q}_s(A_k)} \in \text{BE}_{\alpha_k, m+2}(\tilde{Q}_s(A_k); \varepsilon_k)$, and compute via [G⁺19, Lemma 53] the block-encoding of the product $U_{T(s, A)} \in \text{BE}_{\alpha, m+4}(T(s, A); \varepsilon)$ (called $U_{s, A}$ in Algorithm 2). Finally,

$$U_{g(A)} = \text{LCU}_\varepsilon\left((T(s, A))_{s \in [[D]]^r}, \beta\right) \in \text{BE}_{\alpha \|\beta\|_1, m+4+d}(g(A); \varepsilon).$$

3.4.2 Analysis of the algorithm

Recall that α , m , and ε are defined in (3.4.3). In the previous section, we kept a running tally of the ancilla count; the final tally was $2m + 4 + d$. We also track precision— ε —and subnormalization factor— $\alpha \|\beta\|$ —above.

We can track the number of applications of U_{A_ℓ} and $U_{A_\ell}^\dagger$. Applying QSP for degree at-most- D polynomials multiplies this quantity by at most $2D$. Computing products and running LCU adds; the final LCU call is on D^r unitaries. Thus there are at most $2(D-1)rD^r \in O(rD^{r+1})$ instances.

3.4.3 Subnormalization factor

By enforcing $P_s^{(k)} := T_{s_k}$ (the s_k th Chebyshev polynomial) and, as in Proposition 3.2.4,

$$Q_s(x_r) := \prod_{1 \leq k \leq r} \frac{2 - \delta_{s_k, 0}}{\pi} \int_{-1}^1 g(x) \prod_{1 \leq k \leq r} T_{s_k}(x_k) \frac{dx_k}{\sqrt{1 - x_k^2}},$$

we see that $\|P_s^{(k)}\|_\infty = 1$ and $\|Q_s\|_\infty \leq 2$. This leaves us with $\beta_s \leq 2$. (Also as before, again using Proposition B.4.4 we can efficiently compute Q_s given g .) Thus, $\|\beta\| = \sum_{s \in [[D]]^r} 2^{z(s)}$ where $z(s)$ is the number of zeroes appearing in the vector s . This rewrites to $\sum_{0 \leq \ell \leq r} 2^\ell \binom{r}{\ell} D^{r-\ell} = (D+2)^r$, the subnormalization factor asserted in the beginning.

Algorithm 2 MQET for several commuting Hermitian matrices

Input: $A_k \in \mathbb{C}^{N \times N}$ Hermitian and $U_k \in \text{BE}_{\alpha_k, m}(A_k; \varepsilon_k)$ for $0 \leq k \leq r$; $[A_{k_1}, A_{k_2}] = 0$ for all $0 \leq k_1, k_2 \leq r$; $f : [-1, 1]^{r+1} \rightarrow \mathbb{D}$; $g \in \mathbb{C}[x_0, \dots, x_r]$ with degree less than D such that $\|f - g\|_\infty$ is small

Output: $U_{f(A)} \in \text{BE}_{\alpha', m'}(f(A), \varepsilon')$

- 1: **for** $s \in [[D]]^r$ **do**
 - 2: $(s_0, \dots, s_{r-1}) \leftarrow s$
 - 3: $P_s^{(k)} \leftarrow T_{s_k}$ ▷ for $0 \leq k < r$
 - 4: $Q_s(x_r) \leftarrow \prod_{0 \leq k < r} \frac{2 - \delta_{s_k, 0}}{\pi} \int_{-1}^1 g(\mathbf{x}) \prod_{0 \leq k < r} T_{s_k}(x_k) \frac{dx_k}{\sqrt{1 - x_k^2}}$
 - 5: ▷ These satisfy $g(\mathbf{x}) = \sum_{s \in [[D]]^r} Q_s(x_r) \prod_{0 \leq k < r} P_k(x_k)$
 - 6: $\beta_s \leftarrow \|Q_s\|_\infty \prod_{0 \leq k < r} \|P_s^{(k)}\|_\infty$
 - 7: $\tilde{P}_s^{(k)} \leftarrow \frac{P_s^{(k)}}{\|P_s^{(k)}\|_\infty}$ ▷ for $0 \leq k < r$
 - 8: $\tilde{Q}_s \leftarrow \frac{Q_s}{\|Q_s\|_\infty}$
 - 9: Compute $U_{\tilde{P}_s^{(k)}(A_k)} \in \text{BE}_{\alpha_k, m+2}(\tilde{P}_s^{(k)}(A_k); \varepsilon_k)$ ▷ for $0 \leq k < r$, using QSP
 - 10: Compute $U_{\tilde{Q}_s(A_r)} \in \text{BE}_{\alpha_r, m+2}(\tilde{Q}_s(A_r); \varepsilon_r)$ ▷ Using QSP
 - 11: Compute $U_{s,A} \in \text{BE}_{\alpha, m+4} \left(Q_s(A_r) \prod_{0 \leq k < r} P_s^{(k)}(A_k); \varepsilon \right)$
 - 12: **end for**
 - 13: **return** $\text{LCU}_\varepsilon \left((U_{s,A})_{s \in [[D]]^r}, \beta \right)$
-

Chapter 4

On quantum to classical comparison for Davies generators

For a Hamiltonian H and inverse temperature parameter $\beta \geq 0$, the goal of quantum Gibbs sampling is to prepare the quantum Gibbs state ρ_β . One of the leading approaches for quantum Gibbs sampling is based on quantum Markov processes, with the Davies Lindbladian generator being a prominent model (see Definition 4.0.1). Despite significant recent progress, our understanding of the Davies generator remains incomplete. Existing works are able to prove mixing for commuting Hamiltonians [T⁺14; KB⁺16; K⁺24], in the high temperature regime [RFA⁺25], or under strong assumptions on either the Hamiltonian or the Davies generator [B⁺21; RW⁺24; CB⁺21; C⁺24]. Yet, there is a lack of tools for proving mixing outside these settings. This stands in stark contrast to the classical setting, where recent advances have yielded sharp convergence bounds across the full range of temperatures for several important models [ALO⁺20; CLV⁺21; A⁺22; CE⁺22; C⁺22].

A natural question arises: to what extent can techniques from classical Markov chains be leveraged to understand their quantum counterparts? In [T⁺13], it is observed that when H has simple spectrum, the Davies Lindbladian contains a canonical embedded classical Markov chain. Specifically, its action on observables that are diagonal in the Hamiltonian eigenbasis is isomorphic to that of a classical Markov generator called the Pauli master equation. This implies that the spectral gap λ_Q of the quantum process is upper-bounded by the spectral gap λ_{cl} of the embedded classical one, i.e., $\lambda_Q \leq \lambda_{cl}$.

Temme asked the insightful question: when does the opposite comparison $\lambda_Q \gtrsim \lambda_{cl}$ hold? He gave two kinds of answers to this question. First, he constructed a single-body Hamiltonian with simple spectrum and $\lambda_Q = O(1/N)\lambda_{cl}$, where N is the Hilbert space dimension, showing that an efficient comparison is impossible in general. Second, he outlined an approach to proving bounds on λ_Q on a case-by-case basis by considering the matrix entries of the Davies generator (in a basis induced by the Hamiltonian eigenbasis) and applying matrix-analytic tools such as Gershgorin's theorem

to control the off-diagonal part (i.e., its action orthogonal to the diagonal). Since then, this approach was used in several works (e.g. [CB'21; RW'24; C+'24]) to separately control the diagonal and off-diagonal parts of a Davies generator, and conclude that in those cases λ_Q is not much smaller than λ_{cl} . Temme concluded his paper by asking whether his approach could be extended to Hamiltonians with degenerate spectrum, which many interesting examples such as classical Ising models, toric code Hamiltonians, and Heisenberg models exhibit.

In this chapter, we address Temme's question by identifying a condition on the spectrum of H which we call the *short arithmetic progression property*, which implies that the quantum and classical convergence rates of the Davies generator are comparable. Our main result (Theorem 4.0.11) shows that if the length of the longest arithmetic progression in $\text{spec}(H)$ is D , then $\lambda_Q \geq \Omega(1/D)\lambda_{\text{cl}}$. This is consistent with Temme's work, as his example was a Hamiltonian whose spectrum contained a long arithmetic progression of length N ; our result thus shows that long arithmetic progressions are the only obstacle to such a comparison. Notably, the short arithmetic progression property we introduce is considerably weaker than having simple spectrum and simple Bohr spectrum¹ (see Remark 4.0.18), and includes many natural examples which have degeneracy, including the ones mentioned above (see Remark 4.0.19).

Secondarily, we show that several families of generic Hamiltonians — notably, any many-body Hamiltonian with a random local external field (Theorem 4.0.16) — obey the short arithmetic progression property with $D = 2$, implying that for such models the convergence rates of the quantum and classical processes are within a constant factor of each other. As an interesting consequence (Corollary 4.0.15), by applying the classical Cheeger inequality we conclude that for short arithmetic progression Hamiltonians the Davies Lindbladian spectral gap is approximately witnessed by a projection which commutes with the Hamiltonian.

Before stating our results in detail, we recall the definition of the Davies generator and record some preliminary facts.

Definition 4.0.1 (Davies generator). Let n be the number of qubits and $N = 2^n$ be the dimension of the corresponding Hilbert space. Given a Hamiltonian $H \in \mathbb{C}^{N \times N}$, an inverse temperature $\beta \geq 0$,² a set of jump operators S , and a transition rate function $G : \mathbb{R} \rightarrow \mathbb{R}$, the Davies generator³ is a linear operator $\mathcal{L}$ on $\mathbb{C}^{N \times N}$ defined for observable

¹The simple Bohr spectrum condition means that for any eigenvalues $E_j \neq E_k$ and $E_{j'} \neq E_{k'}$ of H , the corresponding energy differences are distinct; that is, $E_j - E_k \neq E_{j'} - E_{k'}$.

²Our results hold for all $\beta \in \mathbb{R}$, but we assume $\beta \geq 0$ to align with standard convention.

³The general formulation of the Davies generator includes a coherent term, writing the generator as

$$\tilde{\mathcal{L}}(f) = i[H, f] + \mathcal{L}(f),$$

where $\mathcal{L}$ is as Definition 4.0.1. This is important for detailed balance [CKG'23] and hypocoercivity arguments [FLT'25], but the inclusion of the coherent term $i[H, f]$ does not change the spectral gap so we will omit it in this paper.

f as

$$\mathcal{L}(f) := \sum_{\omega, S \in \mathcal{S}} \mathcal{L}_{\omega, S}(f),$$

where for Bohr frequency ω and jump operator $S \in \mathcal{S}$,

$$\begin{aligned} \mathcal{L}_{\omega, S}(f) &:= G(\omega)(S(\omega)^\dagger f S(\omega) - \frac{1}{2}\{S(\omega)^\dagger S(\omega), f\}) \\ &= \frac{G(\omega)}{2} \left(S(\omega)^\dagger [f, S(\omega)] + [S(\omega)^\dagger, f] S(\omega) \right), \end{aligned}$$

where $[A, B] = AB - BA$ denotes the commutator, $\{A, B\} = AB + BA$ denotes the anti-commutator, and

$$S(\omega) := \sum_{\lambda_1, \lambda_2: \lambda_1 - \lambda_2 = \omega} \Pi_{\lambda_1} S \Pi_{\lambda_2},$$

where Π_λ is the projector to the eigenspace of H of eigenvalue λ .

Throughout the chapter, we make the standard assumption that:

- the jump operator set is self-adjoint,

$$\mathcal{S}^\dagger = \{S^\dagger : S \in \mathcal{S}\} = \mathcal{S},$$

- and the transition rate function satisfies the detailed balance condition: $G(\omega) = G(-\omega)e^{-\beta\omega}$.

Under these assumptions, $\mathcal{L}$ satisfies the following Kubo–Martin–Schwinger (KMS) reversibility condition (see [BCL'24, Fact 3.1]), which can be viewed as an analog of the time-reversibility condition of classical Markov chains:

$$\langle \mathcal{L}(f), g \rangle_\rho = \langle f, \mathcal{L}(g) \rangle_\rho, \quad (4.0.2)$$

where recall that $\rho \equiv \rho_\beta = \frac{\exp(-\beta H)}{\text{tr}(\exp(-\beta H))}$ is the Gibbs state, and the inner product $\langle \cdot, \cdot \rangle_\rho$ is the KMS inner product, i.e. $\langle A, B \rangle_\rho := \text{tr}(\rho^{1/2} A \rho^{1/2} B^\dagger)$.

For an observable f , define its

- Variance: $\text{Var}(f) := \langle f, f \rangle_\rho - \text{tr}(\rho f) \text{tr}(\rho f^\dagger)$;
- Dirichlet form: $\mathcal{E}_\mathcal{L}(f) := -\langle \mathcal{L}(f), f \rangle_\rho$; and
- Rayleigh quotient of f : $\lambda_\mathcal{L}(f) := \frac{\mathcal{E}_\mathcal{L}(f)}{\text{Var}(f)}$.¹

¹with the convention that $\lambda_\mathcal{L}(f) = +\infty$ if $\text{Var}(f) = 0$.

We define the spectral gap of $\mathcal{L}$ by

$$\lambda_{\mathcal{L}} := \min_f \lambda_{\mathcal{L}}(f). \quad (4.0.3)$$

The spectral gap of $\mathcal{L}$ controls the convergence rate of the Lindbladian dynamics $e^{t\mathcal{L}}$, specifically its χ^2 mixing time (see [KT'12]). We will restrict our attention exclusively to the spectral gap in this chapter, and will not discuss other approaches to establishing mixing. An important role is played by the following invariant subspaces of the Davies generator.

Definition 4.0.4 (V_ω subspaces). For a Hamiltonian H and Bohr frequency ω , let V_ω be the image of the map

$$f \mapsto f(\omega) := \sum_{\lambda_1, \lambda_2: \lambda_1 - \lambda_2 = \omega} \Pi_{\lambda_1} f \Pi_{\lambda_2}.$$

where Π_λ is the projector to eigenspace of H of eigenvalue λ .

We note that V_0 is precisely the set of observables that commute with H (see Proposition 4.1.8).

As shown in [T'13], each V_ω is an invariant subspace under $\mathcal{L}$, that is, $\mathcal{L}(V_\omega) \subseteq V_\omega$. Define the spectral gap in V_ω as

$$\lambda_{\mathcal{L}, \omega} = \min_{f \in V_\omega} \lambda_{\mathcal{L}}(f). \quad (4.0.5)$$

Then

$$\lambda_{\mathcal{L}} = \min_{\omega} \lambda_{\mathcal{L}, \omega}. \quad (4.0.6)$$

We also define the spectral gap with respect to Hermitian observables:

$$\lambda_{\mathcal{L}}^H := \min_{f: f=f^\dagger} \lambda_{\mathcal{L}}(f) \quad \text{and} \quad \lambda_{\mathcal{L},0}^H := \min_{f \in V_0: f=f^\dagger} \lambda_{\mathcal{L}}(f). \quad (4.0.7)$$

In Propositions 4.1.10 and 4.1.12, we show that $\lambda_{\mathcal{L}}^H = \lambda_{\mathcal{L}}$ and $\lambda_{\mathcal{L},0}^H = \lambda_{\mathcal{L},0}$. When the Hamiltonian H has simple spectrum, [T'13] relates the spectral gap $\lambda_{\mathcal{L},0}^H$ of Hermitian observables in V_0 to the spectral gap of a certain classical Markov generator.

Definition 4.0.8 (classical Markov generator, [T'13]). Consider a Hamiltonian H with an orthonormal eigenbasis $U = \{u_i\}_{i=1}^N$. Given a Davies generator $\mathcal{L}$ associated with H , inverse temperature $\beta \geq 0$, transition rate function $G : \mathbb{R} \rightarrow \mathbb{R}$, and jump operator set S , we define a corresponding classical Markov generator with state space U , using the following construction.

Let $\rho = \rho_\beta := \frac{\exp(-\beta H)}{\text{tr}(\exp(-\beta H))}$, and define the corresponding stationary distribution $\pi : U \rightarrow \mathbb{R}_{\geq 0}$ by

$$\pi(u_i) := \langle u_i | \rho | u_i \rangle.$$

The transition rate of the classical Markov generator from state u_i to u_j (for $i \neq j$) is given by

$$P_{\mathcal{L},U}[u_i \rightarrow u_j] = \langle \mathcal{L}(|u_j\rangle\langle u_j|), |u_i\rangle\langle u_i| \rangle = G(E_j - E_i) \sum_{S \in \mathcal{S}} |\langle u_i | S | u_j \rangle|^2 \geq 0,$$

where $E_i = \langle u_i | H | u_i \rangle$ is the eigenvalue of H corresponding to u_i .

Given a test function $F : U \rightarrow \mathbb{R}$, the Dirichlet form of F is defined as

$$\mathcal{E}_{\mathcal{L},U,\text{cl}}(F) = \frac{1}{2} \sum_{i,j} (F(u_i) - F(u_j))^2 \pi(u_i) P_U[u_i \rightarrow u_j].$$

The *classical spectral gap* associated with $\mathcal{L}$ and the basis U is then given by

$$\lambda_{\mathcal{L},U,\text{cl}} := \min_{F:U \rightarrow \mathbb{R}} \frac{\mathcal{E}_{\mathcal{L},U,\text{cl}}(F)}{\text{Var}_{\pi}(F)},^1$$

where $\text{Var}_{\pi}(F) := \mathbb{E}_{\pi}[F^2] - \mathbb{E}_{\pi}[F]^2$ denotes the variance of F with respect to the distribution π .

The reversibility condition of the Lindbladian $\mathcal{L}$ with respect to the KMS inner product $\langle \cdot, \cdot \rangle_{\rho}$ implies that the associated classical Markov generators are reversible with respect to π .

Proposition 4.0.9 ([T13]). *For a Hamiltonian H that has simple spectrum and a Davies generator $\mathcal{L}$ associated with H , let the classical spectral gap of $\mathcal{L}$, denoted by $\lambda_{\mathcal{L},\text{cl}}$, be the spectral gap of the classical Markov generator associated with $\mathcal{L}$ and the unique orthonormal eigenbasis of H (see Definition 4.0.8). Then,*

$$\lambda_{\mathcal{L},0}^H = \lambda_{\mathcal{L},\text{cl}}.$$

4.0.1 Our results

Comparison under the short arithmetic progression property. Our main theorem is a comparison between the spectral gap of the Davies generator $\mathcal{L}$ and the spectral gap of observables in V_0 . Before stating the result, we introduce the following definition.

Definition 4.0.10 (arithmetic progression). A multiset S is said to contain an *arithmetic progression of length D* (or a *D -arithmetic progression*) if there exist a, b such that

$$\{a, a + b, a + 2b, \dots, a + (D - 1)b\} \subseteq S.$$

If $b \neq 0$, we say that S contains a *proper D -arithmetic progression*.

¹with the convention that $\frac{\mathcal{E}_{\mathcal{L},U,\text{cl}}(F)}{\text{Var}_{\pi}(F)} = +\infty$ if $\text{Var}_{\pi}(F) = 0$.

Theorem 4.0.11 (general spectral gap comparison theorem). *If the spectrum of a Hamiltonian H does not contain any proper $(D+1)$ -arithmetic progression, then for any Davies generator $\mathcal{L}$ associated with H ,*

$$\lambda_{\mathcal{L},0} \geq \lambda_{\mathcal{L}} \geq \frac{1}{2D} \lambda_{\mathcal{L},0}.$$

Theorem 4.0.11 is the consequence of (4.0.6) and the following theorem, which compares the spectral gap in V_ω for $\omega \neq 0$ with the spectral gap in V_0 .

Theorem 4.0.12. *Consider a Hamiltonian H and a Bohr frequency $\omega \neq 0$. If the spectrum of H does not contain a $(D+1)$ -arithmetic progression where the difference between consecutive terms is ω , then for any Davies generator $\mathcal{L}$ associated with H ,*

$$\lambda_{\mathcal{L},\omega} \geq \frac{1}{2} \max\left(\frac{1}{D}, 1 - e^{-|\beta\omega|}\right) \lambda_{\mathcal{L},0}.$$

When H has simple spectrum, observables in V_0 are referred to as *diagonal* observables, since they are diagonal in the unique eigenbasis of H , while observables in $\bigcup_{\omega \neq 0} V_\omega$ are called *non-diagonal* observables. In this case, the quantity $\min_{\omega \neq 0} \lambda_{\mathcal{L},\omega}$ characterizes the decay rate of non-diagonal observables, and is known as the *quantum decoherence rate*. Theorem 4.0.12 provides a comparison between the quantum decoherence rate and the decay rate of diagonal observables.

We also obtain the following generalization of [T13], which relates the spectral gap in V_0 to the spectral gap of a certain classical Markov generator, even when H has repeated eigenvalues. The key observation is that any Hermitian observables in V_0 is diagonal in some orthonormal eigenbasis of H .

Proposition 4.0.13. *For any Hamiltonian H and any Davies generator $\mathcal{L}$ associated with H , there exists an orthonormal eigenbasis $\hat{U}$ of H such that the spectral gap of (Hermitian or arbitrary) observables in V_0 is equal to the spectral gap of the classical Markov operator associated with $\mathcal{L}$ and $\hat{U}$, i.e.,*

$$\lambda_{\mathcal{L},0} = \lambda_{\mathcal{L},0}^H = \lambda_{\mathcal{L},\hat{U},\text{cl}}.$$

Moreover, $\hat{U}$ minimizes the associated classical spectral gap over all orthonormal eigenbases of H , i.e.,

$$\hat{U} = \arg \min_U \lambda_{\mathcal{L},U,\text{cl}}.$$

The following is a direct corollary of Theorem 4.0.11 and Proposition 4.0.13.

Theorem 4.0.14. *If the spectrum of a Hamiltonian H does not contain a proper $(D+1)$ -arithmetic progression, then for any Davies generator $\mathcal{L}$, there exists an orthonormal eigenbasis $\hat{U}$ of H such that the spectral gap of $\mathcal{L}$ can be compared to the spectral gap*

$\lambda_{\mathcal{L}, \hat{U}, \text{cl}}$ of the classical Markov operator associated with $\mathcal{L}$ and $\hat{U}$ via the following sandwiching inequality:

$$\lambda_{\mathcal{L}, \hat{U}, \text{cl}} \geq \lambda_{\mathcal{L}} \geq \frac{1}{2D} \lambda_{\mathcal{L}, \hat{U}, \text{cl}}.$$

Moreover, $\hat{U}$ minimizes the classical spectral gap over all orthonormal eigenbases of H , i.e.,

$$\hat{U} = \arg \min_U \lambda_{\mathcal{L}, U, \text{cl}}.$$

If additionally, the Hamiltonian H has a simple spectrum, then H has a unique orthonormal eigenbasis, and the quantum spectral gap $\lambda_{\mathcal{L}}$ can be compared to the spectral gap of the Markov generator associated with this unique eigenbasis, that is, the classical spectral gap $\lambda_{\mathcal{L}, \text{cl}}$ of $\mathcal{L}$, via

$$\lambda_{\mathcal{L}, \text{cl}} \geq \lambda_{\mathcal{L}} \geq \frac{1}{2D} \lambda_{\mathcal{L}, \text{cl}}.$$

Another corollary of Theorem 4.0.11 and Proposition 4.0.13 is that for short arithmetic progression Hamiltonians, the spectral gap of a Davies Lindbladian is approximately witnessed by a projection which commutes with the Hamiltonian. The proof is to apply the classical Cheeger inequality to the classical Markov generator associated with $\mathcal{L}$ and $\hat{U}$.

Corollary 4.0.15. *Suppose the spectrum of H does not contain a proper arithmetic progression of length $D + 1$. Let $\mathcal{L}$ be a Davies generator associated with H , inverse temperature $\beta \geq 0$, transition rate function $G : \mathbb{R} \rightarrow \mathbb{R}$, and jump operator set S . Then, there is a projection P (i.e. $P^2 = P$) such that $[P, H] = 0$ and*

$$\mathcal{E}_{\mathcal{L}}(P) \leq 4\sqrt{DM\lambda_{\mathcal{L}}} \cdot \text{Var}_{\rho}(P),$$

where

$$M = \|G\|_{\infty} \cdot \left\| \sum_{S \in \mathcal{S}} SS^{\dagger} \right\|,$$

and $\|\cdot\|$ denotes the matrix spectral norm and $\|G\|_{\infty} := \sup_{\omega} |G(\omega)|$.

The factor M comes from $\max_u \sum_{v \neq u} P_{\mathcal{L}, \hat{U}}[u \rightarrow v]$, i.e., the maximum total transition rate out of any state of the associated classical Markov operator. For common choices of G , such as $G(\omega) = \frac{1}{1 + \exp(\beta\omega)}$ and $G(\omega) = \min(1, \exp(-\frac{\beta\omega}{2}))$, which corresponds to the transition rates Glauber dynamics and Metropolis-Hasting dynamics respectively, we have $\|G\|_{\infty} \leq 1$, and with jump operators normalized so that $\sum_{S \in \mathcal{S}} SS^{\dagger} \leq I$, it follows that $M \leq 1$.

Short arithmetic progression property for generic Hamiltonians. Next, we show that the absence of 3-term arithmetic progressions is a *generic* property under various models of perturbation. For instance, we prove that adding a generic external field to any fixed Hamiltonian typically yields a new Hamiltonian whose spectrum contains neither 3-term arithmetic progressions nor repeated eigenvalues.¹

Theorem 4.0.16. *For a fixed Hamiltonian $H_0 \in \mathbb{C}^{N \times N}$ and Hermitian $P \in \mathbb{C}^{2 \times 2}$ where P is not a multiple of the identity, let $P_i := I^{\otimes(i-1)} \otimes P \otimes I^{\otimes(n-i)}$, and let $H := H_0 + \sum_{i=1}^n h_i P_i$ where $\mathbf{h} = (h_i)_{i=1}^n \in \mathbb{R}^n$. Then the set of $\mathbf{h}$ such that H has a 3-arithmetic progression or repeated eigenvalues has Lebesgue measure 0.*

Our technique also applies to other types of perturbations (see Section 4.4). A direct consequence of our main comparison theorems (Theorems 4.0.11 and 4.0.14), together with the absence of long proper arithmetic progressions in the spectra of generic Hamiltonians (Theorem 4.0.16 and Propositions 4.4.1, 4.4.2 and 4.4.5), is that for such Hamiltonians, the quantum spectral gap is within a constant factor of the spectral gap of a corresponding classical Markov generator. In particular, we obtain the following:

Corollary 4.0.17. *For a fixed Hamiltonian $H_0 \in \mathbb{C}^{N \times N}$ and Hermitian $P \in \mathbb{C}^{2 \times 2}$ where P is not a multiple of the identity, let $P_i := I^{\otimes(i-1)} \otimes P \otimes I^{\otimes(n-i)}$, and let $H := H_0 + \sum_{i=1}^n h_i P_i$ where $\mathbf{h} = (h_i)_{i=1}^n \in \mathbb{R}^n$. Then, for Lebesgue almost every $\mathbf{h}$, the Hamiltonian H has simple spectrum and for every Davies generator $\mathcal{L}$ associated with H : $\lambda_{\mathcal{L}} = \Theta(1)\lambda_{\mathcal{L},\text{cl}}$.*

Remark 4.0.18. The assumption that the spectrum of H contains no long proper arithmetic progressions is related to, but significantly weaker than, the condition that H has simple spectrum and simple Bohr spectrum—also known as the non-degenerate energy gaps condition (see, e.g., [HH’25, Definition 7]).

Clearly, when H has simple spectrum and simple Bohr spectrum, Theorem 4.0.11 applies with $D = 2$. Moreover, under this assumption, each subspace V_ω with $\omega \neq 0$ has dimension 1. In that case, it is easy to see that $\lambda_{\mathcal{L}} \geq \frac{1}{2}\lambda_{\mathcal{L},\text{cl}}$ (see Proposition 4.0.20).

On the other hand, the assumption that H does not contain long (proper) arithmetic progressions allows for a much richer structure. In particular, it permits the subspaces V_ω (for $\omega \neq 0$) to have dimension as large as $\Omega(2^n)$.

As an example, consider $H = \sum_{i=1}^n h_i Z_i$, or the 1D XY+Z model with periodic boundary conditions (see e.g. [KLW’15]):

$$H = J \sum_{i=1}^n X_i Y_{i+1} + h \sum_{i=1}^n Z_i \quad \text{where} \quad Y_{n+1} = Y_1.$$

¹A special case was considered in [MBS’08], focusing on the restricted setting $H = H_0 + \sum_i h_i Z_i$ with $H_0 = \sum_{i,j} J_{ij} Z_i Z_j$; our result applies more generally to arbitrary H_0 and to broader classes of perturbation (see Theorem 4.0.16 and Section 4.4).

The spectra of these Hamiltonians generically do not contain 3-term arithmetic progressions or repeated eigenvalues (see the proofs of Theorem 4.0.16 and Proposition 4.4.5). The eigenvalues of H in these cases take the form

$$\lambda_z = \sum_{i=1}^n z_i m_i, \quad \text{with } z \in \{\pm 1\}^n,$$

where the m_i are fixed parameters.

Take $\omega = 2m_n$. Then for any pair of configurations z, z' satisfying $z_i = z'_i$ for all $i \in [n-1]$ and $z_n = -z'_n$, we have:

$$\lambda_z - \lambda_{z'} = \omega.$$

There are 2^{n-1} such pairs, implying that the dimension of V_ω is at least 2^{n-1} . This illustrates that the assumption of avoiding long arithmetic progressions still allows V_ω with $\omega \neq 0$ to have exponentially large dimension and complex structure.

Remark 4.0.19. It is important to distinguish between *proper* and *non-proper* arithmetic progressions when applying Theorem 4.0.11.

For instance, consider the Hamiltonian $H = \sum_{i,j} Z_i Z_j$. This Hamiltonian has only $O(n)$ distinct eigenvalues, so by the pigeonhole principle, at least one eigenvalue must have multiplicity $2^{n-o(1)}$. As a result, the spectrum of H contains a *non-proper* arithmetic progression of length $2^{n-o(1)}$, where all terms are equal.

However, the longest proper arithmetic progression in the spectrum has length only $O(n)$. Thus, despite the presence of a long non-proper progression due to the existence of an eigenvalue with high multiplicity, the structure relevant to Theorem 4.0.11, the proper arithmetic progressions, remains short.

4.0.2 Related work and significance

Besides [T13], several works considered the question of comparing quantum and classical spectral gaps [CB'21; RS'24; C+'24], albeit under strong assumptions on the Hamiltonians and/or the Lindbladian dynamics. In particular, [CB'21; RS'24; C+'24] study Lindbladian dynamics defined by random jumps. While [C+'24] studies the Davies generator, [CB'21; RS'24] study a physically-motivated variant of the Davies generator known as the CKG generator [CKG'23]. The analysis in [CB'21; RS'24] relies on bounding the spectral gap for the expected CKG generator over all random jumps; under their assumptions, the expected CKG generator coincides with the expected Davies generator up to a rescaling of the transition rate. For any eigenbasis U of H , the expected Davies generator studied in [CB'21; RS'24; C+'24] preserves the subspace of observables that are diagonal in U , while factoring the non-diagonal observables into invariant subspaces of dimension 1, i.e. $\mathcal{L}(|u_i\rangle\langle u_j|) \subseteq \text{span}(|u_i\rangle\langle u_j|)$ for

any distinct eigenvectors u_i, u_j in U .¹ In each case, these non-diagonal invariant subspaces are then analyzed directly via adaptations of Temme’s argument and shown to be comparable to the associated classical spectral gap. In hindsight this comparison follows more conceptually from the following straightforward extension of Temme’s result (specifically [T13, Theorem 11] combined with a standard bottleneck argument applied to the classical Markov generator). Although the argument is elementary, it appears to have gone unnoticed in the literature. We include a proof in Chapter C.

Proposition 4.0.20 (extension of [T13]). *Consider a Hamiltonian H , orthonormal eigenbasis $U = \{u_i\}$ of H and a Davies generator $\mathcal{L}$ associated with H . Suppose that for all $i \neq j$, $\mathcal{L}(|u_i\rangle\langle u_j|) \subseteq \text{span}(|u_i\rangle\langle u_j|)$. Then, $\lambda_{\mathcal{L}} \geq \frac{1}{2}\lambda_{\mathcal{L},U,\text{cl}}$. In particular, when H has simple spectrum and simple Bohr spectrum, $\lambda_{\mathcal{L}} \geq \frac{1}{2}\lambda_{\mathcal{L},\text{cl}}$.*

In this context, the algorithmic significance of our work is threefold. First, one can imagine scenarios where the expected Lindbladian in the above scheme is a Davies generator corresponding to a Hamiltonian with degenerate Bohr spectrum, but satisfying the short arithmetic progression property. Second, Davies generators of commuting Hamiltonians are local and efficiently implementable, and can be directly used as algorithms; the short arithmetic progression property may be easy to verify and useful in their analysis. Finally, the ability to restrict attention to observables in V_0 may be useful in developing further proof techniques to understand the mixing of Davies and related generators.

In the context of physics, there is a physical intuition which suggests that in this context “dephasing” (corresponding to off-diagonal decay) should happen at time scales shorter than mixing [Z02; Z03; S07]. Theorem 4.0.11 shows that this is true up to a small factor for Davies generators of short arithmetic progression Hamiltonians, clarifying the mechanism of thermalization for such generators. Equivalently, the failure of rapid dephasing can always be blamed on a long arithmetic progression in the spectrum of the Hamiltonian.

While it is easy to construct artificial local Hamiltonians with long arithmetic progressions (e.g. $H = \sum_{i \leq n} 2^{-i} Z_i$), we are not aware of any natural physical n -qubit local Hamiltonians whose spectra contain arithmetic progressions of length longer than $O(n)$, with this bound being met by the Toric code [K03].² A natural research direction is to show that some large class of physical Hamiltonians — such as the transverse field Ising model on arbitrary graphs — satisfy the short arithmetic progression prop-

¹[RW24] presents several results, of which the above assertion applies to all except the cyclic graph and hypercube Hamiltonians. In the cyclic graph case, the expected CKG generator coincides with the expected Davies generator, allowing our comparison result for Davies generators to streamline their analysis. The hypercube case, however, relies on a different approach tailored to their specific choice of jump operators.

²Since the toric code Hamiltonian consists of pairwise commuting terms where each has eigenvalues ± 1 , it is easy to explicitly compute the Hamiltonian spectrum.

erty with $D = \text{poly}(n)$, which would yield a useful quantum to classical comparison for this class.

4.0.3 Techniques

Proof of main comparison theorem (Theorem 4.0.12 and Theorem 4.0.11)

We outline the strategy for proving Theorem 4.0.12. Given an observable $f \in V_\omega$ where $\omega \neq 0$, we construct two observables:

$$g = e^{\frac{\beta\omega}{4}}(ff^\dagger)^{1/2}, \quad \text{and} \quad h = e^{-\frac{\beta\omega}{4}}(f^\dagger f)^{1/2}.$$

Both g and h are Hermitian and belong to V_0 . Our strategy is to

1. lower bound the Dirichlet form of f by the Dirichlet forms of g and h , then
2. relate the Dirichlet forms of g and h to their variances, and
3. lower bound the variance of g and h by the variance of f .

For (1), we rewrite the Dirichlet form of f , g , and h as weighted commutator norms $\|\rho^{1/4}[S(\tilde{\omega}), f]\rho^{1/4}\|_F^2$ then apply a novel trace inequality whose proof utilizes the spectral decomposition of f, g, h . (2) is simply due to g and h being observables in V_0 and the definition of $\lambda_{\mathcal{L},0}$. For (3), we again use a spectral decomposition of f, g , and h and further note that the vectors appearing in this decomposition are eigenvectors of the Hamiltonian H , and use the relationship between their eigenvalues to lower bound the variances of g and h .

Specifically, $g = e^{\frac{\beta\omega}{4}} \sum \lambda_i |u_i\rangle\langle u_i|$, $f = \sum \lambda_i |u_i\rangle\langle v_i|$, and $h = e^{-\frac{\beta\omega}{4}} \sum \lambda_i |v_i\rangle\langle v_i|$, with u_i, v_i being eigenvectors of both H and ρ and the corresponding eigenvalues of H satisfy $H(v_i) = H(u_i) - \omega$.

Intuitively, lower-bounding the variances of g and h is equivalent to showing that g and h are far from (multiples of) the identity. More concretely, using the above decomposition and Hölder's inequality, and letting $\rho(v_i), \rho(u_i)$ be the corresponding eigenvalues of ρ , we get:

$$\text{Var}(g) \geq \left(1 - \sum \rho(u_i)\right) \text{Var}(f), \quad \text{Var}(h) \geq \left(1 - \sum \rho(v_i)\right) \text{Var}(f)$$

and since $1 \geq \sum \rho(v_i) = e^{\beta\omega} \sum \rho(u_i)$ and $1 \geq \sum \rho(u_i) = e^{-\beta\omega} \sum \rho(v_i)$, we conclude:

$$\max(\text{Var}(g), \text{Var}(h)) \geq (1 - e^{-|\beta\omega|}) \text{Var}(f).$$

Next, we provide some intuition on why $\max(\text{Var}(g), \text{Var}(h)) \geq \frac{1}{D} \text{Var}(f)$ where D is the length of the longest arithmetic progression with common difference ω in the spectrum of H . Take a sequence of eigenvectors $u_{i_1}, \dots, u_{i_k}$, where $H(u_{i_{j+1}}) = H(u_{i_j}) - \omega$

for all j ; then $H(v_{i_k}) = H(u_{i_k}) - \omega$ is also in $\text{spec}(H)$. By assuming that this chain is maximal, we conclude that v_{i_k} is not in the spectral decomposition of g . We can assume WLOG that $\beta\omega \geq 0$, in which case we can argue that $\rho(v_{i_k}) \geq \frac{1}{k} \sum_{j=1}^k \rho(u_{i_j})$. Additionally, note that the $k+1 \leq D$ since $H(u_{i_1}), \dots, H(u_{i_k}), H(v_{i_k})$ forms a $(k+1)$ -arithmetic progression with common difference ω in $\text{spec}(H)$; these observations are enough to conclude that $1 - \sum \rho(u_i) \geq 1/D$ and thus $\text{Var}(g) \geq \frac{1}{D} \text{Var}(f)$.

Proof that generic Hamiltonians do not have long arithmetic progressions

Finally, we sketch the proof of Theorem 4.0.16, following the approach in [HH'25]. Let $\lambda_1, \dots, \lambda_N$ be the eigenvalues of a Hamiltonian H . Observe that H has no 3-term arithmetic progression in its spectrum if and only if

$$F := \prod_{i,j,k \text{ distinct}} (\lambda_i + \lambda_j - 2\lambda_k)$$

is nonzero.

Since F is symmetric in the eigenvalues, it can be viewed as a polynomial in the entries of H . Thus, the set of Hamiltonians for which $F = 0$ has measure zero, unless $F \equiv 0$ identically. To rule this out, it suffices to exhibit a single instance where $F \neq 0$; such an example can be constructed for each type of perturbation we consider.

