ABSTRACT

Title of Dissertation: THE COMPLEXITY OF SIMULATING QUANTUM PHYSICS: DYNAMICS AND EQUILIBRIUM

Abhinav Deshpande
Doctor of Philosophy, 2021

Dissertation Directed by: Professor Alexey V. Gorshkov
Department of Physics
University of Maryland

Professor Bill Fefferman
Department of Computer Science
University of Chicago

Quantum computing is the offspring of quantum mechanics and computer science, two great scientific fields founded in the 20th century. Quantum computing is a relatively young field and is recognized as having the potential to revolutionize science and technology in the coming century. The primary question in this field is essentially to ask which problems are feasible with potential quantum computers and which are not. In this dissertation, we study this question with a physical bent of mind. We apply tools from computer science and mathematical physics to study the complexity of simulating quantum systems. In general, our goal is to identify parameter regimes under which simulating quantum systems is *easy* (efficiently solvable) or *hard* (not efficiently solvable). This study leads to an understanding of the features that make certain problems easy or hard to solve. We also get physical insight into the behavior of the system being simulated.

In the first part of this dissertation, we study the classical complexity of simulating quantum dynamics. In general, the systems we study transition from being easy to simulate at short times to being harder to simulate at later times. We argue that the transition timescale is a useful measure for various Hamiltonians and is indicative of the physics behind the change in complexity. We illustrate this idea for a specific bosonic system, obtaining a complexity phase diagram that delineates the system into easy or hard for simulation. We also prove that the phase diagram is robust, supporting our statement that the phase diagram is indicative of the underlying physics.

In the next part, we study open quantum systems from the point of view of their potential to encode hard computational problems. We study a class of fermionic Hamiltonians subject to Markovian noise described by Lindblad jump operators and illustrate how, sometimes, certain Lindblad operators can induce computational complexity into the problem. Specifically, we show that these operators can implement entangling gates, which can be used for universal quantum computation. We also study a system of bosons with Gaussian initial states subject to photon loss and detected using photon-number-resolving measurements. We show that such systems can remain hard to simulate exactly and retain a relic of the "quantumness" present in the lossless system.

Finally, in the last part of this dissertation, we study the complexity of simulating a class of equilibrium states, namely ground states. We give complexity-theoretic evidence to identify two structural properties that can make ground states easier to simulate. These are the existence of a spectral gap and the existence of a classical description of the ground state. Our findings complement and guide efforts in the search for efficient algorithms.

THE COMPLEXITY OF SIMULATING QUANTUM PHYSICS: DYNAMICS AND EQUILIBRIUM

by

Abhinav Deshpande

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2021

Advisory Committee:
Professor Mohammad Hafezi, Chair
Professor Alexey V. Gorshkov, Co-chair/Advisor
Professor Bill Fefferman, Co-advisor
Professor Andrew M. Childs, Dean's Representative
Professor Christopher Monroe

Dedication

To my family, ಅಮ್ಮ (Amma), ಪಪ್ಪಾ (Pappa), ಅಜ್ಜಿ (Ajji), ಅಕ್ಕ (Akka), and ಪೂರ್ವಿ (Poorvi).

Acknowledgments

I would like to first of all thank my advisor Alexey Gorshkov for his constant support throughout my Ph.D. life. I truly enjoy working with him and have learned a lot from him over these years. His infectious enthusiasm and the constant flow of new ideas have always been motivating.

Bill Fefferman, my co-advisor, has also been unwavering in his support throughout my Ph.D. life. His enormous patience when helping me learn is admirable. I have also been fortunate to learn through example his philosophy on writing and presentations, something I wish to emulate one day.

Thanks also go to other professors with whom I have had the fortune to collaborate, such as members of my committee Andrew Childs, Mohammad Hafezi, and Christopher Monroe. I have also benefited immensely from the research environment at QuICS, JQI and CMTC thanks to Brian Swingle, Gorjan Alagic, Jacob Taylor, Michael Gullans, Paul Julienne, Stephen Jordan, and Yi-Kai Liu, and benefited from a similar environment at IQIM during my visit there thanks to Fernando Brandão, John Preskill, Thomas Vidick and the graduate students and postdocs there. I am very grateful to my teachers at UMD such as Andrew Childs, Aravind Srinivasan, Ian Appelbaum, Ian Spielman, Paulo Bedaque, and Ted Jacobson, whose dedication toward teaching is readily apparent in their lectures. I am also indebted to postdocs Aarthi Sundaram, Cedric Lin, Chris Baldwin, Dominik Hangleiter, Jim Garrison, Michael Foss-Feig, Oles Shtanko, Paraj Titum, Przemek Bienias, Shelby Kimmel, and

Zhexuan Gong, and graduate students such as Aaron Ostrander, Ali Hamed Moosavian, Ali Lavasani, Alireza Seif, Andrew Glaudell, Arthur Mehta, Chiao-Hsuan Wang, Eddie Schoute, Gina Quan, Nishad Maskara, Steve Ragole, Tongyang Li, Troy Sewell, Wen-Lin Tan, Yiming Huang, and Yuan Su for being great friends and teaching me many things along the way. I am also grateful to the members of the Gorshkov group for the camaraderie: Adam Ehrenberg, Andrew Guo, Ani Bapat, Dhruv Devulapalli, Fangli Liu, Igor Boettcher, Jacob Bringewatt, Jeremy Young, Joe Iosue, Lucas Brady, Luispe García-Pintos, Minh Tran, Pradeep Niroula, Seth Whitsitt, Su-Kuan Chu, Yidan Wang, and Zach Eldredge.

Life at UMD was very smooth thanks to the tireless efforts of Andrea Svejda, Donna Hammer, Javiera Caceres, Jessica Crosby, Josiland Chambers, Kelly Phillips, Melissa Britton and the numerous other staff at UMD who kept it running. I would like to single out Josiland for her commitment and willingness to always help with anything.

No amount of thanks suffices to express my gratitude to my friends who have been there throughout: my roommates and circle of friends Ankan Bansal, Ankit Gargava, Atul Singh, Faez Ahmed, Kiran Burra, Lovlesh Kaushik, Sarthak Chandra, Saurabh Saxena, Shashank Ganesh, Wei Chen, and Yidan Wang, and my friends from undergrad who've also been in graduate school Abhishek Gupta, Akshansh Singh, Anupam Kumar, Reevu Maity, Shamreen Iram, and Sumit Sinha.

I would also like to acknowledge the schooling system in place in India, my school's library where I first encountered and fell in love with physics, and the encouragement given to me by my teachers at every level. Institutions and programs such as the National Initiative for Undergraduate Science and the olympiads at the Homi Bhabha Centre for Science Education, the Visvesvaraya Industrial & Technological Museum, and the Kishore Vaigyanik Protsahan

Yojana of the Department of Science and Technology furthered my interest in science and helped me in my decision to study physics as a major. My alma mater the Indian Institute of Technology Kanpur gave me a lifetime of experience and learning that I will always cherish.

Finally, my family is the sole reason I am here today. This dissertation is dedicated to them. My parents have made innumerable sacrifices to their own well-being in order to prioritize my sister's and my education. My grandmother and sister have been as important as my parents at shaping me and teaching me to be a good person. I also thank my aunts, cousins, brother-in-law, and other relatives and well-wishers for their support and encouragement. Special thanks and adulation to my one-year-old niece Poorvi, whose presence has lent added color and melody to life.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

AA	Aaronson and Arkhipov
BPP	Bounded-error Probabilistic Polynomial time
BQP	Bounded-error Quantum Polynomial time
CH	Counting Hierarchy
DMRG	Density Matrix Renormalization Group
EC1	Easy Class 1
EC2	Easy Class 2
EC3	Easy Class 3
GBS	Gaussian Boson Sampling
HHKL	Haah-Hastings-Kothari-Low
HOG	Heavy Output Generation
IQP	Instantaneous Quantum Polynomial time
MPS	Matrix Product States
NP	Nondeterministic Polynomial Time
P	Polynomial Time
#P	Sharp P
PGQMA	Polynomially Gapped QMA
PH	Polynomial Hierarchy
PMP	Perfect Matching Permutation
PNR	Photon Number Resolving
PP	Probabilistic Polynomial Time
QCMA	Quantum Classical Merlin Arthur
QCS	Quantum computational supremacy
QMA	Quantum Merlin Arthur
RCS	Random circuit sampling
VQE	Variational Quantum Eigensolver

Citations to Previously Published Work

Most of the work appearing in this thesis is either published or has appeared on the preprint server arXiv. We mention these here.

- Chapter 2: "Dynamical Phase Transitions in Sampling Complexity", A. Deshpande, B. Fefferman, M. C. Tran, M. Foss-Feig, and A. V. Gorshkov, Phys. Rev. Lett. 121, 030501 (2018).
  - Further work related to this problem appeared in "Hierarchy of Linear Light Cones with Long-Range Interactions", M. C. Tran, C.-F. Chen, A. Ehrenberg, A. Y. Guo, A. Deshpande, Y. Hong, Z.-X. Gong, A. V. Gorshkov, and A. Lucas, Phys. Rev. X 10, 031009 (2020).
- Chapter 3: "Complexity Phase Diagram for Interacting and Long-Range Bosonic Hamiltonians", N. Maskara, A. Deshpande, A. Ehrenberg, M. C. Tran, B. Fefferman, and A. V. Gorshkov, arXiv:1906.04178 (2019).
- Chapter 4: "Limits on Classical Simulation of Free Fermions with Dissipation", O. Shtanko, A. Deshpande, P. S. Julienne, and A. V. Gorshkov, arXiv:2005.10840 (2020).
- Chapter 5: "Quantum Computational Supremacy via High-Dimensional Gaussian Boson Sampling", A. Deshpande, A. Mehta, T. Vincent, N. Quesada, M. Hinsche, M. Ioannou, L. Madsen, J. Lavoie, H. Qi, J. Eisert, D. Hangleiter, B. Fefferman, and I. Dhand, arXiv:2102.12474 (2021).
  Other work relating to quantum advantage in experimental contexts has appeared in
  - "Interference of Temporally Distinguishable Photons Using Frequency-Resolved Detection", V. V. Orre, E. A. Goldschmidt, A. Deshpande, A. V. Gorshkov, V. Tamma, M. Hafezi, and S. Mittal, Phys. Rev. Lett. 123, 123603 (2019).
  - "Quantum Approximate Optimization of the Long-Range Ising Model with a Trapped-Ion Quantum Simulator", G. Pagano, A. Bapat, P. Becker, K. S. Collins, A. De, P. W. Hess, H. B. Kaplan, A. Kyprianidis, W. L. Tan, C. Baldwin, L. T. Brady, A. Deshpande, F. Liu, S. Jordan, A. V. Gorshkov, and C. Monroe, Proc. Natl. Acad. Sci. USA 202006373 (2020).
- Chapter 6: "The Importance of the Spectral Gap in Estimating Ground-State Energies", A. Deshpande, A. V. Gorshkov, and B. Fefferman, arXiv:2007.11582 (2020).

Chapter 1: Introduction

> "The theory of computation has traditionally been studied almost entirely in the abstract, as a topic in pure mathematics. This is to miss the point of it. Computers are physical objects, and computations are physical processes. What computers can or cannot compute is determined by the laws of physics alone, and not by pure mathematics."
>
> — David Deutsch [1]

Physics, to slightly overgeneralize, is humanity's quest to understand the natural world. Numerous times in the history of this quest, ideas and techniques from disparate fields in the natural sciences have found their way into the physicist's toolkit because of their utility. Conversely, important mathematical tools have been invented primarily in order to address problems arising in physical contexts. This cross-fertilization of ideas between physics and the other sciences is in of itself a valuable cultural activity, notwithstanding its scientific, technological, and ultimately economic value.

Similarly overgeneralizing, computer science is the quest to understand what is (efficiently) computable in our world. A cornerstone of theoretical computer science is the extended Church-Turing thesis. This is the observation that various models of computation modeled on the physical world seem able to simulate each other with small overhead,

and implies that a problem efficiently solvable according to one computational model is also efficiently solvable according to any other model. This observation enables computer scientists to reason about the efficiency of computation by abstracting away the details of the computational model at hand. Ultimately, the thesis also lends meaning to the question of what is efficiently computable or not. This is because if the notion of efficiently computable depended heavily on the computational model, it would cease to be a universal one.

The goals of physics and computer science are very much aligned. The question of what can be efficiently computed in our world depends on what kinds of devices are physically possible to realize in our world, which ultimately must obey the laws of physics. Therefore, this question is as much a question in physics as it is one of interest to computer science, as is elegantly argued for in Refs. [1; 2], for example.

When quantum mechanics enters this picture, the relation between physics and computer science becomes much more interesting. Despite quantum mechanics already having been formalized by 1930, it was not realized until the 1980s that quantum-mechanical computers could potentially be faster than classical (i.e., ordinary) computers at certain tasks, such as simulating physics [3]. This realization stemmed from the observation that tasks such as simulating quantum systems seem hard on classical computers. Therefore, the existence of quantum computing changes the meaning of what is efficiently computable in our physical world, which is an exciting development for both computer science and physics. We elaborate on this point of view in the coming sections.

1.1 Why Complexity theory

One branch of computer science that acquires new life on account of quantum computing is complexity theory. This is the branch of theoretical computer science that deals with quantifying the resources required to solve various problems and classifying problems into various complexity classes. The efficiency of an algorithm is measured in terms of the asymptotic scaling of the algorithm's runtime as a function of the size of the problem. It is of interest to determine the optimal algorithm for solving a particular problem, i.e. the algorithm with minimum asymptotic scaling of the runtime. However, the runtime for a problem can depend on the model of computation used by virtue of the cost associated with elementary operations. This issue is addressed by the extended Church-Turing thesis, which asserts that differing models of computation can simulate each other with an overhead at most polynomial in the problem size. Due to this reason and the fact that polynomials are closed under composition, we take a problem to be efficiently solvable if the optimal algorithm has runtime polynomial in the problem size (under any model of computation).

Classifying the complexity of problems becomes even more interesting in the context of quantum computing, since quantum computing can change the meaning of what is efficiently computable in our world. This is because of the possibility that quantum computers are not efficiently simulable by classical computers, which would falsify the extended Church-Turing thesis in the quantum setting, meaning that tasks efficiently solvable by quantum computers are not necessarily efficiently solvable on classical computers. A major undertaking in the field of quantum computer science is thus to precisely characterize the class of problems efficiently solvable by quantum computers. We turn to the potential benefits such an undertaking can

have.

First, as has been emphasized above, understanding the power and limitations of quantum computers is tantamount to understanding the fundamental limitations of computing according to the laws of physics. This is because, according to our best guess for how the laws of physics would work, quantum field theories and as-yet-undiscovered theories governing quantum gravity might be efficiently simulable with quantum computers[1] [4], which can be regarded as a statement of the *quantum* extended Church-Turing thesis. Apart from being an extremely interesting physics question in itself, it also gives rise to other questions in physics. An example is that of black hole physics, where a complexity-theoretic point of view has led to several creative ideas and proposals in the field, see for instance Refs. [5–8].

Second, this quest can lead to the identification of new, practical problems that quantum computers can efficiently solve. Indeed, the whole premise of the quest is that the set of problems efficiently solvable on quantum computers (BQP) is potentially larger than that solvable on randomized classical computers (BPP)[2]. In order to understand the power of BQP, it is necessary to identify these problems. It is possible that at least a few of these are of practical use, paving the way for applications of quantum computers.

Third, another practical benefit from understanding the boundaries of BQP is knowing the limitations of quantum computers. As a "meta-benefit", knowing the limitations of quantum computing saves time for practitioners and quantum algorithms researchers. Other applications are in designing protocols for classical communication and verification of quantum computing that are secure against quantum-capable adversaries [9].

[1] Meaning that a hypothetical "quantum gravity computer" is possibly no more powerful than a quantum computer.

[2] For an overview of some of the complexity classes discussed in this dissertation, refer to Appendix B.

This dissertation attempts at capturing some of the rich physics coming out of examining quantum systems from a complexity-theoretic point of view. Specifically, we will consider the problem of simulating quantum systems in the sense of both dynamics and equilibrium. The complexity of solving this problem depends on features of the quantum system and can be classified into various classes. Studying and classifying the dependence of the complexity on system parameters often yields a "complexity phase diagram", which we elaborate upon later in the dissertation. These phase diagrams are also associated with complexity phase *transitions*, illustrative of the physics underlying the quantum system at hand.

1.2 Sampling Complexity

When we study the classical complexity of simulating quantum systems, it is prudent to pay attention to the precise notion of the task being considered. First off, we note that we cannot demand that the classical computer output an explicit representation of the quantum state of the system, since it is infeasible even to write out 2^n complex amplitudes of an n-qubit system for large $n \gtrsim 60$. Moreover, for large system sizes, an experiment cannot feasibly access any particular amplitude or probability, let alone 2^n of them. Therefore, we should phrase the problem in a way that captures what an experimentalist can access. An alternative is to demand that a classical computer be able to efficiently compute (to a suitable accuracy) all k-body correlation functions (k-local observables) for any constant k. This task can be efficiently solved by running the experiment. However, this does not capture everything about a quantum experiment, as we are about to elucidate. Such correlation functions cannot

capture correlations present across all n particles of a system that are vital to the physics.

One way to put the classical simulator and a real quantum experiment on the same footing is to notice that every such quantum experiment is inherently outputting a sample from an underlying distribution describing the probabilities of the 2^n different outcomes. We also allow for the classical algorithm to make some slight error in the distribution it is sampling from, measured in terms of the total variation distance between distributions[3]. This defines the task of *approximate sampling*. Approximate sampling is a nice way of characterizing the task of simulating quantum systems because it satisfies both desiderata of being not too powerful and being able to capture all that an experiment can capture. This is because the total variation distance characterizes the ability of a referee to distinguish between two distributions. Therefore, a small total variation distance between the simulated and the actual distributions implies that the two cannot be reliably distinguished, constituting a successful simulation. Moreover, approximate sampling from the output distribution when measuring in a certain basis also enables one to compute local observables diagonal in that basis[4]. Therefore, it is appropriate to study the classical complexity of approximate sampling[5].

Fortunately, for the notion of sampling from a distribution, there is some evidence that classical computers cannot efficiently simulate certain quantum systems that are otherwise simulable on quantum computers. This statement constitutes evidence against the extended Church-Turing thesis. For the case of sampling exactly from the target distribution, early work by Terhal and DiVincenzo [11] gave some complexity-theoretic evidence

[3]For two distributions $\mathcal{D}_1$ and $\mathcal{D}_2$ over a space of outcomes $\mathcal{X}$ with associated probabilities $p_1(x)$ and $p_2(x)$, the total variation distance is given by $\sum_{x\in\mathcal{X}} |p_1(x) - p_2(x)|$.

[4]In fact, sampling from a distribution is exactly what many experiments are set up to do, and local observables are computed a posteriori from the observed samples.

[5]There is, however, one significant drawback to the total variation distance measure. *Verifying* that two distributions are close under this measure can take exponentially many samples [10].

for the hardness of this task. Subsequent work by Bremner *et al.* [12] on *Instantaneous Quantum Polynomial time (IQP) circuits* and by Aaronson and Arkhipov [13; 14] on *boson sampling* strengthened this complexity-theoretic evidence. They showed that the existence of a classical algorithm for exact sampling implies the falsity of a widely-believed conjecture in complexity theory, namely the non-collapse of the so-called Polynomial Hierarchy (PH). The non-collapse of the PH is a conjecture that generalizes the $\mathsf{P} \neq \mathsf{NP}$ conjecture. Briefly speaking, the polynomial hierarchy is defined to be an infinite tower of complexity classes generalizing NP, with "levels" starting from P, NP, and so on, each level more complex than the previous. The conjecture asserts that the tower is truly infinite, in that no two levels are equivalent to each other. Aaronson and Arkhipov also proved, modulo a few other conjectures, that a classical algorithm to *approximately* sample from certain quantumly sampleable distributions would imply the collapse of the PH.

At a high level, in order to prove the hardness of exact sampling, it suffices to show that computing any single output probability to multiplicative precision[6] is #P-hard. If there is a classical sampling algorithm, then one can, using the tool of Stockmeyer counting [15; 16], compute a multiplicative approximation to the output probability within the third level of the polynomial hierarchy. This thereby solves any $\mathsf{P}^{\#\mathsf{P}}$ problem[7] using access to resources in the third level of the polynomial hierarchy. The contradiction then follows from the fact that the entirety of the polynomial hierarchy (i.e. each level of the PH) is contained in $\mathsf{P}^{\#\mathsf{P}}$ [17], which causes the PH to collapse.

The proof of Aaronson and Arkhipov for the hardness of *approximate* sampling is along

[6] Refer to Appendix A for an overview of these mathematical definitions.

[7] For two complexity classes A and O, the notation A^{O} refers to the class of problems solvable in the class A with access to an oracle for any problem in the class O. An oracle for a class correctly solves problems in that class in one time step.

the same lines but involves new ideas. They show that hardness of approximate sampling follows from the #P-hardness of approximately computing a *random* output probability to a certain additive error. In other words, showing the average-case hardness of estimating output probabilities (as opposed to worst-case hardness in the previous case) is enough. This step necessitates considering a family of random instances, which serves to define the ensemble of output probabilities. The additional conjectures made by Aaronson and Arkhipov serve to give evidence for this statement on average-case hardness. The first conjecture, namely anti-concentration, asserts that a randomly selected output probability is not too small compared to the average output probability. The second conjecture states that there is an equivalence between average-case and worst-case hardness of computing to multiplicative error the output probability of the class of quantum circuits. These conjectures have since been ported to other contexts and are by now standard assumptions in the so-called "quantum computational supremacy" literature [18–24]. While anticoncentration has been proven in some contexts (but not yet for boson sampling) [19; 25–29], the average-case hardness conjecture has remained unproven, although recent work has come tantalizingly close to a proof [30–32].

In order to conclude that the extended Church-Turing thesis is violated, we need to demonstrate that it *is* possible to quantumly sample from the distributions considered above. Assuming that quantum mechanics as we know it, with its exponentially large Hilbert space and the possibility for quantum error correction, is valid, we can theoretically infer that the extended Church-Turing thesis is violated. However, a skeptic may (justifiably) raise the issue that quantum mechanics is not sufficiently validated in the so-called "high complexity" regime, and that the exponentiality of Hilbert space has not been proven to

be *necessary*. Therefore, only by performing an actual experiment and demonstrating that quantum mechanics allows one to sample from the distributions mentioned above can we satisfactorily claim to have refuted the extended Church-Turing thesis. This is what recent notable experiments have claimed to do [33; 34]. The situation is very reminiscent of Bell tests, which aim to give experimental evidence for quantum nonlocality and only recently were performed in a loophole-free manner.

We now turn to noise, an ever-present feature of any quantum experiment. In order to implement an experiment that achieves a small total variation distance as the system size increases, one should either depend on quantum error correction or ensure that the noise strength *decreases* with system size [35], an unrealistic assumption. Recent developments in the field have nevertheless attempted to demonstrate a quantum advantage over classical computers via different means. One route is to move away from the task of approximate sampling to the task of scoring high on a test intended to capture the similarity of the sampled distribution and the original distribution via the so-called (linear) cross-entropy measure in the context of random circuit sampling [33; 36]. Outputting samples that achieve high scores on this test is conjectured to be classically hard [37; 38], although it can be potentially simpler than the task of approximate sampling [39]. There is another nice feature of the linear cross-entropy measure. Under the (strong) assumption that the noise model is given exactly by global depolarization, the linear cross-entropy measures the fidelity of the experiment [36]. Under the same assumption, sampling exactly from the noisy distribution with any nonzero fidelity is also classically hard [33].

1.3 Dynamics

We now focus on the problem of simulating quantum dynamics. By the quantum extended Church-Turing thesis, efficiently simulating quantum dynamics seen in nature should always be possible on a universal quantum computer. This is indeed true by virtue of impressive advances in quantum algorithms for Hamiltonian simulation over the last decade [40–48]. Therefore, studying the classical complexity of this problem is perhaps a more interesting venture if one is only concerned with the existence of (in)efficiency and not its degree.

In this dissertation, we study the classical complexity of simulating quantum dynamics for some quantum systems and investigate when the problem is either solvable in time polynomial in the problem size (*easy*) or not (*hard*). Studying the boundary between easy and hard with regard to the problem of simulating quantum dynamics gives a characterization of when the problem is classical or has non-classical, quantum features. In other words, the classical simulability of a quantum system is a measure of non-classicality, or "quantumness".

Several quantum experiments, at their heart, involve preparing a fixed initial quantum state $|\psi_0\rangle$, evolving it under a potentially time-dependent system Hamiltonian H for time t, and making measurements on the final state $|\psi_t\rangle = \mathcal{T}\exp\left\{[-i\int_0^t \mathrm{d}s H(s)]\right\}|\psi_0\rangle$. The initial state is often simple to prepare, which we take to be a computational basis state for simplicity. Similarly, the measurement at the end is also a simple-to-implement one, and we take the measurement to be in the computational basis. These choices only mildly affect the analysis of the complexity, if at all, since other initial states and choices of measurement

bases can be accounted for with small-depth unitaries[8].

With this setup, we consider the question of how hard it is to approximately sample from the state[9] $|\psi_t\rangle$ as a function of time t, which may be a function of the system size n. Depending on the quantum system and Hamiltonian, for long enough times the dynamics can prepare states that are classically hard to sample from, assuming the conjectures in the previous section. Since, as we have defined it, the notions of "easy" and "hard" are mutually exclusive, at each time $t(n)$ the system is either easy or hard to simulate. Further, since the system evolves from an easy-to-sample-from state to a hard-to-sample-from one, it must have transitioned from easy to hard at some timescale $t_*(n)$. This timescale $t_*(n)$ is what we call the *transition timescale*, and we call the transition between easy and hard a *dynamical phase transition* in sampling complexity.

We propose using the transition timescale as a measure of how "quantum" or how "complex" a class of Hamiltonians is. This makes sense because different Hamiltonians may have different timescales for the transition from easy to hard. This transition timescale may be compared to other indicative timescales of quantum dynamics such as the Ehrenfest time, scrambling time, time to reach maximal entanglement, time to reach maximal circuit complexity, and equilibration time. The transition timescale defined here may differ or coincide with some of these other timescales.

Is the phase transition between classically easy and classically hard cases a "physical" one? In other words, is there a physically interesting characterization of the timescale or some

[8] In restricted models of quantum computation, these choices can and do affect the complexity. An example is that of IQP circuits [12], where there is a strong basis-dependence in the classical complexity. When this happens, we take care to clarify the set of initial states and measurement bases our result applies to.

[9] We have in mind an infinite *family* of initial states and Hamiltonians defined for an infinite number of system sizes.

physics underlying the transition? We seek to answer these questions in this dissertation by taking as example a system of bosonic particles evolving under lattice Hamiltonians. In Chapter 2, we study the transition timescale as a function of system parameters for a system of free bosons hopping on a lattice. We also study the case where the Hamiltonian is promised to be Anderson-localized [49], and show that the system stays classically simulable for all times, or in other words, the transition timescale diverges. In Chapter 3, we continue this study by considering interactions and long-range hopping. For a wide class of interactions, we see that there is almost no dependence of the transition timescale on the presence of interactions or their strength. This illustrates that the complexity phase transition is robust to perturbations and has physical meaning.

The examples we have seen here do not explicitly model possible experimental noise; so far, we have only accounted for noise in terms of the total variation distance error between the target distribution and the actual distribution. As mentioned above, in order to experimentally realize approximate sampling with small total variation distance, one needs quantum error correction. One can make strong assumptions on the noise model in order to weaken the experimental requirement of achieving a high fidelity. However, this state of affairs may be unsatisfactory if the intention of the experiments is to provide a refutation of the extended Church-Turing thesis. This is because the assumption of the noise model is justified by quantum mechanical models of the experiment, however it is these very models whose validity is being tested in an experiment purporting to demonstrate a quantum advantage, as argued in the previous section. Nevertheless, our purpose here is not to invalidate the extended Church-Turing thesis but merely to study, assuming the validity of quantum mechanics in all regimes, what features make a quantum system easy or hard to simulate.

Therefore, considering the hardness of sampling from the noisy distribution under a suitable noise model is acceptable. Moreover, such a study is more realistic than that of the idealized noiseless case.

We take such an approach in Chapters 4 and 5. In Chapter 4, we consider the complexity of simulating otherwise-free fermions subject to Markovian noise described by linear/quadratic Lindbladian jump operators. We classify the complexity of simulating the dynamics according to the type of jump operators present. Interestingly, in some cases it is the dissipation that leads to classical hardness of simulation, since free fermions without dissipation are known to be classically simulable [50–52]. In Chapter 5, we study the complexity of *Gaussian boson sampling*, the problem of sampling from the output distribution of a linear-optical network when certain Gaussian states are sent in as input [53; 54]. Our main aim in this chapter is to gather evidence for the persistence of a computationally complex quantum signal in the output distribution. To this end, we consider the problem of simulating lossy Gaussian boson sampling instances and prove that there is a sign of computational hardness in the output distribution. On the way, we also improve the theoretical analysis for idealized (lossless) Gaussian boson sampling and bring it up to the same standard as that of boson sampling.

1.4 Equilibrium

Simulating and predicting dynamics is far from being the only activity of interest to a physicist; predicting equilibrium properties of a wide variety of physical systems is an equally important task. This task involves computing properties of Hamiltonians with respect to the

thermal (Gibbs) state at temperature T, $\rho = e^{-\beta H} / \operatorname{Tr}\left(e^{-\beta H}\right)$. Specifically, tasks such as computing correlation functions, the partition function $Z = \operatorname{Tr}\left(e^{-\beta H}\right)$, the free energy and entropy, and Gibbs sampling are all of interest in various fields of physics. The last of these, Gibbs sampling, or sampling from the thermal state ρ, is a task whose solution allows for solving the other tasks as well [55; 56].

In computer science as well, Gibbs sampling is of interest in several contexts. One of these is in machine learning, through the use of restricted Boltzmann machines. Another is the matrix multiplicative weights update method, which is used in conjunction with Gibbs sampling [57] and is a useful subroutine in numerous algorithms (see [58] for a survey).

A special case of Gibbs states is that of states at zero temperature, or ground states. These are the minimal-energy eigenstates of the Hamiltonian. It is not an overstatement to claim that a majority of condensed matter physics deals with ground states, ground-state properties, excitations above the ground state and ground-state phase diagrams. Even in high-energy physics, finding ground-state (vacuum) properties and the spectral gap (mass gap) is among the first steps in analyzing a theory.

In computer science, constraint satisfaction problems can be mapped to classical Hamiltonians so that the ground state of the Hamiltonian encodes the answer to the constraint satisfaction problem. In turn, constraint satisfaction problems encode a wide variety of optimization problems across science, industry, and more generally, all aspects of human life. Hence, developing and improving algorithms for finding ground states, or low-energy states, of Hamiltonians would be immensely useful.

In this dissertation, we seek to understand the complexity of computing ground-state properties and preparing ground states. Unlike the case of quantum dynamics, this problem

can also be hard for quantum computers since it is not necessarily the case that nature can efficiently find ground states. There can be three different "phases" for this problem, namely i) classically (and quantumly) easy, ii) classically hard but quantumly easy, and iii) quantumly (and classically) hard. Understanding the boundaries between these phases would help give insight into what features of a Hamiltonian enable efficient preparation of its ground states. Additionally, akin to our studies in Chapters 2 to 4, this endeavor would result in a complexity classification.

The easiness/hardness of preparing ground states is studied in the field of Hamiltonian complexity, which is at the intersection of quantum information, many-body physics, and complexity theory. As a proxy for the hardness of preparing the ground state, we study the problem of finding ground-state energies[10]. Kitaev [59] showed that finding the ground-state energy of a local Hamiltonian is QMA-hard in general. Since then, there has been an effort to understand the complexity of this problem in more physically natural settings [60–63]. There has also been some effort on the algorithms side to provide heuristic quantum algorithms for the problem, such as the variational quantum eigensolver (VQE) [64]. It is an open question to rigorously formulate conditions under which these algorithms can be successful.

In Chapter 6, we take some initial steps toward finding conditions that provably affect the complexity of computing ground-state properties. Specifically, we identify two structural properties of Hamiltonians that affect the complexity of finding their ground-state energies. These are the spectral gap and the existence of a useful classical description of the ground

[10] Since it is possible to efficiently find the ground-state energy with access to the ground state, the hardness of finding ground-state energies implies the hardness of preparing ground states.

state. Our results lead us to make an interesting conjecture regarding the preparation of gapped ground states. This conjecture, if true, would explain the efficacy of the variational quantum eigensolver at preparing ground states of Hamiltonians.

1.5 Outline

The outline of this dissertation is as follows.

In Chapter 2 titled "Dynamical phase transitions in sampling complexity", we study the dynamics of n bosonic particles evolving under a class of nearest-neighbor bosonic Hamiltonians. We identify upper and lower bounds on the transition timescale by giving a sampling algorithm for short times and proving the hardness of the dynamics at longer times. As a corollary, we also obtain results for Anderson localized systems, for which the transition timescale diverges.

In Chapter 3 titled "Complexity phase diagram for interacting and long-range bosonic Hamiltonians", we improve upon these results in two ways. First, we tighten the bounds on the transition timescale from the previous chapter using a new technique. The same technique also allows us to generalize to the case of interacting bosons with long-range hopping. We observe that the phase transition is stable to perturbations such as the addition of interactions and change in their strength. Considering the transition timescale as a function of how long-range the system is also allows us to draw a complexity phase diagram.

In Chapter 4 titled "Complexity of fermionic dissipative interactions and applications to quantum computing", we classify the complexity of simulating the dynamics of free fermions subject to dissipation described by linear and quadratic Lindbladian jump operators. We

classify different jump operators based on the classical complexity of the dynamics, which is either classically easy or universal for quantum computing. We elaborate on the design of an entangling gate using dissipation by giving an example with fermionic atoms.

In Chapter 5 titled "Quantum computational supremacy via high-dimensional Gaussian boson sampling", we improve the available evidence for hardness of Gaussian boson sampling. We also show that computing output probabilities of a lossy Gaussian boson sampling instance is hard on average. We propose an architecture for Gaussian boson sampling designed to balance the competition between loss and connectivity.

Finally, in Chapter 6 titled "The importance of the spectral gap in estimating ground-state energies", we give a complexity classification of the problem of precisely estimating ground-state energies of Hamiltonians with or without a spectral gap. We obtain evidence that the spectral gap can make the problem easier. We also make an interesting conjecture on the relation between spectral gaps and the existence of polynomial-sized quantum circuits to prepare ground states.

The appendices contain relevant mathematical definitions and notation used in the dissertation and an overview of the basics of complexity theory.

Chapter 2: Dynamical phase transitions in sampling complexity

In the quest towards building scalable and fault-tolerant quantum computers, demonstration of a quantum speedup over the best possible classical computers is an important milestone and is termed *quantum computational supremacy* [65]. There are several candidates for tasks where such a speedup could be demonstrated [11–13; 18–22; 37; 51; 66; 67], where the problem is to simulate a quantum system in the sense of approximate sampling. However, there has also been some debate about the required system size before one can claim quantum computational supremacy, due to improved simulation techniques and algorithms [68–70]. In this Chapter, we consider the impact of the field of quantum computational supremacy on other areas of physics and show that studying the complexity of simulating physical systems is useful for understanding phase transitions.

Here we consider the classical complexity of approximate sampling, referred to as "sampling complexity". This is the task of producing samples from a distribution close to the probability distribution occurring in a quantum system upon measurement in a standard basis. This task is a good notion of what it means to simulate physics on a classical computer since it captures how well a computer can mimic an experiment in which one can measure the output at several sites. When we consider sampling complexity as a function of system parameters, the system can be classified as easy in some regimes and hard in some others.

Since the designations "easy" and "hard" are exhaustive and there is no smooth way to go from one regime to another, we posit that this transition from easy to hard happens abruptly, a phenomenon very reminiscent of phase transitions. Just like order parameters are zero on one side of a phase transition and nonzero on another, sampling complexity is different on either side of the transition, and can be used to draw phase boundaries as a function of time and other system parameters. These boundaries can be different from those drawn by more conventional order parameters, signifying new physics in otherwise well-studied systems. Indeed, phase transitions in average-case complexity have been studied both in the classical [71] and quantum regime [72].

In this Chapter, we show that transitions in sampling complexity [73] are indicative of physical transitions. We consider a system of n bosons hopping from one site to another on a lattice of m sites and study the sampling complexity as a function of time for an initial product state. We show that it goes from easy to hard as the scaling of evolution time t with the number of bosons n increases, thereby exhibiting a dynamical phase transition [74; 75]. We find that the timescale at which the complexity changes is the timescale when interference effects start becoming relevant. We conjecture that in general, this is linked to the Ehrenfest timescale at which quantum effects in a system become considerable [76]. We also show that systems in the Anderson-localized phase are always easy to simulate. We use the Lieb-Robinson bound [77] as an ingredient in our proof of easiness.

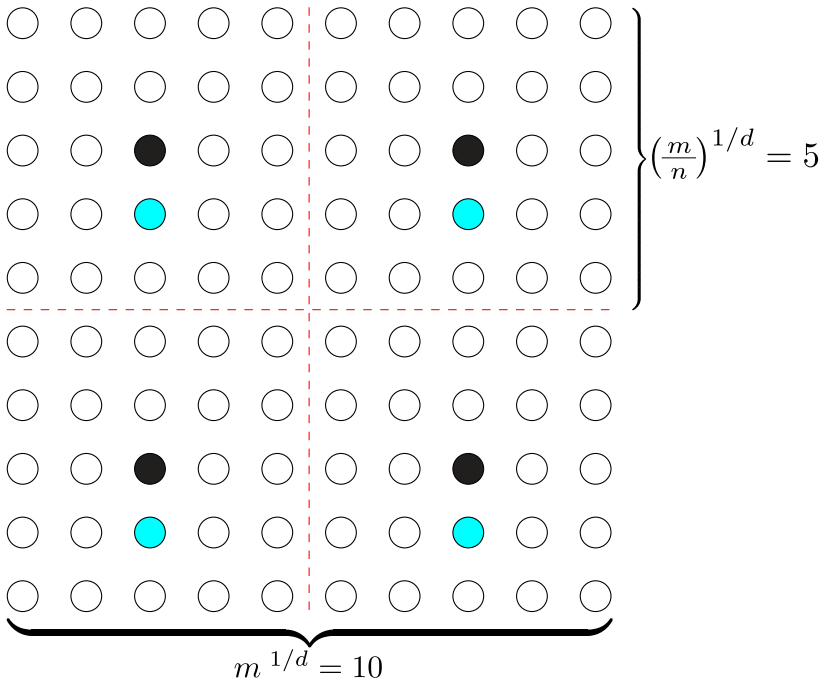

Figure 2.1: An example of the initial state in $d = 2$ dimensions. Here $m = 96$, $n = 4$, $\beta = 3$ and $c_1 = 3/2$. The black circles represent sites with a single boson. The cyan circles represent the ancillas.

2.1 Setup

The model consists of free bosons hopping on a lattice in d dimensions (denoted dD) with sites labeled by indices i, j. Our results in this Chapter can be applied to linear optics as well, with the bosonic sites being replaced by photonic modes. The Hamiltonian is given by $H = \sum_{i,j=1}^{m} J_{ij}(t) a_i^\dagger a_j$, where $a_i^\dagger$ is the creation operator of a boson at the i'th site. $J(t)$, which can be time-dependent in general, is an $m \times m$ Hermitian matrix that encodes the connectivity of the lattice. One way to show hardness in the boson sampling proposal by Aaronson and Arkhipov (AA) [13] is to generate any linear optical unitary U acting on the bosonic sites, in particular, Haar-random unitaries. Any U can be generated through a free boson Hamiltonian by taking $H = i \log U$ and evolving it for unit time. However, this

Hamiltonian can require arbitrarily long-range hops on the lattice in general. Since such long-ranged Hamiltonians may not be realistic, we consider the sampling complexity of a bosonic Hamiltonian with nearest-neighbor hops. J_{ij} is nonzero only if $i = j$ or i and j label adjacent sites. We further restrict J_{ij} to satisfy $|J_{ij}| \leq 1$ in order to set an energy scale. Our proofs remain valid even if we allow the diagonal terms J_{ii} to be unbounded.

One can efficiently solve the equations of motion $i\dot{a_i}^\dagger(t) = [a_i^\dagger(t), H(t)]$ on a classical computer to obtain $a_i^\dagger(t) = \sum_k a_k^\dagger(0) R_{ki}(t)$ for some transformation matrix R. From here onward, we shall take $R(t)$ to be the input to the problem, since it can be determined from the input Hamiltonian H and time t in time poly$(m, \log t)$.

The m sites in the problem are numbered from 1 to m, and together with n ancilla sites, are arranged in a lattice of side length $(m+n)^{1/d}$ in d dimensions. The initial state has n bosons equally spaced in the lattice as shown in Fig. 2.1. We take $m = c_1 n^\beta$, where β controls the sparsity of occupied sites in the lattice and can be set to 5 as required for the hardness of boson sampling [13]. The minimum spacing between any two bosons in the initial state is $2L = \left(\frac{m+n}{n}\right)^{1/d} > c_1^{1/d} n^{\frac{\beta-1}{d}}$. The quantity L is an important length scale in the problem. The ancillas in the lattice, marked in cyan, are not counted as part of the m sites and are present in order to accelerate the time required to construct an arbitrary unitary, which is useful for the hardness result. Their presence does not change the scaling of quantities like L with n.

The input states are described by vectors of the form $r = (r_1, \ldots, r_m)$, specifying the number of bosons on each site, so that $r_1 + \ldots r_m = n$. Measurement in the boson number basis defines a distribution $\mathcal{D}_U$, which we aim to sample from. The probability of finding an

output state $s = (s_1, s_2, \ldots, s_m)$ is given by

$$\Pr_{\mathcal{D}_U}[s] = \frac{1}{r!s!}|\text{Per}(A)|^2, \tag{2.1}$$

where $r! := r_1! \ldots r_m!$ (with $s!$ defined similarly), $A_{n\times n}$ is a matrix formed by taking s_i copies of the i'th column and r_j copies of the j'th row of R in any order, and $\text{Per}(A)$ denotes the permanent of A (see Ref. [78] or Section 2.5 for details).

For the particular choice of initial states described in Fig. 2.1, the task is to sample from a distribution that is close to $\mathcal{D}_U$ in variation distance when given a description of the unitary $R(t)$. We now formalize the notion of efficient sampling.

Definition 1. Efficient sampler: An efficient sampler is a classical randomized algorithm that takes as input the unitary R_{ij} and outputs a sample s from a distribution $\mathcal{D}_\mathcal{O}$ such that the variation distance between the distributions $\epsilon = \|\mathcal{D}_\mathcal{O} - \mathcal{D}_U\| \leq O(\frac{1}{\text{poly}(n)})$, in runtime $\text{poly}(n)$[1].

We call the sampling problem *easy* if there exists an efficient sampler for the problem in the stated regime. Conversely, the problem is hard if there cannot be an efficient sampler. Since a negative statement such as the inexistence of an algorithm is difficult to prove, our practical definition of hardness of a sampling problem is if it is at least as hard as boson sampling. This enables us to use the results of AA [13] to claim the hardness of sampling in some regimes. In doing so, our hardness results ultimately rely on the truth of AA's conjectures. One of these conjectures concerns the hardness of additively approximating $|\text{Per}(G)|^2$

[1]Our easiness results also apply to a stronger notion of approximate sampling, namely that the algorithm samples from an ϵ-close distribution in runtime $\text{poly}(n, 1/\epsilon)$.

for Gaussian-random G with high probability (for which AA give reasonable evidence). The second, more widely believed conjecture is that the polynomial hierarchy, an infinite tower of complexity classes, does not "collapse", i.e. is truly infinite.

We restrict our attention to two special cases where we can show the existence of an efficient sampling algorithm: i) when the system evolves for a time smaller than the system timescale L/v (where v is the Lieb-Robinson velocity of information spreading in the lattice, defined more precisely in Eq. (2.2)), and ii) when there is Anderson localization in the system [49]. These two cases correspond to a promise on the input unitary R. We now state our main results.

Theorem 1.A (Easiness of simulation at short times). *For $\beta > 1$ and for all dimensions d, the sampling problem is easy for all $t \leq 0.9L/v$, i.e. $\forall t \leq c_2 n^{(\beta-1)/d}$ for some constant c_2.*

The intuition behind this Theorem is that when the time is smaller than the Lieb-Robinson timescale of particle interference L/v, the dynamics is approximately classical (in the sense that the particles are distinguishable).

An important technical achievement of this Chapter is to prove rigorously that this intuition is correct. This is done in Lemma 3 by showing that the approximation works.

Theorem 1.B (Hardness of simulation at longer times). *(based on Theorem 3 of Ref. [13]) When $t = \Omega(n^{1+\beta/d})$, the sampling problem is hard in general.*

For this result, we show that we can apply any unitary R after the stated time. Therefore, if an efficient sampler exists for this problem, then we have an efficient boson sampler, which cannot exist by our assumption.

The result for the case of Anderson localization comes out as a corollary from Theorem 1.A:

Corollary 2 (Easiness of Anderson-localized systems). *For Anderson-localized systems in any dimension d, the sampling problem is easy for all times.*

The easiness of sampling for Anderson-localized systems is analogous to results showing efficient simulation of localized systems according to various definitions [79–83].

2.2 Easiness at short times

In this section, we prove Theorem 1.A and Corollary 2. First, let us examine the promise we have on the unitary R in both cases. We use the Lieb-Robinson bound [77] on the speed of information propagation in a system. Applying the bound to our Hamiltonian, we get

$$|[a_i(t), a_j^\dagger(0)]| = |R_{ij}(t)| \leq \min\left(1, \exp\left(\frac{vt - \ell_{ij}}{\xi}\right)\right), \tag{2.2}$$

where ℓ_{ij} is the distance between two sites i and j, v is the upper bound to the velocity of information propagation called the Lieb-Robinson velocity and ξ is a length scale). Note that the results from Ref. [84] do not apply since we have free bosons here and we work in the single-particle subspace. The Lieb-Robinson velocity is at most $4(1 + 2de)$ [85] when $|J_{ij}| \leq 1$ and $\xi = 1$.

When the Hamiltonian is Anderson-localized, the unitary R satisfies the following

promise at all times [86]:

$$|R_{ij}| \leq \exp\left(\frac{-\ell_{ij}}{\xi}\right). \tag{2.3}$$

Here, ξ is the maximum localization length among all eigenvectors. Equation (2.3) can be viewed as a consequence of a Lieb-Robinson bound with zero velocity [87]. On account of the zero-velocity Lieb-Robinson bound, all results for the time-dependent case can be ported to the Anderson localized case, setting $v = 0$.

We give an algorithm that efficiently samples from the output distribution for short times $t < \frac{0.9L}{v} = O(n^{(\beta-1)/d})$, given the promise in Eq. (2.2). The algorithm outputs a sample from a distribution $\mathcal{D}_{DP}$, obtained by assuming that the bosons are distinguishable particles, ignoring the effects of interference. The algorithm is described in more detail in Section 2.6.

2.3 Analysis

We prove the correctness of the algorithm by showing that the variation distance between the distributions is upper bounded by an inverse exponential in n. When the bosons are distinguishable, their dynamics is given by a Markov process, described by the matrix $\mathcal{P}_{kl} = |R(t)|^2_{kl}$. The probability of getting an outcome s is given by

$$\Pr_{\mathcal{D}_{DP}}[s] = \sum_{\sigma} \frac{1}{s!} \mathcal{P}_{\mathsf{in}_1,\mathsf{out}_{\sigma(1)}} \mathcal{P}_{\mathsf{in}_2,\mathsf{out}_{\sigma(2)}} \ldots \mathcal{P}_{\mathsf{in}_n,\mathsf{out}_{\sigma(n)}}, \tag{2.4}$$

where the sum is over all permutations σ mapping the input bosons to the output ones. In the above equations, in_i is the site index of the i'th boson in nondecreasing order in the input and out is defined similarly at the output (see Section 2.5 for an example). We now state a result on how close $\mathcal{D}_{DP}$ is to the true distribution $\mathcal{D}_U$.

Lemma 3. *When $\beta > 1$ and $t \leq 0.9L/v$, the variation distance satisfies $\|\mathcal{D}_{DP} - \mathcal{D}_U\| = O(\exp\left[2\frac{vt-L}{\xi} + 2(d-1)\log L\right])$.*

This Lemma makes intuitive sense: because of the Lieb-Robinson bound Eq. (2.2) and the fact that the initially occupied sites are separated by a minimum distance $\Theta(L)$ from each other, it takes a time $t = \Theta(L/v)$ for the bosons to start interfering considerably. Therefore, the classical and quantum distributions agree exponentially closely in L when $t \leq 0.9L/v$. For a proof of Lemma 3, see Section 2.6. Assuming this Lemma, we now show Theorem 1.A.

Proof of Theorem 1.A. Lemma 3 shows that the algorithm samples from a distribution with exponentially small error in n, since $L = \Theta(n^{(\beta-1)/d})$. To complete the proof of Theorem 1.A, we need to show that the runtime of the algorithm is polynomial in n. This is true because the corresponding Markov process of n distinguishable bosons walking on m sites for one step is efficiently simulable: for each of the n particles, we select one among the m sites to walk to, based on the matrix elements of $\mathcal{P}$. This takes time $O(n \cdot \mathrm{poly}(m))$ = $O(\mathrm{poly}(n))$ to simulate on a classical computer. □

2.4 Hardness at longer times

If we allow the system to evolve for a longer amount of time, we can use the time-dependent control to effect any arbitrary unitary and implement any boson sampling instance in the system. We can perform phase gates on a site k by setting J_{kk} to be nonzero for a particular time, with the hopping terms and other diagonal terms set to zero. We can apply a nontrivial two-site gate between adjacent sites, for example the balanced beamsplitter unitary on the sites 1 and 2, $U = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$, by setting $H = -i(a_1^\dagger a_2 - a_1 a_2^\dagger)$ for time $t = \frac{\pi}{4}$. One can also apply arbitrary unitaries using arbitrary on-site control $J_{ii}(t)$ and fixed, time-independent nearest-neighbor hopping $J_{ij}(t) = 1$.

Using the constructions in Refs. [88; 89], we can effect any arbitrary $m \times m$ unitary on the m sites with $O(m^2)$ depth from a nontrivial beamsplitter and arbitrary single-site phase gates. A construction of AA that employs ancillas to obtain the desired final state rather than applying the full unitary on all the sites can be used to achieve a depth $O(nm^{1/d})$. Each of the n columns of the unitary are implemented in time $O(m^{1/d})$, which corresponds to the timescale set by the Lieb-Robinson velocity and the distance between the furthest two sites in the system[2].

Proof of Theorem 1.B. From the above, when $t = \Omega(n^{1+\frac{\beta}{d}})$, we see that we can effect any arbitrary unitary. This implies that an efficient sampler for this regime can also be used as an efficient boson sampler, which is widely believed not to exist because of AA's results [13]. □

[2] AA's construction applied each column of the unitary in $O(\log m)$-depth, whereas we can only apply it in $O(m^{1/d})$ depth because of the spatial locality of the Hamiltonian and the Lieb-Robinson bounds that follow.

In Chapter 3, we improve upon this bound by applying a different technique.

2.5 Expression for output probabilities

In this section, we describe the standard boson sampling set-up and derive an expression for the output probabilities of a boson sampling experiment that define the distribution $\mathcal{D}_U$. First, let us represent the input and output states pictorially and develop some notation.

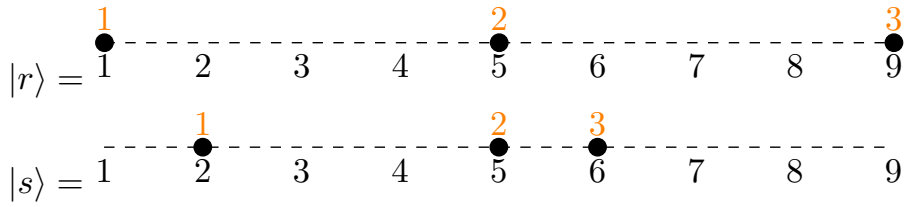

Figure 2.2: A representation of input and output basis states in 1D.

In Fig. 2.2, the top line denotes the input state $|r\rangle$ and the bottom line the output $|s\rangle$. Each filled circle denotes a boson occupying the corresponding lattice site, which is labeled below the circles. The numbers marked in orange above each boson label the bosons from left to right (more generally, this is in a nondecreasing order of the site index). We will call this the boson index.

A given configuration (basis state) is completely specified by specifying the boson number in each site, such as $r = (1, 0, 0, 0, 1, 0, 0, 0, 1)$ and $s = (0, 1, 0, 0, 1, 1, 0, 0, 0)$ in the above. It can also be specified by listing the site index for every boson index, i.e. the occupied sites. Thus the input state can be represented as $\mathsf{in} = (1, 5, 9)$, the output state as $\mathsf{out} = (2, 5, 6)$.

All $n!$ permutations of the boson indices represent valid paths that the bosons can take to the output state, and correspond to the $n!$ terms in the permanent of the matrix. In cases

where there are two or more bosons in a particular site at the input or output, there are $\frac{n!}{r!s!}$ paths (and terms in the amplitude). Here, $r! := r_1!r_2!\ldots r_m!$ and similarly $s!$. By taking repeated rows and columns of R, this has the effect of still giving $n!$ terms in total, which we identify with the $n!$ permutations in the boson indices. The expression for the probability of an outcome s is (here, $b_i := a_i(t)$):

$$\Pr_{\mathcal{D}_U}[s] = \frac{1}{r_1!r_2!\ldots r_m!s_1!s_2!\ldots s_m!}\left|\langle \text{vac}|\, b_1^{s_1}b_2^{s_2}\ldots b_m^{s_m}a_1^{\dagger r_1}a_2^{\dagger r_2}\ldots a_m^{\dagger r_m}\,|\text{vac}\rangle\right|^2 \tag{2.5}$$

$$= \frac{1}{r!s!}\left|\langle \text{vac}|\, (\hat{U}^\dagger a_1^{s_1}\hat{U})(\hat{U}^\dagger a_2^{s_2}\hat{U})\ldots(\hat{U}^\dagger a_m^{s_m}\hat{U})a_1^{\dagger r_1}a_2^{\dagger r_2}\ldots a_m^{\dagger r_m}\,|\text{vac}\rangle\right|^2 \tag{2.6}$$

$$= \frac{1}{r!s!}\left|\langle \text{vac}|\left(\sum_{k_1=1}^{m} R^\dagger_{1k_1}a_{k_1}\right)^{s_1}\ldots\left(\sum_{k_m=1}^{m} R^\dagger_{mk_m}a_{k_m}\right)^{s_m} a_1^{\dagger r_1}a_2^{\dagger r_2}\ldots a_m^{\dagger r_m}\,|\text{vac}\rangle\right|^2, \tag{2.7}$$

where R_{ij} describes the action of $\hat{U}$ on the annihilation operators at a site: $b_i = a_i(t) = \sum_k R^\dagger_{ik}(t)a_k(0)$. Now define the matrix $A^\dagger$ to be the one obtained by taking s_i copies of the i'th row and r_j copies of the j'th column of $R^\dagger$. For concreteness, this can be done by first considering the rows and repeating a row i of $R^\dagger$ whenever $s_i > 1$, or not picking it if $s_i = 0$, to convert it into an $n \times m$ matrix. We can then do the same with columns to convert it into an $n \times n$ matrix. However, the order ultimately does not matter since the quantity that emerges, the permanent, is symmetric under exchange of rows or columns. We have

$$\Pr_{\mathcal{D}_U}[s] = \frac{1}{r!s!}\left|\sum_\sigma\prod_i R^\dagger_{\mathsf{out}_{\sigma(i)},\mathsf{in}_i}\right|^2 \tag{2.8}$$

$$= \frac{1}{r!s!}\left|\sum_\sigma\prod_i A^\dagger_{\sigma(i),i}\right|^2, \tag{2.9}$$

where the sum is over all permutations σ. This finally gives us

$$\Pr_{\mathcal{D}_U}[s] = \frac{1}{r!s!}\left|\mathrm{Per}(A^\dagger)\right|^2 = \frac{1}{r!s!}|\mathrm{Per}(A)|^2, \tag{2.10}$$

where $\mathrm{Per}(A)$ is the *permanent* of A.

2.6 Algorithm

The sampling algorithm is given below. It is easy to see that it implements one step of a Markov process of n distinguishable bosons walking on a lattice.

Algorithm 1: Sampling algorithm

Input: Unitary $R(t)$, tolerance ϵ

Output: Sample s drawn from $\mathcal{D}_{DP}$, a distribution that is close to $\mathcal{D}_U$.

1 $\mathcal{P}_{kl} = |R(t)|^2_{kl}$

2 **for** $i \in \{1, 2, \ldots n\}$**, do**

3 Select site l from the distribution $\mathcal{P}_{\mathsf{in}_i,l}$ for the boson at in_i to hop to.

4 Increment output boson number of site l by 1: $s_l \to s_l + 1$ (or equivalently, assign $\mathsf{out}_i = l$)

5 **end**

6 **return** configuration s (or out), a sample from $\mathcal{D}_{DP}$.

Note that $\mathcal{P}$ from line 1 is a doubly stochastic matrix describing the classical Markov process. To see that the runtime is polynomial in n, note that the loop is over n boson indices. Line 3 takes time $O(m \log m) = \tilde{O}(n^\beta)$, giving a total runtime of $\tilde{O}(nm) = \tilde{O}(n^{1+\beta})$. The notation $\tilde{O}$ suppresses factors of $\log n$.

2.7 Bound on variation distance

Here we derive a bound on the variation distance $\|\mathcal{D}_U - \mathcal{D}_{DP}\| = \frac{1}{2}\sum_s |\Pr_{\mathcal{D}_U}(s) - \Pr_{\mathcal{D}_{DP}}(s)|$. Rewriting the actual probability in terms of the amplitudes, we have

$$\Pr_{\mathcal{D}_U}(s) = |\phi|^2, \quad \text{with} \tag{2.11}$$

$$\phi^* = \frac{1}{\sqrt{r!s!}} \sum_\sigma R_{\mathsf{in}_1,\mathsf{out}_{\sigma(1)}} R_{\mathsf{in}_2,\mathsf{out}_{\sigma(2)}} \ldots R_{\mathsf{in}_n,\mathsf{out}_{\sigma(n)}} \tag{2.12}$$

$$= \frac{1}{\sqrt{s!}} \sum_\sigma A_{1,\sigma(1)} A_{2,\sigma(2)} \ldots A_{n,\sigma(n)}, \tag{2.13}$$

where A is the $n \times n$ matrix formed by taking the appropriate number of copies of each row and column of $R_{m\times m}$. We have set $r! = 1$ since our input state has bosons in distinct sites. Continuing,

$$\Pr_{\mathcal{D}_U}(s) = \frac{1}{s!} \sum_\sigma |R_{\mathsf{in}_1,\mathsf{out}_{\sigma(1)}}|^2 |R_{\mathsf{in}_2,\mathsf{out}_{\sigma(2)}}|^2 \ldots |R_{\mathsf{in}_n,\mathsf{out}_{\sigma(n)}}|^2 +$$

$$\frac{1}{s!} \sum_{\sigma\neq\tau} R_{\mathsf{in}_1,\mathsf{out}_{\sigma(1)}} R_{\mathsf{in}_2,\mathsf{out}_{\sigma(2)}} \ldots R_{\mathsf{in}_n,\mathsf{out}_{\sigma(n)}} (R_{\mathsf{in}_1,\mathsf{out}_{\tau(1)}} R_{\mathsf{in}_2,\mathsf{out}_{\tau(2)}} \ldots R_{\mathsf{in}_n,\mathsf{out}_{\tau(n)}})^*. \tag{2.14}$$

The probability distribution $\mathcal{D}_{DP}$ that the algorithm samples from is given by the first line of Eq. (2.14):

$$\Pr_{\mathcal{D}_{DP}}(s) = \frac{1}{s!} \sum_\sigma \mathcal{P}_{\mathsf{in}_1,\mathsf{out}_{\sigma(1)}} \mathcal{P}_{\mathsf{in}_2,\mathsf{out}_{\sigma(2)}} \ldots \mathcal{P}_{\mathsf{in}_n,\mathsf{out}_{\sigma(n)}}, \tag{2.15}$$

where the sum is over all the $n!$ ways of assigning the n input states to the n output states. As before, the $s!$ is to account for overcounting when two distinct permutations in the boson

index refer to the same site index in the output state.

The expression for the probability is proportional to the permanent of the matrix with the positive entries $\mathcal{P}_{\mathsf{in}_i,\mathsf{out}_j}$, and can hence be efficiently approximated [90]. Note that the algorithm does not explicitly calculate this probability but only samples from the distribution. We can now prove Lemma 3.

Proof of Lemma 3. The variation distance is given by

$$\varepsilon = \sum_s \frac{1}{2s!} \left| \sum_{\sigma \neq \tau} R_{\mathsf{in}_1,\mathsf{out}_{\sigma(1)}} \dots R_{\mathsf{in}_n,\mathsf{out}_{\sigma(n)}} (R_{\mathsf{in}_1,\mathsf{out}_{\tau(1)}} \dots R_{\mathsf{in}_n,\mathsf{out}_{\tau(n)}})^* \right| \tag{2.16}$$

$$\leq \sum_s \frac{1}{2} \sum_{\sigma \neq \tau} |R_{\mathsf{in}_1,\mathsf{out}_{\sigma(1)}} \dots R_{\mathsf{in}_n,\mathsf{out}_{\sigma(n)}}| |R_{\mathsf{out}_{\tau(1)},\mathsf{in}_1} \dots R_{\mathsf{out}_{\tau(n)},\mathsf{in}_n}| \tag{2.17}$$

$$= \sum_s \frac{1}{2} \sum_{\sigma,\rho} |R_{\mathsf{in}_1,\mathsf{out}_{\sigma(1)}} R_{\mathsf{out}_{\sigma(1)},\mathsf{in}_{\rho(1)}}| \dots |R_{\mathsf{in}_n,\mathsf{out}_{\sigma(n)}} R_{\mathsf{out}_{\sigma(n)},\mathsf{in}_{\rho(n)}}|, \tag{2.18}$$

where $\rho = \tau^{-1} \circ \sigma \neq \mathsf{Id}$, the identity permutation. The last equality comes from rearranging the terms in $|R_{\mathsf{out}_{\tau(1)},\mathsf{in}_1} \dots R_{\mathsf{out}_{\tau(n)},\mathsf{in}_n}|$ so that the terms involving $R_{\mathsf{in}_i,\mathsf{out}_{\sigma(i)}}$ and $R_{\mathsf{out}_{\sigma(i)},j}$ (for some j) are together:

$$\sum_\sigma \sum_\tau \prod_i |R_{\mathsf{out}_{\tau(i)},\mathsf{in}_i}| = \sum_\sigma \sum_\tau \prod_i |R_{\mathsf{out}_i,\mathsf{in}_{\tau^{-1}(i)}}| \xrightarrow[i \to \sigma(i)]{\text{rearrange}} \sum_\sigma \sum_\tau \prod_i |R_{\mathsf{out}_{\sigma(i)},\mathsf{in}_{\tau^{-1}(\sigma(i))}}|. \tag{2.19}$$

Summing over all outcomes s (or configurations out), Eq. (2.18) is equivalent to

$$\varepsilon \leq \frac{1}{2} \sum_j \sum_{\rho \neq \mathsf{Id}} \prod_i |R_{\mathsf{in}_i,j_i}| |R_{j_i,\mathsf{in}_{\rho(i)}}|, \tag{2.20}$$

where the sum j is over ordered tuples $(j_1, \dots j_n)$, representing the intermediate lattice

sites that the bosons in positions $(\mathsf{in}_1, \ldots \mathsf{in}_n)$ jump to, before jumping back to positions $(\mathsf{in}_{\rho(1)}, \ldots \mathsf{in}_{\rho(n)})$. We can proceed to break the sum in Eq. (2.20) based on the number of fixed points of the permutation ρ, that is, the number of indices i such that $\rho(i) = i$. We bound these quantities separately as follows:

$$\sum_{\substack{j_i,\rho \\ \rho(i)\neq i}} |R_{\mathsf{in}_i,j_i}||R_{j_i,\mathsf{in}_{\rho(i)}}| = C_i \leq c \;\forall\; i \;\text{ and}$$

$$\sum_{\substack{j_i,\rho \\ \rho(i)= i}} |R_{\mathsf{in}_i,j_i}||R_{j_i,\mathsf{in}_{\rho(i)}}| = \sum_{j_i} |R_{\mathsf{in}_i,j_i}|^2 = D_i = 1 \;\forall\; i. \tag{2.21}$$

The variation distance is therefore bounded above:

$$\varepsilon \leq \frac{1}{2} \sum \prod_{i\in\mathcal{I}_C} C_i \prod_{k\in\mathcal{I}_D} D_k, \tag{2.22}$$

where the sum is over subsets $\mathcal{I}_C$ of the indices representing the input state, in. $\mathcal{I}_D$ is the complement of $\mathcal{I}_C$ and $|\mathcal{I}_D|$ is the number of fixed points. Suppose we find an upper bound c for C_i in Eq. (2.21), we then have

$$\varepsilon \leq \frac{1}{2} \sum_{l=|\mathcal{I}_C|=2}^{n} \binom{n}{l} c^l = (c+1)^n - nc - 1. \tag{2.23}$$

In Lemma 4, we show that $c = \eta L^{d-1} \mathrm{e}^{(vt-L)/\xi}$ for some constant η. Continuing from

Eq. (2.23),

$$\varepsilon \leq \frac{1}{2}[(c+1)^n - nc - 1] \tag{2.24}$$

$$= \binom{n}{2}(1+h)^{n-2}c^2 \text{ for some } h \in [0,c] \text{ (by Taylor's theorem)} \tag{2.25}$$

$$\varepsilon \leq \exp\left[2\log n + (n-2)\log(1+c) + 2\log c\right]. \tag{2.26}$$

Now, plugging in the value of c and assuming that $vt \leq 0.9L$ and $\beta > 1$, we get

$$\varepsilon \leq O\left(\exp\left[(n-2) \times \eta L^{d-1}\mathrm{e}^{(vt-L)/\xi} + 2\frac{vt-L}{\xi} + 2(d-1)\log L\right]\right) \tag{2.27}$$

$$\leq O\left(\exp\left[2\frac{vt-L}{\xi} + 2(d-1)\log L\right]\right). \tag{2.28}$$

In the first line, we use the inequality $\log(1+x) \leq x$. In the second, we use

$$\exp\left[(n-2) \times \eta L^{d-1}\mathrm{e}^{(vt-L)/\xi}\right] = O(1) \tag{2.29}$$

since $vt < L$ and $|vt - L| = \Omega(n^{\beta-1})$. This completes the proof of Lemma 3. □

Lemma 4. *For all constant dimensions d, $c = \eta L^{d-1}\mathrm{e}^{(vt-L)/\xi}$.*

Proof. Recall that

$$C_i = \sum_{j_i} \sum_{\rho(i) \neq i} |R_{\mathsf{in}_i, j_i}| \left|R_{j_i, \mathsf{in}_{\rho(i)}}\right|. \tag{2.30}$$

Since we are looking for a bound that applies for all i, let us, for convenience, make the following changes in notation: $\mathsf{in}_i \to i, j_i \to j, \mathsf{in}_{\rho(i)} \to k$, denoting a boson starting at

position i, jumping to j and then to k, where i and $k \neq i$ are site indices belonging to in. We split the sum in C_i based on the distance between i and j, $\ell_{ij} =: \ell$.

$$C_i = \sum_{\substack{j \\ \ell \leq L}} \sum_{\substack{k \in \mathsf{in} \\ k \neq i}} |R_{i,j}| \, |R_{j,k}| + \sum_{\substack{j \\ \ell > L}} \sum_{\substack{k \in \mathsf{in} \\ k \neq i}} |R_{i,j}| \, |R_{j,k}| \, . \tag{2.31}$$

Consider the first term:

$$\sum_{\substack{j \\ \ell \leq L}} |R_{i,j}| \sum_{\substack{k \in \mathsf{in} \\ k \neq i}} |R_{j,k}| \leq \sum_{\substack{j \\ \ell \leq L}} |R_{i,j}| \sum_{\substack{k \in \mathsf{in} \\ k \neq i}} \mathrm{e}^{(vt - \ell_{kj})/\xi} \tag{2.32}$$

$$\leq \sqrt{\left(\sum_{\substack{j \\ \ell \leq L}} 1^2 \right) \left(\sum_{\substack{j \\ \ell \leq L}} |R_{i,j}|^2 \right)} \mathrm{e}^{vt/\xi} \sum_{k : \|\vec{x}_k\| \geq 1} \mathrm{e}^{(-2L\|\vec{x}_k\| + L)/\xi} \tag{2.33}$$

$$\leq a L^{d/2} \mathrm{e}^{(vt+L)/\xi} \sum_{k : ||\vec{x}_k|| \geq 1} \mathrm{e}^{-2L||\vec{x}_k||/\xi} \tag{2.34}$$

$$\leq ab \mathrm{e}^{(vt-L)/\xi} L^{d/2}. \tag{2.35}$$

Here in the first line, we have used the Lieb-Robinson bound Eq. (2.2). In the second line, we use the Cauchy-Schwarz inequality and the fact that $\ell = \ell_{ij} \leq L$. In the second line, $\vec{x}_k$ is the position vector of site k relative to site i, re-scaled by $2L$. Therefore the sum over $\vec{x}_k$ is over all vectors with integer coordinates. In the last line, we use Lemma 6, to be proven below. a and b are constants independent of n that depend on the dimension d and the length scale ξ.

