Fine-Grained Complexity for Quantum Problems from Size-Preserving Circuit-To-Hamiltonian Constructions

Nai-Hui Chia
Department of Computer Science, Rice University, Houston, TX, USA
Ken Kennedy Institute, Rice University, Houston, TX, USA
Smalley-Curl Institute, Rice University, Houston, TX, USA

Atsuya Hasegawa
Graduate School of Mathematics, Nagoya University, Japan

François Le Gall
Graduate School of Mathematics, Nagoya University, Japan

Yu-Ching Shen
Department of Computer Science, Houston, Rice University, Houston, TX, USA

Abstract

The local Hamiltonian (LH) problem is the canonical QMA-complete problem introduced by Kitaev. In this paper, we show its hardness in a very strong sense: we show that the 3-local Hamiltonian problem on n qubits cannot be solved classically in time $O(2^{(1-\varepsilon)n})$ for any $\varepsilon > 0$ under the Strong Exponential-Time Hypothesis (SETH), and cannot be solved quantumly in time $O(2^{(1-\varepsilon)n/2})$ for any $\varepsilon > 0$ under the Quantum Strong Exponential-Time Hypothesis (QSETH). These lower bounds give evidence that the currently known classical and quantum algorithms for LH cannot be significantly improved.

Furthermore, we are able to demonstrate fine-grained complexity lower bounds for approximating the quantum partition function (QPF) with an arbitrary constant relative error. Approximating QPF with relative error is known to be equivalent to approximately counting the dimension of the solution subspace of QMA problems. We show the SETH and QSETH hardness to estimate QPF with constant relative error. We then provide a quantum algorithm that runs in $O(\sqrt{2^n})$ time for an arbitrary $1/\mathrm{poly}(n)$ relative error, matching our lower bounds and improving the state-of-the-art algorithm by Bravyi, Chowdhury, Gosset, and Wocjan (Nature Physics 2022) in the low-temperature regime.

To prove our fine-grained lower bounds, we introduce the first size-preserving circuit-to-Hamiltonian construction that encodes the computation of a T-time quantum circuit acting on N qubits into a $(d+1)$-local Hamiltonian acting on $N + O(T^{1/d})$ qubits. This improves the standard construction based on the unary clock, which uses $N + O(T)$ qubits.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory

Keywords and phrases Fine-grain complexity, SETH, QSETH, Local Hamiltonian problem, Quantum partition problem

Digital Object Identifier 10.4230/LIPIcs.CCC.2026.12

Related Version *arXiv version*: https://arxiv.org/abs/2602.14379

Funding *Nai-Hui Chia*: NSF Award FET-2535028, NSF Award FET-2243659, NSF Career Award FET-2339116, Google Scholar Award, and DOE Quantum Testbed Finder Award DE-SC0024301.
Atsuya Hasegawa: JSPS KAKENHI grant No. 24H00071, 25K24674, 25K2446.
François Le Gall: JSPS KAKENHI grant No. 24H00071, 25K24674, 25K2446, MEXT Q-LEAP grant No. JPMXS0120319794, JST ASPIRE grant No. JPMJAP2302 and JST CREST grant No. JPMJCR24I4.
Yu-Ching Shen: NSF Award FET-2339116 and NSF Award FET-2243659.

41st Computational Complexity Conference (CCC 2026).
Editor: Dana Moshkovitz; Article No. 12; pp. 12:1–12:35

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Acknowledgements NHC and YCS thank Chunhao Wang and Christopher Ye for valuable discussions. FLG thanks Suguru Tamaki for discussions. The authors thank the anonymous referees for their comments.

1 Introduction

1.1 Background

Local Hamiltonian (LH) problem

Kitaev [34] introduced the k-local Hamiltonian (k-LH) problem, which asks to estimate the ground-state energy of a k-local Hamiltonian

$$H = \sum_{i=1}^{m} H_i \,,$$

acting on n qubits, where $m = \text{poly}(n)$ and $\|H\| \leq \text{poly}(n)$, with constant additive error.[1] Kitaev, at the same time, showed that the problem is QMA-complete, i.e., any efficient quantum verification procedure can be embedded into the local Hamiltonian problem. The embedding procedure is called the circuit-to-Hamiltonian construction, which is a quantum analog of the Cook-Levin theorem [16, 38]. The local Hamiltonian problem is in the interface between physics and computer science, and has been the central problem in quantum complexity theory. Subsequent work showed that, when we reduce the locality to 2, restrict the interaction of local terms to geometrically local, and consider some specific type of Hamiltonians, the problem is still QMA-hard [7, 17, 32, 45, 46].