We remark that this approach generalizes to the case of longer D -term arithmetic progressions. In that case, the role of the factor $(\lambda_i + \lambda_j - 2\lambda_k)$ is replaced by a suitable expression such as

$$\sum_{j=1}^{D-2} \left((\lambda_{i_j} - \lambda_{i_{j+1}}) - (\lambda_{i_{j+1}} - \lambda_{i_{j+2}}) \right)^2,$$

which similarly vanishes when a D -term arithmetic progression is present. We can also prove that H does not have repeated eigenvalues using a similar strategy, where we replace F with $G = \prod_{i \neq j} (\lambda_i - \lambda_j)$.

4.0.4 Organization

In Section 4.2, we prove Theorem 4.0.12 and Theorem 4.0.11. In Section 4.3, we prove Proposition 4.0.13, Theorem 4.0.14, and Corollary 4.0.15. In Section 4.4, we prove Theorem 4.0.16 and related results.

4.1 Preliminaries

4.1.1 Notation

Let n be the number of qubits, and let $N = 2^n$ denote the dimension of the Hilbert space. We write ρ for the quantum Gibbs state, use $\langle \cdot, \cdot \rangle_\rho$ to denote the KMS inner

product, and define the associated norm by

$$\|A\|_\rho^2 := \langle A, A \rangle_\rho = \|\rho^{1/4} A \rho^{1/4}\|_{\mathbb{F}}^2.$$

4.1.2 Linear algebra

Proposition 4.1.1. *Let A, B be arbitrary matrices. Then the nonzero spectrum of AB and BA are the same, i.e., $\text{spec}(AB) \cap \mathbb{R}_{\neq 0} = \text{spec}(BA) \cap \mathbb{R}_{\neq 0}$. In other words, if $\lambda \neq 0$ is an eigenvalue of AB with multiplicity m , then it is also an eigenvalue of BA with multiplicity m , thus $\text{tr}(AB) = \text{tr}(BA)$.*

Proof. It's enough to check that, if $v_1, \dots, v_m$ are linearly independent eigenvectors of AB corresponding to an eigenvalue $\lambda \neq 0$, then $Bv_1, \dots, Bv_m$ are linearly independent eigenvectors of BA corresponding to the eigenvalue λ . First, observe that, for each v_i , $BA(Bv_i) = B(ABv_i) = \lambda(Bv_i)$. Suppose $\{Bv_i\}_i$ are not linearly independent; then there exist nontrivial coefficients $\{c_i\}_i$ such that $\sum_i c_i(Bv_i) = 0$, but this means $\lambda \sum_i c_i v_i = AB \sum_i c_i v_i = A(B \sum_i c_i v_i) = 0$, thus $\sum_i c_i v_i = 0$, contradicting the assumption that $\{v_i\}_{i=1}^m$ are linearly independent. ■

4.1.3 Dirichlet form and spectral gap of the Davies generator

We begin by reviewing some basic properties of the subspace of observables V_ω (as defined in Definition 4.0.4), and then establish the relationship between the Lindbladian spectral gap $\lambda_{\mathcal{L}}$ and the spectral gap $\lambda_{\mathcal{L}, \omega}$ in V_ω , as stated in (4.0.6). These properties have been observed in [T13], but we include the proofs for the sake of completeness.

Proposition 4.1.2. *If $f \in V_\omega$ and $g \in V_{\omega'}$ then $fg \in V_{\omega+\omega'}$.*

Proof. Note that $f = \sum_\lambda \Pi_{\lambda+\omega} f \Pi_\lambda$ and $g = \sum_{\lambda'} \Pi_{\lambda'+\omega'} g \Pi_{\lambda'}$ and $\Pi_\lambda \Pi_{\lambda'} = \Pi_\lambda \delta_{\lambda=\lambda'}$, then

$$fg = \left(\sum_\lambda \Pi_{\lambda+\omega} f \Pi_\lambda \right) \left(\sum_{\lambda'} \Pi_{\lambda'+\omega'} g \Pi_{\lambda'} \right) = \sum_\lambda \Pi_{\lambda+\omega} f \Pi_\lambda g \Pi_{\lambda-\omega'} \in V_{\omega+\omega'}. \quad \blacksquare$$

Proposition 4.1.3. *If $h \in V_\omega$ with $\omega \neq 0$ then $\text{tr}(h) = 0$. Consequently, $f \in V_\omega$ and $g \in V_{\omega'}$ with $\omega - \omega' \neq 0$ then $\langle f, g \rangle_\rho = 0$.*

Proof. For any $h \in V_{\tilde{\omega}}$ with $\tilde{\omega} \neq 0$, we can write $h = \sum_\lambda \Pi_{\lambda+\tilde{\omega}} h \Pi_\lambda$ and since $\Pi_\lambda \Pi_{\lambda+\tilde{\omega}} = 0$,

$$\text{tr}(h) = \sum_\lambda \text{tr}(\Pi_{\lambda+\tilde{\omega}} h \Pi_\lambda) = \sum_\lambda \text{tr}(\Pi_\lambda \Pi_{\lambda+\tilde{\omega}} h) = 0.$$

By Proposition 4.1.2, $f \rho^{1/2} g^\dagger \rho^{1/2} \in V_{\omega-\omega'}$, and $\omega - \omega' \neq 0$ so $\langle f, g \rangle_\rho = \text{tr}(f \rho^{1/2} g^\dagger \rho^{1/2}) = 0$. ■

Proposition 4.1.4. *Consider a Hamiltonian H and a Davies generator $\mathcal{L}$ associated with H . Then, V_ω is an invariant subspace of $\mathcal{L}$; i.e., $\mathcal{L}(V_\omega) \subseteq V_\omega$.*

Proof. Fix $f \in V_\omega$. We only need to show that for any Bohr frequency ω' , $\mathcal{L}_{\omega',S}(f) \in V_\omega$, which is true by Proposition 4.1.2 and the fact that $S(\omega') \in V_{\omega'}, S(\omega')^\dagger \in V_{-\omega'}$. ■

The following relate the spectral gap of the Lindbladian to the spectral gap on the invariant subspaces V_ω .

Proposition 4.1.5. *Let $\lambda_{\mathcal{L},\omega}$ be as defined in (4.0.5). Then,*

$$\lambda_{\mathcal{L}} = \min_{\omega} \lambda_{\mathcal{L},\omega}.$$

To prove Proposition 4.1.5, we need the following proposition.

Proposition 4.1.6. *For any arbitrary observable $f \in \mathbb{C}^{N \times N}$, we have*

$$f = \sum_{\omega} f(\omega)$$

where the sum is over Bohr frequencies ω and $f(\omega)$ is as defined in Definition 4.0.4. Moreover,

$$\mathcal{E}_{\mathcal{L}}(f) = \sum_{\omega} \mathcal{E}_{\mathcal{L}}(f(\omega)) \quad \text{and} \quad \text{Var}(f) = \sum_{\omega} \text{Var}(f(\omega)).$$

Proof. For the first statement, since $I = \sum_{\lambda} \Pi_{\lambda}$, we can write

$$f = \left(\sum_{\lambda_1} \Pi_{\lambda_1} \right) f \left(\sum_{\lambda_2} \Pi_{\lambda_2} \right) = \sum_{\lambda_1, \lambda_2} \Pi_{\lambda_1} f \Pi_{\lambda_2} = \sum_{\omega} \sum_{\lambda_1, \lambda_2: \lambda_1 - \lambda_2 = \omega} \Pi_{\lambda_1} f \Pi_{\lambda_2} = \sum_{\omega} f(\omega)$$

By Propositions 4.1.3 and 4.1.4,

$$\mathcal{E}_{\mathcal{L}}(f) = - \left\langle \sum_{\omega} f(\omega), \sum_{\omega} \mathcal{L}(f(\omega)) \right\rangle_{\rho} = - \sum_{\omega} \langle f(\omega), \mathcal{L}(f(\omega)) \rangle_{\rho} = \sum_{\omega} \mathcal{E}_{\mathcal{L}}(f(\omega)).$$

Note that by Proposition 4.1.3, $\text{tr}(\rho f) = \text{tr}(\rho f(0))$ since $\text{tr}(\rho f(\omega)) = 0$ for all $\omega \neq 0$, and $\langle f(\omega), f(\omega') \rangle_{\rho} = 0$ if $\omega \neq \omega'$. We thus obtain:

$$\begin{aligned} \text{Var}(f) &= \langle f, f \rangle_{\rho} - \text{tr}(\rho f) \text{tr}(\rho f^\dagger) = \langle f, f \rangle_{\rho} - \text{tr}(\rho f(0)) \text{tr}(\rho(f(0))^\dagger) \\ &= \sum_{\omega} \langle f(\omega), f(\omega) \rangle_{\rho} - \text{tr}(\rho f(0)) \text{tr}(\rho(f(0))^\dagger) \\ &= \sum_{\omega} \text{Var}(f(\omega)). \end{aligned} \quad \blacksquare$$

Proof of Proposition 4.1.5. Obviously, for all ω ,

$$\lambda_{\mathcal{L},\omega} = \min_{f \in V_\omega} \lambda(f) \geq \min_f \lambda(f) = \lambda_{\mathcal{L}}, \quad \text{thus} \quad \min_\omega \lambda_{\mathcal{L},\omega} \geq \lambda_{\mathcal{L}}.$$

Let f be such that $\lambda(f) = \frac{\mathcal{E}_{\mathcal{L}}(f)}{\text{Var}(f)} = \lambda_{\mathcal{L}}$. Using Proposition 4.1.6, write $f = \sum_\omega f(\omega)$ with $f(\omega) \in V_\omega$. Then

$$\lambda_{\mathcal{L}} = \frac{\sum_\omega \mathcal{E}_{\mathcal{L}}(f(\omega))}{\sum_\omega \text{Var}(f(\omega))} \geq \min_\omega \lambda_{\mathcal{L}}(f(\omega)) \geq \min_\omega \lambda_{\mathcal{L},\omega}. \quad \blacksquare$$

Proposition 4.1.7. Recalling that $\rho \equiv \rho_\beta = \frac{\exp(-\beta H)}{\text{tr}(\exp(-\beta H))}$, for any observable $f \in V_\omega$ and $m \in \mathbb{R}$,

$$\rho^m f = f \rho^m \exp(-\beta m \omega).$$

Proof. For $f \in V_\omega$, we have $f = f(\omega) = \sum_\lambda \Pi_{\lambda+\omega} f \Pi_\lambda$, so that

$$\begin{aligned} e^{-\beta m H} f &= \sum_\lambda e^{-\beta m H} \Pi_{\lambda+\omega} f \Pi_\lambda = \sum_\lambda e^{-\beta m(\lambda+\omega)} \Pi_{\lambda+\omega} f \Pi_\lambda \\ &= e^{-\beta m \omega} \sum_\lambda \Pi_{\lambda+\omega} f \Pi_\lambda e^{-\beta m \lambda} = e^{-\beta m \omega} \sum_\lambda \Pi_{\lambda+\omega} f \Pi_\lambda e^{-\beta m H} = e^{-\beta m \omega} f e^{-\beta m H}; \end{aligned}$$

thus

$$\rho^m f = e^{-\beta m \omega} f \rho^m. \quad \blacksquare$$

Proposition 4.1.8. $f \in V_0$ if and only if f commutes with H .

Proof. Consider $f \in V_0$, then $f = \sum_\lambda \Pi_\lambda f \Pi_\lambda$ and $H \Pi_\lambda = \lambda \Pi_\lambda = \Pi_\lambda H$ thus

$$f H = \sum_\lambda \Pi_\lambda f \Pi_\lambda H = \sum_\lambda \lambda \Pi_\lambda f \Pi_\lambda = \sum_\lambda H \Pi_\lambda f \Pi_\lambda = H f.$$

Consider f that commutes with H . Easy to see that $f^\dagger$ also commutes with H , since $f^\dagger H = (H f)^\dagger = (f H)^\dagger = H f^\dagger$. Thus, $\frac{f+f^\dagger}{2}$ and $\frac{i(f-f^\dagger)}{2}$ also commute with H . Since $\frac{f+f^\dagger}{2}$ is Hermitian and commutes with H , $\frac{f+f^\dagger}{2}$ and H are simultaneously diagonalizable by some eigenbasis $\{u_j\}$ of H , so

$$\frac{f+f^\dagger}{2} = \sum_j \gamma_j |u_j\rangle\langle u_j|.$$

Similarly, for some eigenbasis $\{u'_i\}$ of H ,

$$\frac{i(f - f^\dagger)}{2} = \sum_j \gamma'_j |u'_j\rangle\langle u'_j|,$$

thus $f = \sum_j \gamma_j |u_j\rangle\langle u_j| - i \sum_j \gamma'_j |u'_j\rangle\langle u'_j|$ so $f \in V_0$. ■

We note that the Dirichlet form can be written as a sum of norms of commutators. Previously, [CM'18, Proposition 2.5] showed a similar result for the special case $G(\omega) = \exp(-\frac{\beta\omega}{2})$. Related identities for other specific choices of the transition rate function G can be found in [RFA'25; CR'25]. For completeness, we include a proof in Chapter C.

Proposition 4.1.9 (Dirichlet in divergence form). *For any Hamiltonian H , and Davies generator $\mathcal{L}$ as defined in Equation (4.0.1), the Dirichlet form $\mathcal{E}_{\mathcal{L}}(f) = -\langle \mathcal{L}(f), f \rangle_\rho$ can be rewritten as*

$$\mathcal{E}_{\mathcal{L}}(f) = \frac{1}{2} \sum_{S \in \mathcal{S}, \omega} \tilde{G}(\omega) \| [S(\omega), f] \|_\rho^2,$$

where $\tilde{G}(\omega) = G(\omega) e^{\frac{\beta\omega}{2}}$.

We can now show that the spectral gap is unchanged if restricted to Hermitian observables.

Proposition 4.1.10. *Let $\lambda_{\mathcal{L},0}, \lambda_{\mathcal{L},0}^H$ be as defined in (4.0.5) and (4.0.7) respectively. Then*

$$\lambda_{\mathcal{L},0} = \lambda_{\mathcal{L},0}^H.$$

We need the following proposition.

Proposition 4.1.11. *With the same setup as Proposition 4.1.9,*

$$\mathcal{E}_{\mathcal{L}}(f) = \mathcal{E}_{\mathcal{L}}(f^\dagger) \quad \text{and} \quad \text{Var}(f) = \text{Var}(f^\dagger).$$

Proof. Since $S^\dagger(-\omega) = S(\omega)^\dagger$, we have

$$[S^\dagger(-\omega), f^\dagger] = -[S(\omega), f]^\dagger$$

and thus for $A := \rho^{1/4} [S(\omega), f] \rho^{1/4}$,

$$\| [S(\omega), f] \|_\rho^2 = \| A \|_{\mathbb{F}}^2 = \| A^\dagger \|_{\mathbb{F}}^2 = \| [S^\dagger(-\omega), f^\dagger] \|_\rho^2.$$

Combining this with $\tilde{G}(\omega) = \tilde{G}(-\omega)$, we have

$$\mathcal{E}_{\mathcal{L}}(f) = \frac{1}{2} \sum_{S \in \mathcal{S}, \omega} \tilde{G}(\omega) \| [S(\omega), f] \|_\rho^2 = \frac{1}{2} \sum_{S \in \mathcal{S}, \omega} \tilde{G}(-\omega) \| [S^\dagger(-\omega), f^\dagger] \|_\rho^2 = \mathcal{E}_{\mathcal{L}}(f^\dagger).$$

Next,

$$\text{Var}(f) = \text{tr}\left(f\rho^{1/2}f^\dagger\rho^{1/2}\right) - \text{tr}(\rho f)\text{tr}(\rho f^\dagger) = \text{tr}\left(f^\dagger\rho^{1/2}f\rho^{1/2}\right) - \text{tr}(\rho f^\dagger)\text{tr}(\rho f) = \text{Var}(f^\dagger). \quad \blacksquare$$

Proof of Proposition 4.1.10. Obviously $\lambda_{\mathcal{L},0}^H \geq \lambda_{\mathcal{L},0}$. We prove $\lambda_{\mathcal{L},0} \geq \lambda_{\mathcal{L},0}^H$. Fix $f \in V_0$ such that $\frac{\mathcal{E}_{\mathcal{L}}(f)}{\text{Var}(f)} = \lambda_{\mathcal{L},0}$. We show that

$$\mathcal{E}_{\mathcal{L}}(f) = \mathcal{E}_{\mathcal{L}}\left(\frac{f+f^\dagger}{2}\right) + \mathcal{E}_{\mathcal{L}}\left(\frac{i(f-f^\dagger)}{2}\right).$$

Indeed, by Propositions 4.1.9 and 4.1.11, the LHS can be rewritten as

$$\begin{aligned} \mathcal{E}_{\mathcal{L}}(f) &= \frac{\mathcal{E}_{\mathcal{L}}(f) + \mathcal{E}_{\mathcal{L}}(f^\dagger)}{2} = \frac{1}{2} \sum_{S \in \mathcal{S}, \omega} \tilde{G}(\omega) \frac{\| [S(\omega), f] \|_\rho^2 + \| [S(\omega), f^\dagger] \|_\rho^2}{2} \\ &=_{(1)} \frac{1}{2} \sum_{S \in \mathcal{S}, \omega} \tilde{G}(\omega) \left(\left\| \left[S(\omega), \frac{f+f^\dagger}{2} \right] \right\|_\rho^2 + \left\| \left[S(\omega), \frac{i(f-f^\dagger)}{2} \right] \right\|_\rho^2 \right) \\ &= \mathcal{E}_{\mathcal{L}}\left(\frac{f+f^\dagger}{2}\right) + \mathcal{E}_{\mathcal{L}}\left(\frac{i(f-f^\dagger)}{2}\right) \end{aligned}$$

where (1) is due to the identity $\text{tr}(AA^\dagger) + \text{tr}(BB^\dagger) = \frac{1}{2}(\text{tr}((A+B)(A+B)^\dagger) + \text{tr}((A-B)(A-B)^\dagger))$, where $A = \rho^{1/4}[S(\omega), f]\rho^{1/4}$ and $B = \rho^{1/4}[S(\omega), f^\dagger]\rho^{1/4}$.

Next, we show

$$\text{Var}\left(\frac{f+f^\dagger}{2}\right) + \text{Var}\left(\frac{i(f-f^\dagger)}{2}\right) = \text{Var}(f)$$

and using the fact that $\frac{f+f^\dagger}{2}, \frac{i(f-f^\dagger)}{2}$ are Hermitian and in V_0 , conclude that

$$\begin{aligned} \mathcal{E}_{\mathcal{L}}(f) &= \mathcal{E}_{\mathcal{L}}\left(\frac{f+f^\dagger}{2}\right) + \mathcal{E}_{\mathcal{L}}\left(\frac{i(f-f^\dagger)}{2}\right) \\ &\geq \lambda_0^H \left(\text{Var}\left(\frac{f+f^\dagger}{2}\right) + \text{Var}\left(\frac{i(f-f^\dagger)}{2}\right) \right) = \lambda_{\mathcal{L},0}^H \text{Var}(f), \end{aligned}$$

which implies $\lambda_{\mathcal{L},0} \geq \lambda_{\mathcal{L},0}^H$ due to the choice of f .

Indeed,

$$\text{Var}\left(\frac{f+f^\dagger}{2}\right) + \text{Var}\left(\frac{i(f-f^\dagger)}{2}\right)$$

$$\begin{aligned}
 &= \frac{1}{4} \left(\langle f + f^\dagger, f + f^\dagger \rangle_\rho - \text{tr}(\rho(f + f^\dagger))^2 + \langle f - f^\dagger, f - f^\dagger \rangle_\rho - \text{tr}(\rho(f - f^\dagger)) \text{tr}(\rho(f^\dagger - f)) \right) \\
 &= \frac{1}{2} (\langle f, f \rangle_\rho + \langle f^\dagger, f^\dagger \rangle_\rho) - \text{tr}(\rho f) \text{tr}(\rho f^\dagger) \\
 &= \text{Var}(f). \quad \blacksquare
 \end{aligned}$$

With a similar proof,¹ we obtain the following.

Proposition 4.1.12. *Let $\lambda_{\mathcal{L}}$, $\lambda_{\mathcal{L}}^{\text{H}}$ be as defined in (4.0.3) and (4.0.7) respectively. Then*

$$\lambda_{\mathcal{L}} = \lambda_{\mathcal{L}}^{\text{H}}.$$

Remark 4.1.13. An alternative proof of Propositions 4.1.10 and 4.1.12 is by viewing $\lambda_{\mathcal{L}}$ as the second smallest eigenvalue of $\mathcal{L}$, and $\lambda_{\mathcal{L},0}$ as the second smallest eigenvalue of the linear operator obtained by restricted $\mathcal{L}$ to the invariant subspace V_0 . If f is an eigenvector of $\mathcal{L}$ with eigenvalue $\lambda_{\mathcal{L}}$, then so is $f^\dagger$ by Proposition 4.1.11 and the min-max principle. The linear combination $\frac{f+f^\dagger}{2}$ is Hermitian, and is also an eigenvector of $\mathcal{L}$ with eigenvalue $\lambda_{\mathcal{L}}$, yielding that $\lambda_{\mathcal{L}} = \lambda_{\mathcal{L}}^{\text{H}}$. Similarly, if $f \in V_0$ is an eigenvector of $\mathcal{L}$ with eigenvalue $\lambda_{\mathcal{L},0}$, then the Hermitian observable $\frac{f+f^\dagger}{2} \in V_0$ is also an eigenvector of $\mathcal{L}$ with eigenvalue $\lambda_{\mathcal{L},0}$, implying that $\lambda_{\mathcal{L},0} = \lambda_{\mathcal{L},0}^{\text{H}}$.

Some formulations of the Davies generator include the coherent Hamiltonian evolution explicitly, writing the generator as

$$\tilde{\mathcal{L}}(f) = i[H, f] + \mathcal{L}(f),$$

where $\mathcal{L}$ is as Definition 4.0.1. The following proposition clarifies that the inclusion of the coherent term $i[H, f]$ has no effect on the Dirichlet form, and thus does not affect the spectral gap.

Proposition 4.1.14. *Let $\tilde{\mathcal{L}}(f) = i[H, f] + \mathcal{L}(f)$, where $\mathcal{L}$ is the Davies generator as defined in Definition 4.0.1, and let ρ be the stationary state. Then:*

1. *For all observables f , the real part of the Dirichlet form is unaffected by the coherent term:*

$$\text{Re}(\langle f, \tilde{\mathcal{L}}(f) \rangle_\rho) = \langle f, \mathcal{L}(f) \rangle_\rho.$$

2. *In particular, for any Hermitian observable $f = f^\dagger$, the Dirichlet forms of $\tilde{\mathcal{L}}$ and $\mathcal{L}$ coincide:*

$$\langle f, \tilde{\mathcal{L}}(f) \rangle_\rho = \langle f, \mathcal{L}(f) \rangle_\rho.$$

¹Take arbitrary f such that $\lambda(f) = \lambda_{\mathcal{L}}$ instead of $f \in V_0$ such that $\lambda(f) = \lambda_{\mathcal{L},0}$.

Proof. We only need to check that

$$\langle [H, f], f \rangle_\rho \in \mathbb{R} \quad \text{and} \quad \text{for all } f \text{ such that } f = f^\dagger, \langle [H, f], f \rangle_\rho = 0.$$

Let $\tilde{f} = \rho^{1/4} f \rho^{1/4}$. Since ρ commutes with H , we can rewrite

$$\begin{aligned} \langle [H, f], f \rangle_\rho &= \text{tr}(\rho^{1/4} H f \rho^{1/2} f^\dagger \rho^{1/4}) - \text{tr}(\rho^{1/4} f H \rho^{1/2} f^\dagger \rho^{1/4}) \\ &= \text{tr}(\tilde{f}^\dagger H \tilde{f}) - \text{tr}(\tilde{f} H \tilde{f}^\dagger) \\ &= \text{tr}(\tilde{f}^\dagger H \tilde{f} - \tilde{f} H \tilde{f}^\dagger) \in \mathbb{R} \end{aligned}$$

since $\tilde{f}^\dagger H \tilde{f} - \tilde{f} H \tilde{f}^\dagger$ is Hermitian, thus has real trace. If $f = f^\dagger$ then $\tilde{f} = \tilde{f}^\dagger$, so the above yields

$$\langle [H, f], f \rangle_\rho = 0. \quad \blacksquare$$

4.2 Proof of main comparison theorems

In this section, we prove Theorems 4.0.11 and 4.0.12.

For a Bohr frequency ω , and an operator $f \in V_\omega$, we let

$$g := e^{\frac{\beta\omega}{4}} (f f^\dagger)^{1/2} \quad \text{and} \quad h := e^{-\frac{\beta\omega}{4}} (f^\dagger f)^{1/2} \quad (4.2.1)$$

Theorem 4.2.2 (comparing Dirichlet forms). *For any $f \in V_\omega$, and g, h as defined in (4.2.1). For any Davies generator $\mathcal{L}$,*

$$2\mathcal{E}_{\mathcal{L}}(f) \geq \mathcal{E}_{\mathcal{L}}(g) + \mathcal{E}_{\mathcal{L}}(h).$$

To prove Theorem 4.2.2, we need the following propositions.

Proposition 4.2.3. *Fix an arbitrary observable ϕ . Let $\tilde{\phi} := \rho^{1/2} \phi \rho^{1/2}$. Then for any Bohr frequency $\tilde{\omega}$,*

$$\begin{aligned} ||[S(\tilde{\omega}), \phi]||_\rho^2 &= e^{-\frac{\beta\tilde{\omega}}{2}} \text{tr}(S(\tilde{\omega}) \tilde{\phi} \tilde{\phi}^\dagger S(\tilde{\omega})^\dagger) + e^{\frac{\beta\tilde{\omega}}{2}} \text{tr}(S(\tilde{\omega})^\dagger \tilde{\phi}^\dagger \tilde{\phi} S(\tilde{\omega})) \\ &\quad - \text{tr}(S(\tilde{\omega}) \tilde{\phi} S(\tilde{\omega})^\dagger \tilde{\phi}^\dagger) - \text{tr}(S(\tilde{\omega}) \tilde{\phi}^\dagger S(\tilde{\omega})^\dagger \tilde{\phi}). \end{aligned}$$

Proof. We begin by writing

$$||[S(\tilde{\omega}), f]||_\rho^2 = \text{tr}(\rho^{1/4} (S(\tilde{\omega}) \phi - \phi S(\tilde{\omega})) \rho^{1/2} (\phi^\dagger S(\tilde{\omega})^\dagger - S(\tilde{\omega})^\dagger \phi^\dagger) \rho^{1/4})$$

$$\begin{aligned}
 &= \operatorname{tr}\left(\rho^{1/4}S(\tilde{\omega})\phi\rho^{1/2}\phi^\dagger S(\tilde{\omega})^\dagger\rho^{1/4}\right) + \operatorname{tr}\left(\rho^{1/4}\phi S(\tilde{\omega})\rho^{1/2}S(\tilde{\omega})^\dagger\phi^\dagger\rho^{1/4}\right) \\
 &\quad - \operatorname{tr}\left(\rho^{1/4}S(\tilde{\omega})\phi\rho^{1/2}S(\tilde{\omega})^\dagger\phi^\dagger\rho^{1/4}\right) - \operatorname{tr}\left(\rho^{1/4}\phi S(\tilde{\omega})\rho^{1/2}\phi^\dagger S(\tilde{\omega})^\dagger\rho^{1/4}\right).
 \end{aligned} \tag{4.2.4}$$

Using Proposition 4.1.7 for $S(\tilde{\omega}) \in V_{\tilde{\omega}}$ and $S(\tilde{\omega})^\dagger \in V_{-\tilde{\omega}}$ and the fact that $\operatorname{tr}(AB) = \operatorname{tr}(BA)$ from Proposition 4.1.1, we have

$$\begin{aligned}
 \operatorname{tr}\left(\rho^{1/4}S(\tilde{\omega})\phi\rho^{1/2}\phi^\dagger S(\tilde{\omega})^\dagger\rho^{1/4}\right) &= e^{-\frac{\beta\tilde{\omega}}{2}} \operatorname{tr}\left(S(\tilde{\omega})\rho^{1/4}\phi\rho^{1/2}\phi^\dagger\rho^{1/4}S(\tilde{\omega})^\dagger\right) \\
 &= e^{-\frac{\beta\tilde{\omega}}{2}} \operatorname{tr}\left(S(\tilde{\omega})\tilde{\phi}\tilde{\phi}^\dagger S(\tilde{\omega})^\dagger\right), \\
 \operatorname{tr}\left(\rho^{1/4}\phi S(\tilde{\omega})\rho^{1/2}S(\tilde{\omega})^\dagger\phi^\dagger\rho^{1/4}\right) &= e^{\frac{\beta\tilde{\omega}}{2}} \operatorname{tr}\left(\rho^{1/4}\phi\rho^{1/4}S(\tilde{\omega})S(\tilde{\omega})^\dagger\rho^{1/4}\phi^\dagger\rho^{1/4}\right) \\
 &= e^{\frac{\beta\tilde{\omega}}{2}} \operatorname{tr}\left(S(\tilde{\omega})^\dagger\tilde{\phi}^\dagger\tilde{\phi}S(\tilde{\omega})\right), \\
 \operatorname{tr}\left(\rho^{1/4}S(\tilde{\omega})\phi\rho^{1/2}S(\tilde{\omega})^\dagger\phi^\dagger\rho^{1/4}\right) &= \operatorname{tr}\left(S(\tilde{\omega})\rho^{1/4}\phi\rho^{1/4}S(\tilde{\omega})^\dagger\rho^{1/4}\phi^\dagger\rho^{1/4}\right) \\
 &= \operatorname{tr}\left(S(\tilde{\omega})\tilde{\phi}S(\tilde{\omega})^\dagger\tilde{\phi}^\dagger\right), \\
 \operatorname{tr}\left(\rho^{1/4}\phi S(\tilde{\omega})\rho^{1/2}\phi^\dagger S(\tilde{\omega})^\dagger\rho^{1/4}\right) &= \operatorname{tr}\left(\rho^{1/4}\phi\rho^{1/4}S(\tilde{\omega})\rho^{1/4}\phi^\dagger\rho^{1/4}S(\tilde{\omega})^\dagger\right) \\
 &= \operatorname{tr}\left(\tilde{\phi}S(\tilde{\omega})\tilde{\phi}^\dagger S(\tilde{\omega})^\dagger\right) \\
 &= \operatorname{tr}\left(S(\tilde{\omega})\tilde{\phi}^\dagger S(\tilde{\omega})^\dagger\tilde{\phi}\right).
 \end{aligned}$$

Plugging each of the four equalities above into (4.2.4) gives the result. $\blacksquare$

Proposition 4.2.5. *Let A, B be arbitrary matrices. Let f be an arbitrary matrix, and let $g := (ff^\dagger)^{1/2}$ and $h := (f^\dagger f)^{1/2}$. Then, $\operatorname{tr}(AgA^\dagger g) + \operatorname{tr}(BhB^\dagger h) \geq \operatorname{tr}(AfB^\dagger f^\dagger) + \operatorname{tr}(Bf^\dagger A^\dagger f)$.*

Proof. Pick an orthonormal eigenbasis $u_1, \dots, u_m$ of $ff^\dagger = g^2$ with corresponding eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_k > 0 = \lambda_{k+1} = \dots = \lambda_m = 0$. Then,

$$g = \sum_{i=1}^k \sqrt{\lambda_i} u_i u_i^\dagger.$$

Let $v_i = f^\dagger u_i$. If $i \geq k+1$, then $f^\dagger u_i = 0$ since $\|f^\dagger u_i\|_2^2 = u_i^\dagger f f^\dagger u_i = 0$. Thus

$$f = \sum_{i=1}^m u_i u_i^\dagger f = \sum_{i=1}^k u_i v_i^\dagger$$

and since $\{u_i\}_{i=1}^m$ is an orthonormal eigenbasis, $\sum_{i=1}^m u_i u_i^\dagger = I$ and

$$f^\dagger f = f^\dagger \left(\sum_{i=1}^m u_i u_i^\dagger \right) f = \sum_{i=1}^k v_i v_i^\dagger.$$

For $i \in [k]$, $\|v_i\|_2^2 = u_i^\dagger f f^\dagger u_i = \lambda_i > 0$, so we can rewrite $f^\dagger f = \sum_{i=1}^k \lambda_i \cdot \frac{v_i}{\sqrt{\lambda_i}} \cdot \frac{v_i^\dagger}{\sqrt{\lambda_i}}$ and thus

$$h = (f^\dagger f)^{1/2} = \sum_{i=1}^k \sqrt{\lambda_i} \cdot \frac{v_i}{\sqrt{\lambda_i}} \cdot \frac{v_i^\dagger}{\sqrt{\lambda_i}}.$$

Hence, for $i \in [k]$, letting $\tilde{v}_i = \frac{v_i}{\sqrt{\lambda_i}}$:

$$g = \sum_{i=1}^k \sqrt{\lambda_i} u_i u_i^\dagger, \quad f = \sum_{i=1}^k \sqrt{\lambda_i} u_i \tilde{v}_i^\dagger, \quad h = \sum_{i=1}^k \sqrt{\lambda_i} \tilde{v}_i \tilde{v}_i^\dagger. \quad (4.2.6)$$

Note also that the $\tilde{v}_i$ have unit norm and are pairwise orthogonal, since $v_i^\dagger v_j = u_i^\dagger f f^\dagger u_j = 0$ if $i \neq j$. We can rewrite

$$\begin{aligned} & \text{tr}(A g A^\dagger g) + \text{tr}(B h B^\dagger h) - \text{tr}(A f B^\dagger f^\dagger) - \text{tr}(B f^\dagger A^\dagger f) \\ &= \sum_{i,j} \sqrt{\lambda_i \lambda_j} u_i^\dagger A u_j u_j^\dagger A^\dagger u_i + \sum_{i,j} \sqrt{\lambda_i \lambda_j} \tilde{v}_i^\dagger B \tilde{v}_j \tilde{v}_j^\dagger B^\dagger \tilde{v}_i \\ & \quad - \sum_{i,j} \sqrt{\lambda_i \lambda_j} u_i^\dagger A u_j \tilde{v}_j^\dagger B^\dagger \tilde{v}_i - \sum_{i,j} \sqrt{\lambda_i \lambda_j} \tilde{v}_i^\dagger B \tilde{v}_j u_j^\dagger A^\dagger u_i \\ &= \sum_{i,j} \sqrt{\lambda_i \lambda_j} |u_i^\dagger A u_j - \tilde{v}_i^\dagger B \tilde{v}_j|^2 \geq 0. \end{aligned} \quad \blacksquare$$

Proof of Theorem 4.2.2. We write the Dirichlet forms of f, g, h as sum of commutators as in Proposition 4.1.9. For each Bohr frequency $\tilde{\omega}$, introduce

$$\begin{aligned} T(\tilde{\omega}) &:= \|[S(\tilde{\omega}), f]\|_\rho^2 + \|[S^\dagger(-\tilde{\omega}), f]\|_\rho^2 \\ & \quad - \frac{1}{2} (\|[S(\tilde{\omega}), g]\|_\rho^2 + \|[S^\dagger(-\tilde{\omega}), g]\|_\rho^2 + \|[S(\tilde{\omega}), h]\|_\rho^2 + \|[S^\dagger(-\tilde{\omega}), h]\|_\rho^2). \end{aligned}$$

Using Proposition 4.1.9 and the fact that $\tilde{G}(\tilde{\omega}) = G(\tilde{\omega}) e^{\frac{\beta \tilde{\omega}}{2}} = \tilde{G}(-\tilde{\omega})$, we can rewrite

$$2\mathcal{E}_{\mathcal{L}}(f) = \sum_{S \in \mathcal{S}, \tilde{\omega}} \tilde{G}(\tilde{\omega}) \|[S(\tilde{\omega}), f]\|_\rho^2$$

$$= \frac{1}{2} \left(\sum_{S \in \mathcal{S}, \tilde{\omega}} \tilde{G}(\tilde{\omega}) \| [S(\tilde{\omega}), f] \|_{\rho}^2 + \sum_{S^+ \in \mathcal{S}^+, -\tilde{\omega}} \tilde{G}(\tilde{\omega}) \| [S^+(-\tilde{\omega}), f] \|_{\rho}^2 \right)$$

hence

$$2\mathcal{E}_{\mathcal{L}}(f) - (\mathcal{E}_{\mathcal{L}}(g) + \mathcal{E}_{\mathcal{L}}(h)) = \frac{1}{2} \sum_{\tilde{\omega}} \tilde{G}(\tilde{\omega}) T(\tilde{\omega})$$

We will show that $T(\tilde{\omega}) \geq 0$ and then conclude that

$$2\mathcal{E}_{\mathcal{L}}(f) - (\mathcal{E}_{\mathcal{L}}(g) + \mathcal{E}_{\mathcal{L}}(h)) \geq 0$$

Let $\tilde{f} = \rho^{1/4} f \rho^{1/4}$, $\tilde{g} = \rho^{1/4} g \rho^{1/4}$, $\tilde{h} = \rho^{1/4} h \rho^{1/4}$ as in Proposition 4.2.3. By Proposition 4.1.7, since $f \in V_{\omega}$ and $g, h \in V_0$ and are Hermitian,

$$\begin{aligned} \tilde{g}^2 &= \tilde{g} \tilde{g}^{\dagger} = \rho^{1/4} g \rho^{1/2} g^{\dagger} \rho^{1/4} = \rho^{1/2} g^2 \rho^{1/2} = e^{-\frac{\beta\omega}{2}} \rho^{1/2} f f^{\dagger} \rho^{1/2} = \rho^{1/4} f \rho^{1/2} f^{\dagger} \rho^{1/2} = \tilde{f} \tilde{f}^{\dagger}, \\ \tilde{h}^2 &= \tilde{h} \tilde{h}^{\dagger} = \rho^{1/4} h \rho^{1/2} h^{\dagger} \rho^{1/4} = \rho^{1/2} h^2 \rho^{1/2} = e^{\frac{\beta\omega}{2}} \rho^{1/2} f^{\dagger} f \rho^{1/2} = \rho^{1/4} f^{\dagger} \rho^{1/2} f \rho^{1/2} = \tilde{f}^{\dagger} \tilde{f}. \end{aligned}$$