Now, in the second term for C_i in Eq. (2.31), the intermediate site j is not necessarily close to i. Therefore, there are terms where j is close to $k \neq i$ and one has to treat these terms carefully. For these terms, we use the trivial Lieb-Robinson bound of 1 in Eq. (2.2)

rather than $\exp\left(\frac{vt-\ell_{jk}}{\xi}\right) > 1$.

$$\sum_{\substack{j \\ \ell > L}} |R_{i,j}| \sum_{\substack{k \in \mathsf{in} \\ k \neq i}} |R_{j,k}| \leq \sum_{\substack{j \\ \ell > L}} |R_{i,j}| \sum_{\substack{k \in \mathsf{in} \\ k \neq i}} \min(1, \mathrm{e}^{(vt-\ell_{k,j})/\xi}) \tag{2.36}$$

$$\leq \sum_{\substack{j \\ \ell > L}} |R_{i,j}| \left(1 + \mathrm{e}^{vt/\xi} \sum_{k:||\vec{x}_k|| \geq 1} \mathrm{e}^{(L-2L\|\vec{x}_k\|)/\xi}\right) \tag{2.37}$$

$$\leq \left(1 + b\mathrm{e}^{(vt-L)/\xi}\right) \sum_{\substack{j \\ \ell > L}} |R_{i,j}| \tag{2.38}$$

$$\leq \left(1 + b\mathrm{e}^{(vt-L)/\xi}\right) \mathrm{e}^{vt/\xi} \sum_{\substack{j \\ \ell > L}} \mathrm{e}^{-\ell/\xi} \tag{2.39}$$

$$\leq \left(1 + b\mathrm{e}^{(vt-L)/\xi}\right) \tilde{b} L^{d-1} \mathrm{e}^{(vt-L)/\xi}. \tag{2.40}$$

In the second line, we use 1 as a Lieb-Robinson bound for $|R_{j,k}|$ when $k = k^*$, the site belonging to in that is closest to j. All other k's have distances from j bounded below by $2L\|\vec{x}_k\| - L$, where $\vec{x}_k$ is now the re-scaled position vector of a site k with the origin at k^*. We apply Lemma 6 in the third line and Lemma 5 in the fifth. Collecting everything, we have

$$C_i \leq ab\mathrm{e}^{(vt-L)/\xi} L^{d/2} + \left(1 + b\mathrm{e}^{(vt-L)/\xi}\right) \tilde{b} L^{d-1} \mathrm{e}^{(vt-L)/\xi} \tag{2.41}$$

$$\leq \mathrm{e}^{(vt-L)/\xi} \eta L^{d-1} \text{ for large enough } L. \tag{2.42}$$

□

Lemma 5 (d-dimensional sum)**.** *The sum* $\sum_{\|\vec{x}\| \geq L} \mathrm{e}^{-\|\vec{x}\|/\xi}$ *over points* $\vec{x}$ *with integer coordinates*

is upper bounded by $a_d \mathrm{e}^{-L/\xi}\left(\xi L^{d-1}\right)$ *for large enough* $\frac{L}{\xi}$ *for some dimension-dependent constant* a_d.

Proof. We can view the sum over a lattice of vectors with integer coordinates as a Riemann sum and bound the corresponding d-dimensional integral. Consider the quantity

$$g = \int_{\|\vec{x}\| \geq L} \mathrm{e}^{-\|\vec{x}\|/\xi} \mathrm{d}^d \vec{x} = \frac{2\pi^{d/2}}{\Gamma(\frac{d}{2})} \xi^d \Gamma\left(d, \frac{L}{\xi}\right), \tag{2.43}$$

where $\Gamma(d,x) = \int_x^\infty y^{d-1}\mathrm{e}^{-y}\mathrm{d}y$ is the incomplete Γ function. We can lower bound the integral by the Riemann sum $\sum_\Delta V_\Delta \mathrm{e}^{-\|\vec{y}\|/\xi}$, where the sum is over cells Δ with volume V_Δ centered at lattice points $\vec{x}$. $\vec{y}$ is the point in the cell Δ with the highest norm $\|\vec{y}\|$. Further, the point with the highest norm is not too distant from the one at the center: $\|\vec{y}\| \leq \|\vec{x}\| + \frac{\sqrt{d}}{2}$. Therefore, we have

$$f := \sum_{\|\vec{x}\| \geq L} \mathrm{e}^{\|\vec{x}\|/\xi} \leq g \times \mathrm{e}^{\sqrt{d}/(2\xi)}. \tag{2.44}$$

We now need an upper bound on the incomplete Γ function $\Gamma(d,x)$ for large x [91, Eq. 8.11.i]:

$$\Gamma(d,x) \to x^{d-1}\mathrm{e}^{-x}\left(1 + O\left(\frac{1}{x}\right)\right) \text{ as } x \to \infty. \tag{2.45}$$

Combining Eq. (2.43) and Eq. (2.44), we get:

$$f \leq O\left(\xi L^{d-1} \exp\left[-\frac{L}{\xi}\right]\right). \tag{2.46}$$

□

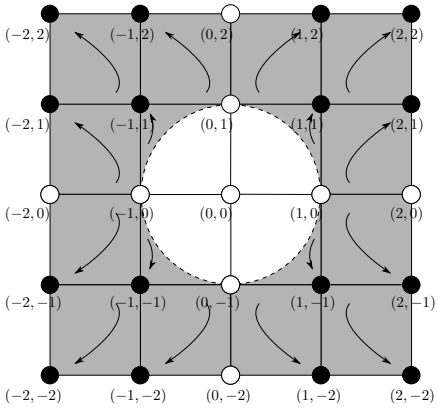

Figure 2.3: Part of the lattice of vectors with integer coordinates. The black dots are the points in the cell with the maximum norm $\|\vec{y}\|$ and the exponential is evaluated at these points. The white ones do not enter the Riemann sum and are related to f_{d-1}, the corresponding quantity in one lower dimension. The arrows show which point in the cell is picked to lower bound the Riemann sum.

Using a similar method, we also get bounds on a related sum.

Lemma 6. *For* $\vec{x} \in \mathbb{Z}^d$, $f_d := \sum_{\|\vec{x}\|\geq 1} \mathrm{e}^{-2L\|\vec{x}\|/\xi} \leq b_d \exp\left[-\frac{2L}{\xi}\right]$ *for some dimension-dependent constant* b_d.

Proof. We prove the statement by induction on the dimension d. For $d = 1$, the statement is seen to be true since the sum evaluates exactly:

$$f_1 = \sum_{|x|\geq 1} \mathrm{e}^{-2L|x|/\xi} = 2\sum_{x=1}^{\infty} \mathrm{e}^{-2Lx/\xi} = \frac{2\mathrm{e}^{-2L/\xi}}{1-\mathrm{e}^{-2L/\xi}} \leq 2.1 \times \mathrm{e}^{-2L/\xi}. \tag{2.47}$$

For the inductive step, consider the integral $g(d) = \int_{\|\vec{y}\|\geq 1}^{\infty} \mathrm{e}^{-2L\|\vec{y}\|/\xi} \mathrm{d}^d\vec{y}$. This is lower bounded by the Riemann sum represented in Fig. 2.3. The white dots represent vectors with at least one zero coordinate and do not enter the Riemann sum according to this way of dividing the region of integration into cells. In the following, the set of points with at least

one zero coordinate is denoted $_{cc}$. We have:

$$g(d) = \int_{\|\vec{y}\| \geq 1}^{\infty} \mathrm{e}^{-2L\|\vec{y}\|/\xi} \mathrm{d}^d \vec{y} \geq \sum_{\vec{x} \notin {}_{cc}} \mathrm{e}^{-2L\|\vec{x}\|/\xi} \Delta_{\vec{x}}, \tag{2.48}$$

where $\Delta_{\vec{x}}$ is the volume of the cell associated with the lattice vector $\vec{x}$. In Fig. 2.3, the volume of most cells (whose center is at distance 1.5 or beyond from the origin) is 1. The cells near the origin have some volume $\alpha_d < 1$ that depends on the dimension. Lower bounding all volumes $\Delta_{\vec{x}}$ by α_d,

$$\sum_{\vec{x} \notin {}_{cc}} \mathrm{e}^{-2L\|\vec{x}\|/\xi} < \frac{g(d)}{\alpha_d} \tag{2.49}$$

$$= \frac{2\pi^{d/2}}{\alpha_d \Gamma(\frac{d}{2})} \left(\frac{\xi}{2L}\right)^d \Gamma\left(d, \frac{2L}{\xi}\right). \tag{2.50}$$

Now it remains to upper bound the contribution from summing over the points $_{cc}$. Notice that the sum over these points is upper bounded by the sum over d hyperplanes of dimension $d-1$. From the inductive hypothesis,

$$\sum_{\vec{x} \in {}_{cc}} \mathrm{e}^{-2L\|\vec{x}\|/\xi} \leq d f_{d-1} \tag{2.51}$$

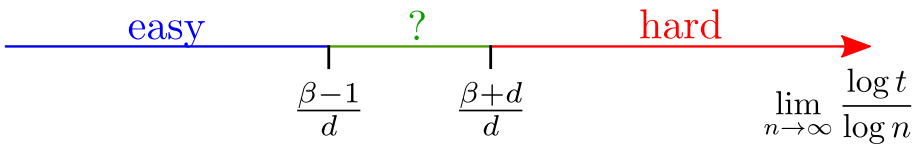

Figure 2.4: Complexity phase diagram of free bosons, illustrating that sampling complexity can delineate boundaries of a physical system as a function of system parameters, including time. Our results indicate that $\frac{\beta-1}{d} \leq c \leq \frac{\beta+d}{d}$, where c is the transition point for the scaling exponent of t with n.

since there are d hyperplanes of dimension $d-1$. Adding Eqs. (2.50) and (2.51), we get

$$\sum_{\vec{x}} \mathrm{e}^{-2L\|\vec{x}\|/\xi} = f_d < d f_{d-1} + \frac{2\pi^{d/2}}{\alpha_d \Gamma(\frac{d}{2})} \left(\frac{\xi}{2L}\right)^d \Gamma\left(d, \frac{2L}{\xi}\right). \tag{2.52}$$

$$\leq d b_{d-1} \exp\left[-\frac{2L}{\xi}\right] + \frac{2\pi^{d/2}}{\alpha_d \Gamma(\frac{d}{2})} \left(\frac{\xi}{2L}\right) \exp\left[-\frac{2L}{\xi}\right] \tag{2.53}$$

$$f_d < b_d \exp\left[-\frac{2L}{\xi}\right], \tag{2.54}$$

proving the lemma. In the second line we have expanded the incomplete Γ function for large L/ξ. □

2.8 Outlook

We have defined the sampling problem for local Hamiltonian dynamics and given upper and lower bounds for the scaling of time $t(n)$ with the number of bosons n for which the problem is efficiently simulable or hard to classically simulate, respectively. Our results are captured in Fig. 2.4 that illustrates the complexity phase diagram of the system. For time-independent systems, we observe that boson sampling is classically easy for all times if the system is Anderson-localized. In a future work [92], we show that there is a class of static, local Hamiltonians that generate a hard-to-sample output distribution at some time that is

not formally infinity. This means that sampling complexity distinguishes Anderson-localized and delocalized systems, which makes it similar to an order parameter that distinguishes different phases.

The case with nonzero Lieb-Robinson velocity [Eq. (2.2)] shows two regimes of the scaling of t with n where sampling is provably easy/hard. We have shown that sampling is easy when $t \leq t_{\text{easy}} = \Theta(n^{\frac{\beta-1}{d}})$ and hard when $t \geq t_{\text{hard}} = \Theta(n^{1+\frac{\beta}{d}})$. Since our definitions of easiness and hardness are exhaustive, we argue that there must exist a constant c such that sampling is efficient for $t < \Theta(n^c)$ and hard otherwise, illustrating a dynamical phase transition. Our proof implies that $c \in [\frac{\beta-1}{d}, \frac{\beta+d}{d}]$ and we show in the next chapter (based on Ref. [93]) that c is given by $\frac{\beta-1}{d}$. The transition is between two regimes, one for short times in which the system's dynamics is essentially indistinguishable from classical dynamics; and the other in which quantum mechanical effects dominate to such an extent as to forbid an efficient classical simulation.

This result may be viewed as a generalization of a similar result in Ref. [94], where it was shown that exact boson sampling for depth-4 circuits is hard. The results there are not directly comparable to ours, since Ref. [94] assumes $\beta = 1$, whereas our results need $\beta > 1$ (easiness) and $\beta \geq 2$ or 5 (hardness). The reason we get easiness even after polynomial time is that we deal with approximate sampling, a less stringent notion of simulation. In addition, adapting the hardness proof of Ref. [94] into our setup, we obtain that the system is hard to exactly sample from at timescale $\Omega(8L/v) = \Omega(n^{\frac{\beta-1}{d}})$ [92], partially answering what happens in the subregion with the question mark in Fig. 2.4. Another work [95] roughly concurrent to ours [96] showed that if the initial state is not well separated ($L = 1$ in our language), the transition timescale is logarithmic in the number of bosons.

Recent studies [76; 97–99] have also studied transitions based on time-dependent features of the out-of-time-ordered correlator. This raises the question of the connection between sampling complexity and scrambling time [100; 101] in quantum many-body systems [102–105] and fast scramblers like the Sachdev-Ye-Kitaev [97; 106–109] models, and black holes [110–114], where one can explore the connection to recent conjectures on complexity in the dual conformal field theory [113–116]. Further, it would be interesting to study sampling complexity in various other physical settings [117], like systems with topological or many-body-localized phases and systems in their ground state to explore the connection with Hamiltonian complexity.

We therefore propose another motivation for considering sampling complexity beyond the field of quantum computational supremacy: complexity provides a natural way to classify phases of matter that is complementary to traditional approaches based on symmetries and topology. This is akin to how the study of entanglement in many-body physics has helped us understand phases of matter [118; 119] and characterize thermalization and localization [120; 121].

Coming to experimental implementations, platforms such as ultracold atoms in optical lattices [122; 123] and superconducting circuits [124; 125] are ideal for experimentally studying the transition by comparing the distribution sampled by the algorithm and the experimental distribution. The ingredients required, which have been realized in several groups [126–131], are: 1. Preparation of the initial state [132; 133] of the type shown in Fig. 2.1[3]. 2. Evolution under a Hamiltonian with either arbitrary time-dependent nearest-

[3]All of our arguments go through even if the initial state is not regular as in Fig. 2.1 but still has the property that the initially occupied sites are separated by a minimum distance of $\Theta(L)$ from each other.

neighbor hopping strength or fixed nearest-neighbor hopping strength $J_{ij}(t) = 1$ together with arbitrary time-dependent on-site potential $J_{ii}(t)$ [132; 134], and 3. Single-site resolved measurement of occupation number of the sites [127; 128]. Cold atoms in quantum gas microscopes can be controlled at the single site level, enabling all three ingredients above. To maintain integrability, we can turn off the Hubbard interaction for the bosonic atoms by tuning to a Feshbach resonance [135–138].

New architectures like optical tweezers [139] are also promising since they allow for deterministic creation of desired initial states and feature tunable interactions in time [140; 141]. Similarly, superconducting circuits have been proposed for quantum simulation of quantum walks [142] and the Bose-Hubbard model [143–145]. The gmon qubit architecture, which was used in Ref. [145], naturally allows for time-dependent variation of coupling strengths [146], and raises the prospect of an experimental study of the transition.

Chapter 3: Complexity phase diagram for interacting and long-range bosonic Hamiltonians

A major effort in quantum computing is to find examples of quantum speedups over classical algorithms, despite the absence of general principles characterizing such a speedup. The study of classical simulability of quantum systems evolving in time allows one to identify features underlying a quantum advantage. Studying the classical simulability of both quantum circuits [11–13; 18; 19; 50–52; 147–152] and Hamiltonians [153; 154], especially under restrictions such as spatial locality [24; 27; 95; 96; 155], allows one to understand the classical-quantum divide in terms of their respective computational complexity. Computational complexity has been closely linked to phases of matter in contexts such as dynamical phase transitions [96], measurement based quantum computing [156], thermal phase transitions [157], and entanglement phase transitions [158].

In this Chapter, we characterize the worst-case computational complexity of simulating time evolution under bosonic Hamiltonians and study a dynamical phase transition in approximate sampling complexity [95; 96]. Previous work [96] studied free bosons with nearest-neighbor hopping but did not consider the robustness of the transition to perturbations in the Hamiltonian, a crucial question in the study of any phase transition. We generalize Ref. [96] to include number-conserving interactions and long-range hops and conclude

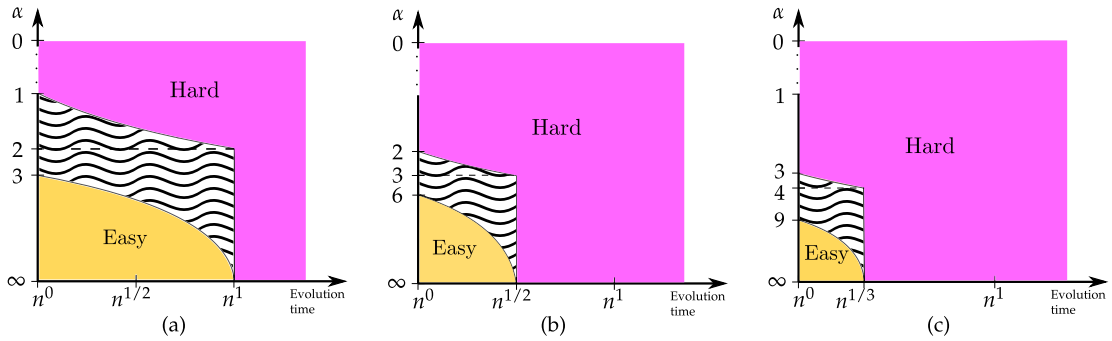

Figure 3.1: Slices of the complexity phase diagram for the long-range bosonic Hamiltonian in a) 1D, b) 2D, and c) 3D with n bosons when the number of sites is $m = \Theta(n^2)$. Colors represent whether the sampling problem is easy (yellow), hard (magenta), or not currently known (hatched). The X-axis parametrizes the evolution time as a polynomial function of n, and the Y-axis is α, the exponent characterizing the long-range nature of the hopping Hamiltonian (with scale $y = 1/\sqrt{\alpha}$ except for the point $\alpha = 0$).

that the phase transition survives under perturbations. These kinds of interactions are ubiquitous in experimental implementations of hopping Hamiltonians with ultracold atoms and superconducting circuits [159; 160]. Long-range hops that fall off as a power law are also native to several architectures [161–165]. We study the location of the phase transition and its dependence on various system parameters, constructing a *complexity phase diagram*, a slice of which is presented in Fig. 3.1.

3.1 Setup and summary of results

Consider a system of n bosons hopping on a cubic lattice of m sites in D dimensions with real-space bosonic operators a_j. We let $m = \Theta(n^\beta)$ and assume sparse filling: $\beta \geq 1$. The Hamiltonian $H = \sum_{i,j} J_{ij}(t) a_i^\dagger a_j + \text{h.c.} + \sum_i f(n_i)$ has on-site interactions $f(n_i)$ and time-dependent hopping terms bounded by a power-law in the distance $d(i,j)$ as $|J_{ij}(t)| \leq 1/d(i,j)^\alpha$. The parameter α governs the degree of locality. When $\alpha = 0$, the system has all-to-all couplings, while $\alpha \to \infty$ corresponds to nearest-neighbor

hops. The on-site terms $J_{ii}(t)$ can be large, and the interaction strength is parametrized by V. For concreteness, our hardness results are derived using a Bose-Hubbard interaction $f(n_i) = Vn_i(n_i - 1)/2$, but the timescales we present are valid for generic on-site interactions [166]. The bosons in the initial states considered are sparse and well-separated. Specifically, partition the lattice into K clusters $C_1, \ldots, C_K$ containing $b_1, \ldots, b_K$ initial bosons, respectively, such that $b := \max b_i = O(1)$ does not scale with lattice size. Define the width L_i of a cluster C_i as the minimum distance between a site outside the cluster and an initially occupied site inside the cluster and let $L = \min_i L_i$. As in Ref. [96], we consider states with the bosons roughly equally spaced throughout the lattice, so that $m = \Theta(nL^D)$, giving $L = \Theta(n^{(\beta-1)/D})$.

The computational task of approximate sampling is to simulate projective measurements of the time-evolved state in the local boson-number basis. The *approximate sampling complexity* measures the classical resources needed to produce samples from a distribution $\tilde{\mathcal{D}}$ that is $\epsilon = O(1/\mathsf{poly}(n))$-close in total variation distance to the target distribution $\mathcal{D}$[1]. Sampling from a distribution $\tilde{\mathcal{D}}$ satisfying the above takes runtime $T(n, t)$ in the worst case on a classical computer, where t is the evolution time. Like thermodynamic quantities, the complexity is defined asymptotically as $n \to \infty$, so we consider the scaling of T along a curve $t(n)$. Along any curve $t(n) = cn^\gamma$, sampling is *easy* if there exists a polynomial-runtime classical algorithm for all n, or *hard* if such an algorithm cannot exist. Since the problem is either easy or hard for a particular function $t(n)$, there is always a transition in complexity as opposed to a smooth crossover. We prove upper and lower bounds on the

[1]This is a weaker requirement than demanding a classical sampler that works for arbitrary ϵ and has runtime $\mathrm{poly}(n, 1/\epsilon)$.

α	Easiness exponent	Hardness exponent[4]
$\alpha < \frac{D}{2}$	?	$\frac{\beta}{D}\left(\alpha - \frac{D}{2}\right)$
$\frac{D}{2} \leq \alpha \leq D$		0
$D \leq \alpha \leq D+1$		$\frac{\beta-1}{D}(\alpha - D)$
$D+1 \leq \alpha \leq 2D + \frac{D}{\beta-1}$	0[5]	$\frac{\beta-1}{D}$
$2D + \frac{D}{\beta-1} \leq \alpha < \infty$	$\frac{\beta-1}{D}\left(\frac{\alpha-2D}{\alpha-D}\right) - \frac{1}{\alpha-D}$	
$\alpha \to \infty$	$\frac{\beta-1}{D}$	$\frac{\beta-1}{D}$ if $D \geq 2$ or $V < \infty$ ∞ otherwise.

Table 3.1: Leading exponents γ_{easy} and γ_{hard} in the easiness and hardness timescales for various regimes of α. Some of these transitions are likely sharp and some coarse.

transition timescale by presenting sampling algorithms on the easiness side, and performing reductions to quantum supremacy proposals on the hardness side. Specifically, we show that approximate sampling is easy for all times $t < c_{\text{easy}} n^{\gamma_{\text{easy}}}$ and hard for all times $t > c_{\text{hard}} n^{\gamma_{\text{easy}}}$.

We find that the transition comes in two types, which we call "sharp" and "coarse". For sharp transitions, these bounds coincide in the exponent $\gamma_{\text{easy}} = \gamma_{\text{hard}}$ and the transition occurs in the coefficient[2] $c_{\text{easy}} \leq c_{\text{hard}}$. For coarse transitions, however, the transition occurs in the exponent. In our results, we will show that sampling is easy for any time $t = O(n^{\gamma_{\text{easy}}})$, but hard along any curve with exponent $\gamma_{\text{hard}} > \gamma_{\text{easy}}$[3] (see [167] for more precise definitions). An example of a sharp transition is when the transition timescale is $t_* = 2n$, so that the problem is easy for all times $t \leq 1.99n$ and hard for all times $t \geq 2.01n$. An example of a coarse transition is when the transition timescale is $t_* = \Theta(n \log n)$, so that the problem is easy for all times $t \leq cn$ and hard for all times $t \geq cn^{1.01}$.

We summarize our main result in Table 3.1. The easiness result comes from applying

[2] More precisely, for sharp transitions, we have $t_{\text{hard}} = \Theta(t_{\text{easy}})$, while for coarse transitions, $t_{\text{hard}} = \omega(t_{\text{easy}})$.

[3] In principle, there can be other kinds of coarse transitions where sampling is hard for any time $t = \Omega(n^{\gamma_{\text{hard}}})$ and easy for any $\gamma_{\text{easy}} < \gamma_{\text{hard}}$.

[4] Up to an additive constant $\delta > 0$ that is present when indicated in the text.

[5] The easiness timescale for this case is $t_{\text{easy}} = \log n$.

classical algorithms for quantum simulation, and depend on Lieb-Robinson bounds on information transport [77; 168–171]. The hardness results come from reductions to families of quantum circuits for which efficient approximate samplers cannot exist, modulo widely believed conjectures in complexity theory [13; 24; 27; 155], and from fast protocols to transmit quantum information across long distances [172; 173].

Note that the hardness exponents in Table 3.1 sometimes come with an arbitrarily small additive term $\delta > 0$. This happens whenever at least one of the following cases hold: $\alpha < D/2$, $V = o(1)$, or $D = 1$. When the easiness and hardness timescales coincide, we interpret the presence of this term δ as signifying a coarse transition, since it ensures $\gamma_{\text{hard}} > \gamma_{\text{easy}}$.

We examine the various limits: $\alpha \to \infty$ (nearest-neighbor), $\alpha \to 0$ (all-to-all connectivity), $V \to 0$ (free bosons), and $V \to \infty$ (hardcore bosons). First, when $\alpha \to \infty$, the hardness timescale upper bound is $O(L)$ in all cases except when $V \to \infty$ and $D = 1$, which we discuss later. This matches the easiness timescale $t = \Omega(L)$, which corresponds to the distance L between clusters. We therefore pin down the transition timescale to $\Theta(L)$, which is when interference between clusters become relevant [96]. In the opposite limit when the model is sufficiently long-range ($\alpha < D/2$), the role of the dimension is unimportant, giving $\gamma_{\text{hard}} < 0$ in all cases. This suggests a hardness timescale close to 0, signifying the immediate onset of hardness.

Next, we observe that the transition timescale does not heavily depend on the interaction strength V, showing that the location of the complexity phase transition is robust to the presence of interactions. The only exception is the limit of hardcore interactions and nearest-neighbor hops ($V, \alpha \to \infty$) in 1D. In this case, there is no hardness regime. This is

α	Easiness exponent	Hardness exponent
$\alpha < \frac{D}{2}$	?	$\frac{1}{D}\left(\alpha - \frac{D}{2}\right)$
$\frac{D}{2} \leq \alpha \leq D+1$		0
$D+1 \leq \alpha < \infty$	$\frac{-1}{\alpha - D}$	
$\alpha \to \infty$	0	0 if $D \geq 2$ or $V < \infty$ ∞ otherwise.

Table 3.2: Easiness and hardness timescale exponents when $\beta = 1$.

because the model maps to that of free fermions, or equivalently, matchgate circuits, which are easy to simulate at all times [50; 51].

The results for the regime $\beta = 1$, or $m = \Theta(n)$, can be experimentally relevant and are shown in Table 3.2. In this case, the separation between clusters is a constant, and when $\alpha \geq D+1$, the algorithm works only for time $t = O(n^{-1/(\alpha - D)})$, which asymptotes to 0 as $n \to \infty$. For $\alpha \to \infty$, this is (almost) matched by the hardness timescale being $t_{\text{hard}} = O(1)$, since the leading exponents in the timescales are zero in both cases. We now outline the proofs of our results, whose details may be found in Sections 3.6, 3.7, 3.9 and 3.10.

3.2 Easy-sampling timescale

To derive t_{easy}, we give an efficient sampling algorithm. The algorithm performs time evolution on each cluster C_i separately. This takes polynomial time in the number of basis states, which is $\binom{|C_i|+b_i-1}{b_i} = O(|C_i|^{b_i})$ and hence polynomial in n when $b_i = O(1)$. This product-state approximation of the exact time-evolved state $|\psi(t)\rangle = U_t |\psi(0)\rangle$ is achieved by decomposing the propagator U_t via a spatial decomposition scheme for quantum simulation [44; 171] that we call the HHKL decomposition for the authors of Ref. [44]. We complete the derivation of the easiness timescale by showing that the approximation is good for times

$t < O(t_{\text{easy}})$.

We briefly present the HHKL decomposition, which is powerful but remarkably simple. Let H_R be the sum over all terms in the Hamiltonian supported completely in region R and implicitly let $XY = X \cup Y$ represent the union of regions. The decomposition scheme approximates the time evolution unitary acting on region XYZ (where Y separates regions X and Z) by forward evolution on YZ, backward evolution on Y, and forward evolution on XY: $U_{XYZ} \approx U_{XY}(U_Y)^\dagger U_{YZ}$. The operator norm error made by this approximation is [171] $O\left((e^{vt}-1)\Phi(X)(\ell^{-\alpha+D+1}+e^{-\ell})\right)$, where $v>0$ is a characteristic velocity, $\Phi(X)$ is the area of the boundary of X, and ℓ is the minimum distance between any pair of sites in X and Z. The error is small for times t shorter than the time it takes for information to propagate from X to Z.

The velocity v of information propagation is also known as a Lieb-Robinson velocity and is determined by the operator norm of terms in the Hamiltonian which couple different sites [168]. Since bosonic operators have unbounded operator norm, this could result in an unbounded velocity [84]. However, because of boson number conservation under the Hamiltonian, the dynamics is fully contained in the n-boson subspace, within which the operator norm of each term is $O(n)$. While free bosons ($V=0$) behave as in the single-particle subspace, implying the Lieb-Robinson velocity is $O(1)$, in the interacting case, an $O(n)$ Lieb-Robinson velocity would cause the asymptotic easiness timescale to vanish ($t_{\text{easy}} \to 0$).

Nevertheless, we can derive an easiness timescale independent of V for a clustered initial state. Intuitively, at short times each boson is well-localized within its original cluster. Therefore, the relevant subspace has at most b bosons in each cluster C_i. Truncating the Hilbert space to allow only $b+1$ bosons per cluster is therefore a good approximation at

short times [174], and the truncation error vanishes in the asymptotic limit. The modified Hamiltonian H' after truncation has terms with norm only $O(b)$, giving an effective Lieb-Robinson velocity $v = O(b) = O(1)$ for states close to the initial state[6]. For this modified Hamiltonian, we apply the HHKL decomposition to bound the error caused by simulating each cluster separately. Once the error has been calculated, the timescale immediately follows by solving $\epsilon(t) = O(1)$ for $t = t_{\text{easy}}$, which is a lower bound on the transition timescale t_*. In Section 3.7, we give the full dependence of t_{easy} on various system parameters, including the filling fraction of bosons.

3.3 Sampling hardness timescale

To derive t_{hard}, we give protocols to simulate quantum circuits by setting the time dependent parameters $J_{ij}(t)$ of the long-range bosonic Hamiltonian. This implies sampling is worst-case hard after time t_{hard}. Specifically, if a general sampling algorithm exists for times $t \geq t_{\text{hard}}$, we prove this algorithm can also simulate hard instances of boson sampling [13] when interactions are weak, and quantum circuits that are hard to simulate [24] when interactions are strong.

In the interacting case, our reduction from universal quantum computation to a long-range Hamiltonian hinges on implementing a universal gate set. Using a dual-rail encoding to encode a qubit in two modes of each cluster C_i, we show in Section 3.9 how to implement arbitrary single-qubit operations in $O(1)$ time and controlled-phase gates [175] between adjacent clusters in a time that depends on their spacing L. The two-qubit gate uses free particle state-transfer as a subroutine [172; 173] to bring adjacent logical qubits near each

[6] In other words, this is a state-dependent Lieb-Robinson velocity, or a butterfly velocity.

other. We implement the constant-depth circuit of Ref. [24], which consists only of nearest-neighbor gates between qubits in a 2D grid. The total time for hardness under this scheme takes time $O(\min[L\,, L^{\alpha-D}])$ when $\alpha > D$ and $O(1)$ when $\alpha \in [D/2, D]$. In 1D, simulating a 2D circuit introduces extra overhead. Nevertheless, we can recover the same timescale up to an infinitesimal $\delta > 0$ in the exponent by only encoding n^δ logical qubits. For hardcore bosons, the above scheme mentioned does not work and the entangling gate is constructed differently, and features an easiness result for the 1D nearest-neighbor case. Lastly, when $\alpha < D\,/\,2$, state transfer takes time $o(1)$ but the time for an entangling gate is $O(1)$. We can still achieve coarse hardness for time $o(1)$ by mapping the system onto free bosons, which we now come to.

In the noninteracting case, we implement the boson sampling scheme of Ref. [13], which showed that a Haar-random unitary applied to m sites containing n bosons gives a hard-to-sample state. It also gave an $O(n \log m)$-depth decomposition of a linear-optical unitary in the circuit model without spatial locality. In Section 3.10, we give a faster implementation for the continuous-time Hamiltonian model, which can include simultaneous noncommuting terms but imposes spatial locality, a result of independent interest. Specifically, we show that most linear-optical states of n bosons on m sites can be constructed in time $\min[O(nm^{1/D})\,, \tilde{O}(nm^{\alpha/D-1/2}])$, which is faster than the circuit model when $\alpha < D\,/\,2$. This result also uses free-particle state transfer as a subroutine. As in the 1D interacting case, we can implement the reduction on a polynomially growing number of bosons n^δ, resulting in the timescale of Table 3.1 for free bosons. This result resolves an important conceptual question posed by Ref. [96] for the noninteracting, nearest-neighbor case by closing the gap between t_{easy} and t_{hard}. In this limit, the transition timescale is at $\Theta(L/v)$,

both with and without interactions, showing that the algorithm of Ref. [96] is optimal and that the presence of interactions does not change the phase diagram.

3.4 Sharp and coarse transitions

In the nearest-neighbor limit $\alpha \to \infty$, the exponents on our hardness and easiness timescales match up to an infinitesimal (δ), so we can make precise statements about the nature of the transition. In the presence of interactions and in two dimensions and above, the bounds on the timescale in the nearest-neighbor limit coincide up to a multiplicative constant at $t_{\text{easy}} = t_{\text{hard}} = \Theta(L)$, proving the transition is sharp. However, in 1D, the hardness timescale only matches up to an infinitesimal $\delta > 0$ in the exponent $c_{\text{hard}} = c_{\text{easy}} + \delta$. In 1D, we show that c_{hard} cannot be improved any further, proving that this is a coarse transition.

To understand the physics behind the two kinds of transitions, it is illuminating to study the approach to the transition point from both sides. On the easiness side, the important quantity is the many-body entanglement. Recall that at short times, the wavefunction is approximately separable, implying easiness of classical simulation. The separable state is computed using an HHKL decomposition, whose errors grow in time until they become $O(1)$ at the transition timescale. These errors upper bound the amount of entanglement present across any cut, so the easy phase corresponds to states with no entanglement, and the complexity transition occurs as the entanglement grows from zero entanglement to area-law entanglement.

However, sampling complexity in one dimension is special because area-law entanglement is classically simulable using matrix product states [176; 177]. Specifically in 1D, we

prove an extended easiness timescale of $t_{\text{easy}} = cL$ for any constant c. Sampling is easy for all times $O(L)$, implying that the δ in the hardness exponent cannot be removed, and the transition is coarse. We leave open the question of entanglement scaling between the area-law and volume-law regimes. However, if $D \geq 2$, the argument based on entanglement breaks down because tensor-network contraction takes time exponential in the system size in both the worst case and average case [178; 179], and there are known examples of constant-depth 2D circuits that are hard to simulate [24].

On the hardness side, many-body entanglement is necessary but not sufficient for sampling hardness [50; 51; 180]. Since our hardness results in the interacting case rely on mapping bosons to qubits via a dual-rail encoding, we understand the transition by counting the number of encoded logical qubits. For coarse transitions, as the evolution time approaches the transition timescale t^*, the number of encoded logical qubits shrinks as n^δ, where $\delta \to 0$ as $t \to t_*$. This illustrates that while the problem is still asymptotically hard as $n \to \infty$, one needs to go to higher boson numbers n to achieve the same computational complexity. For sharp transitions on the other hand, the number of encoded logical qubits seems to suddenly jump to $O(n^c)$ for some constant c. In Section 3.11, we elaborate on how to use number of effective logical qubits as an order parameter for the phase transition. Whether or not such an order parameter is a universal way to characterize complexity phase transitions deserves closer attention.

Finally, in the noninteracting case, we cannot definitively prove that the transition is coarse in $D \geq 2$. However, the infinitesimal in the hardness timescale is suggestive of a coarse transition. More precisely, our work implies that either the transition is coarse, or there exists a constant-depth boson sampling circuit on a nearest-neighbor architecture for

which approximate sampling is classically hard. It has been proved by Brod [94] that exact sampling of constant-depth boson sampling is classically hard. Nevertheless, Brod points out that it is very unclear if this exact sampling hardness can be extended to the approximate sampling hardness we consider in this Chapter. As such, in $D \geq 2$ without interactions, the type of transition is an open problem.

3.5 Related Models

We first show that our results can be easily adapted to a wide range of experimentally and theoretically interesting Hamiltonians. First, in the hardcore limit, spin Hamiltonians naturally map onto our model. Fermionic systems with nearest-neighbor interactions can also be incorporated by performing the mapping described in Ref. [181]. Our model is also relevant to cold atom experiments that have been proposed as candidates for observing quantum computational supremacy [24; 95; 159; 160], especially in the nearest-neighbor limit.

The power-law hopping $1/r^{\alpha}$ can be engineered to directly implement the classes of Hamiltonians we study. This can be done by virtually coupling the band of interest to another with a quadratic band edge to implement exponentially decaying hopping [165; 182; 183]. Doing this simultaneously with multiple detunings approximates a power-law with high accuracy as a sum of exponentials [184].

In the hardcore limit, the long-range hops translate to long-range interactions between spins, which model quantum-computing platforms such as Rydberg atoms and trapped ions [141; 161; 185–187]. Therefore, the Hamiltonian we study models various physically inter-

esting situations, both in the several limiting cases ($\alpha \to \infty$, $V \to 0$, $V \to \infty$) as well as in the general case of finite nonzero α and V.

Our methods are fairly general for lattice models with long-range power law decaying interactions. However, we do require number-conservation so the local Hilbert space dimension can be bounded at short times. As an example of how to extend our results to different models, we can straightforwardly incorporate long-range density-density interactions $K_{ij}(t)\, n_i n_j$. The only effect on the easiness times is to modify the Lieb-Robinson velocity to $v = O(b^2)$. This is an overall constant that does not affect the exponent.

Our model can also be used to describe a distributed modular quantum network when the Hubbard interaction V can vary spatially. Specifically, a module of qubits can be represented by hardcore bosons ($V \to \infty$), while photonic communication channels linking distant modules can be represented by sites with $V = 0$ separating the modules. As in quantum networks, our hardness times in the nearest-neighbor regime are dominated by gates between nodes, while operations within a single node are free.

3.6 Approximation error under HHKL decomposition

We first argue why it is possible to apply the HHKL decomposition lemma to H' with a Lieb-Robinson velocity of order $O(1)$. As mentioned in Section 3.2, H' is a Hamiltonian that lives in the truncated Hilbert space of at most $b+1$ bosons per cluster. Let Q be a projector onto this subspace. Then $H' = QHQ$. Time-evolution under this modified Hamiltonian H' keeps a state within the subspace since $\left[e^{-iQHQt}, Q\right] = 0$.

The Lieb-Robinson velocity only depends on the norm of terms in the Hamiltonian

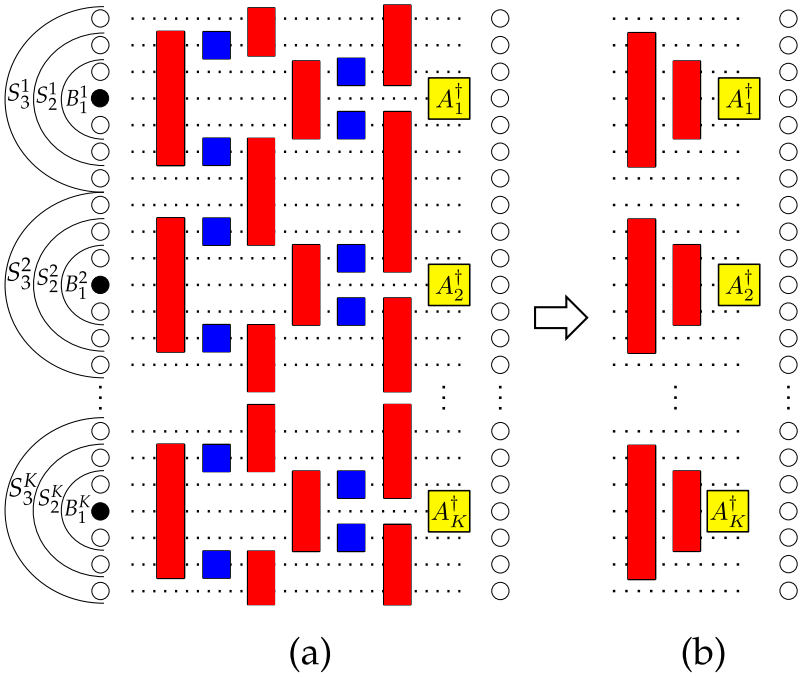

Figure 3.2: (a) Decomposition of the first two steps of the unitary evolution followed by (b) pushing the commuting terms past $A_i^\dagger$ (the product of all initial creation operators in a cluster i) to the vacuum. Red boxes represent forward evolution and blue boxes backward evolution in time.

which couple lattice sites. On-site terms do not contribute, which can be seen by moving to an interaction picture [168; 170]. Therefore, since no state has more than $b+1$ bosons on any site within the image of Q, the maximum norm of coupling terms in H' is $\left\| Q a_i^\dagger a_j Q \right\| \leq b+1$. Therefore, the Lieb-Robinson velocity is at most $O(b)$ instead of $O(n)$, and we can apply the HHKL decomposition to the evolution generated by the truncated Hamiltonian H'. We now prove that the error made by decomposing the evolution due to H' is small.

Lemma 7 (Decomposition error for H')**.** *For all V and $\alpha > D+1$, the error incurred (in 2-norm) by decomposing the evolution due to H' for time t is*

$$\epsilon(t) \leq O\left(K(e^{vt_1} - 1)(\ell^{-\alpha+D+1} + e^{-\ell}) \sum_{j=0}^{N-1} (r_0 + j\ell)^{D-1} \right), \tag{3.1}$$

where $N = t/t_1$ and $\ell \leq L/N$ can be chosen to minimize the error, and r_0 is the radius of the smallest sphere containing the initially occupied bosons in a cluster.

The sketch of the proof is as follows: recall that within each cluster C_i, there is a group of bosons initially separated from the edge of the cluster by a region of width L_i. Naive application of the HHKL decomposition for the long-range case results in a timescale $t_{\text{easy}} \sim \log(L)$, because of the exponential factor $(e^{vt} - 1)$ in the error. To counter this, we apply the HHKL decomposition in small time-steps t_1. Thus, within each time-step, the exponential factor can be approximated as $e^{vt_1} - 1 \approx vt_1$, turning this exponential dependence into a polynomial one at the cost of an increased number of time-steps.

The first two time-steps are depicted pictorially in Fig. 3.2, and illustrate the main ideas. The full propagator acting on the entire lattice is decomposed by applying the HHKL decomposition K times, such that two of every three forward and reverse time-evolution operators commute with all previous operators by virtue of being spatially disjoint, allowing them to be pushed through and act identically on the vacuum. The remaining forward evolution operator effectively spreads out the bosonic operators by distance ℓ. The error per time-step is polynomially suppressed by $O(\ell^{-\alpha+D+1} + e^{-\gamma\ell})$.

While it reduces the exponential factor to a polynomial one, using time-slices comes at the cost of extra polynomial factors, originating from the sum over boundary terms $\sum_{j=0}^{N-1}(r_0 + j\ell)^{D-1}$.

Proof of Lemma 7. Let the initial positions of the bosons be denoted by $(\mathsf{in}_1, \ldots, \mathsf{in}_n)$. The initial state is $|\psi(0)\rangle = a^\dagger_{\mathsf{in}_1} \ldots a^\dagger_{\mathsf{in}_n} |0\rangle$ As before, the first two time-steps are illustrated in Fig. 3.2. Within each cluster C_i, there is a group of bosons initially separated from

the edge of the cluster by a region of width L_i. Let $A_i^\dagger(0) = \prod_{\mathsf{in}_j \in C_i} a_{\mathsf{in}_j}^\dagger$ be the creation operator for the group of bosons in the cluster C_i, so that the initial state can also be expressed as $|\psi(0)\rangle = \prod_{i=1}^K A_i^\dagger(0)\,|0\rangle$. The forward-time propagator on a region R is $U_{t_0,t_1}^R = \mathcal{T}\exp\left(-i\int_{t_0}^{t_1} H_R(s)ds\right)$. When evolved for short times, each creation operator $a_{\mathsf{in}_i}^\dagger(t)$ is mostly supported over a small region around its initial position. Therefore, as long as these regions do not overlap, each operator approximately commutes, and the state is approximately separable.

Let A^i be the smallest ball upon which $A_i^\dagger(0)$ is supported. Let $B_0^i = A^i$ and denote its radius r_0^i, and define $r_0 = \max r_0^i$. B_k^i is a ball of radius $r_0^i + k\ell$ containing A^i, where ℓ will be chosen to minimize the error. S_k^i is the shell $B_k^i \setminus B_{k-1}^i$ (see Fig. 3.2). The complement of a set X is denoted as X^c. We divide the evolution into N time steps between $t_0 = 0$ and $t_N = t$, and first show that the evolution is well-controlled for one time step from 0 to $t_1 = t/N$. We apply this decomposition K times, once for each cluster, letting $X = B_0^i$, $Y = S_1^i$ and Z be everything else:

$$U_{0,t_1} \approx U_{0,t_1}^{B_1^1}(U_{0,t_1}^{S_1^1})^\dagger U_{0,t_1}^{(B_0^1)^c} \tag{3.2}$$

$$\approx U_{0,t_1}^{B_1^1}(U_{0,t_1}^{S_1^1})^\dagger U_{0,t_1}^{B_1^2}(U_{0,t_1}^{S_1^2})^\dagger U_{0,t_1}^{(B_0^1 B_0^2)^c} \tag{3.3}$$

$$\approx U_{0,t_1}^{B_1^1}(U_{0,t_1}^{S_1^1})^\dagger \ldots U_{0,t_1}^{B_1^K}(U_{0,t_1}^{S_1^K})^\dagger U_{0,t_1}^{(B_0^1 \ldots B_0^K)^c}. \tag{3.4}$$

The total error is $O\left(\sum_{i=1}^K (e^{vt_1}-1)\Phi(B_0^i)(\ell^{-\alpha+D+1}+e^{-\gamma\ell})\right) = O\left(K(e^{vt_1}-1)r_0^{D-1}(\ell^{-\alpha+D+1}+e^{-\gamma\ell})\right)$. Applying the decomposed unitary to the initial state and pushing commuting terms through

to the vacuum state, we get

$$U_{0,t_1} |\psi(0)\rangle \approx U^{B_1^1}_{0,t_1} A_1^\dagger \dots U^{B_1^K}_{0,t_1} A_K^\dagger |0\rangle = \left(\prod_{i=1}^{K} U^{B_1^i}_{0,t_1} A_i^\dagger \right) |0\rangle .$$

We can repeat the procedure for the unitary U_{t_1,t_2}, where $t_2 = 2t_1$. Now, the separating region Y will be S_2^i, so that $S_2^i \cap B_1^i = \emptyset$. Each such region still has width ℓ, but now the boundary of the interior is $\Phi(B_1^i) = O((r_0 + \ell)^{D-1})$. We get

$$U_{t_1,t_2} \approx \left(\prod_{i=1}^{K} U^{B_2^i}_{t_1,t_2} (U^{S_2^i}_{t_1,t_2})^\dagger \right) U^{(B_1^1 \dots B_1^K)^c}_{t_1,t_2}, \tag{3.5}$$

with error $O(K(e^{vt_1} - 1)(r_0 + l)^{D-1}(\ell^{-\alpha+D+1} + e^{-\gamma\ell}))$. The unitaries supported on S_2^i and $(B_1^1 \dots B_1^K)^c$ commute with all the creation operators supported on sites B_1^i, giving $|\psi(t_2)\rangle \approx U^{B_2^1}_{t_1,t_2} U^{B_1^1}_{0,t_1} \dots U^{B_2^K}_{t_1,t_2} U^{B_1^K}_{0,t_1} |\psi(0)\rangle$. By applying this procedure a total of N times, once for each time step, we get the approximation

$U_{0,t_N} |\psi(0)\rangle \approx U^{B_N^1}_{t_{N-1},t_N} \dots U^{B_1^1}_{0,t_1} \dots U^{B_N^K}_{t_{N-1},t_N} \dots U^{B_1^K}_{0,t_1} |\psi(0)\rangle$. The total error in the state (in 2-norm) is

$$\epsilon \leq O\left(K(e^{vt_1} - 1)(\ell^{-\alpha+D+1} + e^{-\gamma\ell}) \sum_{j=0}^{N-1} (r_0 + j\ell)^{D-1} \right) \tag{3.6}$$

$$= O\left(n(e^{vt_1} - 1)(\ell^{-\alpha+D+1} + e^{-\gamma\ell}) N L^{D-1} \right), \tag{3.7}$$

proving Lemma 7. The last inequality comes from the fact that $K \leq n$ and that $r_0 + (N - 1)\ell \leq \min L_i = L$. The latter condition ensures that the decomposition of the full unitary is separable on the clusters.

In the regime $\alpha > 2D + D/(\beta - 1)$, t_{easy} is optimized by choosing a fixed time-step size $t_1 = O(1)$. Then, the number of steps N scales as the evolution time $N = t/t_1$. By the last few time-steps, the bosonic operators have spread out and have a boundary of size L^{D-1}, so the boundary terms contribute $O(NL^{D-1})$ in total. In the regime $D + 1 < \alpha \leq 2D + D/(\beta - 1)$, the boundary contribution outweighs the suppression $\ell^{-\alpha+D+1}$. Instead, we use a single time-step in this regime, resulting in $t_{\text{easy}} = \Omega(\log n)$ when $\beta > 1$.

□

3.7 Closeness of evolution under H and H'.

Next, we show that the states evolving due to H and H' are close, owing to the way the truncation works. This will enable us to prove that the easiness timescale for H is the same as that of H'. Suppose that an initial state $|\psi(0)\rangle$ evolves under two different Hamiltonians $H(t)$ and $H'(t)$ for time t, giving the states $|\psi(t)\rangle = U_t |\psi(0)\rangle$ and $|\psi'(t)\rangle = U'_t |\psi(0)\rangle$, respectively. Define $|\delta(t)\rangle = |\psi(t)\rangle - |\psi'(t)\rangle$ and switch to the rotating frame, $|\delta^r(t)\rangle = U_t^\dagger |\delta(t)\rangle = |\psi(0)\rangle - U_t^\dagger U'_t |\psi(0)\rangle$. Now taking the derivative,

$$i\partial_t |\delta^r(t)\rangle = 0 + U_t^\dagger H(t) U'_t |\psi(0)\rangle - U_t^\dagger H'(t) U'_t |\psi(0)\rangle \tag{3.8}$$

$$= U_t^\dagger (H(t) - H'(t)) |\psi'(t)\rangle . \tag{3.9}$$

The first line comes about because $i\partial_t U'_t = H'(t)U'_t$ and $i\partial_t U_t^\dagger = -U_t^\dagger H(t)$, owing to the time-ordered form of U_t.

Now, we can bound the norm of the distance, $\delta(t) := \||\delta(t)\rangle\| = \||\delta^r(t)\rangle\|$.

$$\delta(t) \leq \delta(0) + \int_0^t d\tau \|(H(\tau) - H'(\tau)) |\psi'(\tau)\rangle\| \tag{3.10}$$

$$= \int_0^t d\tau \|(H(\tau) - H'(\tau)) |\psi'(\tau)\rangle\|, \tag{3.11}$$

since $\delta(0) = 0$.

The next step is to bound the norm of $(H - H') |\psi'(\tau)\rangle$ (we suppress the time label τ in the argument of H and H' here and below). We use the HHKL decomposition: $|\psi'(\tau)\rangle = |\phi(\tau)\rangle + |\epsilon(\tau)\rangle$, where the state $|\phi(\tau)\rangle$ is a product state over clusters, and $|\epsilon(\tau)\rangle$ is the error induced by the decomposition. We first show that $(H - H') |\phi(\tau)\rangle = 0$. Since $|\phi(\tau)\rangle$ is a product state of clusters, each of which is time-evolved separately, boson number is conserved within each cluster. Therefore, each cluster has at most b bosons, and $Q |\phi(\tau)\rangle = |\phi(\tau)\rangle$. Furthermore, only the hopping terms in H can change the boson number distribution among the different clusters, and these terms move single bosons. This implies that $H |\phi(\tau)\rangle$ has at most $b+1$ bosons per cluster, and remains within the image of Q, denoted im Q. Combining these observations, we get $H' |\phi(\tau)\rangle = QHQ |\phi(\tau)\rangle = H |\phi(\tau)\rangle$. This enables us to say that $(H - H') |\phi(\tau)\rangle = (H - QHQ) |\phi(\tau)\rangle = 0$. Equation (3.11) gives us

$$\delta(t) \leq \int_0^t d\tau \|(H(\tau) - H'(\tau))(|\phi(\tau)\rangle + |\epsilon(\tau)\rangle)\| \tag{3.12}$$

$$= \int_0^t d\tau \|(H(\tau) - H'(\tau)) |\epsilon(\tau)\rangle\|, \tag{3.13}$$

$$\leq \max_{\tau, |\eta\rangle \in \text{im } Q} \|(H(\tau) - H'(\tau)) |\eta\rangle\| \int_0^t d\tau \||\epsilon(\tau)\rangle\|. \tag{3.14}$$

In the last inequality, we have upper bounded $\|(H(\tau) - H'(\tau)) |\epsilon(\tau)\rangle\|$ by $\max_{|\eta\rangle \in \text{im } Q} \|(H - H') |\eta\rangle\| \times$

$\epsilon(\tau)$, where $\epsilon(\tau) := \||\epsilon(\tau)\rangle\|$. The quantity $\max_{|\eta\rangle\in \mathrm{im}\ Q} \|(H - H')\,|\eta\rangle\|$ can be thought of as an operator norm of $H - H'$, restricted to the image of Q. It is enough to consider a maximization over states $|\eta\rangle$ in the image of Q because we know that the error term $|\epsilon(\tau)\rangle$ also belongs to this subspace, as $|\psi'(\tau)\rangle$ belongs to this subspace. Further, we give a uniform (time-independent) bound on this operator norm, which accounts for the maximization over times τ.

Lemma 8. $\max_{|\eta\rangle} \|(H - QHQ)\,|\eta\rangle\| \leq \left\|\sum_{i\in C_k, j\in C_l} J_{ij} a_i^\dagger a_j\right\| \leq O(bL^{D-\alpha})$.

Proof. Notice that for each term H_i in the Hamiltonian, the operator $H - QHQ$ contains $H_i - QH_iQ$, where the rightmost Q can be neglected since $Q\,|\eta\rangle = |\eta\rangle$. The on-site terms $\sum_i J_{ii} a_i^\dagger a_i + V n_i(n_i - 1)/2$ do not change the boson number. Therefore, they cannot take $|\eta\rangle$ outside the image of Q, and do not contribute to $(H - QHQ)\,|\eta\rangle$. The only contribution comes from hopping terms that change boson number, which we bound by

$$\left\| \sum_{i\in C_k, j\in C_l} J_{ij} a_i^\dagger a_j \right\|, \tag{3.15}$$

where the sum is over sites i and j in distinct clusters C_k and C_l, respectively. This is because only hopping terms that connect different clusters can bring $|\eta\rangle$ outside the image of Q, since hopping terms within a single cluster maintain the number of bosons per cluster.

For illustration, let us focus on terms that couple two clusters C_1 and C_2. The distance between these two clusters is denoted L_{12}. For any coupling J_{ij} with $i \in C_1$ and $j \in C_2$,

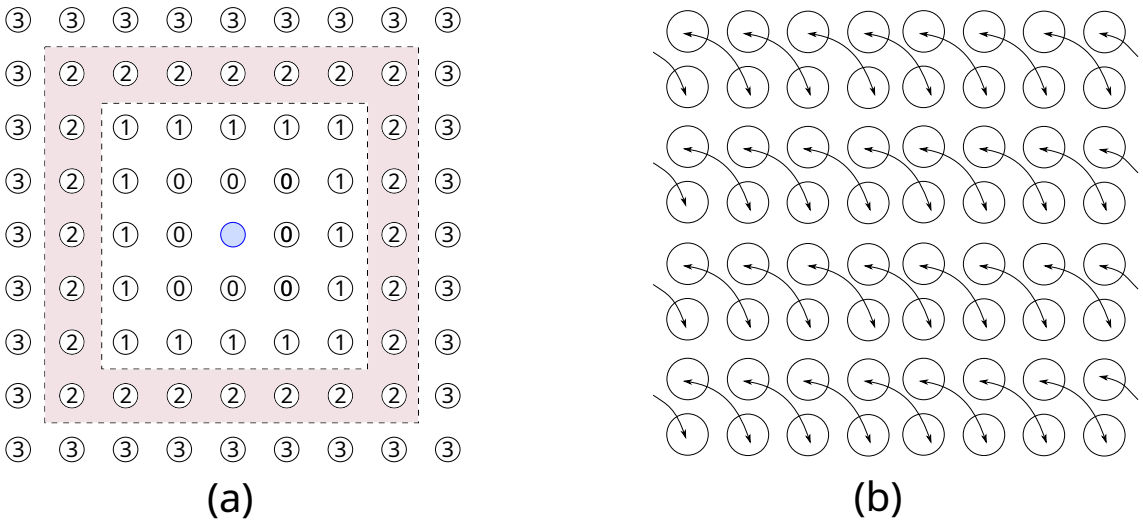

Figure 3.3: (a) Cluster distances between the blue cluster in the center and nearby clusters. The total number of clusters at cluster distance $l = 2$ (pink background) is given by $(2l + 3)^2 - (2l+1)^2 = 24$. (b) Non-overlapping pairing between clusters separated by a diagonal. The distance between these clusters is $l = 0$, since they share a boundary and contain adjacent sites.

we can bound $|J_{ij}| \leq L_{12}^{-\alpha}$ by assumption. Let

$$H_{12}^{\text{hop}} = \sum_{i \in C_1} \sum_{j \in C_2} J_{ij} a_i^\dagger a_j + \text{h.c.} \tag{3.16}$$

denote the sum over all such pairs of sites. Then, we can bound $\left\| H_{12}^{\text{hop}} |\eta\rangle \right\| \leq O(b)$. To see this, diagonalize $H_{12}^{\text{hop}} = \sum_i w_i b_i^\dagger b_i$. Since H_{12}^{hop} only acts on two clusters, each normal mode contains up to $2b$ bosons. The maximum eigenvalue of H_{12}^{hop} is bounded by $2b \max_i w_i$, where w_i is the maximum normal mode frequency, given by the eigenvalue of the matrix $J_{ij} : i \in C_1, j \in C_2$. We now apply the Gershgorin circle theorem, which states that the maximum eigenvalue of J is bounded by the quantity $\max_i(\sum_j |J_{ij}|) \leq L^D L_{12}^{-\alpha}$.

Taking advantage of the fact that the clusters form a cubic lattice in D dimensions, we can group pairs of clusters by their relative distances. If we label clusters i by their D-dimensional coordinate $i_1, i_2, ..., i_D$, then we can define the cluster distance l between i

and j as $l+1 = \max_d |i_d - j_d|$. Cluster distance l corresponds to a minimum separation $l \times L$ between sites in different clusters. With this definition, there are $((2l+3)^D - (2l+1)^D) \approx 2^D D l^{D-1}$ clusters at a cluster distance l from any given cluster (Fig. 3.3(a)), and $K \times 2^D D l^{D-1}$ total pairs of clusters at cluster distance l. Notice that for a given separation vector, $K/2$ pairs of clusters (K total) can be simultaneously coupled without overlap (Fig. 3.3(b)). Therefore, there are approximately $2^{D+1} D l^{D-1}$ non-overlapping groupings per distance l. The sum over these non-overlapping Hamiltonians $H^{\text{hop}}_{a_1 b_1} + ... + H^{\text{hop}}_{a_{K/2} b_{K/2}}$ for each grouping is block diagonal. Therefore, the spectral norm (maximum eigenvalue) of the total Hamiltonian is equal to the maximum of the spectral norm over all irreducible blocks. Putting all this together, as long as $D - 1 - \alpha < -1$ the bound becomes

$$\max_{|\eta\rangle \in \text{im } Q} \|(H - QHQ)|\eta\rangle\| \leq \sum_{l=0}^{l_{\max}} O(2^{D+1} D l^{D-1})(2b) L^D (lL)^{-\alpha} = O(bL^{D-\alpha}). \quad (3.17)$$

□

We are now in a position to prove our main easiness result.

Theorem 9 (Easiness result). *For $\alpha > D+1$, and for all V, including $V = o(1)$ and $V = \omega(1)$, we have $t_{\text{easy}} = \Omega(n^{\gamma_{\text{easy}}})$, with*

$$\gamma_{\text{easy}} = \frac{\beta - 1}{D} \times \frac{\alpha - 2D}{\alpha - D} - \frac{1}{\alpha - D}, \quad (3.18)$$

and $t_{\text{easy}} = \Omega(\log n)$ if $\gamma_{\text{easy}} < 0$.

Proof. There are two error contributions, ϵ and δ, to the total error. The HHKL error ϵ is given by evaluation of Eq. (3.7), which is minimized by either choosing $N = 1$ or $N = t/t_1$

with t_1 a small fixed constant. This leads to three regimes with errors

$$\epsilon \leq O(1) \times \begin{cases} ne^{vt-L}, & \alpha \to \infty \\ \frac{nt^{\alpha-D}}{L^{\alpha-2D}}, & 2D + \frac{D}{\beta-1} < \alpha < \infty \\ \frac{n(e^{vt}-1)}{L^{\alpha-D-1}}, & D+1 < \alpha \leq 2D + \frac{D}{\beta-1}. \end{cases} \tag{3.19}$$

The truncation error, arising from using H' rather than H in the first step, is given by

$$\delta(t) \leq O(bL^{D-\alpha}) \int_0^t d\tau \epsilon(\tau). \tag{3.20}$$

Therefore, we can upper bound $\delta(t)$ by ϵ times an additional factor. This factor is $bL^{D-\alpha}t$, $\delta(t) = bL^{D-\alpha}t\epsilon$, when $\epsilon(\tau) = \mathsf{poly}(\tau)$ ($\alpha < \infty$), and it is $L^{D-\alpha}$, $\delta(t) = L^{D-\alpha}\epsilon$ when $\epsilon(\tau) = \exp(v\tau)$ ($\alpha \to \infty$). Our easiness results only hold for $\alpha > D+1$, so the L-dependent factor serves to suppress the truncation error in the asymptotic limit. Although the additional factor of t could cause $\delta(t) > \epsilon$ at late times, by this time, $\epsilon > \Omega(1)$ and we are no longer in the easy regime. Therefore, the errors presented in Eq. (3.19) can be immediately applied to calculate the timescales in Section 3.2. □

The resulting timescales are summarized in Table 3.3, which highlights the scaling of the timescale with respect to different physical parameters. We also consider the scaling of the easiness timescales when the density of the bosons increases by a factor ρ. In our setting, we implement this by scaling the number of bosons by ρ while keeping the number of lattice sites and the number of clusters (and their size) fixed. The effect of this is to increase the

Regime	Error	$t_{\text{easy}}(n, L)$	$t_{\text{easy}}(\rho)/t_{\text{easy}}(\rho = 1)$
$\alpha \to \infty$	ne^{vt-L}	L	$1/\rho$
$2D + \frac{D}{\beta-1} < \alpha < \infty$	$\frac{nvt^{\alpha-D}}{L^{\alpha-2D}}$	$n^{\frac{-1}{\alpha-D}} L^{\frac{\alpha-2D}{\alpha-D}}$	$\rho^{\frac{-2}{\alpha-D}}$
$D + 1 < \alpha \leq 2D + \frac{D}{\beta-1}$	$\frac{n(e^{vt}-1)}{L^{\alpha-D-1}}$	$(\alpha - D - 1)\log L - \log n$	$1/\rho$

Table 3.3: Summary of easiness timescales in the different regimes. Timescales follow from the error and are presented first as a function of n and L, which are the relevant physical scales of the problem. We study the effect of the density by performing the scaling $n \to \rho n, L \to L, b \to \rho b, m \to m$. The last column shows the timescale as a function of ρ in terms of the timescale when $\rho = 1$, namely $t_{\text{easy}}(\rho = 1)$.

Lieb-Robinson velocity: $v \to v\rho$. For all three cases, the net effect of increasing the density by a factor ρ is to decrease the easiness timescale.

3.8 Extended Easiness timescale for 1D

In this section, we prove that 1D systems with nearest-neighbor interactions ($\alpha \to \infty$) can be simulated for longer times by using matrix product states (MPS). Specifically, we can show that approximate sampling up to time $t = cL$ for any constant c is easy. This shows sampling is easy for any timescale $t = O(L)$. When combined with our hardness result for timescales $t = O(L^{1+\delta})$, this proves the transition is coarse.

The first step is to show how time evolution up to time $t = cL$ can be approximated with low error using the HHKL decomposition. Earlier, we proved that for $t = L/v$, we can self-consistently truncate the local Hilbert space dimension while incurring little error. Then, the sampling algorithm relies on the separability of the approximate wavefunction.

For $c > 1$, this separability no longer holds. Nevertheless, we can use a similar argument to self-consistently truncate the local Hilbert space. However, due to the lack of

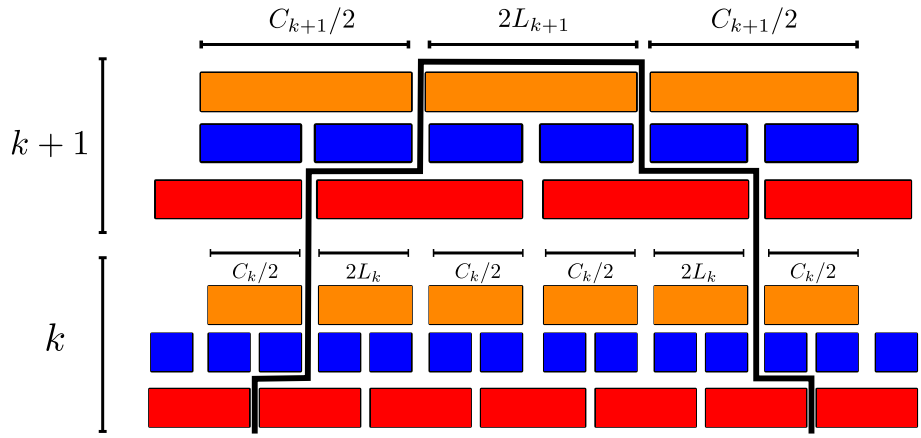

Figure 3.4: Schematic of the HHKL decomposition used to extend the classical simulation algorithm to longer times. Red and orange blocks denote forward time evolution, and blue blocks denote backwards time evolution. Two timesteps are depicted, k and $k+1$. The past causal lightcone of one the blocks is outlined in black. The lightcone dimensions and HHKL block size double during subsequent timesteps, leading to an exponentially large causal region.

separability, we only have an efficient simulation algorithm in $D = 1$, where tensor network representations of low-entanglement (area law) states can be efficiently contracted.

We divide the time evolution into time intervals, where we have already proved sampling is easy in the first interval up to time $t = t_1 = L/v_0$. Recall that during this first time interval, we applied the HHKL decomposition with a block size L. For the next time interval, we decompose the unitary U_{t_2,t_1} choosing an HHKL block size $2L$. By looking at the past causal lightcone of a single site, it is clear that the boson number per site is bounded by $4b$. Therefore, the Lieb-Robinson velocity during this time interval is four times larger that the first time interval, $v_1 = 4v_0$. Setting the approximation error $O\left(ne^{v_1(t_2-t_1)-2L}\right)$ to be $o(1)$, we can determine that for $t_2 - t_1 = O(\frac{L}{2v_0})$ the approximation errors are well controlled. Next, we need to generalize the argument to k time evolution steps.

We prove that we can define a self-consistent Hilbert space truncation scheme and an HHKL decomposition of the time-evolution unitary. During the k-th time interval,

we choose the HHKL block size to be $L_k = 2^k L$. To estimate the maximum local boson number, we look at the spatial extent of the past causal lightcone. Let $C_k + 2L_k$ denote the spatial extent of the past causal lightcone. Here C_k measures how far sideways the lightcone extends (see Fig. 3.4). Looking at the lightcone, we can write down a recurrence relation for C_k.

$$2L_{k+1} + C_{k+1} = 4L_{k+1} + C_k \tag{3.21}$$

$$C_{k+1} = 2L_{k+1} + C_k = 2^{k+2}L + C_k \tag{3.22}$$

$$\implies C_k = 4(2^k - 1)L, \tag{3.23}$$

where we have applied the initial condition $C_0 = 0$. Combined with the boson density $\rho = \frac{1}{2L}$, the maximum boson number after step k is proportional to the size of the lightcone:

$$N_k = \rho(C_k + 2L_k) \tag{3.24}$$

$$= \rho L(6 \times 2^k - 4) = 3 \times 2^k - 2 \tag{3.25}$$

This tells us we can choose timesteps $t_k - t_{k-1} = O(L/3v_0)$ for the error to be $o(1)$.

To complete the proof, we need to argue that the HHKL time-evolution can be simulated efficiently. The simplest way to do this is to fix the evolution time $t = cL/3v_0$. Then, we can think of the HHKL decomposed unitary as a finite-depth unitary circuit, which generates constant (system-size independent) entanglement across any cut. Therefore, standard MPS algorithms can be applied to with cost polynomial in the system size, the local Hilbert space dimension (physical index dimension), and the entanglement across any cut

(virtual index dimension) [176]. Although both the physical and virtual dimensions scale exponentially or faster with c, since we have fixed c to be system-size independent, we have an efficient sampling algorithm. Furthermore, this algorithm works for any fixed c, proving that sampling cannot be hard for evolution times $t = O(L)$ and that the transition is coarse.

3.9 Hardness timescale for interacting bosons

In this section, we provide more details about how to achieve the timescales in our hardness results, which we state in more detail first.