Using Proposition 4.2.3, and noting that $S^+(-\tilde{\omega}) = S(\tilde{\omega})^{\dagger}$, we can rewrite

$$\begin{aligned} \| [S(\tilde{\omega}), f] \|_{\rho}^2 &= e^{-\frac{\beta\tilde{\omega}}{2}} \operatorname{tr} \left(S(\tilde{\omega}) \tilde{f} \tilde{f}^{\dagger} S(\tilde{\omega})^{\dagger} \right) + e^{\frac{\beta\tilde{\omega}}{2}} \operatorname{tr} \left(S(\tilde{\omega})^{\dagger} \tilde{f}^{\dagger} \tilde{f} S(\tilde{\omega}) \right) \\ &\quad - \operatorname{tr} \left(S(\tilde{\omega}) \tilde{f} S(\tilde{\omega})^{\dagger} \tilde{f}^{\dagger} \right) - \operatorname{tr} \left(S(\tilde{\omega}) \tilde{f}^{\dagger} S(\tilde{\omega})^{\dagger} \tilde{f} \right). \\ \| [S^+(-\tilde{\omega}), f] \|_{\rho}^2 &= e^{\frac{\beta\tilde{\omega}}{2}} \operatorname{tr} \left(S(\tilde{\omega})^{\dagger} \tilde{f} \tilde{f}^{\dagger} S(\tilde{\omega}) \right) + e^{-\frac{\beta\tilde{\omega}}{2}} \operatorname{tr} \left(S(\tilde{\omega}) \tilde{f}^{\dagger} \tilde{f} S(\tilde{\omega})^{\dagger} \right) \\ &\quad - \operatorname{tr} \left(S(\tilde{\omega})^{\dagger} \tilde{f} S(\tilde{\omega}) \tilde{f}^{\dagger} \right) - \operatorname{tr} \left(S(\tilde{\omega})^{\dagger} \tilde{f}^{\dagger} S(\tilde{\omega}) \tilde{f} \right). \\ \| [S(\tilde{\omega}), g] \|_{\rho}^2 &= e^{-\frac{\beta\tilde{\omega}}{2}} \operatorname{tr} \left(S(\tilde{\omega}) \tilde{g} \tilde{g}^{\dagger} S(\tilde{\omega})^{\dagger} \right) + e^{\frac{\beta\tilde{\omega}}{2}} \operatorname{tr} \left(S(\tilde{\omega})^{\dagger} \tilde{g}^{\dagger} \tilde{g} S(\tilde{\omega}) \right) \\ &\quad - \operatorname{tr} \left(S(\tilde{\omega}) \tilde{g} S(\tilde{\omega})^{\dagger} \tilde{g}^{\dagger} \right) - \operatorname{tr} \left(S(\tilde{\omega}) \tilde{g}^{\dagger} S(\tilde{\omega})^{\dagger} \tilde{g} \right) \\ &= e^{-\frac{\beta\tilde{\omega}}{2}} \operatorname{tr} \left(S(\tilde{\omega}) \tilde{f} \tilde{f}^{\dagger} S(\tilde{\omega})^{\dagger} \right) + e^{\frac{\beta\tilde{\omega}}{2}} \operatorname{tr} \left(S(\tilde{\omega})^{\dagger} \tilde{f} \tilde{f}^{\dagger} S(\tilde{\omega}) \right) \\ &\quad - \operatorname{tr} \left(S(\tilde{\omega}) \tilde{g} S(\tilde{\omega})^{\dagger} \tilde{g} \right) - \operatorname{tr} \left(S(\tilde{\omega}) \tilde{g}^{\dagger} S(\tilde{\omega})^{\dagger} \tilde{g} \right) \\ &= e^{-\frac{\beta\tilde{\omega}}{2}} \operatorname{tr} \left(S(\tilde{\omega}) \tilde{f} \tilde{f}^{\dagger} S(\tilde{\omega})^{\dagger} \right) + e^{\frac{\beta\tilde{\omega}}{2}} \operatorname{tr} \left(S(\tilde{\omega})^{\dagger} \tilde{f} \tilde{f}^{\dagger} S(\tilde{\omega}) \right) - 2 \operatorname{tr} \left(S(\tilde{\omega}) \tilde{g} S(\tilde{\omega})^{\dagger} \tilde{g} \right) \end{aligned}$$

where we use $\tilde{g} = \tilde{g}^{\dagger}$ and $\tilde{g}^2 = \tilde{f} \tilde{f}^{\dagger}$. Similarly,

$$\begin{aligned} \| [S^+(-\tilde{\omega}), g] \|_{\rho}^2 &= e^{\frac{\beta\tilde{\omega}}{2}} \operatorname{tr} \left(S(\tilde{\omega})^{\dagger} \tilde{g} \tilde{g}^{\dagger} S(\tilde{\omega}) \right) + e^{-\frac{\beta\tilde{\omega}}{2}} \operatorname{tr} \left(S(\tilde{\omega}) \tilde{g}^{\dagger} \tilde{g} S(\tilde{\omega})^{\dagger} \right) \\ &\quad - \operatorname{tr} \left(S(\tilde{\omega})^{\dagger} \tilde{g} S(\tilde{\omega}) \tilde{g}^{\dagger} \right) - \operatorname{tr} \left(S(\tilde{\omega})^{\dagger} \tilde{g}^{\dagger} S(\tilde{\omega}) \tilde{g} \right) \end{aligned}$$

$$= e^{\frac{\beta\tilde{\omega}}{2}} \operatorname{tr}\left(S(\tilde{\omega})^\dagger \tilde{f} \tilde{f}^\dagger S(\tilde{\omega})\right) + e^{-\frac{\beta\tilde{\omega}}{2}} \operatorname{tr}\left(S(\tilde{\omega}) \tilde{f} \tilde{f}^\dagger S(\tilde{\omega})^\dagger\right) - 2 \operatorname{tr}\left(S(\tilde{\omega}) \tilde{g} S(\tilde{\omega})^\dagger \tilde{g}\right)$$

since by Proposition 4.1.1, $\operatorname{tr}((S(\tilde{\omega})\tilde{g})(S(\tilde{\omega})^\dagger\tilde{g})) = \operatorname{tr}((S(\tilde{\omega})^\dagger\tilde{g})(S(\tilde{\omega})\tilde{g}))$. Analogously, using $\tilde{h} = \tilde{h}^\dagger$ and $\tilde{h}^2 = \tilde{f}^\dagger \tilde{f}$, we obtain:

$$\begin{aligned} \|[S(\tilde{\omega}), h]\|_\rho^2 &= e^{-\frac{\beta\tilde{\omega}}{2}} \operatorname{tr}\left(S(\tilde{\omega}) \tilde{f}^\dagger \tilde{f} S(\tilde{\omega})^\dagger\right) + e^{\frac{\beta\tilde{\omega}}{2}} \operatorname{tr}\left(S(\tilde{\omega})^\dagger \tilde{f}^\dagger \tilde{f} S(\tilde{\omega})\right) - 2 \operatorname{tr}\left(S(\tilde{\omega}) \tilde{h} S(\tilde{\omega})^\dagger \tilde{h}\right) \\ \|[S^\dagger(-\tilde{\omega}), h]\|_\rho^2 &= e^{\frac{\beta\tilde{\omega}}{2}} \operatorname{tr}\left(S(\tilde{\omega})^\dagger \tilde{f}^\dagger \tilde{f} S(\tilde{\omega})\right) + e^{-\frac{\beta\tilde{\omega}}{2}} \operatorname{tr}\left(S(\tilde{\omega}) \tilde{f}^\dagger \tilde{f} S(\tilde{\omega})^\dagger\right) - 2 \operatorname{tr}\left(S(\tilde{\omega}) \tilde{h} S(\tilde{\omega})^\dagger \tilde{h}\right) \end{aligned}$$

Let

$$\begin{aligned} M_1 &:= e^{-\frac{\beta\tilde{\omega}}{2}} \operatorname{tr}\left(S(\tilde{\omega}) \tilde{f} \tilde{f}^\dagger S(\tilde{\omega})^\dagger\right) + e^{\frac{\beta\tilde{\omega}}{2}} \operatorname{tr}\left(S(\tilde{\omega})^\dagger \tilde{f} \tilde{f}^\dagger S(\tilde{\omega})\right) \\ M_2 &:= e^{\frac{\beta\tilde{\omega}}{2}} \operatorname{tr}\left(S(\tilde{\omega})^\dagger \tilde{f}^\dagger \tilde{f} S(\tilde{\omega})\right) + e^{-\frac{\beta\tilde{\omega}}{2}} \operatorname{tr}\left(S(\tilde{\omega}) \tilde{f}^\dagger \tilde{f} S(\tilde{\omega})^\dagger\right) \end{aligned}$$

Observe that by Proposition 4.1.1,

$$\operatorname{tr}\left(S(\tilde{\omega}) \tilde{f} S(\tilde{\omega})^\dagger \tilde{f}^\dagger\right) = \operatorname{tr}\left((S(\tilde{\omega}) \tilde{f})(S(\tilde{\omega})^\dagger \tilde{f}^\dagger)\right) = \operatorname{tr}\left((S(\tilde{\omega})^\dagger \tilde{f}^\dagger)(S(\tilde{\omega}) \tilde{f})\right) = \operatorname{tr}\left(S(\tilde{\omega})^\dagger \tilde{f}^\dagger S(\tilde{\omega}) \tilde{f}\right),$$

and similarly, $\operatorname{tr}\left(S(\tilde{\omega}) \tilde{f}^\dagger S(\tilde{\omega})^\dagger \tilde{f}\right) = \operatorname{tr}\left(S(\tilde{\omega})^\dagger \tilde{f} S(\tilde{\omega}) \tilde{f}^\dagger\right)$, so we can rewrite

$$\begin{aligned} \|[S(\tilde{\omega}), f]\|_\rho^2 + \|[S^\dagger(-\tilde{\omega}), f]\|_\rho^2 &= M_1 + M_2 - 2 \operatorname{tr}\left(S(\tilde{\omega}) \tilde{f} S(\tilde{\omega})^\dagger \tilde{f}^\dagger\right) - 2 \operatorname{tr}\left(S(\tilde{\omega}) \tilde{f}^\dagger S(\tilde{\omega})^\dagger \tilde{f}\right) \\ \|[S(\tilde{\omega}), g]\|_\rho^2 + \|[S^\dagger(-\tilde{\omega}), g]\|_\rho^2 &= 2M_1 - 4 \operatorname{tr}\left(S(\tilde{\omega}) \tilde{g} S(\tilde{\omega})^\dagger \tilde{g}\right) \\ \|[S(\tilde{\omega}), h]\|_\rho^2 + \|[S^\dagger(-\tilde{\omega}), h]\|_\rho^2 &= 2M_2 - 4 \operatorname{tr}\left(S(\tilde{\omega}) \tilde{h} S(\tilde{\omega})^\dagger \tilde{h}\right) \end{aligned}$$

thus

$$T(\tilde{\omega}) = 2(\operatorname{tr}\left(S(\tilde{\omega}) \tilde{g} S(\tilde{\omega})^\dagger \tilde{g}\right) + \operatorname{tr}\left(S(\tilde{\omega}) \tilde{h} S(\tilde{\omega})^\dagger \tilde{h}\right) - \operatorname{tr}\left(S(\tilde{\omega}) \tilde{f} S(\tilde{\omega})^\dagger \tilde{f}^\dagger\right) - \operatorname{tr}\left(S(\tilde{\omega}) \tilde{f}^\dagger S(\tilde{\omega})^\dagger \tilde{f}\right)) \geq 0$$

where the inequality is due to Proposition 4.2.5, and the fact that $\tilde{g} = c$ and $\tilde{h} = (\tilde{f}^\dagger \tilde{f})^{1/2}$. ■

Theorem 4.2.7 (comparing variances). *Take $\omega \neq 0$ and let $D \equiv D_{H,\omega}$ be the length of the longest difference- ω arithmetic sequence in the spectrum of the Hamiltonian H . For any $f \in V_\omega$, and g, h as defined in (4.2.1). Then,*

$$\operatorname{Var}(g) + \operatorname{Var}(h) \geq \frac{1}{D} \operatorname{Var}(f).$$

Proof. Note that $g \in V_0$ and thus g commutes with H (see Proposition 4.1.8) and is Hermitian, so using the same argument as in (4.2.6) in the proof of Proposition 4.2.5, there exists an orthonormal eigenbasis $\{u_i\}$ of H with corresponding eigenvalues $H(u_i) = \langle u_i | H | u_i \rangle$ such that

$$g = e^{\frac{\beta\omega}{4}} \sum_{i=1}^k \lambda_i |u_i\rangle\langle u_i|, \quad f = \sum_{i=1}^k \lambda_i |u_i\rangle\langle v_i|, \quad h = e^{-\frac{\beta\omega}{4}} \sum_{i=1}^k \lambda_i |v_i\rangle\langle v_i| \quad (4.2.8)$$

where $\lambda_i \in \mathbb{R}_{>0}$ for each $i \in [k]$ and $|v_i\rangle = \frac{f^\dagger |u_i\rangle}{\lambda_i}$. Observe that the vectors v_i are pairwise orthogonal, since $v_i^\dagger v_j = (\lambda_i \lambda_j)^{-1} \cdot u_i^\dagger f f^\dagger u_j = (\lambda_i \lambda_j)^{-1} u_i^\dagger g^2 u_j = 0$ if $i \neq j$. In addition, for each $i \in [k]$, the vector v_i has unit norm. Note also that v_i is an eigenvector of H with eigenvalues $H(v_i) = \langle v_i | H | v_i \rangle = H(u_i) - \omega$, since $f = f(\omega) = \sum_{E_1, E_2: E_1 - E_2 = \omega} \Pi_{E_1} f \Pi_{E_2}$ and $\Pi_{E_1} |u_i\rangle = 0$ unless $E_1 = H(u_i)$, we have

$$|v_i\rangle = \lambda_i^{-1} f^\dagger |u_i\rangle = \lambda_i^{-1} \Pi_{H(u_i) - \omega} f^\dagger \Pi_{H(u_i)} |u_i\rangle = \Pi_{H(u_i) - \omega} |v_i\rangle$$

so

$$H|v_i\rangle = (H(u_i) - \omega)|v_i\rangle.$$

Let $\rho(u_i) = \langle u_i | \rho | u_i \rangle$ and $\rho(v_i) = \langle v_i | \rho | v_i \rangle$ be the eigenvalue of ρ corresponding to eigenvectors u_i and v_i . Then, $\rho(v_i) = \rho(u_i) e^{\beta\omega}$. We write

$$\text{Var}(f) = \sum_{i=1}^k \sqrt{\rho(u_i) \rho(v_i)} \lambda_i^2 = e^{\frac{\beta\omega}{2}} \sum_{i=1}^k \rho(u_i) \lambda_i^2 = e^{-\frac{\beta\omega}{2}} \sum_{i=1}^k \rho(v_i) \lambda_i^2 \quad (4.2.9)$$

$$\begin{aligned} \text{Var}(g) &= e^{\frac{\beta\omega}{2}} \left(\sum_{i=1}^k \rho(u_i) \lambda_i^2 - \left(\sum_{i=1}^k \rho(u_i) \lambda_i \right)^2 \right) \\ &\geq \left(1 - \sum_{i=1}^k \rho(u_i) \right) \left(e^{\frac{\beta\omega}{2}} \sum_{i=1}^k \rho(u_i) \lambda_i^2 \right) = \left(1 - \sum_{i=1}^k \rho(u_i) \right) \text{Var}(f) \end{aligned} \quad (4.2.10)$$

where the inequality is due to Hölder's inequality. Similarly, we have

$$\text{Var}(h) = e^{-\frac{\beta\omega}{2}} \left(\sum_{i=1}^k \rho(v_i) \lambda_i^2 - \left(\sum_{i=1}^k \rho(v_i) \lambda_i \right)^2 \right) \geq \left(1 - \sum_{i=1}^k \rho(v_i) \right) \text{Var}(f). \quad (4.2.11)$$

We first consider the case $\beta\omega \geq 0$. We show that

$$1 - \sum_{i=1}^k \rho(u_i) \geq \frac{1}{D} \quad (4.2.12)$$

which together with (4.2.10) gives the desired conclusion that $\text{Var}(g) \geq \frac{1}{D} \text{Var}(f)$.

Let the distinct values appearing in the sequence $(H(u_i))_{i=1}^k$ be $\alpha_1, \dots, \alpha_q$. Partition the set $Q := \{\alpha_1, \dots, \alpha_q\}$ into a family $\mathcal{I}$ of (disjoint) maximally long arithmetic sequences where in each sequence, the difference between two consecutive terms is ω . Such a decomposition can be obtained by iteratively removing the maximally long arithmetic progression from the set Q .

We rewrite

$$Q = \bigcup_{i \in \mathcal{I}} \{\alpha_{i,1}, \dots, \alpha_{i,d_i}\}$$

where for each $i \in \mathcal{I}$, $\alpha_{i,1}, \dots, \alpha_{i,d_i}$ satisfies $\alpha_{i,j} = \alpha_{i,1} - \omega(j-1)$ for all $j \in [d_i]$. Note that $\alpha_{i,d_i+1} := \alpha_{i,d_i} - \omega$ is also an eigenvalue of ρ , since if $u_{\hat{j}}$ is an eigenvector of H with eigenvalue α_{i,d_i} that appears in (4.2.8), then $f^+ u_{\hat{j}}$ is an eigenvector of H with eigenvalue $\alpha_{i,d_i} - \omega$. Thus the spectrum of H contains the length- $(d_i + 1)$ arithmetic progression $\alpha_{i,1}, \dots, \alpha_{i,d_i+1}$ with difference between consecutive terms being ω . Note also that $\alpha_{i,d_i+1} \notin Q$, otherwise this contradicts the maximality of $\alpha_{i,1}, \dots, \alpha_{i,d_i}$.

For an eigenvalue α of H , let $\hat{\alpha} := \frac{\exp(-\beta\alpha)}{\text{tr}(\exp(-\beta H))} \geq 0$ be the corresponding eigenvalue of ρ , $m(\alpha)$ be its multiplicity in the spectrum of H , and $r(\alpha)$ be its multiplicity in the sequence $(H(u_i))_{i=1}^k$. Obviously, $r(\alpha) \leq m(\alpha)$.

We show, for each $i \in \mathcal{I}$,

$$m(\alpha_{i,d_i+1})\hat{\alpha}_{i,d_i+1} + \sum_{j=1}^{d_i} (m(\alpha_{i,j}) - r(\alpha_{i,j}))\hat{\alpha}_{i,j} \geq \frac{1}{d_i + 1} \sum_{j=1}^{d_i+1} m(\alpha_{i,j})\hat{\alpha}_{i,j} \geq \frac{1}{D} \sum_{j=1}^{d_i+1} m(\alpha_{i,j})\hat{\alpha}_{i,j} \quad (4.2.13)$$

where the last inequality is because $d_i + 1 \leq D$, due to the assumption on the spectrum of H and the fact that H contains the length- $(d_i + 1)$ arithmetic progression $\alpha_{i,1}, \dots, \alpha_{i,d_i+1}$ with difference between consecutive terms being ω .

Before showing (4.2.13), we check that it implies the desired inequality (4.2.12).

$$\begin{aligned} 1 - \sum_{j=1}^k \rho(u_j) &=_{(1)} \sum_{\alpha} m(\alpha)\hat{\alpha} - \sum_{i \in \mathcal{I}, j \in [d_i]} r(\alpha_{i,j})\hat{\alpha}_{i,j} \\ &= \sum_{i \in \mathcal{I}} \sum_{j=1}^{d_i+1} m(\alpha_{i,j})\hat{\alpha}_{i,j} + \sum_{\alpha \notin \{\alpha_{i,j}: i \in \mathcal{I}, j \in [d_i+1]\}} m(\alpha)\hat{\alpha} - \sum_{i \in \mathcal{I}} \sum_{j=1}^{d_i} r(\alpha_{i,j})\hat{\alpha}_{i,j} \\ &= \sum_{i \in \mathcal{I}} \left(m(\alpha_{i,d_i+1})\hat{\alpha}_{i,d_i+1} + \sum_{j=1}^{d_i} (m(\alpha_{i,j}) - r(\alpha_{i,j}))\hat{\alpha}_{i,j} \right) + \sum_{\alpha \notin \{\alpha_{i,j}: i \in \mathcal{I}, j \in [d_i+1]\}} m(\alpha)\hat{\alpha} \\ &\geq_{(2)} \frac{1}{D} \sum_{\alpha} m(\alpha)\hat{\alpha} \end{aligned}$$

$$= \frac{1}{D}$$

where (1) and the final equality is by $1 = \text{tr}(\rho) = \sum_{\alpha} m(\alpha) \hat{\alpha}$, and (2) is by (4.2.13).

Now, we show (4.2.13). Recall that for all $j \in [d_i]$, $r(\alpha_{i,j}) \leq m(\alpha_{i,j})$. Furthermore, for all $j \in [d_i]$, $r(\alpha_{i,j}) \leq m(\alpha_{i,j+1})$. Indeed, fix one such j . Let $u_{h_1}, \dots, u_{h_r}$ be the eigenvectors of H from (4.2.8) with eigenvalue $\alpha_{i,j}$. Then $f^+ u_{h_1}, \dots, f^+ u_{h_r}$ are pairwise orthogonal (and thus linearly independent) eigenvectors of H with eigenvalue $\alpha_{i,j} - \omega = \alpha_{i,j+1}$, hence $r(\alpha_{i,j}) \leq m(\alpha_{i,j+1})$.

For $j^* = \arg \max_{j \in [d_i+1]} m(\alpha_{i,j}) \hat{\alpha}_{i,j}$, we have

$$\begin{aligned} & m(\alpha_{i,d_i+1}) \hat{\alpha}_{i,d_i+1} + \sum_{j=1}^{d_i} (m(\alpha_{i,j}) - r(\alpha_{i,j})) \hat{\alpha}_{i,j} \\ &= \sum_{j=1}^{j^*-1} (m(\alpha_{i,j}) - r(\alpha_{i,j})) \hat{\alpha}_{i,j} + m(\alpha_{i,j^*}) \hat{\alpha}_{i,j^*} + \sum_{j=j^*}^{d_i} (m(\alpha_{i,j+1}) \hat{\alpha}_{i,j+1} - r(\alpha_{i,j}) \hat{\alpha}_{i,j}) \\ &\geq m(\alpha_{i,j^*}) \hat{\alpha}_{i,j^*} \end{aligned}$$

where the inequality is due to $(m(\alpha_{i,j}) - r(\alpha_{i,j})) \hat{\alpha}_{i,j} \geq 0$, and

$$m(\alpha_{i,j+1}) \hat{\alpha}_{i,j+1} - r(\alpha_{i,j}) \hat{\alpha}_{i,j} = \hat{\alpha}_{i,j} (e^{\beta\omega} m(\alpha_{i,j+1}) - r(\alpha_{i,j})) \geq \hat{\alpha}_{i,j} (m(\alpha_{i,j+1}) - r(\alpha_{i,j})) \geq 0,$$

where we use the assumption $\beta\omega \geq 0$ and the fact that $m(\alpha_{i,j+1}) \geq r(\alpha_{i,j})$. Thus we have obtained

$$m(\alpha_{i,d_i+1}) \hat{\alpha}_{i,d_i+1} + \sum_{j=1}^{d_i} (m(\alpha_{i,j}) - r(\alpha_{i,j})) \hat{\alpha}_{i,j} \geq m(\alpha_{i,j^*}) \hat{\alpha}_{i,j^*} \geq \frac{1}{d_i+1} \sum_{j=1}^{d_i+1} m(\alpha_{i,j}) \hat{\alpha}_{i,j}.$$

and conclude that $\text{Var}(g) \geq \frac{1}{D} \text{Var}(f)$, which implies $\text{Var}(g) + \text{Var}(h) \geq \frac{1}{D} \text{Var}(f)$.

The case $\beta\omega \leq 0$ follows similarly, where we prove $\text{Var}(h) \geq \frac{1}{D} \text{Var}(f)$. Thus, in either case, we have

$$\text{Var}(g) + \text{Var}(h) \geq \frac{1}{D} \text{Var}(f). \quad \blacksquare$$

We have the following alternative comparison between the variances.

Lemma 4.2.14. *For any $f \in V_{\omega}$, and g, h as defined in (4.2.1),*

$$\text{Var}(g) + \text{Var}(h) \geq (1 - e^{-|\beta\omega|}) \text{Var}(f).$$

Proof. From the proof of Theorem 4.2.7, we have

$$\text{Var}(g) \geq \left(1 - \sum_{i=1}^k \rho(u_i)\right) \text{Var}(f) \quad \text{and} \quad \text{Var}(h) \geq \left(1 - \sum_{i=1}^k \rho(v_i)\right) \text{Var}(f)$$

where for $i \in [k]$, u_i, v_i are eigenvectors of the stationary state ρ with corresponding eigenvalues $\rho(u_i), \rho(v_i)$ respectively, and $\rho(v_i) = \rho(u_i)e^{\beta\omega}$. Moreover, $\{u_i\}_{i=1}^k$ are linearly independent and $\{v_i\}_{i=1}^k$ are linearly independent, so $\max\{\sum_{i=1}^k \rho(u_i), \sum_{i=1}^k \rho(v_i)\} \leq \text{tr}(\rho) = 1$. If $\beta\omega \geq 0$ then

$$e^{\beta\omega} \sum_{i=1}^k \rho(u_i) = \sum_{i=1}^k \rho(v_i) \leq 1$$

so

$$\text{Var}(g) = \left(1 - \sum_{i=1}^k \rho(u_i)\right) \text{Var}(f) \geq (1 - e^{-|\beta\omega|}) \text{Var}(f).$$

Similarly, if $\beta\omega \leq 0$ then

$$e^{-\beta\omega} \sum_{i=1}^k \rho(v_i) = \sum_{i=1}^k \rho(u_i) \leq 1 \Rightarrow \text{Var}(h) = \left(1 - \sum_{i=1}^k \rho(v_i)\right) \text{Var}(f) \geq (1 - e^{-|\beta\omega|}) \text{Var}(f). \quad \blacksquare$$

Proof of Theorem 4.0.12. Fix an arbitrary $\omega \neq 0$ and an arbitrary $f \in V_\omega$. Let g, h be as defined in Theorem 4.2.2; then, $g, h \in V_0$. We have

$$2\mathcal{E}(f) \geq_{(1)} \mathcal{E}(g) + \mathcal{E}(h) \geq_{(2)} \lambda_{\mathcal{L},0}(\text{Var}(g) + \text{Var}(h)) \geq_{(3)} \lambda_{\mathcal{L},0} \max\left(\frac{1}{D}, 1 - e^{-|\beta\omega|}\right) \text{Var}(f)$$

where (1) is by Theorem 4.2.2, (2) by definition of $\lambda_{\mathcal{L},0}$, and (3) by Theorem 4.2.7 and Lemma 4.2.14. Thus

$$\lambda_{\mathcal{L},\omega} = \min_{f \in V_\omega} \frac{\mathcal{E}(f)}{\text{Var}(f)} \geq \frac{1}{2} \max\left(\frac{1}{D}, 1 - e^{-|\beta\omega|}\right) \lambda_{\mathcal{L},0}. \quad \blacksquare$$

Proof of Theorem 4.0.11. By Proposition 4.1.5,

$$\lambda_{\mathcal{L},0} \geq \lambda_{\mathcal{L}} = \min_{\omega} \lambda_{\mathcal{L},\omega} \geq \frac{1}{2D} \lambda_{\mathcal{L},0}$$

where the last inequality follows from: for all $\omega \neq 0$, $\lambda_{\mathcal{L},\omega} \geq \frac{1}{2D} \lambda_{\mathcal{L},0}$, by Theorem 4.0.12. \blacksquare

4.3 Relating quantum to classical spectral gap

In this section, we prove Proposition 4.0.13 and Theorem 4.0.14, which relate the quantum spectral gap to the spectral gap of certain classical Markov generators.

First, we prove Proposition 4.0.13. By Proposition 4.1.10, $\lambda_{\mathcal{L},0} = \lambda_{\mathcal{L},0}^H$, so we only need to show the following.

Proposition 4.3.1. *For any Hamiltonian H and Davies generator $\mathcal{L}$,*

$$\lambda_{\mathcal{L},0}^H = \min_U \lambda_{\mathcal{L},U,\text{cl}} \quad (4.3.2)$$

where the minimum is taken over the orthonormal eigenbasis U of the Hamiltonian H , and $\lambda_{\mathcal{L},U,\text{cl}}$ is the classical spectral gap associated with $\mathcal{L}$ and the orthonormal eigenbasis U , as defined in Definition 4.0.8.

Proof. Consider an arbitrary Hermitian $f \in V_0$. Note that f commutes with H and is Hermitian, so f and H are simultaneously diagonalizable by a certain orthonormal eigenbasis U of H . Let u_i be the i -th eigenvector in U , $\pi(u_i) = \langle u_i | \rho | u_i \rangle$ and write $f = \sum f_i u_i u_i^\dagger$ with $f_i \in \mathbb{R}$. Define $F : U \rightarrow \mathbb{R}$ where $F(u_i) = f_i$. We will show that the Dirichlet form of f with respect to $\mathcal{L}$ (respectively, the variance of f) is equal to the Dirichlet form of F (respectively, the variance of F) with respect to the classical chain P_U defined in Definition 4.0.8. Concretely,

$$\mathcal{E}_{\mathcal{L}}(f) = \mathcal{E}_{\mathcal{L},U,\text{cl}}(F) \quad \text{and} \quad \text{Var}_{\rho}(f) = \text{Var}_{\pi}(F).$$

Hence,

$$\lambda_{\mathcal{L},0}^H = \min_U \min_{f = \sum_i f_i u_i u_i^\dagger, f_i \in \mathbb{R} \forall i} \frac{\mathcal{E}_{\mathcal{L}}(f)}{\text{Var}_{\rho}(f)} = \min_U \min_{F:U \rightarrow \mathbb{R}} \frac{\mathcal{E}_{\mathcal{L},U,\text{cl}}(F)}{\text{Var}_{\pi}(F)} = \min_U \lambda_{\mathcal{L},U,\text{cl}}$$

where the outside minimum is taken over orthonormal eigenbases U of H .

Note that $\{|u_i\rangle\langle u_j|\}_{i,j \in [N]}$ forms an orthonormal basis of $\mathbb{C}^{N \times N}$ with respect to the inner product $\langle A, B \rangle = \text{tr}(AB^\dagger)$. Note that this is distinct from the KMS-product $\langle A, B \rangle_{\rho} = \text{tr}(\rho^{1/2} A \rho^{1/2} B^\dagger)$. Since $\mathcal{L}$ is a linear map, we can write

$$\mathcal{L}(f) = \sum_{i,j,k} f_i \langle \mathcal{L}(|u_i\rangle\langle u_i|), |u_j\rangle\langle u_k| \rangle |u_j\rangle\langle u_k|.$$

We can also directly compute that for $i \neq j$, and $E_i := \langle u_i | H | u_i \rangle$, $E_j := \langle u_j | H | u_j \rangle$,

$$\langle \mathcal{L}(|u_j\rangle\langle u_j|), |u_i\rangle\langle u_i| \rangle = G(E_j - E_i) \sum_{S \in \mathcal{S}} (\langle u_i | S | u_j \rangle)^2 = P_{\mathcal{L},U}[u_i \rightarrow u_j]$$

and

$$-\langle \mathcal{L}(|u_i\rangle\langle u_i|), |u_i\rangle\langle u_i| \rangle = \sum_{j:j \neq i} \langle \mathcal{L}(|u_j\rangle\langle u_j|), |u_i\rangle\langle u_i| \rangle = \sum_{j:j \neq i} P_{\mathcal{L},U}[u_i \rightarrow u_j]. \quad (4.3.3)$$

$$\begin{aligned} -\mathcal{E}_{\mathcal{L}}(f) &= \langle \mathcal{L}(f), f \rangle_{\rho} = \left\langle \sum_{i,j,k} f_i \langle \mathcal{L}(|u_i\rangle\langle u_i|), |u_j\rangle\langle u_k| \rangle |u_j\rangle\langle u_k|, \sum_r f_r |u_r\rangle\langle u_r| \right\rangle_{\rho} \\ &= \sum_{i,j,k,r} f_i f_r \langle \mathcal{L}(|u_i\rangle\langle u_i|), |u_j\rangle\langle u_k| \rangle \langle |u_j\rangle\langle u_k|, |u_r\rangle\langle u_r| \rangle_{\rho} \end{aligned} \quad (4.3.4)$$

The reversibility condition of the classical Markov generator follows from the reversibility of $\mathcal{L}$ (see (4.0.2)). Indeed,

$$\begin{aligned} \pi(u_i) \langle \mathcal{L}(|u_j\rangle\langle u_j|), |u_i\rangle\langle u_i| \rangle &= \langle \mathcal{L}(|u_j\rangle\langle u_j|), |u_i\rangle\langle u_i| \rangle_{\rho} \\ &= \langle \mathcal{L}(|u_i\rangle\langle u_i|), |u_j\rangle\langle u_j| \rangle_{\rho} \\ &= \pi(u_j) \langle \mathcal{L}(|u_i\rangle\langle u_i|), |u_j\rangle\langle u_j| \rangle. \end{aligned} \quad (4.3.5)$$

Since u_r is an eigenvector of ρ with corresponding eigenvalue $\rho(u_r) = \langle u_r | \rho | u_r \rangle = \pi(u_r)$, we obtain:

$$\begin{aligned} \langle |u_j\rangle\langle u_k|, |u_r\rangle\langle u_r| \rangle_{\rho} &= \text{tr}(|u_j\rangle\langle u_k| \rho^{1/2} |u_r\rangle\langle u_r| \rho^{1/2}) = \text{tr}(|u_j\rangle\langle u_k| (\rho(u_r)^{1/2} |u_r\rangle\langle u_r| \rho(u_r)^{1/2})) \\ &= \rho(u_r) \text{tr}(|u_j\rangle\langle u_k| |u_r\rangle\langle u_r|) = \rho(u_r) \langle u_k | u_r \rangle \langle u_r | u_j \rangle = \pi(u_r) \delta_{k=r} \delta_{j=r} \end{aligned}$$

Substituting into (4.3.4) gives

$$\begin{aligned} \mathcal{E}_{\mathcal{L}}(f) &= - \sum_{i,j} f_i f_j \pi(u_j) \langle \mathcal{L}(|u_i\rangle\langle u_i|), |u_j\rangle\langle u_j| \rangle \\ &= \sum_i f_i^2 \pi(u_i) \sum_{j:j \neq i} P_{\mathcal{L},U}[u_i \rightarrow u_j] - \sum_{i \neq j} f_i f_j \pi(u_i) P_{\mathcal{L},U}[u_i \rightarrow u_j] \\ &= \frac{1}{2} \sum_{i \neq j} (f_i^2 + f_j^2 - 2f_i f_j) \pi(u_i) P_{\mathcal{L},U}[u_i \rightarrow u_j] \\ &= \frac{1}{2} \sum_{i \neq j} (f_i - f_j)^2 \pi(u_i) P_{\mathcal{L},U}[u_i \rightarrow u_j] = \mathcal{E}_{\mathcal{L},U,\text{cl}}(F) \end{aligned}$$

where in the second equality we use (4.3.3) and in the third we use the reversibility of the classical Markov generator. Next,

$$\text{Var}_{\rho}(f) = \text{tr}(\rho^{1/2} f \rho^{1/2} f^{\dagger}) - \text{tr}(\rho f)^2 = \sum_i f_i^2 \langle u_i | \rho | u_i \rangle - \left(\sum_i f_i \langle u_i | \rho | u_i \rangle \right)^2$$

$$\begin{aligned}
 &= \sum_i F(u_i)^2 \pi(u_i) - \left(\sum_i F(u_i) \pi(u_i) \right)^2 \\
 &= \text{Var}_\pi(F).
 \end{aligned}$$

■

Proof of Theorem 4.0.14. Let $\hat{U}$ be an orthonormal eigenbasis of H satisfying $\hat{U} = \arg \min_U \lambda_{\mathcal{L}, U, \text{cl}}$. Then, by Proposition 4.0.13,

$$\lambda_{\mathcal{L}, 0} = \lambda_{\mathcal{L}, 0}^H = \lambda_{\mathcal{L}, \hat{U}, \text{cl}},$$

and by Theorem 4.0.11

$$\lambda_{\mathcal{L}, \hat{U}, \text{cl}} \geq \lambda_{\mathcal{L}} \geq \frac{1}{2D} \lambda_{\mathcal{L}, \hat{U}, \text{cl}}.$$

■

To prove Corollary 4.0.15, we use a generalized version of Cheeger's inequality that applies to general Markov generators. Although the proof is essentially the same as in the standard case, we include it in Chapter C for completeness.

Proposition 4.3.6. *Consider a classical Markov generator on a state space X , with transition rates $P[x \rightarrow y] \geq 0$ for all $x \neq y \in X$. Assume that the generator is reversible with respect to a probability distribution π on X ,¹ that is, for all $x \neq y$,*

$$\pi(x)P[x \rightarrow y] = \pi(y)P[y \rightarrow x].$$

Recall that the Dirichlet form is given by $\mathcal{E}(f) = \frac{1}{2} \sum_{x \neq y} (f(x) - f(y))^2 \pi(x)P[x \rightarrow y]$, and the spectral gap is $\lambda = \inf_{f: X \rightarrow \mathbb{R}} \frac{\mathcal{E}(f)}{\text{Var}_\pi(f)}$.

Define the bottleneck ratio $\Phi := \min_{S \subseteq X, \pi(S) \leq 1/2} \frac{Q(S, S^c)}{\pi(S)}$, where

$$Q(S, S^c) = \sum_{x \in S, y \notin S} \pi(x)P[x \rightarrow y].$$

Let $\Lambda := \sup_{x \in X} \sum_{y \neq x} P[x \rightarrow y]$ denote the maximum total transition rate out of any state. Then the following inequality holds:

$$\Phi^2 \leq 2\Lambda\lambda.$$

Proof of Corollary 4.0.15. By Theorem 4.0.14, for a certain orthonormal eigenbasis $\hat{U} = \{u_i\}_i$ of H ,

$$\lambda_{\mathcal{L}, \hat{U}, \text{cl}} \leq 2D\lambda_{\mathcal{L}}. \quad (4.3.7)$$

¹ π is not assumed to be the unique stationary distribution of the process.

We adopt the same notation as in Definition 4.0.8 and Proposition 4.0.13.

Let $\Lambda := \max_{u_i \in \hat{U}} \sum_{u_j: j \neq i} P_{\mathcal{L}, \hat{U}}[u_i \rightarrow u_j]$. We prove that

$$\Lambda \leq M = \|G\|_\infty \cdot \left\| \sum_{S \in \mathcal{S}} SS^\dagger \right\|. \quad (4.3.8)$$

Indeed,

$$\begin{aligned} \sum_{u_j: j \neq i} P_{\mathcal{L}, \hat{U}}[u_i \rightarrow u_j] &= \sum_{u_j: j \neq i} \left(G(E_j - E_i) \sum_{S \in \mathcal{S}} \langle u_i | S | u_j \rangle \langle u_j | S^\dagger | u_i \rangle \right) \\ &\leq \|G\|_\infty \sum_{u_j} \sum_{S \in \mathcal{S}} \langle u_i | S | u_j \rangle \langle u_j | S^\dagger | u_i \rangle \\ &= \|G\|_\infty \langle u_i | \left(\sum_{S \in \mathcal{S}} S \left(\sum_{u_j} |u_j\rangle \langle u_j| \right) S^\dagger \right) | u_i \rangle \\ &= \|G\|_\infty \langle u_i | \left(\sum_{S \in \mathcal{S}} SS^\dagger \right) | u_i \rangle \\ &\leq \|G\|_\infty \cdot \left\| \sum_{S \in \mathcal{S}} SS^\dagger \right\| \end{aligned}$$

As argued in Proposition 4.0.13, the classical Markov generator $P_{\mathcal{L}, \hat{U}}$ is reversible wrt π . By Cheeger's inequality (see Proposition 4.3.6), there exists a set $S \subseteq \hat{U}$ such that:

$$\pi(S) \leq 1/2 \quad \text{and} \quad \sum_{u_i \in S, u_j \notin S} \pi(u_i) P_{\mathcal{L}, \hat{U}}[u_i \rightarrow u_j] \leq \pi(S) \sqrt{2\Lambda \lambda_{\mathcal{L}, \hat{U}, \text{cl}}}. \quad (4.3.9)$$

Let $F : \hat{U} \rightarrow \mathbb{R}$ be defined by $F(u_i) = [u_i \in S]$ and let $f = \sum F(u_i) |u_i\rangle \langle u_i|$. Observe that f is a Hermitian projector that commutes with H i.e. $f = f^\dagger$, $[f, H] = 0$ and $f^2 = f$. By the same argument as in the proof of Proposition 4.0.13 and a direct calculation, we have:

$$\mathcal{E}_{\mathcal{L}}(f) = \mathcal{E}_{\mathcal{L}, U, \text{cl}}(F) = \sum_{u_i \in S, u_j \notin S} \pi(u_i) P_{\mathcal{L}, \hat{U}}[u_i \rightarrow u_j] \quad \text{and} \quad \text{Var}_\rho(f) = \text{Var}_\pi(F) = (1 - \pi(S))\pi(S).$$

Combining (4.3.7), (4.3.9), and (4.3.8), we get

$$\begin{aligned} \mathcal{E}_{\mathcal{L}}(f) &\leq \sqrt{2\Lambda \lambda_{\mathcal{L}, \hat{U}, \text{cl}}} \cdot \pi(S) \leq 2\sqrt{2M \lambda_{\mathcal{L}, \hat{U}, \text{cl}}} \cdot \pi(S)(1 - \pi(S)) \\ &= 2\sqrt{2M \lambda_{\mathcal{L}, \hat{U}, \text{cl}}} \cdot \text{Var}_\rho(f) \\ &\leq 4\sqrt{DM \lambda_{\mathcal{L}}} \cdot \text{Var}_\rho(f). \end{aligned} \quad \blacksquare$$

4.4 Absence of long proper arithmetic progressions

In this section, we prove Theorem 4.0.16, i.e., perturbing any fixed Hamiltonian by a generic external field results in a new Hamiltonian whose spectrum contains no 3-arithmetic progressions or repeated eigenvalues. We can also show similar results for other types of perturbations. For example, we show that perturbing any fixed Hamiltonian by a generic degree- k perturbation yields a new Hamiltonian whose spectrum contains no 3-arithmetic progression.