Theorem 10 (Hardness result). *When $\alpha \geq D/2$, $V = \Omega(1)$, and $D \geq 2$, the hardness timescale is $t_{\text{hard}} = O(n^{\gamma^{\text{I}}_{\text{hard}}})$, where*

$$\gamma^{\text{I}}_{\text{hard}} = \begin{cases} \frac{\beta-1}{D} \min[1, \alpha - D], & \alpha > D \\ 0, & \alpha \in [\frac{D}{2}, D]. \end{cases} \tag{3.26}$$

In all other cases, i.e. when at least one of the following cases hold: $\alpha < D/2$, $V = o(1)$, or $D = 1$, the timescale is $t_{\text{hard}} = O(n^{\gamma^{\text{II}}_{\text{hard}}})$, where

$$\gamma^{\text{II}}_{\text{hard}} = \delta + \begin{cases} \frac{\beta-1}{D} \min\left[1 + \frac{O(\log(V+1))}{\log n}, \alpha - D\right], & \alpha > D \\ 0, \quad \alpha \in [\frac{D}{2}, D] \\ \frac{\beta}{D}\left(\alpha - \frac{D}{2}\right), & \alpha < D/2 \end{cases} \tag{3.27}$$

for an arbitrarily small $\delta > 0$.

Almost any bosonic interaction is universal for BQP [166] and hence these results are

applicable to general on-site interactions $f(n_i)$. Reference [188] also answers the questions of what additional gates or Hamiltonians can make linear optics universal. We first describe how a bosonic system with fully controllable local fields $J_{ii}(t)$, hoppings $J_{ij}(t)$, and a fixed Hubbard interaction $\frac{V}{2}\sum_i \hat{n}_i(\hat{n}_i - 1)$ can implement a universal quantum gate set. To simulate quantum circuits, which act on two-state spins, we use a dual-rail encoding. Using $2n$ bosonic modes, and n bosons, n logical qubits are defined by partitioning the lattice into pairs of adjacent modes, and a boson is placed in each pair. Each logical qubit spans a subspace of the two-mode Hilbert space. Specifically, $|0\rangle_L = |10\rangle, |1\rangle_L = |01\rangle$. We can implement any single qubit (2-mode) unitary by turning on a hopping between the two sites (X-rotations) or by applying a local on-site field (Z-rotations). To complete a universal gate set, we need a two-qubit entangling gate. This can be done, say, by applying a hopping term between two sites that belong to different logical qubits [175]. All these gates are achievable in $O(1)$ time when $V = \Theta(1)$. In the limit of large Hubbard interaction $V \to \infty$, the entangling power of the gate decreases as $1/V$ [175] and one needs $O(V)$ repetitions of the gate in order to implement a standard entangling gate such as the CNOT.

For hardness proofs that employ postselection gadgets, we must ensure that the gate set we work with comes equipped with a Solovay-Kitaev theorem. This is the case if the gate set is closed under inverse, or contains an irreducible representation of a non-Abelian group [189]. In our case, the gate set contains single-qubit Paulis and hence has a Solovay-Kitaev theorem, which is important for the postselection gadgets to work as intended.

We will specifically deal with the scheme proposed in Ref. [24]. It applies a constant-depth circuit on a grid of $\sqrt{n} \times \sqrt{n}$ qubits in order to implement a random IQP circuit [12; 19] on $\sqrt{n}$ effective qubits. This comes about because the cluster state, which is a

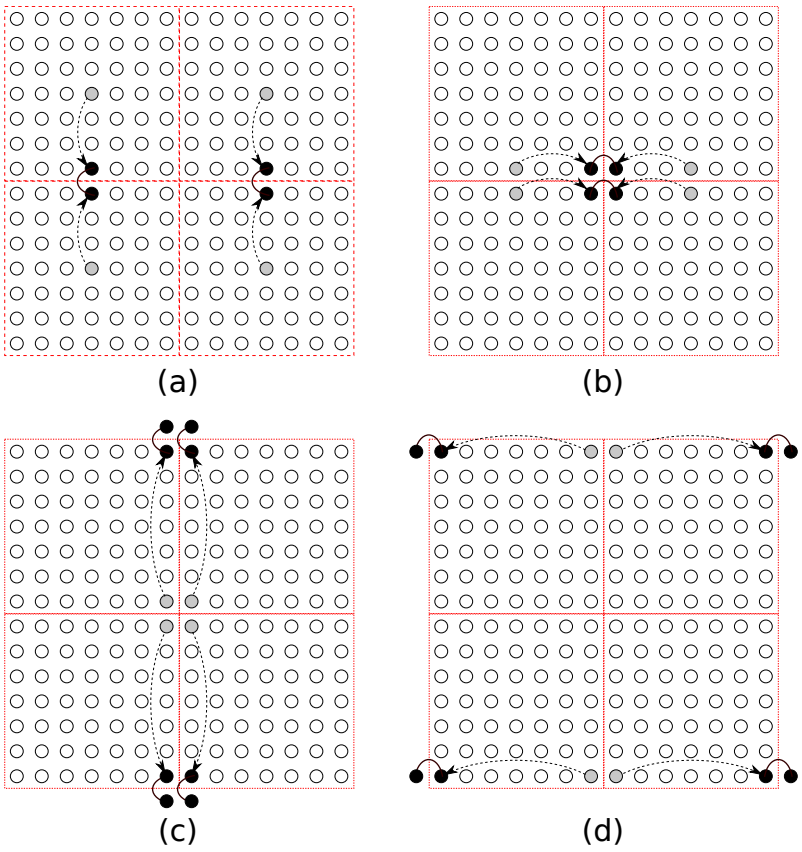

Figure 3.5: A protocol that implements the logical circuit of Ref. [24]. Each subfigure shows the location of the site that previously encoded the $|1\rangle$ state in gray. The current site that encodes the $|1\rangle$ state is in black. The site that encodes $|0\rangle$ is not shown but moves similarly as the $|1\rangle$ state. The distance traversed by each qubit is $L + L + 2L + 2L = 6L$.

universal resource for measurement-based quantum computation, can be made with constant depth on a two-dimensional grid.

For short-range hops ($\alpha \to \infty$), we implement the scheme in four steps as shown in Fig. 3.5. In each step, we move the logical qubits to bring them near each other and make them interact in order to effect an entangling gate. For short-range hopping, the time taken to move a boson to a far-off site distance L away dominates the time taken for an entangling gate. The total time for an entangling gate is thus $O(L) + O(1) = O(L)$.

For long-range hopping, we use the same scheme as in Fig. 3.5, but we use the long-range hopping to speed up the movement of the logical qubits. This is precisely the question of state transfer using long-range interactions/hops [172; 173; 190]. In the following we give an overview of the best known protocol for state transfer, but first we should clarify the assumptions in the model. The Hamiltonian is a sum of $O(m^2)$ terms, each of which has norm bounded by at most $1/d(i,j)^\alpha$. Since we assume we can apply any Hamiltonian subject to these constraints, in particular, we may choose to apply hopping terms across all possible edges. This model makes it possible to go faster than the circuit model if we compare the time in the Hamiltonian model with depth in the circuit model. This power comes about because of the possibility of allowing simultaneous noncommuting terms to be applied in the Hamiltonian model.

The state transfer protocols in Ref. [172; 173] show such a speedup for state transfer. The broad idea in both protocols is to apply a map $|1\rangle_1 \to |1\rangle_A := \sum_{j\in A} \frac{1}{\sqrt{|A|}} |1\rangle_j$, followed by the steps $|1\rangle_A \to |1\rangle_B$ and $|1\rangle_B \to |1\rangle_2$, where A and B are regions of the lattice to be specified. In the protocol of Ref. [172], which is faster than that of Ref. [173] for $\alpha \leq D/2$, $A = B = \{j : j \neq 1, 2\}$ and each step takes time $O(L^\alpha/\sqrt{N-2})$, where $N-2$ is the

number of ancillas used and L is the distance between the two furthest sites. In the protocol of Ref. [173], which is faster for $\alpha \in (D/2, D+1]$, A and B are large regions around the initial and final sites, respectively. This protocol takes time $O(1)$ when $\alpha < D$, $O(\log L)$ when $\alpha = D$, and $O(L^{\alpha-D})$ when $\alpha > D$.

In our setting, we use the state transfer protocols to move the logical qubit faster than time $O(L)$ in each step of the scheme depicted in Fig. 3.5. If $\alpha < D/2$, we use all the ancillas in the entire system, giving a state transfer time of $O(m^{\alpha/D-1/2}) = O(n^{\beta(\frac{\alpha}{D}-\frac{1}{2})})$. If $\alpha > D/2$, we only use the empty sites in a cluster as ancillas in the protocol of Ref. [173], giving the state transfer time mentioned above. This time is faster than $O(L)$, the time it would take for the nearest-neighbor case, when $\alpha < D+1$. Therefore, for 2D or higher and $\alpha \geq D/2$, the total time it takes to implement a hard-to-simulate circuit is $\min[L, L^{\alpha-D}\log L]+O(1)$, proving Theorem 10 for interacting bosons. When $\alpha < D/2$, the limiting step is dominated by the entangling gate, which takes time $O(1)$. Therefore for this case we only get fast hardness through boson sampling, which is discussed in Section IV. Note that when $t = o(1)$ and interaction strength is $V = \Theta(1)$, the effect of the interaction is governed by $Vt = o(1)$, which justifies treating the problem for short times as a free-boson problem.

3.9.1 One dimension

In 1D with nearest-neighbor hopping, we cannot hope to get a hardness result for simulating constant depth circuits, which is related to the fact that one cannot have universal measurement-based quantum computing in one dimension. We change our strategy here. The overall goal in 1D is to still be able to simulate the scheme in Ref. [24] since it provides

a faster hardness time (at the cost of an overhead in the qubits). The way this is done is to either (i) implement $O(n)$ SWAPs in 1D in order to implement an IQP circuit [12], or (ii) use the long-range hops to directly implement gates between logical qubits at a distance L away.

For the first method, we use state transfer to implement a SWAP by moving each boson within a cluster a distance $\Theta(L)$. This takes time $O(t_s(L))$, where $t_s(L)$ is the time taken for state transfer over a distance L and is given by

$$t_s(L) = c \times \begin{cases} L, & \alpha > 2 \\ L^{\alpha-1}, & \alpha \in (1,2] \\ \log L, & \alpha = 1 \\ 1, & \alpha \in [\frac{1}{2}, 1) \\ L^{\alpha-1/2}, & \alpha < \frac{1}{2}. \end{cases} \tag{3.28}$$

We write this succinctly as $O\left(\min[L, L^{\alpha-1}\log L + 1, L^{\alpha-1/2}]\right)$. The total time for n SWAPs is therefore $O\left(n \times \min[L, L^{\alpha-1}\log L + 1, L^{\alpha-1/2}]\right)$.

The second method relies on the observation that when $\alpha \to 0$, the distinction between 1D and 2D becomes less clear, since at $\alpha = 0$, the connectivity is described by a complete graph and all hopping strengths are equal. Let us give some intuition for the $\alpha \to 0$ case. One would directly "sculpt" a 2D grid from the available graph, which is a complete graph on n vertices (one for every logical qubit) with weights w_{ij} given by $d(i,j)^{-\alpha}$. If we want to arrange qubits on a 1D path, we can assign an indexing to qubits in the 2D

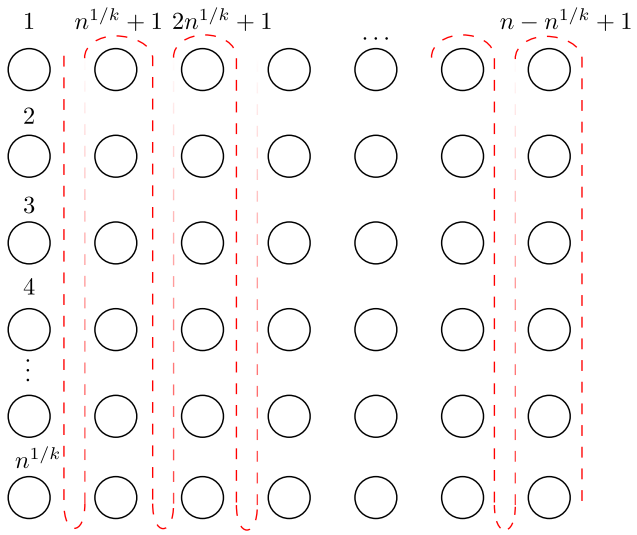

Figure 3.6: A snaking scheme to assign indices to qubits in 2D for a $n^{1/k} \times n^{1-1/k}$ grid, which is used in mapping to 1D.

grid and place them in the 1D path in increasing order of their index. One may, in particular, choose a "snake-like indexing" depicted in Fig. 3.6. This ensures that nearest-neighbor gates along one axis of the 2D grid map to nearest-neighbor gates in 1D. Gates along the other axis, however, correspond to nonlocal gates in 1D. Suppose that the equivalent grid in 2D is of size $n^{1/k} \times n^{1-1/k}$. The distance between two qubits that have to participate in a gate is now marginally larger ($O(Ln^{1/k})$ instead of $O(L)$), but the depth is greatly reduced: it is now $O(n^{1/k})$ instead of $O(n)$. We again use state transfer to move close to a far-off qubit and then perform a nearest-neighbor entangling gate. This time is set by the state transfer protocol, and is now $t_s(n^{1/k}L) = O\left(n^{1/k} \times \min[L, L^{\alpha-1}\log L + 1, L^{\alpha-1/2}]\right)$. For large $k = \Theta(1)$, this gives us the bound $O\left(\min[L^{1+\delta}, L^{\alpha-1+\delta} + L^{\Theta(\delta)}, L^{\alpha-1/2+\delta}]\right)$ for any $\delta > 0$, giving a coarse transition. Notice, however, that faster hardness in 1D comes at a high cost– the effective number of qubits on which we implement a hard circuit is only $\Theta(n^{1/k}) = n^{\Theta(\delta)}$, which approaches a constant as $\delta \to 0$.

This example of 1D is very instructive– it exhibits one particular way in which the

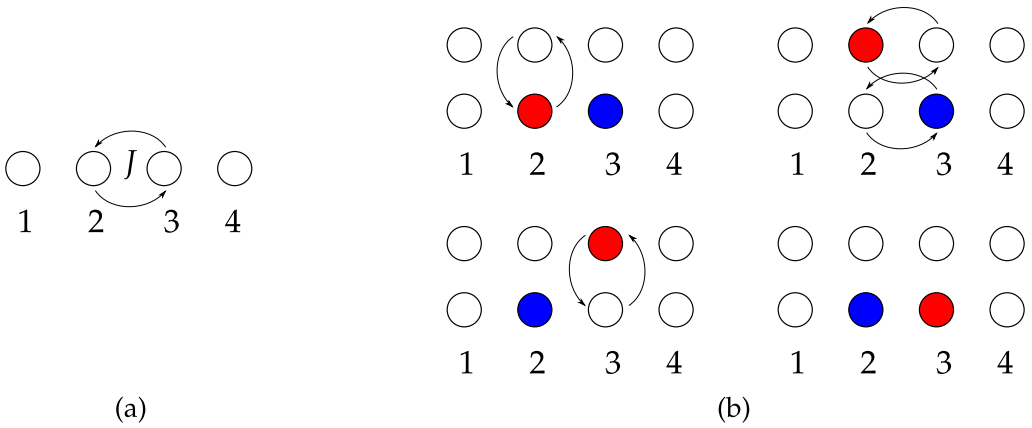

Figure 3.7: (a) A hopping between sites 2 and 3 that implements the mode unitary $\begin{pmatrix} \cos(|J|t) & -i\sin(|J|t)J/|J| \\ -i\sin(Jt)J^*/|J| & \cos(|J|t) \end{pmatrix} = e^{-it(\mathrm{Re}\{J\}X+\mathrm{Im}\{J\}Y)}$. When $|J|t = \pi$, this is a SWAP between two modes with phases $(-iJ/|J|, -iJ^*/|J|)$ that depends on $\arg J$, the argument of J. (b) A "physical" SWAP between sites 2 and 3 by using ancilla sites available whenever the system is not nearest-neighbor in 1D. The colors are used to label the modes and how they move, and do not mean that both sites are occupied. The total hopping phase incurred when performing the physical SWAP can be set to be $(+i, -i)$, which cannot be achieved with just the hopping term shown in (a).

complexity phase transition can happen. As we take higher and higher values of k, the hardness time would decrease, coming at the cost of a decreased number of effective qubits. This smoothly morphs into the easiness regime when $\alpha \to \infty$ since in this regime both transitions happen at $t = \Theta(L)$.

If the definition of hardness is more stringent (in order to link it to fine-grained complexity measures such as explicit quantitative lower bound conjectures), then the above mentioned overhead is undesirable. In this case we would adopt the first strategy to implement SWAPs and directly implement a random IQP circuit on all the n qubits. This would increase the hardness time by a factor n.

3.9.2 Hardcore limit

In the hardcore limit $V \to \infty$, the strategy is modified. Let us consider a physical qubit to represent the presence ($|1\rangle$) or absence ($|0\rangle$) of a boson at a site. A nearest-neighbor hop translates to a term in the Hamiltonian that can be written in terms of the Pauli operators as $XX + YY$. Further, an on-site field $J_{ii}a_i^\dagger a_i$ translates to a term $\propto Z$. There are no other terms available, in particular single-qubit rotations about other axes X or Y. This is because the total boson number is conserved, which in the spin basis corresponds to the conservation of $\sum_i Z_i$. This operator indeed commutes with both the allowed Hamiltonian terms specified above.

Let us now discuss the computational power of this model. When the physical qubits are constrained to have nearest-neighbor interactions in 1D, this model is nonuniversal and classically simulable. This can be interpreted due to the fact that this model is equivalent to matchgates on a path (i.e. a 1D nearest-neighbor graph), which is nonuniversal for quantum computing without access to a SWAP gate. Alternatively, one can apply the Jordan-Wigner transformation to map the spin model onto free fermions. One may then use the fact that fermion sampling is simulable on a classical computer [51].

When the connectivity of the qubit interactions is different, the model is computationally universal for BQP. In the matchgate picture, this result follows from Ref. [191], which shows that matchgates on any graph apart from the path or the cycle are universal for BQP in an encoded sense. In the fermion picture, the Jordan-Wigner transformation on any graph other than a path graph would typically result in nonlocal interacting terms that are not quadratic in general. Thus, the model cannot be mapped to free, quadratic fermions

and the simulability proof from Ref. [51] breaks down.

Alternatively, a constructive way of seeing how we can recover universality is as follows. Consider again the dual rail encoding and two logical qubits placed next to each other as in Fig. 3.7. Apply a coupling $J(a_2^\dagger a_3 + a_3^\dagger a_2)$ on the modes 2 and 3 for time $t = \frac{\pi}{2J}$. This effects the transition $|10\rangle_{23} \to -i\,|01\rangle_{23}$ and $|01\rangle_{23} \to -i\,|10\rangle_{23}$, while leaving the state $|11\rangle_{23}$ the same. Now we swap the modes 2 and 3 using an ancilla mode that is available by virtue of having either long-range hopping or having $D > 1$. This returns the system back to the logical subspace of exactly one boson in modes 1 & 2, and one boson in modes 3 & 4, and effects the unitary $\mathrm{diag}\{1, 1, 1, -1\}$ in the (logical) computational basis. This is an entangling gate that can be implemented in $O(1)$ time and thus the hardness timescale for hardcore interactions is the same as that of Hubbard interactions with $V = \Theta(1)$.

We finally discuss the case when V is polynomially large. Using the dual-rail encoding and implementing the same protocol as the non-hardcore case now takes the state $|11\rangle_{23}$ to $\lambda\,|11\rangle_{23} + \mu\frac{|20\rangle_{23}+|02\rangle_{23}}{\sqrt{2}}$, with $\mu \propto \frac{J}{\sqrt{8J^2+V^2}} \sin\left(\frac{t\sqrt{8J^2+V^2}}{2}\right)$. When $|\mu| \neq 0$, we get an error because the state is outside the logical subspace. The probability with which this action happens is suppressed by $1/V^2$, however, which is polynomially small when $V = \mathsf{poly}(n)$.

However, one can do better: by carefully tuning the hopping strength $J \in [0, 1]$ and the evolution time t, one can always achieve the goal of getting $\mu = 0$ exactly and implementing an operation $\exp\left[-i\frac{\pi}{2}X\right]$ in the $|10\rangle_{23}, |01\rangle_{23}$ subspace. This requires setting $t\sqrt{2J^2 + \frac{V^2}{4}} = m\pi$ and $t = \frac{2\pi}{J}$ for integer m. This can be solved as follows: set $m = \lceil\sqrt{8+V^2}\rceil$, and $J = \frac{V}{\sqrt{m^2-8}}$ (which is ≤ 1 since $m \geq \sqrt{8+V^2}$). The time is set by the condition $t = \frac{2\pi}{J}$, which is $\Theta(1)$. This effects a logical CPHASE$[\phi]$ gate with angle $\phi = -\pi V/J$.

Finally, the above parameters that set μ exactly to zero work even for exponentially large $V = \Omega(\exp(n))$, but this requires exponentially precise control of the parameters J and t, which may not be physically feasible. In this case, we simply observe that $|\mu|^2$, the probability of going outside the logical subspace and hence making an error, is $O(1/V^2)$, which is exponentially small in n. Therefore, in this limit, the gate we implement is exponentially close to perfect, and the complete circuit has a very small infidelity as well.

3.10 Hardness timescale for free bosons

In this section, we review Aaronson and Arkhipov's method of creating a linear optical state that is hard to sample from [13]. We then give a way to construct such states in time $\tilde{O}\left(nm^{\alpha/D-1/2}\right)$ with high probability in the Hamiltonian model, and prove Theorem 10 for free bosons.

For free bosons, in order to get a state that is hard to sample from, we need to apply a Haar-random linear-optical unitary on m modes to the state $|1, 1, \ldots 1, 0, 0, \ldots 0\rangle$. Aaronson and Arkhipov gave a method of preparing the resulting state in $O(n \log m)$ depth in the circuit model. Their method involves the use of ancillas and can be thought of as implementing each column of the Haar-random unitary separately in $O(\log m)$-depth. Here we mean that we apply the map $|1\rangle_j \rightarrow \sum_{i\in\Lambda} U_{ij} |1\rangle_i$ to "implement" the column i of the linear-optical unitary U. In the Hamiltonian model, we can apply simultaneous, non-commuting terms of a Hamiltonian involving a common site. The only constraint is that each term of the Hamiltonian should have a bounded norm of $1/d(i,j)^\alpha$. In this model, when α is small, it is possible to implement each unitary in a time much smaller than $O(\log m)$– indeed, we

show the following:

Lemma 11. *Let U be a Haar-random unitary on m modes. Then with probability $1 - \frac{1}{\mathsf{poly}(m)}$ over the Haar measure, each of the first n columns of U can be implemented in time $O\left(\frac{\sqrt{\log m}}{m^{1/2-\alpha/D}}\right)$.*

To prove this, we will need an algorithm that implements columns of the unitary. For convenience, let us first consider the case $\alpha = 0$. The algorithm involves two subroutines, which we call the single-shot and state-transfer protocols. Both protocols depend on the following observation. If we implement a Hamiltonian that couples a site i to all other sites $j \neq i$ through coupling strengths J_{ij}, then the effective dynamics is that of two coupled modes $a_i^\dagger$ and $b^\dagger = \frac{1}{\omega}\sum_{j\neq i} J_{ij} a_j^\dagger$, where $\omega = \sqrt{\sum_{j\neq i} J_{ij}^2}$. The effective speed of the dynamics is given by ω– for instance, the time period of the system is $\frac{2\pi}{\omega}$.

The single-shot protocol implements a map $a_i^\dagger \to \gamma_i a_i^\dagger + \sum_{j\neq i}\gamma_j a_j^\dagger$. This is done by simply applying the Hamiltonian $H \propto a_i^\dagger(\sum_{j\neq i}\gamma_j a_j) + \text{h.c.}$ for time $t = \frac{1}{\omega}\cos^{-1}|\gamma_i|$. In the case $\alpha = 0$, we can set the proportionality factor equal to $1/\max|\gamma_j|$. This choice means that the coupling strength between i and the site k with maximum $|\gamma_k|$ is set to 1 (the maximum), and all other couplings are equal to $|\frac{\gamma_j}{\gamma_k}|$.

The other subroutine, the state-transfer protocol is also an application of the above observation and appears in Ref. [172]. It achieves the map $a_i^\dagger \to \gamma_i a_i^\dagger + \gamma_j a_j^\dagger$ via two rounds of the previous protocol. This is done by first mapping site i to the uniform superposition over all sites except i and j, and then coupling this uniform superposition mode to site j. The time taken for this is $\frac{1}{\omega}\left(\frac{\pi}{2} + \cos^{-1}|\gamma_i|\right)$. Since $\omega = \sqrt{m-2}$ (all $m-2$ modes are coupled with equal strength to modes i or j), this takes time $O\left(\frac{1}{\sqrt{m}}\right)$.

These subroutines form part of Algorithm 2. It can be seen that Algorithm 2 imple-

Algorithm 2: Algorithm for implementing one column of a unitary

Input: Unitary U, column index j

1 Reassign the mode labels for modes $i \neq j$ in nonincreasing order of $|U_{ij}|$.

2 Implement the state-transfer protocol to map the state $a_j^\dagger |\text{vac}\rangle$ to $U_{jj} a_j^\dagger |\text{vac}\rangle + \sqrt{1 - |U_{jj}|^2} a_1^\dagger |\text{vac}\rangle$. Skip this step if $|U_{jj}| \geq |U_{j1}|$ already.

3 Use the single-shot protocol between site 1 and the rest ($i \neq 1, j$) to map $a_1^\dagger \to \frac{U_{1j}}{\sqrt{1-|U_{jj}|^2}} a_1^\dagger + \sum_{i \neq 1,j} \frac{U_{ij}}{\sqrt{1-|U_{jj}|^2}} a_i^\dagger$.

ments a map $a_j^\dagger \to U_{jj} a_j^\dagger + \sum_{i \neq j} U_{ij} a_i^\dagger$, as desired. To prove Lemma 11 we need to examine the runtime of the algorithm when U is drawn from a Haar-random distribution.

Proof of Lemma 11. First, notice that since the Haar measure is invariant under the action of any unitary, we can in particular apply a permutation map to argue that the elements of the i'th column are drawn from the same distribution as the first column. Next, recall that one may generate a Haar-random unitary by first generating m uniform random vectors in $\mathbb{C}^m$ and then performing a Gram-Schmidt orthogonalization. In particular, this means that the first column of a Haar-random unitary may be generated by generating a uniform random vector with unit norm. This implies that the marginal distribution over any column of a unitary drawn from the Haar measure is simply the uniform distribution over unit vectors, since we argued above that all columns are drawn from the same distribution.

Now, let us examine the runtime. The first step (line 2 of the algorithm) requires time $t = O\left(\frac{1}{\sqrt{m}}\right)$ irrespective of U_{jj} because the total time for state-transfer is $\frac{1}{\omega}\left(\frac{\pi}{2} + \cos^{-1} U_{jj}\right) \leq$

$\frac{\pi}{\omega} = \frac{\pi}{\sqrt{m-2}}$. Next, the second step takes time $t = \frac{1}{\omega}\cos^{-1}\left(\frac{U_{1j}}{\sqrt{1-|U_{1j}^2|}}\right) = O(\frac{1}{\omega})$. Now,

$$\omega = \sqrt{1^2 + \frac{|U_{3j}|^2/(1-|U_{jj}|^2)}{|U_{2j}|^2/(1-|U_{jj}|^2)} + \frac{|U_{4j}|^2}{|U_{2j}|^2} + \dots} \tag{3.29}$$

$$= \sqrt{\frac{\sum_{i=2, i\neq j}^{m} |U_{ij}|^2}{|U_{2j}|^2}} = \sqrt{\frac{1-|U_{1j}|^2-|U_{jj}|^2}{|U_{2j}|^2}} \tag{3.30}$$

Now in cases where $|U_{jj}| \leq |U_{1j}|$ (where $|U_{1j}|$ is the maximum absolute value of the column entry among all other modes $i \neq j$), which happens with probability $1 - \frac{1}{m}$, we will have $\omega^2 \geq \frac{1-2|U_{1j}|^2}{|U_{2j}|^2}$. In the other case when $|U_{jj}| \geq |U_{1j}|$, meaning that the maximum absolute value among all entries of column j is in row j itself, we again have $\omega^2 \geq \frac{1-2|U_{jj}|^2}{|U_{2j}|^2}$. Both these cases can be written together as $\omega^2 \geq \frac{1-2|U_{1j}|^2}{|U_{2j}|^2}$, where we now denote U_{1j} as the entry with maximum absolute value among *all* elements of column j. The analysis completely hinges on the typical ω we have, which in turn depends on $|U_{1j}|$. We will show $\Pr\left(\omega^2 \geq \frac{cm}{\log m}\right) \geq 1 - \frac{1}{\mathsf{poly}(m)}$, which will prove the claim for $\alpha = 0$.

$$\Pr\left(\omega^2 \geq \frac{cm}{\log m}\right) \geq \Pr\left(1 - 2|U_{1j}|^2 \geq c_1 \;\&\; |U_{2j}|^2 \leq \frac{c_1 \log m}{cm}\right) \tag{3.31}$$

since the two events on the right hand side suffice for the first event to hold. Further,

$$\Pr\left(1 - 2|U_{1i}|^2 \geq c_1 \;\&\; |U_{2j}|^2 \leq \frac{c_1 \log m}{cm}\right) \geq \Pr\left(|U_{1j}|^2 \leq \frac{c_1 \log m}{cm}\right) \tag{3.32}$$

for large enough m with some fixed $c_1 = 0.99$ (say), since $|U_{2j}|^2 \leq |U_{1j}|^2$ and $1 - 1.98 \log m/m \geq 0.99$ for large enough m.

To this end, we refer to the literature on order statistics of uniform random unit

vectors $(z_1, z_2, \ldots z_m) \in \mathbb{C}^m$ [192]. This Chapter gives an explicit formula for $F(x, m)$, the probability that all $|z_j|^2 \leq x$. We are interested in this quantity at $x = c_1 \log m/(cm)$, since this gives us the probability of the desired event ($\omega^2 \geq cm/\log m$). We have

$$\Pr\left(\frac{1}{k+1} \leq x \leq \frac{1}{k}\right) = \sum_{l=0}^{k} \binom{m}{l} (-1)^l (1-lx)^{m-1}. \tag{3.33}$$

It is also argued in Ref. [192] that the terms of the series successively underestimate or overestimate the desired probability. Therefore we can expand the series and terminate it at the first two terms, giving us an inequality:

$$\Pr\left(\frac{1}{k+1} \leq x \leq \frac{1}{k}\right) = 1 - m(1-x)^{m-1} + \frac{m^2}{2}(1-2x)^{m-1} - \ldots \tag{3.34}$$

$$\geq 1 - m(1-x)^{m-1}. \tag{3.35}$$

Choosing $c = c_1/4 = 0.2475$, we are interested in the quantity when $k = \lfloor \frac{m}{4 \log m} \rfloor$:

$$\Pr(x \leq 4\log m/m) \geq 1 - m(1 - 4\log m/m)^{m-1} \geq 1 - \frac{1}{m^{3-4/m}}, \text{since} \tag{3.36}$$

$$(1 - 4\log m/m)^{m-1} = \exp\left[(m-1)\log\left(1 - \frac{4\log m}{m}\right)\right] \leq \exp\left[-4(m-1)\frac{\log m}{m}\right] = m^{-4(1-1/m)}. \tag{3.37}$$

This implies that the time for the single-shot protocol is also $t = O(\frac{1}{\omega}) = O(\sqrt{\frac{\log m}{m}})$ for a single column. Notice that we can make the polynomial appearing in $\Pr(\omega^2 \geq cm/\log m) \geq 1 - 1/\mathsf{poly}(m)$ as small as possible by suitably reducing c. To extend the proof to all columns,

we use the union bound. In the following, let t_j denote the time to implement column j.

$$\Pr\left(\exists j : t_j > \sqrt{\frac{\log m}{cm}}\right) \leq \sum_j \Pr\left(t_j > \sqrt{\frac{\log m}{cm}}\right) \tag{3.38}$$

$$\leq m \times \frac{1}{\mathsf{poly}(m)} = \frac{1}{\mathsf{poly}(m)} \tag{3.39}$$

when the degree in the polynomial is larger than 1, just as we have chosen by setting $c = 0.2475$. This implies

$$\Pr\left(\forall j : t_j \leq \sqrt{\frac{\log m}{cm}}\right) = 1 - \Pr\left(\exists j : t_j > \sqrt{\frac{\log m}{cm}}\right) \geq 1 - \frac{1}{\mathsf{poly}(m)}. \tag{3.40}$$

This completes the proof in the case $\alpha = 0$. When $\alpha \neq 0$, we can in the worst-case set each coupling constant to a maximum of $O(m^{-\alpha/D})$, which is the maximum coupling strength of the furthest two sites separated by a distance $O(m^{1/D})$. This factor appears in the total time for both the state-transfer [172] and single-shot protocols, and simply multiplies the required time, making it $O\left(\sqrt{\frac{\log m}{m}} \times m^{\alpha/D}\right) = O\left(\frac{\sqrt{\log m}}{m^{1/2-\alpha/D}}\right)$. Finally, if there are any phase shifts that need to be applied, they can be achieved through an on-site term $J_{ii}a_i^\dagger a_i$, whose strength is unbounded by assumption and can thus take arbitrarily short time. □

The total time for implementing boson sampling on n bosons is therefore $O\left(n\frac{\sqrt{\log m}}{m^{1/2-\alpha/D}}\right) = \tilde{O}\left(n^{1+\beta(\frac{\alpha}{D}-\frac{1}{2})}\right)$, since we should implement n columns in total.

3.10.1 Optimizing hardness time

We can optimize the hardness time by implementing boson sampling not on n bosons, but on n^δ of them, for any $\delta \in (0, 1]$. The explicit lower bounds on running time of classical algorithms we would get assuming fine-grained complexity-theoretic conjectures is again something like $\exp\left[n^{\mathsf{poly}(\delta)}\right]$ for any $\delta \in (0, 1]$. This grows very slowly with n, but it still qualifies as subexponential, which is not polynomial or quasipolynomial (and, by our definition, would fall in the category "hard"). This choice of parameters allows us to achieve a smaller hardness timescale at the cost of getting a coarse (type-II) transition. We analyze this idea in three cases: $\alpha \leq D/2$, $\alpha \in (\frac{D}{2}, D]$ and $\alpha > D$.

When $\alpha \leq D/2$, we perform boson sampling on the nearest set of n^δ bosons with the rest of the empty sites in the lattice as target sites. In terms of the linear optical unitary, the unitary acts on $m - n^\delta = \Theta(m)$ sites in the lattice, although only the n^δ columns corresponding to initially occupied sites are relevant. Using the protocol in Lemma 11, the total time to implement n^δ columns of an $m \times m$ linear optical unitary is $O(n^\delta m^{\alpha/D-1/2} \log n) = \tilde{O}(n^\delta n^{\frac{\beta}{D}(\alpha-D/2)})$.

When $\alpha \in (\frac{D}{2}, D]$, the strategy is modified. We first move the nearest set of n^δ bosons into a contiguous set of sites within a single cluster. This takes time $O(n^\delta)$, since each boson may be transferred in time $O(1)$. We now perform boson sampling on these n^δ bosons with the surrounding $n^{2\delta}$ sites as targets, meaning that the effective number of total sites is $m_{\text{eff}} = O(n^{2\delta})$, as required for the hardness of boson sampling. Applying Lemma 11, the time required to perform hard instances of boson sampling is now $O(n^\delta n^{2\delta(\alpha/D-1/2)} \log n) = n^{O(\delta)}$ for arbitrarily small $\delta > 0$.

Lastly, when $\alpha > D$, we use the same protocol as above. The time taken for the state transfer is now $n^\delta \times \min[L, L^{\alpha-D}]$. Once state transfer has been achieved, we use nearest-neighbor hops instead of Lemma 11 to create an instance of boson sampling in time $O(n^{2\delta/D})$. Since state transfer is the limiting step, the total time is $n^\delta \times \min[L, L^{\alpha-D}]$. The hardness timescale is obtained by taking the optimum strategy in each case, giving the hardness timescale $t_{\text{hard}} = \tilde{O}(n^{\gamma^{\text{II}}_{\text{hard}}})$, where

$$\gamma^{\text{II}}_{\text{hard}} = \delta + \begin{cases} \frac{\beta-1}{D} \min[1, \alpha - D] & \alpha > D \\ 0 & \alpha \in (\frac{D}{2}, D] \\ \frac{\beta}{D}\left(\alpha - \frac{D}{2}\right) & \alpha < \frac{D}{2} \end{cases} \tag{3.41}$$

for an arbitrarily small $\delta > 0$. This proves Theorem 10 for free bosons and for interacting bosons in the case $\alpha < D/2$. When we compare with Ref. [96], which states a hardness result for $\alpha \to \infty$, we see that we have almost removed a factor of n from the timescale coming from implementing n columns of the linear optical unitary. Our result here gives a coarse hardness timescale of $\Theta(L)$ that matches the easiness timescale of L. More importantly, this makes the noninteracting hardness timescale the same as the interacting one.

3.10.2 Almost free bosons, $V = o(1)$

When the interaction strength satisfies $V = o(1)$ or when the bosons are almost free, we can treat the evolution as being close to that of free bosons. The total variation distance error between the actual distribution and the distribution modeled by free bosons can be

upper bounded by

$$\varepsilon \leq \left\| |\psi(t)\rangle - \left|\tilde{\psi}(t)\right\rangle \right\|_2 =: \delta(t), \tag{3.42}$$

where we take H and $|\psi\rangle$ to be the actual Hamiltonian and state and $\tilde{H}$ and $\left|\tilde{\psi}\right\rangle$ to be their respective free-bosonic approximations. Therefore, by the same logic leading up to Eq. (3.11), we have the same expression here for $\delta(t)$:

$$\delta(t) \leq \int_0^t \mathrm{d}\tau \left\| \left(H(\tau) - \tilde{H}(\tau)\right) \left|\tilde{\psi}(\tau)\right\rangle \right\| \tag{3.43}$$

$$\leq \int_0^t \mathrm{d}\tau \left\| V \sum_i f(n_i) \left|\tilde{\psi}(\tau)\right\rangle \right\|. \tag{3.44}$$

Just as before, we use the fact that at short times, the boson number in each cluster (and hence on each site) is bounded. Specifying to the case of Bose-Hubbard interactions, we have

$$\varepsilon \leq \int_0^t \mathrm{d}\tau \frac{V}{2} \left\| \sum_i n_i(n_i - 1) \left|\tilde{\psi}(\tau)\right\rangle \right\|. \tag{3.45}$$

3.11 Types of transitions

In this section, we study in more detail the number of encoded logical qubits in the system inherent in the hardness proof, which we define more carefully below. Our definition is valid for any family of Hamiltonians and the definition for circuit architectures is analogous. Given constraints on the evolution time t (taken to be the depth in case of circuits), we would

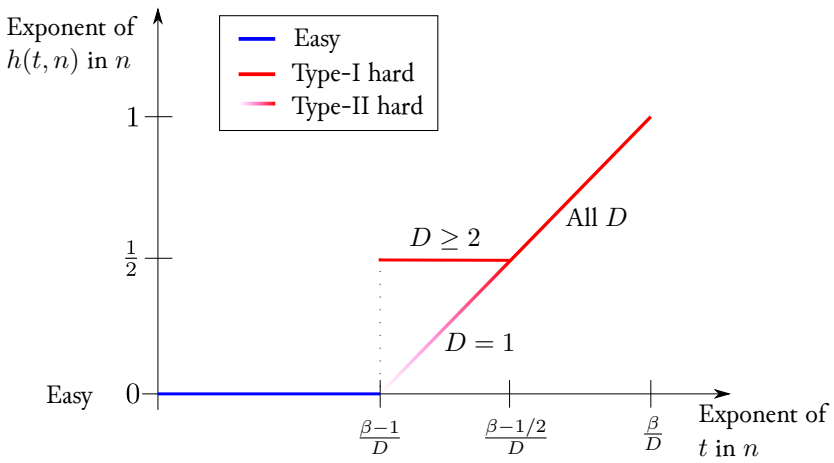

Figure 3.8: Schematic of types-I and II complexity phase transitions. The points on the X-axis, $(\beta - 1)/D$ and β/D, are different for different dimensions and therefore the lines do not really intersect in the way depicted here.

like to reduce the problem of simulating arbitrary circuits of nearest-neighbor gates in one dimension acting on m qubits for depth m to the problem of simulating postselected Hamiltonian evolution for time t under the given Hamiltonian. We call this reduction exploiting postselection an "embedding". For a given reduction, we define the quantity $h(t, n)$ to be the number of qubits m that we can embed into a Hamiltonian evolution on n qubits/particles evolving for time t. The motivation for defining this quantity is that if the function $h(t, n)$ is at least some (possibly sublinear) polynomial in n, then simulating time evolution for time t under the Hamiltonian is hard modulo some conjectures.

3.11.1 Defining phase transitions

The function $h(t, n)$ with respect to a certain reduction is an increasing function of the time t for a fixed n, since given more time, the number of qubits that can be embedded cannot decrease. At the transition timescale t_*, the quantity h transitions from being subpolynomial

in n to polynomial in n as the system transitions from being easy to simulate to being hard to simulate. The 2D and 1D phase diagrams in Fig. 3.1 can also be viewed in terms of the function h of the corresponding family of Hamiltonians.

Figure 3.8 shows the behavior of $h(t, n)$ under our hardness reductions for the two types of transitions. We can see from the figure that if these lines are correct, sharp transitions are akin to a I-order phase transition, and coarse transitions like II-order transitions. Specifically, if the exponent of n in $h(t, n)$ has a discontinuity with respect to the exponent of n in the evolution time t, we have a I-order transition. For II-order transitions, on the other hand, there is no discontinuity in the exponent of n in $h(t, n)$. The exponent may be defined as $\lim_{n\to\infty}(\log h(t, n)/\log n)$. Therefore, the exponent and hence the transition is only well-defined in the thermodynamic limit $n \to \infty$, just as in regular phase transitions where non-analyticities are only visible in the thermodynamic limit.

Let us look at how we obtained the lines. First, the region before exponent $\frac{\beta-1}{D}$ corresponding to time $o(L)$ is easy in all dimensions from our easiness results, meaning h necessarily scales smaller than any polynomial in n in the blue region. Also, we know from our hardness results that near β/D (i.e. time $t = O(Ln^{1/D})$) we have enough time to "touch" all n qubits, giving $h(t, n) = n$ for all dimensions. This can be extended to hardness for $h(t, n) = n^\delta$ for all $\delta \in (0, 1]$ in all dimensions, with the corresponding time $Ln^{\delta/D} = n^{(\beta-1+\delta)/D}$. For $D \geq 2$ on the other hand, we can do better: we can always encode $\sqrt{n}$ qubits even if the time is $\Theta(L)$. This strategy is better than the first strategy for time $t < n^{(\beta-1/2)/D}$. It is not known if these curves for the function h are optimal, since for optimality we should rule out other reductions that might achieve a better scaling.

To sum up, considering the h-index gives a complementary way of looking at complex-

ity phase transitions of the sort we study. This is a fine-grained view of looking at complexity phase transitions, since we care not just about whether the h-index is polylog (easy) or poly (hard), but how exactly it scales with n as a function of the evolution time.

3.12 Outlook

We have mapped out the complexity of the long-range Bose-Hubbard model as a function of the particle density β, the degree of locality α, the dimensionality D, and the evolution time t. A particularly interesting open question concerns the regions of the phase diagram without definitive easiness/hardness results. These gaps are closely related to open problems in other areas of many-body physics and quantum computing. In the nearest-neighbor limit, there is no gap between t_{easy} and t_{hard}. When α is finite, closing the gap is closely tied to finding state-transfer protocols which saturate Lieb-Robinson bounds. Stronger Lieb-Robinson bounds can increase t_{easy}, and faster state-transfer can reduce t_{hard}, as evidenced by the improvement over the original version of the manuscript [93] due to results from Ref. [173]. These observations show that studying complexity phase transitions provides a nice testbed for, and gives an alternative perspective on results pertaining to the locality of quantum systems.

Our results can be easily adapted to a wide range of Hamiltonians in atomic physics and quantum information. Several interesting Hamiltonians are special cases of the one that we study, or straightforward extensions. For example, some experimental platforms such as cold atoms, neutral atoms, and trapped ions can be naturally mapped onto our model. This model is also ideal for studying interesting phenomena in quantum information such as dy-

namics in long-range interacting systems, quantum computational supremacy, entanglement phase transitions, and models of modular networks. We elaborate on these connections in Section 3.5.

If the qualitative features of the phase diagram we have derived for the Bose-Hubbard model hold more generally, our results may hint at a notion of universality present in transitions between complexity phases. In 1D, we have proved that the transition is always coarse. However in 2D and higher, when there are interactions, the transition is sharp. In contrast, in 2D and higher for non-interacting transitions, it is unknown whether the transition is coarse or sharp, and the classification depends on approximate sampling hardness of constant-depth boson sampling. This dependence on the dimension and possible dependence on whether the system is interacting or noninteracting suggests the possibility of classifying complexity phases of matter, and the transitions between them, based on generic features such as connectivity, dimensionality, and kinds of interactions.

Along this line, it would be interesting to study whether similar features occur for different kinds of complexity phase transitions. In our work, the transition occurs in the dynamics of a many-body Hamiltonian. However, different approaches are possible. For example, one could consider open quantum systems, where decoherence might drive dynamical transitions from hard to easy. A particularly rich class of open quantum systems are random quantum circuits with interspersed measurements [193–195]. These have been used recently to study new non-equilibrium phases, entanglement phase transitions, and could be a promising platform to study complexity phase transitions.

Chapter 4: Complexity of Fermionic Dissipative Interactions and Applications to Quantum Computing

Understanding whether a particular quantum system is easy or hard to simulate from the perspective of classical computation is a crucial task serving several goals. The first goal, as a primary step of many numerical studies, is to find efficient classical algorithms describing the desired quantum phenomena. Another goal arises in quantum computing, where finding many-body systems lacking an effective classical description may be worthwhile for constructing quantum computation [196] and simulation [3; 197] devices. The versatility of the classical simulability problem can be illustrated by considering the sampling problem for noninteracting and interacting fermions [14; 50; 51; 198]. There are efficient classical algorithms to simulate fermions described by a quadratic Hamiltonian: the amplitudes of time-evolved many-body configurations are expressed by an efficiently-computable analytical formula [51; 199]. The existence of an efficient algorithm makes the free-fermion approximation a numerically accessible and valuable method with applications to condensed-matter systems. At the same time, simulating interacting fermions is believed to be classically intractable. Indeed, simulating general interacting fermions is as hard as simulating the output of a universal quantum computer (see, for example, Ref. [166]). A similar practical differentiation between easy and hard problems can be applied to other systems [18; 19; 66; 95; 96].

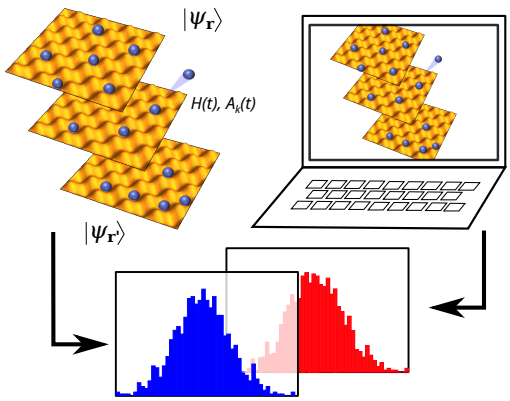

Figure 4.1: **Classical simulability**. We look for the existence of an efficient algorithm running on a classical computer and producing (sampling) the many-body configurations with the probability distribution close to the physical system after measurement in some basis. We show that, for fermionic systems with Hamiltonian $H(t)$ an with dissipation described by quadratic Lindblad jump operators $A_k(t)$, such an algorithm exists for at least a restricted number of problems, while the worst-case scenario requires a quantum computer in order to be solved efficiently. The three optical lattices illustrate the state of the system at initial, intermediate, and final times.

In this Chapter, we study the fate of classical simulability of fermionic systems in the presence of dissipation, both for computing local observables and for sampling from the many-body output distribution (to be defined shortly). To obtain a classification of the complexity of simulating free fermions with dissipation, we consider a general class of Markovian processes, i.e. dynamics that depends only on the instantaneous system state and is independent of the preceding evolution [200]. In previous studies, it was shown that Markovian single-fermion loss or gain terms keep the noninteracting system classically tractable [201; 202]. As a step forward, we consider a much wider class of *quadratic-linear* Lindblad jump operators. Using the method of stochastic trajectories [203; 204], we establish a wide subclass of problems that can be simulated classically. At the same time, we demonstrate that, in general, quadratic Lindblad jump operators are at least as hard to simulate as most elastically interacting systems. More precisely, we establish a connection between dissipative

interactions and fault-tolerant universal quantum computation exploiting the quantum Zeno effect [205–209]. Therefore, simulating evolution under quadratic Lindblad jump operators is equivalent in power to quantum computation. The tractability and intractability results together show that simulation of quartic dissipative Liouvillian operators is a problem whose complexity can be changed from hard to easy by varying one or more parameters in the system [96].

One motivation behind this Chapter is the existence of a variety of accessible fermionic physical systems involving inelastic processes. Examples of dissipative processes described by quadratic Lindblad jump operators include two-body loss in trapped alkali atoms [210–212], alkaline-earth atoms [213–218], and cold molecules [164; 219]. As we will show, Feshbach resonances [220; 221] can be used to significantly suppress coherent interactions between cold atoms, simultaneously increasing the rate of atom-pair trap losses. More general types of dissipation can be created by adding a source of atoms [222–224] or inelastic photon scattering [225–227]. In condensed matter physics, an example of a process that can potentially be described by a Lindblad equation in the Markovian approximation is Cooper-pair loss [228; 229] and dephasing [230]. Recent progress in the control of dissipative electronic systems has brought them into focus in condensed matter physics. Some of the novel effects in noninteracting and mean-field fermionic systems include dissipation-induced magnetism [231–233], dissipative superfluids and superconductors [234; 235], dissipative Kondo effect [236; 237], non-Hermitian topological phases [238–245], and non-Hermitian localization [246–248].

We provide a classification of dissipative fermionic processes into *easy* (efficiently simulable) and *hard* (not efficiently simulable) classes according to their worst-case computational

complexity. The classical simulability problem may be phrased in two ways, either in terms of evaluation of few-body observables or sampling from the full probability distribution on many-body outcomes. In the first task (*few-body observables*), a classical computer is required to output the expectation value of an observable supported on k sites, where k does not grow with the system size. In the second task (*sampling*), a classical computer is tasked with producing samples from the same distribution as the one obtained by measuring the time-evolved state in some canonical basis (see Fig. 4.1). Both tasks allow for the computer to make a small error ϵ, measured appropriately in each case[1]. The task of sampling is computationally harder; an algorithm producing samples in some product-state basis can also be used to obtain expectation values of few-body observables in the same basis. Therefore, in this Chapter, we focus mainly on the easiness of sampling in arbitrary product-state bases as a criterion for overall easiness and on the hardness of computing few-body observables as a criterion for overall hardness. This choice gives the stronger of the two results for both easiness and hardness.

We note that a limited version of classical simulability for some models below was also studied in previous works [249–251]. In particular, it was shown that two-point correlation functions in such models can be evaluated via solving a closed set of equations. This result establishes classical simulability of local observables and can be used in various problems such as dissipative transport or optical response. However, this result alone is not sufficient to establish the simulability of sampling. In contrast to local observables, simulated sampling requires the full knowledge of the many-body output probabilities, therefore the sampling

[1]To be precise, for the first task (estimating few-body observables), the error is measured by the maximum difference in the estimated and ideal expectation values of a unit-norm observable O. For the second task (sampling), the error is measured through the maximum variation distance between the ideal and the sampled probability distributions. The definitions are elaborated upon in Appendix A.

complexity of systems with simulable low-order correlations remains unclear. As we revisit below, Gaussian systems are the exception that allow reducing these output probabilities to two-point correlation functions via Wick's theorem; other systems we study below do not have such simple reduction. To overcome this problem, we develop the easiness proof that does not require applicability of Wick's theorem. In conclusion, sampling is a stronger notion of simulation compared to two-point correlators in a sense that any local observables can be efficiently obtained using an oracle producing sampling outcomes

Let us emphasize the importance of the provided complexity analysis. While established easy dissipation classes are limited to certain fine-tuned processes, such limited simulable models have an important application for quantum computing. For example, classical models can be used in calibration of quantum computers, simulation of the impact of noise on sampling, and analysis of fermionic quantum error-correcting codes [252]. More fundamentally, identifying easy classes is an important first step to analyze easy-hard transitions in open fermionic systems, as we analyze below. At the same time, the hardness result we obtain in this Chapter establishes utilizing dissipative interactions as an alternative path toward building a universal quantum computer. This conclusion is surprising, since dissipative interactions generally produce mixed states. Therefore, dissipative interactions should be used only in a manner utilizing a blockade mechanism induced by the quantum Zeno effect, as we show below. In cold atomic systems, controlling dissipative interactions differs from photonic systems studied before [207; 208] and can be achieved using an atomic Feshbach resonance. In Ref. [253], we analyze in detail a scheme for universal quantum computation with ^{40}K atoms and illustrate that, with realistic experimental parameters, an entangling gate with low error-rate of roughly 10^{-4} is possible. The existence of both easy and hard

classes for two-body dissipation establishes it as a valuable model for physical analysis of noisy intermediate-scale quantum devices.

4.1 Model

We consider dynamics generated by the Lindblad master equation [200; 254]

$$\frac{d\rho}{dt} = -i[H(t), \rho] + \sum_{k=1}^{k_A} A_k(t)\rho A_k^\dagger(t) - \frac{1}{2}\{A_k^\dagger(t)A_k(t), \rho\}, \tag{4.1}$$

where $\{X, Y\} \equiv XY + YX$ is the anticommutator, $\rho(t)$ is the density matrix of the system, $H(t)$ is a noninteracting Hamiltonian, and $A_k(t) \in \mathcal{A}(t)$ form a set of k_A Lindblad jump operators. We set $\hbar = 1$ throughout the Chapter unless specified otherwise. Both the Hamiltonian and the Lindblad jump operators may depend explicitly on the time but not on the state itself. The corresponding map $\rho(t_1) = \mathcal{V}(t_1, t_2)\rho(t_2)$ between arbitrary times t_1 and $t_2 \geq t_1$ satisfies $\mathcal{V}(t_2, t_1) = \mathcal{V}(t_2, \tau)\mathcal{V}(\tau, t_1)$ for any $t_2 \geq \tau \geq t_1$. This divisibility condition is commonly referred to as the most general definition of Markovian dynamics [255]. The master equation in Eq. (4.1) is invariant under certain transformations of the set of Lindblad jump operators $\mathcal{A}(t)$, such as operator permutations, multiplying any Lindblad operator by a phase factor, or splitting/merging of identical operators as $A_k \rightleftarrows \{\sqrt{p}A_k, \sqrt{1-p}A_k\}$, $0 \leq p \leq 1$.

As a physical system of interest, we consider a fermionic many-body problem where $N \leq L$ spinless fermions initially occupy L available levels (modes). Such systems are commonly described by the second quantization method, which expresses any operator, including the Hamiltonian and Lindblad jump operators, in terms of fermionic Fock operators $c_n^\dagger$ and

Type	Examples of A_k	Complexity
Dephasing	$c_1^\dagger c_1$	Easy (EC1)
Particle shuffle	$c_1^\dagger c_2$ & $c_2^\dagger c_1$	
Classical fluctuations	$c_1^\dagger$ & c_1	
Classical pair fluctuations	$c_1^\dagger c_2^\dagger$ & $c_1 c_2$	
Mixing unitaries	$2c_1^\dagger c_1 - 1 + i(c_2^\dagger + c_2)$	Easy (EC2)
Single-particle loss/gain	c_1 OR $c_1^\dagger$	Easy (EC3)
Incoherent hopping	$c_1^\dagger c_2$	Hard
Pair loss/gain	$c_1 c_1$ OR $c_1^\dagger c_1^\dagger$	

Table 4.1: Comparison between different types of noninteracting fermion dynamics with additional dissipation. For simplicity, we provide examples for two modes out of L, denoted by numbers 1 and 2. The symbol & means that both operators are present in the set $\mathcal{A}(t)$ with factors equal in absolute value. Abbreviations EC1, EC2, EC3 stand for Easy Class 1, 2, and 3 described in the text.

c_n, $n \in \{0, 1, \dots L-1\}$. Fock operators create and annihilate a single fermion in a particular mode and satisfy the canonical commutation relations $\{c_n, c_m\} = 0$, $\{c_n^\dagger, c_m\} = \delta_{nm}$. Though the conventional fermion operators are suitable in most physical problems, in the absence of fermion number conservation it is convenient to use the $2L$ Hermitian Majorana fermion operators $\gamma_{2n} = c_n + c_n^\dagger$ and $\gamma_{2n+1} = -i(c_n - c_n^\dagger)$, due to their simple anticommutation relations $\{\gamma_i, \gamma_j\} = 2\delta_{ij}$, $i, j \in \{0, 1, \dots, 2L-1\}$. We consider the most general form of a noninteracting Hamiltonian [256]

$$H(t) = \frac{i}{2} \sum_{i,j=0}^{2L-1} \alpha_{ij}(t) \gamma_i \gamma_j + \sum_{i=0}^{2L-1} \beta_i(t) \gamma_i, \tag{4.2}$$

where $\alpha(t)$ is a real-valued antisymmetric $2L \times 2L$ matrix and $\beta(t)$ is a real $2L$ vector. We assume that the magnitude of all entries of $\alpha(t)$ and $\beta(t)$ and their time derivatives scale at most polynomially with system size.

In this Chapter, we focus on the classical resources needed to approximately sample

from the fermion distribution at time t,

$$P(\mathbf{r}|\mathbf{r}') = \langle\psi_{\mathbf{r}}|\rho(t)|\psi_{\mathbf{r}}\rangle, \rho(0) = |\psi_{\mathbf{r}'}\rangle\langle\psi_{\mathbf{r}'}|, \tag{4.3}$$

where $\mathbf{r}'$ and $\mathbf{r}$ denote the positions of occupied modes in the initial and final (measured) product-state configurations, respectively, and $|\psi_{\mathbf{r}}\rangle$ is a product state defined as $|\psi_{\mathbf{r}}\rangle = c^\dagger_{r_1} \dots c^\dagger_{r_N}|0\rangle = \gamma_{2r_1} \dots \gamma_{2r_N}|0\rangle$. Importantly, because the dynamics may not conserve the total fermion number, the final number of fermions $\tilde{N}$ can, in general, be different from the initial number: $N \neq \tilde{N}$.

We establish the sufficiency of polynomial resources for classically simulating dynamics due to arbitrary noninteracting Hamiltonians in Eq. (4.2) and a limited set of Lindblad jump operators $A_k(t) \in \mathcal{A}(t)$ in the worst case. In order to prove polynomial-time simulability (also called easiness) for limited classes of dissipative dynamics, we reduce the problem to that of simulating unitary noninteracting fermionic evolution, an easy problem for a classical computer. In order to prove hardness for more general Lindblad jump operators, we exploit the ability of dissipative dynamics to perform arbitrary quantum computation (i.e. we prove that simulating universal quantum computation reduces to simulating Lindbladian dynamics).

The results of this Chapter are briefly illustrated in Table 4.1. First of all, we define three classically tractable classes of Lindblad jump operators (defined as Easy Classes 1, 2, and 3). All of these cases allow for polynomial-time sampling of any Hamiltonian and Lindblad jump operators from the given class on a classical computer, with error scaling inverse-polynomially with L. Easy Class 1 (EC1) allows for simulation of self-adjoint sets of quadratic

Lindblad jump operators: all Lindblad jump operators in the set $\mathcal{A}(t)$ come with their Hermitian conjugate. This class includes such widely used examples as dephasing, incoherent particle shuffle, and classical fluctuations of the number of fermions and of the number of fermion pairs. Easy Class 2 (EC2) works with unitary quadratic Lindblad jump operators. Finally, Easy Class 3 (EC3) describes the loss or gain of a single particle in the system and can be used in combination with EC1 and/or EC2. At the same time, there exists a class of Lindblad jump operators with a nonzero measure that is hard to classically simulate. Examples from this class include pair loss/gain and incoherent fermion hopping. Below we explore each class separately.

We focus on quadratic-linear Lindblad jump operators of the form

$$A_k(t) = \frac{i}{2} \sum_{i,j=0}^{2L-1} a_k^{ij}(t)\gamma_i\gamma_j + \sum_{i=0}^{2L-1} b_k^i(t)\gamma_i + d_k(t), \tag{4.4}$$

where $a_k(t)$ and $b_k(t)$ are arbitrary complex-valued $2L \times 2L$ matrices and $2L$-vectors respectively, and $d_k(t)$ is a number. In this problem, we assume that the number k_A of nontrivial Lindblad jump operators from this class is at most $L(L+1)$. In fact, any instance where $\mathcal{A}$ has a larger number of operators can be reduced to a smaller set through a linear transformation [200]. Also, as with the Hamiltonian, we assume that the magnitude of the entries of $a_k(t), b_k(t)$, and $d_k(t)$ and their time derivatives grow at most polynomially with the system size.

This Chapter is organized as follows. In Section 4.2, we provide a brief introduction to free-fermion sampling, recalling established results in the literature and connecting them to the most general form of quadratic-linear Hamiltonians. In Section 4.3, we derive three new

algorithms allowing us to solve distinct classes of fermionic problems involving quadratic Lindblad jump operators and prove that these algorithms run in time that is polynomial in both L and the inverse of the distance from the exact distribution. In Section 4.4, we establish generic class of systems that belong to the hard class and show their robustness to the presence of minor imperfections.

4.2 Free-fermion sampling

In this Section, we discuss the noninteracting fermion problem in the absence of dissipation. We recap the work of Terhal and DiVincenzo [51], which shows that all output probabilities $P(\mathbf{r}|\mathbf{r}')$ in Eq. (4.3) and the marginal probabilities can be obtained using a classically tractable analytical formula. Before referring to this result, we need to incorporate the linear terms present in Eq. (4.2) into effective quadratic dynamics. In order to do so, we consider a slightly larger system containing an extra ancilla $(L+1)$th mode [256], labeled as $n = L$. Next, we choose new effective dynamics such that the ancilla mode remains in the state $|+\rangle \equiv (|0\rangle + |1\rangle)/\sqrt{2}$ during the entire evolution, including the initial and final times, i.e.

$$|\psi_{\mathbf{r}'}\rangle \to |\psi_{\mathbf{r}'}\rangle \otimes |+\rangle, \qquad |\psi_{\mathbf{r}}\rangle \to |\psi_{\mathbf{r}}\rangle \otimes |+\rangle. \tag{4.5}$$

To construct such dynamics, we consider a new Hamiltonian by replacing $\gamma_i \to i\gamma_i\gamma_{2L}$, where γ_{2L} and γ_{2L+1} are Majorana operators acting on the ancilla mode. It is straightforward to check that such a transformation results in a new purely quadratic Hamiltonian (without any linear terms) that keeps the state of the ancilla stationary and does not modify the dynamics

of the original Hamiltonian. The new coefficients in Eq. (4.2) are

$$\alpha_{ij} \to \tilde{\alpha}_{ij} = \alpha_{ij} + \delta_{i2L+2}\beta_j - \delta_{j2L+2}\beta_i, \tag{4.6}$$

where we use by default $\beta_{2L} = \beta_{2L+1} = 0$. Given that the modified initial and final conditions for the system and the ancilla are $\{\mathbf{r}'\} \to \{\mathbf{r}', s'\}$, $\{\mathbf{r}\} \to \{\mathbf{r}, s\}$, $s, s' \in \{0, 1\}$, the probability $P(\mathbf{r}|\mathbf{r}')$ of obtaining outcome $\mathbf{r}$ for the original system can be computed from the probability $P(\{\mathbf{r}, s\}|\{\mathbf{r}', s'\})$ for the system with the ancilla as follows:

$$P(\mathbf{r}|\mathbf{r}') = \frac{1}{2} \sum_{s,s' \in \{0,1\}} P(\{\mathbf{r}, s\}|\{\mathbf{r}', s'\}). \tag{4.7}$$

Summarizing, this method ensures that the dynamics of a linear-quadratic Hamiltonian can always be reduced to the dynamics of a quadratic one by expanding the system size by one mode. Therefore, we henceforth consider only quadratic Hamiltonians.

Let us derive the formula for the sampling probability. We start from a (backwards) time-evolved Majorana fermion operator $\gamma_i(t) = U_t \gamma_i U_t^\dagger$, where $U_t = \mathcal{T} \exp\left(-i \int_0^t H(t')dt'\right)$. Here $\mathcal{T}$ exp is the standard time-ordered exponential. Given the quadratic structure of the Hamiltonian, this evolution is a linear transformation $\gamma_i(t) = \sum_i R_{ij}(t)\gamma_j$, where $R = \mathcal{T} \exp\left(-2 \int_0^t \alpha(t')dt'\right)$ is a unitary $2L \times 2L$ matrix. One can use this expression to derive the time evolution of a fermion operator as

$$U_t c_n U_t^\dagger = \frac{1}{2} U_t(\gamma_{2n} + i\gamma_{2n+1})U_t^\dagger = \sum_j T_{nj}\gamma_j, \tag{4.8}$$

where $T_{nj} \equiv R_{2n,j} + iR_{2n+1,j}$ are elements of a $L \times 2L$ transformation matrix T. Labeling

the initially empty sites as l'_i and recalling that the initial fermion positions are r'_i and that the final positions are r_i, the linearity allows to write the output probability in Eq. (4.3) at any time as

$$P(\mathbf{r}|\mathbf{r}') = \langle\psi_{\mathbf{r}}|\, U_t\, |\psi_{\mathbf{r}'}\rangle\langle\psi_{\mathbf{r}'}|\, U_t^\dagger\, |\psi_{\mathbf{r}}\rangle \tag{4.9}$$

$$= \langle\psi_{\mathbf{r}}|U_t c^\dagger_{r'_1} c_{r'_1} \dots c_{l'_{L-N}} c^\dagger_{l'_{L-N}} U_t^\dagger|\psi_{\mathbf{r}}\rangle \tag{4.10}$$

$$= \sum_{n_1,\dots n_L;m_1,\dots m_{L-N}} T^*_{r'_1 n_1} T_{r'_1 m_1} \dots T_{l'_{L-N} m_L} T^*_{l'_{L-N} n_L} \times \tag{4.11}$$

$$\langle 0|\gamma_{2r_N} \dots \gamma_{2r_1}\gamma_{n_1}\gamma_{m_1} \dots \gamma_{m_L}\gamma_{n_L}\gamma_{2r_1} \dots \gamma_{2r_N}|0\rangle. \tag{4.12}$$

This expression can be computed efficiently using Wick's theorem and written in a compact form. Let $\mathcal{I}$ be a subset of indices with increasing order and $A[\mathcal{I}]$ be the matrix whose elements satisfy $A[\mathcal{I}]_{ij} \equiv A_{\mathcal{I}_i,\mathcal{I}_j}$. Consider the set $\mathcal{I} = \{r'_i, L + l'_j, 2L + 2r_k\}$, where $i \in \{1, 2, \dots N\}$, $j \in \{1, 2, \dots L - N\}$, and $k \in \{1, 2, \dots \tilde{N}\}$ take all possible values. Then the result can be written as [51]

$$P(\mathbf{r}|\mathbf{r}') = \mathrm{Pf}\, M[\mathcal{I}], \tag{4.13}$$

where Pf is the Pfaffian, and M is a $4L \times 4L$ matrix

$$M = \begin{pmatrix} T\Lambda T^T & T\Lambda T^\dagger & T\Lambda \\ T^*\Lambda T^T & T^*\Lambda T^\dagger & T^*\Lambda \\ \Lambda T^T & \Lambda T^\dagger & I \end{pmatrix}, \tag{4.14}$$

where, in turn, the $2L \times 2L$ matrix Λ is

$$\Lambda = I_{L\times L} \otimes \begin{pmatrix} 1 & i \\ -i & 1 \end{pmatrix}. \tag{4.15}$$

The expression in Eq. (4.13) can be efficiently evaluated on a classical computer using existing polynomial-time algorithms for computing Pfaffians [199]. Similarly, marginal probabilities can be efficiently computed conditioning on the output of a given fraction of sites, as in Ref. [51], which is enough to efficiently sample from the output probability distribution.

4.3 Easy Classes

Here we present three algorithms that allow simulating specific fermionic problems involving quadratic Hamiltonians and quadratic-linear Lindblad jump operators. All methods are based on stochastic unraveling, i.e., replacing dissipative dynamics by a stochastic free-fermion Hamiltonian without changing the outcome distribution (see also Ref. [204]). Since each stochastic realization can be simulated efficiently by a classical computer, as established in the previous section, a classical computer can serve as a black box sampler that reproduces measured outcomes. In this Section, we demonstrate that the classes of problems belonging to the aforementioned Easy Classes 1–3 are efficiently simulable. In particular, we show that these problems require computation resources C (number of operations) bounded as $C \leq \text{poly}\,(L, t^2/\epsilon)$ to sample from a distribution that is ϵ-close to the target distribution. Therefore, we establish efficient classical algorithms for approximate dissipative fermion sampling in the presence of certain classes of quadratic-linear Lindblad jump operators.

4.3.1 Efficient classical algorithms

Let us define *Easy Class 1* (EC1) as problems that involve quadratic-linear self-adjoint Lindblad sets $\mathcal{A}(t)$ defined as follows. We assume that at any time one can divide the set as a union of two equal-size subsets, $\mathcal{A} = \mathcal{A}_1 \cup \mathcal{A}_2$, where the Hermitian conjugate of every Lindblad operator in $\mathcal{A}_1$ returns an operator in $\mathcal{A}_2$ (and vice versa). Under this division, any Hermitian Lindblad operator must be included in both subsets $\mathcal{A}_1$ and $\mathcal{A}_2$ with normalization factor $1/\sqrt{2}$. The latter splitting can be seen as a transformation that keeps the Lindblad equation invariant, as defined earlier below Eq. (4.1). Examples from EC1 include several important physical models such as dephasing and classical fluctuations (see examples of sets in lines 1–4 in Table 4.1).

In previous works, it was shown that such systems have two-point correlation functions that are classically simulable by solving a closed set of linear equations [249–251]. This is indeed a strong indication that the system can be simulable in the broader context of sampling complexity. However, as we noted previously, Wick's theorem is not applicable to non-Gaussian states. This means that the scheme we utilized to obtain Eq. (4.13) does not work any more. We now show an alternative scheme using stochastic unraveling that leads us to the easiness result.

To efficiently simulate dynamics from EC1, we consider a stochastic Hamiltonian

$$H'(t) = H(t) + \sum_{A_k \in \mathcal{A}_1} \theta_k(t) A_k(t) + \theta_k^\dagger(t) A_k^\dagger(t), \tag{4.16}$$

where $\theta_k(t)$ is a complex random variable taking constant values $\theta_k(t) = \xi_{nk}/\sqrt{\Delta\tau}$ during

short time intervals $t \in [n\Delta\tau, (n+1)\Delta\tau]$. The discrete complex Gaussian variables ξ_{nk} satisfy $\mathbb{E}\xi_{nk} = 0$, $\mathbb{E}\xi^*_{nk}\xi_{n'k'} = \delta_{nn'}\delta_{kk'}\delta_{ab}$, where $\mathbb{E}$ denotes the expectation value taken over the random variables. Then, given a stochastic Hamiltonian of the form in Eq. (4.16), the original dynamical map $\mathcal{V}(t_2, t_1)$ generated by the Lindblad equation can be approximated as

$$\mathcal{V}(t_2, t_1) = \mathbb{E}\mathcal{V}_{\text{st}}(t_2, t_1) + \delta\mathcal{V}(t_2, t_1)\Delta\tau + O(\Delta\tau^2), \tag{4.17}$$

where $\delta\mathcal{V}(t_2, t_1)\Delta\tau$ is the lowest-order correction (to be explicitly derived below) and $\mathcal{V}_{\text{st}}$ is a stochastic unitary map

$$\mathcal{V}_{\text{st}}(t_2, t_1)\rho = U(t_2, t_1)\rho U^\dagger(t_2, t_1). \tag{4.18}$$

In the above, $U(t_2, t_1) = \mathcal{T}\exp\left(-i\int_{t_1}^{t_2} dt' H'(t')\right)$ encodes the time-evolution due to $H'(t)$ in Eq. (4.16). The average $\mathbb{E}$ in Eq. (4.17) is taken over the stochastic fields $\theta_k(t)$. The resulting output probabilities satisfy

$$P(\mathbf{r}|\mathbf{r}') = \mathbb{E}\, P_{\text{st}}(\mathbf{r}|\mathbf{r}') + O(\Delta\tau), \tag{4.19}$$

where $P_{\text{st}}(\mathbf{r}|\mathbf{r}')$ is the output probability for the unitary dynamics in Eq. (4.18) obtained via the formula in Eq. (4.13). Therefore, a computer programmed to sample from the distribution for a randomly chosen set of unitary trajectories will produce outcomes with the same probabilities as the physical process following Lindbladian evolution, up to $O(\Delta\tau)$ error. The cost of suppression of this error in terms of computational resources will be discussed

later in this section. Here we just specify that the correction to the dynamical map, which we treat as an error, can be expressed as

$$\delta\mathcal{V}(t_2,t_1) = \mathbb{E}\int_{t_1}^{t_2} dt' \mathcal{V}_{\text{st}}(t_2,t')\mathcal{D}(t')\mathcal{V}_{\text{st}}(t',t_1), \tag{4.20}$$

where $\mathcal{D}(t)$ is a time-local superoperator

$$\mathcal{D}(t)\rho = \sum_{\alpha} D_{\alpha}^{(1)}(t)\rho D_{\alpha}^{(2)}(t). \tag{4.21}$$

Here, the operators $D_{\alpha}^{(i)}(t) = \text{poly}_4(H(t), A_k(t))$ can be expressed as polynomials of degree four in terms of the Hamiltonian and Lindblad jump operators at time t. Therefore, $D_{\alpha}^{(i)}(t)$ can always be presented as a sum of terms, each being a product of no more than eight Majorana operators. The specific form of these operators and the derivation of Eq. (4.17) can be found in Section 4.5.

Although the proposed unraveling scheme represents dynamics in terms of stochastic trajectories for Gaussian pure states, the resulting averaged mixed state is non-Gaussian, in contrast to previously studied problems [201; 202]. Therefore, the overall dynamics of EC1 represents dissipative interactions of fermions, while the proposed method can be seen as a good choice of time-dependent density matrix decomposition.

Let us consider another class of problems, *Easy Class 2* (EC2), that include unitary quadratic Lindblad jump operators $A_k = \sqrt{\Gamma_k(t)}Y_k(t)$, where $\Gamma_k(t) \geq 0$ are time-dependent rates and $Y_k(t) = \exp(-iG_k(t))$ are unitary operators generated by quadratic-linear Hamiltonians $G_k(t)$ of the form in Eq. (4.2). To classically simulate dynamics under

EC2, we also consider discretized time evolution with sufficiently small timesteps $\Delta\tau$ and set the unitary transformation $U(t_1, t_2) = \prod_{n=n_1}^{n_2} U_n$, where the timestep unitaries U_n are generated randomly according to the rule

$$U_n = U_n^0 \times \begin{cases} Y_k(t_n), & p_k = \Gamma_k(t_n)\Delta\tau, \\ I, & p_0 = 1 - \sum_k \Gamma_k(t_n)\Delta\tau. \end{cases} \tag{4.22}$$

Here p_k are probabilities that are used to generate the respective outcomes, $U_n^0 = \mathcal{T} \exp\left(-i \int_{n\Delta\tau}^{(n+1)\Delta\tau} H(t)dt\right)$, and $t_n \in [n\Delta\tau, (n+1)\Delta\tau]$ are random times generated from the uniform distribution.

Notwithstanding the slightly different stochastic unraveling, the procedure for approximating EC2 is the same as for EC1. In particular, the system dynamics is described by the expression in Eq. (4.17) leading to the distribution in Eq. (4.19), with the average taken over stochastic unitary realizations. The correction term has the form in Eq. (4.21), but the operators $D_\alpha^{(i)}(t)$ here are degree-two polynomials in the Hamiltonian and Lindblad jump operators. The detailed form of the operators along with the derivation can be found in Section 4.6.

Finally, let us consider *Easy Class 3* described by generic linear Lindblad jump operators $A_k(t) = \sum_i b_k^i \gamma_i + d_k I_{L\times L}$, which can be obtained by setting $a_k = 0$ in Eq. (4.4) without assuming any additional restrictions on the set $\mathcal{A}(t)$. Previous works had already shown that linear jump operators can be simulated classically [201; 202]. However, this proof applies only to Gaussian states and cannot be extended to, for example, Lindblad equations that also contain easy quadratic jump operators. Below we propose a different way of simulating linear

jump operators similar to one we used for EC1. This technique would allow us to combine EC1, EC2, and EC3 into a single easy class of Lindblad equations, including both quadratic and linear processes.

Now let us show that simulation of linear jump operators is equivalent to simulating a unitary system extended by a number of ancilla modes. In particular, we require $L_a = t/\Delta\tau$ ancilla fermion modes equal to the number of time steps after discretization. Let us enumerate the ancilla modes described by Majorana fermion operators γ_{2n} and γ_{2n+1} using indices $n = L, \ldots, L + L_a - 1$. We also assume that the ancilla modes are initialized in the vacuum state and traced out after performing the evolution. The dynamics of both the system and the ancillas can be described as unitary evolution with the Hamiltonian in Eq. (4.16), with one important difference. Now, the quantities $\theta_k(t)$ in the time interval $t \in [n\Delta\tau, (n+1)\Delta\tau]$ are operators instead of numbers, and are given by

$$\theta_k(t) = \xi_{nk}\Delta\tau^{-1/2}(\gamma_{2(L+n-1)} + i\gamma_{2(L+n)-1}). \tag{4.23}$$

The random variables ξ_{nk} are the same as in EC1. The idea is that we pair a loss (gain) term on the system with a gain (loss) term on the ancilla to make the overall Hamiltonian quadratic. After discarding the ancilla modes, the evolution becomes equivalent to the target dissipative dynamics, up to a discretization error that originates from the approximation in Eq. (4.17) and Eq. (4.21), with $D_\alpha^{(i)}(t)$ expressed as degree-four polynomials in the Hamiltonian and Lindblad jump operators, as shown in Section 4.7.

4.3.2 Performance of the classical algorithms

Let us quantify the error of the method of quantum trajectories used for Easy Classes 1–3, and then show that the sampled distribution can be made arbitrarily close to the exact one with an appropriate choice of the timestep $\Delta\tau$. In order to characterize the approximation error ϵ associated with sampling from a distribution $\tilde{P}(\mathbf{r}|\mathbf{r}')$ different from the ideal distribution $P(\mathbf{r}|\mathbf{r}')$, we utilize the total variation distance

$$\epsilon = \frac{1}{2} \max_{\mathbf{r}'} \sum_{\mathbf{r}} \left| \tilde{P}(\mathbf{r}|\mathbf{r}') - P(\mathbf{r}|\mathbf{r}') \right|, \tag{4.24}$$

where the maximization is over all possible initial product-state configurations $\mathbf{r}'$.

Using convexity of the absolute value and the expression for the correction in Eqs. (4.20)–(4.21), the error can be bounded as (see Section 4.8),

$$\epsilon \leq \frac{\Delta\tau}{2} \max_{\mathbf{r}'} \sum_{\alpha} \int_0^t dt' C^{\alpha}_{\mathbf{r}'}(t,t') + O(\Delta\tau^2), \tag{4.25}$$

where

$$C^{\alpha}_{\mathbf{r}'}(t,t') = \mathbb{E} \sum_{\mathbf{r}} \left| \langle \mathbf{r}| D^{(1)}_{\alpha}(t,t') \rho_{\mathbf{r}'}(t) D^{(2)}_{\alpha}(t,t') |\mathbf{r}\rangle \right|. \tag{4.26}$$

Here $D^{(i)}_{\alpha}(t,t') = \mathcal{V}_{\text{st}}(t,t') D^{(i)}_{\alpha}(t')$ and $\rho_{\mathbf{r}'}(t) = \mathcal{V}_{\text{st}}(t,0)\rho_{\mathbf{r}'}$ are operators transformed according to unitary evolution for a single stochastic trajectory, and the average $\mathbb{E}$ is taken over all trajectories. We now use the following lemma to further bound this expression.

Lemma 12. *Consider two sparse operators O_1 and O_2 whose matrix elements satisfy*

$$\langle \mathbf{r}|O_\alpha|\mathbf{r}'\rangle = 0 \quad \text{if} \quad d_H(\mathbf{r},\mathbf{r}') \geq k_\alpha, \quad \alpha \in \{1,2\}, \tag{4.27}$$

where d_H is the Hamming distance, and $\mathbf{r}$, $\mathbf{r}'$ are binary strings of length L representing computational basis states. Let ρ be a normalized positive semidefinite operator, $\rho \geq 0$, $\operatorname{Tr}\rho = 1$, then

$$\sum_{\mathbf{r}} |\langle \mathbf{r}|O_1 \rho O_2|\mathbf{r}\rangle| \leq \frac{L^{k_1+k_2}}{k_1! k_2!} \|O_1\|_{\max}\|O_2\|_{\max}, \tag{4.28}$$

where $\|O_\alpha\|_{\max}$ is the max-norm.

The proof of the lemma can be found in Section 4.8. The result of the Lemma allows us to simplify Eq. (4.26) as

$$C^{\alpha}_{\mathbf{r}'}(t,t') \leq \frac{L^{k_{1\alpha}+k_{2\alpha}}}{k_{1\alpha}! k_{2\alpha}!} \mathbb{E}\|D^{(1)}_\alpha(t,t')\|_{\max}\|D^{(2)}_\alpha(t,t')\|_{\max}, \tag{4.29}$$

where $k_{i\alpha}$ is the locality of the operator $D^{(i)}_\alpha(t,t')$, i.e. the maximum number of Majorana operators in its decomposition. Because $\mathcal{V}_{\text{st}}$ is a map describing free-fermion evolution, the locality $k_{i\alpha}$ of the operator $D^{(i)}_\alpha(t,t')$ is equal to the locality of $D^{(i)}_\alpha(t')$. At the same time, as analyzed in the previous section, the localities of operators $D^{(i)}_\alpha(t')$ satisfy $k_{\mu\alpha} \leq k_{\text{m}}$, where $k_{\text{m}} = 8$ for EC1/EC3, and $k_{\text{m}} = 4$ for EC2. We can also bound the max-norm by the

(spectral) operator norm

$$\|D_\alpha^{(i)}(t,t')\|_{\max} \leq \|D_\alpha^{(i)}(t,t')\| = \|D_\alpha^{(i)}(t')\|. \tag{4.30}$$

As a result, the error bound is given by

$$\epsilon \leq \frac{\Delta\tau}{2}\frac{L^{2k_{\rm m}}}{(k_{\rm m}!)^2}\mathbb{E}\sum_\alpha \int_0^t dt' \|D_\alpha^{(1)}(t')\|\|D_\alpha^{(2)}(t')\|. \tag{4.31}$$

Since the matrices $D_\alpha^{(i)} = \text{poly}(H, A_k)$ are generated by a quadratic-linear Hamiltonian H and set of Lindbladians A_k, we can always find a polynomially large bound for the norm $\|D_\alpha^{(i)}(t')\| \leq \text{poly}(L)$. Thus, there always exists a discretization step

$$\Delta\tau \leq \frac{\epsilon}{t\ \text{poly}(L)} \tag{4.32}$$

that keeps the error in Eq. (4.31) arbitrarily small, suppressed at least polynomially with the number of modes L.

Let us now estimate the amount of computational resources required to perform the above sampling procedure. For each sample, the algorithm randomly chooses a unitary trajectory according to the given prescription for each class EC1–EC3 and, according to the Terhal-DiVincenzo algorithm, samples outputs from the free-fermion distribution in Eq. (4.13). In particular, it samples the output at site i conditioned upon the outcomes sampled at sites $j < i$, for which the marginal probabilities should also be computed. Consider cases of EC1 and EC2 that do not require ancillas. Once the matrix T is obtained in Eq. (4.8), the number of steps to compute the distribution is equal to $C' = L \times O(L^3) = O(L^4)$, where the factor

$O(L^3)$ is the upper bound on the time it takes to compute a Pfaffian of an $O(L) \times O(L)$ matrix. Further, the runtime for obtaining the matrix T is proportional to $t/\Delta\tau \times M(2L)$, where $M(n) \lesssim O(n^3)$ is the time for $n \times n$ matrix multiplication. In sum, the total bound on the runtime for each trajectory is bounded as $C \sim O(L^4) + O\left(L^3 t/\Delta\tau\right)$. Choosing $\Delta\tau = \epsilon/(t \times \text{poly}(L))$, the runtime is

$$C \leq \text{poly}\left(L, \frac{t^2}{\epsilon}\right). \tag{4.33}$$

For EC3, the derivation is the same up to adjusting the system size to include the ancilla modes, $L \to L + t/\Delta\tau$. This case also has a similar polynomial upper bound on the classical runtime in the form of Eq. (4.33) as long as the evolution time t is polynomial.

Finally, let us analyze the case when the conditions of Class 1–3 are violated. Strictly speaking, then the stochastic method fails as it generally maps the problem to a non-quadratic fermionic evolution, which is not believed to be simulable for arbitrarily long time. However, we can still efficiently simulate the system after this mapping if the product of evolution time t and the correction rate $\delta\Gamma$ (of processes violating easiness conditions) remains small. In particular, if the product is bounded as $\delta\Gamma t < c/L^2$ for some constant c, the dynamics remains classically easy. To obtain this result, we consider a more general form of stochastic unraveling in Eq. (4.18) with nonunitary unraveling. This formula can be Taylor-expanded as $\mathcal{V}_{\text{st}} \to \mathcal{V}'_{\text{st}} = \mathcal{V}_{\text{st}} + \delta\Gamma\tau K_1 + (\delta\Gamma\tau)^2 K_2 + \ldots$, where K_n are local correction superoperators and τ is the evolution time. Therefore, we can update the bound for the operator norms in Eq. (4.30) as $\|D_\alpha^{(i)}(t,t')\| \leq \|D_\alpha^{(i)}(t')\| + \delta\Gamma\tau\|K_1 D_\alpha^{(i)}(t')\| + (\delta\Gamma\tau)^2\|K_2 D_\alpha^{(i)}(t')\| + \ldots$. Since the operator $D_\alpha^{(i)}$ involves at most k_{m} fermion Fock operators and the action of K_n

involves at most four fermion operators, we see that $\|K_n D_\alpha^{(i)}(t)\| \leq O(L^{k_m+2n})$. Therefore, the norm $\|D_\alpha^{(i)}(t,t')\|$ is always bounded by a poly(L) value if $\delta\Gamma < c/(L^2\tau)$, where $\tau = t - t' \leq t$. This result leads to Eq. (4.33). As a result, if the dissipation is close enough to the symmetric point, the evolution remains classically easy. This result may be helpful for analyzing the precision needed for implementing this dynamics in intermediate-scale quantum devices.

4.4 Hard class

We have so far demonstrated cases when the probability distribution generated by the Lindblad equation is efficiently simulable on a classical computer. Can we extend these proofs to the most general case of quadratic A_k's? Since quadratic operators A_k correspond to single-fermion jumps in many cases, one may expect that the problem can be solved in the single-particle sector, similar to unitary free-fermion dynamics. However, such an intuition is incomplete. A simple explanation can be obtained using the Fermi exclusion principle that requires the transition between two modes to depend on the occupation of the target mode; thus a quadratic Lindbladian jump operator can induce many-body correlations in the system that quickly become classically intractable.

4.4.1 Reduction from a generic quantum circuit

We now provide a rigorous argument for worst-case hardness based on the equivalence of dynamics under classes $H(t)$ and $A_k(t)$ in Eq. (4.1) on the one hand and universal quantum computing on the other. Let us start with the simplest map utilizing quadratic Hamiltonians.

We distribute all modes into $L/2$ pairs, each pair corresponding to a logical qubit in the state $|0\rangle_L = |01\rangle$ or $|1\rangle_L = |10\rangle$. Then, utilizing only quadratic Hamiltonians and Lindblad jump operators, we can implement any quantum circuit with arbitrary precision. Thus, by showing the equivalence of the dynamics to universal quantum computation, we obtain hardness results for both estimating time-evolved local observables $\langle O(t)\rangle$ and sampling from the time-evolved state in any local basis. The obtained hardness result is therefore on par with the best complexity-theoretic evidence that simulating quantum circuits (in both senses) is hard.

First, using single-fermion hopping between the two sites of a qubit, we can reproduce arbitrary single-qubit operations [175]. Second, to approximately generate a desired two-qubit gate, we can use hopping combined with a quadratic Lindblad operator. In particular, assigning the two-qubit logical states $|00\rangle_L = |0101\rangle$, $|01\rangle_L = |0110\rangle$, $|10\rangle_L = |1001\rangle$, and $|11\rangle_L = |1010\rangle$, the control-Z gate can be implemented by simultaneously applying the hopping Hamiltonian $H = J(c_2^\dagger c_3 + c_3^\dagger c_2)$ and pair-loss operator $A = \Gamma c_3 c_4$ for time $t = \pi/J$, in the limit $\Gamma \gg J$. This type of dynamics can be analyzed as follows. The logical states $|01\rangle_L$ and $|10\rangle_L$ remain invariant in the course of the evolution. At the same time, in the limit $\gamma \equiv \Gamma/J \to \infty$, due to the quantum Zeno effect, the Lindblad operator's action disallows any coherent transition involving states where qubits 3 and 4 are both occupied (i.e. $|\cdots 11\rangle$). As a result, the logical state $|00\rangle_L$ is unaffected by the evolution. Therefore, the only evolving logical state is $|11\rangle_L$, which acquires a phase factor $\exp(i\pi) = -1$ after time $t = \pi/J$. As a result, the effective transformation on the two logical qubits is the control-Z

gate

$$|\psi\rangle \to U_\pi|\psi\rangle, \qquad U_\pi = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}. \tag{4.34}$$

Together with arbitrary single-qubit operations, the control-Z gate is enough to obtain dynamics universal for quantum computing and hence hard to approximately sample from, assuming standard conjectures in complexity theory [14; 30].

Importantly, the performance of the dissipative gate relies on the Zeno-effect blockade effective for $\gamma \to \infty$. In the limit of large but finite γ, the two-qubit system has the probability $\epsilon = 2\pi/\gamma + O(\gamma^{-2})$ of ending up in states $|0011\rangle$ or $|0000\rangle$, which could result in an error in the gate. To avoid computational error, we can choose the ratio γ to be arbitrarily large by taking vanishing $J \to \text{poly}^{-1}(L)$ for any given $\Gamma < 0$. Therefore, we can keep the error below any given threshold at the cost of increased overall computation time, which remains polynomial in system size.

The proposed architecture is not unique and allows for modified/generalized realizations of logical qubits and gates. For example, if the pair decay is always present on any two neighboring modes, one may introduce an empty ancilla mode between two logical qubits in order to ensure that logical states don't decay. As another example, if the control Hamiltonian is linear in terms of Majorana operators, a logical qubit can be encoded using just a single mode. Moreover, for a reader focused on applications, we discuss below a practical

modification of qubit encoding implementable in cold atoms.

Now let us show that pair loss is not the unique dissipation present in the hard class. In fact, this class also includes any quadratic dissipation connected to pair loss by a time-dependent linear Bogoliubov transformation, $A'(t) = \Gamma Y(t)c_1c_2Y^\dagger(t)$, where $Y(t) = \exp(-iG(t))$ is a free-fermion unitary transformation and $G(t)$ is a Hermitian operator from the quadratic-linear class in Eq. (4.2). To demonstrate this equivalence, we consider the pair loss scheme described above but simultaneously replace all pair losses A with $A'(t)$, the Hamiltonian $H(t)$ with $H'(t) = Y^\dagger(t)H(t)Y(t)$, and instead of the initial and final states, choose states transformed by $Y(0)^\dagger$ and $Y(t)$, respectively. The resulting process has the same probability distribution; thus its complexity would be the same. As a result, Lindbladians such as incoherent transitions $A = \Gamma c_1^\dagger c_2$ or pair gains $A = \Gamma c_1^\dagger c_2^\dagger$ are also classically hard in combination with free-fermion dynamics (see Table 4.1).

4.4.2 Robustness of the hardness result

The error associated with imperfect Zeno blockade cannot be arbitrarily suppressed by slowing down the computation if there are small generic corrections to the dissipative dynamics. These corrections can be viewed as the presence of additional Lindblad jump operators with total rate Γ'. Such terms generate additional transitions with the probability $\epsilon' \sim \pi\Gamma'/J$, where Γ' is the combined rate of added operators A' and/or other errors. In contrast to the imperfect-Zeno-blockade error, this type of error diverges for small J. Therefore, there is an optimal value $J \sim \sqrt{\Gamma'/\Gamma}$ that minimizes the overall gate error to $\epsilon + \epsilon' \sim O(1)$, including, besides standard errors, *leakage* into states outside of the logi-

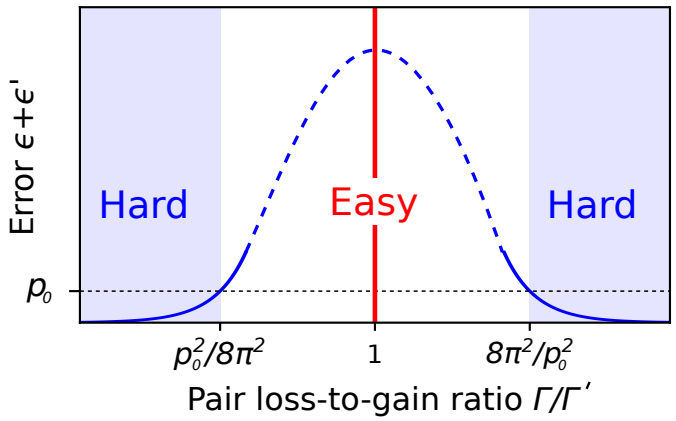

Figure 4.2: **Complexity phase diagram for a fermionic system with simultaneous pair losses and gain.** The plot illustrates the connection between complexity of simulation and the hypothetical dissipative control-Z gate error $\epsilon + \epsilon'$ (both axes use log scale). When the error is smaller than the best-known two-qubit error-correction threshold p_0, the worst-case system dynamics is equivalent to that of a fault-tolerant quantum computer (blue shaded regions) and, according to existing complexity conjectures, is classically computationally hard to simulate. In contrast, when the rates of gain and loss are exactly equal, the problem belongs to EC1 (vertical red line) with the effective classical algorithm provided in the text. The result for the unshaded region remains inconclusive. The dashed line represents qualitative extrapolation.

cal Hilbert space. For fixed Γ, there always exists a choice of $\Gamma' \sim O(1)$ that keeps the error below any provided threshold, $\epsilon + \epsilon' < p_0$, where $p_0 < 0$. According to the leakage threshold theorem in Ref. [257], which is a generalization of earlier standard threshold results [258–260], a universal set of such gates can be used to implement fault-tolerant quantum computing. Therefore, there are instances of Lindblad evolutions that remain hard to simulate for arbitrarily long times.

One particular example of a dissipative correction to ideal dynamics is the presence of pair gain $A'_{ij} = \Gamma' c_i^\dagger c_j^\dagger$ that acts on exactly the same sites as pair loss A_{ij}. In this case, the minimum error is $\epsilon + \epsilon' = \sqrt{8\pi^2 \Gamma'/\Gamma}$ and the problem remains hard for a classical computer if $\Gamma' \leq p_0^2/8\pi^2\Gamma$. Since the entangling gate is also implementable using pair gain instead of loss, this inequality also works after replacing Γ by Γ'. Thus, the problem of simulating

the evolution in the regions $\Gamma'/\Gamma \leq p_0^2/8\pi^2$ and $\Gamma'/\Gamma \geq 8\pi^2/p_0^2$ is classically hard. The complexity for the rest of parameter space remains an open problem. Notably, there exists at least one point in this range, $\Gamma = \Gamma'$, that is easily simulable by a classical computer since it is in EC1. Therefore, by changing the ratio Γ'/Γ, we can potentially induce a complexity phase transition. Figure 4.2 illustrates the connection between gate error and sampling complexity.

Summarizing, we established quantum computational universality of quadratic dissipation combined with free fermion dynamics, where dissipation replaces the unitary interactions between fermions. This result open a possibility of using simple dissipation processes as a resource for quantum computing. In the following section, we illustrate the feasibility of this proposal by considering a system of cold atoms.

4.5 Easy Class 1

In this section, we analyze the convergence of the average unitary stochastic evolution to the exact Lindblad dynamics in the case of Easy Class 1 (EC1). First, we set the initial time to be zero and consider the final time t being an integer multiple of timestep $\Delta\tau$. This assumption holds without loss of generality since $\Delta\tau$ may be adjusted appropriately to capture any particular final time. Then the overall evolution of unitary can be written as a product

$$U(t) = \prod_{n=0}^{t/\Delta\tau} U_n, \tag{4.35}$$

where the timestep unitary U_n is expressed in terms of a time-ordered exponential

$$U_n = \mathcal{T} \exp\Big(-i \int_{n\Delta\tau}^{(n+1)\Delta\tau} dt H'(t)\Big) \tag{4.36}$$

generated by the stochastic Hamiltonian $H'(t)$ in Eq. (4.16),

$$H'(t) = H(t) + \frac{R(t)}{\sqrt{\Delta\tau}}. \tag{4.37}$$

Here $H(t)$ is the original time-dependent Hamiltonian, and $R(t) = \sum_k \xi_{nk} A_k(t) + \xi^*_{nk} A^\dagger_k(t)$ is the normalized stochastic part, where ξ_{nk} are independent complex Gaussian variables defined for times $n\Delta\tau \le t \le (n+1)\Delta\tau$.

Let us consider the ordered exponential expansion of the timestep unitary in Eq. (4.36):

$$\begin{aligned} U_n = I &- i\Delta\tau^{1/2} R_n - \Delta\tau\Big(iH_n + \frac{1}{2}R_n^2\Big) \\ &- \Delta\tau^{3/2}\Big(\mathcal{P}H_n R_n - \frac{i}{6}R_n^3\Big) \\ &- \frac{1}{2}\Delta\tau^2\Big(H_n^2 - \frac{i}{3}\mathcal{P}H_n R_n R_n - \frac{1}{12}R_n^4\Big) \\ &+ O(\Delta\tau^{5/2}), \end{aligned} \tag{4.38}$$

where we denote the discretized value of an operator O_n and permutation sum respectively as

$$\begin{aligned} O_n &= \frac{1}{\Delta\tau}\int_n dt O(t) \equiv \frac{1}{\Delta\tau}\int_{n\Delta\tau}^{(n+1)\Delta\tau} dt O(t), \\ \mathcal{P}O_1 \dots O_m &= \sum_{\sigma\in S_m} O_{\sigma(1)} \dots O_{\sigma(m)}. \end{aligned} \tag{4.39}$$

The average over the stochastic field can be taken for each timestep independently. Therefore, the effect of the timestep unitary in Eq. (4.38) is

$$\begin{aligned}\mathbb{E}U_n\rho U_n^\dagger = \Big(I + \mathcal{L}_n\Delta\tau + \frac{1}{2}\mathcal{L}_n^2\Delta\tau^2\Big)\rho \\ + \mathcal{D}_n\rho\Delta\tau^2 + O(\Delta\tau^3).\end{aligned} \tag{4.40}$$

In the equation above, $\mathcal{L}_n$ is the generator of the original Lindblad equation, $\lim_{\Delta\tau\to 0}\mathcal{L}_n = \mathcal{L}(n\Delta\tau)$, expressed as

$$\begin{aligned}\mathcal{L}_n\rho = -i[H_n,\rho] + \sum_k\Big(A_{kn}\rho A_{kn}^\dagger - \frac{1}{2}\{A_{kn}^\dagger A_{kn},\rho\}\Big) \\ + \sum_k\Big(A_{kn}^\dagger\rho A_{kn} - \frac{1}{2}\{A_{kn}A_{kn}^\dagger,\rho\}\Big),\end{aligned} \tag{4.41}$$

and $\mathcal{D}_n$ represents the lowest-order correction occurring due to the timestep being nonzero:

$$\begin{aligned}\mathcal{D}_n\rho =& \frac{1}{4}\sum_{kk'}\Big(A_{k'n}^\dagger A_{kn}^\dagger\rho[A_{k'n},A_k] + A_{k'n}^\dagger A_{kn}\rho[A_{k'n},A_{kn}^\dagger] \\ &+ A_{k'n}A_{kn}^\dagger\rho[A_{k'n}^\dagger,A_k] + A_{k'n}A_{kn}\rho[A_{k'n}^\dagger,A_{kn}^\dagger]\Big) \\ &+ \sum_k\Big(A_{kn}\rho V_{kn} + V_{kn}\rho A_{kn} + V_{kn}^\dagger\rho A_{kn}^\dagger + A_{kn}^\dagger\rho V_{kn}^\dagger\Big) \\ &+ W_n\rho + \rho W_n^\dagger.\end{aligned} \tag{4.42}$$

Here we have used the notation

$$
\begin{aligned}
V_{kn} &= \sum_{k'} \frac{1}{4}\{A_{kn}^\dagger, \{A_{k'n}^\dagger A_{k'n}\}\} - \frac{1}{6}\mathcal{P}A_{kn}^\dagger A_{k'n}^\dagger A_{k'n} A_{k'n}, \\
W_n &= -\frac{i}{6}\sum_k \Big(\Big[[H_n, A_{kn}], A_{kn}^\dagger \Big] + \{H_n, A_{kn}^\dagger A_{kn}\} \Big) \\
&\quad - \frac{1}{8}\Big(\sum_k \{A_{kn}^\dagger, A_{kn}\}\Big)^2 + \frac{1}{48}\sum_{kk'} \mathcal{P}A_{kn}^\dagger A_{kn} A_{k'n}^\dagger A_{k'n}.
\end{aligned}
\tag{4.43}
$$

The overall expression in Eq. (4.42) can be written in a compact form,

$$
\mathcal{D}_n \rho = \sum_\alpha D_{\alpha n}^{(1)} \rho D_{\alpha n}^{(2)}, \tag{4.44}
$$

where $D_{\alpha n}^{(i)} = \text{poly}(H_n, A_{kn})$ are polynomials of degree less than four.

The averaged stochastic map in Eq. (4.40) can be rewritten as a continuous evolution and then decomposed using Dyson series for the small parameter $\Delta\tau$,

$$
\begin{aligned}
\mathbb{E}\mathcal{V}_{\text{st}}(t_2, t_1) &= \mathcal{T}\exp\Big(\int_{t_1}^{t_2} dt' (\mathcal{L}(t') + \mathcal{D}(t')\Delta\tau)\Big) + O(\Delta\tau^2) \\
&= \mathcal{V}(t_2, t_1) + \int_{t_1}^{t_2} dt' \mathcal{V}(t_2, t')\mathcal{D}(t')\mathcal{V}(t', t_1)\Delta\tau + O(\Delta\tau^2),
\end{aligned}
\tag{4.45}
$$

where the generators $\mathcal{L}(t)$ and $\mathcal{D}(t)$ are continuous versions of the operators in Eq. (4.41) and Eq. (4.42), in which the $\Delta\tau$-averaged operators A_{kn} and H_n are replaced by the corresponding instantaneous values at time t, i.e. $A(t)$ and $H(t)$, respectively. To obtain the expression in Eq. (4.20), we recursively replace $\mathcal{V}(t_2, t')$ and $\mathcal{V}(t', t_1)$ on the right-hand side by their stochastic average and collect all $O(\Delta\tau^2)$ terms.

4.6 Easy Class 2

In this section, we analyze the convergence of the average stochastic unitary evolution to the Lindblad dynamics in the case of Easy Class 2 (EC2). The single timestep evolution averaged over stochastic unitaries in Eq. (4.22) is equivalent to the map

$$\begin{aligned}\mathbb{E}U_n\rho U_n^\dagger &= U_n^0\Big(\rho + \int_n dt \sum_k \Gamma_k(t)\Big(Y_k(t)\rho Y_k^\dagger(t) - \rho\Big)\Big)U_n^{0\dagger} \\ &= \Big(I + \mathcal{L}_n + \frac{1}{2}\mathcal{L}_n^2\Big)\rho + \mathcal{D}_n\rho\Delta\tau^2 + O(\Delta\tau^3),\end{aligned} \tag{4.46}$$

where the target Liouville operator is

$$\mathcal{L}_n\rho = -i[H_n,\rho] + \sum_k A_{kn}\rho A_{kn}^\dagger - \Gamma_{kn}\rho. \tag{4.47}$$

The correction now takes the form

$$\begin{aligned}\mathcal{D}_n\rho = &\sum_k\Big(A_{kn}\rho C_{kn} + C_{kn}^\dagger\rho A_{kn}^\dagger\Big) \\ &- \frac{1}{2}\sum_{kk'} A_{kn}A_{k'n}\rho A_{k'n}^\dagger A_{kn}^\dagger - \frac{1}{2}\Gamma_n^2\rho,\end{aligned} \tag{4.48}$$

denoting $C_{kn} = \Gamma_n A_{kn}^\dagger + \frac{i}{2}[A_{kn}^\dagger, H_n]$ and $\Gamma_n = \Delta\tau^{-1}\int_n dt\sum_k \Gamma_k(t)$. This expression has the form of Eq. (4.44) with operators $D_{\alpha n}^{(i)}$ being a sum of products of at most four Majorana fermion operators.

4.7 Easy Class 3

In this section, we analyze Easy Class 3 (EC3) and show the convergence of the system-ancilla stochastic evolution under the Hamiltonian in Eq. (4.16) using the stochastic operators in Eq. (4.23) to the dissipative dynamics with linear Lindblad jump operators. Let us start from a many-body pure state of the fermions occupying L modes of the system and L_a ancilla modes at time $t = n\Delta\tau$, denoting it as $|\Psi_n\rangle$. At the nth timestep, the evolution acts on the system and the nth ancilla mode only. Thus, the state at time $t = n\Delta\tau$ is a product state of subsystem states: (1) correlated state of L system modes together with the first n ancilla modes and (2) the product states of the remaining $L_a - n$ ancilla modes, i.e.

$$|\Psi_n\rangle = |\phi_n\rangle_{L+n} \otimes |0\rangle_{L_a-n}. \tag{4.49}$$

The evolution is governed by the Hamiltonian

$$H'(t) = H(t) \otimes I_A + \frac{1}{\sqrt{\Delta\tau}}\Big(K(t) + K^\dagger(t)\Big), \tag{4.50}$$

where the stochastic terms are

$$K(t) = \sum_k f_{nk} A_k(t)(\gamma_{2(L+n-1)} + i\gamma_{2(L+n)-1}) \tag{4.51}$$

at times $n\Delta\tau \le t \le (n+1)\Delta\tau$, and f_{nk} are independent real Gaussian variables. Then, at the $(n+1)$th step, the system-ancilla state $|\Psi_n\rangle = U_n|\Psi_n\rangle$ is

$$\begin{aligned}
|\Psi_{n+1}\rangle &= |\Psi_n\rangle - i\Delta\tau^{1/2}K_n|\phi_n\rangle|1\rangle|0\rangle_{T-n-1} \\
&- \Delta\tau\Big(iH_n + \frac{1}{2}K_n^\dagger K_n\Big)|\phi_n\rangle|0\rangle_{L_a-n} \\
&- \Delta\tau^{3/2}\Big(\frac{1}{2}\{H_n, K_n\} - \frac{i}{6}K_nK_n^\dagger K_n\Big)|\phi_n\rangle|1\rangle|0\rangle_{L_a-n-1} \\
&- \frac{1}{2}\Delta\tau^2\Big(H_n^2 - \frac{i}{3}\{H_n, K_n^\dagger K_n\} \\
&\qquad\qquad - \frac{i}{3}R_n^\dagger H_n K_n - \frac{1}{12}(K_n^\dagger K_n)^2\Big)|\phi_n\rangle|0\rangle_{L_a-n} \\
&+ O(\Delta\tau^{5/2}),
\end{aligned} \tag{4.52}$$

where we used the discrete-time operator values H_n and K_n obtained as in Eq. (4.38).

The interpolated continuous-time evolution for the density matrix of the system can be presented in the form

$$\begin{aligned}
\frac{d}{dt}\rho &= \frac{1}{\Delta\tau}\mathbb{E}\ \mathrm{Tr}_A\Big(|\Psi_{n+1}\rangle\langle\Psi_{n+1}| - |\Psi_n\rangle\langle\Psi_n|\Big)\Big|_{n=\lfloor t/\Delta\tau\rfloor} \\
&= \Big(I + \mathcal{L}_n + \frac{1}{2}\mathcal{L}_n^2\Big)\rho + \mathcal{D}_n\rho\Delta\tau^2 + O(\Delta\tau^3),
\end{aligned} \tag{4.53}$$

where $\lfloor x\rfloor$ is the floor function. The target Liouville operator is

$$\mathcal{L}_n\rho = -i[H_n, \rho] + \sum_k A_{kn}\rho A_{kn}^\dagger - \frac{1}{2}\{A_{kn}^\dagger A_{kn}, \rho\} \tag{4.54}$$

and the correction is

$$\begin{aligned}\mathcal{D}_n\rho = &\sum_{kk'}\Big(\frac{1}{4}A_k^\dagger A_{k'}\rho A_{k'}^\dagger A_k - \frac{1}{2}A_k A_{k'}\rho A_{k'}^\dagger A_k^\dagger\Big)\\ &+\sum_k\Big(A_k\rho Q_k + Q_k^\dagger\rho A_k^\dagger\Big)+M\rho+\rho M^\dagger,\end{aligned} \tag{4.55}$$

where

$$\begin{aligned}Q_k &= \frac{1}{12}\{A_k^\dagger, A_{k'}^\dagger A_{k'}\}\\ M &= \frac{i}{6}\sum_k\Big(A_k^\dagger H A_k - \frac{1}{2}\{H, A_k^\dagger A_k\}\Big)\\ &\quad -\frac{1}{12}\sum_{kk'}\Big(A_{k'}^\dagger A_{k'}A_k^\dagger A_k - \frac{1}{2}A_{k'}^\dagger A_k A_k^\dagger A_{k'}\Big).\end{aligned} \tag{4.56}$$

As is the case for EC1 and EC2, the correction is described by Eq. (4.44) with operators $D_{\alpha n}^{(i)}$ being a sum of products of at most eight Majorana fermion operators.

4.8 Error analysis

In this section, we first derive Eq. (4.25) and then provide the proof of Lemma 12 in Section 4.3.1. The error can be formally expressed in terms of evolution superoperators as

$$\epsilon = \frac{1}{2}\max_{\mathbf{r}'}\sum_{\mathbf{r}}\big|\langle\mathbf{r}|\mathbb{E}\mathcal{V}_{\text{st}}(t,0)\rho_{\mathbf{r}'}(0) - \mathcal{V}(t,0)\rho_{\mathbf{r}'}(0)|\mathbf{r}\rangle\big|, \tag{4.57}$$

where $\mathcal{V}(t_2,t_1)$ is the Markovian map generated by Eq. (4.1) in Section 4.1, $\mathcal{V}_{\text{st}}(\xi,t_2,t_1)$ is a unitary trajectory map depending on either a realization of the discrete stochastic field ξ_{kn} (EC1 and EC3) or a random choice of unitaries (EC2). We use the Dyson-like expansion in

Eq. (4.20) and the convexity of the absolute value to upper bound the error as

$$\begin{aligned}\epsilon \leq & \frac{\Delta\tau}{2}\mathbb{E}\max_{\mathbf{r}'}\int_0^t dt' \sum_{\mathbf{r}}\left|\langle\mathbf{r}|\mathcal{V}_{\text{st}}(t,t')\mathcal{D}(t')\mathcal{V}_{\text{st}}(t',0)\rho_{\mathbf{r}'}|\mathbf{r}\rangle\right| \\ & + O(\Delta\tau^2).\end{aligned} \tag{4.58}$$

Using the fact that $\mathcal{V}_{\text{st}}$ is a unitary map, we can rewrite

$$\mathcal{V}_{\text{st}}(t',0) = \mathcal{V}_{\text{st}}^{-1}(t',t)\mathcal{V}_{\text{st}}(t,0), \tag{4.59}$$

where the inverse of a unitary map is well-defined through the inverse unitary transformations. This expression leads directly to Eq. (4.25), taking into account that

$$\mathcal{V}_{\text{st}}^{-1}(t',t)\mathcal{D}(t)\rho\mathcal{V}_{\text{st}}^{-1}(t',t) = \sum_{\alpha} D_{\alpha}^{(1)}(t,t')\rho D_{\alpha}^{(2)}(t,t'), \tag{4.60}$$

where $D_{\alpha}^{(i)}(t,t') = \mathcal{V}_{\text{st}}(t,t')D_{\alpha}^{(i)}(t')$.

4.8.1 Proof of the Lemma

Let us rewrite the left-hand side of Eq. (4.28) using the spectral decomposition $\rho = \sum_\mu p_\mu |\psi_\mu\rangle\langle\psi_\mu|$ and triangle inequality as

$$\begin{aligned}\sum_{\mathbf{r}} |\langle \mathbf{r}|O_1 \rho O_2|\mathbf{r}\rangle| &= \sum_{\mathbf{r}} \Big| \sum_{\mu,\mathbf{r}_1\mathbf{r}_2} p_\mu \psi^\mu_{\mathbf{r}_1} \psi^{\mu *}_{\mathbf{r}_2} \langle \mathbf{r}|O|\mathbf{r}_1\rangle\langle \mathbf{r}_2|O|\mathbf{r}\rangle \Big| \\ &\leq \sum_\mu p_\mu \sum_{\mathbf{r}} \sum_{\mathbf{r}_1\mathbf{r}_2} |\psi^\mu_{\mathbf{r}_1}||\psi^\mu_{\mathbf{r}_2}||\langle \mathbf{r}|O|\mathbf{r}_1\rangle||\langle \mathbf{r}_2|O|\mathbf{r}\rangle| \\ &\leq \|O_1\|_{\max}\|O_2\|_{\max} \sum_{\mu,\mathbf{r}} \sum_{\mathbf{r}_1\in D(k_1,\mathbf{r})} \sum_{\mathbf{r}_2\in D(k_2,\mathbf{r})} p_\mu |\psi^\mu_{\mathbf{r}_1}||\psi^\mu_{\mathbf{r}_2}|,\end{aligned} \tag{4.61}$$

where we denote $\|O\|_{\max} = \max_{ij}|O_{ij}|$ to be the max-norm of the matrix O, and $D_k(\mathbf{r})$ is a sphere with radius k with respect to Hamming distance. Using the inequality

$$|\psi^\mu_{\mathbf{r}_1}||\psi^\mu_{\mathbf{r}_2}| \leq \frac{1}{2}\Big(|\psi^\mu_{\mathbf{r}_1}|^2 + |\psi^\mu_{\mathbf{r}_2}|^2\Big) \tag{4.62}$$

and the property that the sphere $D(k,\mathbf{r})$ contains $\binom{L}{k} \leq L^k/k!$ states, we obtain

$$\sum_{\mathbf{r}} |\langle \mathbf{r}|O_1\rho O_2|\mathbf{r}\rangle| \leq \frac{1}{k_1!k_2!}\|O_1\|_{\max}\|O_2\|_{\max} L^{k_1+k_2}, \tag{4.63}$$

where we use the fact that the density matrix is properly normalized, $\mathrm{Tr}\,\rho = 1$.

4.9 Discussion

In this Chapter, we have demonstrated how simple forms of dissipation affect the complexity of simulation of noninteracting fermions. In particular, focusing on linear-quadratic Lindblad jump operators, we have shown the existence of two complementary complexity classes of Lindblad jump operators, easy and hard for simulation on a classical computer. Using the error-correction formalism, we showed that the hard class has a finite volume in the parameter space and tolerates the presence of small arbitrary corrections. At the same time, the easy classes may have small measure and could become hard even as a result of arbitrarily small corrections to the master equation.

We have expanded the region of classical simulability of free-fermions in the presence of Markovian errors from single-qubit loss/gain to more general quadratic-linear Lindblad jump operators. The algorithms we devise for EC1–EC3 based on the stochastic unraveling approach provably work in polynomial time. This shows that a large class of dissipation processes such as dephasing or single-fermion decay can be treated with the help of efficient classical algorithms.

At the same time, more complex processes are BQP-complete, which we show by explicitly constructing an entangling gate and showing the equivalence of the problem with universal quantum computation. We thus place limitations on the extent to which the simulability result may be extended, since we believe quantum computation is strictly more powerful than classical computation. Our detailed analysis shows that it is within the range of experimental feasibility to implement with cold atoms a quantum computer with purely dissipative atom-atom interactions, an exciting possibility for experiments in quantum com-

puting. For example, dissipative quantum systems such as alkaline-earth atoms may serve in the next generation of quantum supremacy experiments. Also, our result suggests that simulating fermion dynamics may be hard for quantum particles experiencing dissipation, for example, quasiparticles in solid-state systems. Future work can explore the hardness of simulation of electronic systems with quasiparticle dynamics approximated with quadratic-linear Lindblad jump operators that include the effects of electron-electron, electron-phonon, and electron-impurity scattering processes. Alternatively, physical systems following such dynamics with high accuracy may be a future platform for quantum computing experiments.

It may be interesting to explore the connection of our results with the theory of matchgate (free-fermion) computations and the role played by non-Gaussianity. Quadratic fermionic Hamiltonians and single-fermion loss give rise to Gaussian operations and are hence easily simulable [202]. It is known [261] that any non-Gaussian fermionic state is a resource for fermionic computation, boosting the computational power of free fermions from being classically simulable to being universal for quantum computation. Our results suggest that quadratic-linear Lindblad jump operators are non-Gaussian in general. Therefore, it would be interesting to quantify the amount of non-Gaussianity (or "magic") for the Lindblad operations we study here.

Along the same lines, one can quantify a different resource for non-classicality, such as a suitable measure of entanglement for open fermionic systems. Efficient sampling from the full output distribution in arbitrary bases can allow for efficient computation of certain measures of entanglement such as Rényi entropies [262]. Relatedly, Ref. [263] has studied the logarithmic negativity for free fermions with gain/loss Lindblad terms.

Further, one may also consider how the complexity of simulating dynamics under

quadratic Lindbladians changes with time. Since the system starts off in a Fock state that is easy to sample from, and dynamics under quadratic Hamiltonians with quadratic Lindblad jump operators can generate states that are hard to sample from, one can see a dynamical transition in sampling complexity [93; 96]. It is worthwhile to investigate whether these transitions are sharp or coarse (as defined in Ref. [93]) since this can identify what "universality class" free fermions with noise belong to. Techniques such as Lieb-Robinson-like bounds for the evolution of free particles with dissipation [264; 265] would be relevant here.

Another exciting direction is the study of worst-to-average-case equivalence in complexity, which seeks to understand the complexity of typical instances as opposed to worst-case instances [30; 31]. It would be interesting to see if the Cayley path technique of Ref. [31] can be adapted to argue for average-case hardness of dissipative fermionic dynamics.

Chapter 5: Quantum Computational Supremacy via High-Dimensional Gaussian Boson Sampling

5.1 Introduction

We are arriving at an exciting era for quantum computing in which quantum experiments are pushing the limits of what is efficiently computable by the most powerful classical supercomputers. The first major goal for this era is the demonstration of a scalable quantum advantage or *quantum computational supremacy (QCS)* over classical computers. QCS is important as a probe of the foundations of computer science, where it can be seen as an experimental violation of the extended Church-Turing thesis, and it also serves as an important benchmarking tool for comparing near-term experiments on different platforms in a fair and consistent manner. The recent groundbreaking demonstrations of QCS [33; 34] constitute the first significant experimental evidence against the extended Church-Turing thesis.

Notwithstanding, multiple potential loopholes have been pointed out [266–270]. Indeed, QCS will not be marked by a single isolated experiment but rather will be established by gradually improving and scaling up "high complexity" experiments run over the course of many years, which improving classical algorithms will try to simulate. Our confidence that we have arrived in this new era will grow as multiple experiments, performed in different

physical architectures, independently reach this conclusion in a comparable fashion. In this way, the goal may be seen as being analogous to Bell inequality violations, which were originally conducted in landmark experiments starting in the 1970s [271; 272] performed on a variety of different platforms but only much later were loopholes closed in a recent series of impressive experiments [273–276].

Among different approaches to demonstrating QCS [14; 33; 34; 36; 277; 278], photonics provides a promising path as it enables room-temperature operation, fast gate speeds and remarkable potential for scalability [279; 280]. Arguably, the most feasible approach to demonstrating QCS with photonics is to perform the *Gaussian boson sampling (GBS)* protocol [53; 54; 281]. Indeed, this protocol is at the heart of the recent QCS demonstration performed by a team from USTC [34], which employed a GBS device with 100 modes and an average of around 45 photons. However, GBS has several important limitations. On the experimental side, current implementations of GBS either lack programmability [34] or have high loss rates, which could render the system classically simulable [177; 282]. Also, from a theoretical standpoint, there is a comparative lack of complexity-theoretic evidence for the hardness of GBS [268–270].

In this Chapter, we overcome these dual challenges. We close important theoretical loopholes in the hardness argument for GBS and provide evidence for the hardness of classically simulating GBS even in the presence of loss.

We first address the open theoretical questions about GBS, namely, hardness in the regime with little overall noise in the form of optical loss (Section 5.2). More specifically, to provide complexity-theoretic evidence for the hardness of approximately simulating GBS, we prove average-case hardness of computing output probabilities in the noise-free case as

well as the so-called 'hiding property' [14]. These results bring GBS to the level of evidence shared by other QCS proposals such as *random circuit sampling* (RCS) and conventional *boson sampling* (see e.g., Refs. [14; 30; 36]). We then show that average-case hardness of computing output probabilities still holds in a regime of high loss rates, building on recent results [32], and discuss the implications of this result on the noise-regimes in which one may still expect GBS to be hard to simulate on a classical computer. These results bolster the evidence for QCS in the USTC experiment and also any future GBS experiments.

5.2 Hardness of approximate GBS

We begin by reviewing and significantly strengthening the hardness argument for the task of simulating GBS as introduced in Refs. [53; 54]. The structure of this section is as follows. We first introduce the model of Gaussian boson sampling and then examine the evidence for the hardness of approximate boson sampling. In order to bring the hardness evidence for GBS to the same standard as that of boson sampling, we identify two properties required for establishing hardness using the standard QCS arguments, namely *hiding* and *average-case hardness of approximating probabilities*, and we provide strong evidence for these properties in GBS. Specifically, we reduce the hiding property to a highly plausible conjecture in random matrix theory, for which we provide analytical and numerical evidence. Additionally, we provide evidence for approximate average-case hardness by proving approximate worst-case hardness and near-exact average-case hardness of computing the output probabilities. We then extend the latter results to the case of computing output probabilities of noisy GBS, which can be well-motivated when the noise model describing the experimen-

tal data is trusted. These results show that the evidence of a quantum "signal" remains in the output distribution even in the presence of noise. Finally, we discuss the implications of these results on the complexity of simulating GBS in the presence of noise.

5.2.1 Recap: Gaussian boson sampling

GBS is the computational task of sampling the photon number statistics of a Gaussian state. Obtaining a sample from a typical GBS experiment involves the following steps. First, M single-mode squeezed vacuum states are prepared[1]. These states are then interfered on an M-mode linear/optical interferometer containing beam-splitters and phase shifters. Finally, the Gaussian state at the output of the interferometer is impinged on M *photon-number-resolving (PNR)* detectors. The resulting pattern of photon number outcomes from the detectors is the required sample. Because single-mode squeezed states can be generated and interfered deterministically and at room temperatures and with high rates, GBS is experimentally feasible on large scales already today, as evidenced by the recent experiment from USTC [34].

In more detail, a typical GBS experiment involves interfering M single-mode squeezed vacuum states with squeezing parameters $\{r_i\}_{i=1}^{M}$ at an interferometer specified by an $M \times M$ linear-optical unitary matrix U. Note that some of the modes can be optionally prepared in the vacuum state, and these can be specified by setting their squeezing parameter to zero.

The probability of detecting n_1 photons in the first mode, n_2 in the second, and so

[1] In the general setup, a GBS experiment can have a general Gaussian state at the input.

on, denoted by $\mathbf{n} = (n_1, \ldots, n_M)$, is

$$\Pr(\mathbf{n}) = \frac{|\mathrm{Haf}(A_{\mathbf{n},\mathbf{n}})|^2}{\prod_{j=1}^{M} n_j! \cosh r_j}. \tag{5.1}$$

Here, $A = A^T = U \left(\oplus_{i=1}^{M} \tanh(r_i)\right) U^T$ is the so-called adjacency matrix of the (pure, zero-displacement) Gaussian state [53], and $A_{\mathbf{n},\mathbf{n}}$ is the symmetric matrix of size $N = \sum_{i=1}^{M} n_i$ (i.e the total photon number) obtained by repeating the i^{th} column and row of A a total of n_i times. In particular, if $n_i = 0$ then the corresponding row and column is deleted. Finally, the Hafnian Haf$(\cdot)$ of a symmetric $N \times N$ matrix B is given by

$$\mathrm{Haf}(B) = \sum_{\mu \in \mathrm{PMP}(N)} \prod_{(i,j)\in\mu} B_{i,j}, \tag{5.2}$$

where PMP(N) is the set of perfect matching permutations of N elements, i.e., permutations $\mu : [N] \to [N]$ satisfying $\mu(2k-1) < \mu(2k)$, $\mu(2k-1) < \mu(2k+1)$. The Hafnian of a 0×0 matrix is defined to be 1 and the one of an odd-size matrix is defined to be 0, which is a manifestation of the fact that squeezed states are supported on even photon number states only. By allowing arbitrary linear-optical unitaries and arbitrary squeezing parameters of each squeezer, an arbitrary symmetric matrix A can be encoded (up to scaling pre-factors) into a Gaussian state. For generic instances, the best-known algorithms to calculate Hafnians scale like $N^3 2^{N/2}$ where N is the size of the matrix [283].

5.2.2 Recap: Approximate sampling hardness of boson sampling

Before we state our technical results, we review the main steps of the hardness argument for conventional boson sampling as given by Aaronson and Arkhipov [14]. These steps provide context for the hardness results of GBS that we present in the remainder of this section.

In a standard boson sampling experiment, instead of interfering single-mode squeezed states at an interferometer as done in Gaussian boson sampling, an N-photon M-mode Fock state is prepared and evolved under a linear-optical unitary and then measured in the photon-number basis. The boson sampling task is to, given a linear-optical unitary as an input, output samples from the output distribution of a corresponding boson sampling experiment.

Aaronson and Arkhipov showed that it is not possible for a classical computer to efficiently do this task unless certain complexity-theoretic conjectures are false. In particular, they reduce the task of approximating the probabilities of outputs to sampling using an efficient classical algorithm, making use of an approximate counting algorithm due to Stockmeyer [15]. This probability estimation can in turn be related to approximating the permanent of a certain sub-matrix of the linear-optical unitary, which is is provably hard for a class known as #P. While the Stockmeyer reduction is not efficient, the existence of a classical efficient sampling algorithm would imply that #P-hard problems could be solved using fewer computational resources than expected, amounting to an argument by contradiction.

The main difficulty in the hardness argument for boson sampling arises when extending it to the setting of *approximate sampling*. Here, the task is to sample from from any distribution that is within a constant-size total-variation distance from a given ideal bo-

son sampling distribution. This additional constraint takes into account that actual devices are bound to achieve only some finite and typically additive precision. In this setting, one may therefore argue for a separation of computational power between quantum and classical devices.

Given this constraint, the hardness argument for the task of approximate sampling must take into account that the constant error budget on the distribution can be distributed *arbitrarily* across all outcome probabilities. In particular, this means that any specific outcome probability of the actually sampled distribution might have a large (constant-size) error when compared to the ideal distribution, which would imply that the sampler cannot be used to estimate the true outcome probabilities. To get around this issue, the argument is extended to random problem instances: via a property of the distribution over problem instances called *hiding*, one can then translate typical outcomes of fixed instances to fixed outcomes of random instances. This enforces that with high probability, the overall constant error budget for the entire distribution is manifest in small errors on the individual probabilities that are proportional to the inverse size of the sample space, that is, $\propto 1/\binom{M}{N}$. Technically, in standard boson sampling, showing the hiding property boils down to showing that the distribution of any small enough sub-matrix of a Haar-random unitary is approximately (in total-variation distance) an entry-wise complex normal distribution. In particular, Aaronson and Arkhipov show that when $M \in \omega(N^5)$, we can "hide" a random Gaussian matrix in a small enough sub-matrix of the large Haar-random unitary by an appropriate procedure [14, Lemma 5.7] because all of these sub-matrices are indistinguishable random Gaussian matrices.

For the approximate sampling task to remain computationally intractable, it remains

to show that estimating the outcome probabilities up to inverse-exponentially small error is #P-hard for any fraction of the problem instances on which the obtained error bound holds—a property called *approximate average-case hardness*. More precisely, given a random problem instance, estimating the probability of a given outcome must be #P-hard with high probability. As evidence toward this property, it has been shown that *exactly* computing those output probabilities is in fact #P-hard on average (and this was a motivation for boson sampling in the first place), and it is known that estimating them to the required robustness level is worst-case hard. However, the hardness of computing those probabilities to a sufficiently large robustness level on average is still unknown.

5.2.3 Average-case hardness of computing GBS probabilities

As outlined in Section 5.2.2, the question of hardness of approximate sampling boils down to whether it is #P-hard to approximate most output probabilities. In this subsection, we show the average-case hardness of this task when the allowed additive approximation error is exponentially small[2].

The output probabilities of GBS are given in terms of $\left|\mathrm{Haf}((UI_KU^T)_{\mathbf{n},\mathbf{n}})\right|^2$. In Ref. [284], we show that the distribution over the $N \times N$ matrices $(UI_KU^T)_{\mathbf{n},\mathbf{n}}$ for Haar random U is well approximated by complex, symmetric Gaussian matrices XX^T. Hence, to show the average-case hardness of computing output probabilities of GBS, it suffices to consider the following problem:

[2]Note, however, that this is a different exponential from what is desired in order to show a polynomial hierarchy collapse.

(δ, ϵ)-Squared-Hafnians-of-Gaussians

Input A matrix XX^T with $X \sim \mathcal{G}_{N,K}(0, 1/M)$.

Output $\left|\text{Haf}(XX^T)\right|^2$ to additive error ϵ, with probability $\geq \delta$ over the distribution $\mathcal{G}_{N,K}(0, 1/M)$.

To complete the argument that an efficient classical approximate sampling algorithm for GBS cannot exist, it remains to prove the #P-hardness of (δ, ϵ)-Squared-Hafnians-of-Gaussians as formalized by the following approximate average-case hardness conjecture.

Conjecture 13. *The (δ, ϵ)-Squared-Hafnians-of-Gaussians problem is #P-hard for any $\epsilon = O\left(N! \tanh^N(r)/(\cosh^K(r) M^N)\right)$ and any constant $\delta > 3/4$.*

As in all other known proposals for demonstrating QCS, this approximate average-case hardness conjecture remains open. Nonetheless, just like in other proposals, it turns out that one can give evidence for Conjecture 13. Namely, we can prove a weaker version of the conjecture with a smaller robustness level $\epsilon = O\left(\exp[-6N \log N - \Omega(N)]\right)$ as opposed to $\epsilon = O\left(\exp[-N \log N - \Omega(N)]\right)$ in Conjecture 13.

Theorem 14. *The (δ, ϵ)-Squared-Hafnians-of-Gaussians problem is #P-hard under PH reductions for any $\epsilon \leq O\left(\exp[-6N \log N - \Omega(N)]\right)$ and any constant $\delta > 3/4$.*

We provide a detailed proof of Theorem 14 in Section 5.3. The technique we employ in the proof is a worst-to-average-case reduction [14; 285]. That is, by assuming access to an oracle for the (δ, ϵ)-Squared-Hafnians-of-Gaussians problem, we show that one in fact approximate $\text{Haf}(XX^T)$ for any matrix $X \in \mathbb{C}^{N \times K}$. This latter task is #P-hard in the

worst-case as we show in Section 5.3.1. At a high level, the worst-to-average-case reduction relies on the fact that $\left|\mathrm{Haf}(XX^T)\right|^2$ is a low degree (of degree $2N$) polynomial over the entries of the matrix X. This allows us to use the oracle to perform polynomial interpolation. Therefore, by combining this observation with the techniques of Refs. [14; 32; 285; 286], we obtain a worst-to-average case reduction for exactly computing the output probabilities.

Together, our results on the hiding property and the approximate average-case conjecture in GBS, significantly strengthen the evidence for the hardness of approximately simulating GBS in terms of the total-variation distance to the ideal output distribution. Given our results, GBS is now on par with the other leading QCS proposals in terms of complexity-theoretic evidence for approximate sampling hardness [14; 30–32; 36; 277; 286], up to a plausible conjecture in random matrix theory—for which we provided theoretical and numerical evidence in Ref. [284]. To achieve a demonstration covered by those complexity-theoretic results, however, the loss rate at every element of the linear-optical circuit, must scale inversely with the total number of such elements—a daunting challenge from an experimental perspective.

5.2.4 Hardness of computation of output probabilities for noisy GBS

In the remainder of this section, we will go one step further and assess how the complexity-theoretic argument for sampling hardness is affected by more realistic noise levels, in particular, in terms of photon loss. In terms of scaling, any constant loss rate of the individual optical elements will lead to an exponential suppression of quantum effects in the output distribution. We now show that, nonetheless and surprisingly, an evidence of a quan-

tum signal remains even in the presence of such significant loss. We then discuss to what extent and in which regimes such a quantum signal might lead to the hardness of simulating a lossy GBS experiment.

Our main result in this section is the average-case hardness of computing the noisy output probability of a random GBS instance, which we obtain by using similar arguments to recent work of Bouland et al. [32], but now extended to the GBS setting. Our results are valid for any noise model that is local, stochastic, and is error-detectable using linear optics. More specifically, we consider a setting where the noise acts locally after every gate, and is of the form

$$\mathcal{N}_i[\rho] = (1 - \eta_i)\rho + \eta_i \mathcal{E}_i[\rho], \tag{5.3}$$

where stochasticity requires $\mathcal{E}_i$ to itself be a valid channel (i.e. a completely positive trace preserving map) with no identity component.

Consider the following problem.

(ϵ, η)-NoisyGBS-Probability

Input A noisy GBS instance, consisting of the linear-optical unitary U on M modes chosen from the Haar measure $\mathcal{H}$, the squeezing parameters at the input, and a description of the noise channels with parameters η_i. Let $\eta = \max_i \eta_i$.

Output With probability δ over instances, an estimate of the quantity $\Pr(\mathbf{n})$ to additive error ϵ, where $\Pr(\mathbf{n})$ is the probability of obtaining an outcome $\mathbf{n}$ that is collision-free and has a total number of photons N scaling as $N = \mathsf{poly}(M)$.

With probability $1 - \delta$, an arbitrary output.

In the above definition, we take $\delta = 1$ to mean the worst-case problem. We prove the following statement of average-case hardness of computing noisy probabilities.

Theorem 15. *There exists a noise threshold η_* and a sufficiently large polynomial such that the problem (ϵ, η)-NoisyGBS-Probability is #P-hard under PH reductions for any constant $\delta > 3/4$, $\eta \leq \eta_*$, and $\epsilon \leq 2^{-\mathsf{poly}}$.*

There are two parts to the proof. The first part is a proof of worst-case hardness of the problem (when $\delta = 1$), and the second a worst-to-average-case equivalence. For worst-case hardness, it turns out that it suffices for the noise to be stochastic and to be able to error-detect it [287]. These conditions are both met for optical loss. Optical loss is stochastic by virtue of being a convex combination of the channels corresponding to no photon loss, single-photon loss, and so on [288]. Moreover, optical loss can also be detected and corrected using only linear-optical operations and photo-detection with high thresholds [280].

These two properties, namely stochasticity and error-detectability, allows us to apply

a result of Fujii [287] to prove worst-case hardness of estimating the noisy probabilities. For the worst-to-average-case equivalence, all we need is for the polynomial structure in the problem to be preserved. This can be satisfied for any local noise model. Preserving the polynomial structure of the output probability enables us to continue to use the same proof techniques as in the previous Section 5.2.3, as we elaborate in Section 5.4.

We close out this section by reminding the reader again that, in this section, we considered the hardness of computing output probabilities. While these are not tasks that are feasible for any realistic quantum device, our results nevertheless indicate that there is a computationally intractable (but exponentially small) "quantum signal" present in the system.

5.2.5 The complexity of noisy and approximate GBS

We now discuss the implications of the hardness result for computing noisy GBS probabilities on the complexity of sampling from the output distribution of noisy GBS. An immediate implication of this result is that it is classically hard to *exactly sample* from the noisy distribution of a worst-case GBS experiment. This is because the quantum signal is still present in the distribution, so the argument based on Stockmeyer's algorithm (and sketched in Section 5.2.2) is valid. Thus, in the idealized situation in which loss is the only source of noise of an experimental system and the exact loss rate is known, simulating a worst-case GBS experiment is classically intractable. Note that loss rates can be inferred from standard optical tomography procedures such as that of Ref. [289]. Given that this result links the hardness of simulating the noisy experiment to an exponentially small quantum signal in the form of output probabilities, it is crucial that the noise model accurately captures the working

of the device.

We now discuss the more realistic situation in which loss is the predominant, but not the sole, source of noise in a photonic experimental system. What can we say about the hardness of approximate sampling in such a situation? To begin with, let us draw on some intuition from RCS schemes acting on n qubits. Here, the additive error incurred in estimating output probabilities using the Stockmeyer algorithm is $O(2^{-n})$ with high probability (since this is the size of a typical output probability in an RCS experiment). In the presence of uncorrected noise, an error of $O(2^{-n})$ in the noisy output probability can be too large for hardness. For example, there is evidence that with gate-wise depolarizing noise, the probabilities will deviate from uniform by merely $O(2^{-m})$, where m is the total number of gates[3] [290]. This means that approximate-sampling hardness cannot be shown using these techniques, since it is not hard to approximate the noisy probabilities any more. Indeed, in this regime, the noisy distribution is exponentially close in total-variation distance to the uniform distribution, rendering the approximate sampling task for the noisy distribution classically simulable.

In the case of noisy GBS, the dominant noise model, namely loss, leads to the vacuum state for a sufficiently deep network, which is again a distribution that is easy to classically sample from (similar to the uniform distribution in qubit RCS schemes). However, if we post-select on a certain minimal number of photons surviving, the distribution need not be easy to simulate. This post-selection is efficient when the depth of the circuit scales poly-logarithmically in the number of modes. In this case, the quantum signal will be large enough so that even with an inverse exponential error, deviations from the easy distribution

[3]Here we are considering the size m to satisfy $m = \omega(n)$.

can be detected.

This excludes the simulation algorithm that samples from the trivial distribution[4] and is a necessary condition for approximate average-case hardness to hold. In summary, our results indicate that there might be 'room in the middle' in terms of gate depth and noise rates, where hardness of sampling might hold. In fact, this intuition lies at the heart of the high-dimensional architecture presented in Ref. [284]. This architecture is designed in such a way that only as few gate applications as necessary for hardness are executed, so that the leeway for noise to ruin the hardness of sampling is minimized. We stress, however, that at the moment, existing proof techniques do not suffice to make a claim of this nature. In fact, in certain regimes of noisy GBS, approximate sampling is known to be classically efficient [282].

5.2.6 Hardness for computing noisy probabilities in high-dimensional GBS

In this subsection, we now argue for the hardness of computing output probabilities for the noisy, high-dimensional GBS setup. In particular, we show that hardness is present even in shallow depth noisy high-dimensional GBS architectures. This is in contrast to the results in Section 5.2.4, where no restriction is made on the depth.