Proposition 4.4.1. *For any Hamiltonian H_0 and $k \in \mathbb{N}$, the spectrum of the Hamiltonian $H := H_0 + \sum_{S \in \binom{[n]}{k}} A_S \prod_{i \in S} X_i$ almost surely (with respect to the randomness of $A = (A_S)_{S \in \binom{[n]}{k}} \in \mathbb{R}^{\binom{[n]}{k}}$) does not have any 3-arithmetic progressions, in the sense that the set of $A \in \mathbb{R}^{\binom{[n]}{k}}$ such that H has a 3-arithmetic progressions has Lebesgue measure 0.*

For quadratic perturbation, we can also restrict the domain of the interactions to be supported on an arbitrary connected graph.

Proposition 4.4.2. *For any given Hamiltonian H_0 and (simple) connected graph $G = G([n], E)$, the spectrum of the Hamiltonian $H := H_0 + \sum_{\{i,j\} \in E(G)} A_{ij} X_i X_j$ almost surely (with respect to the randomness of $A = (A_{ij})_{\{i,j\} \in E(G)} \in \mathbb{R}^{E(G)}$) does not have any 3-arithmetic progressions, in the sense that the set of $A \in \mathbb{R}^{E(G)}$ such that H has a 3-arithmetic progressions has Lebesgue measure 0.*

As corollaries, we obtain the following results regarding the spectrum of the transverse field Ising model.

Corollary 4.4.3. *For any simple connected graph $G = G([n], E)$ and fixed external field H , the spectrum of the Hamiltonian $H = \sum_{\{i,j\} \in E(G)} A_{ij} Z_i Z_j + h \sum_i X_i$ almost surely (with respect to the randomness of $A = (A_{ij})_{\{i,j\} \in E(G)}$) does not have any 3-arithmetic progressions, in the sense that the set of $A \in \mathbb{R}^{E(G)}$ such that the spectrum of H has a 3-arithmetic progressions has Lebesgue measure 0.*

Corollary 4.4.4. *For any interaction matrix $A = (A_{ij})_{i,j}$, the spectrum of the Hamiltonian $H = \sum_{i,j} A_{ij} Z_i Z_j + \sum_i h_i X_i$ almost surely (with respect to the randomness of $h = (h_i)_{i=1}^n \in \mathbb{R}^n$) does not have any 3-arithmetic progressions or repeated eigenvalues, in the sense that the set of $h \in \mathbb{R}^n$ such that the spectrum of H has any 3-arithmetic progressions or repeated eigenvalues has Lebesgue measure 0.*

The above shows that the spectrum of Hamiltonians with *varying* interactions or external fields generically does not contain a long arithmetic progression. We can also show that this property holds for some Hamiltonians with uniform interaction and external field.

Proposition 4.4.5. *Let H be the 1D XY+Z Hamiltonian with periodic boundary condition, i.e.,*

$$H = J \sum_{i=1}^n X_i Y_{i+1} + h \sum_{i=1}^n Z_i \quad \text{where} \quad Y_{n+1} = Y_1.$$

For an odd prime n , the spectrum of H almost surely (with respect to the randomness of J and h) does not contain any 3-arithmetic progressions or repeated eigenvalues, in the sense that the set of $(J, h) \in \mathbb{R}^2$ such that the spectrum of H has 3-arithmetic progression or repeated eigenvalues or repeated eigenvalues has Lebesgue measure 0.¹

We can also show that perturbing any fixed Hamiltonian H_0 by the 1D XY+Z Hamiltonian yields a Hamiltonian whose spectrum does not contain any 3-arithmetic progression or repeated eigenvalues. This class includes several translationally invariant 1D models studied in the literature (see [KLW'15]).

Proposition 4.4.6. *For any given Hamiltonian H_0 , the spectrum of the Hamiltonian*

$$H := H_0 + J \sum_{i=1}^n X_i Y_{i+1} + h \sum_{i=1}^n Z_i \quad \text{where} \quad Y_{n+1} = Y_1$$

almost surely (with respect to the randomness of J and h) does not contain any 3-arithmetic progression or repeated eigenvalues, in the sense that the set of $(J, h) \in \mathbb{R}^2$ such that the spectrum H has any 3-arithmetic progressions or repeated eigenvalues has Lebesgue measure 0.

We first prove Theorem 4.0.16, which will be a model for the other proofs in this section.

Proof of Theorem 4.0.16. We use a similar strategy as in [HH'25]. Let the eigenvalues of H be $\lambda_1, \lambda_2, \dots, \lambda_N$. Let

$$F(h_1, \dots, h_n) := \prod_{i,j,k \text{ distinct}} (\lambda_i + \lambda_j - 2\lambda_k), \quad G(h_1, \dots, h_n) := \prod_{i,j:i \neq j} (\lambda_i - \lambda_j).$$

Note that the spectrum of H has a 3-term arithmetic progression if and only if $F = 0$, and similarly, the spectrum H has eigenvalues with multiplicities if and only if $G = 0$. Since F and G are symmetric under permutation of the eigenvalues of H , they are symmetric polynomials in the eigenvalues $\lambda_1, \dots, \lambda_N$, and thus by the fundamental theorem of symmetric polynomials, can be written as a polynomial in terms of $\sum_{i=1}^N \lambda_i^k = \text{tr}(H^k)$ for $k \in [N]$. Since $\text{tr}(H^k)$ are polynomials in $h_1, \dots, h_n$ (treating the entries of H_0 as fixed constants), F and G are polynomials in $h_1, \dots, h_n$. Since the zeros

¹For odd primes n , H almost surely does not have repeated eigenvalues; this was proved in [KLW'15, Lemma 1].

of a multivariate polynomial are of measure zero unless the polynomial is identically zero, it suffices to find a particular $h = (h_i)_{i=1}^n$ such that $F(h)$ and $G(h)$ are nonzero. Pick $R > 0$, and let $\tilde{h}_i := R^{-1}h_i$ and $\tilde{H}_{\tilde{h}} := R^{-1}H_0 + \sum_{i=1}^n \tilde{h}_i P_i$. Then, the eigenvalues of $\tilde{H}_{\tilde{h}}$ are $\tilde{\lambda}_i := R^{-1}\lambda_i$ for $i \in [N]$. We can rewrite

$$F(h) = R^{N(N-1)(N-2)} \prod_{i,j,k \text{ distinct}} (\tilde{\lambda}_i + \tilde{\lambda}_j - 2\tilde{\lambda}_k),$$

$$G(h) = R^{N(N-1)} \prod_{i,j:i \neq j} (\tilde{\lambda}_i - \tilde{\lambda}_j).$$

Let the eigenvalues of the Hermitian matrix $P \in \mathbb{C}^{2 \times 2}$ be $c_1, c_2 \in \mathbb{R}$, where $c_1 \neq c_2$, since P is not a multiple of the identity.

Note that the eigenvalues of $\sum_{i=1}^n \tilde{h}_i P_i$ are $\lambda_j^* = \frac{c_1 - c_2}{2} \sum_{i=1}^n (2x_i - 1) \tilde{h}_i + C$ where $x_1 \dots x_n \in \{0, 1\}^n$ is the binary representation of j , and $C = \frac{\lambda_1 + \lambda_2}{2} \sum_i \tilde{h}_i$. Keeping $\tilde{h}$ fixed while taking $R \rightarrow +\infty$, $\prod_{i,j,k \text{ distinct}} (\tilde{\lambda}_i + \tilde{\lambda}_j - 2\tilde{\lambda}_k)$ (and $\prod_{i,j:i \neq j} (\tilde{\lambda}_i - \tilde{\lambda}_j)$, respectively) converges to $\prod_{i,j,k \text{ distinct}} (\lambda_i^* + \lambda_j^* - 2\lambda_k^*)$ ($\prod_{i,j:i \neq j} (\lambda_i^* - \lambda_j^*)$ respectively), so it suffices to exhibit a particular $\tilde{h}$ such that $\prod_{i,j,k \text{ distinct}} (\lambda_i^* + \lambda_j^* - 2\lambda_k^*) \neq 0$ and $\prod_{i,j:i \neq j} (\lambda_i^* - \lambda_j^*) \neq 0$, which is equivalent to $\sum_{i=1}^n \tilde{h}_i P_i$ not having any 3-term arithmetic progressions or repeated eigenvalues in its spectrum. The existence of such an $\tilde{h}$ is guaranteed by the following proposition. ■

Proposition 4.4.7. *If $\tilde{h}_1, \dots, \tilde{h}_n$ are linearly independent over $\mathbb{Z}$, then $\sum_{i=1}^n \tilde{h}_i P_i$ does not have any 3-term arithmetic progressions or repeated eigenvalues in its spectrum.*

Proof. The eigenvalues of $\sum_{i=1}^n \tilde{h}_i P_i$ are $\frac{c_1 - c_2}{2} \sum_{i=1}^n z_i \tilde{h}_i + C$ for $z \in \{\pm 1\}^n$ and $C = \frac{c_1 + c_2}{2} \sum_i \tilde{h}_i$, where $c_1, c_2 \in \mathbb{R}$ are the eigenvalues of P , and note that $c_1 \neq c_2$ since P is not a multiple of the identity.

Suppose for contradiction that $\sum_{i=1}^n P_i \tilde{h}_i$ has repeated eigenvalues. Then there exist distinct $v, z \in \{\pm 1\}^n$ such that

$$0 = \sum_{i=1}^n v_i \tilde{h}_i - \sum_{i=1}^n z_i \tilde{h}_i = \sum_{i=1}^n (v_i - z_i) \tilde{h}_i.$$

Since v, z are distinct vectors in $\{\pm 1\}^n$, $(v_i - z_i)_{i=1}^n \in \{0, \pm 2\}^n \subseteq \mathbb{Z}^n$ and is not the zero vector, so $(\tilde{h}_i)_{i=1}^n$ are not linearly independent over $\mathbb{Z}$, a contradiction.

Suppose for contradiction that there exists a 3-term arithmetic progression, i.e., there exist (pairwise) distinct $v, w, z \in \{\pm 1\}^n$ such that

$$0 = \sum_{i=1}^n v_i \tilde{h}_i + \sum_{i=1}^n w_i \tilde{h}_i - 2 \sum_{i=1}^n z_i \tilde{h}_i = \sum_{i=1}^n (v_i + w_i - 2z_i) \tilde{h}_i.$$

Since v and w are distinct, there exists i such that $v_i \neq w_i$, and thus $v_i + w_i - 2z_i = 1 + (-1) - 2z_i = -2z_i \neq 0$. So $(v_i + w_i - 2z_i)_{i=1}^n \in \{0, \pm 2, \pm 4\}^n \subseteq \mathbb{Z}^n$ and is not the zero vector, contradicting that $(\tilde{h}_i)_{i=1}^n$ are linearly independent over $\mathbb{Z}$. ■

Proof of Proposition 4.4.2. By the same argument as in proof of Theorem 4.0.16, the statement reduces to the following proposition. ■

Proposition 4.4.8. *If G is connected and $(A_{ij})_{\{i,j\} \in E(G)}$ are linearly independent over $\mathbb{Z}$, then the spectrum of $\sum_{\{i,j\} \in E(G)} A_{ij} X_i X_j$ does not have any 3-term arithmetic progression.*

Proof. Note that $\sum_{\{i,j\} \in E(G)} A_{ij} X_i X_j$ is equivalent to $\sum_{\{i,j\} \in E(G)} A_{ij} Z_i Z_j$ by a unitary transformation, so we only need to work with $\sum_{\{i,j\} \in E(G)} A_{ij} Z_i Z_j$.

The computational basis forms an eigenbasis of $\sum_{\{i,j\} \in E(G)} A_{ij} Z_i Z_j$ where the eigenvalue corresponding to $y \in \{0, 1\}^n$ is $\sum_{\{i,j\} \in E(G)} A_{ij} (2y_i - 1)(2y_j - 1)$. The eigenvalues of $\sum_{\{i,j\} \in E(G)} A_{ij} Z_i Z_j$ are thus $\sum_{\{i,j\} \in E(G)} A_{ij} z_i z_j$ for $z \in \{\pm 1\}^n$. Suppose for contradiction that there exists a 3-term arithmetic progression, i.e. there exist pairwise distinct $v, w, z \in \{\pm 1\}^n$ such that

$$0 = \sum_{\{i,j\} \in E(G)} v_i v_j A_{ij} + \sum_{\{i,j\} \in E(G)} w_i w_j A_{ij} - 2 \sum_{\{i,j\} \in E(G)} z_i z_j A_{ij} = \sum_{\{i,j\} \in E(G)} (v_i v_j + w_i w_j - 2z_i z_j) A_{ij}.$$

It is easy to see that $v_i v_j + w_i w_j - 2z_i z_j \in \{0, \pm 2, \pm 4\}$, so for linearly independent (A_{ij}) , the above is true iff $v_i v_j + w_i w_j - 2z_i z_j = 0$ for all $\{i, j\} \in E(G)$. This implies that $v_i v_j = w_i w_j = z_i z_j$ for all $\{i, j\} \in E(G)$.

Indeed, suppose there exists $\{i, j\} \in E(G)$ such that $v_i v_j \neq w_i w_j$ then since $v_i v_j, w_i w_j \in \{\pm 1\}$, $v_i v_j + w_i w_j - 2z_i z_j = 0 - 2z_i z_j \neq 0$, a contradiction. Thus for all $\{i, j\} \in E(G)$: $v_i v_j = w_i w_j$, which then implies $v_i v_j = w_i w_j = z_i z_j$ since $v_i v_j + w_i w_j - 2z_i z_j = 0$.

Fix an arbitrary $i \in [n]$. By the pigeonhole principle, since $v_i, w_i, z_i \in \{\pm 1\}$, either $v_i = w_i$ or $v_i = z_i$. Wlog assume $v_i = w_i$. Then for any j such that $\{i, j\} \in E(G)$, $v_j = w_j$ (since $v_i v_j = w_i w_j$). By this same argument, any j such that there is a path in G from i to j has $v_j = w_j$, and this includes all vertices in G since G is connected. Thus v is identical to w , contradicting that v and w are distinct vectors in $\{\pm 1\}^n$. ■

Proof of Proposition 4.4.1. By the same argument as in proof of Theorem 4.0.16, the statement reduces to the following proposition. ■

Proposition 4.4.9. *If $(A_S)_{S \in \binom{[n]}{k}}$ are linearly independent over $\mathbb{Z}$, then the spectrum of*

$$\sum_{S \in \binom{[n]}{k}} A_S \prod_{i \in S} X_i$$

does not have any 3-term arithmetic progression.

Proof. Note that $\sum_{S \in \binom{[n]}{k}} A_S \prod_{i \in S} X_i$ is unitarily equivalent to $\sum_{S \in \binom{[n]}{k}} A_S \prod_{i \in S} Z_i$, so it suffices to work with the latter. The computational basis forms an eigenbasis for $\sum_{S \in \binom{[n]}{k}} A_S \prod_{i \in S} Z_i$, with eigenvalues indexed by $y \in \{0, 1\}^n$ given by $\sum_{S \in \binom{[n]}{k}} A_S \prod_{i \in S} (2y_i - 1)$. Equivalently, the eigenvalues are $\sum_{S \in \binom{[n]}{k}} A_S \prod_{i \in S} z_i$ for $z \in \{\pm 1\}^n$.

Suppose for contradiction that there exists a 3-term arithmetic progression, i.e. then by the same argument as in the proof of Proposition 4.4.2, there exists pairwise distinct $v, w, z \in \{\pm 1\}^n$ such that for all $S \in \binom{[n]}{k}$, $\prod_{i \in S} v_i = \prod_{i \in S} w_i = \prod_{i \in S} z_i$. We prove the following claim by induction on $k \geq 1$, and concludes that the above implies either $v = w$ or $v = z$, contradicting that v, w, z are pairwise distinct.

Claim 4.4.10. Let $v, w \in \{\pm 1\}^n$ such that for all $S \in \binom{[n]}{k}$, $\prod_{i \in S} v_i = \prod_{i \in S} w_i$. If k is odd then $v_i = w_i$ for all $i \in [n]$. If k is even then either $v = w$ or $v = -w$.

The base case $k = 1$ is trivially true. Consider $k \geq 2$, and suppose we have proved the claim for all $k' \leq k - 1$. There are two cases, k is odd and k is even.

Suppose k is odd. If $v_i \neq w_i$ for all $i \in [n]$, then since $v_i, w_i \in \{\pm 1\}$, it must be that $v_i = -w_i$ for all $i \in [n]$, but this implies $\prod_{i \in S} v_i = (-1)^k \prod_{i \in S} w_i = -\prod_{i \in S} w_i \neq \prod_{i \in S} w_i$, a contradiction. Hence, there exists i^* such that $v_{i^*} = w_{i^*}$. This means for any $T \subseteq \binom{[n] \setminus \{i^*\}}{k-1}$, by factoring out $v_{i^*} = w_{i^*}$ from $\prod_{i \in T \cup \{i^*\}} v_i = \prod_{i \in T \cup \{i^*\}} w_i$, we have $\prod_{i \in T} v_i = \prod_{i \in T} w_i$. Apply the induction hypothesis for $k - 1$, and note that $k - 1$ is even, we have that either $v_i = w_i$ for all $i \in [n]$ or $v_i = -w_i$ for all $i \in [n] \setminus \{i^*\}$. In the second case, for any $S \subseteq \binom{[n] \setminus \{i^*\}}{k}$, $\prod_{i \in S} v_i = -\prod_{i \in S} w_i \neq \prod_{i \in S} w_i$ (a contradiction). So the first case i.e. $v_i = w_i$ for all $i \in [n]$ holds.

Suppose k is even. Fix an arbitrary index i^* , then there are two cases: either $v_{i^*} = w_{i^*}$ or $v_{i^*} = -w_{i^*}$. If $v_{i^*} = w_{i^*}$ then by the same argument as above, $\prod_{i \in T} v_i = \prod_{i \in T} w_i$ for any $T \in \binom{[n] \setminus \{i^*\}}{k-1}$, which, by the inductive hypothesis, implies that $v_i = w_i$ for all $i \in [n]$. For the second case, $v_{i^*} = -w_{i^*}$. Let $\tilde{v}_i = -v_i$ for all i , then $\tilde{v}_{i^*} = w_{i^*}$ and $\prod_{i \in T} \tilde{v}_i = \prod_{i \in T} w_i$ for any $T \in \binom{[n] \setminus \{i^*\}}{k-1}$, thus $\tilde{v}_i = w_i$ for all $i \in [n]$ by the induction hypothesis. ■

Proof of Proposition 4.4.5. We normalize the Hamiltonian so that $h = 1$. By [KLW'15], the eigenvalues of H are $\lambda_z(J) = \sum_{k=1}^n z_k m_k(J)$ where $m_k(J) = J\mu_k - \sqrt{J^2\mu_k^2 + 1}$ and $\mu_k = \sin\left(\frac{2\pi k}{n}\right)$, for $z \in \{\pm 1\}^n$. Let n be an odd prime.

The set of $J \in \mathbb{R}$ for which H contains a 3-term arithmetic progression is

$$\bigcup_{\substack{v, w, z \in \{\pm 1\}^n \\ \text{distinct}}} \{J \in \mathbb{R} : \lambda_v(J) + \lambda_w(J) - 2\lambda_z(J) = 0\}.$$

Thus, it suffices to show that for each distinct triple $v, w, z \in \{\pm 1\}^n$, the set

$$\{J \in \mathbb{R} : \lambda_v(J) + \lambda_w(J) - 2\lambda_z(J) = 0\}$$

has measure zero. Since $\lambda_v(J) + \lambda_w(J) - 2\lambda_z(J)$ is a real-analytic function of J , its zero set has measure zero unless the function is identically zero.

Suppose that $\lambda_v(J) + \lambda_z(J) - 2\lambda_w(J)$ is identically 0; then for $d_k = \frac{v_k + w_k - z_k}{2} \in \{0, \pm 1, \pm 2\}$, we have, for all J , that

$$\sum_{k=1}^n d_k m_k(J) = 0. \quad (4.4.11)$$

We will show the following claim:

Claim 4.4.12. If (4.4.11) holds for some $d_1, \dots, d_n \in \mathbb{Z}$ then $d_1 = \dots = d_n = 0$.

Given this claim, we conclude that $0 = 2d_k = v_k + w_k - z_k$ for all $k \in [n]$. But, since v, w are distinct, there exists $k \in [n]$ such that $v_k \neq w_k$, in which case $2d_k = v_k + w_k - 2z_k = -2z_k \neq 0$, a contradiction.

Now, we verify the claim.

First, observe that for odd prime n , $\mu_1, \dots, \mu_{(n-1)/2}$ are linearly independent over $\mathbb{Z}$. Suppose for contradiction that there exists $a_1, \dots, a_{(n-1)/2} \in \mathbb{Z}$ such that $0 = \sum_{j=1}^{(n-1)/2} a_j \mu_j$. Since each μ_j can be rewritten as $\mu_j = \frac{\omega^j - \omega^{-j}}{2i}$ where $\omega = e^{2i\pi/n}$, this would imply that $\{\omega_j\}_{j=1}^{n-1}$ is linearly dependent over $\mathbb{Z}$, a contradiction (see [L'79, Lemma 2.11]).

For small J , the Taylor series around $J = 0$ for (4.4.11) gives

$$0 = -2 \sum_k d_k + 2J \sum_k \mu_k d_k - J^2 \sum_k \mu_k^2 d_k + O(J^3).$$

This implies:

$$0 = \sum_{k=1}^n \mu_k d_k$$

Using the facts that

$$\mu_n = \sin(2\pi) = 0 \quad \text{and} \quad \text{for all } k \in [n-1], \mu_k = \sin\left(\frac{2\pi k}{n}\right) = -\sin\left(\frac{2\pi(n-k)}{n}\right) = -\mu_{n-k},$$

we can rewrite

$$0 = \sum_{k=1}^n \mu_k d_k = \sum_{k=1}^{(n-1)/2} \mu_k (d_k - d_{n-k})$$

Since the set $\{\mu_k\}_{k=1}^{(n-1)/2}$ is linearly independent over $\mathbb{Z}$, this implies

$$d_k = d_{n-k} = 0$$

$k \in \{1, \dots, (n-1)/2\}$. Substituting this into (4.4.11), and using $m_n(J) = -1$ and $\mu_k = -\mu_{n-k}$, we get:

$$0 = \sum_{k=1}^n d_k m_k(J)$$

$$\begin{aligned}
 &= -d_n + \sum_{k=1}^{n-1} J\mu_k d_k - \sum_{k=1}^{n-1} d_k \sqrt{J^2 \mu_k^2 + 1} \\
 &= -d_n + J \sum_{k=1}^{(n-1)/2} d_k (\mu_k + \mu_{n-k}) - 2 \sum_{k=1}^{(n-1)/2} d_k \sqrt{J^2 \mu_k^2 + 1} \\
 &= -d_n - 2 \sum_{k=1}^{(n-1)/2} d_k \sqrt{J^2 \mu_k^2 + 1}.
 \end{aligned}$$

This equation holds for all J . Dividing both sides by J yields:

$$0 = -\frac{d_n}{J} - 2 \sum_{k=1}^{(n-1)/2} d_k \sqrt{\mu_k^2 + (1/J)^2}.$$

Taking the limit $J \rightarrow +\infty$, and using the fact that $\mu_k = \sin(\frac{2\pi k}{n}) \geq 0$ for all $k \in \{1, \dots, (n-1)/2\}$, we obtain:

$$0 = -2 \sum_{k=1}^{(n-1)/2} d_k \mu_k.$$

Again, using the linear independence of $\{\mu_k\}_{k=1}^{(n-1)/2}$ over $\mathbb{Z}$, we conclude $d_k = 0$ for all $k \in \{1, \dots, (n-1)/2\}$. Since $d_k = d_{n-k}$, this implies $d_k = 0$. Finally, substituting into (4.4.11) and using $m_n(J) = -1$ yields $d_n = 0$, which finishes the claim.

Now, we prove that the set of J s.t. H contains repeated eigenvalues has measure 0.¹ By the same argument above, we only need to verify that for $v, w \in \{\pm 1\}^n$, if $\lambda_v(J) - \lambda_w(J)$ is identically 0 then $v = w$. Indeed, letting $d'_k = \frac{v_k - w_k}{2} \in \{0, \pm 1\} \subseteq \mathbb{Z}$, if

$$\lambda_v(J) - \lambda_w(J) = 2 \sum_{k=1}^n d'_k m_k(J)$$

is identically 0 then the above claim implies $d'_k = 0$ for all $k \in [n]$ and thus $v = w$. ■

Proof of Proposition 4.4.6. By the same argument as in proof of Theorem 4.0.16, the statement reduces to Proposition 4.4.5. ■

¹This was shown in [KLW¹⁵], but we include a proof for the sake of completeness.

Chapter 5

On the ground state energies of discrete and semiclassical Schrödinger operators

Let $V : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ be a smooth $\mathbb{Z}^d$ -periodic potential and let $\theta \in (0, 1)$ be a parameter. We study the family of discrete Schrödinger operators $A_\theta : \ell_2(\theta\mathbb{Z}^d) \rightarrow \ell_2(\theta\mathbb{Z}^d)$ defined by

$$A_\theta f(k\theta) := \left(2df(k\theta) - \sum_{j \sim k} f(j\theta) \right) + V(k\theta)f(k\theta) = (L + V)f(k\theta), \quad \text{for } k \in \mathbb{Z}^d \quad (5.0.1)$$

where $j \sim k$ indicates that $j, k \in \mathbb{Z}^d$ differ in exactly one coordinate and L is the discrete Laplacian of the corresponding graph on $\theta\mathbb{Z}^d$. This family includes several well-studied operators, such as the almost-Mathieu operator which corresponds to the special case

$$V(x) := \lambda(2 - 2\cos(2\pi x)), \quad \text{for } \lambda > 0 \quad (5.0.2)$$

in $d = 1$. The operator A_θ is bounded and self-adjoint with $\text{spec}(A_\theta) \subseteq [0, 4d + \|V\|_\infty]$. We refer to the infimum of its spectrum as its *ground state energy*, denoted $\mu(A_\theta)$, which admits the well-known variational characterization

$$\mu(A_\theta) = \inf_{\substack{f \in \ell_2(\theta\mathbb{Z}^d) \\ \|f\|=1}} \langle f, A_\theta f \rangle = \inf_{\substack{f \in \ell_2(\theta\mathbb{Z}^d) \\ \|f\|=1}} \left(\sum_{j \sim k \in \mathbb{Z}^d} (f(k\theta) - f(j\theta))^2 + \sum_{k \in \mathbb{Z}^d} V(k\theta)f(k\theta)^2 \right). \quad (5.0.3)$$

Motivated by questions in geometric group theory, there has been interest [BVZ'97; BZ'05; B⁺'17a; B⁺'17b; O'22] in proving good upper and lower bounds on $\mu(A_\theta)$, $\theta \in [0, \frac{1}{2}]$ for the special case (5.0.2) — roughly speaking, this quantity arises naturally

in the representation theory of the Heisenberg group and controls the spectral gap of the random walk on that group. More recently, a line of work initiated by Ozawa [O'16; KNO'19; KKN'21] introduced a method for producing semidefinite programming proofs of Kazhdan's Property (T) for various groups including $SL_n(\mathbb{Z})$, in which effective bounds on certain $\mu(A_\theta)$ play a role. The best known bounds on $\mu(A_\theta)$ for the almost-Mathieu operator due to [BZ'05] rely on delicate trigonometric calculations, and as far as we are aware there is no systematic method for proving such bounds in general.

In contrast, in many cases the ground state energy of a corresponding periodic Schrödinger operator on $\mathbb{R}^d$ can be estimated using well-established tools of spectral theory — see e.g. [K'52; GGS'92] for examples of such bounds in 1 dimension, and Proposition 5.0.10 for coarse bounds in any dimension. In this context, consider the d -dimensional Schrödinger operator with potential V and semiclassical parameter θ , $B_\theta : H^2(\mathbb{R}^d) \rightarrow L^2(\mathbb{R}^d)$, defined as

$$B_\theta f(x) := -\theta^2 \sum_{i=1}^d \frac{d^2}{dx_i^2} f(x) + V(x)f(x) = (-\theta^2 \Delta + V)f(x). \quad (5.0.4)$$

B_θ is an unbounded selfadjoint nonnegative operator. Let $\mu(B_\theta)$ denote the ground state energy of B_θ . We then have the variational characterization (see e.g. [RS'81, Chapter VIII])

$$\mu(B_\theta) = \inf_{\substack{g \in H^2(\mathbb{R}^d) \\ \|g\|=1}} \langle g, B_\theta g \rangle = \inf_{\substack{g \in H^2(\mathbb{R}^d) \\ \|g\|=1}} \int_{\mathbb{R}^d} \theta^2 \|\nabla g(x)\|^2 + V(x)g(x)^2 dx. \quad (5.0.5)$$

The purpose of this chapter is to establish a relationship between $\mu(A_\theta)$ and $\mu(B_\theta)$ for a wide class of V , allowing bounds in one setting to be transferred to the other. To our knowledge, this is the first work to relate the two quantities. Our results are general, in that they apply for a wide class of potentials, and effective, in that they provide concrete bounds for fixed θ as is required in the applications to group theory mentioned above.

We prove the following three results which articulate that while the spectra of A_θ and B_θ are clearly different, their ground state energies are in many cases comparable. With the exception of Theorem 5.0.6, we do not expect the bounds obtained here to be sharp in terms of their dependence on θ .

Upper bound on $\mu(A_\theta)$ for continuous periodic V , irrational θ .

Theorem 5.0.6. *Let $V : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ be continuous and $\mathbb{Z}^d$ -periodic. For θ irrational,*

$$\mu(A_\theta) \leq \mu(B_\theta).$$

The irrationality assumption above cannot be removed in general — see Remark 5.2.11 for a discussion. However the above result can be extended to all θ by instead considering the infimum of the *union spectrum* of A_θ (see Definition 5.2.8) instead of $\mu(A_\theta)$. The proof of Theorem 5.0.6 appears in Section 5.2.

Lower bound on $\mu(A_\theta)$ for smooth periodic V , small θ . Under stronger assumptions on V , we establish the reverse inequality $\mu(A_\theta) \geq \mu(B_\theta)(1 - o(1))$ as $\theta \rightarrow 0$. To state our result we introduce the following definitions.

Definition 5.0.7 (critical point). For a multi-index $\alpha = (\alpha_1, \dots, \alpha_d) \in \mathbb{N}^d$, let $D^\alpha := \left(\frac{\partial}{\partial x_1}\right)^{\alpha_1} \dots \left(\frac{\partial}{\partial x_d}\right)^{\alpha_d}$. A point $m \in \mathbb{R}^d$ is a *critical point of order p* of V if $D^\alpha V(m) = 0$ for all α satisfying $\|\alpha\|_1 \leq p - 1$, and $D^\alpha V(m) \neq 0$ for some α satisfying $\|\alpha\|_1 = p$.

Definition 5.0.8 (coercive). A potential V is *P -coercive* with parameters (t_0, K, P) , $K \geq 1$, if for every $t \leq t_0$ and $y \in \mathbb{R}^d$, $V(y) \leq t + \inf V$ implies that there exists a critical point m of order p with $2 \leq p \leq P$ such that $\|y - m\| \leq Kt^{\frac{1}{p}}$. We say that V is *uniformly p -coercive* if it is p -coercive and all of its critical points have order exactly p .

Coercivity guarantees that if a function is close to its minimum value, then its input must be polynomially close to a critical point. Note that any real analytic periodic potential with finitely many local minima satisfies this property. The bound below depends on the parameters of coercivity as well as V 's smoothness.

Theorem 5.0.9. Let $V : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ be smooth, $\mathbb{Z}^d$ -periodic, and P -coercive with bounded directional derivative in every sequence of P directions — i.e., for each $x \in \mathbb{R}^d$ and $\alpha \in \mathbb{N}^d$ with $\|\alpha\|_1 \leq P$, there is an upper bound on $D^\alpha V(x)$ independent of α and x . Then for θ sufficiently small,

$$\left(1 + O\left(\theta^{\frac{1}{2(P+1)}}\right)\right)\mu(A_\theta) \geq \mu(B_\theta) - O\left(\theta^{2 - \frac{3P+1}{2P(P+1)}}\right).$$

The proof of Theorem 5.0.9 appears in Section 5.3, where a more precise version (Theorem 5.3.1) specifying the constants in the asymptotic notation is provided. The estimate above is nontrivial at least for V with isolated critical points: for such potentials, it can be shown that the additive $O(\theta^{2 - \frac{3P+1}{2P(P+1)}})$ term on the right hand side is of smaller order than $\mu(B_\theta)$ as $\theta \rightarrow 0$. In particular we prove the following estimates in Chapter D using simple arguments which we suspect are known but could not find in the literature.

Proposition 5.0.10. Suppose $V : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ is a smooth $\mathbb{Z}^d$ -periodic potential which is uniformly p -coercive for some $p \geq 2$ and has only isolated critical points. Then as $\theta \rightarrow 0$ we have

$$(a) \quad \mu(B_\theta) = \Omega\left(\theta^{\frac{2p}{p+1}}\right) \text{ if } d = 1, 2,$$

(b) $\mu(B_\theta) = \Omega\left(\theta^{\frac{2p}{p+2}}\right)$ if $d \geq 3$, and

(c) $\mu(B_\theta) = O\left(\theta^{\frac{2p}{p+2}}\right)$ for all $d \geq 1$,

where the implicit constants depend on V .

Lower bound on $\mu(A_\theta)$ for bounded V and any θ . For $y \in \theta\mathbb{Z}^d$, let $C(y) := y + \{0, \theta\}^d$ and $\bar{C}(y) := y + [0, \theta]^d$. We then have the following comparison between $\mu(A_\theta)$ and $\mu(B_\theta)$ under the less restrictive assumptions that V is bounded, nonnegative, and its local averages on the cubes $C(y)$ and $\bar{C}(y)$ are comparable for all $y \in \theta\mathbb{Z}^d$.

Theorem 5.0.11. *Let $V : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, not necessarily periodic with $\|V\|_\infty \leq 1$, and let $\theta \in (0, 1)$. Suppose that there exists $a = a(\theta) > 0$ such that for all $y \in \theta\mathbb{Z}^d$,*

$$\frac{1}{\theta^d} \int_{\bar{C}(y)} V(x) dx \leq a \cdot \frac{1}{2^d} \sum_{k \in C(y)} V(k). \quad (5.0.12)$$

Then,

$$c_d \mu(A_\theta) \geq \mu(B_\theta)$$

where $c_d = \max\left\{\frac{3^d}{2^d} + 5 \cdot 3^d, 4 \cdot \frac{3^d}{2^d} \cdot a\right\}$.

Combined with Theorem 5.0.6 this establishes that for V satisfying the conditions of both theorems and θ irrational, $\mu(A_\theta)$ and $\mu(B_\theta)$ agree up to a multiplicative factor. The quantity c_d arises as the worst-case relationship between local Rayleigh quotients on individual cubes $C(y)$ and $\bar{C}(y)$. The requirement that $a > 0$ in particular rules out the possibility that $V(y) = 0$ for all $y \in \theta\mathbb{Z}^d$ while having $V \neq 0$ elsewhere, in which case we could have $0 = \mu(A_\theta) < \mu(B_\theta)$. The proof of Theorem 5.0.11 appears in Section 5.4.

Example. As an example, consider the almost-Mathieu operator. The potential (5.0.2) is coercive with parameters $\left(4\lambda, \frac{1}{2\pi\sqrt{\lambda}}, 2\right)$, its Lipschitz constant is $4\pi\lambda$, and its second derivative is bounded by $8\pi^2\lambda$. Therefore we find from Theorem 5.3.1 that

$$\mu(B_\theta) \leq \mu(A_\theta) + \frac{1}{1 - \mu(A_\theta)} \left(8\pi\sqrt{\lambda}\theta^{1+q/2} + \left(\mu(A_\theta) + \frac{8\pi\lambda\theta^{1-q}}{1 - 4\pi\lambda\theta^{1-q}} \right) \mu(A_\theta) \right)$$

for all $0 < q < 1$. This choice of V also obeys the conditions of Proposition 5.0.10, so from the choice $q = \frac{5}{6}$ and the constant in (D.1) being $\sqrt[4]{8\pi}\sqrt{\lambda}$, we obtain the multiplicative upper bound

$$\frac{\mu(B_\theta)}{\mu(A_\theta)} \leq 1 + \frac{1}{1 - \sqrt[4]{8\pi}\sqrt{\lambda}\theta^2} \left(8\pi\sqrt{\lambda}\theta^{\frac{1}{12}} + \sqrt[4]{8\pi}\sqrt{\lambda}\theta^2 + \frac{8\pi\lambda\theta^{\frac{1}{6}}}{1 - 4\pi\lambda\theta^{\frac{1}{6}}} \right).$$

The lower bound of 1 follows from Theorem 5.0.6. We note that the large error bound of $O\left(\theta^{\frac{1}{12}}\right)$ (in comparison to [BWZ'26]'s (5.0.13); see the following section for further discussion) is a lossy consequence of applying a Lipschitz bound for the potential on its entire domain.

In the setting of Theorems 5.0.11/5.4.1, when $\theta < \frac{1}{2}$, V has $a \leq \frac{2}{3}$; note that if $\theta = \frac{1}{2}$ such an a does not exist. Thus c_1 as in the Theorem statement is $\frac{33}{2}$ (we note that this can be improved to $c_1 = 9$ with modified choices of parameters in the proof which are preferable in $d = 1$).

Related work

The question of quantitative bounds on $\mu(A_\theta)$ in the case of the almost-Mathieu operator (5.0.2) was first investigated by Béguin, Valette, and Zuk [BVZ'97] in the context of proving mixing time bounds for random walks on the Heisenberg group; they proved $\mu(A_\theta) = \Omega(1)$ for θ in a neighborhood of $\frac{1}{2}$ and conjectured a constant lower bound for $\theta \in [\frac{1}{4}, \frac{1}{2}]$. This was significantly improved by Boca and Zaharescu [BZ'05] who proved nonasymptotic bounds in the full range $\theta \in [0, \frac{1}{2}]$ via an elementary trigonometric calculation (see also Ozawa [O'22] for a simplified proof of a slightly weaker result). Bump, Diaconis, Hicks, Miclo, and Widom [B⁺'17a; B⁺'17b] gave alternative proofs of similar bounds based on an uncertainty principle for circulant matrices and by a comparison to the harmonic oscillator spectrum via a limiting argument.¹ Our results do not improve the best known bound of [BZ'05], but provide an alternate and more conceptual route to proving such bounds.

There has been an extensive study of the fine structure of the spectrum of discrete Schrödinger operators in the mathematical physics community, but surprisingly few works giving quantitative bounds on the edge of the spectrum. Davies [D'99] studied the numerical problem of estimating $\mu(A_\theta)$ and observed that it is related to subtle dynamical phenomena. Akkouché [A'10] gave a necessary and sufficient condition for the ground state energy of a discrete Schrödinger operator $-L + \lambda V$ on $\mathbb{Z}^d$ to be positive for a large (not necessarily periodic or nonnegative) class of potentials V , for small λ : V merely needs to have mean bounded away from 0 on all sufficiently large intervals. This is related to the condition (5.0.12) but holds only on some fixed macroscopic scale, whereas $a(\theta)$ is posited to exist for a chosen scale θ . Milatovic [M'13] generalized this result to other bounded degree graphs.

We will not attempt to survey the vast literature on the spectrum of Schrödinger operators on $\mathbb{R}^d$; most relevant for this chapter are the works by Kato [K'52], Gesztesy, Graf, and Simon [GGS'92], and Arendt and Barry [AB'96; AB'97], who proved bounds on $\mu(B_\theta)$ in $d = 1$ for various classes of potentials. More broadly, there appear to be no results in the literature relating spectral bounds for discrete and continuous

¹This latter argument only gave asymptotic bounds as $\theta \rightarrow 0$.

Schrödinger operators. Establishing such bounds, as is done here, could provide a fruitful bridge between the two areas.

Motivated by this work, Becker, Wittsten, and Zworski [BWZ'26] used semiclassical methods to significantly improve some aspects of Theorem 5.0.9. They establish

$$\frac{\mu(A_\theta)}{\mu(B_\theta)} = 1 + O_d(\theta) \quad \text{as } \theta \rightarrow 0 \quad (5.0.13)$$

for all periodic smooth confining potentials $V : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, with a full asymptotic expansion in θ when V satisfies further nondegeneracy assumptions. Their result is stronger than Theorem 5.0.9 in that it obtains a smaller ratio as a function of θ , but weaker in that it does not provide explicit estimates for fixed θ . The classes of potentials considered in the two theorems are incomparable: [BWZ'26] requires V to be confining, whereas Theorem 5.0.9 requires V to be coercive with critical points of finite order. Theorem 5.0.11 proves a much weaker bound than (5.0.13), but requires fewer assumptions on V , and Theorem 5.0.6 proves a stronger upper bound of 1 for irrational θ .

5.1 Preliminaries

ε -pseudo-ground-states. For $f \in D(T)$, we say that f is a ε -pseudo-ground-state or ε -p.g.s. if $\mathcal{R}_T(f) \leq \mu(T) + \varepsilon$. When T is either A_θ or B_θ and f is an ε -p.g.s. of T , we may assume that f is non-negative by considering $|f|$, which is easily seen to have a smaller Rayleigh quotient than f . We will sometimes want to assume f is positive; in this case we add a small everywhere-positive perturbation from $D(T)$ which increases the Rayleigh quotient by an arbitrarily small amount.

Fourier coefficients on a hypercube. We will often consider functions restricted to $C(y)$ for some choice of $y \in \theta\mathbb{Z}^d$; it will be useful to view these as functions defined on a hypercube. To that end, we recall some facts from Fourier analysis on a hypercube: consider a function F defined on the d -dimensional hypercube $\{-1, 1\}^d$ with the standard edge set. The functions $\chi_S(x) := \prod_{i \in S} x_i$ for $S \subseteq [d]$ form an orthogonal basis for the space of functions on $\{-1, 1\}^d$ with respect to the inner product

$$\langle f, g \rangle := \sum_{x \in \{-1, 1\}^d} f(x)g(x).$$

We can then write $F(x) = \sum_{S \subseteq [d]} \widehat{F}(S) \chi_S(x)$. From [O'14] we have

$$\mathbb{E}[F(x)^2] = \sum_{S \subseteq [d]} \widehat{F}(S)^2, \quad \text{and} \quad \mathbb{E}[F(x) \cdot LF(x)] = 2 \sum_{S \subseteq [d]} |S| \widehat{F}(S)^2$$

$$\text{so } \langle F, LF \rangle = 2^{d+1} \sum_{S \subseteq [d]} |S| \widehat{F}(S)^2 \quad (5.1.1)$$

where the expectation is taken over x drawn uniformly at random from $\{-1, 1\}^d$ and L is the unnormalized graph Laplacian: $Lf(x) := 2df(x) - \sum_{y \sim x} f(y)$.

5.2 Constructing a discrete p.g.s. via sampling

In this section we prove Theorem 5.0.6. Our strategy is: for any $g : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, we exhibit $f : \theta\mathbb{Z}^d \rightarrow \mathbb{R}_{\geq 0}$ such that $\mathcal{R}_{A_\theta}(f) \leq \mathcal{R}_{B_\theta}(g)$. The naive construction is to let f take g 's values on $\theta\mathbb{Z}^d$. The observation is that for θ irrational, ergodicity implies there is *a priori* no reason to prefer $\theta\mathbb{Z}^d \subset \mathbb{R}^d$ over any other shift $\theta\mathbb{Z}^d + \eta \subset \mathbb{R}^d$, and we show by averaging over all such translations that there exists a translation giving rise to a satisfactory f .

Proof of Theorem 5.0.6. Let $g : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$. Let $\eta \in [0, \theta]^d$, drawn uniformly at random with each coordinate independent. We define the test function f_η as $f_\eta(k\theta) := g(k\theta + \eta)$.

Since we sample g at $\theta\mathbb{Z}^d + \eta$ to construct f_η , it is natural to sample V on the same lattice. To this end define $A_{\theta, \eta}$, a discrete operator on functions defined on $\theta\mathbb{Z}^d$, via

$$A_{\theta, \eta} f(k\theta) := 2df(k\theta) - \sum_{\substack{j, k \in \mathbb{Z}^d \\ j \sim k}} f(j\theta) + \sum_{k \in \mathbb{Z}^d} V(k\theta + \eta) f(k\theta).$$

We aim to show that there exists an η so $\mathcal{R}_{A_{\theta, \eta}}(f_\eta) \leq \mathcal{R}_{B_\theta}(g)$. First we have the following lemma, relating $\mu(A_{\theta, \eta})$ and $\mu(A_\theta)$, which is known (appearing in e.g. [J21] for the almost-Mathieu operator), but we include a proof in Chapter D for completeness.

Lemma 5.2.1. *If $V : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ is continuous and periodic and θ is irrational, then $\mu(A_{\theta, \eta})$ is independent of η .*

It is thus sufficient to relate $\mu(A_{\theta, \eta})$ and $\mu(B_\theta)$. We will look at

$$\frac{\mathbb{E}[\langle f_\eta, A_{\theta, \eta} f_\eta \rangle]}{\mathbb{E}[\|f_\eta\|^2]}. \quad (5.2.2)$$

Note that if X and Y are non-negative random variables with $\mathbb{E}[X] < \infty$ and $X \neq 0$ almost surely, then with nonzero probability

$$\frac{Y}{X} \leq \frac{\mathbb{E}[Y]}{\mathbb{E}[X]}. \quad (5.2.3)$$




Figure 5.1: In $d = 2$: $I_{k,1}^\eta$ (purple, dashed), $J_{k',1}^\eta$ (purple, dotted), $I_{k,2}^\eta$ (blue, solid), and $J_{k',2}^\eta$ (blue, dotted) for $k = (0, 1)$ and $\eta = (\frac{1}{3}\theta, \frac{7}{10}\theta)$.




Figure 5.2: In $d = 2$, $J_{k',1}^\eta$ (purple, dotted) and $\bigsqcup_{0 \leq \omega_2 < \theta} J_{k',1}^\omega$ (magenta hatching) for $k' = k_2 = 1$. Note that this union is independent of $\eta = (\frac{1}{3}\theta, \frac{7}{10}\theta)$.

By Section 5.1, we can assume that $f > 0$, so $\|f_\eta\|^2 > 0$ almost surely and it will suffice to bound the ratio of expectations in (5.2.2).

The following segmentation will be useful in our analysis when we want to vary only one coordinate: for $k \in \mathbb{Z}^d$, $1 \leq i \leq d$, $\eta \in [0, \theta)^d$, let

$$I_{k,i}^\eta := \{x : x_j = (k\theta + \eta)_j \ \forall j \neq i, \ k_i\theta + \eta_i \leq x_i < (k_i + 1)\theta + \eta_i\}.$$

Letting $\{e_i\}_{i=1}^d$ be the coordinate basis of $\mathbb{R}^d$, we can also express these sets as $I_{k,i}^\eta = (k\theta + \eta) + [0, \theta)e_i$.

Given k and i , let $k_{-i} := (k_1, \dots, k_{i-1}, k_{i+1}, \dots, k_d) \in \mathbb{Z}^{d-1}$. For $k' \in \mathbb{Z}^{d-1}$ define

$$J_{k',i}^\eta := \bigcup_{\substack{k \in \mathbb{Z}^d \\ k_{-i} = k'}} I_{k,i}^\eta.$$

Immediately from this we have

$$\bigsqcup_{\substack{0 \leq \eta_j < \theta \\ j \neq i}} J_{k',i}^\eta = \bigcup_{\substack{k \in \mathbb{Z}^d \\ k_{-i} = k'}} \overline{C}(\theta k), \quad \text{therefore} \quad \bigsqcup_{k' \in \mathbb{Z}^{d-1}} \bigsqcup_{\substack{0 \leq \eta_j < \theta \\ j \neq i}} J_{k',i}^\eta = \mathbb{R}^d. \quad (5.2.4)$$

Note $J_{k',i}^\eta$ is independent of η_i but is dependent on each η_j , $j \neq i$. See Figure 5.1 for examples of $I_{k,i}^\eta$ and $J_{k',i}^\eta$ and Figure 5.2 for an illustration of (5.2.4).

We evaluate each quantity involved for the main result.

Claim 5.2.5. The norms of f_η and g satisfy $\mathbb{E}_\eta[\|f_\eta\|^2] = \frac{1}{\theta^d}\|g\|^2$.

Proof. This follows from directly evaluating this expectation:

$$\begin{aligned}\mathbb{E}_\eta[\|f_\eta\|^2] &= \mathbb{E}_\eta\left[\sum_{k \in \mathbb{Z}^d} f_\eta(k\theta)^2\right] = \sum_{k \in \mathbb{Z}^d} \mathbb{E}_\eta[g(k\theta + \eta)^2] \\ &= \sum_{k \in \mathbb{Z}^d} \mathbb{E}_{\eta_1}[\dots \mathbb{E}_{\eta_d}[g(k\theta + \eta)^2] \dots] \\ &= \frac{1}{\theta^d}\|g\|^2.\end{aligned}$$

■

Claim 5.2.6. The quadratic forms of f_η and g with the respective Laplacian operators satisfy

$$\mathbb{E}_\eta[\langle f_\eta, L f_\eta \rangle] \leq \frac{1}{\theta^{d-2}}\|\nabla g\|^2.$$

Proof. First we evaluate $\langle f_\eta, L f_\eta \rangle$:

$$\begin{aligned}\langle f_\eta, L f_\eta \rangle &= \sum_{k \in \mathbb{Z}^d} \sum_{i=1}^d (f_\eta((k + e_i)\theta) - f_\eta(k\theta))^2 \\ &= \sum_{i=1}^d \sum_{k \in \mathbb{Z}^d} (g((k + e_i)\theta + \eta) - g(k\theta + \eta))^2 && \text{construction of } f_\eta \\ &= \sum_{i=1}^d \sum_{k \in \mathbb{Z}^d} \left(\int_{I_{k,i}^\eta} \frac{\partial}{\partial x_i} g(x) dx \right)^2 \\ &\leq \sum_{i=1}^d \sum_{k \in \mathbb{Z}^d} \theta \int_{I_{k,i}^\eta} \left(\frac{\partial}{\partial x_i} g(x) \right)^2 dx && \text{Cauchy-Schwarz.}\end{aligned}$$

Then,

$$\mathbb{E}_\eta[\langle f_\eta, L f_\eta \rangle] \leq \theta \sum_{i=1}^d \sum_{k' \in \mathbb{Z}^{d-1}} \mathbb{E}_\eta \left[\int_{J_{k',i}^\eta} \left(\frac{\partial}{\partial x_i} g(x) \right)^2 dx \right]$$

$$\begin{aligned}
 &= \theta \sum_{i=1}^d \mathbb{E}_\eta \left[\sum_{k' \in \mathbb{Z}^{d-1}} \int_{J_{k',i}^\eta} \left(\frac{\partial}{\partial x_i} g(x) \right)^2 dx \right] && \text{Tonelli} \\
 &= \theta \sum_{i=1}^d \mathbb{E}_{(\eta_1, \dots, \eta_{i-1}, \eta_{i+1}, \dots, \eta_d)} \left[\sum_{k' \in \mathbb{Z}^{d-1}} \int_{J_{k',i}^\eta} \left(\frac{\partial}{\partial x_i} g(x) \right)^2 dx \right] && \text{invariance of } J_{k',i}^\eta \text{ w.r.t. } \eta_i \\
 &= \frac{\theta}{\theta^{d-1}} \sum_{i=1}^d \int_{\mathbb{R}^d} \left(\frac{\partial}{\partial x_i} g(x) \right)^2 dx && \text{from (5.2.4)} \\
 &= \frac{1}{\theta^{d-2}} \|\nabla g\|^2.
 \end{aligned}$$

■

Claim 5.2.7. The quadratic forms of f_η and g with the respective potential operators satisfy

$$\mathbb{E}_\eta [\langle f_\eta, V f_\eta \rangle] = \frac{1}{\theta^d} \langle g, V g \rangle.$$

Proof. Again, we directly evaluate this expectation

$$\mathbb{E} [\langle f_\eta, V f_\eta \rangle] = \mathbb{E}_\eta \left[\sum_{k \in \mathbb{Z}^d} V(k\theta + \eta) f_\eta(k\theta)^2 \right] = \sum_{k \in \mathbb{Z}^d} \mathbb{E}_\eta [V(k\theta + \eta) g(k\theta + \eta)^2] = \frac{1}{\theta^d} \langle g, V g \rangle.$$

■

Altogether then we see

$$\frac{\mathbb{E}_\eta [\langle f_\eta, A_{\theta,\eta} f_\eta \rangle]}{\mathbb{E}_\eta [\langle f_\eta, f_\eta \rangle]} \leq \left(\frac{1}{\theta^{d-2}} \|\nabla g\|^2 + \frac{1}{\theta^d} \langle g, V g \rangle \right) \frac{\theta^d}{\|g\|^2} = \frac{\theta^2 \|\nabla g\|^2 + \langle g, V g \rangle}{\|g\|^2}.$$

Minimizing over choices of g , by (5.2.3) there is η such that $\mu(A_{\theta,\eta}) \leq \mu(B_\theta)$ for irrational θ . Finally by Lemma 5.2.1, $\mu(A_\theta) \leq \mu(B_\theta)$. ■

The above proof considers the infimum of the spectrum of the operator A_θ . However, we can instead consider its *union spectrum*, which is more natural in some contexts.

Definition 5.2.8 (union spectrum, cf. [MJ'17]). The *union spectrum* of A_θ is $S_+(A_\theta) := \bigcup_{\eta} \sigma(A_{\theta,\eta})$.

Note that if θ is irrational, then by Lemma 5.2.1, $\inf S_+(A_\theta) = \mu(A_\theta)$. By considering the infimum of the union spectrum, we can extend Theorem 5.0.6 to all θ :

Corollary 5.2.9. *For all $\theta > 0$, $\inf S_+(A_\theta) \leq \mu(B_\theta)$.*

Proof. By the above proof we know that there exists η such that

$$\mu(A_{\theta,\eta}) \leq \frac{\langle f_\eta, A_{\theta,\eta} f_\eta \rangle}{\langle f_\eta, f_\eta \rangle} \leq \frac{\langle g, B_\theta g \rangle}{\langle g, g \rangle}.$$

Minimizing over g and recalling the definition of $S_+(A_\theta)$, the claim is shown. $\blacksquare$

For choices of V such that $\mu(A_\theta)$ is continuous with respect to θ , this result extends immediately to rational θ by continuity.

Corollary 5.2.10. *For V such that $\mu(A_\theta)$ is continuous with respect to θ , $\mu(A_\theta) \leq \mu(B_\theta)$ for all θ .*

Remark 5.2.11. Choices of V satisfying Corollary 5.2.10 are not immediately apparent. A notable counterexample is the almost-Mathieu operator. Let $V_\eta(x) = 2 - 2 \cos(2\pi x + \eta)$ and $V = V_0$. Then for $\theta = 1$ and any η , we have $V_\eta(k\theta) = V(\eta)$ for all k , so $\mu(A_{\theta,\eta}) = V(\eta)$. Now consider $\theta' = 1 - \varepsilon$ where ε is a small positive irrational (so θ' is irrational). Then for $\eta = 0$, we have $V(k\theta') = V(k(1 - \varepsilon)) = V(-k\varepsilon) = V(k\varepsilon)$, so $\mu(A_{1-\varepsilon}) = \mu(A_\varepsilon)$. By Proposition 5.0.10 and Theorem 5.0.6, $\mu(A_\varepsilon) \leq O(\varepsilon)$, so as $\varepsilon \rightarrow 0$ along a sequence of irrationals, $\mu(A_{\theta'}) \rightarrow 0$. Since θ' is irrational, $\mu(A_{\theta',\eta}) \rightarrow 0$ for all η , so we see $\mu(A_{\theta,\eta})$ is not continuous in θ at $\theta = 1$ for any η such that $V(\eta) > 0$ (as $\mu(A_{1,\eta}) = V(\eta)$). By considering e.g. $2 - 2 \cos(2\pi kx)$ for integers $k \geq 1$, one can also create discontinuities at $\theta = \frac{1}{k}$.

5.3 Constructing a continuous p.g.s. via multilinear interpolation

In this section we prove Theorem 5.0.9. Whereas the previous direction of the proof used the probabilistic method to find a pseudo-ground-state for A_θ from one for B_θ , here we will deterministically construct a pseudo-ground-state for B_θ from one for A_θ by multilinear interpolation. In $d = 1$ this amounts just to “connecting” the function values by line segments. In higher dimensions, on each hypercube in $\theta\mathbb{Z}^d$, the interpolating function is the unique multilinear function in each variable which equals the given function on the corners of the hypercube. The techniques brought to bear are first Boolean Fourier-analytic, in terms of understanding both local and global behavior of the functions, and then a careful consideration of the potential’s polynomial behavior near its minimizers.

The precise result that we prove, an explicit form of Theorem 5.0.9, is:

Theorem 5.3.1. *Let $V : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ be $\mathbb{Z}^d$ -periodic, sufficiently smooth, and (t_0, K, P) coercive. Let L be the Lipschitz constant of V and let $\frac{D_v^k D_u V}{k!} \leq M_P$ for any unit vectors u, v , and $k \leq P - 1$ where D_v is the directional derivative with respect to v . For any $q \in (0, 1)$ and θ small enough that $\theta^q \leq t_0$,*

$$\left(1 + \frac{1}{1 - \mu(A_\theta)} \left(\mu(A_\theta) + \frac{2L\theta^{1-q}}{1 - L\theta^{1-q}} \right) \right) \mu(A_\theta) \geq \mu(B_\theta) - \frac{1}{1 - \mu(A_\theta)} \left(2M_P K \sqrt{d} \theta^{1+q(P-1)/P} \right).$$

5.3.1 Defining the test function

First we set up some notation and basic facts. For $y \in \theta\mathbb{Z}^d$, $C(y)$ identifies naturally with the vertices of the hypercube $\{-1, 1\}^d$, and we let $E(C(y))$ be the edge set of the corresponding graph.

Let f be a positive bounded function on $\theta\mathbb{Z}^d$. Let g be the piecewise multilinear extension of f , defined on $\mathbb{R}^d$ as

$$g(x_1, \dots, x_d) := \mathbb{E}_{x_1, \dots, x_d} [f(k_{x_1}\theta, \dots, k_{x_d}\theta)] \quad (5.3.2)$$

where each k_{x_i} is independent with the following distribution for $k_i\theta \leq x_i \leq (k_i + 1)\theta$:

$$\begin{aligned} \mathbb{P}(k_{x_i} = k_i) &= \frac{(k_i + 1)\theta - x_i}{\theta}, \\ \mathbb{P}(k_{x_i} = k_i + 1) &= \frac{x_i - k_i\theta}{\theta} = 1 - \mathbb{P}(k_{x_i} = k_i). \end{aligned}$$

The function g obtained in this way is continuous on each closed hypercube, and is continuous on all of $\mathbb{R}^d$ since g is defined uniquely on the intersection of any two adjacent cubes: if $x_i = k_i\theta$ then $\mathbb{P}(k_{x_i} = k_i - 1) = \mathbb{P}(k_{x_i} = k_i + 1) = 0$ and so the expectation is taken over the same set of lattice points, namely those forming the codimension 1 hypercube of the intersection.

It will be convenient to express g in terms of the Boolean Fourier coefficients of f on each hypercube $C(y)$. To that end, given $y \in \theta\mathbb{Z}^d$, let v_y be a vector such that $x \mapsto \frac{\theta}{2}x - v_y$ sends $C(y) \rightarrow \{-1, 1\}^d$ and $y \mapsto -1$. Define the following functions:

$$F(x) := f\left(\frac{\theta}{2}x\right), \text{ a dilation of } f, \quad \text{and} \quad F_y(x) := f\left(\frac{\theta}{2}x - v_y\right), \text{ a dilation and translation of } f.$$

Note that F is independent of y , and F_y is simply a shifted version of F . We will only ever consider $F_y(x)$ for $x \in \{-1, 1\}^d$, so we restrict it to this domain. As introduced in Section 5.1 we can write F_y as

$$F_y(x) = \sum_{S \subseteq [d]} \widehat{F}_y(S) \chi_S(x).$$

Similarly, we define $G(x) := g(\frac{\theta}{2}x)$ and $G_y := g(\frac{\theta}{2}x - v_y)$. Since there is a unique multilinear function agreeing with F_y on $\{-1, 1\}^d$, we have the expansion

$$G_y(x) = \sum_{S \subseteq [d]} \widehat{F}_y(S) \chi_S(x)$$

for $x \in [-1, 1]^d$.

It will be convenient to relate $\mathbb{E}[G_y(x)^2]$ (where the expectation is taken over x drawn uniformly at random from $[-1, 1]^d$) and $\langle G, -\Delta G \rangle$ to $\{\widehat{F}_y(S)\}$. This simply amounts to integrating the characters with respect to the Lebesgue measure and differentiating. For any χ_S , we have

$$\int_{[-1, 1]^d} (\chi_S(x))^2 dx = \int_{[-1, 1]^d} \prod_{i \in S} x_i^2 dx = 2^{d-|S|} \prod_{i \in S} \int_{-1}^1 x_i^2 dx_i = 2^{d-|S|} \left(\frac{2}{3}\right)^{|S|} = 2^d \left(\frac{1}{3}\right)^{|S|}.$$

We then have for every $y \in \theta\mathbb{Z}^d$:

$$\mathbb{E}[G_y(x)^2] = \sum_{S \subseteq [d]} \left(\frac{1}{3}\right)^{|S|} \widehat{F}_y(S)^2, \quad \text{and} \quad (5.3.3)$$

$$\mathbb{E}\left[\left(\frac{\partial}{\partial x_i} G_y(x)\right)^2\right] = \sum_{S: i \in S} \left(\frac{1}{3}\right)^{|S|-1} \widehat{F}_y(S)^2 \quad \text{so} \quad \langle G_y, -\Delta G_y \rangle = 2^d \sum_{S \subseteq [d]} \left(\frac{1}{3}\right)^{|S|-1} |S| \widehat{F}_y(S)^2. \quad (5.3.4)$$

To prove Theorem 5.3.1, we will analyze the contributions from the Laplacian terms and the potential terms to the Rayleigh quotient separately.

5.3.2 Laplacian term analysis

We want to find expressions for $\|f\|^2$, $\|g\|^2$, $\langle f, Lf \rangle$, and $\langle g, -\Delta g \rangle$ with respect to $\widehat{F}_y$ for $y \in \theta\mathbb{Z}^d$. Since we won't be looking at a single cube in these settings, we define

$$m_k = \sum_{S: |S|=k} \sum_{y \in \mathbb{Z}^d} \widehat{F}_y(S)^2.$$

Proposition 5.3.5. *For f, g as above,*

$$\|f\|^2 = \sum_{k=0}^d m_k, \quad \langle f, Lf \rangle = 4 \sum_{k=0}^d k m_k, \quad \|g\|^2 = \theta^d \sum_{k=0}^d \left(\frac{1}{3}\right)^k m_k, \quad \langle g, -\Delta g \rangle = 4\theta^{d-2} \sum_{k=0}^d \left(\frac{1}{3}\right)^{k-1} k m_k. \quad (5.3.6)$$

Proof. Recall from (5.1.1) that for every $y \in \theta\mathbb{Z}^d$,

$$\mathbb{E}[F_y(x)^2] = \sum_{S \subseteq [d]} \widehat{F}_y(S)^2 \quad \text{and} \quad \langle F_y, LF_y \rangle = 2^{d+1} \sum_{S \subseteq [d]} |S| \widehat{F}_y(S)^2.$$

We relate the norm and Laplacian terms on a single cube to those on the whole space:

$$\|f\|^2 = \frac{1}{2^d} \sum_y \|F_y\|^2 \quad \text{and} \quad \langle f, Lf \rangle = \frac{1}{2^{d-1}} \sum_y \langle F_y, LF_y \rangle.$$

Combining these we have

$$\|f\|^2 = \sum_{k=0}^d m_k \quad \text{and} \quad \langle f, Lf \rangle = 4 \sum_{k=0}^d k m_k,$$

as desired. Recall from (5.3.3) and (5.3.4) that

$$\mathbb{E}[G_y(x)^2] = \sum_{S \subseteq [d]} \left(\frac{1}{3}\right)^{|S|} \widehat{F}_y(S)^2 \quad \text{and} \quad \mathbb{E}\left[\left(\frac{\partial}{\partial x_i} G(x)\right)^2\right] = \sum_{S: i \in S} \left(\frac{1}{3}\right)^{|S|-1} \widehat{F}_y(S)^2.$$

Taking into account the dilation and the change of variables, we see that

$$\|g\|^2 = \left(\frac{\theta}{2}\right)^d \sum_y \|G_y\|^2 \quad \text{and} \quad \langle g, -\Delta g \rangle = \frac{2^2}{\theta^2} \left(\frac{\theta}{2}\right)^d \langle G, -\Delta G \rangle.$$

Combining these we have

$$\|g\|^2 = \theta^d \sum_{k=0}^d \left(\frac{1}{3}\right)^k m_k \quad \text{and} \quad \langle g, -\Delta g \rangle = 4\theta^{d-2} \sum_{k=0}^d \left(\frac{1}{3}\right)^{k-1} k m_k,$$

as desired. ■

Given (5.3.6), we now prove the following.

Lemma 5.3.7. *Let f be an ε -pseudo-ground-state for A_θ and let g be its multilinear extension. Then*

$$\frac{\mathcal{R}_{-\theta^2\Delta}(g)}{\mathcal{R}_L(f)} \leq \frac{1}{1 - \mu(A_\theta) - \varepsilon}.$$

Proof. In the following, keep in mind that m_k is dependent on θ . Assume without loss of generality that $\|f\|^2 = \sum_{k=0}^d m_k = 1$. From (5.3.6) we have

$$\mathcal{R}_{-\theta^2\Delta}(g) = -\theta^2 \frac{\langle g, \Delta g \rangle}{\|g\|^2} = \frac{4 \sum_{k=0}^d \left(\frac{1}{3}\right)^{k-1} k m_k}{\sum_{k=0}^d \left(\frac{1}{3}\right)^k m_k} \quad \text{and} \quad 4 \sum_{k=0}^d \left(\frac{1}{3}\right)^{k-1} k m_k \leq 4 \sum_{k=0}^d k m_k = \langle f, Lf \rangle$$

so that

$$\mathcal{R}_{-\theta^2\Delta}(g) \leq \frac{\langle f, Lf \rangle}{\sum_{k=0}^d \left(\frac{1}{3}\right)^k m_k} = \frac{\mathcal{R}_L(f)}{\sum_{k=0}^d \left(\frac{1}{3}\right)^k m_k}.$$

We also note that if $d = 1$, this becomes an equality.

Now we want to lower bound $\sum_{k=0}^d \left(\frac{1}{3}\right)^k m_k$. It will suffice to just lower bound m_0 . Since $\sum_{k=0}^d m_k = 1$, we can view $\langle f, Lf \rangle$ as a weighted average of elements of $[d]$ (times 4). Since f is an ε -p.g.s., we know $\mathcal{R}_{A_\theta}(f) \leq \mu(A_\theta) + \varepsilon$. From this and (5.3.6) we must have $m_0 \geq 1 - (\mu(A_\theta) + \varepsilon) = 1 - \mu(A_\theta) - \varepsilon$, so

$$\sum_{k=0}^d \left(\frac{1}{3}\right)^k m_k = m_0 + \frac{1}{3}m_1 + \dots \geq 1 - \mu(A_\theta) - \varepsilon.$$

Altogether we have

$$\frac{\mathcal{R}_{-\theta^2\Delta}(g)}{\mathcal{R}_L(f)} \leq \frac{1}{\sum_{k=0}^d \left(\frac{1}{3}\right)^k m_k} \leq \frac{1}{1 - \mu(A_\theta) - \varepsilon}$$

as desired. ■

5.3.3 Potential term analysis

Now all that remains is to analyze the potential terms. We will analyze the Rayleigh quotient of the potential terms restricted to each $C(y)$ (or $\bar{C}(y)$ for the continuous case). We define the following quantities:

$$\begin{aligned} \|g\|_{\bar{C}(y)}^2 &:= \int_{\bar{C}(y)} g(x)^2 dx, & \langle g, Vg \rangle_{\bar{C}(y)} &:= \int_{\bar{C}(y)} V(x)g(x)^2 dx, \\ \|f\|_{C(y)}^2 &:= \sum_{x \in C(y)} f(x)^2, & \langle f, Vf \rangle_{C(y)} &:= \sum_{x \in C(y)} V(x)f(x)^2. \end{aligned}$$

If we do a “cube by cube” analysis, we then need some way to relate this to the whole Rayleigh quotient. Observe that we can relate the “local” Rayleigh quotients to the “global” Rayleigh quotients via

$$\frac{\langle g, Vg \rangle}{\|g\|^2} = \sum_{y \in \theta\mathbb{Z}^d} \frac{\|g\|_{\bar{C}(y)}^2}{\|g\|^2} \frac{\langle g, Vg \rangle_{\bar{C}(y)}}{\|g\|_{\bar{C}(y)}^2} \quad \text{and} \quad \frac{\langle f, Vf \rangle}{\|f\|^2} = \sum_{y \in \theta\mathbb{Z}^d} \frac{\|f\|_{C(y)}^2}{2^d \|f\|^2} \frac{\langle f, Vf \rangle_{C(y)}}{\|f\|_{C(y)}^2} \quad (5.3.8)$$

where the 2^d factor corrects for overcounting, since each vertex x is contained in 2^d cubes. It is natural to define the following probability measures on $\theta\mathbb{Z}^d$:

$$\nu_g(y) := \frac{\|g\|_{\bar{C}(y)}^2}{\|g\|^2} \quad \text{and} \quad \nu_f(y) := \frac{\|f\|_{\bar{C}(y)}^2}{2^d \|f\|^2} = \frac{\mathbb{E}_{k \in \bar{C}(y)} f(k)^2}{\|f\|^2}. \quad (5.3.9)$$

Lemma 5.3.10. *For f and g as above, for all $y \in \theta\mathbb{Z}^d$,*

$$\nu_g(y) \leq \frac{1}{1 - \mu(A_\theta) - \varepsilon} \nu_f(y).$$

Proof. Assume $\|f\|^2 = \sum_{k=0}^d m_k = 1$. Recall from (5.1.1) and (5.3.6) that

$$\|g\|^2 = \theta^d \sum_{k=0}^d \left(\frac{1}{3}\right)^k m_k \quad \text{and} \quad \|g\|_{\bar{C}(y)}^2 = \theta^d \sum_{S \subseteq [d]} \left(\frac{1}{3}\right)^{|S|} \widehat{F}_y(S)^2.$$

Then we have

$$\nu_g(y) = \frac{\sum_{S \subseteq [d]} \left(\frac{1}{3}\right)^{|S|} \widehat{F}_y(S)^2}{\sum_{k=0}^d \left(\frac{1}{3}\right)^k m_k}.$$

From Lemma 5.3.7 we must have $m_0 \geq 1 - \mu(A_\theta) - \varepsilon$. We see $\sum_{k=0}^d \left(\frac{1}{3}\right)^k m_k \geq m_0 \geq 1 - \mu(A_\theta) - \varepsilon$, so

$$\nu_g(y) \leq \frac{1}{1 - \mu(A_\theta) - \varepsilon} \left(\sum_{S \subseteq [d]} \left(\frac{1}{3}\right)^{|S|} \widehat{F}_y(S)^2 \right).$$

From (5.1.1), we can write $\nu_f(y)$ as

$$\nu_f(y) = \sum_{S \subseteq [d]} \widehat{F}_y(S)^2.$$

Clearly $\sum_{S \subseteq [d]} \left(\frac{1}{3}\right)^{|S|} \widehat{F}_y(S)^2 \leq \sum_{S \subseteq [d]} \widehat{F}_y(S)^2$, so the claim is shown. $\blacksquare$

Now we arrive at the main estimate.

Lemma 5.3.11. *Let $V : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ be smooth, $\mathbb{Z}^d$ -periodic and (t_0, K, P) -coercive. Let L be the Lipschitz constant of V and let $\frac{D_v^k D_u V}{k!} \leq M_P$ for any unit vectors u, v , and $k \leq P-1$. For any $q \in (0, 1)$ and θ small enough that $\theta^q \leq t_0$,*

$$|\mathcal{R}_V(g) - \mathcal{R}_V(f)| \leq \frac{1}{1 - \mu(A_\theta) - \varepsilon} \left(2M_P K \sqrt{d} \theta^{1+q(P-1)/P} + \frac{2L\theta^{1-q}}{1 - L\theta^{1-q}} \mathcal{R}_V(f) \right).$$

Proof. Suppose $y \in \theta\mathbb{Z}^d$. Let $q \in (0, 1)$ and θ small enough that $\theta^q \leq t_0$.

Case 1. $V(y) \geq \theta^q$. Then, for $x \in \bar{C}(y)$,

$$1 - L\theta^{1-q} \leq \frac{V(x)}{V(y)} \leq 1 + L\theta^{1-q},$$

so that we have

$$\begin{aligned} \mathcal{R}_{V, \bar{C}(y)}(g) &\leq V(y)(1 + L\theta^{1-q}), \\ \mathcal{R}_{V, C(y)}(f) &\geq V(y)(1 - L\theta^{1-q}), \\ \frac{\mathcal{R}_{V, \bar{C}(y)}(g)}{\mathcal{R}_{V, C(y)}(f)} &\leq \frac{1 + L\theta^{1-q}}{1 - L\theta^{1-q}}, \\ \left| \mathcal{R}_{V, \bar{C}(y)}(g) - \mathcal{R}_{V, C(y)}(f) \right| &\leq \frac{2L\theta^{1-q}}{1 - L\theta^{1-q}} \mathcal{R}_{V, C(y)}(f). \end{aligned}$$

Case 2. $V(y) \leq \theta^q$.

Consider the difference $\left| \mathcal{R}_{V, \bar{C}(y)}(g) - \mathcal{R}_{V, C(y)}(f) \right|$. Being the difference of two random variables supported in $\bar{C}(y)$, it is at most

$$\sup_{x \in \bar{C}(y)} V(x) - \inf_{x \in \bar{C}(y)} V(x) \leq 2 \sup_{x \in \bar{C}(y)} |V(x) - V(y)|.$$

By the fundamental theorem of calculus,

$$\sup_{x \in \bar{C}(y)} |V(x) - V(y)| \leq \sqrt{d}\theta \cdot \sup_{\substack{z \in \bar{C}(y) \\ u \in S^{d-1}}} D_u V(z).$$

Let z and u maximize this expression. By coercivity, m is a critical point of order $p \leq P$. We know $D_{z-m}^k D_u V(m) = 0$ for any $k < p - 1$. By Taylor's remainder theorem we can bound

$$D_u V(z) \leq \sup_{\|\xi - m\|_\infty < \sqrt{d}\theta} D_{z-m}^{p-1} D_u V(\xi) \cdot \frac{\|z - m\|_2^{p-1}}{(p-1)!}.$$

By coercivity, we know that $\|z - m\|_2^{p-1} \leq K\theta^{q(p-1)/p}$. Using the uniform bound M_P we have

$$\max_{x \in \bar{C}(y)} |V(x) - V(y)| \leq M_P K \sqrt{d}\theta \cdot \theta^{q(p-1)/p} = M_P K \sqrt{d}\theta^{1+q(p-1)/p} \leq M_P K \sqrt{d}\theta^{1+q(P-1)/P}. \quad (5.3.12)$$

To relate the global Rayleigh quotients we recall (5.3.8). By Lemma 5.3.10, $v_g \leq \frac{1}{1-\mu(A_\theta)-\varepsilon} v_f$, so by (5.3.12) we have in both cases:

$$\left| \mathcal{R}_V(g) - \mathcal{R}_V(f) \right| \leq \frac{1}{1-\mu(A_\theta)-\varepsilon} \left(2M_P K \sqrt{d}\theta^{1+q(P-1)/P} + \frac{2L\theta^{1-q}}{1-L\theta^{1-q}} \mathcal{R}_V(f) \right). \quad \blacksquare \quad (5.3.13)$$

5.3.4 Proof of Theorem 5.3.1

Proof of Theorem 5.3.1. We first verify that the g constructed in (5.3.2) is a valid test function for (5.0.5). Note that g is continuous and equal to a polynomial of degree at most d on each cube $\bar{C}(y)$. Since $f \in L^\infty(\mathbb{R}^d)$, g has all derivatives $\partial_i g$ uniformly bounded on $\mathbb{R}^d = \bigcup_{y \in \theta\mathbb{Z}^d} \text{int}(\bar{C}(y))$. Thus, g is Lipschitz. By e.g. [E'22, pp. 279], we see

that $g \in W^{1,\infty}(\mathbb{R}^d)$. By Jensen's inequality we further have $\|g\|_{\bar{C}(y)}^2 \leq O(\|f\|_{C(y)}^2)$ for every y , so $g \in H^1(\mathbb{R}^d)$, as desired.