To do this, we simply observe that the previous argument for worst-case hardness, which depends on the noise being local, stochastic, and error-detectable, continues to hold for the limited-depth setup. For average-case hardness of computing noisy probabilities, we again use a worst-to-average-case reduction. However, the polynomial interpolation in this

[4] This distribution is uniform on every photon number sector and every such sector is sampled according to the ideal photon number distribution.

case is different, since a random instance is not Haar distributed any more but rather according to $\mathcal{U}$, the distribution over random instances of high-dimensional GBS. To explain further, consider the usual interpolation $X(t) = (1-t)X + tY$, where $X(0) = X$ is drawn from $\mathcal{U}$ and $X(1) = Y$ is the matrix corresponding to a worst-case high-dimensional GBS instance. In this case, there is no guarantee that the interpolated matrices $X(t)$ also correspond to high-dimensional GBS instances of small depth. We get around this issue by choosing a gate-wise interpolation that is similar to that seen in RCS [30; 31].

We first define the problem of computing output probabilities of a restricted-depth high-dimensional GBS architecture.

(ϵ, η)-HighDimensional-NoisyGBS-Probability

Input A noisy GBS instance drawn from $\mathcal{U}$ that can be implemented in D dimensions with a constant number of cycles $C = O(1)$ with noise parameter η.

Output With probability δ over instances, an estimate of $\Pr(\mathbf{n})$, a collision-free outcome with $N = \mathsf{poly}(M)$ photons, to additive error ϵ.

With probability $1 - \delta$, an arbitrary output.

Similar to the previous sections, we can again obtain an average-case hardness result, that we state here and prove in Section 5.5:

Theorem 16. *There exists a noise threshold η_* and a sufficiently large polynomial such that the problem (ϵ, η)-HighDimensional-NoisyGBS-Probability is #P-hard under PH reductions for any constant $\delta > 3/4$, $\eta \leq \eta_*$, and $\epsilon \leq 2^{-\mathsf{poly}}$.*

5.3 Average-case hardness of computing GBS output probabilities

In this section, we show average-case hardness of computing GBS output probabilities. As explained in Section 5.2.3, this amounts to showing that the following problem is #P-hard.

(δ, ϵ)-Squared-Hafnians-of-Gaussians

Input A matrix XX^T with $X \sim \mathcal{G}_{N,K}(0, 1/M)$.

Output $\left|\text{Haf}(XX^T)\right|^2$ to additive error ϵ, with probability $\geq \delta$ over the distribution $\mathcal{G}_{N,K}(0, 1/M)$.

The proof will proceed in two steps: First, we will show that an oracle for the (δ, ϵ)-Squared-Hafnians-of-Gaussians problem allows one to approximate $\left|\text{Haf}(YY^T)\right|^2$ for arbitrary $Y \in \mathbb{C}^{2N \times 2K}$. This first part of the proof constitutes the worst-to-average-case reduction. Second, we will show that approximating $\left|\text{Haf}(YY^T)\right|^2$ for arbitrary $Y \in \mathbb{C}^{2N \times 2K}$ is actually #P-hard in the worst-case. We show this by reducing the task of approximating the permanent of an arbitrary complex $N \times N$ matrix to the task of approximating $\left|\text{Haf}(YY^T)\right|^2$.

5.3.1 Worst-case hardness

Consider the following problem:

ϵ-Squared-Hafnians

Input A matrix YY^T with $Y \in \mathbb{C}^{N\times K}$ for $K \in \mathbb{N}$, $N \in 2\mathbb{N}$ such that the entries of Y are of the form $(x+iy)/\sqrt{M}$ for $|x|, |y|$ some $O(1)$-bounded integers and additive-error tolerance $\epsilon > 0$.

Output An estimate h s.t. $\left|h - \left|\text{Haf}(YY^T)\right|^2\right| \leq \epsilon$.

We prove the following Lemma.

Lemma 17. *The problem ϵ-Squared-Hafnians is worst-case #P-hard for any additive error $\epsilon \leq 1/(2M^N)$.*

Proof. Without loss of generality, we restrict to $N \leq K$. We begin the proof by noting that the permanent of any square matrix G can be expressed as the Hafnian of a corresponding block matrix twice the size of G [54],

$$\text{Per}(G) = \text{Haf}\left[\begin{pmatrix} 0 & G \\ G^T & 0 \end{pmatrix}\right]. \tag{5.4}$$

Hence, computing the squared permanent of any complex $N/2\times N/2$ matrix $G \in \mathbb{C}^{N/2\times N/2}$ reduces to computing the squared Hafnian of a corresponding block matrix

$$B(G) = \begin{pmatrix} 0 & G \\ G^T & 0 \end{pmatrix}. \tag{5.5}$$

Computing the squared permanent exactly is known to be worst-case #P-hard even over $0/1$-matrices [14; 291].

Next we note that any matrix $B(G)$ for $G \in \mathbb{C}^{N/2\times N/2}$ can be decomposed as XX^T in terms of some complex matrix $X \in \mathbb{C}^{N\times K}$. Indeed the block matrix $B(G)$ is a complex, symmetric matrix, so we can decompose it using the Takagi decomposition as WDW^T, where $W \in U(N)$ is a unitary matrix and $D \in \mathbb{R}^{N\times N}$ is a nonnegative diagonal matrix. We now define $X' = (WD^{1/2})$ and X by appending $(K - N)$ all-0-columns to X'. This gives rise to a decomposition of $B(G) = XX^T$ with $X \in \mathbb{C}^{N\times K}$. Hence it is #P-hard to exactly compute the Hafnian of matrices of the form XX^T in the worst case. Additionally, since the Hafnian is a continuous function, we can compute Haf(XX^T) to an arbitrary level of precision by considering Haf(YY^T) with the entries of Y being of the form $x + iy$, with x and y integers (by suitably rescaling the entries of the matrix). Finally, we note that by normalization we can assume that the entries of the matrix Y are of the form $(x+iy)/\sqrt{M}$ with x and y $O(1)$ bounded integers. Then the squared Hafnian of YY^T is an integer multiple of $1/M^N$. Therefore, computing the Hafnian of YY^T up to additive error of $1/(2M^N)$ serves to compute the squared Hafnian of $B(G)$ exactly, which is #P-hard. This concludes the proof.

The proof holds equally for $N \in \mathsf{poly}(K)$: in this case we embed a square matrix in $\mathbb{C}^{K\times K}$ and append 0 rows instead of columns. □

5.3.2 Worst-to-average equivalence

In this section, we prove the average-case hardness of computing GBS output probabilities. That is, we prove the following Lemma:

Theorem 18 (Theorem 14 restated). *The (δ, ϵ)-Squared-Hafnians-of-Gaussians problem*

is #P*-hard under* PH *reductions for any* $\epsilon \leq O\left(\exp[-6N \log N - \Omega(N)]\right)$ *and any constant* $\delta > 3/4$.

We first sketch the proof idea and elaborate on the technique used. The overall idea is to give a worst-to-average-case reduction from the problem ϵ-Squared-Hafnians to the problem (δ, ϵ)-Squared-Hafnians-of-Gaussians. The worst-case #P-hardness of problem ϵ-Squared-Hafnians has been established in the previous section.

We use the same technique as Refs. [14; 32] to establish this reduction. Assume that we are given an oracle O that solves (δ, ϵ)-Squared-Hafnians-of-Gaussians, meaning that with probability at least δ over the input X, it outputs a squared Hafnian of XX^T to additive error ϵ. The rest of the time, it may output an incorrect value, with no guarantees whatsoever on how close the output is to the desired output. In the following, we will show how to use the oracle O to obtain the squared Hafnian of an arbitrary worst-case matrix YY^T with high probability (this latter probability is over the choice of the random variables instantiated in the algorithm).

The key idea is that for $X \in \mathbb{C}^{N \times K}$, the quantity $\left|\mathrm{Haf}\left(XX^T\right)\right|^2$ is a degree $2N$ polynomial over the entries of the matrix X. This allows for the use of polynomial interpolation to recover the squared Hafnian of an arbitrary worst-case matrix YY^T. An important technique we use in this proof is the robust Berlekamp-Welch algorithm due to Ref. [32], which is important for polynomial interpolation over $\mathbb{R}$ as opposed to a finite field. Polynomial interpolation over the reals is a technique often used for the problem of average-case hardness of computing output probabilities of random quantum circuits [30; 31]. The Berlekamp-Welch algorithm cannot be used as is for the reals, and therefore, recent works [30; 31]

use techniques like Lagrange interpolation. The new robust Berlekamp-Welch algorithm of Ref. [32] allows for improved robustness of the worst-to-average-case reduction.

As an example, in the context of random quantum circuits over n qubits and m gates, Lagrange interpolation can only give average-case #P-hardness of computing output probabilities to error $2^{-O(m^3)}$ rather than the $O(2^{-n})$ that suffices for proving the hardness of approximate sampling (see [14; 31]). The modified Berlekamp-Welch algorithm of Ref. [32], which is boosted with an NP oracle, can sidestep the need for Lagrange interpolation and obtain average-case #P-hardness with $2^{-O(m \log m)}$ error (see also, the recent work of Kondo et al. [286] which also obtains this robustness error).

Theorem 19 (Robust Berlekamp-Welch algorithm [32])**.** *Let p be a univariate polynomial of degree d over the reals. Suppose that we have $k \geq 100d^2$ points (x_i, y_i), with $\{x_i\}$ uniformly spaced in the interval $[0, \kappa]$ and obeying the promise*

$$\Pr[|y_i - p(x_i)| \geq \Delta] \leq \eta < \frac{1}{4}. \tag{5.6}$$

Then there is a P^{NP} algorithm that can estimate $p(1)$ to additive error $\Delta \exp[d \log \kappa^{-1} + O(d)]$ with probability at least 2/3.

Proof of Theorem 14. The polynomial interpolation procedure is as follows. Let $X(t)$ be the matrix obtained by drawing a random $X \sim \mathcal{G}_{N,K}(0, 1/M)$ and setting

$$X(t) := (1-t)X + tY, \tag{5.7}$$

where Y is the matrix corresponding to the worst-case instance. Now, the quantity

$$p(t) := \left|\text{Haf}(X(t)X^T(t))\right|^2 \tag{5.8}$$

is a polynomial of degree $2N$ over the entries of $X(t)$, and consequently, over t itself. For t close to 0, $X(t)$ is close to Gaussian distributed, while when t is close to 1, the distribution is close to being deterministic. We select k points in the range $[0, \kappa]$ and query the oracle O for the value of $p(t)$ for these points. By the promise, the oracle outputs the correct value of $p(t)$ for most values of t with high probability. Conditioned on this event, the robust Berlekamp-Welch algorithm stated in Theorem 19 allows one to reconstruct the polynomial in the second level of the polynomial hierarchy. The polynomial can then be evaluated at the point $t = 1$ to obtain an estimate of the squared Hafnian of the worst-case matrix YY^T.

We now check that the conditions of Theorem 19 are met. We say that a call to the oracle O is successful if it outputs the squared Hafnian of a matrix to additive error ϵ. By assumption, for X drawn at random from $\mathcal{G}_{N,K}(0, 1/M)$, the oracle is successful with probability at least δ. Note however that the matrix $X(t) = (1-t)X + tY$ is not exactly distributed according to $\mathcal{G}_{N,K}(0, 1/M)$. Instead, for small t, due to the rescaling by $(1-t)$ and the shift by tY, $X(t)$ is distributed according to a slightly different distribution $\mathcal{G}'$. If we query the oracle for the value of $p(t)$ with matrices drawn from this different distribution $\mathcal{G}'$, the probability of success can, in the worst case, decrease. By definition, the success probability can decrease at most by the variation distance between the two distributions $\mathcal{G}_{N,K}(0, 1/M)$ and $\mathcal{G}'$, which is $O(t \max(N, K)^2)$. Therefore, for $K \geq N$, the probability of success is at least $\delta - O(\kappa K^2)$. We choose κ to be $O(c/K^2)$ with some small enough c

so that the success probability is at least $\delta - O(c) > 3/4$. This ensures that the conditions of the theorem are met.

We finally conclude by examining the additive error to which we can compute, using the $\mathsf{BPP}^{\mathsf{NP}}$ reduction, the squared Hafnian of the worst-case matrix YY^T. If the additive error for successful queries to the oracle is at most ϵ, Theorem 19 implies that the error in computing $p(1)$ is $\epsilon \exp[d \log \kappa^{-1} + O(d)]$. Plugging in $d = 2N$ and $\kappa = c/N^2$, we get the total additive error in estimating $p(1)$ to be $\epsilon \exp[4N \log N + O(N)]$. Finally, we note that the squared Hafnian is shown to be worst-case hard for additive error $O(1/M^N)$. Therefore, we make the choice

$$\epsilon \exp[4N \log N + O(N)] = O\left(\frac{1}{M^N}\right), \tag{5.9}$$

or

$$\begin{aligned} \epsilon &= O\left(\exp[-4N \log N - \Omega(N) - 2N \log N]\right) \\ &= O\left(\exp[-6N \log N - \Omega(N)]\right), \end{aligned} \tag{5.10}$$

where we have assumed $M = \Theta(N^2)$. This choice ensures that we can, with probability at least 2/3, compute the squared Hafnian of an arbitrary matrix with bounded entries of the form YY^T to additive error $O(1/M^N)$. As shown in Lemma 17, this task is #P-hard. This completes our proof. □

5.4 Average-case hardness of computing noisy GBS output probabilities

In this section, we argue that computing the output probabilities for a *noisy* random GBS experiment is #P-hard on average. That is, we show the following lemma.

Lemma 20. *There exists a polynomial $p(N)$ and a loss threshold η_* such that (ϵ, η)-NoisyGBS-Probability with $\eta \leq \eta_*$, $\delta > 3/4$, and $\epsilon \leq 2^{-p(N)}$ is #P-hard under PH reductions.*

Proof. For worst-case hardness despite the presence of noise, we follow the proof technique in Refs. [32; 287]. At a high level, the worst-case hardness follows from the error-detection property of the system. In particular, the error-detection property implies that as long as the noise η is smaller than a certain threshold η_*, there is a fixed outcome on a subset of the modes, say $\mathbf{m}$, such that conditioned on this outcome, the probability distribution on the rest of the modes is exponentially close to the target noiseless distribution. In other words, we have

$$\left| \Pr_{\text{noisy}}[\mathbf{n}|\mathbf{m}] - \Pr_{\text{ideal}}[\mathbf{n}] \right| \leq 2^{-\text{poly}(N)} \tag{5.11}$$

for any desired polynomial on the right hand side. Since $\Pr_{\text{ideal}}[\mathbf{n}]$ is #P-hard to approximate in the worst case by virtue of Lemma 17, so is computing the conditional probability

$$\Pr_{\text{noisy}}[\mathbf{n}|\mathbf{m}] = \frac{\Pr_{\text{noisy}}[\mathbf{n}, \mathbf{m}]}{\Pr_{\text{noisy}}[\mathbf{m}]}. \tag{5.12}$$

The denominator here is the probability of seeing the outcome $\mathbf{m}$, which flags the no-error

event. The probability of this can be exponentially small, and satisfies [32; 287]

$$\left| \Pr_{\text{noisy}}[\mathbf{m}] - (1-\eta)^{O(Md)} \right| \leq \Pr_{\text{noisy}}[\mathbf{m}] 2^{-\mathsf{poly}(N)}, \tag{5.13}$$

where η is the maximum noise parameter as in Eq. (5.3). In other words, for an error-detected circuit, the probability that the outcome on the subset of heralding modes is in the state $\mathbf{m}$ is exponentially close to the probability that no error occurred, which is given by $(1-\eta)^{O(Md)}$.

Therefore, approximating $\Pr_{\text{noisy}}[\mathbf{n}, \mathbf{m}]$ is also #P-hard:

$$\left| \Pr_{\text{noisy}}[\mathbf{n}, \mathbf{m}] - \Pr_{\text{noisy}}[\mathbf{n}|\mathbf{m}](1-\eta)^{O(Md)} \right| \leq \Pr_{\text{noisy}}[\mathbf{n}, \mathbf{m}] 2^{-\mathsf{poly}(N)} \tag{5.14}$$

$$\Rightarrow \left| \Pr_{\text{noisy}}[\mathbf{n}, \mathbf{m}] - \Pr_{\text{ideal}}[\mathbf{n}](1-\eta)^{O(Md)} \right| \leq \Pr_{\text{noisy}}[\mathbf{n}, \mathbf{m}] 2^{-\mathsf{poly}(N)} + 2^{-\mathsf{poly}(N)}. \tag{5.15}$$

Since computing $\Pr_{\text{ideal}}[\mathbf{n}]$ to additive error $\pm O(2^{-\mathsf{poly}(N)})$ is #P-hard, so is computing $\Pr_{\text{noisy}}[\mathbf{n}, \mathbf{m}]$ to additive error $O(2^{-\mathsf{poly}(N)})(1-\eta)^{O(Md)}$. A similar analysis in Ref. [32] shows that it is $\mathsf{coC}_{=}\mathsf{P}$-hard to compute a noisy probability in the worst case to additive error $2^{-O(m \log m)}$ in the context of RCS. This proves the worst-case hardness.

For the worst-to-average-case reduction, we again use the technique of polynomial interpolation in conjunction with a robust Berlekamp-Welch algorithm. We observe that any noisy output probability for a local noise model can still be written as a polynomial in the gate entries of the circuit, using the Feynman sum-over-paths idea. As before, we perform interpolation from a random instance from the ensemble to the worst-case-hard

instance. This is achieved now using the Cayley path interpolation technique of Ref. [31] instead of the direct interpolation between two matrices[5]. The full interpolation involves interpolating every gate of a circuit implementation from the average-case instance A_i to the worst-case instance W_i along the Cayley path

$$C_i(t) = \left(t\mathbb{1} + (2-t)A_iW_i^{-1}\right)\left((2-t)\mathbb{1} + tA_iW_i^{-1}\right)^{-1} \cdot W_i, \tag{5.16}$$

which satisfies $C_i(0) = A_i$ and $C_i(1) = W_i$. Using this interpolation and the fact that any local noise can be "purified" gate-wise by introducing ancillary systems of finite dimension, we can again write the noisy probability $\Pr_{\text{noisy}}[\mathbf{n}, \mathbf{m}][t]$ as a polynomial in t. The rest of the proof follows from before. □

5.5 Average-case hardness of computing noisy probabilities in high-dimensional GBS

For the worst-case hardness of computing noisy probabilities of the high-dimensional GBS architecture, we mainly use the previous results on error-detection of noise. The additional ingredient used is the fact that a constant-depth linear-optical architecture in two dimensions (and higher) has been shown by Brod [94] to be hard to exactly sample from.

The proof of Ref. [94] uses post-selection to argue for exact sampling hardness. Note that the post-selection result does not, by itself, imply the #P-hardness of computing output probabilities[6]. However, we note that the post-selection proof can often be "opened up" in

[5] This is because the noisy output probability is no longer a simple function of only the linear-optical unitary (like the Hafnian), but is also a function of the circuit implementation.

[6] It implies the PP-hardness of strong simulation, which involves computing both the output probabilities

order to directly argue about the hardness of computing output probabilities. This is done by giving an amplitude-preserving reduction from a BQP circuit to the circuit family in question (here, high-dimensional GBS). Since computing output amplitudes of BQP circuits is #P-hard, so is computing that of the circuit family in question. Using the results in the previous section, so is computing the *noisy* output probability in the worst case for an error-detected circuit as long as the noise level is smaller than some (constant) threshold η_*.

The average case hardness again essentially follows by observing that there is a polynomial structure in the output probability, to prove Theorem 16. We again use the Cayley technique of Ref. [31] to set up the polynomial interpolation in this case, and use results from Ref. [32] to strengthen it, such as using a variable rescaling and applying a robust version of the Berlekamp-Welch algorithm (Theorem 19).

5.6 Open problems and discussions

5.6.1 Open problems

In this Chapter, we have provided strong evidence for the hardness of Gaussian boson sampling, putting it on eye level with other known schemes to show QCS such as random circuit sampling and have substantially advanced the theory of GBS. Still, some theoretical questions are outstanding.

1. In Ref. [284], we have been able to show that two plausible conjectures in random matrix theory allow us to obtain the hiding property for a noiseless GBS set up, without restrictions on the number of active modes. Can we obtain a similar hiding property

and the marginals.

for the high-dimensional GBS set-up introduced in Ref. [284]? Is this also possible in the presence of noise? Answering these questions is crucial for extending the hardness of computing output probabilities to the hardness of approximate sampling from experimentally realizable distributions.

2. Informally, the anti-concentration conjecture for boson sampling (or GBS) states that the output probability of a random instance is unlikely to be very small. If this conjecture was true, then now-standard arguments can show that the output probability corresponding to an approximate sampler is, with high probability, a good multiplicative estimate to the ideal output probability. Proving this conjecture true, in either the case of boson sampling or GBS, would give increased evidence to support the goal of proving QCS via photonics. A proof of such a conjecture is challenged due to the fact that tools of unitary designs [155] are presumably unavailable in the bosonic setting [14; 29].

3. Notwithstanding, it would be insightful to compute the second moments $\mathbb{E}_{X\sim\mathcal{G}(0,1/M)}|\operatorname{Haf}(XX^T)|^4$ for the distribution we have found to characterize GBS problem instances. These moments thus characterize the so-called *collision probability* of seeing the same outcome twice in an experiment, which in turn can be related not only to anticoncentration but also the verifiability of approximate GBS from samples [10; 292], thus shedding some light on the structure of the GBS output distribution.

4. An important task in demonstrating QCS is to verify that the performed experiment indeed contains a non-trivial quantum signal that cannot be efficiently spoofed. The Google QCS demonstration relied on linear cross-entropy benchmarking fidelity, and

the USTC experiment used a heavy-output generation (HOG) ratio test as an alternative path to verifiable hardness. Whether the HOG-ratio test can be spoofed efficiently by a classical adversary such as the algorithms considered in Refs. [267; 268] is an open problem.

5. The recent result in Ref. [293] presents a classical algorithm for the simulation of high-dimensional boson sampling experiments in certain regimes. As described, this algorithm is not applicable to the architecture as described in Ref. [284]. Extending the algorithm to be relevant to the present architecture is an open problem.

6. With current optical technology, loss is the dominant source of noise in any GBS experiment. Consequently we were motivated to obtain hardness results for computing the output probabilities of a GBS experiment in the presence of significant photon loss. It is natural to investigate if similar hardness results can be obtained in the presence of other possible sources of experimental noise such as such as mode mismatch, multiple Schmidt modes, interferometer phase drift and detector dark counts.

7. It is a challenge to the community, after all, to relate boson sampling closer to practically important computational tasks, following up on the applications reviewed in Ref. [281] and to identify new applications.

5.6.2 Conclusion and outlook

In summary, this Chapter brings the convincing demonstration of QCS on a programmable photonics device substantially closer to reality. It overcomes previously outstanding major theoretical challenges in the field by providing stronger evidence for the

hardness of GBS. Crucially, we have presented a novel architecture for high-dimensional GBS using optical delay lines that promises low levels of noise without compromising on its programmability. We benchmarked this architecture against the best available classical simulation algorithms and found that already experiments involving a moderate number of modes are far beyond reach for those algorithms.

We close by briefly commenting on the experimental prospects of realizing high-dimensional GBS. Since high-dimensional GBS can be implemented in the time domain according to the scheme presented in Ref. [284], only a single squeezer and a single detector are required. If multiple detectors are available, these can be de-multiplexed using optical switches in order to increase the effective repetition rate of the experiment and reduce the length of the delay lines. Especially promising is the case of $D = 3$, $a = 6$, $C = 1$, which can be implemented with only three optical delay lines and three each of re-programmable beam-splitters and phase shifters. Assuming reasonable values of squeezer out-coupling losses, free-space to fiber coupling loss and detector efficiency, we estimate that such a setup can be built using current optical technology with around 40% transmission, higher than that enabled by the ultra-low non-programmable loss interferometer in the USTC experiment. Such a setup would enable the largest demonstration of QCS yet with a mean detected photon number of 80 in a programmable device with 216 total modes.

Chapter 6: The importance of the spectral gap in estimating ground-state energies

6.1 Introduction

Several exotic phenomena in our world occur only at very low temperatures, most notably, those occurring in condensed matter such as superconductivity, superfluidity and the fractional and integer quantum Hall effects. Beyond these examples in condensed-matter physics, the low-energy physics of systems of several interacting particles is of interest in several fields such as particle physics, atomic, molecular, and optical physics, chemistry, and quantum computing. Accordingly, finding effective descriptions of ground states of many-body Hamiltonians is a very natural and important task in physics.

Given the prevalence and importance of this task, a natural question is that of the computational difficulty of solving this task in naturally occurring situations. Questions such as the hardness of solving a computational task belong to the domain of computational complexity theory. A good proxy for the difficulty of obtaining ground-state descriptions is the difficulty of solving a weaker problem, namely that of computing ground-state energies of many-body Hamiltonians. This question is studied in the domain known as "Hamiltonian complexity" (see e.g. Ref. [294]), an area of research at the intersection of quantum many-

body physics and computational complexity theory.

This area of research originated from Kitaev's result that the LocalHamiltonian problem, which is the problem of computing the ground-state energy of a local Hamiltonian, is QMA-complete [59] (we refer a reader unfamiliar with complexity-theoretic language to Appendix B). The complexity class QMA is the quantum generalization of NP. Kitaev's result may be viewed as an analogue of the seminal Cook-Levin theorem [295; 296] in computer science, generalized to the setting of quantum constraint satisfaction problems. Despite the tremendous amount of progress in understanding the power of local Hamiltonians, many important questions remain, such as whether the task remains hard under less-demanding notions of approximation error [297–299] and whether there exist short classical descriptions of ground states of local Hamiltonians (see e.g., Refs. [297; 300; 301]), among others.

One important question about LocalHamiltonian is the role played by the spectral gap. The spectral gap is a traditionally important quantity in the context of ground-state properties of any physical system and is defined as the difference between the smallest two eigenvalues of the Hamiltonian. Many important families of Hamiltonians in physics have the "gap property", meaning that the spectral gap in the limit of large system size $n \to \infty$ is lower-bounded by a constant. Important conjectures in physics are concerned with the existence of the gap property for certain Hamiltonians [302; 303], a problem that is known to be undecidable in general [304]. Furthermore, the existence of a spectral gap implies various tractability results for the ground states of Hamiltonians. For instance, in one dimension, the gap property significantly restricts the entanglement structure of ground states through the area law of entanglement, implying efficient classical representations of the same [305], and further, classically efficient algorithms to compute the ground-state energy [306; 307].

It is not known whether these properties hold for higher dimensions.

Despite the physical importance of the spectral gap, its role in the context of the LocalHamiltonian problem itself is much less clear. In particular, it is not known whether LocalHamiltonian is QMA-complete in the presence of nontrivial lower bounds on the spectral gaps, even when the lower bound is $\Omega(1/\mathsf{poly}(n))$ [308; 309]. Meanwhile, if the spectral gap is promised to be lower bounded by a constant, there are no-go results [310; 311] that rule out any QMA-hardness proof that proceeds by generalizing the clock construction technique. This technique underlies all known QMA-completeness results, in analogy with the theory of NP-completeness, where the Cook-Levin theorem plays a foundational role. Therefore, Hamiltonians with any nontrivial lower bounds on the spectral gap can be less complex than the general case.

In this Chapter, we take an initial step towards answering the question of the role played by the spectral gap in the LocalHamiltonian problem. To do so, we study QMA in the precise setting, i.e. the class PreciseQMA. In the precise setting, the completeness (the minimum probability of accepting a correct statement) and soundness (the maximum probability of accepting an incorrect statement) of the protocol are separated by a quantity called the promise gap that scales inverse-exponentially in the size of the input. For the LocalHamiltonian problem, this translates to computing the ground-state energy to within inverse-exponential precision in the system size.

Computing ground-state energies to inverse-exponential precision is not an artificial task. This task corresponds to computing polynomially many digits of the answer, which is very desirable in some cases [312]. Algorithms whose runtimes scale as $\mathsf{polylog}(1/\epsilon)$ for additive error ϵ can compute quantities to inverse-exponential precision in polynomial

time, and such algorithms have been found for Hamiltonian simulation and linear systems [40; 313; 314]. There are also situations where precise knowledge of the ground-state energy of a Hamiltonian is essential. For example, in quantum chemistry, chemical reactivity rates depend on the Born-Oppenheimer potential-energy surface for the nuclei. Each point on this surface is an electronic ground-state energy for a particular arrangement of the nuclei. Small uncertainties in the ground-state energy can exponentially influence the calculated rate k via Arrhenius's law $k \propto \exp[-\beta\Delta E]$, where ΔE is an energy barrier and β the inverse temperature (see, e.g., Ref. [315]). Another example is in condensed-matter physics, where algorithms such as the density matrix renormalization group (DMRG) routinely compute several digits of the ground-state energy (see, e.g., Ref. [316]). Precise knowledge of the ground-state energy can enable one to identify the locations of quantum phase transitions by identifying non-analyticities [317]. Interestingly, the class of Hamiltonians for which the energy can be precisely measured correspond to Hamiltonians that can be fast-forwarded [318].

Fefferman and Lin [319] studied the complexity of the class PreciseQMA, and showed the mysterious result that it equals PSPACE. This is surprising since QMA $\subseteq$ PP [320–322] (also see Fig. B.1 in Appendix B for reference), and an alternative characterization of the class PP is PreciseBQP, the precise analogue of BQP (the class of problems efficiently solvable on quantum computers). Since PreciseBQP can handle inverse-exponentially small promise gaps and contains QMA, one might have expected that adding the modifier Precise– to QMA would not have changed the power of the class by much.

We provide an explanation for this seemingly unexpected boost in complexity from

QMA, which is a subset of PP, to PreciseQMA, which equals PSPACE[1]. Specifically, we find that in order for the precise version of LocalHamiltonian, i.e. PreciseLocalHamiltonian, to be PSPACE-hard, the spectral gap of the Hamiltonian must necessarily shrink superpolynomially with the size of the system n (measured by the number of qudits in the system). We give strong evidence that if the spectral gap shrinks no faster than a polynomial in the system size, i.e. if the spectral gap is bounded by $\Omega(1/\mathsf{poly})$, the complexity of the problem is strictly less powerful. In particular, we show that this problem characterizes the complexity class PP, which is a subset of PSPACE and is widely believed to be distinct from PSPACE. If the problem were PSPACE-hard, the so-called counting hierarchy, defined as $\mathsf{CH} = \mathsf{PP} \cup \mathsf{PP}^{\mathsf{PP}} \cup \ldots$ [323], would collapse, which is considered an unlikely possibility. Our results therefore bring out the importance of the spectral gap, a quantity not well understood so far in Hamiltonian complexity.

Another main result of ours concerns the existence of polynomial-size quantum circuits to prepare ground states of local Hamiltonians. This is an important question that has implications in circuit-complexity of ground states of natural Hamiltonians and is directly related to whether natural Hamiltonians can be efficiently cooled down to zero temperature. In complexity-theoretic language, the question may be phrased in terms of the power of classical versus quantum witnesses in Merlin-Arthur proof systems, or more formally, the so-called QMA vs. QCMA question. The (in)equivalence of these classes is an important open question in quantum complexity theory and many-body physics, which has remained unsettled despite recent progress in the oracle setting (see e.g., Refs. [300; 301]). The precise version of QCMA, or PreciseQCMA, is known to be equal to $\mathsf{NP}^{\mathsf{PP}}$ (see e.g., Refs. [324;

[1]The class PSPACE contains PP, and is believed to be unequal to and much larger than PP: see Fig. B.1.

325]), indicating a separation between PreciseQCMA and PreciseQMA (= PSPACE) unless the counting hierarchy collapses. Interestingly, we show strong *equivalence* results for the PreciseQMA vs. PreciseQCMA question in the presence of spectral gaps.

Our results and the proof techniques we develop here also have consequences for other areas of complexity theory and many-body physics. Our second main result mentioned earlier roughly says that in the precise regime, the promise of an inverse-polynomial lower bound on the spectral gap is equivalent to the promise that there exists a polynomial-size circuit to prepare the ground state. We make an interesting conjecture in this regard in Section 6.1.3, which could have a bearing on the performance of near-term quantum algorithms for quantum chemistry and on the circuit complexity of various low-energy states, which is an important question in gravitational and high-energy physics [326; 327]. Furthermore, our results can shed light on an attempt to give a quantum-inspired reproof [328; 329] of the celebrated IP = PSPACE result [330] via interactive protocols for the class PreciseQMA. Our results also allow us to rule out sufficiently strong error-reduction techniques for the class postQMA.

This Chapter is structured as follows. In the rest of Section 6.1, we state the main results, give a high-level overview of the proof techniques and their implications, and discuss the relation of our results to other work in the literature. In Section 6.2, we give the definitions of some other complexity classes and define some new classes that appear in this Chapter. We also define natural problems complete for these classes. We then formally state the results pertaining to the class PP in Section 6.3 and PSPACE in Section 6.4. We also consider the complexity of related classes in Section 6.5, after which the Appendices have detailed proofs of our claims.

6.1.1 Results

We describe a general problem we study here, called (δ, Δ)-LocalHamiltonian. Informally, it is the problem of estimating the ground-state energy of a given k-local Hamiltonian acting on n qudits[2] to additive error at most δ, when promised that the spectral gap is at least Δ (see precise definitions in Section 6.2). In the absence of any bound on the spectral gap (i.e. $\Delta = 0$), the problem $(1/\mathsf{poly}(n), 0)$-LocalHamiltonian is, by definition, the same as k-LocalHamiltonian, which is complete for QMA for $k \geq 2$ [59; 331; 332]. Meanwhile, $(1/\exp(n), 0)$-LocalHamiltonian is, by definition, Precise-k-LocalHamiltonian [319], which is complete for PreciseQMA. We henceforth suppress the dependence on the number of qudits n in the notation exp and poly for the rest of the Chapter.

To our knowledge, Aharonov et al. [308] were the first to study the k-LocalHamiltonian problem in the presence of a spectral gap. Specifically, they considered $(1/\mathsf{poly}, 1/\mathsf{poly})$-LocalHamiltonian and showed it to be complete for the class PGQMA (Polynomially Gapped QMA). The definition of PGQMA, which is given in Section 6.2, depends on a notion of a spectral gap for *proof systems*, distinct from that for Hamiltonians. For complexity classes associated with proof systems such as QMA, QCMA and the variants we study in this Chapter, the spectral gap corresponds to the gap in the highest and second-highest accept probabilities of the optimal witness and the next-optimal orthogonal witness. A priori, the two notions of a spectral gap have no relation with each other. We show that the two notions are equivalent for various cases (δ and Δ each behaving as $1/\mathsf{poly}$ or $1/\exp$), by showing that (δ, Δ)-LocalHamiltonian is complete for the appropriate spectral-gapped

[2]The results in this Chapter are applicable generally to qudits of any dimension $d \geq 2$, but we will often work with qubits in our proofs.

QMA class.

To understand the relation between the gapped QMA classes and the regular versions without a spectral gap, we focus on the precise regime, so that $\delta = 1/\exp$ henceforth for the rest of this section. By specifying the spectral gap to be $\Omega(1/\mathsf{poly})$, we get the problem $(1/\exp, 1/\mathsf{poly})$-LocalHamiltonian. We show in Lemma 31 that this problem is in a class we call PrecisePGQMA (Precise Polynomially Gapped QMA), which is the precise analogue of PGQMA. We also show (Lemma 32) that PrecisePGQMA $\subseteq$ PP, implying that PrecisePGQMA is likely different from PreciseQMA, which equals PSPACE. Specifically, assuming that PP $\neq$ PSPACE, there is a separation between PrecisePGQMA and PreciseQMA. The PP upper bound on PrecisePGQMA is optimal: we show that $(1/\exp, 1/\mathsf{poly})$-LocalHamiltonian is PP-hard (Lemma 34). Thus, we tightly characterize the complexity of the class by showing PrecisePGQMA $=$ PP and prove that $(1/\exp, 1/\mathsf{poly})$-LocalHamiltonian is its associated complete problem.

The results in the previous paragraph show that the PSPACE-hardness result of Ref. [319] relies on the fact that the spectral gaps of the associated Hamiltonians can decay rapidly with the system size. This raises the question of the maximum scaling of the spectral gap required in order to retain PSPACE-hardness. This is an important question since if the PSPACE-hardness results only apply when there is no promise whatsoever on the spectral gap, it would indicate that PSPACE-hardness of Precise-k-LocalHamiltonian is artificial. We rule out this possibility by showing that if the spectral gap is bounded below by $1/\exp$, i.e. if we consider the problem $(1/\exp, 1/\exp)$-LocalHamiltonian, the problem remains PSPACE-hard. Specifically, we show in Theorem 23 that this problem is complete for a class called PreciseEGQMA (Precise Exponentially Gapped QMA). Next, we show that

Spectral gap (Δ)	(δ, Δ)-GS-Description-LocalHamiltonian		(δ, Δ)-LocalHamiltonian	
	$\delta = 1/\mathsf{poly}$	$\delta = 1/\exp$	$\delta = 1/\mathsf{poly}$	$\delta = 1/\exp$
$1/\mathsf{poly}$	PGQCMA {26} ($=_R$ QCMA {44})	PrecisePGQCMA {28} (= PP {30})	PGQMA	PrecisePGQMA {22} (= PP {29})
$1/\exp$	EGQCMA ($=_R$ QCMA {44})	PreciseEGQCMA {27} (= $\mathsf{NP}^{\mathsf{PP}}$ {45})	EGQMA(?)	PreciseEGQMA {23} (= PSPACE {38})
0	QCMA {24}	PreciseQCMA {25} (= $\mathsf{NP}^{\mathsf{PP}}$)	QMA	PreciseQMA (= PSPACE)

Table 6.1: Complexity of variants of the LocalHamiltonian problem as a function of the parameters δ, the promise gap, and Δ, the spectral gap. The problem is complete for the class mentioned in each cell. For reference, we mention in curly brackets the theorem number corresponding to the results proved in this Chapter. The question mark corresponding to the entry EGQMA indicates that the result is a conjecture and the notation $=_R$ denotes equivalence under randomized reductions (defined in Section 6.5.3).

PreciseEGQMA equals PSPACE (Theorem 38), implying that instances with $\Omega(1/\exp)$ spectral gaps are no less complex than the general case.

Lastly, we consider the analogues of these classes when the witness is classical, which gives us the classes QCMA (Quantum Classical Merlin Arthur), PreciseQCMA, PrecisePGQCMA (Precise Polynomially Gapped QCMA) and PreciseEGQCMA (Precise Exponentially Gapped QCMA). The complete problems for these classes are the appropriate versions of the Local-Hamiltonian problem under the additional promise that there is an efficient classical description of a circuit to prepare a low-energy state, as we show in Theorems 24 to 27. We define this problem in Section 6.2.1 and denote it (δ, Δ)-GS-Description-LocalHamiltonian, which is the problem of computing the ground-state energy to additive error δ, given the promise that there exists a polynomial-size circuit to prepare a low-energy state and promised that the spectral gap of the Hamiltonian is at least Δ. As stated in Corollary 37, we show that PrecisePGQCMA has the same complexity as PrecisePGQMA, implying that in the precise setting, once there is a $\Omega(1/\mathsf{poly})$ promise on the spectral gap, a further promise that there exists an efficient circuit to prepare a low-energy state is redundant. We comment more on this result in Section 6.1.3.

In Table 6.1, we give an overview of the parameter dependence of the complexity of two main problems studied in this Chapter, namely (δ, Δ)-LocalHamiltonian and (δ, Δ)-GS-Description-LocalHamiltonian. The problems are completely characterized by the appropriately gapped versions of QMA or QCMA, or their precise variants. The complexity class in any cell in the table is a subset of all the classes below it in the same column, since these classes correspond to weaker promises on the spectral gap. Similarly, the complexity class associated with (δ, Δ)-GS-Description-LocalHamiltonian is a subset of that associated with (δ, Δ)-LocalHamiltonian, because the former problem is associated with an extra promise. While we have given evidence that PrecisePGQMA $\neq$ PreciseQMA, it is unknown whether the same holds for the question PGQMA $\stackrel{?}{=}$ QMA. Similarly, while we have proved PreciseEGQMA $=$ PreciseQMA, it would be interesting to see if a similar result holds for EGQMA.

6.1.2 Techniques

Here, we give an overview of the primary techniques used in proving our results.

Imaginary-time evolution and the power method.— To show the containment PrecisePGQMA $\subseteq$ PP, we use a technique called the "power method" [333]. The broad idea behind the algorithm is that if a matrix A is promised to have a spectral gap between the largest two eigenvalues, the behavior of A^d for large d is dominated by the largest eigenvalue. We give a PP algorithm to compute $\mathrm{Tr}\left(A^d\right)$ for an exponentially large matrix A and $d = \mathsf{poly}(n)$ for a wide class of matrices A. This wide class includes sparse matrices and matrices representing local observables as special cases. The PP algorithm uses the Feynman sum-over-paths idea

[334] to express the trace as a sum over 2^{poly} many terms, each of which is a product over quantities of the form $\langle x| R |y\rangle$ for some matrix R whose entries are efficiently computable. A PP algorithm can decide whether the sum over 2^{poly} many terms, each term computable in polynomial time, is above or below a threshold.

The power method is closely related to another technique called the "cooling algorithm", inspired by a brief discussion by Schuch et al. [178]. The idea is that letting a system evolve in imaginary time can produce an unnormalized state close to the ground state. Imaginary-time evolution is a linear, albeit nonunitary, operation and produces an unnormalized state ρ' in general. Schuch et al. relied on a quantum characterization of PP, namely postBQP. The class postBQP [335] is the class of problems solvable in polynomial time on a quantum computer with access to the resource of postselection, which is the ability to condition on exponentially unlikely events. Aaronson [335] showed that any linear operation, even nonunitary ones, may be simulated in postBQP. Schuch et al.'s algorithm [178] works by decomposing the imaginary-time evolution operation $\exp[-\beta H]$ into a series of local operations $\exp[-\beta H_i]$ using Trotterization, and implementing each local operation using the resource of postselection. Unfortunately, the state-of-the-art error bounds for Trotterization of imaginary-time evolution [48] give, at best, a multiplicative error that is exponential in n (see also Refs. [44; 171]), and hence this technique does not work in the precise regime. Using a modification of this idea, we prove a more general statement about precise computation of ground-state local observables for Hamiltonians with a spectral gap. Specifically, we use the idea of imaginary time evolution and give a P^{PP} algorithm that provably works not just for $1/\mathsf{poly}$ precision, but also $1/\exp$ precision in computing local observables in addition to the Hamiltonian. Our technique is closely related to the power method, since the core of

the algorithm is to compute expectation values of powers of the Hamiltonian.

Small-penalty clock construction.— Our second major technical contribution is a modification of the clock construction that we call the small-penalty clock construction. One of the ways this technique is useful is as follows. As mentioned earlier and as will be described in detail in Section 6.2, it is possible to consider spectral-gapped versions of both the LocalHamiltonian problem and the class QMA and their variants. We have already discussed the (natural) notion of a spectral gap for Hamiltonians. For QMA and related classes, the spectral gap is related to the difference in accept probabilities between the optimal and next-optimal witness. Our technique allows us to bridge the notion of spectral gap in both cases by constructing spectral-gap-preserving reductions. In other words, the small-penalty clock construction allows us to prove that the Hamiltonians resulting from the construction inherit a spectral gap related to the gap in accept probabilities in the circuit, for several variants of QMA. This ability is used in the proofs of Theorems 23 to 27. An interesting feature of the modified clock construction is that it also allows us to show that, when there is a classical witness (i.e. a QCMA computation), the resulting Hamiltonian has a classical description for a state with energy close to the ground-state energy. Another related application of the small-penalty clock construction is that it also allows us to show complexity lower bounds like in Lemmas 34 and 36. In these cases, we directly reduce from PP to the appropriate gapped version of the LocalHamiltonian problem instead of a reduction from the corresponding $-$QMA class.

We now spell out what enables the small-penalty clock construction to show the above results. As mentioned before, the clock construction and its variants encompass all current proofs of hardness for QMA and related classes. Typically, this consists of mapping a circuit

to a Hamiltonian $H = H_{\text{input}} + H_{\text{prop}} + H_{\text{clock}} + H_{\text{output}}$. Roughly speaking, each term locally enforces that the computation is a valid step of a QMA protocol by adding energy penalties to undesirable states. The "witness register", where a quantum prover may input any quantum state, is left unpenalized and the Hamiltonian therefore has no terms acting on the witness register. The role of H_{output} is to ensure that witnesses and computations that lead to a *low* accept probability at the output get a *high* energy penalty. In the absence of the penalty term at the output, the ground-state space of the Hamiltonian is well-known and is given by the subspace of the so-called "history states", each with the same energy. The output penalty term H_{output} is what breaks the degeneracy and helps create a promise gap, and we will henceforth refer to this as simply the penalty term without qualification.

However, the addition of the penalty term makes the eigenstates of the Hamiltonian difficult to analyze, since the magnitude of the penalty can be large, i.e. $\Omega(1)$ in strength. In this Chapter, we often choose the output penalty terms to have small strength. This might seem like a strange choice to make since one is typically interested in making the promise gap as large as possible. However, since we are dealing with instances where the promise gap is already exponentially small, our choice is not too costly. The advantage this gives us is that the ground-state energy tracks the effect of the output penalty more faithfully. More concretely, the smallness of the penalty term allows us to use tools like the Schrieffer-Wolff transformation [336; 337], which can be viewed as a rigorous formulation of degenerate perturbation theory. We review the Schrieffer-Wolff transformation in Section 6.6.

Spectral gap in adjacency matrix.— For the proof of Theorem 38, we show a reduction[3] from a natural PSPACE-complete graph problem to an instance of a problem known as

[3]This reduction is inspired by unpublished work by one of us and Cedric Lin [338], and we supplement it

$(1/\exp, 1/\exp)$-SparseHamiltonian[4]. This problem is a generalization of $(1/\exp, 1/\exp)$-k-LocalHamiltonian, allowing for the Hamiltonian to be any sparse Hamiltonian with a spectral gap $\geq 1/\exp$. Sparse Hamiltonians are Hermitian matrices that can be exponentially large, with at most $\mathsf{poly}(n)$ nonzero entries per row in some basis and an efficient algorithm for computing any entry of the matrix. They are a generalization of local Hamiltonians.

The PSPACE-complete graph problem may be described as SuccinctGraphReachability, which is to decide if there is a path from one vertex to another in a succinctly described graph of exponential size (also see Ref. [318]). We show that one can always construct a PSPACE-bounded Turing machine such that the resulting Hamiltonian after the reduction always has a spectral gap that is at least $1/\exp(n)$. We do this through an explicit analysis of the eigenvalues of the Hamiltonian, which are related to the lengths of cycles and paths of the graph constructed from the Turing machine. Next, we give a PreciseEGQMA upper bound to $(1/\exp, 1/\exp)$-SparseHamiltonian, i.e. the problem in the presence of a spectral gap, establishing that PSPACE $\subseteq$ PreciseEGQMA.

6.1.3 Discussion

Our first main result was that the addition of even an inverse-polynomially small spectral gap takes the complexity of precisely estimating the ground-state energy of a local Hamiltonian from PreciseQMA $=$ PSPACE to PrecisePGQMA $=$ PP. Note that this result also implies a difference between the case of no spectral gap and a constant spectral gap. Therefore, we have given a *provable setting* where the difference in complexity between

with a technique to create spectral gaps.

[4]We actually show a reduction to the complement of the problem (where YES and NO instances are reversed), but this turns out not to matter because PSPACE is closed under complement.

two problems is attributable entirely to the spectral gap.

Our second main result concerned a modification of the same problem of precisely estimating the ground-state energy of a local Hamiltonian promised to have an inverse-polynomial spectral gap. When additionally promised that there exists a classical description of a circuit to prepare a state whose energy is exponentially close to the ground-state energy, our results show that the complexity of the problem does not get weaker. Specifically, we show that the class PrecisePGQCMA is equivalent to PrecisePGQMA.

The above equivalence result is in sharp contrast with the belief PreciseQCMA $\neq$ PreciseQMA in the non-spectral-gapped case. This inequality follows from the conjecture that $\mathsf{NP}^{\mathsf{PP}} \neq \mathsf{PSPACE}$, which, if false, would lead to a collapse of the counting hierarchy. The inequality PreciseQCMA $\neq$ PreciseQMA rules out the possibility of there being polynomial-size circuits to prepare ground states of local Hamiltonians to exponential precision, since otherwise the prover could simply supply a description of such a circuit. Our equivalence result that PrecisePGQMA $=$ PrecisePGQCMA is consistent with the following intriguing conjecture about the circuit-complexity of ground states of low-energy Hamiltonians, although it does not imply the conjecture.

Conjecture 21. *Consider any Hamiltonian H on n qubits with ground-state energy E_1 and a $1/\mathsf{poly}$ spectral gap. Then there exists a low-energy state $|\psi\rangle$ satisfying $\langle\psi| H |\psi\rangle \leq E_1 + 2^{-\mathsf{poly}(n)}$ that can be prepared by an efficient quantum circuit, namely a circuit of the form $|\psi\rangle = U|0\rangle^m$, where m and the size of U are both polynomials in n.*

Note that Conjecture 21 implies both PrecisePGQMA $=$ PrecisePGQCMA $=$ PP and PGQMA $=$ PGQCMA$(=_R$ QCMA$)$; however the reverse direction need not hold. Our

result does not imply Conjecture 21 because the reduction does not imply anything about how the PrecisePGQCMA verification protocol looks like. We also note that the quantum circuits referred to in Conjecture 21 may be hard to find–the conjecture is only concerned with the existence of such circuits, and not with whether these circuits can be obtained by an efficient algorithm. In complexity-theoretic language, these circuits may be nonuniform. This is why Conjecture 21 is not in contradiction with Ref. [339], which argues that finding efficient matrix-product-state representations of Hamiltonians with a $\Omega(1/\mathsf{poly})$ spectral gap can be hard.

If Conjecture 21 were true, it would also explain the observed success of quantum algorithms such as the variational quantum eigensolver (VQE) [64; 340], which seek to solve a much simpler problem of preparing low-energy states of translation-invariant many-body Hamiltonians with energy $1/\mathsf{poly}$-close to the ground-state energy. A large class of translation-invariant Hamiltonians have a spectral gap that is either a constant, $\Theta(1)$ (gapped phases), or vanishing in the system size as $\Theta(1/n^{1/D})$ (gapless phases described by conformal field theories in D-dimensions). Therefore, Conjecture 21 applies to both these cases and would imply the existence of polynomial-size circuits to prepare states with high overlap with the ground state. Such circuits are generally found in the VQE algorithm if one optimizes over sufficiently many parameters. This behavior is in line with other instances where a lower bound on the spectral gap implies tractability of the ground state in various senses [305; 307; 341; 342].

Coming to the case of exponentially small spectral gaps, we have shown that $\mathsf{PreciseEGQMA} = \mathsf{PreciseQMA}$. This implies that $\mathsf{PreciseEGQMA} \neq \mathsf{PreciseEGQCMA}$ unless the counting hierarchy collapses. Therefore, we have given a class of local Hamiltonians (in the proof of

Lemma 46) with exponentially small spectral gaps, whose ground states have exponentially large circuit complexity. This is a result of independent interest, and it might be interesting to study the whether these Hamiltonians can be classified as quantum spin glasses, which are believed to be hard to cool down to zero temperature [343].

In another intriguing line of work, Aharonov and Green [328] and Green, Kindler, and Liu [329] have given interactive protocols for precise quantum complexity classes with a computationally bounded prover $\mathcal{P}$ and a computationally bounded verifier $\mathcal{V}$, denoted $\mathsf{IP}[\mathcal{P}, \mathcal{V}]$. A goal of this line of work is to give a quantum-inspired proof of the result $\mathsf{IP} = \mathsf{PSPACE}$ [330] by giving an interactive protocol for PreciseQMA [329] (which equals PSPACE) with a BPP verifier. This has been successful so far with PreciseBQP and PreciseQCMA (which equals $\mathsf{NP}^{\mathsf{PP}}$) but not yet with PreciseQMA. From the result of Ref. [328] and our result that $\mathsf{PrecisePGQMA} = \mathsf{PP}$, there is an $\mathsf{IP}[\mathsf{PreciseBQP}, \mathsf{BPP}]$ protocol for PrecisePGQMA. Our results indicate that the spectral gap might play an important role in extending such an interactive protocol to PSPACE. Namely, such an extension would need to be able to work with inverse-exponentially small spectral gaps.

In addition, the class postQMA [325; 344] is the class where there is a quantum prover and a postBQP verifier, where one may condition (postselect) on exponentially unlikely outcomes. This class has been shown to be equal to PreciseQMA [325], so an alternative approach mentioned by Green et al. [329] to reprove the result $\mathsf{IP} = \mathsf{PSPACE}$ is to exhibit an $\mathsf{IP}[\mathsf{postQMA}, \mathsf{BPP}]$ protocol for postQMA. To complete such a proof, it would suffice to prove a witness-preserving amplification technique like in QMA [322; 345] that additionally handles postselection. Witness-preserving amplification is a technique for improving the promise gap of an interactive protocol by modifying the verifier's strategy while keeping

the witness fixed. We show in Lemma 41 that, assuming PP $\neq$ PSPACE, the soundness of a postQMA protocol cannot be reduced beyond a particular point without requiring the witness to grow larger or requiring the postselection success probability to shrink. Therefore, we obtain evidence that a witness-preserving amplification technique for postQMA should differ significantly from the technique of Marriott and Watrous [322], since in the latter, repeating the verifier's circuit suffices to get any soundness parameter $s \leq 2^{-\mathsf{poly}}$.

So far, we have considered the spectral-gap promise to be applicable to both YES and NO instances of the problems defined. We can also define asymmetric problems where only the YES instances are promised to have a spectral gap. The motivation for considering such asymmetric promises is that they are related to complexity classes where the accepting witness is promised to be unique, such as the class UQMA [308]. The problems with asymmetric promises can only be harder than their symmetric analogues, since the promise is weaker. We show that for both $\Omega(1/\mathsf{poly})$ and $\Omega(1/\exp)$ spectral gaps in the precise setting, there is no difference between symmetric and asymmetric promises on the spectral gaps. Specifically, we show in Theorem 42 that the classes with asymmetric promises are of the same complexity as those with symmetric promises.

We remark here that the promise of a spectral gap above a unique ground state is distinct from assuming that we have a UQMA instance. The reason is that for LocalHamiltonian, the presence of a spectral gap does not imply that there is a unique accepting witness, it only implies a unique ground state. In case the ground-state subspace is polynomially degenerate, the PP algorithm continues to work to produce estimates of the ground-state energy.

Lastly, we add that results shown in the precise regime do not always imply analogous

results in the non-precise regime. For example, our work gives evidence that PrecisePGQCMA $\neq$ PreciseQCMA, but in the non-precise regime we can show PGQCMA $=_R$ QCMA. In this respect, inequivalence results in the high-precision regime resemble oracle separation results in complexity theory, which is a mature area of research with several important results [346–348]. While oracle separations do not constitute strong evidence for the inequivalence of two complexity classes, they are useful in ruling out proof techniques that work relative to oracles, or "relativize". Similarly, inequivalence results in the precise regime can rule out proof techniques from extending to the precise regime. For example, a purported proof that QCMA $=$ QMA must not work in the precise regime, otherwise we would obtain PreciseQCMA $=$ PreciseQMA, or PSPACE $=$ PP, which is believed to be unlikely.

6.1.4 Related work

The study of Hamiltonian complexity [60–63; 294; 331; 332; 349; 350] has given rise to many techniques and important results applicable in quantum many-body physics, such as [304; 339; 351–357]. The clock construction has also been analyzed in detail recently [358–360].

The study of exponentially small promise gaps in the context of quantum classes can be traced to Watrous [361], who defined PQP and showed its equivalence with postBQP, which equals PP [335]. In the precise setting, one can sometimes give far stronger evidence for the (in)equivalence of complexity classes than in the analogous bounded error setting, as is the case for precise versions of the questions of QCMA vs. QMA [319] and QMA(2) vs. QMA [319; 362–364]. There has been work on quantum interactive proof systems with

exponentially small promise gaps, such as in the context of $\mathsf{QMA}(2)$ [364], or with even smaller gaps, such as in Refs. [365–367]. Fefferman and Lin [319; 338] studied the precise regime of QMA, showing it to equal PSPACE, leading to other works concerning precise classes [325; 368]. Gharibian et al. [324] considered quantum generalizations of the polynomial hierarchy, where precise classes and spectral gaps are relevant to the definitions and proof techniques.

Aharonov et al. [308] were the first to consider the complexity of the LocalHamiltonian problem in the presence of spectral gaps, motivated by the question of uniqueness [369] for randomized and quantum classes. They showed the equivalence of UQCMA and QCMA, and that of UQMA and PGQMA, using similar techniques as Valiant and Vazirani [369] in their proof of equivalence of UNP and NP. Jain et al. [309] defined the class FewQMA and showed that it is contained in $\mathsf{P}^{\mathsf{UQMA}}$, giving a technique to reduce the dimension of accepting witnesses.

More recently, González-Guillén and Cubitt [310] studied the spectral gap of a large class of Hamiltonians that encode history states in their ground state and showed that the spectral gap is upper bounded by $O(1/\mathsf{poly})$. A similar result was obtained by Crosson and Bowen [311] using different techniques. These works are mainly concerned with the existence of a $\Theta(1)$ spectral gap, whereas our results distinguish between $1/\mathsf{poly}$ and $1/\exp$ spectral gaps.

Finally, Ambainis [370] studied the problem of estimating spectral gaps and local observables and gave a $\mathsf{P}^{\mathsf{QMA}[\log]}$ upper bound for these problems, while also giving $\mathsf{P}^{\mathsf{QMA}[\log]}$-hardness results (also see Ref. [371]). The class $\mathsf{P}^{\mathsf{QMA}[\log]}$ is the class of problems solvable in polynomial time by making logarithmically many (adaptive) queries to a QMA oracle.

Gharibian and Yirka [371] showed that $\mathsf{P}^{\mathsf{QMA[log]}} \subseteq \mathsf{PP}$ and extended previous hardness results to more natural Hamiltonians. Gharibian, Piddock, and Yirka [372] also gave a very natural complete problem for the class $\mathsf{P}^{\mathsf{QMA[log]}}$ in the context of computing local observables in ground states. Novo et al. [373] have recently studied the closely-related problem of sampling from the distribution obtained by making energy measurements and obtain various interesting hardness results, under different notions of error.

6.2 Definitions and complete problems

We have seen the definition of BQP in terms of the class $\mathsf{BQP}[c, s]$ with general parameters c and s. The Precise- version of BQP can be defined similarly.

Definition 2. $\mathsf{PreciseBQP} = \cup_{c-s \geq 1/\exp} \mathsf{BQP}[c, s]$.

This class is known to be equal to PP (see, e.g., Ref. [324]).

We now give an equivalent definition of QMA in terms of the eigenvalues of an operator called the *accept operator*. We will then define a very general class called Gapped QMA, $\mathsf{GQMA}[c, s, g_1, g_2]$, which has several parameters. By specifying these parameters, we can define the major complexity classes in this Chapter. The complexity classes corresponding to classical witnesses (QCMA and its derivatives) are defined analogously.

The alternative definition of QMA is in terms of the "accept operator" $Q(U_x) = \langle 0|^{\otimes m} U_x^\dagger \Pi_{\text{out}} U_x |0\rangle^{\otimes m}$ on the witness register, where Π_{out} is the projector on to the accept state ($|1\rangle_o$). For any state $|\Psi\rangle$ provided as a witness, the quantity $\langle\Psi| Q_x |\Psi\rangle$ is the accept probability of the circuit. We will henceforth suppress the dependence of Q on the unitary U_x and the instance x. The eigenvalues of Q, $\lambda_1(Q) \geq \lambda_2(Q) \geq \ldots$ are important

quantities to consider since the accept probability of any input proof state is a convex combination of these eigenvalues. The alternative definition of QMA in terms of the operator Q is as follows:

Definition 3 (Alternative definition of $\mathsf{QMA}[c, s]$)**.** $A = (A_{\text{yes}}, A_{\text{no}})$ is a $\mathsf{QMA}[c, s]$ problem iff for every instance x there exists a uniformly generated circuit U_x of size $\mathsf{poly}(n)$ acting on $m + w = \mathsf{poly}(n)$ qubits, with the property that

If $x \in A_{\text{yes}}$: $\lambda_1(Q) \geq c$

If $x \in A_{\text{no}}$: $\lambda_1(Q) \leq s$,

where $Q = Q(U_x)$ is as above.

Note that we are typically interested in the behavior of the maximum accept probability, which equals the largest eigenvalue of Q. We are also interested in the lowest eigenvalue of a Hamiltonian H for the LocalHamiltonian problem and its variants. Therefore, we order eigenvalues in nonincreasing order for accept operators and in nondecreasing order for Hamiltonians. For the same reason, we define the spectral gap differently for accept operators and Hamiltonians. For a Hamiltonian, we define the spectral gap to be the difference in the smallest two eigenvalues $E_2 - E_1$. For accept operators, the spectral gap is the difference between the *highest two eigenvalues* $\lambda_1(Q) - \lambda_2(Q)$. This is equal to the difference in the accept probabilities of the optimal witness and the next-optimal witness orthogonal to it. It will usually be clear from context which spectral gap we are referring to.

Now let us define the class $\mathsf{GQMA}[c, s, g_1, g_2]$. It corresponds to a promise on the operator Q having a spectral gap of at least g_1 in the YES case, and at least g_2 in the NO case:

Definition 4 (Gapped QMA)**.** $\mathsf{GQMA}[c, s, g_1, g_2]$ is the class of promise problems $A =$

$(A_{\text{yes}}, A_{\text{no}})$ such that for every instance x, there exists a polynomial size verifier circuit U_x acting on $\mathsf{poly}(n)$ qubits and its associated accept operator Q such that

If $x \in A_{\text{yes}}$: $\lambda_1(Q) \geq c$ and $\lambda_1(Q) - \lambda_2(Q) \geq g_1$

If $x \in A_{\text{no}}$: $\lambda_1(Q) \leq s$ and $\lambda_1(Q) - \lambda_2(Q) \geq g_2$.

This definition is a generalization of the class PGQMA (Polynomially Gapped QMA) defined by Aharonov et al. in Ref. [308]:

Definition 5. $\mathsf{PGQMA} = \cup_{c-s, g_1, g_2 \geq 1/\mathsf{poly}} \mathsf{GQMA}[c, s, g_1, g_2]$.

To see the relation of this class with QMA, notice that by setting $g_1 = g_2 = 0$, the promise on spectral gaps becomes vacuous, since $\lambda_1(Q) \geq \lambda_2(Q)$ by definition. Therefore, we get the equality $\mathsf{GQMA}[c, s, 0, 0] = \mathsf{QMA}[c, s]$. We also define

Definition 6 (Exponentially Gapped QMA)**.** $\mathsf{EGQMA} = \cup_{\substack{c-s \geq 1/\mathsf{poly} \\ g_1, g_2 \geq 1/\exp}} \mathsf{GQMA}[c, s, g_1, g_2]$.

We now come to precise versions of these classes, where the completeness–soundness gap $c - s$ can be exponentially small, giving us more powerful classes. The first of these is PreciseQMA, which was defined in Ref. [319] and shown to be equal to PSPACE.

Definition 7. $\mathsf{PreciseQMA} = \cup_{c-s \geq 1/\exp} \mathsf{QMA}[c, s]$.

This definition should be compared to the precise version of GQMA, which comes in two varieties: the spectral gaps can either be polynomially small (PrecisePGQMA) or exponentially small (PreciseEGQMA).

Definition 8 (PrecisePGQMA)**.** PrecisePGQMA, short for Precise Polynomially Gapped QMA, is the class with exponentially small promise gaps and polynomially small spectral gaps:

$$\mathsf{PrecisePGQMA} = \cup_{\substack{c-s \geq 1/\exp \\ g_1, g_2 \geq 1/\mathsf{poly}}} \mathsf{GQMA}[c, s, g_1, g_2].$$

Definition 9 (PreciseEGQMA)**.** PreciseEGQMA, short for Precise Exponentially Gapped QMA, has both the promise gap and spectral gap exponentially small:

$$\mathsf{PreciseEGQMA} = \cup_{\substack{c-s\geq 1/\exp \\ g_1,g_2\geq 1/\exp}} \mathsf{GQMA}[c, s, g_1, g_2].$$

We now come to complexity classes in which the prover sends a classical witness but the verifier remains quantum. The classicality of the witness can be enforced by measuring the qubits sent by the prover in the computational basis and interpreting qubits in the computational basis as classical bits. If the verifier is only allowed to make measurements at the end, we use the standard protocol for deferring measurements: we apply a "copy operation" U_c that has CNOTs from the qubits in the witness register to an ancilla register in the state $|0\rangle^w$. We leave the qubits in the witness state unmeasured. This modified circuit has the property that it preserves the accept probabilities of input witness states that are in the computational basis. Further, the eigenstates of the modified accept operator acting on the register can be taken to be computational basis states. This allows us to define QCMA and its derivatives in terms of the accept operator and also allows us to consider a gapped version of QCMA:

Definition 10 ($\mathsf{GQCMA}[c, s, g_1, g_2]$)**.** $A = (A_{\text{yes}}, A_{\text{no}})$ is a $\mathsf{GQCMA}[c, s]$ problem iff for every instance x there exists a uniformly generated circuit U_x of size $\mathsf{poly}(n)$ acting on $m + w = \mathsf{poly}(n)$ qubits, with the property that

If $x \in A_{\text{yes}}$: $\lambda_1(Q) \geq c$ and $\lambda_1(Q) - \lambda_2(Q) \geq g_1$

If $x \in A_{\text{no}}$: $\lambda_1(Q) \leq s$, and $\lambda_1(Q) - \lambda_2(Q) \geq g_2$,

where $Q = Q(U_x U_c)$ is the accept operator of the modified circuit with the copy operation U_c described above.

Definition 11. The derived classes of GQCMA are given by

- $\mathsf{QCMA}[c,s] = \mathsf{GQCMA}[c,s,0,0]$.
- $\mathsf{QCMA} = \cup_{c-s>1/\mathsf{poly}} \mathsf{QCMA}[c,s]$.
- $\mathsf{PreciseQCMA} = \cup_{c-s>1/\exp} \mathsf{QCMA}[c,s]$.
- Polynomially Gapped QCMA: $\mathsf{PGQCMA} = \cup_{\substack{c-s>1/\mathsf{poly} \\ g_1,g_2>1/\mathsf{poly}}} \mathsf{GQCMA}[c,s,g_1,g_2]$.
- Precise Polynomially Gapped QCMA:

  $\mathsf{PrecisePGQCMA} = \cup_{\substack{c-s>1/\exp \\ g_1,g_2>1/\mathsf{poly}}} \mathsf{GQCMA}[c,s,g_1,g_2]$.
- Exponentially Gapped QCMA: $\mathsf{EGQCMA} = \cup_{\substack{c-s>1/\mathsf{poly} \\ g_1,g_2>1/\exp}} \mathsf{GQCMA}[c,s,g_1,g_2]$.
- Precise Exponentially Gapped QCMA:

  $\mathsf{PreciseEGQCMA} = \cup_{\substack{c-s>1/\exp \\ g_1,g_2>1/\exp}} \mathsf{GQCMA}[c,s,g_1,g_2]$.

6.2.1 Complete problems

We now come to the definitions of problems that are complete for these classes. The classic problem complete for the class QMA is the LocalHamiltonian problem [59; 331; 332]. We define a k-local observable to be a Hermitian operator A that can be written as a sum over operators A_i supported on k qudits at most: $A = \sum_i^{\mathsf{poly}(n)} A_i$. We assume that each term has bounded operator norm $\|A_i\| \leq \mathsf{poly}(n)$. The task in the LocalHamiltonian problem is to estimate the ground-state energy of a local Hamiltonian. The decision version of the problem is as follows:

k-LocalHamiltonian$[a, b]$

Input A description of a k-local Hamiltonian $H = \sum_i h_i$ on n qubits with $h_i \succeq 0$, two numbers a and b with $b > a$.

Output YES if the ground-state energy $E_1 \leq a$,

NO if $E_1 \geq b$, promised that one of them is the case.

Henceforth we omit the phrase "promised that one of them is the case" because we will be exclusively considering promise problems unless otherwise specified. Kitaev [59] showed that 5-LocalHamiltonian$[a, b]$ with $b - a = \Omega(1/\mathsf{poly})$ is QMA-complete, which was improved to $k = 3$ and then $k = 2$ in Refs. [331; 332]. The parameter $\delta := b - a$, the promise gap, is a measure of the accuracy to which the solution is desired. We define the problem in terms of δ only, as follows:

Definition 12. δ-k-LocalHamiltonian $:= \cup_{b-a\geq\delta}$ k-LocalHamiltonian$[a, b]$.

We now come to the gapped and precise versions of the problem, which turn out to be complete for their respective $-$QMA variants. We also suppress the notation k in the name of the problem, though there is formally a dependence on k. In this Chapter, our hardness results hold for $k \geq 3$ and it may be possible to improve our results to hold for $k = 2$.

LocalHamiltonian$[a, b, g_1, g_2]$

Input Description of a k-local Hamiltonian $H = \sum_i h_i$ with $h_i \succeq 0$, numbers a, b, g_1, and g_2 with $b > a$.

Output YES if the ground-state energy $E_1 \leq a$ and any state orthogonal to the ground state has energy $\geq E_1 + g_1$,

NO if $E_1 \geq b$ and any state orthogonal to the ground state has energy $\geq E_1 + g_2$.

In both the YES and NO cases above, we see that the Hamiltonian has a unique ground state and a spectral gap of at least g_1 in the YES case and g_2 in the NO case. The above problem with promise gap $\delta = b - a$ and spectral gap $\Delta = \min[g_1, g_2]$ is defined to be:

Definition 13. (δ, Δ)-LocalHamiltonian $:= \cup_{\substack{b-a \geq \delta \\ g_1, g_2 > \Delta}}$ LocalHamiltonian$[a, b, g_1, g_2]$.

In the non-precise regime, the problem $(1/\mathsf{poly}, 1/\mathsf{poly})$-LocalHamiltonian was shown to be complete for PGQMA for $k \geq 2$ [308].