We now aim to bound the difference between

$$\mathcal{R}_{B_\theta}(g) = \mathcal{R}_{-\theta^2\Delta}(g) + \mathcal{R}_V(g) \quad \text{and} \quad \mathcal{R}_{A_\theta}(f) = \mathcal{R}_L(f) + \mathcal{R}_V(f).$$

From Lemma 5.3.7, multiplying through and rearranging yields the following relation on the Laplacian terms:

$$|\mathcal{R}_{-\theta^2\Delta}(g) - \mathcal{R}_L(f)| \leq \frac{\mu(A_\theta) + \varepsilon}{1 - \mu(A_\theta) - \varepsilon} \mathcal{R}_L(f).$$

Combining this with (5.3.13) we have

$$|\mathcal{R}_{B_\theta}(g) - \mathcal{R}_{A_\theta}(f)| \leq \frac{1}{1 - \mu(A_\theta) - \varepsilon} \left(2M_P K \sqrt{d} \theta^{1+q(P-1)/P} + \left(\mu(A_\theta) + \varepsilon + \frac{2L\theta^{1-q}}{1 - L\theta^{1-q}} \right) \mathcal{R}_{A_\theta}(f) \right). \quad (5.3.14)$$

Minimizing over all f we have the following bound on $\mu(B_\theta)$:

$$\mu(B_\theta) \leq \mu(A_\theta) + \frac{1}{1 - \mu(A_\theta)} \left(2M_P K \sqrt{d} \theta^{1+q(P-1)/P} + \left(\mu(A_\theta) + \frac{2L\theta^{1-q}}{1 - L\theta^{1-q}} \right) \mu(A_\theta) \right). \quad \blacksquare \quad (5.3.15)$$

Remark 5.3.16. It is clear from the proof that Theorem 5.0.9 holds for any periodic $V \in C^P$, not necessarily C^∞ .

5.4 A bound for any potential

We give an upper bound on $\mu(B_\theta)$ with respect to $\mu(A_\theta)$ with few constraints on our potential V . The result we prove is the following, a slightly sharper version of Theorem 5.0.11.

Theorem 5.4.1. *Let $V : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, not necessarily periodic with $\|V\|_\infty \leq 1$, and let $\theta \in (0, 1)$. Suppose that there exists $a \in \mathbb{R}$ such that for all $y \in \theta\mathbb{Z}^d$,*

$$\frac{1}{\theta^d} \int_{\bar{C}(y)} V(x)^2 dx \leq a \cdot \frac{1}{2^d} \sum_{k \in C(y)} V(k)^2.$$

Then,

$$c_d \mu(A_\theta) \geq \mu(B_\theta)$$

where $c_d = \max\left\{\frac{3^d}{2^d} + 5 \cdot 3^d, \frac{3^d}{2^d} \frac{a}{r}\right\}$ and r is a dimension-dependent constant between $\frac{1}{4}$ and 1.

Proof. Let f be a positive function on $\theta\mathbb{Z}^d$. We again consider the multilinear extension g as an expectation as defined in (5.3.2). We relate each term in the Rayleigh quotients. The following bound achieves a better constant than Proposition 5.3.5.

Lemma 5.4.2. *The norms of f and g satisfy $\|g\|^2 \geq \left(\frac{2}{3}\right)^d \theta^d \|f\|^2$.*

Proof. We prove this via induction on d . For $d = 1$ we have

$$\begin{aligned} \int_{\mathbb{R}} g(x)^2 dx &= \sum_{k \in \mathbb{Z}} \int_{k\theta}^{(k+1)\theta} \frac{1}{\theta^2} ((x - k\theta)f(k\theta) + ((k+1)\theta - x)f((k+1)\theta))^2 dx \\ &\geq \sum_{k \in \mathbb{Z}} \frac{1}{\theta^2} \int_{k\theta}^{(k+1)\theta} (x - k\theta)^2 f(k\theta)^2 + ((k+1)\theta - x)^2 f((k+1)\theta)^2 dx \\ &= \frac{\theta}{3} \sum_{k \in \mathbb{Z}} (f(k\theta)^2 + f((k+1)\theta)^2) \\ &= \frac{2\theta}{3} \|f\|^2. \end{aligned}$$

Assume that for any function f defined on $\theta\mathbb{Z}^d$, its multilinear extension g satisfies $\|g\|^2 \geq \left(\frac{2}{3}\right)^d \theta^d \|f\|^2$.

Consider f defined on $\theta\mathbb{Z}^{d+1}$ and its multilinear extension g . For any $x_1 \in \mathbb{R}$ with $k\theta \leq x_1 \leq (k+1)\theta$ and for any $y \in \theta\mathbb{Z}^d$, define

$$f_{x_1}(y) := \frac{x_1 - k\theta}{\theta} f(k\theta, y) + \frac{(k+1)\theta - x_1}{\theta} f((k+1)\theta, y).$$

Let g_{x_1} be the multilinear extension of f_{x_1} and observe that by definition $g_{x_1}(y) = g(x_1, y)$. Now we evaluate the norm

$$\begin{aligned} \int_{\mathbb{R}^{d+1}} g(x)^2 dx &= \int_{\mathbb{R}} \int_{\mathbb{R}^d} g(x_1, y)^2 dy dx_1 = \int_{\mathbb{R}} \int_{\mathbb{R}^d} g_{x_1}(y)^2 dy dx_1 \\ &\geq \int_{\mathbb{R}} \left(\frac{2\theta}{3}\right)^d \|f_{x_1}\|^2 dx_1 \\ &= \left(\frac{2\theta}{3}\right)^d \int_{\mathbb{R}} \sum_{y \in \theta\mathbb{Z}^d} f_{x_1}(y)^2 dx_1 \end{aligned}$$

in

$$\begin{aligned}
 &= \left(\frac{2\theta}{3}\right)^d \sum_{y \in \theta\mathbb{Z}^d} \sum_{k \in \mathbb{Z}} \int_{k\theta}^{(k+1)\theta} \frac{1}{\theta^2} ((x_1 - k\theta)f(k\theta, y) + ((k+1)\theta - x_1)f((k+1)\theta, y))^2 dx_1 \\
 &\geq \left(\frac{2\theta}{3}\right)^d \sum_{y \in \theta\mathbb{Z}^d} \sum_{k \in \mathbb{Z}} \frac{1}{\theta^2} \int_{k\theta}^{(k+1)\theta} (x_1 - k\theta)^2 f(k\theta, y)^2 + ((k+1)\theta - x_1)^2 f((k+1)\theta, y)^2 dx_1 \\
 &= \left(\frac{2\theta}{3}\right)^d \cdot \frac{\theta}{3} \sum_{y \in \theta\mathbb{Z}^d} \sum_{k \in \mathbb{Z}} f(k\theta, y)^2 + f((k+1)\theta, y)^2 \\
 &= \left(\frac{2\theta}{3}\right)^{d+1} \|f\|^2,
 \end{aligned}$$

where the fourth line follows by Tonelli's theorem and the fifth by positivity. $\blacksquare$

Lemma 5.4.3. *The quadratic forms of f and g with the respective Laplacian operators satisfy*

$$\langle g, -\Delta g \rangle \leq \theta^{d-2} \langle f, Lf \rangle.$$

Proof. From (5.3.6),

$$\langle f, Lf \rangle = 4 \sum_{k=0}^d km_k \quad \text{and} \quad \langle g, -\Delta g \rangle = 4\theta^{d-2} \sum_{k=0}^d \left(\frac{1}{3}\right)^{k-1} km_k.$$

Immediately we see

$$\langle g, -\Delta g \rangle \leq \theta^{d-2} \langle f, Lf \rangle. \quad \blacksquare$$

Lemma 5.4.4. *Let $V : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, not necessarily periodic with $\|V\|_{\infty} \leq 1$, and let $\theta \in (0, 1)$. Suppose that there exists $a \in \mathbb{R}$ such that for all $y \in \theta\mathbb{Z}^d$,*

$$\frac{1}{\theta^d} \int_{\bar{C}(y)} V(x)^2 dx \leq a \cdot \frac{1}{2^d} \sum_{k \in C(y)} V(k)^2.$$

The quadratic forms of f and g with the respective potential operators satisfy

$$\langle g, Vg \rangle \leq \theta^d \left(10 \cdot 2^{d-1} \langle f, Lf \rangle + \frac{a}{r} \langle f, Vf \rangle \right)$$

where r is a d -dependent constant between $\frac{1}{4}$ and 1.

Proof. We will consider cases based on the size of the local Laplacian Rayleigh quotients. Let $\alpha(d)$ be a function of d to be determined later. Let $y \in \theta\mathbb{Z}^d$ and let $\langle f, Lf \rangle_{C(y)} := \sum_{(u,v) \in E(C(y))} (f(u) - f(v))^2$.

Case 1. $\frac{\langle f, Lf \rangle_{C(y)}}{\|f\|_{C(y)}^2} \geq \alpha(d)$.

We use the naive bound

$$\frac{\langle g, Vg \rangle_{\bar{C}(y)}}{\|g\|_{\bar{C}(y)}^2} \leq 1 \quad \text{so} \quad \frac{\langle g, Vg \rangle_{\bar{C}(y)}}{\|g\|_{\bar{C}(y)}^2} \leq \frac{1}{\alpha(d)} \frac{\langle f, Lf \rangle_{C(y)}}{\|f\|_{C(y)}^2}.$$

Note that $\|g\|_{\bar{C}(y)}^2 = \theta^d \sum_{S \subset [d]} \left(\frac{1}{3}\right)^{|S|} \widehat{F}_y(S)^2$ and $\|f\|_{C(y)}^2 = 2^d \sum_{S \subset [d]} \widehat{F}_y(S)^2$, so we immediately see

$$\frac{\|g\|_{\bar{C}(y)}^2}{\|f\|_{C(y)}^2} \leq \frac{\theta^d}{2^d}.$$

Using this and rearranging gives that

$$\langle g, Vg \rangle_{\bar{C}(y)} \leq \frac{\|g\|_{\bar{C}(y)}^2}{\|f\|_{C(y)}^2} \frac{1}{\alpha(d)} \langle f, Lf \rangle_{C(y)} \leq \frac{\theta^d}{2^d} \frac{1}{\alpha(d)} \langle f, Lf \rangle_{C(y)}. \quad (5.4.5)$$

Case 2. $\frac{\langle f, Lf \rangle_{C(y)}}{\|f\|_{C(y)}^2} \leq \alpha(d)$.

Assume $\|f\|_{C(y)}^2 = 1$. Let $\mu := \frac{\langle f, \mathbb{K} \rangle}{2^d}$. By the Poincaré inequality,

$$\|f - \mu \mathbb{K}\|_{C(y)}^2 \leq \langle f, Lf \rangle \leq \alpha(d).$$

Let $\|f - \mu \mathbb{K}\|_{C(y), \infty}^2$ be the infinity norm of $f - \mu \mathbb{K}$ on $C(y)$. Then

$$\|f - \mu \mathbb{K}\|_{C(y), \infty}^2 \leq \|f - \mu \mathbb{K}\|_{C(y)}^2 \leq \alpha(d).$$

By its definition, $\mu = \sqrt{\mathbb{E}[f^2] - \text{Var}(f)} \geq \sqrt{\frac{1}{2^d} - \frac{\alpha(d)}{2^d}}$. We see that for all $x \in C(y)$,

$$\sqrt{\frac{1}{2^d}(1 - \alpha(d))} - \sqrt{\alpha(d)} \leq f(x) \leq \sqrt{\frac{1}{2^d}} + \sqrt{\alpha(d)}.$$

Now we relate the minimum and maximum values of f on $C(y)$. Define r by

$$\sqrt{r} = \frac{\sqrt{1 - \alpha(d)} - \sqrt{2^d \alpha(d)}}{1 + \sqrt{2^d \alpha(d)}}.$$

Then we have $\min_{k \in C(y)} f(k)^2 \geq r \max_{k \in C(y)} f(k)^2$.

Combining the above,

$$\langle f, Vf \rangle_{C(y)} \geq \min_{k \in C(y)} f(k)^2 \sum_{k \in C(y)} V(k) \geq \frac{2^d}{\theta^d} \frac{r}{a} \max_{k \in C(y)} f(k)^2 \int_{\bar{C}(y)} V(x) dx \geq \frac{2^d}{\theta^d} \frac{r}{a} \langle g, Vg \rangle_{\bar{C}(y)}. \quad (5.4.6)$$

From (5.4.5) and (5.4.6), for any $y \in \theta\mathbb{Z}^d$

$$\langle g, Vg \rangle_{\bar{C}(y)} \leq \frac{\theta^d}{2^d} \frac{a}{r} \langle f, Vf \rangle_{C(y)} + \frac{\theta^d}{2^d} \frac{1}{\alpha(d)} \langle f, Lf \rangle_{C(y)}.$$

Now we sum over all $y \in \theta\mathbb{Z}^d$. Observe that in summing $\langle f, Vf \rangle_{C(y)}$ over y , each vertex is counted 2^d times; similarly, in summing $\langle f, Lf \rangle_{C(y)}$, each edge is counted 2^{d-1} times. Thus,

$$\begin{aligned} \langle g, Vg \rangle &= \sum_{y \in \theta\mathbb{Z}^d} \langle g, Vg \rangle_{\bar{C}(y)} \leq \frac{\theta^d}{2^d} \frac{a}{r} \sum_{y \in \theta\mathbb{Z}^d} \langle f, Vf \rangle_{C(y)} + \frac{\theta^d}{2^d} \frac{1}{\alpha(d)} \sum_{y \in \theta\mathbb{Z}^d} \langle f, Lf \rangle_{C(y)} \\ &= \frac{\theta^d}{2^d} \frac{a}{r} 2^d \langle f, Vf \rangle + \frac{\theta^d}{2^d} \frac{1}{\alpha(d)} 2^{d-1} \langle f, Lf \rangle \\ &= \theta^d \left(\frac{a}{r} \langle f, Vf \rangle + \frac{1}{2\alpha(d)} \langle f, Lf \rangle \right). \end{aligned}$$

Letting $\alpha(d) = \frac{1}{10 \cdot 2^d}$ we have $r \geq \frac{1}{4}$ and we achieve the desired result. ■

From Lemma 5.4.2, Lemma 5.4.3, and Lemma 5.4.4 we have

$$\begin{aligned} \frac{\langle g, B_\theta g \rangle}{\|g\|_2^2} &\leq \frac{3^d}{2^d \theta^d \|f\|^2} \left(\theta^d (1 + 10 \cdot 2^{d-1}) \langle f, Lf \rangle + \frac{a}{r} \theta^d \langle f, Vf \rangle \right) \\ &= \frac{1}{\|f\|^2} \left(\left(\frac{3^d}{2^d} + 5 \cdot 3^d \right) \langle f, Lf \rangle + \frac{3^d}{2^d} \frac{a}{r} \langle f, Vf \rangle \right). \end{aligned}$$

This proves our claim with

$$c_d = \max \left\{ \frac{3^d}{2^d} + 5 \cdot 3^d, \frac{3^d}{2^d} \frac{a}{r} \right\}. \quad \blacksquare$$

Appendix A

Appendix to Chapter 2 (elements in the number ring)

Here, we study irreducible elements and sums of two squares in $\mathbb{Z}[\varphi]$. As in [RS'16, §C], here we shall summarize results about classifying irreducible elements in $\mathbb{Z}[\varphi]$ and an efficient algorithm for writing certain elements of $\mathbb{Z}[\varphi]$ as a sum of two squares. For this section, let $N = N_{\mathbb{Q}(\varphi)/\mathbb{Q}}$. We readily compute $N(a + b\varphi) = a^2 + ab - b^2$. Recall that $\mathbb{Z}[\varphi]$ is a Euclidean domain with respect to $|N|$, and that $\mathbb{Z}[\varphi]^\times = \langle \pm\varphi \rangle$ (cf. [PS'18, §4.1.4]). Also recall that $\varphi^\bullet$ is φ 's Galois conjugate, which happens to equal $1 - \varphi$. Extend $\bullet$ $\mathbb{Q}$ -linearly to all of $\mathbb{Q}(\varphi)$.

Proposition A.1. *Let p be irreducible in $\mathbb{Z}$. If $p \equiv \pm 2 \pmod{5}$ then p is irreducible in $\mathbb{Z}[\varphi]$. If $p \equiv \pm 1 \pmod{5}$ then there is an algorithm (running in time $O(\text{poly log } p)$) to compute a $\mathbb{Z}[\varphi]$ -irreducible element dividing p .*

Proof. If p is reducible in $\mathbb{Z}[\varphi]$ then there exist non-units $u, v \in \mathbb{Z}[\varphi]$ with $uv = p$ and accordingly $N(u) \mid N(p) = p^2$; that is, $N(u) = p$.

- Observe that $a^2 + ab - b^2 \equiv (a - 2b)^2 \pmod{5}$ and that neither of ± 2 are quadratic residues modulo 5. This proves the first part of the proposition, as if $u = a + b\varphi$ then $N(u)$ can never equal p .
- To prove the second part, we first show that if $p \equiv 1 \pmod{5}$ then there exists $x \in \mathbb{Z}$ satisfying $x^2 - x - 1 \equiv 0 \pmod{p}$. Indeed, rewrite the equation to $(x - 2^{-1})^2 \equiv 1 + 4^{-1}$. Now,

$$\left(\frac{1 + 4^{-1}}{p}\right) = \left(\frac{1 + 4^{-1}}{p}\right)\left(\frac{4}{p}\right) = \left(\frac{5}{p}\right) = \left(\frac{p}{5}\right)(-1)^{p-1} = \left(\frac{p}{5}\right) = 1$$

with the third equality following by quadratic reciprocity. Let $y \in \mathbb{Z}$ be any modulo- p square root of $1 + 4^{-1}$. Then $x = y + 2^{-1}$ suffices. Set $z = x - \varphi \in \mathbb{Z}[\varphi]$. Then $p \mid x^2 - x - 1 = zz^\bullet$ but p does not divide z as z 's φ -part is not divisible by

p , so p is not irreducible. Thus $u = \gcd(p, z)$ will be a non-unit divisor of p , as desired. (u cannot be a unit because otherwise $\gcd(p^\bullet, z^\bullet) = \gcd(p, z^\bullet)$ will also be a unit, but then there are no prime divisors of $zz^\bullet$ also dividing p , a contradiction.)

The timing of the above method is due to the extended Euclidean algorithm for modular inverses, and the standard Euclidean algorithm for GCDs in a Euclidean domain, which have respective runtimes $O(\log^2 p)$ and $O(\log p)$.

The only case not covered in the above is $p = 5$, which trivially factors as $(-1 + 2\varphi)^2$.

Algorithm A.2. Assume that there is an efficient blackbox algorithm to factor $n \in \mathbb{Z}$ in time $O(\text{poly } \log n)$. Then there is an algorithm which factors $n \in \mathbb{Z}[\varphi]$ in time $O(\text{poly } \log |N(n)|)$.

Proof. Factor $N(n)$ as $\prod_{\ell=1}^m p_\ell^{e_\ell}$. Each prime p_ℓ is factorizable in time $O(\text{poly } \log p_\ell) \subset O(\text{poly } \log N(n))$, by Proposition A.1. There are $O(\log N(n))$ such primes.

To each irreducible $u \in \mathbb{Z}[\varphi]$, let its *associated prime* be the least (positive) prime factor of $N(u)$; that is, $\sqrt{N(u)}$ when u is a $\mathbb{Z}$ -irreducible times a unit, and $N(u)$ otherwise.¹

Lemma A.3. For any $u \in \mathbb{Z}[\varphi] \setminus \{0\}$, it is not the case that both u and $u\varphi$ can be written as a sum of two squares.

Proof. Suppose not, so write $u = w^2 + x^2$ and $u\varphi = y^2 + z^2$. Then we may take the product

$$u^\bullet u\varphi = N(u)\varphi = (w^\bullet y + x^\bullet z)^2 + (w^\bullet z - x^\bullet y)^2 = (a + b\varphi)^2 + (c + d\varphi)^2$$

for some $a, b, c, d \in \mathbb{Z}$. The right-hand side expands out with 1-part² equal to $a^2 + b^2 + c^2 + d^2$, but $N(u)\varphi$ has 1-part equal to 0, so $a = b = c = d = 0$, and therefore $N(u) = 0$, a contradiction.

Proposition A.4. Let $u \in \mathbb{Z}[\varphi]$ be irreducible, with associated prime p . If $p \equiv 1, 3, 7, 9, 13, 17 \pmod{20}$ then there is an efficient algorithm in time $O(\text{poly } \log p)$ to write either u or $u\varphi$ as a sum of two squares.

Observe that $\phi(20) = 8$ (ϕ here signifying Euler's totient function), so by Chebotarev's density theorem the above-asserted algorithm manages to write $3/4$ of primes as sums of two squares.

¹Let us quickly establish why the associated prime is well-defined. Suppose $p \neq q$ are $\mathbb{Z}$ -primes with $p, q \mid N(u) = uu^\bullet$. Let $p', q' \in \mathbb{Z}[\varphi]$ be irreducibles such that $|N(p')| = p$ if $p \equiv \pm 1 \pmod{5}$ or $p = 5$, and $p' = p$ otherwise; and define q' similarly. Then $p', q' \mid uu^\bullet$, so as irreducibles we have p' dividing either u or $u^\bullet$, and similarly for q' , hence $p = q$ since u and $u^\bullet$ are themselves irreducible. Thus, $N(u)$ is divisible by at most one prime at most twice.

²Viewing $\mathbb{Q}(\varphi) \cong \langle 1, \varphi \rangle_{\mathbb{Q}} \cong \mathbb{Q}^2$ so that the “1-part” and “ φ -part” are the corresponding components in the canonical isomorphism.

Proof. Before beginning in earnest, first remark that 2 factors in $\mathbb{Z}[i, \varphi]$ as $(1+i)(1-i)$, and that $1 \pm i$ are irreducible because $N_{\mathbb{Q}(i, \varphi)/\mathbb{Q}(\varphi)}(1 \pm i) = 2$, a $\mathbb{Z}[\varphi]$ -irreducible.

- Suppose $p \equiv 1 \pmod{4}$ (i.e. $p \equiv 1, 9, 13, 17 \pmod{20}$). Then there exists even $x \in \mathbb{Z}$ satisfying $x^2 + 1 \equiv 0 \pmod{p}$, that is, $u \mid p \mid x^2 + 1$. Let $g = \gcd(u, x + i)$, computed in $\mathbb{Z}[i, \varphi]$, and set $s = \operatorname{re}(g)$ and $t = \operatorname{im}(g)$. We claim that either u or $u\varphi$ equals a squared unit times $s^2 + t^2$. Indeed, suppose $q \in \mathbb{Z}[i, \varphi]$ is an irreducible divisor of u . We wish to show that q divides exactly one of $x \pm i$. (Clearly it divides at least one of them, by $q \mid u \mid x^2 + 1$.)

- Say $1 \pm i \neq q \in \mathbb{Z}[i, \varphi]$ is an irreducible dividing u and both $x \pm i$. Then $q \mid 2i$. By our choice of q this is impossible.
- Suppose now we let q be one of $1 \pm i$. Then in order to have $1 \pm_1 i \mid x \pm_2 i$ (where $\pm_1$ and $\pm_2$ are independent signs) we must have $(1 \pm_1 i)(a + bi) = x \pm_2 i$ where $a, b \in \mathbb{Z}[\varphi]$. Solving for a and b by looking at real and imaginary parts, we have

$$\begin{aligned} a \mp_1 b &= x \\ b \pm_1 a &= \pm_2 1 \end{aligned}$$

and solving yields $2a = x \pm_1 \pm_2 1$, so in order for a to exist we must have x is odd. This is a contradiction to our assumption that x is even, regardless of the choice of $\pm_1$ and $\pm_2$, thus in fact we have neither of $1 \pm i$ dividing either of $x \pm i$.

Thus, if q divides u with multiplicity m_q (so that u factors in $\mathbb{Z}[i, \varphi]$ as $u = \varphi^m \prod_{q \mid u} q^{m_q}$), we must have that q divides $x \pm i$ with multiplicity at least m and

$x \mp i$ with multiplicity 0; and so¹ $g = \prod_{q \mid u, x+i} q^{m_q}$ while $\bar{g} = \gcd(u, x - i) = \prod_{q \mid u, x-i} q^{m_q}$.

Thus $u = \varphi^m g \bar{g} = \varphi^m (\operatorname{re}(g)^2 + \operatorname{im}(g)^2)$. If $2 \mid m$ then $u = (\varphi^{m/2} \operatorname{re}(g))^2 + (\varphi^{m/2} \operatorname{im}(g))^2$.

- Suppose $p \equiv 3 \pmod{4}, \pm 2 \pmod{5}$ (i.e. $p \equiv 3, 7 \pmod{20}$). Then compute

$$\left(\frac{-5}{p}\right) = \left(\frac{-1}{p}\right) \left(\frac{5}{p}\right) = -\left(\frac{p}{5}\right) (-1)^{p-1} = 1$$

so that there exists even $x \in \mathbb{Z}$ which is not a multiple of 5 satisfying $x^2 + 5 \equiv 0 \pmod{p}$, that is, $u = p \mid x^2 + 5$. Let $g = \gcd(g, x + i(-1 + 2\varphi))$, computed in $\mathbb{Z}[i, \varphi]$, and set $s = \operatorname{re}(g)$ and $t = \operatorname{im}(g)$. We claim that $p = s^2 + t^2$. Indeed, suppose $q \in \mathbb{Z}[i, \varphi]$ is an irreducible divisor of p (in $\mathbb{Z}[i, \varphi]$). We wish to show that q divides exactly one of $x \pm i(-1 + 2\varphi)$. (Clearly it divides at least one of them, by $q \mid p \mid x^2 + 5$.)

¹Depending on choice of GCD algorithm, there may be a unit factor in front, but we assume not without loss of generality.

- Say $-1+2\varphi, 1\pm i \neq q \in \mathbb{Z}[i, \varphi]$ is an irreducible dividing u and both $x \pm i(-1+2\varphi)$. Then $q \mid 2(-1+2\varphi)i$. By our choice of q this is impossible.
- Suppose we let $q = -1+2\varphi = \sqrt{5}$. Then as 5 does not divide x , we have

$$\frac{x \pm i(-1+2\varphi)}{q} = \frac{x}{5}(-1+2\varphi) \pm i \in \mathbb{Q}(i, \varphi) \setminus \mathbb{Z}[i, \varphi]$$

and so in fact q divides neither of $x \pm i(-1+2\varphi)$.

- Suppose now we let q be one of $1\pm i$. Then in order to have $1\pm_1 i \mid x\pm_2 i(-1+2\varphi)$ (where $\pm_1$ and $\pm_2$ are independent signs) we must have $(1\pm_1 i)(a+bi) = x\pm_2 i(-1+2\varphi)$ where $a, b \in \mathbb{Z}[\varphi]$. Solving for a and b by looking at real and imaginary parts, we have

$$\begin{aligned} a \mp_1 b &= x \\ b \pm_1 a &= \pm_2(-1+2\varphi) \end{aligned}$$

and solving yields $2a = x\pm_1\pm_2(-1+2\varphi)$, so in order for a to exist we must have x is odd. This is a contradiction to our assumption that x is even, regardless of the choice of $\pm_1$ and $\pm_2$, thus in fact we have neither of $1\pm i$ dividing either of $x \pm i$.

It is here that we require that $\mathbb{Z}[i, \varphi]$ is norm-Euclidean, as this ensures that GCDs are efficiently computable, using an Euclidean algorithm with respect to the Euclidean function (the norm). This is a result of Lemmermeyer [L'11, Theorem 2]; the journal version [BS'23] contains a less general proof of the same fact, and I am grateful to Thomas Browning for alerting me to this reference. As before, the bottlenecks are in the extended and standard Euclidean algorithms, which are both of runtime $O(\log^2 p)$.

Observe that we can also include the $\mathbb{Z}$ -irreducibles $2 = 1^2+1^2$ and $5 = (-1+2\varphi)^2+0^2$.

Theorem A.5. *For $x \in \mathbb{Z}[\varphi]$, factor it as*

$$x = \prod_{\ell=1}^n u_{\ell}^{m_{\ell}}$$

where $u_{\ell} \in \mathbb{Z}[\varphi]$ are all irreducible, with associated primes p_{ℓ} , and $m_{\ell} \in \mathbb{N}$. Then either x or $x\varphi$ may be written as the sum of two squares if, for all $\ell \in [n]$, one of the following holds:

- $p_{\ell} = 2$;
- $p_{\ell} \equiv 1, 3, 7, 9, 13, 17 \pmod{20}$;
- $2 \mid m_{\ell}$;

and, for arbitrary x not a priori satisfying all three criteria, this whole procedure (i.e., either determining a sum of two squares, or rejecting on the basis of some ℓ failing all three conditions) runs in time $O(\text{poly } \log |N(x)|)$.

Proof. For each $\ell \in [n]$, consider s_ℓ and t_ℓ defined as follows:

- If $p_\ell = 2$ then $s_\ell = t_\ell = 1$.
- If $p_\ell \equiv 1, 3, 7, 9, 13, 17 \pmod{20}$ then s_ℓ and t_ℓ are as computed using Proposition A.4.
- If $2 \mid m_\ell$ then $s_\ell = u_\ell^{m_\ell/2}$ and $t_\ell = 0$.

Now, we do the standard (inductive) trick where $(s_\ell^2 + t_\ell^2)(s_{\ell'}^2 + t_{\ell'}^2) = (s_\ell s_{\ell'} + t_\ell t_{\ell'})^2 + (s_\ell t_{\ell'} - t_\ell s_{\ell'})^2$. Since we have that $s_\ell^2 + t_\ell^2$ equals either u_ℓ or $u_\ell \varphi$, the resulting product from performing the trick n times gives a sum of two squares $s^2 + t^2$ equal to $x\varphi^{n'}$ where $0 \leq n' \leq n$. If $2 \mid n'$ then we return the pair $(s\varphi^{-n'/2}, t\varphi^{-n'/2})$, and otherwise we return $(s\varphi^{(1-n')/2}, t\varphi^{(1-n')/2})$.

For the runtime analysis, simply factor x using Algorithm A.2 and then for each resulting factor, apply Proposition A.4 after checking that one of the three criteria holds (immediately halting and rejecting if any factor fails).

Appendix B

Appendix to Chapter 3

B.1 Selecting polynomial approximations

Consider the task of computing a low-degree uniform-precision- ε polynomial approximation p to a given continuous function $f : S \rightarrow \mathbb{C}$ on the square $S := S^2 = [-1, 1]^2$, a special case of the d -hypercube $S^d := [-1, 1]^d$.

First, some definitions. Let $\mathcal{P}_n$ be the complex polynomials of degree at most n . Let $\text{dist}_{[a,b]}(f, \mathcal{P}_n) := \inf_{p \in \mathcal{P}_n} \|f - p\|_{L^\infty[a,b]}$, for f any function taking values on $[a, b]$. We say that the n th Chebyshev points on $[a, b]$ are $T_n = (t_0^{(n)}, \dots, t_n^{(n)}) \in [a, b]^{n+1}$ which attain $\min_{t_\ell \in [a,b]} \max_{x \in [a,b]} \prod_{0 \leq \ell < n} |x - t_\ell|$. For $0 \leq k < n$ and $T = (t_1, \dots, t_n)$, define the interpolating polynomials

$$L_n^{(k)}(x, T) := \prod_{\substack{0 \leq \ell < n \\ \ell \neq k}} \frac{x - t_\ell}{t_k - t_\ell}.$$

Let $\omega(f, E, \delta) := \sup_{\substack{x, y \in E \\ |x, y| < \delta}} |f(x) - f(y)|$, for any compact set $E \subset \mathbb{R}^d$ and any continuous $f : E \rightarrow \mathbb{R}$.

Consider the following result.

Theorem B.1.1 (Pleśniak [P'94, Lemma 2.2]). *Consider the domain $P := \prod_{\ell=1}^d [a_\ell, b_\ell] \subset \mathbb{R}^d$ and any continuous function $f : P \rightarrow \mathbb{R}$. Select positive integers n_ℓ for $1 \leq \ell \leq d$ and*

compute T_{n_ℓ} .¹ Define

$$p(x) := \sum_{k \in \prod_{1 \leq \ell \leq d} [[n_\ell]]} f(t_{k_1}^{(n_1)}, \dots, t_{k_d}^{(n_d)}) \prod_{1 \leq \ell \leq d} L_{n_\ell}^{k_\ell}(x_\ell, T_{n_\ell}). \quad (\text{B.1.2})$$

Then, we have that²

$$\sup_{x \in P} |f(x) - p(x)| \leq \sup_{x \in P} \sum_{1 \leq \ell \leq d} \sum_{j \in \prod_{1 \leq k < \ell} [[n_k]]} (\log n_\ell + 2) \text{dist}_{[a_\ell, b_\ell]} \left(f(t_{j_1}^{(n_1)}, \dots, t_{j_{\ell-1}}^{(n_{\ell-1})}, \cdot, x_{\ell+1}, \dots, x_d), \mathcal{P}_{n_\ell} \right).$$

Notice that $\deg p = \sum_{1 \leq \ell \leq d} n_\ell$. We are interested in the case of $d = 2$ and $a_1 = a_2 = -1$, $b_1 = b_2 = 1$, and $\deg p = n_1 + n_2$. It just so happens that:

Proposition B.1.3 (classical). *The n th Chebyshev points on $[-1, 1]$ are $t_\ell^{(n)} = \cos(\frac{2\pi\ell}{n})$.*

Thus, our task is greatly simplified, as the bound from Theorem B.1.1 reduces to

$$\sum_{0 \leq j < n_1} (\log n_2 + 2) \text{dist}_{[-1, 1]} \left(f(t_j^{(n_1)}, \cdot), \mathcal{P}_{n_2} \right)$$

which is in turn boundable via:

Theorem B.1.4 (Jackson's inequality, cf. [R'69, Theorem 1.4]). *Let $f \in C[-1, 1]$ be continuous. Then,*

$$\text{dist}_{[-1, 1]}(f, \mathcal{P}_n) \leq 6\omega\left(f, [-1, 1], \frac{1}{n}\right).$$

and its explicit, more general version

Theorem B.1.5 (Jackson's inequality, cf. [P'94, Theorem 2.1]). *Let $f \in C^k[a, b]$ and $0 \leq k < n$. Then,*

$$\text{dist}_{[a, b]}(f, \mathcal{P}_n) \leq \frac{\left(\frac{\pi}{4}(b-a)\right)^k}{k! \binom{n+1}{k}} \|f^{(k)}\|_{L^\infty[a, b]}.$$

Notice that the denominator grows as n^k for $n \gg k$, so if $\|f^{(k)}\|_{L^\infty[a, b]} \ll n^k$ then the bound is interesting.

Therefore we can apply the bound to $\text{re } f$ and $\text{im } f$, supposing that each has the appropriate differentiability along each “slice.” Define the quantity

$$]M_{r, k}(f) := \sup_{t \in [-1, 1]} \|\partial_2^k(\text{re } f)(t, \cdot)\|_{L^\infty[-1, 1]}$$

¹The published statement of this theorem elects to work with Fekete (Legendre) points, but Chebyshev points feature an exponentially better Lebesgue constants and so are advantageous in this setting. The proof is unchanged.

²The published statement of this theorem erroneously omits the supremum on the right-hand side.

and similarly $M_{i,k}(f)$ for the imaginary part. Let $M_k(f) := \max\{M_{r,k}(f), M_{i,k}(f)\}$.

We thus have in full:

Theorem B.1.6. *For all $n_1, n_2 \in \mathbb{N}$, let $f \in C[-1, 1]^2$ satisfy $\operatorname{re} f(t, \cdot), \operatorname{im} f(t, \cdot) \in C^k[-1, 1]$ for all $t \in [-1, 1]$ and some $0 \leq k \leq n_1, n_2$. Then p as in (B.1.2) satisfies*

$$\sup_{x \in [-1, 1]^2} |f(x) - p(x)| \leq \frac{(\pi/2)^k M_k(f)}{k!} \frac{n_1 \log n_2}{\binom{n_2+1}{k}} \in \tilde{O}(n^{1-k})$$

for $n_1, n_2 \approx n/2$.

Of course, this is worst-case behavior, so in practice f will converge to its best polynomial approximant even faster.

While it is often the case that we focus on functions defined only on $\mathbb{D} \subset S \subset \mathbb{C}$, in our setting we assume that any such function of focus extends continuously to all of S in such a way that allows us to apply Theorem B.1.6, i.e. that it extends with k -times-differentiable real and imaginary parts. We name this the following:

Property B.1.7 ((γ, k) -extensibility). A function $f : \mathbb{D} \rightarrow \mathbb{C}$ is said to be (γ, k) -extensible if there exists an extension $\tilde{f} : S \rightarrow \mathbb{C}$ satisfying $\tilde{f} = f$ on $\mathbb{D}$, $\|\tilde{f}\|_{L^\infty S} \leq \gamma$, and $\operatorname{re} \tilde{f}, \operatorname{im} \tilde{f} \in C^k[-1, 1]$.

Theorem B.1.6 has a natural version for S^d . For $f \in CS^d$, let

$$M_{r,k}(f) := \max_{2 \leq \ell \leq d} \sup_{t \in S^{d-1}} \left\| \partial_\ell^k (\operatorname{re} f)(t_1, \dots, t_{\ell-1}, \cdot, t_\ell, \dots, t_{d-1}) \right\|_{L^\infty[-1, 1]},$$

$$M_{i,k}(f) := \max_{2 \leq \ell \leq d} \sup_{t \in S^{d-1}} \left\| \partial_\ell^k (\operatorname{im} f)(t_1, \dots, t_{\ell-1}, \cdot, t_\ell, \dots, t_{d-1}) \right\|_{L^\infty[-1, 1]},$$

and $M_k(f) := \max\{M_{r,k}, M_{i,k}\}$. The same reasoning as before gets:

Theorem B.1.8. *For all $n_1, \dots, n_d \in \mathbb{N}$, let $f \in C[-1, 1]^2$ satisfy*

$$\operatorname{re} f(t_1, \dots, t_{\ell-1}, \cdot, t_\ell, \dots, t_{d-1}), \operatorname{im} f(t_1, \dots, t_{\ell-1}, \cdot, t_\ell, \dots, t_{d-1}) \in C^k[-1, 1]$$

for all $t \in S^{d-1}$ and $2 \leq \ell \leq d$ and some $0 \leq k \leq \min_{2 \leq \ell \leq d} n_\ell$. Then p as in (B.1.2) satisfies

$$\sup_{x \in S^d} |f(x) - p(x)| \leq \frac{(\pi/2)^k M_k(f)}{k!} \sum_{2 \leq \ell \leq d} \frac{\log n_\ell}{\binom{n_\ell+1}{k}} \prod_{1 \leq j < \ell} n_j.$$

B.2 Exponentiation of normal matrices

The core observation—that normal M decomposes into two attainable commuting Hermitians—can be used also for ‘Hamiltonian’ simulation, where the goal is to obtain access to a block-encoding of e^{Mt} for $t \in \mathbb{R}$, though as M is not actually Hermitian the resulting exponential is itself not unitary. Nonetheless, it is plausible to desire such a matrix function, and the form (3.2.1) readily resolves this task:

$$e^{Mt} = e^{At} e^{iBt} \quad (\text{B.2.1})$$

precisely because A and B (or really, At and iBt) commute. e^{At} is unitary and achievable by standard methods (e.g. Trotterization, cf. [L’22, Chapter 5]). e^{iBt} is not unitary but is approximable using a block-encoding of B and then either QET as in [L’22, §7.5], by choosing a polynomial approximation to $x \mapsto e^{ixt}$, or as in [T’+21], by choosing a quadrature. It is then a matter of taking the product, adding the ancilla counts, and updating precisions accordingly (Proposition 3.1.1).

B.3 Optimizing the number of block-encoding instances

QSP for block-encodings of Hermitian matrices entails using linear-in-polynomial-degree-many instances of the block-encoding, while our method is quadratic. We use relatively little in our method specific to normalcy, after using the existence of (commuting) $A = \text{re } M$ and $B = \text{im } M$. One might hope to find a modification where this count is $o(D^2)$ without sacrificing too many ancillae. Unfortunately, this turns out to not be possible with any similar scheme:

Proposition B.3.1. *Let $G \in \mathbb{C}[x, y]$ have degree exactly $D-1$ and say $g(x+iy) := G(x, y)$. Suppose $\Theta(D)$ of G ’s “maximal monomials,” i.e. those of the form $x^\ell y^m$, $\ell + m = D-1$, have nonzero coefficients. Then G has no decomposition as*

$$g(x+iy) = G(x, y) = \sum_{1 \leq k \leq E} P_k(x) Q_k(y) \quad (\text{B.3.2})$$

where $P_k, Q_k \in \mathbb{C}[z]$ for $1 \leq k \leq E$, and $E \in o(D)$.

For clarity, we assume that in fact all D of G ’s maximal monomials have nonzero coefficient; however, the result certainly holds for any constant fraction of D .

Notice that this result holds no matter how large the degrees of P_k and Q_k are allowed to be.

Proof. Write

$$G(x, y) = \sum_{\substack{\ell, m \geq 0 \\ \ell + m < D}} c_{\ell, m} x^\ell y^m$$

and suppose that such $P_k(x) = \sum_{0 \leq j < D'} a_{k,j} x^j$ and $Q_k(y) = \sum_{0 \leq j < D'} b_{k,j} y^j$ as in (B.3.2) exist, with $\deg P_k, \deg Q_k < D'$ (any $D' \in \mathbb{N}$ is possible here) for all $1 \leq k \leq E$. The $x^\ell y^m$ coefficient of (B.3.2)'s right-hand side is

$$\sum_{1 \leq k \leq E} a_{k,\ell} b_{k,m} = c_{j,\ell}. \quad (\text{B.3.3})$$

This is just a dot product! So we are inspired to consider the following matrices:

$$\begin{aligned} \mathcal{A} &:= (a_{k,j})_{\substack{1 \leq k \leq E \\ 0 \leq j < D'}} = \begin{pmatrix} a_{1,0} & a_{2,0} & \dots & a_{E,0} \\ a_{1,1} & a_{2,1} & \dots & a_{E,1} \\ \vdots & \vdots & \ddots & \vdots \\ a_{1,D'} & a_{2,D'} & \dots & a_{E,D'} \end{pmatrix} \\ \mathcal{B} &:= (b_{k,j})_{\substack{1 \leq k \leq E \\ 0 \leq j < D'}} = \begin{pmatrix} b_{1,0} & b_{2,0} & \dots & b_{E,0} \\ b_{1,1} & b_{2,1} & \dots & b_{E,1} \\ \vdots & \vdots & \ddots & \vdots \\ b_{1,D'} & b_{2,D'} & \dots & b_{E,D'} \end{pmatrix} \\ C &:= (c_{\ell,m})_{\substack{0 \leq \ell < D \\ 0 \leq m < D}} = \begin{pmatrix} c_{0,0} & c_{0,1} & \dots & c_{0,D-1} \\ c_{1,0} & \ddots & \ddots & \\ \vdots & \ddots & & \\ c_{D-1,0} & & & \end{pmatrix} \end{aligned}$$

where in C , the unspecified entries are all 0. Then, by (B.3.3), $\mathcal{A}\mathcal{B}^\top = C$. We have that $E \ll D$ so $\text{rank } \mathcal{A}, \text{rank } \mathcal{B} \leq E$, thus $\text{rank}(\mathcal{A}\mathcal{B}^\top) \leq E \ll D$.

However, $\det C = (-1)^{\lfloor n/2 \rfloor} \prod_{0 \leq k < D} c_{k,D-1-k} \neq 0$ by hypothesis. Therefore C is full-rank, i.e. $\text{rank } C = D - 1$. This stands in contradiction to $\text{rank } C = \text{rank}(\mathcal{A}\mathcal{B}^\top) \in o(D)$. ■

We interpret this as follows. The sum $\sum_{0 \leq k \leq E} (\deg P_k + \deg Q_k)$ is, up to a small constant factor, the number of calls to the block-encoding of the given normal matrix for implementing the decomposition (B.3.2) using the QSP/LCU approach described herein. If we choose to approximate f by $p(z, \bar{z}) \in \mathbb{C}[z, \bar{z}]$, then p is characterized by $O(D^2)$ complex numbers (one for each term $z^\ell \bar{z}^m$ for $\ell + m < D$), so we end up using about D -many polynomials, many or possibly all of which have degree $\Theta(D)$.

B.4 Chebyshev expansion of Chebyshev interpolants

In Section 3.2.3 and Section 3.4.3 it becomes necessary to evaluate expressions of the form

$$\int_{-1}^1 T_j(x) L_n^{(m)}(x) \frac{dx}{\sqrt{1-x^2}}.$$

Here, we turn this integral into a sum of $n-1$ terms.

First, we compute the Chebyshev coefficients for monomials:

Lemma B.4.1. *Take integers $m, j \geq 0$. Suppose $m \geq j$ and m and j share the same parity. Then,*

$$\int_{-1}^1 T_j(x) x^m \frac{dx}{\sqrt{1-x^2}} = \frac{\pi}{2^m} \binom{m}{\frac{m-j}{2}}.$$

Otherwise, the integral evaluates to 0.

Proof. The “otherwise” case is trivial by parity considerations and since the Chebyshev polynomials span $\mathbb{R}[x]$.

For convenience, define $T_{-n} := T_n$ for $n > 0$ (in contrast to the convention $T_n \equiv 0$ sometimes used). The result then follows by iteratively multiplying T_n by $2x$: $2xT_n(x) = T_{n+1}(x) + T_{n-1}(x)$, which holds regardless of n ’s sign. In particular,

$$(2x)^a T_b(x) = \sum_{0 \leq c < a} \binom{a}{c} T_{a+b-2c}(x). \quad \blacksquare$$

Remark B.4.2. It is a standard fact that

$$T_n(x) = \frac{n}{2} \sum_{k=0}^{\lfloor n/2 \rfloor} (-1)^k \frac{2^{n-2k}}{n-k} \binom{n-k}{k} x^{n-2k}. \quad (\text{B.4.3})$$

Proposition B.4.4. *For the Chebyshev polynomial T_n , if k shares n ’s parity then let $a_{n,k} := \frac{n}{2} (-1)^{\frac{n-k}{2}} \frac{2^k n}{n+k} \binom{\frac{n+k}{2}}{\frac{n-k}{2}}$, otherwise let $a_{n,k} := 0$. Put $f_j(y) := \sum_{0 \leq \ell < j} a_\ell y^{j-\ell}$ (which satisfies $f_j(y) = y f_{j-1}(y) + a_j$ for $j > 0$). Then,*

$$\int_{-1}^1 T_k(x) L_n^{(i)}(x) \frac{dx}{\sqrt{1-x^2}} = \pi \sum_{\substack{k \leq j < n \\ 2|(j-k)}} \frac{f_j(t_{n,i})}{2^j} \binom{j}{\frac{j-k}{2}}.$$

The advantage of this formulation is that it gives the integral in terms of a finite sum, in particular requiring $O(n)$ computations (since $f_k(t_{n,i})$ and $\binom{k}{0} = 1$ can be computed initially as the base case for recursion).




Figure B.1: The circuits for the gate W (left) and W' (right) [GMF²¹, Figure 1].

Proof. From (B.4.3) we perform polynomial long division of $x - t_{n,i}$ into $T_n(x)$, and then use Lemma B.4.1. ■

B.5 Nonlinear transformation of complex amplitudes primitives

In the interest of self-containedness, we present here the definitions of the gates W , $\tilde{G}$, and $\tilde{G}'$ as used in Section 3.3.

In Figure B.1, U is the prepare oracle for the starting state, i.e.

$$U|0\rangle^{\otimes n} = \sum_{0 \leq k < n} c_k |k\rangle,$$

and S and H are the standard 1-qubit gates (in $\text{PU}(2)$)

$$S = \begin{pmatrix} 1 & \\ & i \end{pmatrix}, \quad H = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

(The function of the S in W' but not W is to distinguish between the real and imaginary parts of the amplitudes.) The effect of W is to establish

$$W|k\rangle|0\rangle^{\otimes n}|0\rangle = |k\rangle(|v\rangle|+\rangle + |k\rangle|-\rangle)$$

(using the single-qubit states $|\pm\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$).

In Figure B.2, Z refers to the standard 1-qubit gate (in $\text{PU}(2)$)

$$Z = \begin{pmatrix} 1 & \\ & -1 \end{pmatrix}$$

and S_0 is reflection about the all-zeros state of the bottom two registers $|0\rangle^{\otimes n}|0\rangle$. The effect of G 's construction is to find an operator whose eigenstates include the explicitly-known eigenpairs

$$\left(-\operatorname{re} c_k \pm i\sqrt{1 - (\operatorname{re} c_k)^2}, \frac{1}{\sqrt{2}}|k\rangle \left(\frac{1}{\alpha_k}(|v\rangle + |k\rangle)|0\rangle \pm \frac{1}{\beta_k}(|v\rangle - |k\rangle)|1\rangle \right) \right)$$




Figure B.2: The circuits for the gate G (left) and G' (right) [GMF'21, Figure 2].




Figure B.3: The circuit for the gate $\tilde{G}$ (left) and $\tilde{G}'$ (right) [GMF'21, Figure 3].

(for $\alpha_k, \beta_k > 0$ appropriate normalizing factors).

Finally, in Figure B.3, X is the standard 1-qubit gate (in $\text{PU}(2)$)

$$X = \begin{pmatrix} & 1 \\ 1 & \end{pmatrix};$$

the circuit functions to give a LCU of G and $G^\dagger$ (or G' and $(G')^\dagger$), to eliminate any remaining imaginary parts.

Appendix C

Appendix to Chapter 4 (deferred proofs)

Proof of Proposition 4.0.20. Since $\mathcal{L}(|u_i\rangle\langle u_j|) \subseteq \text{span}(|u_i\rangle\langle u_j|)$, we can write $\mathcal{L}(|u_i\rangle\langle u_j|) = c_{i,j}|u_i\rangle\langle u_j|$ for some $c_{i,j} \in \mathbb{C}$.

We use the same notations as in Definition 4.0.8. Since for $i \neq j$, each $\text{span}(|u_i\rangle\langle u_j|)$ forms an invariant subspace under $\mathcal{L}$, we only need to lower bound the spectral gap in this subspace. Note that for any observable f in this subspace, we can write $f = c|u_i\rangle\langle u_j|$ with $c \in \mathbb{C}$, $\mathcal{E}_{\mathcal{L}}(f) = -\sqrt{\pi(u_i)\pi(u_j)}|c|^2 c_{ij}$ and $\text{Var}(f) = \sqrt{\pi(u_i)\pi(u_j)}|c|^2$, so if $\pi(i)\pi(j) = 0$ then $\lambda(f) = +\infty$ and we are done. From now on, assume $\pi(u_i), \pi(u_j) \neq 0$. In this case, the spectral gap in this subspace is $-c_{ij}$.

Let E_i be the eigenvalue corresponding to eigenvector u_i . One can directly calculate (see also [T13, Theorem 11]):

$$\begin{aligned}
 -c_{ij} &= -\langle |u_i\rangle\langle u_j|, \mathcal{L}(|u_i\rangle\langle u_j|) \rangle \\
 &= \sum_S -\left(G(0)\langle u_i|S^\dagger|u_i\rangle\langle u_j|S|u_j\rangle - \frac{1}{2}\left(\sum_k G(E_k - E_j)|\langle u_j|S|u_k\rangle|^2 + \sum_k G(E_k - E_i)|\langle u_i|S|u_k\rangle|^2\right) \right) \\
 &= \frac{1}{2}G(0) \sum_S |\langle u_i|S|u_i\rangle - \langle u_j|S|u_j\rangle|^2 \\
 &\quad + \frac{1}{2}\left(\sum_{k \neq j} G(E_k - E_j) \sum_S |\langle u_j|S|u_k\rangle|^2 + \sum_{k \neq i} G(E_k - E_i) \sum_S |\langle u_i|S|u_k\rangle|^2\right) \\
 &= \frac{1}{2}G(0) \sum_S |\langle u_i|S|u_i\rangle - \langle u_j|S|u_j\rangle|^2 + \frac{1}{2}\left(\sum_{k \neq j} P_{\mathcal{L},U}[u_j \rightarrow u_k] + \sum_{k \neq i} P_{\mathcal{L},U}[u_i \rightarrow u_k]\right) \\
 &\geq \frac{1}{2}\left(\sum_{k \neq j} P_{\mathcal{L},U}[u_j \rightarrow u_k] + \sum_{k \neq i} P_{\mathcal{L},U}[u_i \rightarrow u_k]\right)
 \end{aligned}$$

$$\geq \frac{1}{2}\lambda_{\mathcal{L},U,\text{cl}},$$

which finishes the claim.

To prove the last inequality, consider the test function $F : U \rightarrow \mathbb{R}$ where $F(u_k) = [k = i]$, and noting that the classical Markov generator is reversible (see proof of Proposition 4.0.13), then

$$\mathcal{E}_{\mathcal{L},U,\text{cl}}(F) = \frac{1}{2} \left(\sum_{k \neq i} \pi(u_i) P_{\mathcal{L},U}(u_i \rightarrow u_k) + \sum_{k \neq i} \pi(u_k) P_{\mathcal{L},U}(u_k \rightarrow u_i) \right) = \pi(u_i) \sum_{k \neq i} P_{\mathcal{L},U}(u_i \rightarrow u_k)$$

and

$$\text{Var}_{\pi}(F) = \pi(u_i)(1 - \pi(u_i));$$

thus

$$\sum_{k \neq i} P_{\mathcal{L},U}(u_i \rightarrow u_k) \geq (1 - \pi(u_i))\lambda_{\mathcal{L},U,\text{cl}}.$$

Similarly,

$$\sum_{k \neq j} P_{\mathcal{L},U}[u_j \rightarrow u_k] \geq (1 - \pi(u_j))\lambda_{\mathcal{L},U,\text{cl}},$$

so

$$\sum_{k \neq j} P_{\mathcal{L},U}[u_j \rightarrow u_k] + \sum_{k \neq i} P_{\mathcal{L},U}[u_i \rightarrow u_k] \geq (2 - \pi(u_i) - \pi(u_j))\lambda_{\mathcal{L},U,\text{cl}} \geq \lambda_{\mathcal{L},U,\text{cl}}.$$

When H has simple spectrum and simple Bohr spectrum, each subspace V_{ω} has dimension 1, so the precondition holds and H has a unique orthonormal eigenbasis, so $\lambda_{\mathcal{L}} \geq \frac{1}{2}\lambda_{\mathcal{L},\text{cl}}$. $\blacksquare$

Proof of Proposition 4.1.9. Note that S is self-adjoint, and $S^{\dagger}(-\omega) = S(\omega)^{\dagger}$, so we can write

$$\begin{aligned} & 2\langle \mathcal{L}_{\omega,S}(f) + \mathcal{L}_{-\omega,S^{\dagger}}(f), f \rangle_{\rho} \\ &= G(\omega) \langle S(\omega)^{\dagger}[f, S(\omega)] + [S(\omega)^{\dagger}, f]S(\omega), f \rangle_{\rho} \\ &\quad + G(-\omega) \langle S^{\dagger}(-\omega)^{\dagger}[f, S^{\dagger}(-\omega)] + [S^{\dagger}(-\omega)^{\dagger}, f]S^{\dagger}(-\omega), f \rangle_{\rho} \\ &= G(\omega) \left(\text{tr} \left(S(\omega)^{\dagger}[f, S(\omega)] \rho^{1/2} f^{\dagger} \rho^{1/2} \right) + \text{tr} \left([S(\omega)^{\dagger}, f] S(\omega) \rho^{1/2} f^{\dagger} \rho^{1/2} \right) \right) \\ &\quad + G(-\omega) \left(\text{tr} \left(S(\omega)[f, S(\omega)^{\dagger}] \rho^{1/2} f^{\dagger} \rho^{1/2} \right) + \text{tr} \left([S(\omega), f] S(\omega)^{\dagger} \rho^{1/2} f^{\dagger} \rho^{1/2} \right) \right) \\ &=_{(1)} G(\omega) e^{\beta\omega/2} \left(\text{tr} \left([f, S(\omega)] \rho^{1/2} f^{\dagger} S(\omega)^{\dagger} \rho^{1/2} \right) + \text{tr} \left([S(\omega)^{\dagger}, f] \rho^{1/2} S(\omega) f^{\dagger} \rho^{1/2} \right) \right) \\ &\quad + G(-\omega) e^{-\beta\omega/2} \left(\text{tr} \left([f, S(\omega)^{\dagger}] \rho^{1/2} f^{\dagger} S(\omega) \rho^{1/2} \right) + \text{tr} \left([S(\omega), f] \rho^{1/2} S(\omega)^{\dagger} f^{\dagger} \rho^{1/2} \right) \right) \end{aligned}$$

$$\begin{aligned}
&= \tilde{G}(\omega) \left(\text{tr} \left([f, S(\omega)] \rho^{1/2} f^\dagger S(\omega)^\dagger \rho^{1/2} \right) + \text{tr} \left([S(\omega), f] \rho^{1/2} S(\omega)^\dagger f^\dagger \rho^{1/2} \right) \right) \\
&\quad + \tilde{G}(-\omega) \left(\text{tr} \left([S(\omega)^\dagger, f] \rho^{1/2} S(\omega) f^\dagger \rho^{1/2} \right) + \text{tr} \left([f, S(\omega)^\dagger] \rho^{1/2} f^\dagger S(\omega) \rho^{1/2} \right) \right) \\
&=_{(2)} -\tilde{G}(\omega) \text{tr} \left([S(\omega), f] \rho^{1/2} [S(\omega), f]^\dagger \rho^{1/2} \right) - \tilde{G}(-\omega) \text{tr} \left([S(\omega)^\dagger, f] \rho^{1/2} [S(\omega)^\dagger, f]^\dagger \rho^{1/2} \right) \\
&=_{(3)} -\tilde{G}(\omega) \| [S(\omega), f] \|_\rho^2 - \tilde{G}(-\omega) \| [S^\dagger(-\omega), f] \|_\rho^2,
\end{aligned}$$

where (1) is due to applying Proposition 4.1.7 to $S(\omega) \in V_\omega$ and $S(\omega)^\dagger \in V_{-\omega}$, (2) is by rearranging terms and the fact that $\tilde{G}(\omega) = G(\omega) e^{\frac{\beta\omega}{2}} = \tilde{G}(-\omega)$, and (3) is by definition of $\| \cdot \|_\rho$ and $S(\omega)^\dagger = S^\dagger(-\omega)$. Hence, we obtain the desired equality

$$\begin{aligned}
\mathcal{E}_\mathcal{L}(f) &= - \left\langle \sum_{S \in \mathcal{S}, \omega} \mathcal{L}_{\omega, S}(f), f \right\rangle_\rho \\
&= -\frac{1}{2} \sum_{S \in \mathcal{S}, \omega} \langle \mathcal{L}_{\omega, S}(f) + \mathcal{L}_{-\omega, S^\dagger}(f), f \rangle_\rho \\
&= \frac{1}{4} \sum_{S \in \mathcal{S}, \omega} (\tilde{G}(\omega) \| [S(\omega), f] \|_\rho^2 + \tilde{G}(-\omega) \| [S^\dagger(-\omega), f] \|_\rho^2) = \frac{1}{2} \sum_{S, \omega} \tilde{G}(\omega) \| [S(\omega), f] \|_\rho^2. \quad \blacksquare
\end{aligned}$$

Proof of Proposition 4.3.6. For $x \neq y$, let $Q(x, y) = \pi(x)P[x \rightarrow y]$. We adapt the standard proof of Cheeger's inequality from [LPW'17]. We first show the following claim:

Claim C.1. Given a function $\psi : X \rightarrow \mathbb{R}_{\geq 0}$, order X so that ψ is non-increasing. If $\pi(\psi > 0) \leq 1/2$ then

$$\mathbb{E}_\pi[\psi] \leq \Phi^{-1} \sum_{\substack{x, y \in X \\ x < y}} (\psi(x) - \psi(y)) Q(x, y)$$

Proof. Let $S = \{x : \psi(x) > t\}$ with $t > 0$. Since $\pi(S) \leq \pi(\psi > 0) \leq 1/2$, we use the definition of the bottleneck ratio Φ to obtain:

$$\Phi \leq \frac{Q(S, S^c)}{\pi(S)} = \frac{\sum_{x, y \in X} Q(x, y) [\psi(x) > t \geq \psi(y)]}{\pi(\psi > t)}$$

Rearranging and noting that $\psi(x) > \psi(y)$ only for $x < y$, we get:

$$\pi(\psi > t) \leq \Phi^{-1} \sum_{x, y \in X} Q(x, y) [\psi(x) > t \geq \psi(y)]$$

Integrating over t , noting that

$$\int_0^\infty [\psi(x) > t \geq \psi(y)] dt = \psi(x) - \psi(y),$$

and that $\mathbb{E}_\pi[\psi] = \int_0^\infty \pi(\psi > t)$, we obtain:

$$\mathbb{E}_\pi[\psi] \leq \Phi^{-1} \sum_{x < y} (\psi(x) - \psi(y))Q(x, y). \quad \blacksquare$$

Define the inner product $\langle \cdot, \cdot \rangle_\pi$ on the space of real-valued functions $f : X \rightarrow \mathbb{R}$ by

$$\langle f, g \rangle_\pi := \sum_{x \in X} \pi(x) f(x) g(x).$$

We view P as a matrix with entries $P[x, y] = P[x \rightarrow y]$, and set $P[x \rightarrow x] := -\sum_{y \neq x} P[x \rightarrow y]$, so that each row sums to zero. Note that the Dirichlet form can be written as $\mathcal{E}(f) = \langle -Pf, f \rangle_\pi$.

Since P is self-adjoint with respect to $\langle \cdot, \cdot \rangle_\pi$, the operator $-P$ has real, nonnegative eigenvalues. Its smallest eigenvalue is 0, and the spectral gap λ is the second smallest eigenvalue. Let $f_2 : X \rightarrow \mathbb{R}$ be an eigenfunction corresponding to λ , so that $-Pf_2 = \lambda f_2$. Without loss of generality, assume $\pi(f_2 > 0) \leq 1/2$; otherwise, we may replace f_2 with $-f_2$. Define a nonnegative function $f : X \rightarrow \mathbb{R}_{\geq 0}$ by

$$f(x) := \max(f_2(x), 0),$$

and let $g := -Pf$.

We show that, for all x ,

$$g(x) \leq \lambda f(x) \tag{C.2}$$

There are two cases, $f(x) = 0$ and $f(x) > 0$.

If $f(x) = 0$: $g(x) = -\sum_{y: y \neq x} P[x \rightarrow y] f(y) \leq 0 = f(x)$.

If $f(x) > 0$: In this case, $f(x) = f_2(x)$ and since $f(y) \geq f_2(y)$ for all y , we have

$$\begin{aligned} g(x) &= -P[x \rightarrow x]f(x) - \sum_{y: y \neq x} P[x \rightarrow y]f(y) = -P[x \rightarrow x]f_2(x) - \sum_{y: y \neq x} P[x \rightarrow y]f(y) \\ &\leq -P[x \rightarrow x]f_2(x) - \sum_{y: y \neq x} P[x \rightarrow y]f_2(y) = (-Pf_2)(x) = \lambda f_2(x) = \lambda f(x). \end{aligned}$$

Since $f(x) \geq 0$ for all x , (C.2) yields

$$\mathcal{E}(f) = \langle -Pf, f \rangle_\pi = \sum_x \pi(x) g(x) f(x) \leq \lambda \sum_x \pi(x) f(x)^2 = \lambda \langle f, f \rangle_\pi. \tag{C.3}$$

Applying Claim C.1 for the function $\psi = f^2$, and use the ordering of X induced by ψ , we get:

$$\langle f, f \rangle_\pi = \mathbb{E}_\pi[\psi] \leq \Phi^{-1} \sum_{x < y} (f^2(x) - f^2(y))Q(x, y)$$

Since both sides are nonnegative, we may square the inequality and then apply the Cauchy–Schwarz inequality to get:

$$\begin{aligned} \langle f, f \rangle_\pi^2 &\leq \Phi^{-2} \left(\sum_{x < y} (f^2(x) - f^2(y)) Q(x, y) \right)^2 \\ &\leq \Phi^{-2} \left(\sum_{x < y} (f(x) - f(y))^2 Q(x, y) \right) \left(\sum_{x < y} (f(x) + f(y))^2 Q(x, y) \right) \end{aligned}$$

Using the identities $(f(x) + f(y))^2 = 2(f(x)^2 + f(y)^2) - (f(x) - f(y))^2$, $\sum_{x < y} (f(x) - f(y))^2 Q(x, y) = \mathcal{E}(f)$, and $Q(x, y) = Q(y, x)$, we can rewrite

$$\begin{aligned} \sum_{x < y} (f(x) + f(y))^2 Q(x, y) &= 2 \sum_{x < y} (f(x)^2 + f(y)^2) \pi(x) P[x \rightarrow y] - \sum_{x < y} (f(x) - f(y))^2 Q(x, y) \\ &= 2 \sum_{x \in X} \left(\pi(x) f(x)^2 \sum_{y \neq x} P[x \rightarrow y] \right) - \mathcal{E}(f) \\ &\leq 2\Lambda \langle f, f \rangle_\pi - \mathcal{E}(f) \end{aligned}$$

thus

$$\langle f, f \rangle_\pi^2 \leq \Phi^{-2} \mathcal{E}(f) (2\Lambda \langle f, f \rangle_\pi - \mathcal{E}(f)) \leq 2\Lambda \Phi^{-2} \mathcal{E}(f) \langle f, f \rangle_\pi$$

Dividing both sides by $\langle f, f \rangle_\pi \geq 0$ and applying (C.3), we obtain

$$\Phi^2 \leq 2\Lambda\lambda. \quad \blacksquare$$

Appendix D

Appendix to Chapter 5 (deferred proofs)

Proof of Proposition 5.0.10. Assume $f : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ with $\|f\|^2 = 1$ satisfies

$$\theta^2 \|\nabla f\|^2 + \int_{\mathbb{R}^d} V(x) f(x)^2 dx \leq \theta^r,$$

for r a constant depending on p and d to be chosen later. Let

$$W := \{x : V(x) \leq 2\theta^r\}$$

and notice by Markov's inequality [A⁺24] (applied to the probability distribution defined by the mass of f) that

$$\sum_C \|f|_C\|^2 \frac{\|f|_{\overline{W} \cap C}\|^2}{\|f|_C\|^2} = \|f|_{\overline{W}}\|^2 \leq \frac{1}{2},$$

where the sum is over unit volume cubes in $\mathbb{R}^d$. Let S be the set of cubes C satisfying

$$\frac{\|f|_{\overline{W} \cap C}\|^2}{\|f|_C\|^2} \leq \frac{3}{4},$$

and notice by Markov's inequality that

$$\sum_{C \in S} \|f|_C\|^2 \geq c$$

for a constant c . Let $q > 2$. For each cube $C \in S$ we have $\|f|_{W \cap C}\|^2 \geq \frac{\|f|_C\|^2}{4}$ as well as

$$\frac{\|f|_{W \cap C}\|_q^2}{\|f|_{W \cap C}\|^2} \geq \text{vol}(W \cap C)^{\frac{2}{q}-1},$$

which together yield

$$\|f\|_q^2 \geq \sum_{C \in \mathcal{S}} \|f|_{W \cap C}\|_q^2 \geq \frac{1}{4} \sum_{C \in \mathcal{S}} \text{vol}(W \cap C)^{\frac{2}{q}-1} \|f|_C\|^2.$$

Since f has isolated critical points and is uniformly p -coercive, we have the uniform estimate

$$\text{vol}(W \cap C) = O\left(\theta^{\frac{rd}{p}}\right)$$

for each cube C . Substituting, we have

$$\|f\|_q^2 \geq \Omega(1) \cdot \theta^{\frac{rd}{p}\left(\frac{2}{q}-1\right)} \sum_{C \in \mathcal{S}} \|f|_C\|^2 = \Omega(1) \cdot \theta^{\frac{rd}{p}\left(\frac{2}{q}-1\right)}.$$

(b) If $d \geq 3$ set $q = \frac{2d}{d-2}$ and apply the Sobolev inequality to obtain

$$\|\nabla f\|^2 = \Omega(1) \cdot \theta^{-\frac{2r}{p}},$$

whence $\theta^2 \|\nabla f\|^2 = \Omega\left(\theta^{2-\frac{2r}{p}}\right)$. Choosing $r = 2 - \frac{2r}{p} = \frac{2p}{p+2}$ yields the claim.

(a) If $d = 1, 2$ we instead use the Sobolev inequalities

$$\|\nabla f\|^2 + \|f\|^2 \geq C_q \|f\|_q^2,$$

for $q = \infty$ and q large, respectively. Since $\|f\|^2 = 1$ and the right hand side is unbounded as $\theta \rightarrow 0$, this term can be ignored and we have $\|f\|_q^2 \geq \Omega(1) \cdot \theta^{-\frac{r}{p}}$ for $d = 1$ and $\|f\|_q^2 \geq \Omega(1) \cdot \theta^{-\frac{2r}{p}+o(1)}$ for $d = 2$. Choosing $r = 2 - \frac{r}{p}$ and $r = 2 - \frac{2r}{p} + o(1)$ yields $r = \frac{2p}{p+1}$ and $r = \frac{2p}{p+2} + o(1)$, as advertised.

Now, we set out to prove the upper bound (c) through a different approach. By boundedness, there is a constant H such that $V(x) \leq H\|x\|^p =: W(x)$ where we assume that $V(0) = 0$ and that 0 is a critical point of order p . For all test functions g , $\mathcal{R}_{B_\theta}(g) \leq \mathcal{R}_{-\theta^2\Delta+W}(g)$. Let

$$g_\sigma(x) = \exp\left(-\frac{\|x\|^2}{2\sigma^2}\right).$$

Then,

$$\begin{aligned} \langle g_\sigma, W g_\sigma \rangle &= H \int_{\mathbb{R}^d} \|x\|^p g_\sigma(x) dx \\ &= HS_{d-1} \int_{\mathbb{R}_{\geq 0}} r^{p+d-1} \exp\left(-\frac{r^2}{2\sigma^2}\right) dr \end{aligned}$$

$$= HS_{d-1} 2^{\frac{p+d-2}{2}} \left(\frac{p+d-2}{2} \right)! \sigma^{p+d}.$$

where $S_{d-1} = \frac{2\sqrt{\pi}^d}{\left(\frac{d-2}{2}\right)!}$ is the surface area of S^{d-1} . We also compute directly that

$$\begin{aligned} \langle g_\sigma, (-\theta^2 \Delta + H) g_\sigma \rangle &= \frac{1}{2} \left(d \frac{\theta^2}{\sigma^2} + \text{tr } H \sigma^2 \right) \\ \|g_\sigma\|^2 &= \left(\sqrt{\pi} \sigma \right)^d \end{aligned}$$

so that all told,

$$\langle g_\sigma, (-\theta^2 \Delta + W) g_\sigma \rangle = \frac{1}{2} d \frac{\theta^2}{\sigma^2} + 2^{\frac{p+d}{2}} H \frac{\left(\frac{p+d-2}{2}\right)!}{\left(\frac{d-2}{2}\right)!} \sigma^p.$$

By weighted AM–GM this is minimized by

$$\frac{2d}{p} \frac{p^{\frac{p}{2}+1}}{2^{\frac{p}{2}+1}} \sqrt{\frac{p}{d} 2^{\frac{d-2}{2}} H \frac{\left(\frac{p+d-2}{2}\right)!}{\left(\frac{d-2}{2}\right)!} \theta^{\frac{2p}{p+2}}} \quad (\text{D.1})$$

where the exponent of θ matches the claim. ■

Proof of Lemma 5.2.1. Observe that $\|A_{\theta,0} - A_{\theta,\eta}\| \leq \sup_{n \in \mathbb{Z}^d} |V(n\theta + \eta) - V(n\theta)| \leq \omega(\eta)$ where $\omega(\cdot)$ is the modulus of continuity of V , so $\mu(A_{\theta,\eta})$ is continuous in η .

Let $\varepsilon > 0$ and $\eta \in \mathbb{R}^d$. Since θ is irrational, ergodicity in each coordinate of $n\theta + \eta$ implies that there is some $\eta' \in \mathbb{R}^d$ and $n_0 \in \mathbb{Z}^d$ such that $|\eta' - \eta| \leq \varepsilon$ and $n_0\theta \equiv \eta' \pmod{1}$ (solve for an appropriate entry of n_0 separately for each component). Observe that $A_{\theta,\eta'} = A_{\theta,0}$ since the Laplacian part of A_θ is translation-invariant and we have $V(n_0\theta + n\theta) = V(n\theta + \eta')$. By continuity we now have $|\mu(A_0) - \mu(A_\eta)| \leq \varepsilon$. Letting $\varepsilon \rightarrow 0$ yields the result. ■

Bibliography

- [A'03] D. Aharonov. “A Simple Proof that Toffoli and Hadamard are Quantum Universal”. 2003. arXiv: [quant-ph/0301040](#) (cit. on p. 43).
- [A'10] S. Akkouché. “The spectral bounds of the discrete Schrödinger operator”. In: *Journal of Functional Analysis* 259.6 (2010), pp. 1443–1465. doi: [10.1016/j.jfa.2010.06.001](#) (cit. on p. 96).
- [ACM'24] P. M. Alsing, C. Cafaro, and S. Mancini. “Feynman’s “Simulating Physics with Computers””. In: *International Journal of Theoretical Physics* 63 (2024). doi: [10.1007/s10773-024-05653-4](#) (cit. on p. 1).
- [A+'23] M. Amy, M. Crawford, A. N. Glaudell, M. L. Macasieb, S. S. Mendelson, and N. J. Ross. “Catalytic Embeddings of Quantum Circuits”. 2023. arXiv: [2305.07720](#) (cit. on p. 37).
- [AGR'20] M. Amy, A. Glaudell, and N. J. Ross. “Number-Theoretic Characterizations of Some Restricted Clifford+T Circuits”. In: *Quantum* 4 (2020). doi: [10.22331/q-2020-04-06-252](#) (cit. on p. 37).
- [A+'22] N. Anari, V. Jain, F. Koehler, H. T. Pham, and T.-D. Vuong. “Entropic independence: optimal mixing of down-up random walks”. In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2022. Rome, Italy: Association for Computing Machinery, 2022, pp. 1418–1430. ISBN: 9781450392648. doi: [10.1145/3519935.3520048](#). URL: <https://doi.org/10.1145/3519935.3520048> (cit. on p. 54).
- [ALO'20] N. Anari, K. Liu, and S. Oveis Gharan. “Spectral Independence in High-Dimensional Expanders and Applications to the Hardcore Model”. In: *IEEE 61st Annual Symposium on Foundations of Computer Science*. IEEE Computer Society, Nov. 2020. doi: [10.1137/20M1367696](#) (cit. on p. 54).
- [A+'24] A. Aouad, A. Amaniamedani, T. Pollner, and A. Saberi. “Improved Approximations for Stationary Bipartite Matching: Beyond Probabilistic Independence”. In: (2024). arXiv: [2411.08218](#) (cit. on p. 132).
- [AB'96] W. Arendt and C. J. K. Batty. “The spectral bound of Schrödinger operators”. In: *Potential Analysis* 5 (1996), pp. 207–230. doi: [10.1007/BF00282361](#) (cit. on p. 96).

- [AB'97] W. Arendt and C. J. K. Batty. "The spectral function and principal eigenvalues for Schrödinger operators". In: *Potential Analysis* 7.1 (1997), pp. 415–436. DOI: [10.1023/A:1017928532615](https://doi.org/10.1023/A:1017928532615) (cit. on p. 96).
- [A+'19] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis. "Quantum supremacy using a programmable superconducting processor". In: *Nature* 574 (2019). DOI: [10.1038/s41586-019-1666-5](https://doi.org/10.1038/s41586-019-1666-5) (cit. on p. 3).
- [AJ'09] A. Avila and S. Jitomirskaya. "The Ten Martini Problem". In: *Annals of Mathematics* 170 (2009), pp. 303–342. DOI: [10.4007/annals.2009.170.303](https://doi.org/10.4007/annals.2009.170.303) (cit. on p. 17).
- [B+'21] Ivan B., Á. Capel, L. Gao, A. Lucia, D. Pérez-García, and C. Rouzé. "Entropy Decay for Davies Semigroups of a One Dimensional Quantum Lattice". In: *Communications in Mathematical Physics* (2021). URL: <https://api.semanticscholar.org/CorpusID:244773159> (cit. on p. 54).
- [B+'25] J. Basso, S. Ganguly, A. Sinclair, N. Srivastava, Z. Stier, and T.-D. Vuong. "On quantum to classical comparison for Davies generators". 2025. arXiv: [2510.07267](https://arxiv.org/abs/2510.07267) (cit. on pp. 1, 19).
- [BWZ'26] S. Becker, J. Wittsten, and M. Zworski. "Discrete vs. Continuous in the Semiclassical Limit". In: *SIAM Journal on Mathematical Analysis* 58 (1 2026). DOI: [10.1137/24M1671256](https://doi.org/10.1137/24M1671256) (cit. on pp. 17, 96, 97).
- [BVZ'97] C. Béguin, A. Valette, and A. Zuk. "On the spectrum of a random walk on the discrete Heisenberg group and the norm of Harper's operator". In: *Journal of Geometry and Physics* 21.4 (1997), pp. 337–356. DOI: [10.1016/S0393-0440\(96\)00024-1](https://doi.org/10.1016/S0393-0440(96)00024-1) (cit. on pp. 92, 96).
- [B'64] J. S. Bell. "On the Einstein Podolsky Rosen paradox". In: *Physique Physique Fizika* 1.195 (1964). DOI: [10.1103/PhysiquePhysiqueFizika.1.195](https://doi.org/10.1103/PhysiquePhysiqueFizika.1.195) (cit. on p. 1).