We now focus on the precise regime, i.e. $\delta = \Omega(1/\exp)$. From the results of Ref. [319], we know that $(1/\exp, 0)$-LocalHamiltonian is PreciseQMA-complete for $k \geq 3$. We show that:

Theorem 22. $(1/\exp, 1/\mathsf{poly})$-*LocalHamiltonian is* PrecisePGQMA-*complete.*

Theorem 23. $(1/\exp, 1/\exp)$-*LocalHamiltonian is* PreciseEGQMA-*complete.*

By virtue of these theorems, we can talk about the complexity of the classes PrecisePGQMA and PreciseEGQMA interchangeably with their complete problems. The proofs of these theorems are given in Sections 6.7 and 6.8. The hardness results rely on the small-penalty clock construction, where the size of the penalty term is either $\Theta(1/\mathsf{poly})$ or $\Theta(1/\exp)$. The

upper bounds are shown in Lemmas 52 and 53 and rely on a modification of the standard phase-estimation protocol used to show k-LocalHamiltonian is in QMA. Specifically, we consider the modified protocol of Ref. [319] used for Precise-k-LocalHamiltonian and observe that the spectral gaps in the energies translate to separations in the accept probabilities.

Finally, we turn to complete problems for QCMA and its derivatives. The first problem, GS-Description-LocalHamiltonian, concerns finding the ground-state energy of a k-local Hamiltonian when there is a polynomial-size circuit to prepare a state close to the ground state (which constitutes a classical description of the ground state).

GS-Description-LocalHamiltonian$[a, b, g_1, g_2]$

Input Description of a k-local Hamiltonian $H = \sum_i h_i$, numbers a, $b \geq a + \delta$, polynomials $T(n), m(n)$, together with the promise that there exists a circuit V of size T such that $V|0^m\rangle = |\psi\rangle$ satisfies $\langle\psi| H |\psi\rangle \leq E_1 + \delta^3/f(n)^2$ for some polynomial $f(n) \geq \|H\|$.

Output YES if the ground-state energy of H satisfies $E_1 \leq a$ and the spectral gap of H is at least g_1,

NO if $E_1 \geq b$ and the spectral gap of H is at least g_2.

Definition 14. (δ, Δ)-GS-Description-LocalHamiltonian $:= \cup_{b-a\geq\delta, g_1, g_2 \geq \Delta}$ GS-Description-LocalHamiltonian$[a, b, g_1, g_2]$

As in the case of (δ, Δ)-LocalHamiltonian, if we take $\Delta = 0$, we get a version without any promise on the spectral gap. This is a close relative of the following problem

proved to be QCMA-complete for $\delta = \Omega(1/\text{poly})$ [374].

δ-LowComplexity-LowEnergyStates

Input Description of a k-local Hamiltonian $H = \sum_i h_i$, numbers a, b and polynomials $T(n)$, $m(n)$, with $b \geq a + \delta$.

Output YES if there exists a circuit of size $\leq T(n)$ that acts on $|0^m\rangle$ to prepare a state $|\psi\rangle$ with energy $\langle\psi| H |\psi\rangle \leq a$,

NO if any state $|\psi\rangle$ obtained by applying a circuit of size $T(n)$ on $|0^m\rangle$ has energy $\langle\psi| H |\psi\rangle \geq b$.

This latter problem has a weaker promise than $(\delta, 0)$-GS-Description-LocalHamiltonian. This is because a NO instance of δ-LowComplexity-LowEnergyStates is automatically a NO instance of $(\delta, 0)$-GS-Description-LocalHamiltonian, since any state necessarily has energy $\geq b$. Meanwhile, a NO instance of $(\delta, 0)$-GS-Description-LocalHamiltonian need not be a NO instance of δ-LowComplexity-LowEnergyStates, since for the latter there is no guarantee of a circuit to prepare a state with energy close to the ground-state energy.

Despite having a stronger promise on $(\delta, 0)$-GS-Description-LocalHamiltonian (which only makes the problem less complex), our small-penalty clock construction allows us to prove the same hardness result for both $\delta = 1/\text{poly}$ and $\delta = 1/\exp$:

Theorem 24. $(1/\text{poly}, 0)$-*GS-Description-LocalHamiltonian is* QCMA-*complete.*

Theorem 25. $(1/\exp, 0)$-*GS-Description-LocalHamiltonian is* PreciseQCMA-*complete.*

For the latter theorem in the precise regime, we use the small-penalty clock construction with an exponentially small energy penalty. Lastly, when we add the promise of spectral

gaps, we have the following results:

Theorem 26. $(1/\mathsf{poly}, 1/\mathsf{poly})$-*GS-Description-Local Hamiltonian is* PGQCMA-*complete.*

Theorem 27. $(1/\exp, 1/\exp)$-*GS-Description-Local Hamiltonian is* PreciseEGQCMA-*complete.*

Theorem 28. $(1/\exp, 1/\mathsf{poly})$-*GS-Description-Local Hamiltonian is* PrecisePGQCMA-*complete.*

The upper bounds in Theorems 24 to 28 follow from a precise version of phase estimation, together with the promise that there is a classical description of a circuit to prepare a low-energy state. The lower bounds either follow directly through a small-penalty clock construction or through a reduction from a class that contains the relevant class.

6.3 Problems characterized by PP

In this section, we discuss the complexity of the classes PrecisePGQMA and PrecisePGQCMA, both of which turn out to equal PP.

Theorem 29. PrecisePGQMA = PP.

Theorem 30. PrecisePGQCMA = PP.

We describe here the overall strategy for proving these results. First, we adapt the one-bit phase estimation circuit in Ref. [319] to show that it is possible to compute ground-state energies of sparse Hamiltonians with a spectral gap in the corresponding GQMA class. In particular, we have

Lemma 31. $(1/\exp, 1/\mathrm{poly})$-*Local Hamiltonian* $\in$ PrecisePGQMA.

Next, we use the "power method" [333] to give a PP algorithm for any problem in PrecisePGQMA.

Lemma 32 (One half of Theorem 29)**.** PrecisePGQMA $\subseteq$ PP.

Proof. Suppose we have a GQMA$[c, s, g_1, g_2]$ instance. Then we should give a PP algorithm to precisely compute the maximum eigenvalue λ_1 of the accept operator Q associated with the instance, under the promise that the spectral gap of Q is bounded below by an inverse polynomial. In particular, the spectral gap of the accept operator, given by $\lambda_1 - \lambda_2$, is at least $\min[g_1, g_2] =: \Delta$. Consider the power method to compute the maximum eigenvalue and eigenvector of a positive semidefinite operator Q. This method relies on the observation that upon taking positive powers of the operator Q and estimating its trace, the quantity is dominated by the maximum eigenvalue of Q. In the following, we suppress the dependence of λ_i on Q:

$$\mathrm{Tr}(Q^q) = \sum_i \lambda_i^q \tag{6.1}$$

$$= \lambda_1^q \left(1 + \left(\frac{\lambda_2}{\lambda_1}\right)^q + \ldots\right) \tag{6.2}$$

$$\leq \lambda_1^q + \lambda_1^q (2^w - 1)\left(1 - \frac{\Delta}{\lambda_1}\right)^q. \tag{6.3}$$

On the other hand, we have $\mathrm{Tr}(Q^q) \geq \lambda_1^q$. Therefore, in the YES case, we have

$$\mathrm{Tr}(Q^q) \geq c^q, \tag{6.4}$$

while in the NO case,

$$\mathrm{Tr}(Q^q) \leq s^q + s^q(2^w - 1)\left(1 - \frac{\Delta}{\lambda_1}\right)^q \tag{6.5}$$

$$\leq s^q + s^q(2^w - 1)\left(1 - \frac{\Delta}{s}\right)^q, \tag{6.6}$$

where w is the size of the witness register and we assume $s \geq \Delta \geq 1/\mathsf{poly}$, since otherwise the promise cannot be satisfied. The difference in the two cases is

$$c^q - s^q - s^q(2^w - 1)\left(1 - \frac{\Delta}{s}\right)^q \tag{6.7}$$

$$= c^q - s^q - s^q(2^w - 1)\exp\left[q\log\left(1 - \frac{\Delta}{s}\right)\right] \tag{6.8}$$

$$\geq c^q - s^q - s^q 2^w \exp\left[-\frac{q\Delta}{s}\right] \tag{6.9}$$

$$= s^q\left(\left(1 + \frac{c-s}{s}\right)^q - 1\right) - s^q 2^w \exp\left[-\frac{q\Delta}{s}\right] \tag{6.10}$$

$$\geq s^q\left(\frac{q(c-s)}{s} - 2^w \exp\left[-\frac{q\Delta}{s}\right]\right). \tag{6.11}$$

$$\geq s^q\left(\frac{c-s}{s} - 2^w \exp\left[-\frac{q\Delta}{s}\right]\right). \tag{6.12}$$

If we pick $q = \left\lceil \frac{s}{\Delta}\log\left(\frac{c-s}{2^{w+1}s}\right)\right\rceil = O(\mathsf{poly})$, we can ensure that the difference in $\mathrm{Tr}(Q^q)$ between the YES and NO cases is at least

$$s^q\frac{c-s}{2s} = \Omega(2^{-\mathsf{poly}}). \tag{6.13}$$

This observation suggests that a PP algorithm can decide between the YES and NO cases by computing $\mathrm{Tr}(Q^q)$ for some large enough polynomial q. This is possible because a

PP algorithm can compute a sum of 2^{poly} terms, where every term is efficiently computable in polynomial time. We justify this more rigorously in Section 6.10 (Lemma 60). □

The above result implies that, since $(1/\exp, 1/\mathsf{poly})$-LocalHamiltonian is in PrecisePGQMA, a PP algorithm can precisely compute ground-state energies of local Hamiltonians with a $\Omega(1/\mathsf{poly})$ spectral gap. A similar technique can also be used to show a slightly more general result:

Lemma 33. *Given a local Hamiltonian H and a local observable A, along with a promise that $\|A\| = O(\mathsf{poly})$ and the spectral gap of H is lower-bounded by $\Omega(1/\mathsf{poly})$, a P^{PP} algorithm can decide if the ground-state local observable $\langle E_1|A|E_1\rangle$ is either $\leq a$ or $\geq b$, for $b-a = \Omega(2^{-\mathsf{poly}})$, where $|E_1\rangle$ is the ground state of H.*

This lemma is proved in Section 6.10. Note that both of these results include the case $\Delta = \Theta(1)$, the important case of constant spectral gaps.

We complete the characterization of the power of PrecisePGQMA with the following result.

Lemma 34. $(1/\exp, 1/\mathsf{poly})$-*LocalHamiltonian is* PP-*hard.*

For this proof, we use the small-penalty clock construction, albeit one for the class PreciseBQP as opposed to the class PreciseQMA. In this aspect, it resembles the clock construction of Aharonov et al. [351], where it was used to show BQP-universality of the model of adiabatic quantum computing. We use the technique of applying $\Theta(1/\mathsf{poly})$ small penalties at the output so as to preserve the lower bound of $\Omega(1/\mathsf{poly})$ on the spectral gap shown in Ref. [351]. In sum, Lemmas 31, 32 and 34 together imply Theorems 22 and 29.

We now come to the class PreciseQCMA and its complete problem, $(1/\exp, 0)$-GS-Description-LocalHamiltonian, where we are promised that there is an efficient circuit to prepare a low-energy state. We know that PreciseQCMA $=$ $\mathsf{NP}^{\mathsf{PP}}$ [325], which lies in the second level of the counting hierarchy. Since PrecisePGQMA is characterized by PP, the promise of having a spectral gap is only slightly stronger than the promise of an efficient circuit to prepare the ground state.

Consider now the gapped version of the problem, $(1/\exp, 1/\mathsf{poly})$-GS-Description-LocalHamiltonian, where there is a $1/\mathsf{poly}$ spectral gap in addition to the promise of an efficient circuit to prepare the ground state. This characterizes the class PrecisePGQCMA, for which the proof technique is similar to PrecisePGQMA.

We first show that the gapped version of GS-Description-LocalHamiltonian is in the corresponding GQCMA class, and in particular,

Lemma 35. $(1/\exp, 1/\mathsf{poly})$-*GS-Description-LocalHamiltonian* $\in$ PrecisePGQCMA.

PP-hardness of the problem follows by the same argument as the proof of Lemma 34:

Lemma 36. $(1/\exp, 1/\mathsf{poly})$-*GS-Description-LocalHamiltonian is* PP-*hard.*

We give a unified proof of Lemmas 34 and 36 in Section 6.7. Since PrecisePGQCMA $\subseteq$ PrecisePGQMA $=$ PP, this implies:

Corollary 37. PrecisePGQMA $=$ PrecisePGQCMA $=$ PP.

6.4 Problems characterized by PSPACE

In this section, we discuss the complexity of the class PreciseEGQMA, which turns out to equal PSPACE.

Theorem 38. PreciseEGQMA $=$ PreciseQMA $(=$ PSPACE$)$.

Proof. The containment PreciseEGQMA $\subseteq$ PreciseQMA follows trivially since any PreciseEGQMA instance is automatically a PreciseQMA instance. We show the other direction, PreciseEGQMA $\supseteq$ PreciseQMA, in two steps. Our proof relies on the complexity of the following problem:

SparseHamiltonian$[a, b, g_1, g_2]$

Input A succinct description of a Hermitian matrix of size $2^{\mathsf{poly}(n)} \times 2^{\mathsf{poly}(n)}$, with at most $d = \mathsf{poly}(n)$ many entries in each row and two numbers a and b, with $b > a$. The magnitude of each entry is bounded by $k = \mathsf{poly}(n)$.

Output YES if the smallest eigenvalue $E_1 \leq a$ and the spectral gap of the matrix is at least g_1,

NO if $E_1 \geq b$, and the spectral gap of the matrix is at least g_2.

We define (δ, Δ)-SparseHamiltonian to be $\cup_{\substack{b-a \geq \delta \\ g_1, g_2 \geq \Delta}}$ SparseHamiltonian$[a, b, g_1, g_2]$ and consider the problem with parameters $\delta, \Delta = \Omega(1/\exp)$. First, in Lemma 39, we prove that $(1/\exp, 1/\exp)$-SparseHamiltonian is PSPACE-hard, or equivalently PreciseQMA-hard. Next, we show in Lemma 40 that $(1/\exp, 1/\exp)$-SparseHamiltonian may be solved in PreciseEGQMA. The theorem then follows. □

Lemma 39. $(1/\exp, 1/\exp)$-*SparseHamiltonian is* PSPACE*-hard.*

The reduction is from any problem in PSPACE to an instance of co-$(1/\exp, 1/\exp)$-Gapped-SparseHamiltonian, which is the complement of the problem, in the sense that the YES and NO instances are reversed. Since PSPACE is closed under complement, this still gives

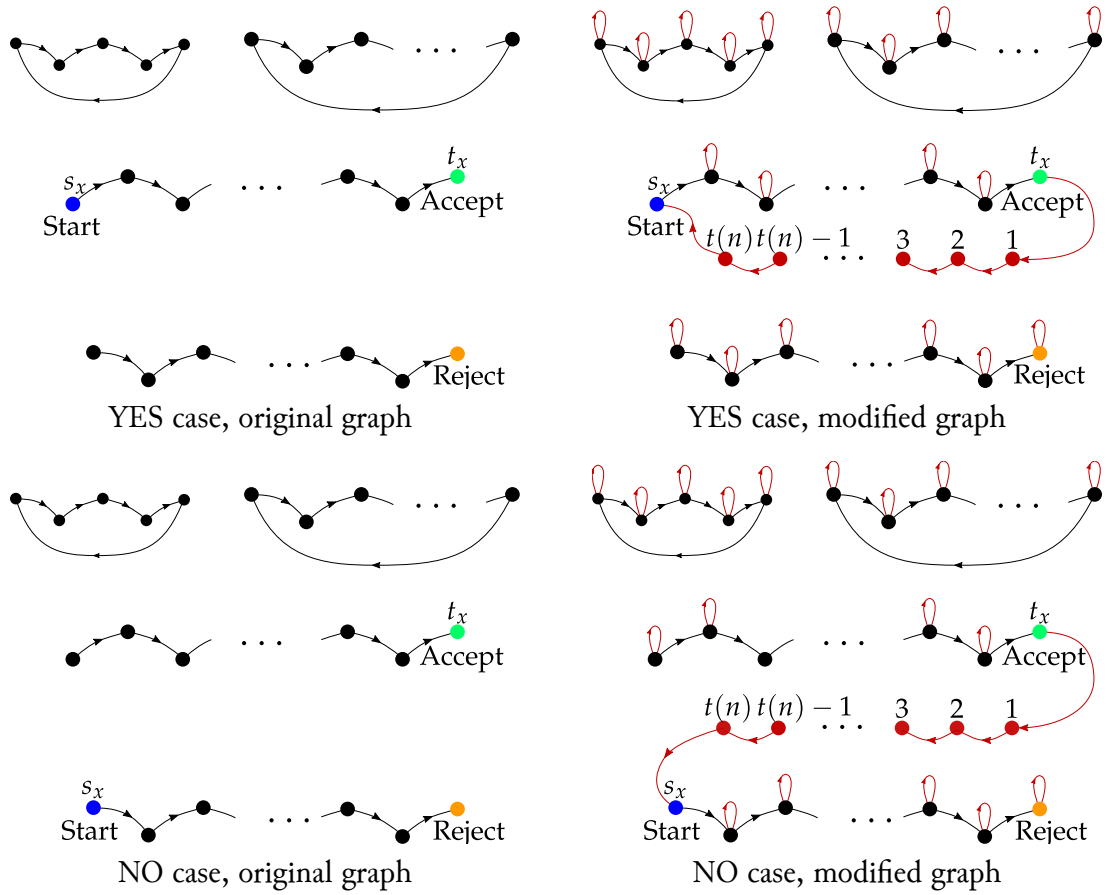

Figure 6.1: Schematic of the original and modified graphs for both YES and NO cases. The original graph in both YES and NO cases consists of vertices with in-degree and out-degree at most 1, due to the fact that the Turing machine is reversible. The start vertex s_x is marked in blue, the accept vertex t_x in green, and the reject vertex in orange. The modified graphs have self-loops on all vertices except the start and the accept vertices. They have additional vertices $1, 2, \ldots t(n)$ without self-loops. All modifications are in maroon.

the desired hardness result. The broad idea is to represent a PSPACE computation as an exponentially large, but sparse, graph. The smallest eigenvalue of the adjacency matrix of this graph encodes information about whether the computation accepts or rejects.

Proof. We use a proof technique adapted from an unpublished manuscript by Fefferman and Lin [338]. First, we use the fact that PSPACE with reversible operations in every step still equals PSPACE: revPSPACE = PSPACE [375]. Indeed, it is known that $\mathsf{SPACE}[\mathfrak{s}(n)] = \mathsf{revSPACE}[\mathfrak{s}(n)]$ [376] with an overhead in time that is exponential in the space, $\mathfrak{s}(n)$. Let $t(n)$ be this upper bound on the running time of the Turing machine, so that we can restrict our attention to the class $\mathsf{revSPACE}[\mathfrak{s}(n)] \cap \mathsf{TIME}[t(n)] = \mathsf{SPACE}[\mathfrak{s}(n)]$. Any computation on a reversible Turing machine may be viewed as traversing a directed *configuration graph*, where each vertex of the graph is determined by the state of the head and the list of symbols on the input and work tapes (Fig. 6.1). When such a Turing machine is restricted to use space polynomial in the input length n, the number of vertices in the graph is upper bounded by an exponential, $2^{\mathsf{poly}(n)}$. Consider the adjacency matrix of the graph, A_x. The description of this exponentially large matrix is succinct because it only requires specifying the input x and the rules of the Turing machine.

We modify the configuration graph $G_x \to G'_x$ so that the smallest eigenvalue of the matrix $A_x^{\dagger\prime} A'_x$ is 0 in the NO case and bounded away by an exponentially small amount in the YES case. We do this modification in a way that ensures the matrix has a spectral gap lower bound of at least $\Omega(1/\exp)$. This is done as follows. First, we modify the configuration graph of the Turing machine by adding self-loops to all vertices except for the start and accept configurations s_x and t_x. We then add a sequence of vertices $\{1, 2, \ldots t(n)\}$ from

the accept configuration t_x, with the directed edges $t_x \to 1 \to 2 \to \ldots \to t(n) \to s_x$, as shown in Fig. 6.1. The adjacency matrix of this modified directed graph G'_x is A'_x, and we are interested in the eigenvalues and spectral gap of $A_x^{\dagger\prime} A'_x$, which is Hermitian and sparse, and also has a succinct representation.

We now analyze this construction. The proof relies on an explicit computation of the eigenvalues for the various subgraphs of the modified configuration graph. In the NO case, the graph G'_x has a path of vertices ending in the reject state (Fig. 6.1). This path contains the starting configuration s_x. Let ℓ be the graph distance between s_x and the reject state. Since we have added the edges $t_x \to 1 \to \ldots t(n) \to s_x$, these vertices and the vertices leading to the accept state are also part of the path (the Turing machine does not explore these vertices in practice). All vertices in this path except for t_x, s_x, and $i : i \in [t(n)]$ have self-loops on them. As we show in Lemma 63, there is a zero eigenvalue in the NO case, with a spectral gap above the zero eigenvalue. The spectral gap is lower bounded by $\Omega(1/\ell_{\text{max}}^2) = \Omega(2^{-\text{poly}})$, where ℓ_{max} is the number of vertices in the longest subgraph.

In the YES case, the subgraph containing the starting vertex is a cycle, with self-loops on all vertices except for t_x, s_x, and the intermediate vertices i. In each case, the eigenvalues for any subgraph are given by $2 - 2\cos\frac{(2k-1)\pi}{2\ell+1} = 4\sin^2\left(\frac{(2k-1)\pi}{4\ell+2}\right)$, $k \in [\ell]$ [338], where ℓ is the number of vertices in the subgraph. The smallest eigenvalue is therefore given by the longest subgraph and this eigenvalue is nondegenerate if no two subgraphs have the same number of vertices. This is why we have added the sequence of edges $t_x \to 1 \to \ldots t(n)$. The role played by these vertices is to elongate the length of the subgraph containing the start and accept configurations by $t(n)$. This ensures that no other subgraph has a length equal to the longest subgraph (since $t(n)$ is the upper bound on the total number of vertices in the graph

before elongation). Therefore, the smallest two eigenvalues are given by $4\sin^2\left(\frac{(2k-1)\pi}{4\ell+2}\right)$, which are separated by $\Theta(t(n)^{-2}) = \Theta(2^{-\mathsf{poly}})$.

To summarize, in the YES case we have $E_1 \geq 2^{-\mathsf{poly}}$ and $E_2 - E_1 \geq 2^{-\mathsf{poly}}$. In the NO case, we have $E_1 = 0$ and $E_2 \geq 2^{-\mathsf{poly}}$. Therefore, we have a promise gap of $2^{-\mathsf{poly}}$ and spectral gap $2^{-\mathsf{poly}}$ in both the YES and NO instances. Furthermore, the matrix $A_x^{\dagger\prime} A'_x$ has entries of magnitude at most 2, and is 3-sparse because of the bounded degree of the configuration graph. Since the minimum eigenvalue is small in the NO case and large in the YES case, we have a reduction to co-$(1/\exp, 1/\exp)$-SparseHamiltonian. Due to the fact that PSPACE is closed under complement, we get PSPACE-hardness of $(1/\exp, 1/\exp)$-SparseHamiltonian. □

Lemma 40. $(1/\exp, 1/\exp)$-*SparseHamiltonian* $\in$ PreciseEGQMA.

The proof of this is mostly the same as the proof of containment of $(1/\exp, 1/\exp)$-LocalHamiltonian in PreciseEGQMA and is also given in Section 6.8. The only difference is that we have a sparse Hamiltonian instead of a local Hamiltonian. This distinction turns out not to matter, however, because of quantum algorithms for Hamiltonian evolution that work well with sparse Hamiltonians [40].

6.5 Other related classes

In this section, we discuss implications of our proof techniques for other complexity classes. The first concerns a technique for amplifying the promise gap in QMA and related classes, called in-place amplification, due to Marriott and Watrous [322]. The second is about the complexity of related classes when the spectral gap promise only applies to one kind of in-

stance (YES instances, for example). We also complete a discussion of the results in Table 6.1 by characterizing the complexity classes PGQCMA, EGQCMA, and PreciseEGQCMA.

6.5.1 Amplification for postQMA

We first define the class postQMA:

Definition 15 (postQMA)**.** postQMA$[c, s]$ is the class of promise problems $A = (A_{\text{yes}}, A_{\text{no}})$ that can be decided in the following way: Apply a uniformly generated quantum circuit U of size poly(n) on a state $|x\rangle$ encoding the input, together with a proof state of size $w(n)$ supplied by an arbitrarily powerful prover. Postselect the first $l = \text{poly}(n)$ qubits at the output onto the $|0\rangle^l$ state, and measure the first qubit of the remaining register at the output, called the decision qubit (o). The postselection probability is $\Omega(2^{-f(n)})$ for a polynomial $f(n)$.

If $x \in A_{\text{yes}}$: $\exists\ |\psi\rangle$ such that $\Pr(o = 1) \geq c$

If $x \in A_{\text{no}}$: $\forall\ |\psi\rangle$, $\Pr(o = 1) \leq s$.

Morimae and Nishimura [325] defined this class and showed that postQMA $:=$ postQMA$[\frac{1}{3}, \frac{2}{3}] =$ PreciseQMA $=$ PSPACE. This result is similar to the result postBQP $=$ PreciseBQP$(=$ PP). They raised the question of whether one can do a Marriott-Watrous type in-place amplification for this class, which, for instance, means boosting the parameters c and s to be $c = 1 - 2^{-\text{poly}}$, $s = 2^{-\text{poly}}$ without changing the size of the witness. If one is allowed to change the witness size, one can simply ask for polynomially many copies of the witness and run the verification in parallel to get the required parameters. The benefit of in-place amplification is that it allows for good completeness and soundness parameters without blowing up the witness size, which turns out to be useful in the proof of QMA $\subseteq$ PP. In-place

amplification for postQMA would also be useful to show IP = PSPACE [328; 329]. Here we give a negative result for a sufficiently strong in-place amplification for postQMA.

Lemma 41 (Upper bound for in-place amplified postQMA). *If* $f(n) = O(w(n))$*, then* $\mathsf{postQMA}[1 - 2^{-t(n)}, 2^{-u(n)}] \subseteq \mathsf{PP}$ *for* $u(n) > w(n) + 1$ *and for any polynomial* $t(n) > 1$.

Proof. Consider a $\mathsf{postQMA}[1 - 2^{-t(n)}, 2^{-u(n)}]$ language. Replace the witness state in the amplified protocol by a maximally mixed state $\frac{\mathbb{1}}{2^w}$. Now, since the overlap of any witness state with the maximally mixed state is 2^{-w}, we have that the postselection success probability is at least $\Omega(2^{-f(n)-w(n)})$. Further, in the YES case, the probability of accepting the string x (conditioned on success) is

$$\Pr(o = 1) \geq 2^{-w(n)} \times (1 - 2^{-t(n)}). \tag{6.14}$$

In the NO case, we have that no matter what state is in the witness register, the accept probability is

$$\Pr(o = 1) \leq 2^{-u(n)}. \tag{6.15}$$

In $\mathsf{PreciseBQP} = \mathsf{PP}$, we can distinguish between these two cases if $2^{-w} - 2^{-t-w} > 2^{-u}$, i.e. if $1 - 2^{-t} > 2^{w-u}$, for which it suffices to have $u(n) > w(n) + 1$ and $t > 1$. □

This result implies that the completeness-soundness gap for postQMA cannot be boosted beyond a point without incurring a blowup in the size of the witness or by reducing the success probability of postselection.

6.5.2 Asymmetric promises on spectral gap and uniqueness

Motivated by a possible connection to the study of unique witnesses for quantum complexity classes, we consider the complexity class $\mathsf{GQMA}[c, s, g_1, 0]$. Here, there is no promise on the spectral gap for NO instances. In the YES case, we have $\lambda_1(Q) \geq c$ and $\lambda_2 \leq \lambda_1 - g_1 \leq 1 - g_1$. If we choose the spectral gap g_1 to be larger than $1 - s$, we see that $\lambda_2 \leq s$, ensuring that in the YES case, there is exactly one accepting witness[5]. The existence of one accepting witness is exactly the promise that defines the class UQMA:

Definition 16 (Unique QMA [308])**.** $\mathsf{UQMA}[c, s]$ is the class of promise problems $A = (A_{\text{yes}}, A_{\text{no}})$ such that for every instance x, there exists a polynomial-size verifier circuit U_x acting on $m = \mathsf{poly}(n)$ qubits and an input quantum proof on $w = \mathsf{poly}(n)$ qubits and the associated accept operator Q has properties

If $x \in A_{\text{yes}}$: $\lambda_1(Q) \geq c$ and $\lambda_2(Q) \leq s$

If $x \in A_{\text{no}}$: $\lambda_1(Q) \leq s$.

Definition 17. $\mathsf{UQMA} := \cup_{c-s \geq 1/\mathsf{poly}} \mathsf{UQMA}[c, s]$.

The earlier statement can be rephrased as "an instance of $\mathsf{GQMA}[c, s, 1 - s, 0]$ is a $\mathsf{UQMA}[c, s]$ instance". In the reverse direction, we can see that a $\mathsf{UQMA}[c, s]$ instance necessarily has a spectral gap $\lambda_1 - \lambda_2 \geq c - s$, and therefore is an instance of $\mathsf{GQMA}[c, s, c - s, 0]$. This hints at, but does not prove, an equivalence between the promise of uniqueness and that of an asymmetric spectral gap of $\Omega(1/\mathsf{poly})$. Aharonov et al. [308] proved a stronger result by showing that the class UQMA is equivalent to the class PGQMA under randomized reductions (defined below), where PGQMA is the class with spectral gaps for *both* the YES

[5] In the sense that any witness orthogonal to the accepting witness rejects with probability at least $1 - s$.

and the NO cases.

In the precise regime, we show the following results for the asymmetric variants of PrecisePGQMA and PreciseEGQMA.

Theorem 42. $\mathsf{PrecisePGQMA} = \cup_{\substack{c-s\geq 1/\exp, \\ g_1 \geq 1/\mathsf{poly}}} \mathsf{GQMA}[c, s, g_1, 0]$.
$\mathsf{PreciseEGQMA} = \cup_{\substack{c-s\geq 1/\exp, \\ g_1 \geq 1/\exp}} \mathsf{GQMA}[c, s, g_1, 0]$.

The proofs are given in Section 6.12 and hinge on the problem of computing ground-state energies when there is a spectral gap only for the YES case, i.e. LocalHamiltonian$[a, b, g_1, 0]$. Since the problem with an asymmetric gap can only be more complex than the symmetric case, the nontrivial part of this lemma is to show that this problem has the same PP upper bound as the symmetric case. This is not straightforward since the power method we described before does not necessarily work for the NO case, since there is no spectral gap. We work around this by making use of Ambainis's technique [370] of identifying spectral gaps, which is possible in PP [371].

6.5.3 Complexity of PGQCMA, EGQCMA, and PreciseEGQCMA

In this subsection we show that the classes PGQCMA and EGQCMA are both equivalent to QCMA under randomized reductions, which we now define.

We say a problem A is random reducible to problem X if every instance a of A can be mapped to a random set of polynomially instances x_i of X, such that

If $a \in A_{\text{yes}}$: $\Pr_i(x_i \in X_{\text{yes}}) \geq 1/\mathsf{poly}$

If $a \in A_{\text{no}}$: $\Pr_i(x_i \in X_{\text{yes}}) = 0$.

A class Y is random reducible to another class Z if every problem in Y is random reducible to some problem in Z (and vice versa), and is denoted $=_R$.

To show PGQCMA $=_R$ QCMA and EGQCMA $=_R$ QCMA, we make use of the class UQCMA (Unique QCMA), which has been defined in Ref. [308], and was shown to be equal to QCMA under randomized reductions.

Definition 18 (UQCMA$[c, s]$ [308])**.** UQMA$[c, s]$ is the class of promise problems $A = (A_{\text{yes}}, A_{\text{no}})$ such that for every instance x, there exists a polynomial-size verifier circuit U_x acting on $m = \text{poly}(n)$ qubits and an input classical proof on $w = \text{poly}(n)$ qubits, whose associated accept operator Q has properties

If $x \in A_{\text{yes}}$: $\lambda_1(Q) \geq c$ and $\lambda_2(Q) \leq s$

If $x \in A_{\text{no}}$: $\lambda_1(Q) \leq s$.

Definition 19. UQCMA $:= \cup_{c-s \geq 1/\text{poly}}$UQCMA$[c, s]$.

Aharonov et al. [308] showed that UQCMA $=_R$ QCMA using generalizations of techniques in Ref. [369] to complexity classes with randomness. In order to show PGQCMA $=_R$ QCMA and EGQCMA $=_R$ QCMA, we show

Lemma 43. PGQCMA $=_R$ UQCMA.

Since PGQCMA $\subseteq$ EGQCMA $\subseteq$ QCMA, the equivalence of EGQCMA with QCMA follows.

To show Lemma 43, we observe that the proof of PGQMA $=_R$ UQMA in Ref. [308] works for classical witnesses. For completeness, we give a self-contained proof here.

Proof of Lemma 43. First, we show the direction UQCMA $\subseteq$ PGQCMA. We observe that in a YES instance of UQCMA$[c, s]$, $\lambda_1 \geq c$ and $\lambda_2 \leq s$. Thus, a YES instance already has a spectral gap of $g_1 \geq c - s$ and is a YES instance of PGQCMA. In the NO case, we modify the verifier's strategy so that it creates a spectral gap. The verifier expects an additional qubit

we call the "flag qubit" from the prover, which is measured in the beginning just like the other qubits of any QCMA proof. The associated accept operator now has twice as many eigenvalues because it acts on a space with one larger qubit.

The verifier's protocol is as follows. If the state of the flag qubit is $|0\rangle$, the verifier continues with the original protocol. This gives the same eigenvalues for the accept operator as the original protocol. If the state of the flag qubit is $|1\rangle$, the verifier accepts with probability $s+(c-s)/\mathsf{poly}$ if the state of the rest of the witness qubits is $|1\rangle^{\otimes w}$. If the state of the rest of the witness register is anything else, the verifier rejects. In the latter case (when the state of the flag qubit is $|1\rangle$), the accept operator has one eigenvalue at $s+(c-s)/\mathsf{poly}$ and 2^w-1 eigenvalues with eigenvalue 0, each case corresponding to some state in the witness. The modified verifier is a PGQCMA instance with completeness c, soundness $s+(c-s)/\mathsf{poly}$ and spectral gaps $g_1 \geq c-s$ and $g_2 \geq (c-s)(1-1/\mathsf{poly})$. Therefore UQCMA $\subseteq$ PGQCMA.

For the other direction, we give a randomized reduction PGQCMA $\subseteq_R$ UQCMA. Consider a YES instance of PGQCMA$[c, s, g_1, g_2]$. We know $\lambda_1 \geq c$ and $\lambda_2 \leq \lambda_1 - g_1$, but we do not know if $\lambda_2 \leq s$, as is required for the instance to be a UQCMA instance. The idea in Ref. [308] is to make a query to a UQCMA$[c_j, s_j]$ oracle with completeness $c_j = c+(j+1)g_1/2$ and soundness $s_j = c + jg_1/2$, for j chosen randomly from $\{0, 1, \ldots \lfloor \frac{2}{(1-c)g_1} \rfloor\}$. In the NO case, all the queries are valid queries to a UQCMA oracle and return the correct answer (NO). In the YES case, since the completeness and soundness in each query differ by $g_1/2$, there is at least one j where $\lambda_1 \geq c_j$ and $\lambda_2 \leq s_j$[6]. Therefore, this is a randomized reduction to UQCMA. □

[6]In the YES case, there could be some queries that are not valid UQCMA instances, and the oracle can answer arbitrarily for such ill-formed queries. This does not, however, hamper the proof, since at single valid query is enough to give a nonzero probability of saying YES.

Therefore, we obtain

Corollary 44. PGQCMA $=_R$ QCMA.

EGQCMA $=_R$ QCMA.

Our final result concerns the class PreciseEGQCMA. Just like we have PreciseEGQMA $=$ PreciseQMA, we can show that exponentially small spectral gaps are no less complex in the case of classical witnesses. We show

Lemma 45. PreciseEGQCMA $=$ PreciseQCMA.

Proof. The direction PreciseEGQCMA $\subseteq$ PreciseQCMA is trivial. For the other direction, we take a PreciseQCMA$[c, s]$ instance and give a PreciseEGQCMA$[c, s, g_1, g_2]$ instance with an exponentially small spectral gap. This is done by modifying the verifier so that no two witnesses y_i and y_j are accepted with the same probability. First, we choose the verifier's gate set so that the accept probability of any witness y is given by $k_{x,y}/2^{l(n)}$, for $k_{x,y} \in [2^{l(n)}]$, where $l(n)$ is the size of the verifier's circuit [377]. The modified verifier rejects the instance straightaway with probability $y_b/2^{\mathsf{poly}}$, where y_b is a number in $[2^w - 1]$ when interpreting the witness y in binary and the polynomial is at least $l(n) + w(n) + \log_2(\frac{1}{c-s})$. If the verifier does not reject at this step, they run the original verification protocol. The overall accept probability when given y is given by $p_y = \frac{k_{x,y}}{2^w}\left(1 - \frac{y_b}{2^{\mathsf{poly}}}\right)$. Since the polynomial satisfies $\mathsf{poly} \geq l(n) + w(n) + \log_2(\frac{1}{c-s})$, the completeness and soundness are given by $c' \geq c - 2^{-w(n)}(c-s)$ and $s' = s$, which are still separated by $2^{-\mathsf{poly}}$.

We now claim that the resulting accept probabilities are distinct for distinct witnesses, and hence separated by an amount $\Omega(2^{-\mathsf{poly}})$. This is easily seen for two distinct y_i and y_j

such that $k_{x,y_i} = k_{x,y_j}$. If $k_{x,y_i} \neq k_{x,y_j}$, then for $p_{y_i} = p_{y_j}$, we need

$$k_{x,y_i} - k_{x,y_j} = \frac{2^w}{2^{l+w+\mathsf{poly}}}(y_{j_b} - y_{i_b}), \tag{6.16}$$

which cannot be satisfied by integers y_{j_b} and y_{i_b} in $[2^l]$. □

The same technique also works to give a more direct proof of $\mathsf{EGQCMA} = \mathsf{QCMA}$.

6.6 The Schrieffer-Wolff transformation

In this section, we give a brief introduction to the Schrieffer-Wolff transformation [336], which is an important tool in some of our subsequent proofs. We follow the exposition in Ref. [337], specialized to our context.

In the context relevant for us, we usually have an "unperturbed" Hamiltonian H_0 and a "perturbation" H_1, together forming the full Hamiltonian $H = H_0 + H_1$. The (possibly degenerate) ground-state subspace of H_0, denoted $\mathcal{S}_0$, has energy λ_0 and is separated from the rest of the spectrum by a gap Δ. We are interested in cases when the Hamiltonian H_1 has small strength relative to the gap Δ, in the sense $\|H_1\| =: \epsilon < \Delta/2$. This ensures that all eigenvalues of H_0 are shifted by an amount smaller than $\Delta/2$ under the perturbation. Therefore, the low-energy subspace of H, given by

$$\mathcal{S} = \left\{ |\psi\rangle : \langle\psi| H |\psi\rangle \in \left[\lambda_0 - \frac{\Delta}{2}, \lambda_0 + \frac{\Delta}{2}\right] \right\}, \tag{6.17}$$

has the same dimension as that of H_0. We denote the the projectors on to $\mathcal{S}_0$ and $\mathcal{S}$ by P_0 and P, respectively. As long as $\epsilon < \Delta/2$, we have $\|P - P_0\| < 1$, which captures the fact

that the dimension of the two subspaces is the same.

Since the dimension of the two subspaces is the same, there exists a unitary U that maps the subspace $\mathcal{S}_0$ to $\mathcal{S}$:

$$UPU^\dagger = P_0, \text{ with} \tag{6.18}$$

$$U = \sqrt{(2P_0 - \mathbb{1})(2P - \mathbb{1})}. \tag{6.19}$$

We are interested in the effective Hamiltonian in the subspace $\mathcal{S}_0$, given by

$$H_{\text{eff}} = P_0 U (H_0 + H_1) U^\dagger P_0. \tag{6.20}$$

The Schrieffer-Wolff transformation allows one to express the generator $V = \log(U)$, and consequently, H_{eff}, as a convergent series in the perturbation H_1. We first write H_1 as $H_1^d + H_1^o$, where H_1^d is block-diagonal in the subspace $\mathcal{S}_0$ and H_1^o is block-off-diagonal. Let the eigenstates of H_0 be given by $\{|i\rangle\}$, with corresponding energies $\{E_i\}$. We denote $\mathcal{I}_0 = \{i : E_i = \lambda_0\}$, which is the set of indices corresponding to the ground-state space. The first few terms of the Schrieffer-Wolff expansion are given by

$$\begin{aligned} H_{\text{eff}} &= H_0 P_0 + P_0 H_1 P_0 + \\ &\frac{1}{2} P_0 \sum_{i \in \mathcal{I}_0, j \notin \mathcal{I}_0} \left(\frac{\langle i | H_1 | j \rangle}{E_i - E_j} |i\rangle\langle j| H_1 + \frac{\langle j | H_1 | i \rangle}{E_i - E_j} H_1 |j\rangle\langle i| \right) P_0 \\ &+ O(\|H_1\|^3). \end{aligned} \tag{6.21}$$

In our work, we use the first-order expansion of the Schrieffer-Wolff series. The series

converges absolutely as long as $\|H_1\| \leq \Delta/16$ [337]. We can upper bound the error caused by truncating the formal series to first order:

$$\begin{aligned}
&\|H_{\text{eff}} - H_0 P_0 - P_0 H_1 P_0\| \leq O(1)\times \\
&\left\| P_0 \sum_{i\in\mathcal{I}_0, j\notin\mathcal{I}_0} \left(\frac{\langle i| H_1 |j\rangle}{E_i - E_j} |i\rangle\langle j| H_1 + \frac{\langle j| H_1 |i\rangle}{E_i - E_j} H_1 |j\rangle\langle i| \right) P_0 \right\| \\
&\leq O(1) \left\| \sum_{i\in\mathcal{I}_0, j\notin\mathcal{I}_0, k\in\mathcal{I}_0} \frac{1}{E_i - E_j} \left(\langle i| H_1 |j\rangle \langle j| H_1 |k\rangle |i\rangle\langle k| + \right.\right. \\
&\left.\left. \langle j| H_1 |i\rangle \langle k| H_1 |j\rangle |k\rangle\langle i| \right) \right\| \qquad (6.22) \\
&\leq O\left(\frac{1}{\Delta}\right) \left\| \sum_{i\in\mathcal{I}_0, k\in\mathcal{I}_0} \left(\langle i| H_1^2 |k\rangle |i\rangle\langle k| + \langle k| H_1^2 |i\rangle |k\rangle\langle i| \right) \right\| \qquad (6.23) \\
&= O\left(\frac{1}{\Delta}\right) \left\| 2 P_0 H_1^2 P_0 \right\| \qquad (6.24) \\
&\leq O\left(\frac{\epsilon^2}{\Delta}\right), \qquad (6.25)
\end{aligned}$$

where we have used $|E_i - E_j| > \Delta$ for states $i \in \mathcal{I}_0, j \notin \mathcal{I}_0$.

6.7 Modified clock constructions with spectral gaps

In this section, we present the small-penalty clock construction and use it to prove the main hardness results in this Chapter. We first illustrate the technique by proving the following lemma.

Lemma 46. $(1/\exp, 1/\exp)$-*Local Hamiltonian is* PreciseEGQMA-*hard.*

Proof. Consider a $\mathsf{GQMA}[c, s, g_1, g_2]$ instance x, where the verifier's circuit U_x acts on $m = \mathsf{poly}(n)$ qubits apart from the proof state. We assume that the circuit has $T = \mathsf{poly}(n)$ gates.

The idea behind the technique is valid generally, but for concreteness we focus on the clock construction of Kempe et al. [331], which proves QMA-hardness of k-LocalHamiltonian for $k \geq 3$. The clock Hamiltonian takes the form

$$H = H_{\text{input}} + H_{\text{prop}} + H_{\text{output}} + H_{\text{clock}}. \tag{6.26}$$

The first term H_{input} ensures that the ground state of H_{input} coincides with input state to the circuit. The term on the proof register is identity, allowing for any witness state given by the prover to be input into the verifier's circuit. It is given by

$$H_{\text{input}} = \sum_{i=1}^{m} |1\rangle\langle 1|_i \otimes \mathbb{1}_{\text{proof}} \otimes H_{\text{clockinit}}. \tag{6.27}$$

In the above, the term $H_{\text{clockinit}}$ ensures that the clock is properly initialized to the $|1\rangle_{\text{clock}}$ state. Next, H_{prop} is a Hamiltonian that ensures the ground state is "propagated" correctly with each gate applied by the verifier:

$$\begin{aligned} H_{\text{prop}} = \sum_{i=0}^{T} & -U_{i+1} \otimes |i+1\rangle\langle i|_{\text{clock}} - U_{i+1}^{\dagger} \otimes |i\rangle\langle i+1|_{\text{clock}} \\ & + \mathbb{1} \otimes (|i\rangle\langle i|_{\text{clock}} + |i+1\rangle\langle i+1|_{\text{clock}}). \end{aligned} \tag{6.28}$$

The ground-state subspace of H_{prop} contains valid "partial" computations until step $i \leq T$, namely $U_i \ldots U_2 U_1 |\psi_0\rangle$ on any initial state $|\psi_0\rangle$ $\forall$ i. The term H_{output} penalizes states that

have any nonzero probability of saying "NO" at the output qubit o of the circuit:

$$H_{\text{output}} = \epsilon \left|0\right\rangle\!\left\langle 0\right|_o \otimes \left|T\right\rangle\!\left\langle T\right|_{\text{clock}}. \tag{6.29}$$

Lastly, H_{clock} ensures that states in the clock register that do not encode a valid time step are penalized. The Hamiltonians H_{clock} and $H_{\text{clockinit}}$ both depend on the details of the particular clock construction. Our analysis does not depend on these details is largely independent of the way the clock register encodes the time. We refer the reader to Ref. [331] for an explanation of their construction.

First consider just the Hamiltonian $H_0 = H_{\text{input}} + H_{\text{prop}} + H_{\text{clock}}$, which is the clock Hamiltonian without a penalty term at the output. The ground-state space of H_0 is exactly given by the subspace $\mathcal{S}_0$ of history states:

$$\begin{aligned} &\mathcal{S}_0 = \text{span}\{|\phi_h\rangle : |\phi\rangle \text{ arbitrary}\}, \text{where} \\ &|\phi_h\rangle := \frac{1}{\sqrt{T+1}} \sum_{i=0}^{T} U_i \ldots U_0 \left|0^m\right\rangle \otimes |\phi\rangle \otimes |i\rangle_{\text{clock}}. \end{aligned} \tag{6.30}$$

where $U_0 = \mathbb{1}$. Any state having zero support on $\mathcal{S}_0$ has an energy at least $\Omega(1/T^3)$ [351], implying that the gap above the zero energy subspace is $\Delta = \Omega(1/T^3)$.

Now, let us add in the term $H_1 = H_{\text{output}}$, with $\left\|H_{\text{output}}\right\| = \epsilon$. We choose $\epsilon < \Delta/16$, unlike the regular clock construction where ϵ is usually taken to be $\Theta(1)$. As long as $\epsilon < \Delta/2$, we can restrict our attention to the zero energy space of H_0, since H_1 can change eigenvalues by at most ϵ. We use the tool of Schrieffer-Wolff transformation as described in Section 6.6 to obtain a description of the Hamiltonian in the low-energy subspace. The

subspace $\mathcal{S}_0$ is the ground-state space of states with energy 0. Since $\|H_1\| = \epsilon$, the associated low-energy subspace of $H = H_0 + H_1$ is

$$\mathcal{S} = \text{span}\{|\Phi\rangle : \langle\Phi| H |\Phi\rangle \in [-\epsilon, \epsilon]\}, \tag{6.31}$$

the subspace with energies in $[-\epsilon, \epsilon]$. In our case $H_0 P_0 = 0$ in the ground subspace spanned by history states $|\phi_h\rangle$, and the matrix elements of $P_0 H_1 P_0$ are given by

$$\langle\phi_h| P_0 H_1 P_0 |\psi_h\rangle = \langle\phi_h| H_1 |\psi_h\rangle \tag{6.32}$$

$$= \frac{1}{T+1} \left(\sum_{i=0}^{T} \langle 0|^m \otimes \langle\phi| \otimes \langle i|_{\text{clock}} U_0^\dagger \dots U_i^\dagger \right) H_1 \times$$

$$\left(\sum_{j=0}^{T} U_j \dots U_0 |0^m\rangle \otimes |\psi\rangle \otimes |j\rangle_{\text{clock}} \right) \tag{6.33}$$

$$= \frac{1}{T+1} \left(\sum_{i=0}^{T} \langle 0|^m \otimes \langle\phi| \otimes \langle i|_{\text{clock}} U_0^\dagger \dots U_i^\dagger \right) \times$$

$$\epsilon |0\rangle\langle 0|_o \otimes |T\rangle\langle T|_{\text{clock}} \left(\sum_{j=0}^{T} U_j \dots U_0 |0^m\rangle \otimes |\psi\rangle \otimes |j\rangle_{\text{clock}} \right) \tag{6.34}$$

$$= \frac{1}{T+1} \langle 0|^m \otimes \langle\phi| \otimes \langle T| U^\dagger \epsilon |0\rangle\langle 0|_o \otimes$$

$$|T\rangle\langle T|_{\text{clock}} U |0^m\rangle \otimes |\psi\rangle |T\rangle_{\text{clock}} \tag{6.35}$$

$$= \frac{\epsilon}{T+1} \langle 0|^m \otimes \langle\phi| U^\dagger |0\rangle\langle 0|_o U |0^m\rangle \otimes |\psi\rangle \tag{6.36}$$

$$= \frac{\epsilon}{T+1} \langle 0|^m \otimes \langle\phi| U^\dagger (\mathbb{1} - \Pi_{\text{out}}) U |0^m\rangle \otimes |\psi\rangle , \tag{6.37}$$

where Π_{out} is the projector onto the accepting state $|1\rangle_o$. Continuing, we have

$$\langle\phi_h| P_0 H_1 P_0 |\psi_h\rangle = \frac{\epsilon}{T+1} (\langle\phi|\psi\rangle - \langle\phi| Q |\psi\rangle), \tag{6.38}$$

meaning that the first order correction $P_0 H_1 P_0$ is simply related to the accept operator Q, which was defined as $Q(U) = \langle 0|^{\otimes m} U^\dagger \Pi_{\text{out}} U |0\rangle^{\otimes m}$. Let the eigenstates of Q be $|\phi_1\rangle, |\phi_2\rangle, \ldots |\phi_{2^w}\rangle$ with eigenvalues $\lambda_1 \geq \lambda_2 \geq \ldots \lambda_{2^w}$. We use the associated history states $|\phi_{i_h}\rangle$ as a basis for the subspace $\mathcal{S}_0$. In this basis, the first order correction $P_0 H_1 P_0$ is diagonal:

$$P_0 H_1 P_0 = \frac{\epsilon}{T+1} \sum_i (1-\lambda_i) |\phi_{i_h}\rangle\langle\phi_{i_h}| . \tag{6.39}$$

We conclude that in the ground space of the original Hamiltonian H_0, the full Hamiltonian H has eigenvalues $\epsilon(1-\lambda_i)/(T+1) \pm O(\epsilon^2/\Delta)$, where the quantity λ_i is the accept probability of the verifier's circuit given $|\phi_i\rangle$ as witness. This is the same conclusion we would obtain by applying degenerate perturbation theory, except that the error bound is rigorous. We now analyze the YES and NO cases to obtain a lower bound on the promise gap. In each case, we also lower bound the spectral gaps in the resulting Hamiltonian.

In the YES case the ground-state energy is $E_1 \leq \epsilon(1-c)/(T+1)$, as can be seen from the fact that the history state $|\phi_h\rangle$ corresponding to an accepting witness $|\phi\rangle$ would have energy $\epsilon(1-\langle\phi| Q |\phi\rangle)/(T+1) \leq \epsilon(1-c)/(T+1)$. Our small-penalty clock construction and the Schrieffer-Wolff transformation comes in handy for the NO case. We see in the NO case that the ground-state energy is at least $E_1 \geq \epsilon(1-s)/(T+1) - O(\epsilon^2/\Delta)$. Therefore, the promise gap is at least $\epsilon(c-s)/(T+1) - O(\epsilon^2/\Delta) = \Omega(1/\text{ exp})$ as long as $\epsilon/\Delta = o((c-s)/(T+1))$.

In the above, if we had chosen $\epsilon = \Theta(1)$ instead of $\epsilon < \Delta/16$, the NO case would have given us a bound $E_1 \geq \Omega(1-s)/T^3$. This would mean that one would have to amplify

the completeness and soundness c, s to near unity in order to get a nontrivial promise gap. However, such an amplification can, in general, shrink the spectral gap of the accept operator. Independently, a large penalty term $\epsilon = \Theta(1)$ could also reorder some eigenvalues, meaning that the spectral properties of the resulting clock Hamiltonian would not faithfully track those of the original accept operator.

The spectral gap in the YES/NO case is $E_2 - E_1 \geq \frac{\epsilon}{T+1}(\lambda_1(Q) - \lambda_2(Q)) - O(\frac{\epsilon^2}{\Delta})$. We take $\epsilon = o(\Delta(c-s)/(T+1)) = o((c-s)/T^4)$, which is exponentially small if $c - s$ is. As long as $\epsilon/\Delta = o(\min[g_1, g_2]/(T+1))$, both the YES and NO cases will have an exponentially small spectral gap. In summary the choice

$$\epsilon = \frac{\min\left[g_1, g_2, (c-s)\right]}{nT^4} = \Theta(1/\exp) \tag{6.40}$$

suffices to have a promise gap and spectral gaps bounded below by $\Omega(1/\exp)$. This proves PreciseEGQMA-hardness of $(1/\exp, 1/\exp)$-LocalHamiltonian and one half of Theorem 23. □

We generalize the above proof technique to the case of GQCMA-hardness of GS-Description-LocalHamiltonian. In addition to showing a promise gap and a spectral gap, we should show that the resulting Hamiltonian has a classical description of a circuit to prepare a low-energy state. We show the following general lemma.

Lemma 47. *(δ, Δ)-GS-Description-LocalHamiltonian is $\mathsf{GQCMA}[c, s, g_1, g_2]$-hard for any δ, Δ satisfying both the following conditions.*

i. $\delta = O((c-s)^2/\mathsf{poly}(n))$ *for some polynomial.*

ii. If $c - s = o(\min[g_1, g_2])$*, then any* Δ *satisfying* $\Delta = O((c-s)\min[g_1, g_2]/\mathsf{poly}(n))$*. Else,* $\Delta = 0$*.*

Proof. To prove GQCMA-hardness, we give a reduction from $\mathsf{GQCMA}[c, s, g_1, g_2]$ to GS-DESCRIPTION-LOCALHAMILTONIAN$[a, b, g_1', g_2']$. We are promised that the input witnesses are computational basis states (this can be assumed without loss of generality), corresponding to the classical witness sent by the prover. We would like to show that there exists a circuit V to prepare a state δ-close in energy to the ground state of the clock Hamiltonian both the YES and NO cases.

Consider again the small-penalty clock construction, with the clock Hamiltonian Eq. (6.26). Let the norm of the penalty term be $\left\|H_{\text{output}}\right\| = \epsilon$. When $\epsilon = 0$, the ground-state space is given by valid history state computations corresponding to computational basis witness states. The spectral gap above this subspace is at least $\Omega(1/T^3)$. The addition of the penalty term changes the energies to $\frac{\epsilon}{T+1}(1-\lambda_k) + O(\epsilon^2 T^3)$, where λ_k is the accept probability upon input computational basis state $|y_k\rangle$ as witness. Consider the history state associated with witness $|y_k\rangle$:

$$|y_{k_h}\rangle := \frac{1}{\sqrt{T+1}} \sum_{i=0}^{T} U_i \ldots U_0 \left|0^m\right\rangle \otimes |y_k\rangle \otimes |i\rangle_{\text{clock}} . \tag{6.41}$$

This state has energy $\langle y_{k_h}| H |y_{k_h}\rangle = \frac{\epsilon}{T+1}(1-\lambda_k)$ and is therefore $O(\epsilon^2 T^3)$-close in energy to the true ground state. Therefore, as long as $\epsilon^2 T^3 < O\left(\frac{(b-a)^3}{f(n)^2}\right)$, a classical description of a circuit that prepares $|y_{k_h}\rangle$ is a valid ground-state description. The circuit may be described by specifying y_k and a circuit that prepares the history state $|\phi_h\rangle$ upon any quantum input $|\phi\rangle$. This latter circuit first prepares the state $\frac{1}{\sqrt{T+1}} \sum_{i=0}^{T} \left|0^m\right\rangle |i\rangle_{\text{clock}}$ and then applies the

unitaries $U_j \dots U_0$ controlled on the clock being in time-step j [374].

The same promise gap and spectral gap analyses as in the proof of Lemma 46 hold. In the YES case, the Hamiltonian has ground-state energy $\leq \frac{\epsilon}{T+1}(1-\lambda_1) \leq \frac{\epsilon}{T+1}(1-c)$. In the NO case, the ground-state energy is at least $\frac{\epsilon}{T+1}(1-\lambda_1) - O(\epsilon^2 T^3) \geq \frac{\epsilon}{T+1}(1-s) - O(\epsilon^2 T^3)$. The promise gap between the ground-state energy for YES and NO cases is $\delta \geq \frac{\epsilon}{T+1}(c-s) - O(\epsilon^2 T^3)$. We make the choice $\epsilon = \Theta(\frac{c-s}{T^4})$ to ensure the promise gap is $\Omega((c-s)^2/T^5)$. This choice is consistent with the choice $\epsilon^2 T^3 \leq O(\frac{(b-a)^3}{f(n)^2})$ made above.

Let us now analyze the spectral gap of the resulting Hamiltonian. Using the Schrieffer-Wolff expansion to obtain the eigenvalues of the Hamiltonian for small ϵ, we have $\Delta \geq \frac{\epsilon}{T+1}(\lambda_1 - \lambda_2) - O(\epsilon^2 T^3)$. The spectral gap is at least $\frac{\epsilon}{T+1}\min[g_1, g_2]$ as long as $\epsilon^2 T^3 = o(\frac{\epsilon}{T+1}\min[g_1, g_2])$. Using the choice of ϵ above, this means the spectral gap is $\Omega((c-s)/T^5 \min[g_1, g_2])$ as long as $c - s = o(\min[g_1, g_2])$. Otherwise, the best bound on the spectral gap is $\Delta \geq 0$. Observing that $T = \mathsf{poly}(n)$ by assumption, we obtain the lemma.

□

The lemma allows us to show the following:

Corollary 48 (Second half of Theorems 24 to 27)**.** $(1/\mathsf{poly}, 0)$-*GS-Description-Local Hamiltonian is* QCMA*-hard.*

$(1/\exp, 0)$-*GS-Description-Local Hamiltonian is* PreciseQCMA*-hard.*

$(1/\mathsf{poly}, 1/\mathsf{poly})$-*GS-Description-Local Hamiltonian is* PGQCMA*-hard.*

$(1/\exp, 1/\exp)$-*GS-Description-Local Hamiltonian is* PreciseEGQCMA*-hard.*

For the problem with $\delta = 1/\exp, \Delta = 1/\mathsf{poly}$, we do not give a direct reduction from a PrecisePGQCMA instance. Instead, we show PP-hardness through the characterization of

PP in terms of the class PreciseBQP. From the PP upper bound to PrecisePGQCMA, we obtain PrecisePGQCMA-completeness of the problem $(1/\exp, 1/\text{poly})$-GS-DESCRIPTION-LOCALHAMILTONIAN. The argument is similar for PrecisePGQMA-hardness of $(1/\exp, 1/\text{poly})$-LOCALHAMILTONIAN.

Lemma 49 (Lemmas 34 and 36 restated)**.**

$(1/\exp, 1/\text{poly})$-*GS-DESCRIPTION-LOCALHAMILTONIAN is* PP-*hard.*

$(1/\exp, 1/\text{poly})$-*LOCALHAMILTONIAN is* PP-*hard.*

Proof. We give a reduction from any problem in PreciseBQP to $(1/\exp, 1/\text{poly})$-GS-DESCRIPTION-LOCALHAMILTONIAN, which is also an instance of $(1/\exp, 1/\text{poly})$-LOCALHAMILTONIAN. Since PreciseBQP is the class of problems that can be decided by quantum circuits with a promise gap $c - s = \Omega(1/\exp)$, it can also be thought of as "PreciseQMA without an input witness". The Hamiltonian is constructed out of the PreciseBQP computation as $H = H_{\text{input}} + H_{\text{prop}} + H_{\text{output}} + H_{\text{clock}}$, where the terms are now

$$H_{\text{input}} = \sum_{i=1}^{m} |0\rangle\langle 0|_i \otimes H_{\text{clockinit}}, \tag{6.42}$$

$$H_{\text{prop}} = \sum_{i=0}^{T} -U_{i+1} \otimes |i+1\rangle\langle i|_{\text{clock}} - U_{i+1}^{\dagger} \otimes |i\rangle\langle i+1|_{\text{clock}}$$

$$+ \mathbb{1} \otimes (|i\rangle\langle i|_{\text{clock}} + |i+1\rangle\langle i+1|_{\text{clock}}), \text{ and} \tag{6.43}$$

$$H_{\text{output}} = \epsilon\, |0\rangle\langle 0|_o \otimes |T\rangle\langle T|_{\text{clock}}\,. \tag{6.44}$$

The only difference from Eqs. (6.27) to (6.29) is that H_{input} does not have support on an unpenalized proof register, since PreciseBQP does not rely on a proof state given as input. This is analogous to the clock construction of Ref. [351], which was instrumental in the

proof that adiabatic quantum computation is universal for BQP.

We again let the Hamiltonian H_0 be $H_{\text{input}} + H_{\text{prop}} + H_{\text{clock}}$ and $H_1 = H_{\text{output}}$. The ground state of H_0 is now nondegenerate (unique) and given by the history state

$$|0_h\rangle := \frac{1}{\sqrt{T+1}} \sum_{i=0}^{T} U_i \dots U_0 |0^m\rangle \otimes |i\rangle_{\text{clock}} . \tag{6.45}$$

Let us denote the ground-state space of H_0 and the projector onto it by Π_0. As for H_1, the ground space Π_1 is spanned by states belonging to subspaces $\mathcal{L}$ and $\mathcal{L}'$, with

$$\mathcal{L} = |1\rangle_o \otimes |T\rangle_{\text{clock}} \tag{6.46}$$

$$\mathcal{L}' = \text{span}\ \{|\psi\rangle\} \otimes \text{span}\ \{|0\rangle_{\text{clock}}, |1\rangle_{\text{clock}}, \dots |T-1\rangle_{\text{clock}}\}, \tag{6.47}$$

with $|\psi\rangle$ arbitrary.

We observe that when $\epsilon = 0$, the Hamiltonian exactly corresponds to Aharonov et al.'s H_{final} [351]. Aharonov et al. [351] showed that this Hamiltonian H_0 has a spectral gap of $\Delta = \Omega(1/T^3)$ in the full Hilbert space. Further, the ground state of H_0 corresponds to the history state of the BQP computation (PreciseBQP in this case), which starts off in a fixed, known state $|0^m\rangle$.

In the YES case, the ground-state energy of $H = H_0 + H_1$ can be bounded above by $\frac{\epsilon}{T+1}(1-c)$. For the NO case, we again use the expression for the perturbed energies in the ground-state space coming from the Schrieffer-Wolff transformation. Specifically, in the NO case, we have $E_1 \geq \epsilon(1-s)/(T+1) - O(\epsilon^2/\Delta)$, where Δ is the spectral gap above

the ground state, just as in the proof of Lemma 46. The promise gap is lower-bounded by

$$\epsilon\frac{1-s}{T+1} - \epsilon\frac{1-c}{T+1} - O\left(\frac{\epsilon^2}{\Delta}\right). \tag{6.48}$$

Therefore, as long as $\epsilon/\Delta = o((c-s)/(T+1))$ and $\epsilon = \Omega(2^{-\mathsf{poly}})$, the promise gap is at least $\Omega(\epsilon(c-s)/(T+1)) = \Omega(2^{-\mathsf{poly}})$. The spectral gap for the unperturbed Hamiltonian H_0, which is the same as the final Hamiltonian in Ref. [308], is at least $\Omega(1/T^3)$. Therefore, we pick $\epsilon = (c-s)/(nT^4)$, which ensures that the conditions above are satisfied.

Coming to the spectral gap of the full Hamiltonian, we observe that since the original Hamiltonian had a spectral gap of $\Omega(1/T^3)$ and the perturbation H_1 is exponentially small, the eigenvalues can change at most by $\|H_1\| = \epsilon$, preserving the spectral gap. So far, we have a reduction from any PreciseBQP instance to an instance of $(1/\exp, 1/\mathsf{poly})$-LocalHamiltonian.

It remains for us to see that there is an efficient circuit that can prepare a state close in energy to the ground state. By the justification in the proof of Lemma 47, we know that choosing the output penalty term to be exponentially small causes the history state of the computation $|0_h\rangle$ to be exponentially close to the ground state in energy. We have also seen the existence of a polynomial size circuit that prepares the history state given a description of the input (which here is $|0^m\rangle$ for PreciseBQP). Note that when $\epsilon = 0$, the ground state is unique and has a $\Omega(1/\mathsf{poly})$ spectral gap above and therefore taking ϵ exponentially small does not pose a problem with spectral gaps. □

The difference between the proof of Lemma 49 and the proof of Lemma 46 is that it is the perturbation ϵ that creates the spectral gap in the proof of Lemma 49, while in the

proof of Lemma 46, the spectral gap already exists in the unperturbed Hamiltonian. This is why we can afford to take ϵ exponentially small here, which is needed to obtain an instance with a promise gap.

Thus, we have seen PP-hardness of $(1/\exp, 1/\mathsf{poly})$-LocalHamiltonian. PrecisePGQMA-hardness of the problem follows from the fact that PrecisePGQMA $\subseteq$ PP (Lemma 32).

Corollary 50. $(1/\exp, 1/\mathsf{poly})$-*LocalHamiltonian is* PrecisePGQMA-*hard.*

Similarly, the PP-hardness of $(1/\exp, 1/\mathsf{poly})$-GS-Description-LocalHamiltonian and the result PrecisePGQCMA $\subseteq$ PrecisePGQMA $=$ PP together imply the following result.

Corollary 51. $(1/\exp, 1/\mathsf{poly})$-*GS-Description-LocalHamiltonian is* PrecisePGQCMA-*hard.*

Lastly, the remaining case is $(1/\mathsf{poly}, 1/\exp)$-GS-Description-LocalHamiltonian with $\delta = 1/\mathsf{poly}$, $\Delta = 1/\exp$, for which we argue that an instance with spectral gap $\Delta = \Omega(1/\mathsf{poly})$ is also an instance with $\Delta = \Omega(1/\exp)$. Therefore, $(1/\mathsf{poly}, 1/\exp)$-GS-Description-LocalHamiltonian is PGQCMA-hard, and, since PGQCMA $=_R$ EGQCMA, EGQCMA-hard under randomized reductions. For the case of $(1/\mathsf{poly}, 1/\exp)$-LocalHamiltonian, we do not currently have a hardness result. This is because, in performing a reduction from EGQMA, we get an instance of $(1/\mathsf{poly}, 0)$-LocalHamiltonian and do not get any promise on the spectral gap that results.

6.8 Precise phase estimation of gapped Hamiltonians

In this section, we show that the $(1/\exp, \Delta)$-LocalHamiltonian problems with either $1/\mathsf{poly}$ or $1/\exp$ spectral gaps defined in Section 6.2.1 are in the corresponding PreciseGQMA class. Together with the results of the previous section, this proves Theorems 22 and 23.

Lemma 52. $(1/\exp, 1/\mathsf{poly})$-*LocalHamiltonian* $\in$ PrecisePGQMA.

Lemma 53. $(1/\exp, 1/\exp)$-*LocalHamiltonian* $\in$ PreciseEGQMA.

The proof relies on phase estimation to infer energies of a local Hamiltonian. The standard phase estimation circuit requires $\exp(n)$ many gates in order to infer the eigenvalues to $1/\exp$ precision. However, since we want to show containment in a Precise- class, we can use the power of being able to distinguish between two cases with exponentially close accept probabilities. It turns out that phase estimation with a single ancillary qubit is enough to distinguish between the YES and NO cases, as shown in Ref. [319]. Moreover, we show that the circuit preserves spectral gaps of the Hamiltonian: if two eigenstates have energies separated by some amount, then the phase estimation circuit also has a gap in the accept probabilities corresponding to these input states.

Below we give a unified proof of Lemmas 40, 52 and 53. Specifically, we show PrecisePGQMA (PreciseEGQMA) containment of the problem $(1/\exp, \Delta)$-GappedSparseHamiltonian with $\Delta = 1/\mathsf{poly}$ ($\Delta = 1/\exp$).

Lemma 54. *GappedSparseHamiltonian*$[a, b, g_1, g_2]$ *has a* GQMA$[c, s, g_1', g_2']$ *protocol with spectral gaps* $g_1' = \Omega(g_1^2/\mathsf{poly})$ *and* $g_2' = \Omega(g_2^2/\mathsf{poly})$ *and promise gap* $c - s = (b-a)^2/\mathsf{poly}$.

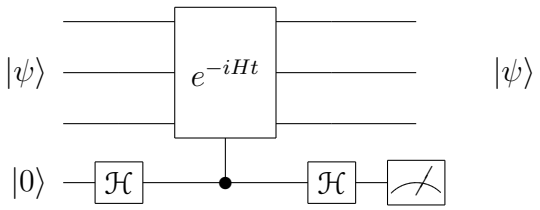

Figure 6.2: One-qubit phase-estimation circuit. The symbol $\mathcal{H}$ on the bottom line denotes the Hadamard gate.

Proof. The strategy is to ask the prover for the ground state of the sparse Hamiltonian. The verifier then performs phase estimation on the witness state with a single ancillary qubit and uses the power to decide between two cases with exponentially close accept probabilities. This power effectively enables computation of the phase of e^{-iHt} to exponential precision, despite having a single ancilla qubit in the phase estimation circuit (see Ref. [319] for more details). The phase estimation circuit is shown in Fig. 6.2. If $t \leq \frac{\pi}{2\|H\|}$, all eigenstates of H would correspond to a unique phase and a unique accept probability for the circuit. We know an upper bound dk on $\|H\|$ through the Gershgorin circle theorem because we are assured that the magnitude of the entries is $\leq k$ and the sparsity is d. Therefore, it suffices to choose $t \leq \frac{\pi}{2dk}$.

In order to perform phase estimation to exponentially small error, we need to apply a controlled-e^{-iHt} rotation to error $\epsilon = 1/\exp$. This is possible due to Hamiltonian simulation algorithms for sparse Hamiltonians, whose circuit size scales as $\mathsf{poly}(n)\log\left(\frac{1}{\epsilon}\right)$ [40], which is polynomial in n, as desired. The accept probability of the circuit upon input an eigenstate $|E_i\rangle$ of the Hamiltonian is $\frac{1+\cos(E_i t)}{2}$. The promise gap can be lower bounded by an inverse exponential, as has been analyzed previously [319].

We can also show a spectral gap in the accept operator, or equivalently, a gap in the accept probabilities of the circuit for the optimal state and any state orthogonal to it. Since

the phase estimation circuit does not apply the exact controlled-e^{-iHt} unitary but a unitary U_x exponentially close to it, the eigenstates of $Q = \langle 0| \Pi_{\text{in}} U_x^\dagger \Pi_{\text{out}} U_x \Pi_{\text{in}} |0\rangle$ are not exactly the eigenstates of e^{-iHt} (or of H). However, since $\left\| e^{-iHt} - U_x \right\| \leq \epsilon$, the eigenvalues of Q are exponentially close to the accept probabilities of the eigenstates $|E_i\rangle$ of H. The difference in accept probabilities can be bounded by ϵ.

The difference in the ideal accept probabilities of the ground state and the first excited state is $\frac{\cos(E_0 t) - \cos(E_1 t)}{2}$. Applying Taylor's theorem to $\cos(E_1 t)$ around the point $E_0 t$, we get

$$\begin{aligned} \cos(E_1 t) = \cos(E_0 t) - \sin(E_0 t) t(E_1 - E_0) - \\ \cos(E_0 t) \frac{t^2 (E_1 - E_0)^2}{2} + \sin(E_0 t) \frac{h^3}{6} \end{aligned} \tag{6.49}$$

for some $h \in [0, (E_1 - E_0)t]$. Therefore,

$$\begin{aligned} \cos(E_0 t) - \cos(E_1 t) &= \sin(E_0 t) t(E_1 - E_0) + \\ &\quad \cos(E_0 t) \frac{t^2 (E_1 - E_0)^2}{2} - \sin(E_0 t) \frac{h^3}{6} && (6.50) \\ &\geq t^2 (E_1 - E_0)^2 / 2 - \frac{t^3 (E_1 - E_0)^3}{6}. && (6.51) \\ &\geq \Omega(t^2 (E_1 - E_0)^2), && (6.52) \end{aligned}$$

where in the second line we use the fact that $E_0 t, E_1 t < \pi/2$ and $(E_1 - E_0)^3 t^3 = O(t(E_1 - E_0))$ in the third. Therefore, the ideal accept probabilities also have a gap of $\Omega((E_1 - E_0)^2 / \|H\|^2) = \Omega(\Delta^2 / \text{poly})$ as long as $\epsilon \leq O(t^2 \Delta^2 / n) = O(\Delta^2 / \text{poly})$. Now, when the applied unitary differs from the ideal one by ϵ in operator norm distance, the gap in the accept probabilities differs from the ideal accept probabilities by 2ϵ. We therefore choose ϵ

sufficiently small, i.e. we choose, say, $\epsilon = \Theta(t^2(E_1 - E_0)^2/2^n)$, which is still $\Omega(2^{-\mathsf{poly}})$, as needed.

To see the existence of a promise gap, notice that $E_0 \leq a$ in the YES case and $E_1 \geq b$ in the NO case, giving $c - s = \Omega(t^2(b-a)^2 - 2\epsilon) = \Omega((b-a)^2/\mathsf{poly})$. This proves the lemma. □

As corollaries, we obtain Lemmas 40, 52 and 53, since a local Hamiltonian is also a sparse Hamiltonian.

6.9 Phase estimation in the presence of efficient circuit descriptions

In this section, we show the problem (δ, Δ)-GS-DESCRIPTION-LOCALHAMILTONIAN is in GQCMA with appropriate bounds on the promise and spectral gaps (Theorems 24 to 28).

We first deal with the case of zero spectral gap:

Lemma 55 (One half of Theorems 24 and 25)**.** $(1/\mathsf{poly}, 0)$-*GS-DESCRIPTION-LOCALHAMILTONIAN* $\in$ QCMA.

$(1/\exp, 0)$-*GS-DESCRIPTION-LOCALHAMILTONIAN* $\in$ PreciseQCMA.

Proof. For the upper bound, we describe a QCMA or PreciseQCMA protocol. We are promised that in both the YES and NO cases, there exists a classical description of a circuit V of polynomial size that will create a state with energy close to the ground-state energy. Specifically, the energy of this state is ϵ-close to the ground-state energy, for $\epsilon < \frac{(b-a)^3}{f(n)^2}$ for a polynomial $f(n) \geq \|H\|$. For QCMA, we have $b-a \geq \Omega(1/\mathsf{poly})$, while for PreciseQCMA, $b - a \geq \Omega(1/\exp)$. The verifier asks the prover to give this description (which is promised to exist). The verifier then creates a state $|\psi\rangle$ with low energy by applying V to $|0^m\rangle$. The

verifier measures the energy of this state via the one-bit phase-estimation protocol outlined in Section 6.8, which involves applying a controlled-e^{-iHt} for time $t \leq \frac{\pi}{2\|H\|}$.

The proof that this verification protocol works is slightly more involved than the $\mathsf{QMA}[c,s]$ case. This is because, in the case of QMA a verifier can assume without loss of generality that the prover sends the optimal eigenstate as a witness. However, in the case of GS-Description-LocalHamiltonian, we are only promised the existence of an efficient circuit to prepare a state close in energy to the ground state, and not the ground state itself[7]. Despite this complication, we can still show that a state close in energy to the ground state behaves similarly with respect to the accept probabilities of the $\mathsf{QCMA}[c,s]$ verifier.

In the YES case, there is a description V that produces a state $|\psi\rangle$ with energy close to the ground-state energy (i.e. with energy $\leq E_1 + \epsilon < a + \frac{(b-a)^3}{\mathsf{poly}(n)}$). We show in Lemma 56 that the accept probability of the verifier upon performing one-bit phase estimation on the state $|\psi\rangle$ is at least $\cos^2\left(\frac{bt}{2}\right) + \Omega((b-a)^2/\mathsf{poly})$. In the NO case, the optimal strategy for the prover is to send the description of a circuit that makes a state as close as possible to the ground state, since the accept probabilities are monotonic in energy and there exists no other state with smaller energy, by definition. Even if the prover sends the verifier a circuit that exactly prepares the ground state $|E_1\rangle$, its energy in the NO case is already $\geq b$. This means that the verifier will accept with probability at most $(1+\cos E_1 t)/2 \leq (1+\cos bt)/2$. Therefore there is a separation in the accept probabilities in the YES and NO cases of $c - s = \Omega((b-a)^2/\mathsf{poly})$, which is $\Omega(1/\mathsf{poly})$ for $b - a = \Omega(1/\mathsf{poly})$ and $\Omega(1/\exp)$ for $b - a = \Omega(1/\exp)$. □

Lemma 56. *If a state $|\psi\rangle$ has energy $\langle\psi|H|\psi\rangle = \langle E\rangle \leq E_1 + \frac{5(b-a)^3}{24f(n)^2}$ for some polynomial*

[7]The weaker promise is more natural since it is more robust.

$f(n) \geq \|H\|$, then in the YES case, the accept probability of the state upon phase estimation with one bit of precision is $\langle p \rangle \geq \cos^2\left(\frac{bt}{2}\right) + \delta$, where $\delta = \Omega\left(\frac{5(b-a)^2}{24f(n)^2}\right)$.

Proof. We are given a state $|\psi\rangle$ with energy $\langle E \rangle$. Let $p_j = |\langle E_j|\psi\rangle|^2$ be the weight of the energy eigenstate E_j. Then we know $p_1 E_1 + p_2 E_2 + \ldots p_{2^n} E_{2^n} = \langle E \rangle$. The probability of accepting $|\psi\rangle$ in the one-bit phase estimation circuit is given by $\langle p \rangle = p_1 \cos^2\left(\frac{E_1 t}{2}\right) + p_2 \cos^2\left(\frac{E_2 t}{2}\right) + \ldots p_{2^n} \cos^2\left(\frac{\|H\|t}{2}\right)$, where $\|H\| = E_{2^n}$. Given the constraint on the energy $\langle E \rangle$, we show in Lemma 57 that $\langle p \rangle \geq \cos^2\left(\frac{E_1 t}{2}\right)(1-x) + \cos^2\left(\frac{\|H\|t}{2}\right) x$, where $x := \frac{\langle E \rangle - E_1}{\|H\| - E_1}$. Now in order to have $\langle p \rangle \geq \cos^2\left(\frac{bt}{2}\right) + \delta$, it suffices to have

$$x \leq \frac{\cos^2\left(\frac{E_1 t}{2}\right) - \cos^2\left(\frac{bt}{2}\right) - \delta}{\cos^2\left(\frac{E_1 t}{2}\right) - \cos^2\left(\frac{\|H\|t}{2}\right)} \tag{6.53}$$

$$= \frac{\cos\left(E_1 t\right) - \cos\left(bt\right) - 2\delta}{\cos\left(E_1 t\right) - \cos\left(\|H\|t\right)}. \tag{6.54}$$

It is therefore sufficient if

$$x \leq \frac{(b-a)t}{2}\left(bt - \frac{b^3 t^3}{6}\right) - \delta, \text{since} \tag{6.55}$$

$$\frac{(b-a)t}{2}\left(bt - \frac{b^3 t^3}{6}\right) - \delta \leq \frac{(b-a)t\sin(bt) - 2\delta}{2} \tag{6.56}$$

$$\leq \frac{(b-a)t\sin(bt) - 2\delta}{\cos(at) - \cos(\|H\|t)} \tag{6.57}$$

$$\leq \frac{\cos(at) - \cos(bt) - 2\delta}{\cos(at) - \cos(\|H\|t)}$$

$$\leq \frac{\cos(at) - \cos(bt) - 2\delta}{\cos(E_1 t) - \cos(\|H\|t)}, \tag{6.58}$$

where we use the inequalities $\sin(bt) \geq bt - \frac{b^3 t^3}{6}$, $E_1 \leq a$, $\cos(at) - \cos(bt) \geq (b-a)t\sin(bt)$, and $2 \leq \cos(at) - \cos(\|H\|t)$. We now require $\delta \geq (b-a)t\sin(bt)/4$, so that

the condition Eq. (6.55) translates to $x \leq \frac{(b-a)t}{4}\left(bt - \frac{b^3t^3}{6}\right)$.

Let us choose $t = \min[1/f(n), 1/b] = 1/f(n)$, since otherwise $b \geq f(n) \geq \|H\|$ and the instance is trivial. We thus know $\|H\|t \leq 1 \leq \pi/2$, $t \geq 1/f(n)$, and $bt < 1$. We also assume that in the YES case, $\|H\| - E_1 \geq b - a$. This is because otherwise a verifier can compute $\frac{\mathrm{Tr}(H)}{2^n}$ efficiently given the Hamiltonian and accept straightaway if $\frac{\mathrm{Tr}(H)}{2^n} \leq b$. This works since $E_1 \leq \frac{\mathrm{Tr}(H)}{2^n}$, and by the promise, $E_1 \leq b \implies E_1 \leq a$. Therefore, without loss of generality, one can assume that the nontrivial instances satisfy $b \leq \frac{\mathrm{Tr}(H)}{2^n} \leq \|H\|$, or $\|H\| - E_1 \geq b - a$.

Therefore, since $\langle E\rangle \leq E_1 + \frac{5(b-a)^3}{24f(n)^2}$, we have

$$\langle E\rangle \leq E_1 + \frac{5(b-a)^2(\|H\| - E_1)}{24f(n)^2} \tag{6.59}$$

$$\implies x \leq \frac{(b-a)^2}{4f(n)^2}\frac{5}{6} = \frac{(b-a)^2}{4f(n)^2}\left(1 - \frac{1}{6}\right) \tag{6.60}$$

$$\leq \frac{(b-a)^2}{4f(n)^2}\left(1 - \frac{b^2t^2}{6}\right) \tag{6.61}$$

$$\leq \frac{(b-a)}{4f(n)}\frac{b}{f(n)}\left(1 - \frac{b^2t^2}{6}\right) \tag{6.62}$$

$$\leq \frac{(b-a)t}{4}\left(bt - \frac{b^3t^3}{6}\right), \tag{6.63}$$

as required. To sum up, we have shown that $\langle E\rangle \leq E_1 + \frac{5(b-a)^3}{24f(n)^2}$ implies $\delta \geq (b-a)t\sin(bt)/4 \geq \frac{(b-a)^2(1-b^2t^2/6)}{4f(n)^2} \geq \frac{5(b-a)^2}{24f(n)^2}$. □

Lemma 57. *For probabilities $p_j : j \in [2^n]$ satisfying $\sum_j p_j E_j \leq \langle E\rangle$ and numbers $E_1 \leq E_2 \leq \ldots E_{2^n}$ satisfying $E_j t \in [0, \pi/2]$, the quantity $\sum_j p_j \cos^2\left(\frac{E_j t}{2}\right)$ is bounded below by $\cos^2\left(\frac{E_1 t}{2}\right)(1-x) + \cos^2\left(\frac{E_{2^n} t}{2}\right)x$, where x is given by $\frac{\langle E\rangle - E_1}{E_{2^n} - E_1}$.*

Proof. Since the function $f(x) = -\cos^2(xt/2)$ is convex for $xt/2 \in [0, \pi/2)$, we have

$$\frac{f(E_1)(E_{2^n} - E_j) + f(E_{2^n})(E_j - E_1)}{E_{2^n} - E_1} \geq f(E_j). \tag{6.64}$$

Therefore,

$$p_j \frac{f(E_1)(E_{2^n} - E_j) + f(E_{2^n})(E_j - E_1)}{E_{2^n} - E_1} \geq p_j f(E_j) \tag{6.65}$$

$$\implies \frac{f(E_1)(E_{2^n} - \langle E \rangle) + f(E_{2^n})(\langle E \rangle - E_1)}{E_{2^n} - E_1} \geq \sum_j p_j f(E_j) \tag{6.66}$$

$$\implies \sum_j p_j \cos^2\left(\frac{E_j t}{2}\right) \geq \cos^2\left(\frac{E_1 t}{2}\right) \frac{E_{2^n} - \langle E \rangle}{E_{2^n} - E_1} + \cos^2\left(\frac{E_{2^n} t}{2}\right) \frac{\langle E \rangle - E_1}{E_{2^n} - E_1}, \tag{6.67}$$

which completes the proof.

□

We now turn to the cases where in addition to the promise of an efficient circuit to prepare a low-energy state, the Hamiltonian is promised to have a spectral gap Δ. For this case, we can show the following:

Lemma 58. *GS-DESCRIPTION-LOCALHAMILTONIAN*$[a, b, g_1, g_2] \in \mathsf{GQCMA}[c, s, g_1', g_2']$ *for* $c - s = \Omega\left(\frac{(b-a)^2}{f(n)^2}\right)$ *and* $\min[g_1', g_2'] \geq \frac{5\Delta^2}{36 f(n)}$*, where* $f(n)$ *is a polynomial upper bound to* $\|H\|$*, and* $\Delta = \min[g_1, g_2] \geq (b-a)^3/f(n)^2$.

Proof. We analyze the same algorithm as the non-gapped case and show that the verification protocol, with slight modifications, preserves the spectral gap. In particular, in the first step of the original protocol, the verifier straightaway accepts if $\mathrm{Tr}(H)/2^n \leq b$ or if the

upper bound to the norm of the Hamiltonian, $f(n)$ satisfies $f(n) \leq b$. We modify this to requiring the verifier to accept only if, in addition to the previous conditions, measurement of the witness register yields the all zeroes string 0^w (where w is the size of the witness register). This has the effect of creating a spectral gap, since in this case only the all-zeroes state is accepted and all other computational-basis states are rejected.

If the first step does not cause the verifier to accept, the verifier assumes that the witness state is a description of the circuit V to prepare a low-energy state $|\psi\rangle$. The verifier then proceeds to prepare this state and measure its energy using the one-bit phase estimation protocol. As shown in the proof of Lemma 55, the protocol has a promise gap $c - s = \Omega\left(\frac{(b-a)^2}{f(n)^2}\right)$.

We now analyze the spectral gap. Let us denote by y the quantity $\frac{\langle E\rangle - E_1}{E_2 - E_1}$ and by x the quantity $\frac{\langle E\rangle - E_1}{E_{2^n} - E_1} \leq y$. Any state with energy $\langle E\rangle := \langle\psi|H|\psi\rangle \leq E_1 + \frac{(b-a)^3}{f(n)^2} \leq E_1 + \Delta$ has a large overlap with the ground state:

$$|\langle\psi|E_1\rangle|^2 \geq 1 - \frac{\langle E\rangle - E_1}{E_2 - E_1} = 1 - y. \tag{6.68}$$

Therefore, any state $|\phi\rangle$ orthogonal to $|\psi\rangle$ must have an overlap with the ground state that satisfies $|\langle\phi|E_1\rangle|^2 \leq y$. This means that the accept probability for any witness orthogonal to the one corresponding to the ground-state description is

$$\langle p_\phi\rangle = \sum_j p_j \cos^2\left(\frac{E_j t}{2}\right) \tag{6.69}$$

$$\leq y\cos^2\left(\frac{E_1 t}{2}\right) + (1-y)\cos^2\left(\frac{E_2 t}{2}\right). \tag{6.70}$$

On the other hand, the accept probability of the optimal witness is at least (Lemma 56)

$$\langle p_\psi \rangle \geq (1-x)\cos^2\left(\frac{E_1 t}{2}\right) + x\cos^2\left(\frac{E_2 t}{2}\right). \tag{6.71}$$

The difference in these two is a lower bound for the spectral gap of the accept operator:

$$\begin{aligned} g_1, g_2 \geq \langle p_\psi \rangle - \langle p_\phi \rangle &\geq \cos^2\left(\frac{E_1 t}{2}\right)(1-x-y) + \\ &\cos^2\left(\frac{E_2 t}{2}\right)(x+y-1) && (6.72)\\ &= \frac{(1-x-y)}{2}(\cos(E_1 t) - \cos(E_2 t)) && (6.73)\\ &\geq \frac{(1-2y)}{2}(E_2 - E_1)\sin(E_2 t). && (6.74) \end{aligned}$$

Now, we know from the promise that $y = \frac{\langle E \rangle - E_1}{E_2 - E_1} \leq \frac{(b-a)^3}{f(n)^2 \Delta} \leq \frac{1}{3}$, and $E_2 \geq E_1 + \Delta \geq \Delta$. Also, we have chosen $t \geq 1/f(n)$ for a polynomial $f(n) \geq \|H\|$. Therefore,

$$\begin{aligned} \min[g_1', g_2'] &\geq \frac{\Delta}{6}\sin(\Delta t) && (6.75)\\ &\geq \frac{\Delta}{6}\left(\Delta t - \frac{\Delta^3 t^3}{6}\right) && (6.76)\\ &\geq \frac{\Delta}{6}\left(\frac{\Delta}{f(n)} - \frac{\Delta^3}{6f(n)^3}\right) && (6.77)\\ &= \frac{\Delta^2}{6f(n)}\left(1 - \frac{\Delta^2}{6f(n)^2}\right) && (6.78)\\ &\geq \frac{5\Delta^2}{36 f(n)}, && (6.79) \end{aligned}$$

since $\Delta \leq \|H\| \leq f(n)$. □

This proves the following results:

Corollary 59 (One half of Theorems 26 to 28)**.**

$(1/\mathsf{poly}, 1/\mathsf{poly})$-*GS-Description-Local Hamiltonian* $\in \mathsf{PGQCMA}$.

$(1/\exp, 1/\exp)$-*GS-Description-Local Hamiltonian* $\in \mathsf{PreciseEGQCMA}$.

$(1/\exp, 1/\mathsf{poly})$-*GS-Description-Local Hamiltonian* $\in \mathsf{PrecisePGQCMA}$.

6.10 Details of PP algorithm

In this section we complete the proof of Lemma 32 by expanding upon the PP algorithm. We also prove Lemma 33 by giving a P^{PP} algorithm to precisely compute ground-state local observables of $\Omega(1/\mathsf{poly})$-spectral-gapped Hamiltonians.

Lemma 60. *A* PP *algorithm can decide whether* $\mathrm{Tr}[Q^q A] \leq a'$ *or* $\geq b'$ *when input thresholds* a' *and* b'*, for matrices* Q *and* A *of size* $2^{\mathsf{poly}(n)} \times 2^{\mathsf{poly}(n)}$ *satisfying the following properties (we use the symbol* R *to denote both matrices* Q *and* A *in the following):*

1. *The norm of the matrix* R *is upper bounded by a polynomial in* n*.*

2. *The matrix* R *may be written as a polynomial of degree* $d = \mathsf{poly}(n)$ *in terms of matrices* $R_i, i \in [m]$ *in the computational basis for* $m = \mathsf{poly}(n)$*, such that:*

   (a) *The matrix elements of each matrix* R_i *are computable to precision* δ *in time polynomial in* n *and* $\log(1/\delta)$*.*

Proof. The quantity $\mathrm{Tr}(Q^q A)$ may be expressed as

$$\sum_x \langle x|Q^q A|x\rangle = \sum_x \sum_{x_1,x_2,\ldots x_q} \langle x| Q |x_1\rangle\langle x_1| Q |x_2\rangle \ldots \langle x_{q-1}| Q |x_q\rangle\langle x_q| A |x\rangle . \tag{6.80}$$

If Q is a polynomial of degree d in terms of matrices $R_1, \ldots R_m$ for $m = \mathsf{poly}(n)$, we can write it as

$$Q = \sum_{\substack{i_1, i_2, \ldots i_m \in [d] \\ i_1 + i_2 + \ldots i_m \leq d}} p_{i_1 i_2 \ldots i_m} R_1^{i_1} R_2^{i_2} \ldots R_m^{i_m}, \tag{6.81}$$

where each tuple $(i_1, \ldots i_m)$ specifies a monomial. The number of terms in the polynomial is bounded above by $(d+1)^m = \exp[m \ \log(d+1)] = O(\exp[\mathsf{poly}(n)])$. We write a term of Eq. (6.92), $\langle x_j | \, Q \, | x_{j+1} \rangle$, as

$$\begin{aligned} \langle x_j | \, Q \, | x_{j+1} \rangle = \sum_{\substack{i_1, i_2, \ldots i_m \in [d] \\ i_1 + i_2 + \ldots i_m \leq d}} & p_{i_1 i_2 \ldots i_m} \, \langle x_j | \, R_1^{i_1} \, | z_{j,1} \rangle \times \\ & \langle z_{j,1} | \, R_2^{i_2} \, | z_{j,2} \rangle \langle z_{j,2} | \ldots \langle z_{j,m-1} | \, R_m^{i_m} \, | x_{j+1} \rangle \, . \end{aligned} \tag{6.82}$$

We can further insert resolutions of the identity in Eq. (6.82) to get a sum over yet more terms. Each term in the resulting sum is a product over polynomially many quantities of the form $\langle w_1 | \, R_s \, | w_2 \rangle$ for some computational basis states $| w_1 \rangle \, , | w_2 \rangle$ and an index $s \in [m]$. Each of these can be computed in polynomial time. The number of terms in the final sum of the form in Eq. (6.92) is still bounded above by 2^{poly}.

From the assumption, the matrix elements of the matrices R_i can be computed to additive error $2^{-g(n)}$ in time scaling as $O(g(n))$. We therefore choose $g(n)$ to be such that the total additive error resulting from the 2^{poly} many paths in Eq. (6.92) is negligible compared to $(b' - a') \times 2^{\mathsf{poly}}$, where the second term (2^{poly}) corresponds to the number of terms in the sum. This can be ensured by taking $g(n)$ to be a sufficiently large polynomial.

Equation (6.92) is a sum over $T = O(2^{\mathsf{poly}})$ many terms f_i, each of which may be computed in polynomial time. Each term of Eq. (6.92) may be interpreted as a path in a Turing machine. Therefore, a PP machine can decide whether $\sum_{i=1}^{T} f_i$ is $\leq a'$ or $\geq b'$ for some thresholds $a', b' \geq a' + \Omega(2^{-\mathsf{poly}})$ input to the PP machine. This is seen as follows. Each term f_i is an efficiently computable real-valued function of the trajectory $x_0^i, x_1^i, \ldots x_K^i$. Let $a_{\max}$ be an upper bound to the norm of A. The PP machine selects a uniformly random trajectory and computes f_i. It accepts with probability $\frac{1}{2} - \frac{f_i}{2^{n+1} a_{\max}} > 0$ and rejects otherwise. The overall acceptance probability is $\frac{1}{T}\sum_i (\frac{1}{2} - \frac{f_i}{2^{n+1} a_{\max}})$. In the YES case, this is at least $\frac{1}{2} - \frac{a'}{2^{n+1} T a_{\max}}$, while in the NO case, it is at most $\frac{1}{2} - \frac{b'}{2^{n+1} T a_{\max}}$. Since we at least have a separation of $2^{-n-1}/T \times \Omega(b' - a') = \Omega(2^{-\mathsf{poly}(n)})$ between the YES and NO instances, this is a valid PP algorithm. □

Lemma 60 applies to the proof of Lemma 32 because the accept operator Q in that proof is a degree $2T + 3$-polynomial in matrices with efficiently computable entries.

For the proof of Lemma 33, we show in Lemma 61 that beginning from the maximally mixed initial state, imaginary time evolution for "time" $-i\beta$ produces a thermal state with high enough overlap with the ground state for a suitable β. Computing local observables in the obtained thermal state then suffices to get exponentially good estimates of ground-state local observables for gapped systems. We make the choice of a maximally mixed initial state in the above because it is guaranteed to have at least overlap 2^{-n} with the ground state.

Lemma 61. *For a Hamiltonian H with spectral gap at least Δ, let ρ_β be the thermal state at temperature $1/(2\beta)$. Also let $|E_1\rangle$ be the ground state of H and let A be any local observable satisfying $\|A\| \leq \mathsf{poly}(n)$. Then for $\beta = \Omega(n\Delta^{-1})$, the thermal expectation value satisfies*

$|\mathrm{Tr}[\rho_\beta A] - \langle E_1| A |E_1\rangle| \leq 2^{-\mathsf{poly}}$.

Proof. Let the eigenstates of the Hamiltonian be given by $|E_i\rangle$, $i \in [2^n]$, with the eigenvalues E_i arranged in nondecreasing order. Consider the initial state $\rho = \mathbb{1}/2^n$ and apply the linear operation $\exp(-\beta H)$, which performs imaginary time evolution for "time" $-i\beta$:

$$\rho \to \rho' = \exp(-\beta H)\rho \ \exp(-\beta H), \tag{6.83}$$

up to normalization. The maximally mixed initial state $\rho = \frac{\mathbb{1}}{2^n} = \sum_i \frac{1}{2^n} |E_i\rangle\langle E_i|$ transforms to the state ρ_β, given by

$$\rho_\beta = \frac{\rho'}{\mathcal{N}} = \frac{1}{\mathcal{N}} e^{-\beta H} \sum_i \frac{1}{2^n} |E_i\rangle\langle E_i| e^{-\beta H} \tag{6.84}$$

$$= \frac{1}{2^n \mathcal{N}} \sum_i e^{-2\beta E_i} |E_i\rangle\langle E_i| . \tag{6.85}$$

This state is the same as the thermal state $e^{-2\beta H}$ at temperature $1/(2\beta)$ up to normalization. The normalization factor $\mathcal{N} = \mathrm{Tr}\,\rho'$ is given by $\sum_i e^{-2\beta E_i}/2^n$. The overlap of the normalized state with the ground state is thus

$$\mathrm{Tr}[\rho_\beta |E_1\rangle\langle E_1|] = \frac{e^{-2\beta E_1}}{2^n \mathcal{N}} \tag{6.86}$$

$$= \frac{e^{-2\beta E_1}}{\sum_i e^{-2\beta E_i}} \tag{6.87}$$

$$= \left(1 + \sum_{i \neq 1} e^{-2\beta(E_i - E_1)}\right)^{-1} . \tag{6.88}$$

Since $E_i - E_1 = \Delta = \Omega(1/n^c)$, if β is taken to be $\Omega(n^d)$ with $d \geq c + 1$, we have that $e^{-2\beta(E_i - E_1)} \leq \exp\left[-2n^{d-c}\right]$. This means that the overlap is at least $1/(1+\exp\left[n \ \log \ 2 - 2n^{d-c}\right]) \geq$

$1 - \exp\left[n \log 2 - 2n^{d-c}\right]$, which means that the trace distance between the normalized states is $\varepsilon = O(\exp\left[-n^{d-c}\right])$. Therefore, the choice $\beta = \Theta(n\Delta^{-1})$ suffices to ensure that the resulting (normalized) state ρ_β is exponentially close to the ground state. Therefore, the thermal expectation value of any local observable A with polynomially bounded spectral norm is also exponentially close to the ground-state expectation value $\langle E_1| A |E_1\rangle$. □

We now show the following lemma about computing unnormalized thermal expectation values, which is the core subroutine of our P^{PP} algorithm.

Lemma 62. *A* PP *algorithm can decide whether* $\mathrm{Tr}\left[e^{-2\beta H} A\right] \leq a'$ *or* $\geq b'$ *when input a sparse Hamiltonian* H*, a number* $\beta \geq 0$*, a local observable* A*, and thresholds* a' *and* b'*.*

Proof. We express the unnormalized thermal expectation value as a sum over several paths as follows:

$$\mathcal{A}_\beta = \mathrm{Tr}\left[e^{-2\beta H} A\right] \tag{6.89}$$

$$= \sum_{x,y} \langle x| e^{-2\beta H} |y\rangle\langle y| A |x\rangle \tag{6.90}$$

$$\approx \sum_{x,y} \langle x| \left(\mathbb{1} - 2\beta H + 2(\beta H)^2 + \ldots \frac{(-2\beta H)^K}{K!}\right) |y\rangle \times$$

$$\langle y| A |x\rangle =: \mathcal{A'}_\beta \tag{6.91}$$

$$= \sum_{k=0}^{K} \frac{1}{k!} \sum_{x_0, x_1, \ldots x_k} \langle x_0| - 2\beta H |x_1\rangle\langle x_1| - 2\beta H |x_2\rangle \ldots$$

$$\langle x_{k-1}| - 2\beta H |x_k\rangle\langle x_k| A |x_0\rangle . \tag{6.92}$$

$$= \sum_{i=1}^{T} f_i. \tag{6.93}$$

This expression is reminiscent of a Euclidean path integral, although there are some differ-

ences. In a Euclidean path integral, one Trotterizes the map $\exp(-\beta H) \approx (\prod_i \exp(-\beta H_i/r))^r$ and use the fact that each term of the Hamiltonian H_i is local in order to compute terms in the series. In contrast, here we have used the Taylor expansion for $\exp(-\beta H)$ and have inserted resolutions of the identity in order to compute the terms $\langle x| H^k |y\rangle$. Using the Taylor series allows us to get exponentially small additive error, which is not guaranteed by Trotterization.

Before we move on, let us analyze the additive error in Eq. (6.91). It is given by:

$$\epsilon \leq \frac{(2\beta\|H\|)^{K+1}}{(K+1)!}\|A\| \times O(1). \tag{6.94}$$

By choosing $K > 2\beta e\|H\| + f(n)$ for some polynomial $f(n) = O(\beta\|H\|/n)$ and $f(n) = \Omega(n)$, we can ensure that the error is bounded above by $\|A\| \exp[-f(n)]$:

$$K+1 \geq 2\beta e\|H\| + f(n) \tag{6.95}$$

$$\begin{aligned}\implies (K+1)\log(K+1) \geq (K+1)\log(2\beta e\|H\|)+ \\ (K+1)\log\left(1+\frac{f(n)}{2\beta e\|H\|}\right)\end{aligned} \tag{6.96}$$

$$\begin{aligned}\geq &(K+1)\log(2\beta e\|H\|) + (K+1)\frac{f(n)}{2\beta e\|H\|} - \\ &\frac{K+1}{2}\left(\frac{f(n)}{2\beta e\|H\|}\right)^2\end{aligned} \tag{6.97}$$

$$\geq (K+1)\log(2\beta e\|H\|) + \Omega(f(n)), \tag{6.98}$$

where we have used the fact that $\log(1+x) \geq x - \frac{x^2}{2}$ for small x and that $f(n) = o(\beta\|H\|)$.

Therefore,

$$\log\left(\frac{(2\beta\|H\|)^{k+1}}{(K+1)!}\right) \leq -\Omega(f(n)), \text{ giving} \tag{6.99}$$

$$\epsilon \leq O\left(\|A\| \exp[-f(n)]\right). \tag{6.100}$$

□

Proof of Lemma 33. From Lemma 61, we know that the normalized state is exponentially close to the true ground state. Therefore, deciding whether the ground state has $\mathrm{Tr}[|\Psi\rangle\langle\Psi| A] \leq a$ or $\geq b$ is equivalent to deciding whether the unnormalized state has expectation value $\mathrm{Tr}[\rho' A] \leq a' = \mathcal{N}_{\text{est}}(a + \|A\|\varepsilon)$ or $\geq b' = \mathcal{N}_{\text{est}}(b - \|A\|\varepsilon)$, where $\mathcal{N}_{\text{est}}$ is an estimate of the normalization of the state and ε the trace distance between the ground state and the thermal state. To maintain a gap between the YES and NO cases, we need $\varepsilon < 2^{-u(n)}/\|A\|$ for some polynomial u, which can be satisfied by taking n^{d-c} in Lemma 61 to be $\geq u(n) + \log\|A\|$. The norm of A is bounded above by a polynomial in n and therefore is a subleading term.

Since the thresholds a' and b' depend on the normalization, we should compute the normalization $\mathcal{N}$ beforehand. Since the normalization is a special case of Eq. (6.89) with $A = \mathbb{1}$, we can use the PP procedure to decide if $\mathcal{N} \leq a_1$ or $\mathcal{N} \geq a_2$ for some a_1, a_2 with $a_2 - a_1 = \Omega(1/\exp)$. Performing binary search over the interval $(0, 1]$ with polynomially many queries to the PP oracle, we can estimate the normalization to exponentially small additive error, giving an estimate $\mathcal{N}_{\text{est}}$.

Therefore, we have shown that a P^{PP} machine can do all the above: compute the normalization and then compute the thermal expectation value for a low-temperature state.

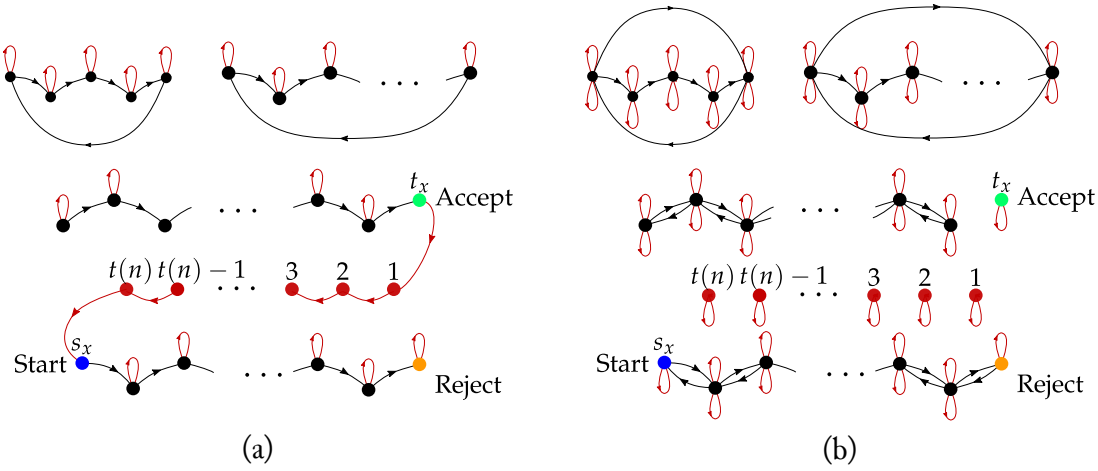

Figure 6.3: (a) Graph G'_x with adjacency matrix A'_x, adapted from Fig. 6.1. (b) Graph with (weighted, directed) adjacency matrix $A_x^{\dagger\prime} A_x$. Vertices with two self-loops can be thought of as a single self-loop with weight 2.

Since we have also shown that setting $\beta = (n/\Delta)$ suffices to get exponentially small error, we have shown that the problem is in P^{PP}. □

This technique is also applicable to Hamiltonians or Hermitian operators that are not necessarily local, or even sparse. For example, it can apply to Hermitian operators of the kind in Lemma 60.

6.11 Turing machine construction for PSPACE-hardness

In this section, we complete the proof of Lemma 39.

Lemma 63 (Lower bound on spectral gap for PSPACE-hard construction)**.** *In the NO case, the construction in the proof of Lemma 39 has a spectral gap of* $\Theta(\ell_{\max}^{-2})$*, where* $\ell_{\max}$ *is the number of vertices in the largest subgraph of* G'_x*.*

Proof. Recall the form of the graph G'_x in the NO case, reproduced here in Fig. 6.3(a). We first restrict our attention to the subgraph of G'_x containing the start and accept con-

figurations. The matrix $A_x^{\dagger\prime}A'_x$, when restricted to this subspace, is further composed of three subspaces, each corresponding to a subgraph, as shown in Fig. 6.3(b). We write $A_x^{\dagger\prime}A'_x = \mathcal{G}_1 \oplus \mathcal{G}_2 \oplus \mathcal{G}_3$. The block $\mathcal{G}_1$ corresponds to the vertices leading to t_x (not including t_x). The block $\mathcal{G}_2$ corresponds to the vertices $\{t_x\} \cup \{1, \dots t(n)\}$. Lastly, $\mathcal{G}_3$ is the block with the vertices starting from s_x and leading to the reject state, which are the configurations visited by the Turing machine. We have

$$\mathcal{G}_1 = \begin{pmatrix} 2 & 1 & & & & \\ 1 & 2 & 1 & & & \\ & 1 & 2 & \ddots & & \\ & & \ddots & \ddots & 1 & \\ & & & 1 & 2 & \end{pmatrix}_{\ell_1 \times \ell_1}, \ \mathcal{G}_2 = \mathbb{1}_{\ell_2 \times \ell_2}, \text{ and}$$

$$\mathcal{G}_3 = \begin{pmatrix} 1 & 1 & & & & & \\ 1 & 2 & 1 & & & & \\ & 1 & 2 & \ddots & & & \\ & & \ddots & \ddots & 1 & & \\ & & & 1 & 2 & 1 & \\ & & & & 1 & 1 & \end{pmatrix}_{\ell_3 \times \ell_3}. \tag{6.101}$$

It may be seen that there is a zero eigenvector $(0, 0, \dots, 0, 1, -1, 1, \dots (-1)^{\ell_3})^T$, with the zeros corresponding to the subspaces $\mathcal{G}_1$ and $\mathcal{G}_2$. We now lower-bound the next-smallest

eigenvalue. Let

$$
\begin{aligned}
r_n(\lambda) &:= \det[\mathcal{G}_3 - \lambda \mathbb{1}_n] \\
&= \det \begin{pmatrix}
1-\lambda & 1 & & & & & \\
1 & 2-\lambda & 1 & & & & \\
& 1 & 2-\lambda & \ddots & & & \\
& & \ddots & \ddots & 1 & & \\
& & & 1 & 2-\lambda & 1 & \\
& & & & 1 & 1-\lambda &
\end{pmatrix}_{n\times n} \qquad (6.102) \\
p_n(\lambda) &:= \det \begin{pmatrix}
2-\lambda & 1 & & & & & \\
1 & 2-\lambda & 1 & & & & \\
& 1 & 2-\lambda & \ddots & & & \\
& & \ddots & \ddots & 1 & & \\
& & & 1 & 2-\lambda & 1 & \\
& & & & 1 & 1-\lambda &
\end{pmatrix}_{n\times n} . \qquad (6.103)
\end{aligned}
$$

The polynomial $p_n(\lambda)$ can be computed exactly [338], and is given by $p_n(2-2\cos\theta) = \frac{\sin((n+1)\theta)-\sin(n\theta)}{\sin\theta} = \frac{\cos\left((n+\frac{1}{2})\theta\right)}{\cos\left(\frac{\theta}{2}\right)}$. We can obtain $r_n(\lambda)$ in terms of $p_n(\lambda)$: $r_n(\lambda) = (1-\lambda)p_{n-1}(\lambda) - p_{n-2}(\lambda)$, giving us

$$
\begin{aligned}
r_n(\lambda) &= f_n(\theta) \\
&= (2\cos\theta - 1)\frac{\cos\left((n-\frac{1}{2})\theta\right)}{\cos\left(\frac{\theta}{2}\right)} - \frac{\cos\left((n-\frac{3}{2})\theta\right)}{\cos\left(\frac{\theta}{2}\right)}, \qquad (6.104)
\end{aligned}
$$

where $\theta = \cos^{-1}(1 - \frac{\lambda}{2})$, or $\lambda = 2 - 2\cos\theta$. The eigenvalues of $\mathcal{G}_3$ are related to the roots of the characteristic polynomial $f_n(\theta) = 0$. We can see that $\theta = 0$ is always a root of the polynomial, giving us the zero eigenvalue ($\lambda = 2 - 2\cos\theta = 0$) for the NO case.

Now, it remains to be shown that the next smallest eigenvalue is bounded away from zero. First consider $\mathcal{G}_1$, whose eigenvalues are the roots of the characteristic equation $\det[\mathcal{G}_1 - \lambda \mathbb{1}_{\ell_1}]$. The eigenvalues of $\mathcal{G}_1$ can be computed in a similar fashion to those of $\mathcal{G}_3$ and are given by $4\sin^2\left(\frac{k\pi}{2(\ell_1+1)}\right), k \in [n]$. The smallest eigenvalue of $\mathcal{G}_1$ is therefore at least $\Omega\left(1/{\ell_1}^2\right)$. It is also easily seen that $\mathcal{G}_2 \succ 0$.

We now come to $\mathcal{G}_3$. As we have seen, $\mathcal{G}_3$ has a zero eigenvalue. In order to show a spectral gap for $\mathcal{G}_3$, we show that the next root of the polynomial $f_{\ell_3}(\theta)$ must occur at least a distance $\Omega({\ell_3}^{-2})$ away. The roots of $\mathcal{G}_3$ are given by [378]

$$\lambda_j = 2 + 2\cos\left(\frac{\pi j}{\ell_3}\right), \quad j \in [\ell_3]. \tag{6.105}$$

Setting $j = \ell_3$ gives the zero eigenvalue and $j = \ell_3 - 1$ the first nonzero eigenvalue. The spectral gap of $\mathcal{G}_3$ is therefore

$$\begin{aligned}
\lambda_{\ell_3 - 1} &= 2 - 2\cos\left(\frac{\pi}{\ell_3}\right) && (6.106)\\
&= 4\sin^2\left(\frac{\pi}{2\ell_3}\right) && (6.107)\\
&\geq \frac{\pi^2}{\ell_3^2} - O\left(\frac{\pi^4}{\ell_3^4}\right) && (6.108)\\
&= \Omega\left(1/\ell_3^2\right). && (6.109)
\end{aligned}$$

Finally, we consider other subgraphs that do not contain the start vertex. Just like the

analysis of the YES case, the eigenvalues for these are bounded away from 0 by ℓ^{-2}, where ℓ is the number of vertices in the subgraph. We have therefore lower bounded the value of the nonzero eigenvalue in each case, showing that the spectral gap is $\Omega(\ell_{\max}^{-2}) = \Omega(2^{-\mathsf{poly}})$. □

6.12 Complexity of PrecisePGQMA and PreciseEGQMA with asymmetric spectral gaps

We show here that the promise of asymmetric spectral gaps does not change the complexity class for both PrecisePGQMA and PreciseEGQMA, proving Theorem 42.

Proof of Theorem 42. It is easy to see that $\mathsf{GQMA}[c, s, g_1, g_2] \subseteq \mathsf{GQMA}[c, s, g_1, 0]$ simply by ignoring the promise on the NO instance. It remains to show that the same upper bounds as the symmetric case hold for the asymmetric case too. For the case of $c - s = \Omega(1/\exp)$, $g_2 = \Omega(1/\exp)$, we observe that one can also ignore the promise on the YES instance and obtain containment in $\mathsf{PreciseQMA} = \mathsf{PSPACE}$, which equals PreciseEGQMA.

It remains to give an upper bound for the class $\cup_{\substack{c-s\geq\Omega(1/\exp)\\ g_1\geq\Omega(1/\mathsf{poly})}}\mathsf{GQMA}[c, s, g_1, 0]$. We give a PP algorithm for any instance from this class, which implies equivalence of the two classes.

We are given a description of a circuit, with the promise that the YES case has $\Omega(1/\mathsf{poly})$ spectral gap for the accept operator Q. We want to decide if $\lambda_1(Q)$ is $\geq c$ (YES) or $\leq s$ (NO). The overall PP algorithm is as follows.

1. Use the $\mathsf{P}^{\mathsf{QMA[log]}}$ algorithm of Ambainis [370] to determine whether an instance has spectral gap $\Delta \geq g_1$ (YES) or $\leq g_1/2$ (NO), for $g_1 = \Omega(1/\mathsf{poly})$.

2. If the spectral gap is g_1 or larger, run the algorithm in Lemma 60 with Hamiltonian $\mathbb{1} - Q$ and accept or reject according to the answer returned by the algorithm.

3. Otherwise reject.

We claim that the algorithm of Ambainis works not just for local Hamiltonians, but also for accept operators like Q. This is because the QMA queries in Ambainis's algorithm pertain to whether the ground-state energy (or the minimum eigenvalue $1 - \lambda_1$ in this case) is smaller or larger than a threshold. A QMA verifier can compute the eigenvalue of the accept operator given an eigenstate, using phase estimation. Therefore, all queries to the oracle about $1 - \lambda_1$ are still valid QMA queries. Also, the final query in Ambainis's algorithm is for the operator $(\mathbb{1} - Q) \otimes \mathbb{1} + \mathbb{1} \otimes (\mathbb{1} - Q)$ on two registers, restricted to the antisymmetric subspace. Since a QMA verifier can also perform a projection onto the antisymmetric subspace, Ambainis's algorithm (i.e. the first step) works to estimate the spectral gap of Q in $\mathsf{P}^{\mathsf{QMA}[\log]}$.

Now, since $\mathsf{P}^{\mathsf{QMA}[\log]} \subseteq \mathsf{PP}$ [371], the overall algorithm is a valid PP algorithm, since the two queries can be made in parallel. To see the correctness, we see that if the instance has a YES answer, then it has a spectral gap of at least g_1 by virtue of the promise. In this case the spectral gap algorithm would return YES. This ensures that the PP algorithm in Lemma 60 works correctly and returns the correct answer $E_1 \leq a$ (YES) or $E_1 \geq b$ (NO). The algorithm outputs YES since the instance has low energy.

In the NO case, there may or may not be a spectral gap. If the spectral gap $\Delta \leq g_1/2$ is not large enough, the spectral gap algorithm returns NO. We reject in this case. If the spectral gap algorithm returns YES, then the spectral gap is at least $\Delta \geq g_1/2$ (this includes the cases when the spectral gap is in the window $[g_1/2, g_1]$, which is outside of the promise

in the spectral gap algorithm). This means that the algorithm in Section 6.10 will work, and return the correct output (NO). Therefore, we see that $\cup_{\substack{c-s \geq \Omega(1/\exp) \\ g_1 \geq \Omega(1/\mathsf{poly})}} \mathsf{GQMA}[c, s, g_1, 0] = \mathsf{PrecisePGQMA}$. □

We remark that it can be seen that $\textsc{LocalHamiltonian}[a, b, g_1, 0]$ with $b - a = \Theta(1/\exp)$ is PrecisePGQMA-complete when the spectral gap g_1 is $1/\mathsf{poly}$ and PreciseEGQMA-complete when g_1 is $1/\exp$.

Appendix A: Mathematical Preliminaries

In this Chapter, we define some notation and some standard terminology in use throughout this dissertation.

A.1 Notation

Definition 20 (Big-Oh)**.** We define $O(g(n))$ as the set of all functions $f(n)$ that satisfy $\forall n \geq n_0, f(n) \leq cg(n)$ for some constant $c > 0$ independent of n.

Definition 21 (Big-Omega)**.** We define $\Omega(g(n))$ as the set of all functions $f(n)$ such that $g(n) \in f(n)$, or equivalently, the set of all functions $f(n)$ that satisfy $\forall n \geq n_0, g(n) \leq cf(n)$ for some constant c.

We have $f(n) \in \Omega(g(n)) \Leftrightarrow g(n) \in O(f(n))$.

Definition 22 (Theta)**.** We define $\Theta(g(n))$ as the set of all functions $f(n)$ satisfying $\forall n \geq n_0, c_1 f(n) \leq g(n) \leq c_2 f(n)$ for some constants $c_2 > c_1 \geq 0$.

In other words, $f(n) \in \Theta(g(n))$ is equivalent to $f(n) \in O(g(n))$ and $f(n) \in \Omega(g(n))$. For all these definitions, we sometimes abuse notation and write $f(n) = O(g(n))$ instead of $f(n) \in O(g(n))$.

Lastly, we also have the o (little-oh) and ω (little-omega) notation.

Definition 23 (Little-oh)**.** We define $o(g(n))$ as the set of all functions $f(n)$ satisfying $\lim_{n\to\infty} \frac{f(n)}{g(n)} = 0$.

Definition 24 (Little-omega)**.** The set $\omega(g(n))$ is the set of functions $f(n)$ satisfying $g(n) \in o(f(n))$, or equivalently, satisfying $\lim_{n\to\infty} \frac{f(n)}{g(n)} \to \infty$.

A.2 Notions of error

Broadly speaking, we have two notions of error: additive and multiplicative. Below, we take the example of computing a target function to some error, but the same notions extend to that of sampling from a target distribution.

Definition 25 (Additive error)**.** An algorithm computes a function $f(x)$ to additive error ϵ if its output $g(x)$ satisfies $g(x) \in [f(x) - \epsilon, f(x) + \epsilon]$.

Definition 26 (Multiplicative error)**.** An algorithm computes a function $f(x)$ to within a multiplicative factor $c > 1$ if its output $g(x)$ satisfies

$$\frac{f(x)}{c} \leq g(x) \leq cg(x). \tag{A.1}$$

We are often interested in the setting $c \approx 1$, where we write $c = 1 + \epsilon$. In this case, we call ϵ the multiplicative error and the output satisfies

$$\frac{f(x)}{1+\epsilon} \leq g(x) \leq (1+\epsilon)f(x). \tag{A.2}$$

$$\implies (1-\epsilon)f(x) \leq g(x) \leq (1+\epsilon)f(x). \tag{A.3}$$

A.3 Notions of simulation

For an n-qudit quantum system and a polynomial-sized description of the quantum state[1], we define the following simulation tasks.

Definition 27 (Strong Simulation)**.** In the task of strong simulation, we are given as input a suitable description of the state $|\psi\rangle$. The task is to output the expectation value of any desired observable in the state, namely $|\langle\psi| O |\psi\rangle|^2$. The output may be exact or approximate, and either accurate to multiplicative or additive error.

This task includes computing quantities such as full output probabilities, marginals of output probabilities, and local observables.

Definition 28 (Weak simulation)**.** In weak simulation, the input is again a suitable description of the state $|\psi\rangle$. The task is to output a sample from the distribution $\mathcal{D}$ obtained by measuring $|\psi\rangle$ in a suitable basis. The output may be exact or approximate up to multiplicative error (measured in terms of the multiplicative error of each output probability).

This task includes the ability to compute local observables diagonal in the same basis up to some small additive error.

Definition 29 (Additive-error approximate sampling)**.** In additive-error approximate sampling, the task is similar to the one above. The desired output is now allowed to be approximate up to small additive error (measured in terms of the total variation distance between the distributions).

[1] We do not mean an *explicit* description of the state, but one that can be implicit. For example, the state can be defined as the output state after Hamiltonian evolution under an efficiently-specifiable Hamiltonian for time t. Or it can be defined as an equilibrium state of such a Hamiltonian. Ultimately, we assume that physically interesting states can be described efficiently in this sense.

Again, this includes the ability to compute local observables up to some small additive error.

Definition 30 (Computing local observables)**.** In this task, the input is the same, a suitable description of the state $|\psi\rangle$. The desired output is the expectation value $|\langle\psi|O|\psi\rangle|^2$ of a k-local observable O, meaning an observable that can be decomposed as $O = \sum_i O_i$, each term of which is supported on k qudits at most.

Appendix B: Complexity-theoretic basics

Here, we give a very brief introduction to the complexity-theoretic definitions and terminology in this Chapter. The reader is referred to a textbook (e.g. Refs. [379; 380]) for a more pedagogical exposition. We are generally concerned with decision problems, where the answer is either "YES" or "NO". These problems can be cast as follows: given an instance x, the task is to decide if it belongs to the class of YES instances ($x \in A_{\text{yes}}$, also called "accept"), or to the class of NO instances ($x \in A_{\text{no}}$, also called "reject"). In principle, there can be problems where certain instances (for example, ill-defined ones) belong neither to A_{yes} or A_{no}. In such cases, we either allow an algorithm to answer arbitrarily, or we supplant the problem with a promise that such instances never occur. These are called *promise problems*.

In complexity theory, one is typically interested in the resources taken to solve various classes of decision problems. Further, one is interested in how the resource cost scales with the size of the problem to be solved, which is quantified in terms of the length of the input, often denoted n. In this dissertation, we use the notation $\mathsf{poly}(n)$ to denote any function that can be upper bounded by $O(n^c)$ for some constant $c = \Theta(1)$. We also denote $\exp(n)$ to be any function $2^{\mathsf{poly}(n)}$. We will omit the dependence on n, which in this dissertation is usually taken to be the number of particles or qudits.

To characterize the complexity of a problem, we give "upper" and "lower" bounds on

the complexity of the problem. Upper bounds are statements of the form "X $\in$ Y", which means that the problem X can be solved with access to a solver for the complexity class Y. For example, Shor [381] proved that FACTORING $\in$ BQP, which means that quantum computers can factor integers in polynomial time (since quantum computers may be viewed as "solvers for the class BQP"). Lower bounds are statements of the form "X is Y-hard". This means that problem X is as hard as any problem in Y. Such statements are often shown via *reductions*. One assumes the existence of an *oracle*, a black box that can solve any instance of the problem X in one timestep. A reduction is a mapping from the complexity class Y to the problem X with the property that any problem in Y can be solved by querying the oracle for X. If such a reduction exists, it implies that the problem X is at least as hard as any problem in the class Y. If a problem X is both in the class Y and is Y-hard, then it means that the upper and lower bounds to the problem match. This means that the problem X is the hardest in the class it belongs to, namely Y. In this case, we say "X is Y-complete." or "X is complete for Y.". We also denote by Y^{Z} the class of problems solvable by a Y machine with access to an oracle for any problem in Z.

Definition 31 (P)**.** The class P (Polynomial time) is defined to be the class of problems efficiently solvable in polynomial time on a classical computer. More formally, it is the class of problems $A = (A_{\text{yes}}, A_{\text{no}} = \{0,1\}^* \backslash A_{\text{yes}})$ such that there exists a polynomial-time Turing machine M that on input a description x of the instance, outputs a bit $M(x)$ satisfying:

$x \in A_{\text{yes}} \Leftrightarrow \quad M(x) = 1$, and

$x \in A_{\text{no}} \Leftrightarrow \quad M(x; y) = 0.$

We interpret the output $M(x) = 1$ as the machine accepting (answering YES), and $M(x) = 0$ as the machine rejecting (answering NO).

We turn to the extremely important class in classical complexity theory, nondeterministic polynomial time (NP). We imagine two parties, Merlin (the prover) and Arthur (the verifier). The prover would like to convince the verifier that a certain problem instance x is a YES instance. The prover, who is computationally unbounded, can supply any bit-string y of length $w = \mathsf{poly}(n)$ to the verifier as a "proof" or "witness". The verifier performs a polynomial-time computation that reads any bit in the bit-string y as many times as desired and either accepts or rejects. NP is the class of problems such that a YES answer can be reliably verified in this way and in case the answer is NO, no matter what string is sent by the (possibly cheating) prover, the verifier rejects.

Definition 32 (NP)**.** NP is the class of problems $A = (A_{\text{yes}}, A_{\text{no}})$ such that there is a polynomial-time Turing machine M that takes as input the instance x and a string y with $|y| = \mathsf{poly}(n)$ such that for every instance,

$x \in A_{\text{yes}} \Leftrightarrow \quad \exists\, y$ such that $M(x; y) = 1$, and

$x \in A_{\text{no}} \Leftrightarrow \quad \forall\, y$, $M(x; y) = 0$.

One of the most important questions in computer science is the P vs. NP problem, which asks whether $\mathsf{P} \stackrel{?}{=} \mathsf{NP}$. The widely believed answer to this question is $\mathsf{P} \neq \mathsf{NP}$.

Conjecture 64 (P vs. NP)**.** $\mathsf{P} \neq \mathsf{NP}$.

We can generalize the class P to allow a classical computer to use randomness and solve problems merely with high probability rather than deterministically. This is the class BPP, which stands for Bounded-error Probabilistic Polynomial time. The error here is measured via the parameters c (minimum probability of saying "YES" if the answer is YES) and s (maximum probability of saying "YES" if the answer is NO).

Definition 33 ($\mathsf{BPP}[c,s]$)**.** $\mathsf{BPP}[c,s]$ is the class of problems $A = (A_{\text{yes}}, A_{\text{no}})$ such that there is a polynomial-time Turing machine M that takes as input the instance x and a random string $y \sim \mathcal{U}$ from the uniform distribution over bit-strings of length $|y| = \mathsf{poly}(n)$ such that for every instance,

$x \in A_{\text{yes}} \Leftrightarrow \quad \Pr_{y\sim\mathcal{U}}[M(x;y) = 1] \geq c$, and

$x \in A_{\text{no}} \Leftrightarrow \quad \Pr_{y\sim\mathcal{U}}[M(x;y) = 1] \leq s$.

The class BPP is simply defined to be $\mathsf{BPP} := \cup_{c,s,:c-s\geq 1/\mathsf{poly}(n)}\mathsf{BPP}[c,s]$. This turns out to be the same as the class $\mathsf{BPP}[\frac{1}{2} + \frac{1}{\mathsf{poly}(n)}, \frac{1}{2} - \frac{1}{\mathsf{poly}(n)}]$, which in turn is the same as $\mathsf{BPP}[2/3, 1/3]$ and $\mathsf{BPP}[1 - 2^{-r}, 2^{-r}]$ for some $r = \mathsf{poly}(n)$. The key to this fact is the Chernoff bound, which allows us to repeat a computation with small *promise gap* $c - s$ and take a majority vote to get exponentially small error $c = 1 - 2^{-\mathsf{poly}(n)}$, $s = 2^{-\mathsf{poly}(n)}$.

If we remove the constraint that there be a small "gap" between the completeness and the soundness, we get a very different class, namely Probabilistic Polynomial time (PP)[1].

Definition 34 (PP)**.** PP is the class of problems $A = (A_{\text{yes}}, A_{\text{no}})$ for which there exists a polynomial Turing machine M that takes as input the instance x and a random string $y \sim \mathcal{U}$ from the uniform distribution over bit-strings of length $|y| = \mathsf{poly}(n)$ and outputs $M(x;y)$ with the following properties:

$x \in A_{\text{yes}} \Leftrightarrow \quad \Pr_{y\sim\mathcal{U}}[M(x;y) = 1] > \frac{1}{2}$, and

$x \in A_{\text{no}} \Leftrightarrow \quad \Pr_{y\sim\mathcal{U}}[M(x;y) = 1] < \frac{1}{2}$.

In this definition, the "$\frac{1}{2}$" can be replaced by any constant in $(0,1)$. The key here is that the promise gap $c-s$ can be exponentially small in n. By virtue of the weak requirement on what constitutes solving a problem, PP is a large class, including, for instance, all of NP

[1] Note the absence of the modifier "Bounded-error"

and coNP.

The largest class we deal with in this dissertation is PSPACE.

Definition 35 (PSPACE)**.** PSPACE is the class of problems that can be deterministically solved by a Turing machine that uses at most a polynomial amount of space.

The space used by a Turing machine is the maximum number of cells in use at any given time step. In the above bound, we do not explicitly upper bound the amount of time the Turing machine is allowed to take. Yet we can prove that no more than exponential time is enough: PSPACE $\subseteq$ EXP, the class of problems solvable deterministically in exponential time.

B.1 Quantum classes

We now define the class BQP (Bounded-error Quantum Polynomial time), which is the class of problems solvable in polynomial time (in n) on a quantum computer with bounded error. More formally,

Definition 36 (BQP$[c, s]$)**.** BQP$[c, s]$ is the class of promise problems $A = (A_{\text{yes}}, A_{\text{no}})$ such that for every instance x, there is a uniformly generated circuit U_x of size poly(n) acting on the state $|0^{\otimes m}\rangle$ for $m = \text{poly}(n)$, with the property that upon measuring the first bit at the output, o, also called the decision qubit, we have

If $x \in A_{\text{yes}}$: $\Pr(o = 1) \geq c$

If $x \in A_{\text{no}}$: $\Pr(o = 1) \leq s$.

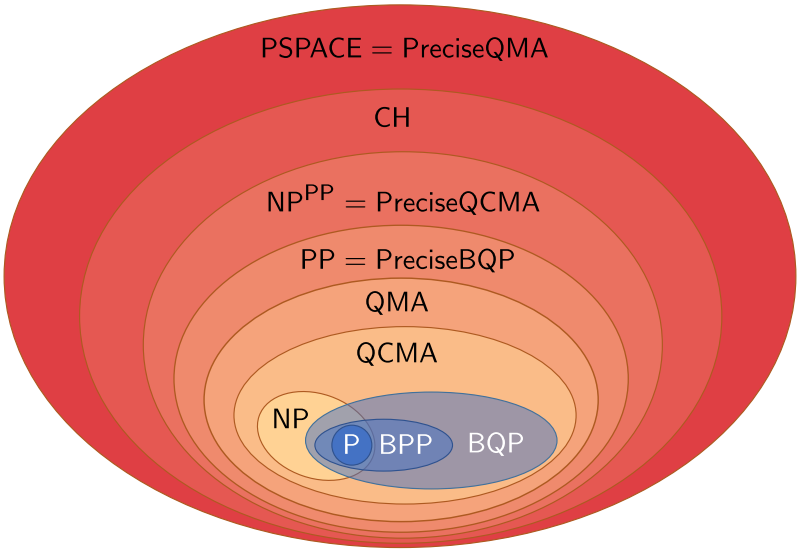

Figure B.1: Major complexity classes featuring in this dissertation, and especially Chapter 6. The classes PSPACE, $\mathsf{NP}^{\mathsf{PP}}$, and PP can be defined purely in terms of quantum computation, and are equal to PreciseQMA, PreciseQCMA, and PreciseBQP, respectively. All inclusions except $\mathsf{P} \subseteq \mathsf{BPP}$ are believed to be strict.

In the above, we imagine that a quantum computer applies a circuit U_x that acts on a standard initial state, measures the first bit at the output, and says YES ("accepts") or NO ("rejects"), depending on whether the bit is measured to be in state $|1\rangle$ or $|0\rangle$. The choice of the bit to measure at the output is arbitrary. The term *uniformly generated circuit* means that given an instance x there is a polynomial-time classical algorithm to generate a description of the circuit U_x to be applied.

Definition 37. $\mathsf{BQP} = \cup_{c-s \geq 1/\mathsf{poly}} \mathsf{BQP}[c, s]$.

The class BQP is the quantum generalization of class BPP (Bounded-Error Probabilistic Polynomial time), the class of problems solvable in polynomial time by a randomized classical computer.

We now come to the class QMA (Quantum Merlin Arthur), which is a quantum generalization of NP. Now the prover can send quantum states to the verifier, and the

verifier can perform quantum computation on top of it. QMA is the class of problems such that a YES answer can be reliably verified in this way and in case the answer is NO, no matter what state is sent by the (possibly cheating) prover, the verifier rejects with high probability. Just like with BQP, QMA is defined with respect to parameters c and s, which are called *completeness* and *soundness*, respectively.

Definition 38 (QMA$[c, s]$)**.** QMA$[c, s]$ is the class of problems $A = (A_{\text{yes}}, A_{\text{no}})$ with the property that, for every instance x, there exists a uniformly generated circuit U_x with the following properties: U_x is of size poly(n) and acts on an input state $|0\rangle^{\otimes m}$, together with a proof (or witness) state $|\Psi\rangle$ of size w supplied by an arbitrarily powerful prover. Both m and w are bounded by polynomials in n. Upon measuring the decision qubit o of the output register, the verifier accepts if $o = 1$, and rejects otherwise. We say $A = (A_{\text{yes}}, A_{\text{no}})$ is a QMA$[c, s]$ problem iff

If $x \in A_{\text{yes}}$: $\exists\ |\Psi\rangle$ such that $\Pr(o = 1) \geq c$

If $x \in A_{\text{no}}$: $\forall\ |\Psi\rangle$, $\Pr(o = 1) \leq s$.

QMA is defined as $\cup_{c-s\geq 1/\text{poly}}QMA[c, s]$.

Lastly, we depict the known inclusions between complexity classes in Fig. B.1. We also describe here the classes not mentioned so far. The class QCMA is analogous to QMA, except that the prover sends a classical witness instead of a quantum one. The class NP^{PP} is a subset of PP^{PP}, since NP $\subseteq$ PP. These classes belong to the counting hierarchy (CH), which is defined as CH $=$ PP $\cup\ \text{PP}^{\text{PP}} \cup \ldots$ [323]. All of these classes are in PSPACE.

Bibliography

[1] D. Deutsch, *The Fabric of Reality* (Penguin UK, 2011).

[2] S. Aaronson, "Guest Column: NP-complete problems and physical reality," ACM SIGACT News **36**, 30–52 (2005).

[3] R. P. Feynman, "Simulating physics with computers," Int. J. Theor. Phys. **21**, 467–488 (1982).

[4] S. P. Jordan, K. S. M. Lee, and J. Preskill, "Quantum Algorithms for Quantum Field Theories," Science **336**, 1130–1133 (2012).

[5] D. Harlow and P. Hayden, "Quantum Computation vs. Firewalls," Journal of High Energy Physics **2013**, 85 (2013).

[6] S. Aaronson, "The Complexity of Quantum States and Transformations: From Quantum Money to Black Holes," (2016), arXiv:1607.05256 [gr-qc, physics:quant-ph].

[7] A. Bouland, B. Fefferman, and U. Vazirani, "Computational Pseudorandomness, the Wormhole Growth Paradox, and Constraints on the AdS/CFT Duality (Abstract)," in *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 151, edited by T. Vidick (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2020) pp. 63:1–63:2.

[8] I. H. Kim, E. Tang, and J. Preskill, "The ghost in the radiation: Robust encodings of the black hole interior," Journal of High Energy Physics **2020**, 31 (2020).

[9] U. Mahadev, "Classical Verification of Quantum Computations," in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (2018) pp. 259–267.

[10] G. Valiant and P. Valiant, "An Automatic Inequality Prover and Instance Optimal Identity Testing," SIAM Journal on Computing **46**, 429–455 (2017).

[11] B. M. Terhal and D. P. DiVincenzo, "Adaptive Quantum Computation, Constant Depth Quantum Circuits and Arthur-Merlin Games," Quantum Inf. Comput. **4**, 134–145 (2002).

[12] M. J. Bremner, R. Jozsa, and D. J. Shepherd, "Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy," Proc. R. Soc. Math. Phys. Eng. Sci. **467**, 459–472 (2011).

[13] S. Aaronson and A. Arkhipov, "The computational complexity of linear optics," in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (ACM Press, New York, New York, USA, 2011) p. 333.

[14] S. Aaronson and A. Arkhipov, "The computational complexity of linear optics," Theory Comput. **9**, 143–252 (2013).

[15] L. Stockmeyer, "The complexity of approximate counting," in *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing*, STOC '83 (Association for Computing Machinery, New York, NY, USA, 1983) pp. 118–126.

[16] L. Stockmeyer, "On Approximation Algorithms for # P," SIAM J. Comput. **14**, 849–861 (1985).

[17] S. Toda, "PP is as Hard as the Polynomial-Time Hierarchy," SIAM J. Comput. **20**, 865–877 (1991).

[18] B. Fefferman and C. Umans, "On the Power of Quantum Fourier Sampling," in *11th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2016)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 61 (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2016) pp. 1:1–1:19.

[19] M. J. Bremner, A. Montanaro, and D. J. Shepherd, "Average-case complexity versus approximate simulation of commuting quantum computations," Phys. Rev. Lett. **117**, 080501 (2016).

[20] E. Farhi and A. W. Harrow, "Quantum Supremacy through the Quantum Approximate Optimization Algorithm," (2016), arXiv:1602.07674 [quant-ph].

[21] B. Fefferman, M. Foss-Feig, and A. V. Gorshkov, "Exact sampling hardness of Ising spin models," Phys. Rev. A **96**, 032324 (2017).