- [B'07] C. M. Bender. "Making sense of non-Hermitian Hamiltonians". In: *Reports on Progress in Physics* 70 (2007). doi: [10.1088/0034-4885/70/6/R03](#) (cit. on p. [11](#)).
- [B'80] P. Benioff. "The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines". In: *Journal of Statistical Physics* 22 (1980), pp. 563–591. doi: [10.1007/BF01011339](#) (cit. on pp. [1](#), [2](#)).
- [BCL'24] T. Bergamaschi, C.-F. Chen, and Y. Liu. "Quantum Computational Advantage with Constant-Temperature Gibbs Sampling". In: *IEEE 65th Annual Symposium on Foundations of Computer Science*. IEEE, 2024, pp. 1063–1085. doi: [10.1109/focs61266.2024.00071](#) (cit. on p. [56](#)).
- [B+'20] M. Beverland, E. Campbell, M. Howard, and V. Kliuchnikov. "Lower bounds on the non-Clifford resources for quantum computations". In: *Quantum Science and Technology* 5.3 (2020). doi: [10.1088/2058-9565/ab8963](#) (cit. on p. [37](#)).
- [BCC'19] E. Bevilacqua, K. Chandran, and Y. Choi. "Ramanujan Congruences for Fractional Partition Functions". 2019. arXiv: [1907.06716](#) (cit. on p. [27](#)).
- [BS'23] T. R. Blackman and Z. Stier. "Fast navigation with Icosahedral Golden Gates". In: *Quantum Information and Computation* 23.11&12 (2023), pp. 901–923. doi: [10.26421/QIC23.11-12-1](#) (cit. on pp. [1](#), [19](#), [117](#)).
- [BZ'05] F. P. Boca and A. Zaharescu. "Norm estimates of almost Mathieu operators". In: *Journal of Functional Analysis* 220.1 (2005), pp. 76–96. doi: [10.1016/j.jfa.2004.09.013](#) (cit. on pp. [17](#), [92](#), [93](#), [96](#)).
- [BRS'15] A. Bocharov, M. Roetteler, and K. M. Svore. "Efficient Synthesis of Universal Repeat-Until-Success Quantum Circuits". In: *Physical Review Letters* 114.8 (2015). doi: [10.1103/PhysRevLett.114.080502](#) (cit. on p. [37](#)).
- [BSS'23] Y. Borns-Weil, T. Saffat, and Z. Stier. "A Quantum Algorithm for Functions of Multiple Commuting Hermitian Matrices". 2023. arXiv: [2302.11139](#) (cit. on pp. [1](#), [11](#), [19](#)).
- [B+'17a] D. Bump, P. Diaconis, A. Hicks, L. Miclo, and H. Widom. "An Exercise (?) in Fourier Analysis on the Heisenberg Group". In: *Annales de la Faculté des sciences de Toulouse: Mathématiques*. Vol. 26. 2017, pp. 263–288. doi: [10.5802/afst.1533](#) (cit. on pp. [92](#), [96](#)).
- [B+'17b] D. Bump, P. Diaconis, A. Hicks, L. Miclo, and H. Widom. "Useful bounds on the extreme eigenvalues and vectors of matrices for Harper's operators". In: *Large Truncated Toeplitz Matrices, Toeplitz Operators, and Related Topics: The Albrecht Böttcher Anniversary Volume* (2017), pp. 235–265. doi: [10.1007/978-3-319-49182-0_13](#) (cit. on pp. [92](#), [96](#)).

- [CM'18] E. A. Carlen and J. Maas. “Non-commutative Calculus, Optimal Transport and Functional Inequalities in Dissipative Quantum Systems”. In: *Journal of Statistical Physics* 178 (2018), pp. 319–378. URL: <https://api.semanticscholar.org/CorpusID:119594279> (cit. on p. 69).
- [CP'18] E. Carvalho Pinto and C. Petit. “Better path-finding algorithms in LPS Ramanujan graphs”. In: *Journal of Mathematical Cryptology* 12.4 (2018). doi: [10.1515/jmc-2017-0051](https://doi.org/10.1515/jmc-2017-0051) (cit. on pp. 8, 24).
- [C+'21] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles. “Variational quantum algorithms”. In: *Nature Reviews Physics* 3 (2021), pp. 625–644. doi: [10.1038/s42254-021-00348-9](https://doi.org/10.1038/s42254-021-00348-9) (cit. on p. 3).
- [CB'21] C.-F. Chen and F. G. S. L. Brandão. “Fast thermalization from the eigenstate thermalization hypothesis”. 2021. arXiv: [2112.07646](https://arxiv.org/abs/2112.07646) (cit. on pp. 54, 55, 62).
- [CKG'23] C.-F. Chen, M. J. Kastoryano, and A. Gilyén. “An efficient and exact non-commutative quantum Gibbs sampler”. 2023. arXiv: [2311.09207](https://arxiv.org/abs/2311.09207) (cit. on pp. 55, 62).
- [CR'25] C.-F. Chen and C. Rouzé. “Quantum Gibbs states are locally Markovian”. In: (2025). arXiv: [2504.02208](https://arxiv.org/abs/2504.02208) (cit. on p. 69).
- [C+'24] H. Chen, B. Li, J. Lu, and L. Ying. “A randomized method for simulating Lindblad equations and thermal state preparation”. 2024. arXiv: [2407.06594](https://arxiv.org/abs/2407.06594) (cit. on pp. 54, 55, 62).
- [C+'22] X. Chen, W. Feng, Y. Yin, and X. Zhang. “Optimal mixing for two-state anti-ferromagnetic spin systems”. In: *IEEE 63rd Annual Symposium on Foundations of Computer Science* (2022), pp. 588–599. doi: [10.1109/FOCS54457.2022.00062](https://doi.org/10.1109/FOCS54457.2022.00062) (cit. on p. 54).
- [CE'22] Y. Chen and R. Eldan. “Localization Schemes: A Framework for Proving Mixing Bounds for Markov Chains (extended abstract)”. In: *IEEE 63rd Annual Symposium on Foundations of Computer Science* (2022), pp. 110–122. URL: <https://api.semanticscholar.org/CorpusID:247315480> (cit. on p. 54).
- [CLV'21] Z. Chen, K. Liu, and E. Vigoda. “Optimal mixing of Glauber dynamics: Entropy factorization via high-dimensional expansion”. In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 1537–1550. doi: [10.1145/3406325.3451035](https://doi.org/10.1145/3406325.3451035) (cit. on p. 54).
- [CW'12] A. M. Childs and N. Wiebe. “Hamiltonian Simulation Using Linear Combinations of Unitary Operations”. 2012. arXiv: [1202.5822](https://arxiv.org/abs/1202.5822) (cit. on pp. 10, 40).

- [CS'17] A. N. Chowdhury and R. D. Somma. “Quantum algorithms for Gibbs sampling and hitting-time estimation”. In: *Quantum Information and Computation* 17.1&2 (2017). doi: [10.26421/QIC17.1-2-3](https://doi.org/10.26421/QIC17.1-2-3) (cit. on p. 13).
- [C'16] H. Cohn. “Factoring may be easier than you think”. 2016. URL: <https://cohn.mit.edu/factoring/> (cit. on p. 2).
- [DEP'25] R. Dalal, S. Evra, and O. Parzanchevski. “Multi-Qubit Golden Gates”. 2025. arXiv: [2509.09047](https://arxiv.org/abs/2509.09047) (cit. on p. 8).
- [D⁺'25] A. M. Dalzell, S. McArdle, M. Berta, P. Bienias, C.-F. Chen, A. Gilyén, C. T. Hann, M. J. Kastoryano, E. T. Khabiboulline, A. Kubica, G. Salton, S. Wang, and F. G. S. L. Brandão. *Quantum Algorithms: A Survey of Applications and End-to-end Complexities*. Cambridge University Press, 2025 (cit. on pp. 3, 13).
- [DSV'03] G. P. Davidoff, P. Sarnak, and A. Valette. *Elementary Number Theory, Group Theory, and Ramanujan Graphs*. Cambridge University Press, 2003 (cit. on p. 26).
- [D'76] E. B. Davies. “Markovian Master Equations. II”. In: *Mathematische Annalen* 219.2 (1976), pp. 147–158. doi: [10.1007/BF01351898](https://doi.org/10.1007/BF01351898) (cit. on p. 13).
- [D'99] E. B. Davies. “Ground state energy of almost periodic Schrödinger operators”. In: *Ergodic Theory and Dynamical Systems* 19.3 (1999), pp. 591–609. doi: [10.1017/S0143385799149058](https://doi.org/10.1017/S0143385799149058) (cit. on p. 96).
- [DN'06] C. M. Dawson and M. A. Nielsen. “The Solovay-Kitaev algorithm”. In: *Quantum Information and Computation* 6.1 (2006), pp. 81–95. doi: [10.26421/QIC6.1-6](https://doi.org/10.26421/QIC6.1-6) (cit. on p. 7).
- [DSS'24] I. Detherage, N. Srivastava, and Z. Stier. “On the Ground State Energies of Discrete and Semiclassical Schrödinger Operators”. In: *Pure and Applied Analysis* 6.4 (2024), pp. 955–976. doi: [10.2140/paa.2024.6.955](https://doi.org/10.2140/paa.2024.6.955) (cit. on pp. 1, 17, 19).
- [D'85] D. Deutsch. “Quantum theory, the Church-Turing principle and the universal quantum computer”. In: *Proceedings of the Royal Society of London A* 400 (1985), pp. 97–117. doi: [10.1098/rspa.1985.0070](https://doi.org/10.1098/rspa.1985.0070) (cit. on p. 2).
- [DJ'92] D. Deutsch and R. Jozsa. “Rapid solutions of problems by quantum computation”. In: *Proceedings of the Royal Society of London A* 439.1907 (1992), pp. 553–558. doi: [10.1098/rspa.1992.0167](https://doi.org/10.1098/rspa.1992.0167) (cit. on p. 2).
- [D'82] D. Dieks. “Communication by EPR devices”. In: *Physics Letter A* 92 (6 1982), pp. 271–272. doi: [10.1016/0375-9601\(82\)90084-6](https://doi.org/10.1016/0375-9601(82)90084-6) (cit. on p. 3).
- [D⁺'21] Y. Dong, X. Meng, K. B. Whaley, and L. Lin. “Efficient phase-factor evaluation in quantum signal processing”. 2021. arXiv: [2002.11649](https://arxiv.org/abs/2002.11649) (cit. on pp. 45, 47, 52).

- [D'47] R. Dubisch. “The Wedderburn structure theorems”. In: *The American Mathematical Monthly* 54.5 (1947), pp. 253–259. doi: [10.2307/2305346](#) (cit. on p. [24](#)).
- [E'25] C. Esparza. “Uniqueness of asymptotically conical shrinking gradient Kähler–Ricci solitons”. 2025. arXiv: [2502.13521](#) (cit. on p. [12](#)).
- [E'22] L. C. Evans. *Partial Differential Equations*. Vol. 19. American Mathematical Society, 2022 (cit. on p. [109](#)).
- [EP'22] S. Evra and O. Parzanchevski. “Ramanujan Complexes and Golden Gates in $PU(3)$ ”. In: *Geometric and Functional Analysis* (2022). doi: [10.1007/s00039-022-00593-9](#) (cit. on p. [8](#)).
- [FLT'25] D. Fang, J. Lu, and Y. Tong. “Mixing time of open quantum systems via hypocoercivity”. In: *Physical Review Letters* 134.14 (2025). doi: [10.1103/PhysRevLett.134.140405](#) (cit. on p. [55](#)).
- [FGG'14] E. Farhi, J. Goldstone, and S. Gutmann. “A Quantum Approximate Optimization Algorithm”. 2014. arXiv: [1411.4028](#) (cit. on p. [3](#)).
- [FEG'17] L. Feng, R. El-Ganainy, and L. Ge. “Non-Hermitian photonics based on parity–time symmetry”. In: *Nature Photonics* 11 (2017), pp. 752–762. doi: [10.1038/s41566-017-0031-1](#) (cit. on p. [11](#)).
- [F'82] R. P. Feynman. “Simulating Physics with Computers”. In: *International Journal of Theoretical Physics* 21.6-7 (1982), pp. 467–488. doi: [10.1007/BF02650179](#) (cit. on p. [1](#)).
- [GGS'92] F. Gesztesy, G. M. Graf, and B. Simon. “The ground state energy of Schrödinger operators”. In: *Communications in Mathematical Physics* 150 (1992), pp. 375–384. URL: <http://math.caltech.edu/SimonPapers/226.pdf> (cit. on pp. [93](#), [96](#)).
- [G'02] J. W. Gibbs. *Elementary Principles of Statistical Mechanics*. Charles Scribner's Sons, 1902. URL: <https://www.gutenberg.org/ebooks/50992> (cit. on p. [12](#)).
- [GS'13] B. Giles and P. Selinger. “Exact synthesis of multiqubit Clifford+ T circuits”. In: *Physical Review A* 87.3 (2013). doi: [10.1103/PhysRevA.87.032332](#) (cit. on p. [37](#)).
- [G+'19] A. Gilyén, Y. Su, G. Low, and N. Wiebe. “Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics”. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. 2019, pp. 193–204. arXiv: [1806.01838](#) (cit. on pp. [9](#), [10](#), [13](#), [40](#), [44](#), [47](#), [52](#)).

- [G'22] A. Glaudell. "Catalytic Embeddings: Linking Gate Teleportation to Circuit Synthesis". Talk at Quantum Information Processing. 2022 (cit. on p. 37).
- [GKS'76] V. Gorini, A. Kossakowski, and E. C. G. Sudarshan. "Completely positive dynamical semigroups of N -level systems". In: *Journal of Mathematical Physics* 17 (5 1976), pp. 821–825. doi: [10.1063/1.522979](https://doi.org/10.1063/1.522979) (cit. on p. 13).
- [G'96] L. K. Grover. "A fast quantum mechanical algorithm for database search". In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. 1996. doi: [10.1145/237814.237866](https://doi.org/10.1145/237814.237866) (cit. on pp. 2, 40).
- [GMF'21] N. Guo, K. Mitarai, and K. Fujii. "Nonlinear transformation of complex amplitudes via quantum singular value transformation". 2021. arXiv: [2107.10764](https://arxiv.org/abs/2107.10764) (cit. on pp. 41, 49, 50, 125, 126).
- [H'19] J. Haah. "Product Decomposition of Periodic Functions in Quantum Signal Processing". In: *Quantum* (2019). doi: [10.22331/q-2019-10-07-190](https://doi.org/10.22331/q-2019-10-07-190) (cit. on pp. 10, 45).
- [H'26] C. Halleck-Dubé. "The cohomological Weighted Fundamental Lemma for nonsplit groups". PhD thesis. University of California, Berkeley, 2026 (cit. on p. 27).
- [H+'21] T. Hammonds, C. Kothari, N. Luntzlar, S. J. Miller, J. Thorner, and H. Wieman. "The Explicit Sato-Tate Conjecture For Primes In Arithmetic Progressions". In: *International Journal of Number Theory* 17.8 (2021). doi: [10.1142/S179304212150069X](https://doi.org/10.1142/S179304212150069X) (cit. on p. 26).
- [HHL'08] A. W. Harrow, A. Hassidim, and S. Lloyd. "Quantum algorithm for solving linear systems of equations". 2008. arXiv: [0811.3171](https://arxiv.org/abs/0811.3171) (cit. on pp. 9, 40).
- [H'25] M. Hartz. "On von Neumann's inequality on the polydisc". In: *Mathematische Annalen* 391 (2025), pp. 5235–5264. doi: [10.1007/s00208-024-03040-2](https://doi.org/10.1007/s00208-024-03040-2) (cit. on p. 10).
- [H'70] W. K. Hastings. "Monte Carlo sampling methods using Markov chains and their applications". In: *Biometrika* 57 (1 1970), pp. 97–109. doi: [10.1093/biomet/57.1.97](https://doi.org/10.1093/biomet/57.1.97) (cit. on p. 13).
- [HH'25] Y. Huang and A. W. Harrow. "Quantum Entropy Thermalization". In: *Communications in Mathematical Physics* 406.6 (May 2025). issn: 1432-0916. doi: [10.1007/s00220-025-05268-8](https://doi.org/10.1007/s00220-025-05268-8) (cit. on pp. 61, 65, 86).
- [J'21] S. Jitomirskaya. "On point spectrum of critical almost Mathieu operators". In: *Advances in Mathematics* (2021). doi: [10.1016/j.aim.2021.107997](https://doi.org/10.1016/j.aim.2021.107997) (cit. on p. 98).
- [J'97] S. Johansson. "A description of quaternion algebras". 1997. url: <http://www.math.chalmers.se/%5C%7Eesj/forskning/structure.ps> (cit. on p. 24).

- [KKN'21] M. Kaluba, D. Kielak, and P. W. Nowak. “On property (T) for $\text{Aut}(F_n)$ and $\text{SL}_n(\mathbb{Z})$ ”. In: *Annals of Mathematics* 193.2 (2021), pp. 539–562. doi: [10.4007/annals.2021.193.2.3](https://doi.org/10.4007/annals.2021.193.2.3) (cit. on p. 93).
- [KNO'19] M. Kaluba, P. W. Nowak, and N. Ozawa. “ $\text{Aut}(F_5)$ has property (T)”. In: *Mathematische Annalen* 375 (2019), pp. 1169–1191. doi: [10.1007/s00208-019-01874-9](https://doi.org/10.1007/s00208-019-01874-9) (cit. on p. 93).
- [KB'16] M. J. Kastoryano and F. G. S. L. Brandao. “Quantum Gibbs samplers: The commuting case”. In: *Communications in Mathematical Physics* 344.3 (2016), pp. 915–957. doi: [10.1007/s00220-016-2641-8](https://doi.org/10.1007/s00220-016-2641-8) (cit. on p. 54).
- [KT'12] M. J. Kastoryano and K. Temme. “Quantum logarithmic Sobolev inequalities and rapid mixing”. In: *Journal of Mathematical Physics* 54 (2012), pp. 052202–052202. URL: <https://api.semanticscholar.org/CorpusID:119682319> (cit. on p. 57).
- [K'52] T. Kato. “Note on the least eigenvalue of the Hill equation”. In: *Quarterly of Applied Mathematics* 10.3 (1952), pp. 292–294. doi: [10.1090/qam/44706](https://doi.org/10.1090/qam/44706) (cit. on pp. 93, 96).
- [KLW'15] J. P. Keating, N. Linden, and H. J. Wells. “Spectra and Eigenstates of Spin Chain Hamiltonians”. In: *Communications in Mathematical Physics* 338.1 (May 2015), pp. 81–102. ISSN: 1432-0916. doi: [10.1007/s00220-015-2366-0](https://doi.org/10.1007/s00220-015-2366-0) (cit. on pp. 61, 86, 89, 91).
- [K+'25] A. D. King, A. Nocera, M. M. Rams, J. Dziarmaga, R. Wiersema, W. Bernoudy, J. Raymond, N. Kaushal, N. Heinsdorf, R. Harris, K. Boothby, F. Altomare, A. J. Berkley, M. Boschnak, K. Chern, H. Christiani, S. Cibere, J. Connor, M. H. Dehn, R. Deshpande, S. Ejtemaee, P. Farré, K. Hamer, E. Hoskinson, S. Huang, M. W. Johnson, S. Kortas, E. Ladizinsky, T. Lai, T. Lanting, R. Li, A. J. R. MacDonald, G. Marsden, C. C. McGeoch, R. Molavi, R. Neufeld, M. Norouzpour, T. Oh, J. Pasvolsky, P. Poitras, G. Poulin-Lamarre, T. Prescott, M. Reis, C. Rich, M. Samani, B. Sheldan, A. Smirnov, E. Sterpka, B. Trullas Clavera, N. Tsai, M. Volkmann, A. Whitar, J. D. Whittaker, W. Wilkinson, J. Yao, T. J. Yi, A. W. Sandvik, G. Alvarez, R. G. Melko, J. Carrasquilla, M. Franz, and M. H. Amin. “Beyond-classical computation in quantum simulation”. 2025. arXiv: [2403.00910v2](https://arxiv.org/abs/2403.00910v2) (cit. on p. 3).
- [K'97] A. Y. Kitaev. “Quantum computations: algorithms and error correction”. In: *Russian Mathematical Surveys* 52.6 (1997), pp. 1191–1249. doi: [10.1070/RM1997v052n06ABEH002155](https://doi.org/10.1070/RM1997v052n06ABEH002155) (cit. on p. 7).
- [K'03] A. Y. Kitaev. “Fault-tolerant quantum computation by anyons”. In: *Annals of Physics* 303.1 (2003), pp. 2–30. doi: [10.1016/S0003-4916\(02\)00018-0](https://doi.org/10.1016/S0003-4916(02)00018-0) (cit. on p. 63).

- [KSV'02] A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation*. Graduate Studies in Mathematics 47. American Mathematical Society, 2002 (cit. on p. 7).
- [KMM'12] V. Kliuchnikov, D. Maslov, and M. Mosca. “Fast and efficient exact synthesis of single qubit unitaries generated by Clifford and T gates”. In: *Quantum Information and Computation* 13.7&8 (2012). URL: <https://www.rintonpress.com/xxqic13/qic-13-78/0607-0630.pdf> (cit. on pp. 7, 23).
- [K⁺'24] J. Kochanowski, Á. M. Alhambra, A. Capel, and C. Rouzé. “Rapid thermalization of dissipative many-body dynamics of commuting Hamiltonians”. 2024. arXiv: [2404.16780](https://arxiv.org/abs/2404.16780) (cit. on p. 54).
- [KT'23] E. Kubischta and I. Teixeira. “Family of Quantum Codes with Exotic Transversal Gates”. In: *Physical Review Letters* 131 (2023). DOI: [10.1103/PhysRevLett.131.240601](https://doi.org/10.1103/PhysRevLett.131.240601) (cit. on p. 8).
- [K'23] G. Kuperberg. “Breaking the cubic barrier in the Solovay–Kitaev algorithm”. 2023. arXiv: [2306.13158](https://arxiv.org/abs/2306.13158) (cit. on p. 7).
- [L'11] F. Lemmermeyer. “Euclid’s Algorithm in quartic CM-fields”. 2011. arXiv: [1108.6215](https://arxiv.org/abs/1108.6215) (cit. on p. 117).
- [L'79] H. W. Lenstra Jr. “Vanishing sums of roots of unity”. In: *Proceedings Bicentennial Congress Wiskundig Genootschap*. Vol. 100/101. Math. Centre Tracts. Amsterdam: Mathematisch Centrum, 1979, pp. 249–268. URL: <https://pub.math.leidenuniv.nl/~lenstrahw/PUBLICATIONS/1979d/art.pdf> (cit. on p. 90).
- [L'83] H. W. Lenstra Jr. “Integer Programming with a Fixed Number of Variables”. In: *Mathematics of Operations Research* 8.4 (1983). URL: <https://people.csail.mit.edu/rrw/presentations/Lenstra81.pdf> (cit. on pp. 8, 23, 32, 34, 35).
- [LPW'17] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov Chains and Mixing Times*. 2nd ed. American Mathematical Society, 2017 (cit. on pp. 14, 129).
- [L'06] D. W. Lewis. “Quaternion algebras and the algebraic legacy of Hamilton’s quaternions”. In: *Irish Mathematical Society Bulletin* 57 (2006), pp. 41–64 (cit. on p. 24).
- [L'22] L. Lin. “Lecture Notes on Quantum Algorithms for Scientific Computation”. https://math.berkeley.edu/~linlin/qasc/qasc_notes.pdf. 2022 (cit. on pp. 43–45, 122).
- [L'76] G. Lindblad. “On the generators of quantum dynamical semigroups”. In: *Communications in Mathematical Physics* 48 (2 1976), pp. 119–130. DOI: [10.1007/BF01608499](https://doi.org/10.1007/BF01608499) (cit. on p. 13).

- [L'26] R. Lopez. “Picard Groups and Brauer Groups of Certain Stacks”. PhD thesis. University of California, Berkeley, 2026 (cit. on p. 27).
- [LC'16] G. H. Low and I. Chuang. “Optimal Hamiltonian Simulation by Quantum Signal Processing”. 2016. arXiv: [1606.02685](#) (cit. on pp. 10, 40).
- [LP'20] A. Lubotzky and O. Parzanchevski. “From Ramanujan graphs to Ramanujan complexes”. In: *Philosophical Transactions of the Royal Society A* 378.2163 (2020). doi: [10.1098/rsta.2018.0445](#) (cit. on pp. 24, 27).
- [LPS'86a] A. Lubotzky, R. Phillips, and P. Sarnak. “Explicit expanders and the Ramanujan conjectures”. In: *Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing*. 1986, pp. 240–246. doi: [10.1145/12130.12154](#) (cit. on pp. 8, 22).
- [LPS'86b] A. Lubotzky, R. Phillips, and P. Sarnak. “Hecke operators and distributing points on the sphere. I”. In: *Communications on Pure and Applied Mathematics* 39.S1 (1986), S149–S186. doi: [10.1002/cpa.3160390710](#) (cit. on pp. 8, 22).
- [LPS'87] A. Lubotzky, R. Phillips, and P. Sarnak. “Hecke operators and distributing points on S^2 . II”. In: *Communications on Pure and Applied Mathematics* 40.4 (1987), pp. 401–420. doi: [10.1002/cpa.3160400402](#) (cit. on pp. 8, 22).
- [LPS'88] A. Lubotzky, R. Phillips, and P. Sarnak. “Ramanujan graphs”. In: *Combinatorica* 8.3 (1988), pp. 261–277. doi: [10.1007/BF02126799](#) (cit. on pp. 8, 22, 24, 26).
- [M'80] Y. Manin. *Computable and Uncomputable*. 1980 (cit. on p. 1).
- [M⁺'21] J. M. Martyn, Z. M. Rossi, A. K. Tan, and I. L. Chuang. “A Grand Unification of Quantum Algorithms”. In: *PRX Quantum* 2 (2021). doi: [10.1103/PRXQuantum.2.040203](#) (cit. on pp. 10, 40).
- [MJ'17] C. A. Marx and S. Jitomirskaya. “Dynamics and spectral theory of quasi-periodic Schrödinger-type operators”. In: *Ergodic Theory and Dynamical Systems* 37.8 (2017), pp. 2353–2393. doi: [10.1017/etds.2016.16](#) (cit. on p. 101).
- [MBS'08] M. Merkli, G. P. Berman, and I. M. Sigal. “Dynamics of collective decoherence and thermalization”. In: *Annals of Physics* 323 (2008), pp. 3091–3112. URL: <https://api.semanticscholar.org/CorpusID:17424313> (cit. on p. 61).
- [M⁺'53] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. “Equation of State Calculations by Fast Computing Machines”. In: *The Journal of Chemical Physics* 21 (6 1953), pp. 1087–1092. doi: [10.1063/1.1699114](#) (cit. on p. 13).

- [M'13] O. Milatovic. "A spectral property of discrete Schrödinger operators with non-negative potentials". In: *Integral Equations and Operator Theory* 76 (2013), pp. 285–300. doi: [10.1007/s00020-013-2060-6](https://doi.org/10.1007/s00020-013-2060-6) (cit. on p. 96).
- [M'06] T. Miyake. *Modular Forms*. Springer, 2006 (cit. on p. 24).
- [M'01] D. W. Morris. "Introduction to arithmetic groups". 2001. arXiv: [math/0106063](https://arxiv.org/abs/math/0106063) (cit. on p. 24).
- [NC'10] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010. doi: [10.1017/CBO9780511976667](https://doi.org/10.1017/CBO9780511976667) (cit. on p. 22).
- [NWD'20] P. Niemann, R. Wille, and R. Drechsler. "Advanced exact synthesis of Clifford+T circuits". In: *Quantum Information Processing* 19.9 (2020), pp. 1–23. url: <https://link.springer.com/article/10.1007/s11128-020-02816-0> (cit. on p. 23).
- [N'91] A. Nilli. "On the second eigenvalue of a graph". In: *Discrete Mathematics* 91 (2 1991), pp. 207–210. doi: [10.1016/0012-365X\(91\)90112-F](https://doi.org/10.1016/0012-365X(91)90112-F) (cit. on p. 8).
- [O'14] R. O'Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014 (cit. on pp. 17, 97).
- [O'16] N. Ozawa. "Noncommutative real algebraic geometry of Kazhdan's property (T)". In: *Journal of the Institute of Mathematics of Jussieu* 15 (2016), pp. 85–90. doi: [10.1017/S1474748014000309](https://doi.org/10.1017/S1474748014000309) (cit. on p. 93).
- [O'22] N. Ozawa. "A substitute for Kazhdan's property (T) for universal non-lattices". In: (2022). arXiv: [2207.05272](https://arxiv.org/abs/2207.05272) (cit. on pp. 92, 96).
- [PS'18] O. Parzanchevski and P. Sarnak. "Super Golden Gates for PU(2)". In: *Advances in Mathematics* 327 (2018), pp. 869–901. doi: [10.1016/j.aim.2017.06.022](https://doi.org/10.1016/j.aim.2017.06.022) (cit. on pp. 7, 8, 23, 24, 26, 28, 30, 33, 114).
- [P'03] V. Paulsen. *Completely Bounded Maps and Operator Algebras*. Cambridge University Press, 2003 (cit. on p. 10).
- [P'83] A. Paz. "A simplified version of H.W. Lenstra's integer programming algorithm and some applications". 1983. url: <https://www2.eecs.berkeley.edu/Pubs/TechRpts/1983/CSD-83-116.pdf> (cit. on pp. 8, 23, 32, 34).
- [P+'19] E. Pednault, J. A. Gunnels, G. Nannicini, L. Horesh, and R. Wisnieff. "Leveraging Secondary Storage to Simulate Deep 54-qubit Sycamore Circuits". 2019. arXiv: [1910.09534](https://arxiv.org/abs/1910.09534) (cit. on p. 3).
- [P+'14] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien. "A variational eigenvalue solver on a photonic quantum processor". In: *Nature Communications* 5.4213 (2014). doi: [10.1038/ncomms5213](https://doi.org/10.1038/ncomms5213) (cit. on p. 3).

- [P'94] W. Pleśniak. “Extension and polynomial approximation of ultradifferentiable functions in $\mathbb{R}^N$ ”. In: *Bulletin de la Société Royale des Sciences de Liège* (1994) (cit. on pp. 119, 120).
- [RW'24] J. Rajakumar and J. D. Watson. “Gibbs Sampling gives Quantum Advantage at Constant Temperatures with $O(1)$ -Local Hamiltonians”. 2024. arXiv: 2408.01516 (cit. on pp. 54, 55, 63).
- [RS'24] A. Ramkumar and M. Soleimanifar. “Mixing time of quantum Gibbs sampling for random sparse Hamiltonians”. In: (2024). arXiv: 2411.04454 (cit. on p. 62).
- [RS'81] M. Reed and B. Simon. *I: Functional Analysis*. Vol. 1. Academic press, 1981 (cit. on p. 93).
- [R'69] T. Rivlin. *An Introduction to the Approximation of Functions*. Blaisdell, 1969 (cit. on p. 120).
- [RFŻ'10] W. Roga, M. Fannes, and K. Życzkowski. “Davies maps for qubits and qutrits”. In: *Reports on Mathematical Physics* 66 (3 2010), pp. 311–329. doi: 10.1016/S0034-4877(11)00003-6 (cit. on p. 13).
- [RS'16] N. J. Ross and P. Selinger. “Optimal ancilla-free Clifford+T approximation of z-rotations”. In: *Quantum Information and Computation* 16.11&12 (2016). URL: <https://www.rintonpress.com/xxqic16/qic-16-1112/0901-0953.pdf> (cit. on pp. 7, 23, 31, 33, 114).
- [RCC'25] Z. M. Rossi, J. L. Ceroni, and I. L. Chuang. “Modular quantum signal processing in many variables”. In: *Quantum* 9 (2025). doi: 10.22331/q-2025-06-18-1776 (cit. on pp. 11, 12).
- [RC'22] Z. M. Rossi and I. L. Chuang. “Multivariable quantum signal processing (M-QSP): prophecies of the two-headed oracle”. In: *Quantum* (2022). doi: 10.22331/q-2022-09-20-811 (cit. on pp. 12, 40, 41).
- [RFA'25] C. Rouzé, D. S. França, and Á. M. Alhambra. “Efficient thermalization and universal quantum computing with quantum Gibbs samplers”. In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. 2025, pp. 1488–1495. doi: 10.1145/3717823.371826 (cit. on pp. 54, 69).
- [S'15a] N. T. Sardari. “Diameter of Ramanujan graphs and random Cayley graphs with numerics”. 2015. arXiv: 1511.09340 (cit. on p. 27).
- [S'19] N. T. Sardari. “Complexity of Strong Approximation on the Sphere”. In: *International Mathematics Research Notices* (2019). doi: 10.1093/imrn/rnz233 (cit. on pp. 29, 35).
- [S'15b] P. Sarnak. “Letter to Aaronson and Pollington on the Solovay–Kitaev Theorem and Golden Gates”. 2015. URL: <https://publications.ias.edu/node/2637> (cit. on pp. 7, 8, 22, 23, 27, 28).

- [SR'18] P. Sarnak and M. T. Rassias. "Interview with Peter Sarnak". In: *EMS Newsletter* (2018). Interview conducted by Michael Th. Rassias, pp. 16–19. URL: <https://ems.press/content/serial-article-files/10192> (cit. on p. 2).
- [S'07] M. Schlosshauer. *Decoherence and the Quantum-To-Classical Transition*. The Frontiers Collection. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. xv, 418. ISBN: 978-3-540-35773-5. DOI: [10.1007/978-3-540-35775-9](https://doi.org/10.1007/978-3-540-35775-9) (cit. on p. 63).
- [S'02] Y. Shi. "Both Toffoli and Controlled-NOT need little help to do universal quantum computation". 2002. arXiv: [quant-ph/0205115](https://arxiv.org/abs/quant-ph/0205115) (cit. on p. 43).
- [S'94] P. W. Shor. "Algorithms for quantum computation: discrete logarithms and factoring". In: *IEEE 35th Annual Symposium on Foundations of Computer Science*. IEEE, 1994. DOI: [10.1109/SFCS.1994.365700](https://doi.org/10.1109/SFCS.1994.365700) (cit. on p. 2).
- [S'95] P. W. Shor. "Scheme for reducing decoherence in quantum computer memory". In: *Physical Review A* 52.4 (1995). DOI: [10.1103/PhysRevA.52.R2493](https://doi.org/10.1103/PhysRevA.52.R2493) (cit. on p. 8).
- [S'96] A. M. Steane. "Error Correcting Codes in Quantum Theory". In: *Physical Review Letters* 77 (1996). DOI: [10.1103/PhysRevLett.77.793](https://doi.org/10.1103/PhysRevLett.77.793) (cit. on p. 8).
- [S'21] Z. Stier. "Short Paths in PU(2)". In: *Quantum Information and Computation* 21.9&10 (2021). DOI: [10.26421/QIC21.9-10-3](https://doi.org/10.26421/QIC21.9-10-3) (cit. on pp. 8, 23, 24, 30, 33, 35).
- [S'00] A. Strömbergsson. "An application of an explicit trace formula to a well-known spectral correspondence on quaternion groups". 2000. URL: <https://www2.math.uu.se/~ast10761/papers/JLny2.ps> (cit. on p. 26).
- [T+'20] S. Takahira, A. Ohashi, T. Sogabe, and T. S. Usuda. "Quantum algorithm for matrix functions by Cauchy's integral formula". In: *Quantum Information & Computation* (2020). arXiv: [2106.08075](https://arxiv.org/abs/2106.08075) (cit. on p. 41).
- [T+'21] S. Takahira, A. Ohashi, T. Sogabe, and T. S. Usuda. "Quantum Algorithms based on the Block-Encoding Framework for Matrix Functions by Contour Integrals". 2021. arXiv: [2106.08076](https://arxiv.org/abs/2106.08076) (cit. on pp. 41, 122).
- [T'13] K. Temme. "Lower bounds to the spectral gap of Davies generators". In: *Journal of Mathematical Physics* 54.12 (2013). DOI: [10.1063/1.4850896](https://doi.org/10.1063/1.4850896) (cit. on pp. 14, 54, 57–59, 62, 63, 66, 127).
- [T'14] K. Temme. "Thermalization Time Bounds for Pauli Stabilizer Hamiltonians". In: *Communications in Mathematical Physics* 350 (2014), pp. 603–637. URL: <https://api.semanticscholar.org/CorpusID:118938674> (cit. on p. 54).
- [VG'18] J. Van Apeldoorn and A. Gilyén. "Improvements in quantum SDP-solving with applications". In: (2018). arXiv: [1804.05058](https://arxiv.org/abs/1804.05058) (cit. on p. 13).

- [V⁺'17] J. Van Apeldoorn, A. Gilyén, S. Gribling, and R. de Wolf. “Quantum SDP-solvers: Better upper and lower bounds”. In: *IEEE 58th Annual Symposium on Foundations of Computer Science*. IEEE. 2017, pp. 403–414 (cit. on p. 13).
- [V'80] M.-F. Vignéras. *Arithmétique des algèbres de quaternions*. Vol. 800. Lecture Notes in Mathematics. Springer, 1980 (cit. on p. 24).
- [V'21] J. Voight. *Quaternion Algebras*. Springer, 2021 (cit. on pp. 24–26).
- [W'24] V. Winstein. “Approximation, Inference, and Sampling of Gibbs Measures”. 2024. URL: https://vilas.us/mathnotes/berkeleytalks/StudentProbability_GibbsMeasures.pdf (cit. on p. 12).
- [W'25] V. Winstein. “Simulating and understanding phase change”. 2025. URL: <https://www.youtube.com/watch?v=itRV2jEtV8Q> (cit. on p. 12).
- [WZ'82] W. K. Wootters and W. H. Zurek. “A single quantum cannot be cloned”. In: *Nature* 299 (1982), pp. 802–803. doi: [10.1038/299802a0](https://doi.org/10.1038/299802a0) (cit. on p. 3).
- [YLC'14] T. J. Yoder, G. H. Low, and I. L. Chuang. “Fixed-point quantum search with an optimal number of queries”. In: *Physical Review Letters* 113 (2014). doi: [10.1103/PhysRevLett.113.210501](https://doi.org/10.1103/PhysRevLett.113.210501) (cit. on p. 13).
- [Z'02] W. H. Zurek. “Decoherence and the transition from quantum to classical-revisited”. In: *Los Alamos Science* 27 (2002), pp. 86–109. URL: <https://www2.unicamp.br/~chibeni/textosdidaticos/zurek-1991.pdf> (cit. on p. 63).
- [Z'03] W. H. Zurek. “Decoherence, einselection, and the quantum origins of the classical”. In: *Reviews of Modern Physics* 75.3 (2003). doi: [10.1103/RevModPhys.75.715](https://doi.org/10.1103/RevModPhys.75.715) (cit. on p. 63).
- [Z'12] M. Zworski. *Semiclassical Analysis*. Vol. 138. Graduate Studies in Mathematics. American Mathematical Society, 2012 (cit. on p. 15).