[22] X. Gao, S.-T. Wang, and L.-M. Duan, "Quantum Supremacy for Simulating A Translation-Invariant Ising Spin Model," Phys. Rev. Lett. **118**, 040502 (2017).

[23] J. Miller, S. Sanders, and A. Miyake, "Quantum supremacy in constant-time measurement-based computation: A unified architecture for sampling and verification," (2017), arXiv:1703.11002.

[24] J. Bermejo-Vega, D. Hangleiter, M. Schwarz, R. Raussendorf, and J. Eisert, "Architectures for Quantum Simulation Showing a Quantum Speedup," Phys. Rev. X **8**, 021010 (2018).

[25] T. Morimae, "Hardness of classically sampling the one-clean-qubit model with constant total variation distance error," Phys. Rev. A **96**, 040302 (2017).

[26] A. Harrow and S. Mehraban, "Approximate unitary t-designs by short random quantum circuits using nearest-neighbor and long-range gates," (2018), arXiv:1809.06957 [quant-ph].

[27] J. Haferkamp, D. Hangleiter, A. Bouland, B. Fefferman, J. Eisert, and J. Bermejo-Vega, "Closing Gaps of a Quantum Advantage with Short-Time Hamiltonian Dynamics," Phys. Rev. Lett. **125**, 250501 (2020).

[28] A. M. Dalzell, N. Hunter-Jones, and F. G. S. L. Brandão, "Random quantum circuits anti-concentrate in log depth," (2020), arXiv:2011.12277 [cond-mat, physics:quant-ph].

[29] S. Nezami, "Permanent of random matrices from representation theory: Moments, numerics, concentration, and comments on hardness of boson-sampling," (2021), arXiv:2104.06423 [quant-ph].

[30] A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, "On the complexity and verification of quantum random circuit sampling," Nat. Phys. **15**, 159–163 (2019).

[31] R. Movassagh, "Quantum supremacy and random circuits," (2019), arXiv:1909.06210 [cond-mat, physics:hep-th, physics:math-ph, physics:quant-ph].

[32] A. Bouland, B. Fefferman, Z. Landau, and Y. Liu, "Noise and the frontier of quantum supremacy," (2021), arXiv:2102.01738 [quant-ph].

[33] 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," Nature **574**, 505–510 (2019).

[34] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, P. Hu, X.-Y. Yang, W.-J. Zhang, H. Li, Y. Li, X. Jiang, L. Gan, G. Yang, L. You, Z. Wang, L. Li, N.-L. Liu, C.-Y. Lu, and J.-W. Pan, "Quantum computational advantage using photons," Science **370**, 1460–1463 (2020).

[35] A. Arkhipov, "Boson Sampling is Robust to Small Errors in the Network Matrix," Phys. Rev. A **92**, 062326 (2015).

[36] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, "Characterizing quantum supremacy in near-term devices," Nat. Phys. **14**, 595–600 (2018).

[37] S. Aaronson and L. Chen, "Complexity-theoretic Foundations of Quantum Supremacy Experiments," in *Proceedings of the 32Nd Computational Complexity Conference*, CCC '17 (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Germany, 2017) pp. 22:1–22:67.

[38] S. Aaronson and S. Gunn, "On the classical hardness of spoofing linear cross-entropy benchmarking," Theory Comput. **16**, 1–8 (2020).

[39] B. Barak, C.-N. Chou, and X. Gao, "Spoofing Linear Cross-Entropy Benchmarking in Shallow Quantum Circuits," (2020), arXiv:2005.02421 [quant-ph].

[40] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, "Exponential improvement in precision for simulating sparse Hamiltonians," Proc. 46th Annu. ACM Symp. Theory Comput. - STOC 14 , 283–292 (2014).

[41] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, "Simulating Hamiltonian Dynamics with a Truncated Taylor Series," Phys. Rev. Lett. **114**, 090502 (2015).

[42] D. W. Berry, A. M. Childs, and R. Kothari, "Hamiltonian Simulation with Nearly Optimal Dependence on all Parameters," in *2015 IEEE 56th Annual Symposium on Foundations of Computer Science* (IEEE, Berkeley, CA, USA, 2015) pp. 792–809.

[43] G. H. Low and I. L. Chuang, "Optimal Hamiltonian Simulation by Quantum Signal Processing," Physical Review Letters **118**, 010501 (2017).

[44] J. Haah, M. B. Hastings, R. Kothari, and G. H. Low, "Quantum algorithm for simulating real time evolution of lattice Hamiltonians," in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, Paris, 2018) pp. 350–360.

[45] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, "Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics," (2018), arXiv:1806.01838 [quant-ph].

[46] G. H. Low and I. L. Chuang, "Hamiltonian Simulation by Qubitization," Quantum **3**, 163 (2019).

[47] A. M. Childs and Y. Su, "Nearly Optimal Lattice Simulation by Product Formulas," Phys. Rev. Lett. **123**, 050503 (2019).

[48] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, "Theory of Trotter Error with Commutator Scaling," Phys. Rev. X **11**, 011020 (2021).

[49] P. W. Anderson, "Absence of Diffusion in Certain Random Lattices," Phys. Rev. **109**, 1492–1505 (1958).

[50] L. Valiant, "Quantum Circuits That Can Be Simulated Classically in Polynomial Time," SIAM J. Comput. **31**, 1229–1254 (2002).

[51] B. M. Terhal and D. P. DiVincenzo, "Classical simulation of noninteracting-fermion quantum circuits," Phys. Rev. A **65**, 032325 (2002).

[52] R. Jozsa and A. Miyake, "Matchgates and classical simulation of quantum circuits," Proc. R. Soc. Math. Phys. Eng. Sci. **464**, 3089–3106 (2008).

[53] C. S. Hamilton, R. Kruse, L. Sansoni, S. Barkhofen, C. Silberhorn, and I. Jex, "Gaussian Boson Sampling," Phys. Rev. Lett. **119**, 170501 (2017).

[54] R. Kruse, C. S. Hamilton, L. Sansoni, S. Barkhofen, C. Silberhorn, and I. Jex, "A detailed study of Gaussian Boson Sampling," Phys. Rev. A **100**, 032326 (2019).

[55] D. Poulin and P. Wocjan, "Sampling from the Thermal Quantum Gibbs State and Evaluating Partition Functions with a Quantum Computer," Phys. Rev. Lett. **103**, 220502 (2009).

[56] C.-F. Chiang and P. Wocjan, "Quantum Algorithm for Preparing Thermal Gibbs States - Detailed Analysis," NATO Sci. Peace Secur. Ser. Inf. Commun. Secur. , 138–147 (2010).

[57] F. G. Brandao and K. M. Svore, "Quantum Speed-Ups for Solving Semidefinite Programs," in *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)* (2017) pp. 415–426.

[58] S. Arora, E. Hazan, and S. Kale, "The Multiplicative Weights Update Method: A Meta-Algorithm and Applications," Theory of Computing **8**, 121–164 (2012).

[59] A. Kitaev, A. Shen, and M. Vyalyi, *Classical and Quantum Computation*, Graduate Studies in Mathematics, Vol. 47 (American Mathematical Society, Providence, Rhode Island, 2002).

[60] R. Oliveira and B. M. Terhal, "The complexity of quantum spin systems on a two-dimensional square lattice," Quantum Inf. Comput. **8**, 0900–0924 (2008).

[61] D. Aharonov, D. Gottesman, S. Irani, and J. Kempe, "The power of quantum systems on a line," Commun. Math. Phys. **287**, 41–65 (2009).

[62] S. Bravyi, D. P. DiVincenzo, R. I. Oliveira, and B. M. Terhal, "The Complexity of Stoquastic Local Hamiltonian Problems," Quantum Inf. Comput. **8**, 0361–0385 (2008).

[63] D. Gottesman and S. Irani, "The Quantum and Classical Complexity of Translationally Invariant Tiling and Hamiltonian Problems," (2009), arXiv:0905.2419 [quant-ph].

[64] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, "A variational eigenvalue solver on a photonic quantum processor," Nat. Commun. **5**, 4213 (2014).

[65] J. Preskill, "Quantum computing and the entanglement frontier," (2012), arXiv:1203.5813 [cond-mat, physics:quant-ph].

[66] R. Jozsa and M. V. den Nest, "Classical simulation complexity of extended Clifford circuits," (2013), arXiv:1305.6190 [quant-ph].

[67] A. W. Harrow and A. Montanaro, "Quantum computational supremacy," Nature **549**, 203–209 (2017).

[68] A. Neville, C. Sparrow, R. Clifford, E. Johnston, P. M. Birchall, A. Montanaro, and A. Laing, "Classical boson sampling algorithms with superior performance to near-term experiments," Nat. Phys. **13**, 1153–1157 (2017).

[69] P. Clifford and R. Clifford, "The Classical Complexity of Boson Sampling," in *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '18 (Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2018) pp. 146–155.

[70] A. M. Dalzell, A. W. Harrow, D. E. Koh, and R. L. La Placa, "How many qubits are needed for quantum computational supremacy?" Quantum **4**, 264 (2020).

[71] S. Kirkpatrick and B. Selman, "Critical Behavior in the Satisfiability of Random Boolean Expressions," Science **264**, 1297–1301 (1994).

[72] C. R. Laumann, R. Moessner, A. Scardicchio, and S. L. Sondhi, "Random Quantum Satisfiabiilty," Quantum Inf. Comput. **10**, 1–15 (2010).

[73] K. P. Seshadreesan, J. P. Olson, K. R. Motes, P. P. Rohde, and J. P. Dowling, "Boson sampling with displaced single-photon Fock states versus single-photon-added coherent states: The quantum-classical divide and computational-complexity transitions in linear optics," Phys. Rev. A **91**, 022334 (2015).

[74] M. Heyl, A. Polkovnikov, and S. Kehrein, "Dynamical Quantum Phase Transitions in the Transverse-Field Ising Model," Phys. Rev. Lett. **110**, 135704 (2013).

[75] M. Heyl, "Dynamical quantum phase transitions: A review," (2017), arXiv:1709.07461 [cond-mat, physics:quant-ph].

[76] E. B. Rozenbaum, S. Ganeshan, and V. Galitski, "Lyapunov Exponent and Out-of-Time-Ordered Correlator's Growth Rate in a Chaotic System," Phys. Rev. Lett. **118**, 086801 (2017).

[77] E. H. Lieb and D. W. Robinson, "The finite group velocity of quantum spin systems," Commun. Math. Phys. **28**, 251–257 (1972).

[78] S. Aaronson, "A linear-optical proof that the permanent is #P-hard," Proc. R. Soc. Math. Phys. Eng. Sci. **467**, 3393–3405 (2011).

[79] M. Friesdorf, A. H. Werner, W. Brown, V. B. Scholz, and J. Eisert, "Many-body localisation implies that eigenvectors are matrix-product states," Phys. Rev. Lett. **114**, 170505 (2015).

[80] Y. Huang, "Efficient simulation of many-body localized systems," (2015), arXiv:1508.04756 [cond-mat, physics:quant-ph].

[81] F. Pollmann, V. Khemani, J. I. Cirac, and S. L. Sondhi, "Efficient variational diagonalization of fully many-body localized Hamiltonians," Phys. Rev. B **94**, 041116 (2016).

[82] X. Yu, D. Pekker, and B. K. Clark, "Finding matrix product state representations of highly-excited eigenstates of many-body localized Hamiltonians," Phys. Rev. Lett. **118**, 017201 (2017).

[83] A. Chapman and A. Miyake, "Classical simulation of quantum circuits by dynamical localization: Analytic results for Pauli-observable propagation in time-dependent disorder," (2017), arXiv:1704.04405 [cond-mat, physics:quant-ph].

[84] J. Eisert and D. Gross, "Supersonic Quantum Communication," Phys. Rev. Lett. **102**, 240501 (2009).

[85] M. B. Hastings, "Locality in Quantum Systems," (2010), arXiv:1008.5137 [math-ph, quant-ph].

[86] J. Fröhlich, F. Martinelli, E. Scoppola, and T. Spencer, "Constructive proof of localization in the Anderson tight binding model," Commun. Math. Phys. **101**, 21–46 (1985).

[87] E. Hamza, R. Sims, and G. Stolz, "Dynamical Localization in Disordered Quantum Spin Systems," Commun. Math. Phys. **315**, 215–239 (2012).

[88] M. Reck, A. Zeilinger, H. J. Bernstein, and P. Bertani, "Experimental realization of any discrete unitary operator," Phys. Rev. Lett. **73**, 58–61 (1994).

[89] W. R. Clements, P. C. Humphreys, B. J. Metcalf, W. S. Kolthammer, and I. A. Walsmley, "Optimal design for universal multiport interferometers," Optica **3**, 1460 (2016).

[90] M. Jerrum, A. Sinclair, and E. Vigoda, "A Polynomial-time Approximation Algorithm for the Permanent of a Matrix with Nonnegative Entries," J. ACM **51**, 671–697 (2004).

[91] F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, and B. V. Saunders, "NIST Digital Library of Mathematical Functions," .

[92] A. Deshpande, B. Fefferman, M. C. Tran, M. Foss-Feig, and A. V. Gorshkov, unpublished .

[93] N. Maskara, A. Deshpande, M. C. Tran, A. Ehrenberg, B. Fefferman, and A. V. Gorshkov, "Complexity phase diagram for interacting and long-range bosonic Hamiltonians," (2019), arXiv:1906.04178.

[94] D. J. Brod, "Complexity of simulating constant-depth BosonSampling," Phys. Rev. A **91**, 042316 (2015).

[95] G. Muraleedharan, A. Miyake, and I. H. Deutsch, "Quantum computational supremacy in the sampling of bosonic random walkers on a one-dimensional lattice," New J. Phys. **21**, 055003 (2018).

[96] A. Deshpande, B. Fefferman, M. C. Tran, M. Foss-Feig, and A. V. Gorshkov, "Dynamical Phase Transitions in Sampling Complexity," Phys. Rev. Lett. **121**, 030501 (2018).

[97] S. Banerjee and E. Altman, "Solvable model for a dynamical quantum phase transition from fast to slow scrambling," Phys. Rev. B **95**, 134302 (2017).

[98] R.-Q. He and Z.-Y. Lu, "Characterizing many-body localization by out-of-time-ordered correlation," Phys. Rev. B **95**, 054201 (2017).

[99] S. V. Syzranov, A. V. Gorshkov, and V. M. Galitski, "Interaction-induced transition in the quantum chaotic dynamics of a disordered metal," (2017), arXiv:1709.09296 [cond-mat, physics:quant-ph].

[100] P. Hosur, X.-L. Qi, D. A. Roberts, and B. Yoshida, "Chaos in quantum channels," J. High Energ. Phys. **2016**, 4 (2016).

[101] D. A. Roberts and B. Yoshida, "Chaos and complexity by design," J. High Energ. Phys. **2017**, 121 (2017).

[102] Y. Huang, Y.-L. Zhang, and X. Chen, "Out-of-time-ordered correlators in many-body localized systems," Ann. Phys. **529**, 1600318 (2017).

[103] R. Fan, P. Zhang, H. Shen, and H. Zhai, "Out-of-time-order correlation for many-body localization," Sci. Bull. **62**, 707–711 (2017).

[104] Y. Chen, "Universal Logarithmic Scrambling in Many Body Localization," (2016), arXiv:1608.02765 [cond-mat].

[105] B. Swingle and D. Chowdhury, "Slow scrambling in disordered quantum systems," Phys. Rev. B **95**, 060201 (2017).

[106] S. Sachdev and J. Ye, "Gapless spin-fluid ground state in a random quantum Heisenberg magnet," Phys. Rev. Lett. **70**, 3339–3342 (1993).

[107] A. Kitaev, "A simple model of quantum holography," (2015).

[108] J. Maldacena and D. Stanford, "Remarks on the Sachdev-Ye-Kitaev model," Phys. Rev. D **94**, 106002 (2016).

[109] J. Maldacena, S. H. Shenker, and D. Stanford, "A bound on chaos," J. High Energy Phys. **2016**, 106 (2016).

[110] Y. Sekino and L. Susskind, "Fast scramblers," J. High Energy Phys. **2008**, 065–065 (2008).

[111] N. Lashkari, D. Stanford, M. Hastings, T. Osborne, and P. Hayden, "Towards the fast scrambling conjecture," J. High Energy Phys. **2013**, 22 (2013).

[112] B. Swingle, "Black holes and the limits of quantum information processing," XRDS Crossroads ACM Mag. Stud. **23**, 52–56 (2016).

[113] A. R. Brown, D. A. Roberts, L. Susskind, B. Swingle, and Y. Zhao, "Holographic Complexity Equals Bulk Action?" Phys. Rev. Lett. **116**, 191301 (2016).

[114] A. R. Brown, D. A. Roberts, L. Susskind, B. Swingle, and Y. Zhao, "Complexity, action, and black holes," Phys. Rev. D **93**, 086006 (2016).

[115] S. Aaronson and D. J. Brod, "BosonSampling with Lost Photons," Phys. Rev. A **93**, 012335 (2016).

[116] A. R. Brown and L. Susskind, "Second law of quantum complexity," Phys. Rev. D **97**, 086015 (2018).

[117] U. Marzolino and T. Prosen, "Computational complexity of nonequilibrium steady states of quantum spin chains," Phys. Rev. A **93**, 032306 (2016).

[118] L. Amico, R. Fazio, A. Osterloh, and V. Vedral, "Entanglement in many-body systems," Rev. Mod. Phys. **80**, 517–576 (2008).

[119] B. Zeng, X. Chen, D.-L. Zhou, and X.-G. Wen, "Quantum Information Meets Quantum Matter – From Quantum Entanglement to Topological Phase in Many-Body Systems," (2015), arXiv:1508.02595 [cond-mat, physics:quant-ph].

[120] J. H. Bardarson, F. Pollmann, and J. E. Moore, "Unbounded Growth of Entanglement in Models of Many-Body Localization," Phys. Rev. Lett. **109**, 017202 (2012).

[121] Z.-C. Yang, A. Hamma, S. M. Giampaolo, E. R. Mucciolo, and C. Chamon, "Entanglement complexity in quantum many-body dynamics, thermalization, and localization," Phys. Rev. B **96**, 020408 (2017).

[122] O. Morsch and M. Oberthaler, "Dynamics of Bose-Einstein condensates in optical lattices," Rev. Mod. Phys. **78**, 179–215 (2006).

[123] I. Bloch, J. Dalibard, and S. Nascimbène, "Quantum simulations with ultracold quantum gases," Nat. Phys. **8**, 267–276 (2012).

[124] I. Buluta and F. Nori, "Quantum Simulators," Science **326**, 108–111 (2009).

[125] M. H. Devoret and R. J. Schoelkopf, "Superconducting Circuits for Quantum Information: An Outlook," Science **339**, 1169–1174 (2013).

[126] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, "A quantum gas microscope for detecting single atoms in a Hubbard-regime optical lattice," Nature **462**, 74 (2009).

[127] W. S. Bakr, A. Peng, M. E. Tai, R. Ma, J. Simon, J. I. Gillen, S. Foelling, L. Pollet, and M. Greiner, "Probing the Superfluid to Mott Insulator Transition at the Single Atom Level," Science **329**, 547–550 (2010).

[128] J. F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr, "Single-atom-resolved fluorescence imaging of an atomic Mott insulator," Nature **467**, 68 (2010).

[129] L. Mazza, D. Rossini, R. Fazio, and M. Endres, "Detecting two-site spin-entanglement in many-body systems with local particle-number fluctuations," New J. Phys. **17**, 013015 (2015).

[130] M. Miranda, R. Inoue, Y. Okuyama, A. Nakamoto, and M. Kozuma, "Site-resolved imaging of ytterbium atoms in a two-dimensional optical lattice," Phys. Rev. A **91**, 063414 (2015).

[131] R. Yamamoto, J. Kobayashi, K. Kato, T. Kuno, Y. Sakura, and Y. Takahashi, "Site-resolved imaging of single atoms with a Faraday quantum gas microscope," Phys. Rev. A **96**, 033610 (2017).

[132] C. Weitenberg, M. Endres, J. F. Sherson, M. Cheneau, P. Schauß, T. Fukuhara, I. Bloch, and S. Kuhr, "Single-spin addressing in an atomic Mott insulator," Nature **471**, 319 (2011).

[133] Y. Wang, X. Zhang, T. A. Corcovilos, A. Kumar, and D. S. Weiss, "Coherent Addressing of Individual Neutral Atoms in a 3D Optical Lattice," Phys. Rev. Lett. **115**, 043003 (2015).

[134] T. Fukuhara, A. Kantian, M. Endres, M. Cheneau, P. Schauß, S. Hild, D. Bellem, U. Schollwöck, T. Giamarchi, C. Gross, I. Bloch, and S. Kuhr, "Quantum dynamics of a single, mobile spin impurity," Nat. Phys. **9**, 235–241 (2013).

[135] S. Paul, P. R. Johnson, and E. Tiesinga, "Hubbard model for ultracold bosonic atoms interacting via zero-point-energy\char21{}induced three-body interactions," Phys. Rev. A **93**, 043616 (2016).

[136] K. Xu, T. Mukaiyama, J. R. Abo-Shaeer, J. K. Chin, D. E. Miller, and W. Ketterle, "Formation of Quantum-Degenerate Sodium Molecules," Phys. Rev. Lett. **91**, 210402 (2003).

[137] J. Herbig, "Preparation of a Pure Molecular Quantum Gas," Science **301**, 1510–1513 (2003).

[138] S. Dürr, T. Volz, A. Marte, and G. Rempe, "Observation of Molecules Produced from a Bose-Einstein Condensate," Phys. Rev. Lett. **92**, 020406 (2004).

[139] A. M. Kaufman, B. J. Lester, C. M. Reynolds, M. L. Wall, M. Foss-Feig, K. R. A. Hazzard, A. M. Rey, and C. A. Regal, "Hong-Ou-Mandel atom interferometry in tunnel-coupled optical tweezers," Science **345**, 306–309 (2014).

[140] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletic, M. Greiner, and M. D. Lukin, "Atom-by-atom assembly of defect-free one-dimensional cold atom arrays," Science **354**, 1024–1027 (2016).

[141] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, "Probing many-body dynamics on a 51-atom quantum simulator," Nature **551**, 579–584 (2017).

[142] E. Flurin, V. V. Ramasesh, S. Hacohen-Gourgy, L. S. Martin, N. Y. Yao, and I. Siddiqi, "Observing Topological Invariants Using Quantum Walk in Superconducting Circuits," Phys. Rev. X **7**, 031023 (2017).

[143] S. Hacohen-Gourgy, V. V. Ramasesh, C. De Grandi, I. Siddiqi, and S. M. Girvin, "Cooling and Autonomous Feedback in a Bose-Hubbard chain with Attractive Interactions," Phys. Rev. Lett. **115**, 240501 (2015).

[144] X.-H. Deng, C.-Y. Lai, and C.-C. Chien, "Superconducting circuit simulator of Bose-Hubbard model with a flat band," Phys. Rev. B **93**, 054116 (2016).

[145] P. Roushan, C. Neill, J. Tangpanitanon, V. M. Bastidas, A. Megrant, R. Barends, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. Fowler, B. Foxen, M. Giustina, E. Jeffrey, J. Kelly, E. Lucero, J. Mutus, M. Neeley, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. White, H. Neven, D. G. Angelakis, and J. Martinis, "Spectral signatures of many-body localization with interacting photons," Science **358**, 1175–1179 (2017).

[146] Y. Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, A. Megrant, J. Y. Mutus, P. J. J. O'Malley, C. M. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, M. R. Geller, A. N. Cleland, and J. M. Martinis, "Qubit architecture with high coherence and fast tunable coupling," Phys. Rev. Lett. **113**, 220502 (2014).

[147] S. Aaronson and D. Gottesman, "Improved Simulation of Stabilizer Circuits," Phys. Rev. A **70**, 052328 (2004).

[148] X. Ni and M. V. den Nest, "Commuting quantum circuits: Efficient classical simulations versus hardness results," (2012), arXiv:1204.4570 [quant-ph].

[149] S. Lloyd, "Almost Any Quantum Logic Gate is Universal," Phys. Rev. Lett. **75**, 346–349 (1995).

[150] D. Deutsch, A. Barenco, and A. Ekert, "Universality in Quantum Computation," Proc. R. Soc. Math. Phys. Eng. Sci. **449**, 669–677 (1995).

[151] M. J. Bremner, C. M. Dawson, J. L. Dodd, A. Gilchrist, A. W. Harrow, D. Mortimer, M. A. Nielsen, and T. J. Osborne, "Practical Scheme for Quantum Computation with Any Two-Qubit Entangling Gate," Phys. Rev. Lett. **89**, 247902 (2002).

[152] M. J. Bremner, A. Montanaro, and D. J. Shepherd, "Achieving quantum supremacy with sparse and noisy commuting quantum computations," Quantum **1**, 8 (2017).

[153] A. M. Childs, D. Leung, L. Mančinska, and M. Ozols, "Characterization of universal two-qubit Hamiltonians," Quantum Inf. Comput. **11**, pp0019–0039 (2011).

[154] A. Bouland, L. Mančinska, and X. Zhang, "Complexity classification of two-qubit commuting hamiltonians," in *31st Conference on Computational Complexity (CCC 2016)*, Leibniz International Proceedings in Informatics (LIPIcs) (2016) pp. 28:1–28:33.

[155] D. Hangleiter, J. Bermejo-Vega, M. Schwarz, and J. Eisert, "Anticoncentration theorems for schemes showing a quantum speedup," Quantum **2**, 65 (2018).

[156] R. Raussendorf, C. Okay, D.-S. Wang, D. T. Stephen, and H. P. Nautrup, "Computationally Universal Phase of Quantum Matter," Phys. Rev. Lett. **122**, 090501 (2019).

[157] E. Mossel, D. Weitz, and N. Wormald, "On the hardness of sampling independent sets beyond the tree threshold," Probab. Theory Relat. Fields **143**, 401–439 (2009).

[158] S. Vijay, "Measurement-Driven Phase Transition within a Volume-Law Entangled Phase," (2020), arXiv:2005.03052 [cond-mat, physics:hep-th, physics:quant-ph].

[159] M. A. Norcia, A. W. Young, and A. M. Kaufman, "Microscopic Control and Detection of Ultracold Strontium in Optical-Tweezer Arrays," Phys. Rev. X **8**, 041054 (2018).

[160] C. Neill, P. Roushan, K. Kechedzhi, S. Boixo, S. V. Isakov, V. Smelyanskiy, A. Megrant, B. Chiaro, A. Dunsworth, K. Arya, R. Barends, B. Burkett, Y. Chen, Z. Chen, A. Fowler, B. Foxen, M. Giustina, R. Graff, E. Jeffrey, T. Huang, J. Kelly, P. Klimov, E. Lucero, J. Mutus, M. Neeley, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, H. Neven, and J. M. Martinis, "A blueprint for demonstrating quantum supremacy with superconducting qubits," Science **360**, 195–199 (2018).

[161] M. Saffman, T. G. Walker, and K. Mølmer, "Quantum information with Rydberg atoms," Rev. Mod. Phys. **82**, 2313–2363 (2010).

[162] J. W. Britton, B. C. Sawyer, A. C. Keith, C.-C. J. Wang, J. K. Freericks, H. Uys, M. J. Biercuk, and J. J. Bollinger, "Engineered two-dimensional Ising interactions in a trapped-ion quantum simulator with hundreds of spins," Nature **484**, 489–492 (2012).

[163] N. Yao, L. Jiang, A. Gorshkov, P. Maurer, G. Giedke, J. Cirac, and M. Lukin, "Scalable architecture for a room temperature solid-state quantum information processor," Nat. Commun. **3**, 1–8 (2012).

[164] B. Yan, S. A. Moses, B. Gadway, J. P. Covey, K. R. A. Hazzard, A. M. Rey, D. S. Jin, and J. Ye, "Observation of dipolar spin-exchange interactions with lattice-confined polar molecules," Nature **501**, 521–525 (2013).

[165] J. S. Douglas, H. Habibian, C.-L. Hung, A. V. Gorshkov, H. J. Kimble, and D. E. Chang, "Quantum many-body models with cold atoms coupled to photonic crystals," Nat. Photonics **9**, 326–331 (2015).

[166] A. M. Childs, D. Gosset, and Z. Webb, "Universal Computation by Multiparticle Quantum Walk," Science **339**, 791–794 (2013).

[167] S. Janson, T. Luczak, and A. Ruciński, *Random Graphs*, Wiley-Interscience Series in Discrete Mathematics and Optimization (John Wiley, New York, 2000).

[168] M. B. Hastings and T. Koma, "Spectral Gap and Exponential Decay of Correlations," Commun. Math. Phys. **265**, 781–804 (2006).

[169] Z.-X. Gong, M. Foss-Feig, S. Michalakis, and A. V. Gorshkov, "Persistence of locality in systems with power-law interactions," Phys. Rev. Lett. **113**, 030602 (2014).

[170] M. Foss-Feig, Z.-X. Gong, C. W. Clark, and A. V. Gorshkov, "Nearly Linear Light Cones in Long-Range Interacting Quantum Systems," Phys. Rev. Lett. **114**, 157201 (2015).

[171] M. C. Tran, A. Y. Guo, Y. Su, J. R. Garrison, Z. Eldredge, M. Foss-Feig, A. M. Childs, and A. V. Gorshkov, "Locality and digital quantum simulation of power-law interactions," Phys. Rev. X **9**, 031006 (2019).

[172] A. Y. Guo, M. C. Tran, A. M. Childs, A. V. Gorshkov, and Z.-X. Gong, "Signaling and Scrambling with Strongly Long-Range Interactions," (2019), arXiv:1906.02662 [quant-ph].

[173] M. C. Tran, C.-F. Chen, A. Ehrenberg, A. Y. Guo, A. Deshpande, Y. Hong, Z.-X. Gong, A. V. Gorshkov, and A. Lucas, "Hierarchy of linear light cones with long-range interactions," Phys. Rev. X **10**, 031009 (2020).

[174] B. Peropadre, A. Aspuru-Guzik, and J. J. García-Ripoll, "Equivalence between spin Hamiltonians and boson sampling," Phys. Rev. A **95**, 032327 (2017).

[175] M. S. Underwood and D. L. Feder, "Bose-Hubbard model for universal quantum walk-based computation," Phys. Rev. A **85**, 052314 (2012).

[176] T. J. Osborne, "Efficient Approximation of the Dynamics of One-Dimensional Quantum Spin Systems," Phys. Rev. Lett. **97**, 157202 (2006).

[177] R. García-Patrón, J. J. Renema, and V. Shchesnovich, "Simulating boson sampling in lossy architectures," Quantum **3**, 169 (2019).

[178] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, "Computational Complexity of Projected Entangled Pair States," Phys. Rev. Lett. **98**, 140506 (2007).

[179] J. Haferkamp, D. Hangleiter, J. Eisert, and M. Gluza, "Contracting projected entangled pair states is average-case hard," Phys. Rev. Research **2**, 013010 (2020).

[180] D. Gottesman, "The Heisenberg Representation of Quantum Computers," (1998), arXiv:quant-ph/9807006.

[181] F. Verstraete and J. I. Cirac, "Mapping local Hamiltonians of fermions to local Hamiltonians of spins," J. Stat. Mech. **2005**, P09012–P09012 (2005).

[182] S.-K. Chu, G. Zhu, J. R. Garrison, Z. Eldredge, A. V. Curiel, P. Bienias, I. B. Spielman, and A. V. Gorshkov, "Scale-Invariant Continuous Entanglement Renormalization of a Chern Insulator," Phys. Rev. Lett. **122**, 120502 (2019).

[183] B. Gadway, "An atom optics approach to studying lattice transport phenomena," Phys. Rev. A **92**, 043606 (2015).

[184] G. M. Crosswhite, A. C. Doherty, and G. Vidal, "Applying matrix product operators to model systems with long-range interactions," Phys. Rev. B **78**, 035116 (2008).

[185] D. Barredo, S. de Léséleuc, V. Lienhard, T. Lahaye, and A. Browaeys, "An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays," Science **354**, 1021–1023 (2016).

[186] S. Korenblit, D. Kafri, W. C. Campbell, R. Islam, E. E. Edwards, Z.-X. Gong, G.-D. Lin, L.-M. Duan, J. Kim, K. Kim, and C. Monroe, "Quantum simulation of spin models on an arbitrary lattice with trapped ions," New J. Phys. **14**, 095024 (2012).

[187] R. Islam, C. Senko, W. C. Campbell, S. Korenblit, J. Smith, A. Lee, E. E. Edwards, C.-C. J. Wang, J. K. Freericks, and C. Monroe, "Emergence and Frustration of Magnetism with Variable-Range Interactions in a Quantum Simulator," Science **340**, 583–587 (2013).

[188] M. Oszmaniec and Z. Zimborás, "Universal extensions of restricted classes of quantum operations," Phys. Rev. Lett. **119**, 220502 (2017).

[189] A. Bouland and M. Ozols, "Trading Inverses for an Irrep in the Solovay-Kitaev Theorem," in *Proceedings of the 13th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2018)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 111 (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2018) pp. 6:1–6:15.

[190] Z. Eldredge, Z.-X. Gong, J. T. Young, A. H. Moosavian, M. Foss-Feig, and A. V. Gorshkov, "Fast Quantum State Transfer and Entanglement Renormalization Using Long-Range Interactions," Phys. Rev. Lett. **119**, 170503 (2017).

[191] D. J. Brod and A. M. Childs, "The computational power of matchgates and the XY interaction on arbitrary graphs," Quantum Inf. Comput. **14**, 0901–0916 (2013).

[192] A. Lakshminarayan, S. Tomsovic, O. Bohigas, and S. N. Majumdar, "Extreme statistics of complex random and quantum chaotic states," Phys. Rev. Lett. **100**, 044103 (2008).

[193] Y. Li, X. Chen, and M. P. A. Fisher, "Quantum Zeno effect and the many-body entanglement transition," Phys. Rev. B **98**, 205136 (2018).

[194] B. Skinner, J. Ruhman, and A. Nahum, "Measurement-Induced Phase Transitions in the Dynamics of Entanglement," Phys. Rev. X **9**, 031009 (2019).

[195] A. Chan, R. M. Nandkishore, M. Pretko, and G. Smith, "Unitary-projective entanglement dynamics," Phys. Rev. B **99**, 224307 (2019).

[196] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, 2010).

[197] S. Lloyd, "Universal Quantum Simulators," Science **273**, 1073–1078 (1996).

[198] E. Knill, "Fermionic linear optics and matchgates," arXiv:quant-ph/0108033 (2001).

[199] M. Wimmer, "Algorithm 923: Efficient numerical computation of the pfaffian for dense and banded skew-symmetric matrices," ACM T. Math. Software **38**, 30 (2012).

[200] H.-P. Breuer and F. Petruccione, *The theory of open quantum systems* (Oxford University Press, 2007).

[201] T. Prosen, "Spectral theorem for the lindblad equation for quadratic open fermionic systems," J. Stat. Mech.-Theory. E. **2010**, P07020 (2010).

[202] S. Bravyi and R. König, "Classical simulation of dissipative fermionic linear optics," Quantum Information & Computation **12**, 925–943 (2012).

[203] C. Gardiner, P. Zoller, and P. Zoller, *Quantum noise*, Vol. 56 (Springer, 2004).

[204] A. Chenu, M. Beau, J. Cao, and A. del Campo, "Quantum simulation of generic many-body open system dynamics using classical noise," Phys. Rev. Lett. **118**, 140403 (2017).

[205] B. Misra and E. G. Sudarshan, "The zeno's paradox in quantum theory," J. Math. Phys. **18**, 756–763 (1977).

[206] A. Degasperis, L. Fonda, and G. C. Ghirardi, "Does the lifetime of an unstable system depend on the measuring apparatus?" Il Nuovo Cimento A (1965-1970) **21**, 471–484 (1974).

[207] J. D. Franson, B. C. Jacobs, and T. B. Pittman, "Quantum computing using single photons and the Zeno effect," Phys. Rev. A **70**, 062302 (2004).

[208] Y.-Z. Sun, Y.-P. Huang, and P. Kumar, "Photonic nonlinearities via quantum zeno blockade," Phys. Rev. Lett. **110**, 223901 (2013).

[209] A. J. Daley, M. M. Boyd, J. Ye, and P. Zoller, "Quantum computing with alkaline-earth-metal atoms," Phys. Rev. Lett. **101**, 170504 (2008).

[210] M. Anderlini, P. J. Lee, B. L. Brown, J. Sebby-Strabley, W. D. Phillips, and J. V. Porto, "Controlled exchange interaction between pairs of neutral atoms in an optical lattice," Nature **448**, 452–456 (2007).

[211] P. J. Lee, M. Anderlini, B. L. Brown, J. Sebby-Strabley, W. D. Phillips, and J. V. Porto, "Sublattice addressing and spin-dependent motion of atoms in a double-well lattice," Phys. Rev. Lett. **99**, 020402 (2007).

[212] C. Zhang, S. L. Rolston, and S. Das Sarma, "Manipulation of single neutral atoms in optical lattices," Phys. Rev. A **74**, 042316 (2006).

[213] A. D. Ludlow, N. D. Lemke, J. A. Sherman, C. W. Oates, G. Quéméner, J. von Stecher, and A. M. Rey, "Cold-collision-shift cancellation and inelastic scattering in a yb optical lattice clock," Phys. Rev. A **84**, 052724 (2011).

[214] R. Zhang, Y. Cheng, H. Zhai, and P. Zhang, "Orbital feshbach resonance in alkali-earth atoms," Phys. Rev. Lett. **115**, 135301 (2015).

[215] L. Riegger, N. Darkwah Oppong, M. Höfer, D. R. Fernandes, I. Bloch, and S. Fölling, "Localized magnetic moments with tunable spin exchange in a gas of ultracold fermions," Phys. Rev. Lett. **120**, 143601 (2018).

[216] F. Scazza, C. Hofrichter, M. Höfer, P. C. De Groot, I. Bloch, and S. Fölling, "Observation of two-orbital spin-exchange interactions with ultracold su(n)-symmetric fermions," Nat. Phys. **10**, 779–784 (2014).

[217] R. Le Targat, X. Baillard, M. Fouché, A. Brusch, O. Tcherbakoff, G. D. Rovera, and P. Lemonde, "Accurate optical lattice clock with ^{87}Sr atoms," Phys. Rev. Lett. **97**, 130801 (2006).

[218] G. Cappellini, L. F. Livi, L. Franchi, D. Tusi, D. Benedicto Orenes, M. Inguscio, J. Catani, and L. Fallani, "Coherent manipulation of orbital feshbach molecules of two-electron atoms," Phys. Rev. X **9**, 011028 (2019).

[219] N. Syassen, D. M. Bauer, M. Lettner, T. Volz, D. Dietze, J. J. García-Ripoll, J. I. Cirac, G. Rempe, and S. Dürr, "Strong dissipation inhibits losses and induces correlations in cold molecular gases," Science **320**, 1329–1331 (2008).

[220] C. Chin, R. Grimm, P. Julienne, and E. Tiesinga, "Feshbach resonances in ultracold gases," Rev. Mod. Phys. **82**, 1225–1286 (2010).

[221] S. Giorgini, L. P. Pitaevskii, and S. Stringari, "Theory of ultracold atomic fermi gases," Rev. Mod. Phys. **80**, 1215–1274 (2008).

[222] G. Barontini, R. Labouvie, F. Stubenrauch, A. Vogler, V. Guarrera, and H. Ott, "Controlling the dynamics of an open many-body quantum system with localized dissipation," Phys. Rev. Lett. **110**, 035302 (2013).

[223] R. Labouvie, B. Santra, S. Heun, S. Wimberger, and H. Ott, "Negative differential conductivity in an interacting quantum gas," Phys. Rev. Lett. **115**, 050601 (2015).

[224] R. Labouvie, B. Santra, S. Heun, and H. Ott, "Bistability in a driven-dissipative superfluid," Phys. Rev. Lett. **116**, 235302 (2016).

[225] H. P. Lüschen, P. Bordia, S. S. Hodgman, M. Schreiber, S. Sarkar, A. J. Daley, M. H. Fischer, E. Altman, I. Bloch, and U. Schneider, "Signatures of many-body localization in a controlled open quantum system," Phys. Rev. X **7**, 011034 (2017).

[226] Y. S. Patil, S. Chakram, and M. Vengalattore, "Measurement-induced localization of an ultracold lattice gas," Phys. Rev. Lett. **115**, 140402 (2015).

[227] J. Li, A. K. Harter, J. Liu, L. de Melo, Y. N. Joglekar, and L. Luo, "Observation of parity-time symmetry breaking transitions in a dissipative floquet system of ultracold atoms," Nat. Commun. **10**, 855 (2019).

[228] R. Bonifacio, F. Casagrande, and M. Milani, "Superradiance and superfluorescence in josephson junction arrays," Phys. Lett. A **101**, 427–431 (1984).

[229] M. C. Cassidy, A. Bruno, S. Rubbert, M. Irfan, J. Kammhuber, R. N. Schouten, A. R. Akhmerov, and L. P. Kouwenhoven, "Demonstration of an ac josephson junction laser," Science **355**, 939–942 (2017).

[230] P. E. Dolgirev, J. Marino, D. Sels, and E. Demler, "Non-gaussian correlations imprinted by local dephasing in fermionic wires," Phys. Rev. B **102**, 100301 (2020).

[231] T. E. Lee and C.-K. Chan, "Heralded magnetism in non-hermitian atomic systems," Phys. Rev. X **4**, 041001 (2014).

[232] T. E. Lee, S. Gopalakrishnan, and M. D. Lukin, "Unconventional magnetism via optical pumping of interacting spin systems," Phys. Rev. Lett. **110**, 257204 (2013).

[233] C. Joshi, F. Nissen, and J. Keeling, "Quantum correlations in the one-dimensional driven dissipative xy model," Phys. Rev. A **88**, 063835 (2013).

[234] S. Diehl, W. Yi, A. J. Daley, and P. Zoller, "Dissipation-induced d-wave pairing of fermionic atoms in an optical lattice," Phys. Rev. Lett. **105**, 227001 (2010).

[235] K. Yamamoto, M. Nakagawa, K. Adachi, K. Takasan, M. Ueda, and N. Kawakami, "Theory of non-hermitian fermionic superfluidity with a complex-valued interaction," Phys. Rev. Lett. **123**, 123601 (2019).

[236] J. A. S. Lourenço, R. L. Eneias, and R. G. Pereira, "Kondo effect in a $\mathcal{PT}$-symmetric non-hermitian hamiltonian," Phys. Rev. B **98**, 085126 (2018).

[237] M. Nakagawa, N. Kawakami, and M. Ueda, "Non-hermitian kondo effect in ultracold alkaline-earth atoms," Phys. Rev. Lett. **121**, 203001 (2018).

[238] M. S. Rudner and L. S. Levitov, "Topological transition in a non-hermitian quantum walk," Phys. Rev. Lett. **102**, 065703 (2009).

[239] T. E. Lee, "Anomalous edge state in a non-hermitian lattice," Phys. Rev. Lett. **116**, 133903 (2016).

[240] D. Leykam, K. Y. Bliokh, C. Huang, Y. D. Chong, and F. Nori, "Edge modes, degeneracies, and topological numbers in non-hermitian systems," Phys. Rev. Lett. **118**, 040401 (2017).

[241] S. Yao and Z. Wang, "Edge states and topological invariants of non-hermitian systems," Phys. Rev. Lett. **121**, 086803 (2018).

[242] Z. Gong, Y. Ashida, K. Kawabata, K. Takasan, S. Higashikawa, and M. Ueda, "Topological phases of non-hermitian systems," Phys. Rev. X **8**, 031079 (2018).

[243] H. Shen, B. Zhen, and L. Fu, "Topological band theory for non-hermitian hamiltonians," Phys. Rev. Lett. **120**, 146402 (2018).

[244] K. Kawabata, K. Shiozaki, M. Ueda, and M. Sato, "Symmetry and topology in non-hermitian physics," Phys. Rev. X **9**, 041015 (2019).

[245] T. Yoshida, K. Kudo, and Y. Hatsugai, "Non-hermitian fractional quantum hall states," Scientific Reports **9**, 16895 (2019).

[246] K. B. Efetov, "Directed quantum chaos," Phys. Rev. Lett. **79**, 491–494 (1997).

[247] I. Y. Goldsheid and B. A. Khoruzhenko, "Distribution of eigenvalues in non-hermitian anderson models," Phys. Rev. Lett. **80**, 2897–2900 (1998).

[248] S. Longhi, D. Gatti, and G. Della Valle, "Non-hermitian transparency and one-way transport in low-dimensional lattices by an imaginary gauge field," Phys. Rev. B **92**, 094204 (2015).

[249] B. Horstmann, J. I. Cirac, and G. Giedke, "Noise-driven dynamics and phase transitions in fermionic systems," Phys. Rev. A **87**, 012108 (2013).

[250] M. Žnidarič, "Exact solution for a diffusive nonequilibrium steady state of an open quantum chain," J. Stat. Mech.–Theory E. **2010**, L05002 (2010).

[251] V. Eisler, "Crossover between ballistic and diffusive transport: the quantum exclusion process," J. Stat. Mech.–Theory E. **2011**, P06007 (2011).

[252] S. Bravyi, B. M. Terhal, and B. Leemhuis, "Majorana fermion codes," New J. Phys. **12**, 083039 (2010).

[253] O. Shtanko, A. Deshpande, P. S. Julienne, and A. V. Gorshkov, "Limits on Classical Simulation of Free Fermions with Dissipation," (2020), arXiv:2005.10840 [cond-mat, physics:physics, physics:quant-ph].

[254] G. Lindblad, "On the generators of quantum dynamical semigroups," Comm. Math. Phys. **48**, 119–130 (1976).

[255] E.-M. Laine, J. Piilo, and H.-P. Breuer, "Measure for the non-markovianity of quantum processes," Phys. Rev. A **81**, 062115 (2010).

[256] J. Colpa, "Diagonalisation of the quadratic fermion hamiltonian with a linear part," J. Phys. A-Math. Gen. **12**, 469 (1979).

[257] P. Aliferis and B. M. Terhal, "Fault-tolerant quantum computation for local leakage faults," Quantum Information & Computation **7**, 139–156 (2007).

[258] D. Aharonov and M. Ben-Or, "Fault-tolerant quantum computation with constant error," in *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing - STOC '97* (ACM Press, El Paso, Texas, United States, 1997) pp. 176–188.

[259] E. Knill, R. Laflamme, and W. H. Zurek, "Resilient quantum computation," Science **279**, 342–345 (1998).

[260] B. M. Terhal, "Quantum error correction for quantum memories," Rev. Mod. Phys. **87**, 307–346 (2015).

[261] M. Hebenstreit, R. Jozsa, B. Kraus, S. Strelchuk, and M. Yoganathan, "All pure fermionic non-Gaussian states are magic states for matchgate computations," Phys. Rev. Lett. **123**, 080503 (2019).

[262] T. Brydges, A. Elben, P. Jurcevic, B. Vermersch, C. Maier, B. P. Lanyon, P. Zoller, R. Blatt, and C. F. Roos, "Probing Rényi entanglement entropy via randomized measurements," Science **364**, 260–263 (2019).

[263] V. Alba and F. Carollo, "Spreading of correlations in Markovian open quantum systems," (2020), arXiv:2002.09527 [cond-mat, physics:hep-th, physics:quant-ph].

[264] D. Poulin, "Lieb-Robinson Bound and Locality for General Markovian Quantum Dynamics," Phys. Rev. Lett. **104**, 190401 (2010).

[265] T. Barthel and M. Kliesch, "Quasilocality and Efficient Simulation of Markovian Quantum Dynamics," Phys. Rev. Lett. **108**, 230504 (2012).

[266] 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 [quant-ph].

[267] G. Kalai and G. Kindler, "Gaussian noise sensitivity and BosonSampling," (2014), arXiv:1409.3093.

[268] J. J. Renema, "Marginal probabilities in boson samplers with arbitrary input states," (2020), arXiv:2012.14917 [quant-ph].

[269] S. Aaronson, "Quantum supremacy, now with bosonsampling," (2020).

[270] S. Aaronson, "Chinese bosonsampling experiment: the gloves are off," (2020).

[271] S. J. Freedman and J. F. Clauser, "Experimental test of local hidden-variable theories," Phys. Rev. Lett. **28**, 938–941 (1972).

[272] A. Aspect, P. Grangier, and G. Roger, "Experimental Realization of Einstein-Podolsky-Rosen-Bohm Gedankenexperiment: A New Violation of Bell's Inequalities," Phys. Rev. Lett. **49**, 91–94 (1982).

[273] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. Vermeulen, R. N. Schouten, C. Abellán, *et al.*, "Loophole-free bell inequality violation using electron spins separated by 1.3 kilometres," Nature **526**, 682–686 (2015).

[274] M. Giustina, M. A. Versteegh, S. Wengerowsky, J. Handsteiner, A. Hochrainer, K. Phelan, F. Steinlechner, J. Kofler, J.-Å. Larsson, C. Abellán, *et al.*, "Significant-loophole-free test of bell's theorem with entangled photons," Phys. Rev. Lett. **115**, 250401 (2015).

[275] L. K. Shalm, E. Meyer-Scott, B. G. Christensen, P. Bierhorst, M. A. Wayne, M. J. Stevens, T. Gerrits, S. Glancy, D. R. Hamel, M. S. Allman, *et al.*, "Strong loophole-free test of local realism," Phys. Rev. Lett. **115**, 250402 (2015).

[276] W. Rosenfeld, D. Burchardt, R. Garthoff, K. Redeker, N. Ortegel, M. Rau, and H. Weinfurter, "Event-ready Bell test using entangled atoms simultaneously closing detection and locality loopholes," Phys. Rev. Lett. **119**, 010402 (2017).

[277] J. Haferkamp, D. Hangleiter, A. Bouland, B. Fefferman, J. Eisert, and J. Bermejo-Vega, "Closing gaps of a quantum advantage with short-time Hamiltonian dynamics," Phys. Rev. Lett. **125**, 250501 (2020).

[278] M. Oszmaniec, N. Dangniam, M. E. S. Morales, and Z. Zimborás, "Fermion Sampling: A robust quantum computational advantage scheme using fermionic linear optics and magic input states," (2020), arXiv:2012.15825 [math-ph, physics:quant-ph].

[279] J. E. Bourassa, R. N. Alexander, M. Vasmer, A. Patil, I. Tzitrin, T. Matsuura, D. Su, B. Q. Baragiola, S. Guha, G. Dauphinais, K. K. Sabapathy, N. C. Menicucci, and I. Dhand, "Blueprint for a scalable photonic fault-tolerant quantum computer," Quantum **5**, 392 (2021).

[280] S. Bartolucci, P. Birchall, H. Bombin, H. Cable, C. Dawson, M. Gimeno-Segovia, E. Johnston, K. Kieling, N. Nickerson, M. Pant, F. Pastawski, T. Rudolph, and C. Sparrow, "Fusion-based quantum computation," (2021), arXiv:2101.09310 [quant-ph].

[281] T. R. Bromley, J. M. Arrazola, S. Jahangiri, J. Izaac, N. Quesada, A. D. Gran, M. Schuld, J. Swinarton, Z. Zabaneh, and N. Killoran, "Applications of near-term

photonic quantum computers: software and algorithms," Quantum Sci. Technol. **5**, 034010 (2020).

[282] H. Qi, D. J. Brod, N. Quesada, and R. García-Patrón, "Regimes of classical simulability for noisy Gaussian boson sampling," Phys. Rev. Lett. **124**, 100502 (2020).

[283] A. Björklund, B. Gupt, and N. Quesada, "A faster hafnian formula for complex matrices and its benchmarking on a supercomputer," ACM J. Exp. Algorithmics **24**, 1–17 (2019).

[284] A. Deshpande, A. Mehta, T. Vincent, N. Quesada, M. Hinsche, M. Ioannou, L. Madsen, J. Lavoie, H. Qi, J. Eisert, D. Hangleiter, B. Fefferman, and I. Dhand, "Quantum Computational Supremacy via High-Dimensional Gaussian Boson Sampling," (2021), arXiv:2102.12474 [quant-ph].

[285] R. J. Lipton, "New directions in testing," in *Distributed Computing and Cryptography, Proceedings of a DIMACS Workshop, Princeton, New Jersey, USA, October 4-6, 1989*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Vol. 2, edited by J. Feigenbaum and M. Merritt (DIMACS/AMS, 1989) pp. 191–202.

[286] Y. Kondo, R. Mori, and R. Movassagh, "Fine-grained analysis and improved robustness of quantum supremacy for random circuit sampling," (2021), arXiv:2102.01960 [quant-ph].

[287] K. Fujii, "Noise Threshold of Quantum Supremacy," (2016), arXiv:1610.03632 [quant-ph].

[288] M. Oszmaniec and D. J. Brod, "Classical simulation of photonic linear optics with lost particles," New J. Phys. **20**, 092002 (2018).

[289] S. Rahimi-Keshari, A. Scherer, A. Mann, A. T. Rezakhani, A. I. Lvovsky, and B. C. Sanders, "Quantum process tomography with coherent states," New Journal of Physics **13**, 013006 (2011).

[290] S. Boixo, V. N. Smelyanskiy, and H. Neven, "Fourier analysis of sampling from noisy chaotic quantum circuits," (2017), arXiv:1708.01875 [quant-ph].

[291] L. Valiant, "The complexity of computing the permanent," Theor. Comput. Sci. **8**, 189–201 (1979).

[292] D. Hangleiter, M. Kliesch, J. Eisert, and C. Gogolin, "Sample Complexity of Device-Independently Certified "Quantum Supremacy"," Phys. Rev. Lett. **122**, 210502 (2019).

[293] C. Oh, Y. Lim, B. Fefferman, and L. Jiang, "Classical simulation of bosonic linear-optical random circuits beyond linear light cone," (2021), arXiv:2102.10083 [quant-ph].

[294] S. Gharibian, Y. Huang, Z. Landau, and S. W. Shin, "Quantum Hamiltonian Complexity," Found. Trends® Theor. Comput. Sci. **10**, 159–282 (2015).

[295] S. A. Cook, "The complexity of theorem-proving procedures," in *Proceedings of the Third Annual ACM Symposium on Theory of Computing - STOC '71* (ACM Press, Shaker Heights, Ohio, United States, 1971) pp. 151–158.

[296] B. Trakhtenbrot, "A Survey of Russian Approaches to Perebor (Brute-Force Searches) Algorithms," IEEE Annals Hist. Comput. **6**, 384–400 (1984).

[297] D. Aharonov and T. Naveh, "Quantum NP - A Survey," (2002), arXiv:quant-ph/0210077.

[298] D. Aharonov, I. Arad, Z. Landau, and U. Vazirani, "The detectability lemma and quantum gap amplification," in *Proceedings of the 41st Annual ACM Symposium on Symposium on Theory of Computing - STOC '09* (ACM Press, Bethesda, MD, USA, 2009) p. 417.

[299] D. Aharonov, I. Arad, and T. Vidick, "Guest column: The quantum PCP conjecture," SIGACT News **44**, 47–79 (2013).

[300] S. Aaronson and G. Kuperberg, "Quantum versus classical proofs and advice," Theory Comput. **3**, 129–157 (2007).

[301] B. Fefferman and S. Kimmel, "Quantum vs. Classical proofs and subset verification," in *43rd International Symposium on Mathematical Foundations of Computer Science, MFCS 2018, August 27-31, 2018, Liverpool, UK*, LIPIcs, Vol. 117, edited by I. Potapov, P. G. Spirakis, and J. Worrell (Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018) pp. 22:1–22:23.

[302] F. D. M. Haldane, "Nonlinear Field Theory of Large-Spin Heisenberg Antiferromagnets: Semiclassically Quantized Solitons of the One-Dimensional Easy-Axis Néel State," Phys. Rev. Lett. **50**, 1153–1156 (1983).

[303] E. Witten, "Physical law and the quest for mathematical understanding," Bull. Am. Math. Soc. **40**, 21–30 (2002).

[304] T. S. Cubitt, D. Perez-Garcia, and M. M. Wolf, "Undecidability of the spectral gap," Nature **528**, 207–211 (2015).

[305] M. B. Hastings, "An area law for one-dimensional quantum systems," J. Stat. Mech. **2007**, P08024 (2007).

[306] Z. Landau, U. Vazirani, and T. Vidick, "A polynomial time algorithm for the ground state of one-dimensional gapped local Hamiltonians," Nat. Phys. **11**, 566–569 (2015).

[307] I. Arad, Z. Landau, U. Vazirani, and T. Vidick, "Rigorous RG algorithms and area laws for low energy eigenstates in 1D," Commun. Math. Phys. **356**, 65–105 (2017).

[308] D. Aharonov, M. Ben-Or, F. G. S. L. Brandao, and O. Sattath, "The Pursuit For Uniqueness: Extending Valiant-Vazirani Theorem to the Probabilistic and Quantum Settings," (2008), arXiv:0810.4840.

[309] R. Jain, I. Kerenidis, G. Kuperberg, M. Santha, O. Sattath, and S. Zhang, "On the Power of a Unique Quantum Witness," Theory Comput. **8**, 375–400 (2012).

[310] C. E. González-Guillén and T. S. Cubitt, "History-state Hamiltonians are critical," (2018), arXiv:1810.06528 [quant-ph].

[311] E. Crosson and J. Bowen, "Quantum ground state isoperimetric inequalities for the energy spectrum of local Hamiltonians," (2018), arXiv:1703.10133 [quant-ph].

[312] A. M. Dalzell and F. G. S. L. Brandao, "Locally accurate MPS approximations for ground states of one-dimensional gapped local Hamiltonians," Quantum **3**, 187 (2019).

[313] R. Babbush, D. W. Berry, I. D. Kivlichan, A. Y. Wei, P. J. Love, and A. Aspuru-Guzik, "Exponentially more precise quantum simulation of fermions in second quantization," New J. Phys. **18**, 033032 (2016).

[314] A. M. Childs, R. Kothari, and R. D. Somma, "Quantum Algorithm for Systems of Linear Equations with Exponentially Improved Dependence on Precision," SIAM J. Comput. **46**, 1920–1950 (2017).

[315] M. Aguilera Sammaritano, M. González Vera, P. Marcelo Cometto, T. Nicola Tejero, G. F. Bauerfeldt, and A. Mellouki, "Temperature dependence of rate coefficients for the gas phase reaction of OH with 3-chloropropene. A theoretical and experimental study," Chem. Phys. Lett. , 137757 (2020).

[316] C. Hubig, J. Haegeman, and U. Schollwöck, "Error estimates for extrapolations with matrix-product states," Phys. Rev. B **97**, 045125 (2018).

[317] S. Sachdev, *Quantum Phase Transitions*, 2nd ed. (Cambridge University Press, Cambridge, 2011).

[318] Y. Atia and D. Aharonov, "Fast-forwarding of Hamiltonians and Exponentially Precise Measurements," Nat. Commun. **8**, 1572 (2017).

[319] B. Fefferman and C. Y.-Y. Lin, "A Complete Characterization of Unitary Quantum Space," in *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 94, edited by A. R. Karlin (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2018) pp. 4:1–4:21.

[320] A. Kitaev and J. Watrous, "Parallelization, amplification, and exponential time simulation of quantum interactive proof systems," in *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing - STOC '00* (ACM Press, Portland, Oregon, United States, 2000) pp. 608–617.

[321] M. N. Vyalyi, "QMA = PP implies that PP contains PH," in *ECCCTR: Electronic Colloquium on Computational Complexity, Technical Reports* (2003).

[322] C. Marriott and J. Watrous, "Quantum Arthur–Merlin games," Comput. Complex. **14**, 122–152 (2005).

[323] E. W. Allender and K. W. Wagner, "Counting Hierarchies: Polynomial Time And Constant Depth Circuits," in *Current Trends in Theoretical Computer Science* (World Scientific, 1993) pp. 469–483.

[324] S. Gharibian, M. Santha, J. Sikora, A. Sundaram, and J. Yirka, "Quantum Generalizations of the Polynomial Hierarchy with Applications to QMA(2)," in *43rd International Symposium on Mathematical Foundations of Computer Science (MFCS 2018)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 117 (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2018) pp. 58:1–58:16.

[325] T. Morimae and H. Nishimura, "Merlinization of complexity classes above BQP," Quantum Inf. Comput. **17**, 959–972 (2017).

[326] R. Jefferson and R. C. Myers, "Circuit complexity in quantum field theory," J. High Energ. Phys. **2017**, 107 (2017).

[327] C. D. White, C. Cao, and B. Swingle, "Conformal field theories are magical," (2020), arXiv:2007.01303 [cond-mat, physics:hep-th, physics:quant-ph].

[328] D. Aharonov and A. Green, "A quantum inspired proof of $\mathsf{P}^{\#\mathsf{P}} \subseteq \mathsf{IP}$," (2017), arXiv:1710.09078 [quant-ph].

[329] A. Green, G. Kindler, and Y. Liu, "Towards a quantum-inspired proof for IP = PSPACE," (2019), arXiv:1912.11611 [quant-ph].

[330] A. Shamir, "IP = PSPACE," J. ACM **39**, 869–877 (1992).

[331] J. Kempe and O. Regev, "3-local Hamitonian is QMA-complete," Quantum Inf. Comput. **3**, 258–264 (2003).

[332] J. Kempe, A. Kitaev, and O. Regev, "The Complexity of the Local Hamiltonian Problem," SIAM J. Comput. **35**, 1070–1097 (2006).

[333] R. V. Mises and H. Pollaczek-Geiringer, "Praktische Verfahren der Gleichungsauflösung," Z. Für Angew. Math. Mech. **9**, 152–164 (1929).

[334] L. M. Adleman, J. DeMarrais, and M.-D. A. Huang, "Quantum Computability," SIAM J. Comput. **26**, 1524–1540 (1997).

[335] S. Aaronson, "Quantum computing, postselection, and probabilistic polynomial-time," Proc. R. Soc. Math. Phys. Eng. Sci. **461**, 3473–3482 (2005).

[336] J. R. Schrieffer and P. A. Wolff, "Relation between the Anderson and Kondo Hamiltonians," Phys. Rev. **149**, 491–492 (1966).

[337] S. Bravyi, D. DiVincenzo, and D. Loss, "Schrieffer-Wolff transformation for quantum many-body systems," Ann. Phys. **326**, 2793–2826 (2011).

[338] B. Fefferman and C. Y.-Y. Lin, "Quantum Merlin Arthur with Exponentially Small Gap," (2016), arXiv:1601.01975 [quant-ph].

[339] N. Schuch, I. Cirac, and F. Verstraete, "Computational Difficulty of Finding Matrix Product Ground States," Phys. Rev. Lett. **100**, 250501 (2008).

[340] T. Jones, S. Endo, S. McArdle, X. Yuan, and S. C. Benjamin, "Variational quantum algorithms for discovering Hamiltonian spectra," Phys. Rev. A **99**, 062304 (2019).

[341] D. Aharonov, I. Arad, and S. Irani, "Efficient algorithm for approximating one-dimensional ground states," Phys. Rev. A **82**, 012315 (2010).

[342] A. Molnar, N. Schuch, F. Verstraete, and J. I. Cirac, "Approximating Gibbs states of local Hamiltonians efficiently with projected entangled pair states," Phys. Rev. B **91**, 045138 (2015).

[343] S. Knysh, "Zero-temperature quantum annealing bottlenecks in the spin-glass phase," Nat. Commun. **7**, 12370 (2016).

[344] N. Usher, M. J. Hoban, and D. E. Browne, "Non-Unitary Quantum Computation in the Ground Space of Local Hamiltonians," Phys. Rev. A **96**, 032321 (2017).

[345] D. Nagaj, P. Wocjan, and Y. Zhang, "Fast amplification of QMA," Quantum Inf. Comput. **9**, 1053–1068 (2009).

[346] S. Aaronson, "Quantum lower bound for the collision problem," in *Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing*, STOC '02 (Association for Computing Machinery, New York, NY, USA, 2002) pp. 635–642.

[347] S. Aaronson, "BQP and the polynomial hierarchy," in *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*, STOC '10 (Association for Computing Machinery, New York, NY, USA, 2010) pp. 141–150.

[348] R. Raz and A. Tal, *Oracle Separation of BQP and PH*, Tech. Rep. 107 (2018).

[349] S. Bravyi, D. P. DiVincenzo, D. Loss, and B. M. Terhal, "Simulation of Many-Body Hamiltonians using Perturbation Theory with Bounded-Strength Interactions," Phys. Rev. Lett. **101**, 070503 (2008).

[350] T. Cubitt and A. Montanaro, "Complexity classification of local Hamiltonian problems," SIAM J. Comput. **45**, 268–316 (2013).

[351] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev, "Adiabatic Quantum Computation is Equivalent to Standard Quantum Computation," SIAM J. Comput. **37**, 166–194 (2007).

[352] N. Schuch and F. Verstraete, "Computational Complexity of interacting electrons and fundamental limitations of Density Functional Theory," Nat. Phys. **5**, 732–735 (2007).

[353] T. S. Cubitt, A. Montanaro, and S. Piddock, "Universal quantum Hamiltonians," Proc. Natl. Acad. Sci. USA **115**, 9497–9502 (2018).

[354] J. Bausch, T. S. Cubitt, A. Lucia, D. Perez-Garcia, and M. M. Wolf, "Size-driven quantum phase transitions," Proc. Natl. Acad. Sci. **115**, 19–23 (2018).

[355] S. Piddock and A. Montanaro, "Universal qudit Hamiltonians," (2018), arXiv:1802.07130 [quant-ph].

[356] T. Kohler and T. Cubitt, "Toy models of holographic duality between local Hamiltonians," J. High Energ. Phys. **2019**, 17 (2019).

[357] T. Kohler, S. Piddock, J. Bausch, and T. Cubitt, "Translationally-Invariant Universal Quantum Hamiltonians in 1D," (2020), arXiv:2003.13753 [quant-ph].

[358] J. Bausch and E. Crosson, "Analysis and limitations of modified circuit-to-Hamiltonian constructions," Quantum **2**, 94 (2018).

[359] L. Caha, Z. Landau, and D. Nagaj, "Clocks in Feynman's computer and Kitaev's local Hamiltonian: Bias, gaps, idling, and pulse tuning," Phys. Rev. A **97**, 062306 (2018).

[360] J. D. Watson, "Detailed Analysis of Circuit-to-Hamiltonian Mappings," (2019), arXiv:1910.01481 [quant-ph].

[361] J. Watrous, "Quantum Computational Complexity," in *Encyclopedia of Complexity and Systems Science*, edited by R. A. Meyers (Springer New York, New York, NY, 2009) pp. 7174–7201.

[362] F. G. S. L. Brandao, "Entanglement Theory and the Quantum Simulation of Many-Body Physics," (2008), arXiv:0810.0026 [quant-ph].

[363] H. Blier and A. Tapp, "All Languages in NP Have Very Short Quantum Proofs," in *Proceedings of the 2009 Third International Conference on Quantum, Nano and Micro Technologies*, ICQNM '09 (IEEE Computer Society, USA, 2009) pp. 34–37.

[364] A. Pereszlényi, "Multi-Prover Quantum Merlin-Arthur Proof Systems with Small Gap," (2012), arXiv:1205.2761 [quant-ph].

[365] T. Ito, H. Kobayashi, and J. Watrous, "Quantum interactive proofs with weak error bounds," in *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS '12 (Association for Computing Machinery, Cambridge, Massachusetts, 2012) pp. 266–275.

[366] Z. Ji, "Compression of quantum multi-prover interactive proofs," in *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing - STOC 2017* (ACM Press, Montreal, Canada, 2017) pp. 289–302.

[367] J. Fitzsimons, Z. Ji, T. Vidick, and H. Yuen, "Quantum proof systems for iterated exponential time, and beyond," in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing - STOC 2019* (ACM Press, Phoenix, AZ, USA, 2019) pp. 473–480.

[368] B. Fefferman, H. Kobayashi, C. Y.-Y. Lin, T. Morimae, and H. Nishimura, "Space-Efficient Error Reduction for Unitary Quantum Computations," in *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, Vol. 55 (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2016) pp. 14:1–14:14.

[369] L. Valiant and V. Vazirani, "NP is as easy as detecting unique solutions," Theor. Comput. Sci. **47**, 85–93 (1986).

[370] A. Ambainis, "On Physical Problems that are Slightly More Difficult than QMA," in *Proceedings of the 2014 IEEE 29th Conference on Computational Complexity*, CCC '14 (IEEE Computer Society, Washington, DC, USA, 2014) pp. 32–43.

[371] S. Gharibian and J. Yirka, "The complexity of simulating local measurements on quantum systems," Quantum **3**, 189 (2019).

[372] S. Gharibian, S. Piddock, and J. Yirka, "Oracle complexity classes and local measurements on physical Hamiltonians," (2019), arXiv:1909.05981 [quant-ph].

[373] L. Novo, J. Bermejo-Vega, and R. García-Patrón, "Quantum advantage from energy measurements of many-body quantum systems," (2019), arXiv:1912.06608 [quant-ph].

[374] P. Wocjan, D. Janzing, and T. Beth, "Two QCMA-Complete problems," Quantum Inf. Comput. **3**, 635–643 (2003).

[375] C. H. Bennett, "Time/Space Trade-Offs for Reversible Computation," SIAM J. Comput. **18**, 766–776 (1989).

[376] K.-J. Lange, P. McKenzie, and A. Tapp, "Reversible Space Equals Deterministic Space," J. Comput. Syst. Sci. **60**, 354–367 (2000).

[377] S. P. Jordan, H. Kobayashi, D. Nagaj, and H. Nishimura, "Achieving perfect completeness in classical-witness quantum Merlin-Arthur proof systems," Quantum Inf. Comput. **12**, 461–471 (2012).

[378] S. Dooley, G. Kells, H. Katsura, and T. C. Dorlas, "Simulating quantum circuits by adiabatic computation: Improved spectral gap bounds," Phys. Rev. A **101**, 042302 (2020).

[379] S. Arora and B. Barak, *Computational Complexity: A Modern Approach* (Cambridge University Press, Cambridge ; New York, 2009).

[380] M. Sipser, *Introduction to the Theory of Computation*, 3rd ed. (Cengage Learning, Boston, MA, 2012).

[381] P. W. Shor, "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer," SIAM J. Comput. **26**, 1484–1509 (1997).