The k-local Hamiltonian problem can be solved classically in $O^*(2^n)$ time using the power method or its variant the Lanczos method [36, 37], for $k = O(\log n)$.[2] On a quantum computer, we can use Grover search [25] combined with other techniques to achieve running time $O^*\left(2^{n/2}\right)$ [3, 22, 24, 33, 39, 42, 47].[3]

Approximating quantum partition function with relative error

Computing the quantum partition function is an even harder problem. The partition function is a function of the temperature and the Hamiltonian of the system. It is defined by

$$Z := \text{tr}[e^{-\beta H}],$$

where β is the inverse of the temperature and H is the Hamiltonian. The *Quantum Partition Function (QPF)* problem asks us to compute the value of Z for a given Hamiltonian at a specified temperature with relative error.[4] Intuitively, QPF is harder than LH because evaluating the partition function requires information about the *entire* spectrum of the Hamiltonian, whereas solving LH only involves the ground-state energy. In fact, we can

[1] Equivalently, we can define the local Hamiltonian problem as $||H|| \leq 1$ and $1/\text{poly}(n)$ additive error. Our definition is the same as [32].

[2] In this paper the notation $O^*(\cdot)$ suppresses $\text{poly}(n)$ factors.

[3] We note that some previous works showed how to obtain larger ("super-Grover") speedups for certain applications in combinatorial optimization [26, 18] or for computing a coarser approximation of the ground-state energy [11].

[4] It is crucial here to consider relative error estimation. When we consider the problem to estimate the (normalized) QPF with additive error $1/\text{poly}(n)$, it is known to be DQC1-complete [8, 15] where DQC1 is the set of problems that can be solved efficiently with only "one-clean qubit" [35].

show that LH reduces to approximating the QPF, even with a constant relative error, and thus approximating QPF with constant-relative error is QMA-hard.[5] Furthermore, Bravyi et al. [9] show that approximating the QPF up to constant-relative error is computationally equivalent to approximately counting the number of witnesses accepted by a QMA verifier.

For algorithms, Poulin and Wocjan [48] demonstrated a quantum algorithm that approximates the partition function Z of an n-qubit system up to relative error δ in $O^*\left(\frac{1}{\delta}\sqrt{\frac{2^n}{Z}}\right)$ time. Later Bravyi, Chowdhury, Gosset, and Wocjan [9] proposed another algorithm achieving the same running time while improving the space complexity.

When the temperature is not too low (i.e., β is not too large), Z is not exponentially small (for example, $Z \approx O(1)$ when $\beta \approx 0.7n$) and the running time is then roughly $O^*(2^{n/2})$.

Our main question

In summary, existing classical and quantum algorithms for LH run in $O^*(2^n)$ and $O^*(2^{n/2})$ time, respectively, suggesting a fundamental barrier to further improvements. The situation for QPF is even more mysterious: no $O(2^n)$-time classical algorithm is currently known to the best of our knowledge, and existing quantum algorithms fail to achieve the expected quadratic speedup in the low-temperature regime. These observations naturally motivate the following questions:

Do there exist classical or quantum algorithms that are significantly faster than $O(2^n)$ or $O(2^{n/2})$ for computing the ground-state energy? What is the best possible algorithms for approximating quantum partition function even in the low-temperature regime?

Fine-grained complexity, SETH and QSETH

Our goal is to provide a $\Omega(2^n)$ classical (resp. $\Omega(2^{n/2})$ quantum) lower bound. We are going to use the language of fine-grained complexity [49, 50, 43], which is a relatively recent field of computational complexity theory that investigates complexity in a more fine-grained way than the conventional class like P and NP.

The k-local Hamiltonian problem is a generalization of k-SAT, the central problem in classical complexity theory and its fine-grained complexity has been well studied. The Strong Exponential Time Hypothesis (SETH) [29] states that k-SAT needs roughly 2^n time for large k (see Conjecture 14).[6] More recently, a quantum analog was proposed in [1, 12] as Quantum Strong Exponential-Time Hypothesis (QSETH). The conjecture is to claim that any quantum algorithm to solve k-SAT needs $\Omega(2^{n/2})$ time (see Conjecture 15). SETH and QSETH are regarded as plausible computational conjectures and serve as a useful tool for deriving conditional lower bounds for various fundamental problems. First, the SAT problem has been studied for decades, and despite this extensive effort, the $\Omega(2^n)$ (resp. $\Omega(2^{n/2})$) lower bound for classical (resp. quantum) algorithms has not been broken; designing quantum algorithms that refute these conjectures would likely require fundamentally new techniques and could lead to major breakthroughs in algorithm design. In addition, (Q)SETH has been used to establish conditional lower bounds for a wide range of problems, such as the

[5] To the best of our knowledge, [9] first mentioned that approximating QPF with a relative error is QMA-hard; we provide a self-contained formal proof in Section 5.1.

[6] We stress that (Q)SETH are about the complexity of k-SAT for large k. For small values of k, an algorithms better than $O(2^{n/2})$ are known. For example, there exists a 1.32793^n-time classical algorithm for 3-SAT [40].

orthogonal vectors problem [6, 1], the closest pair problem [19, 31, 1, 12], the edit distance problem [5, 12], lattice problems [13, 28]. These results highlight the value of (Q)SETH as a powerful framework for studying the hardness of computational problems.

Barriers to fine-grained complexity lower bounds for $O(1)$-local Hamiltonians

We aim to prove the fine-grained complexity lower bound for LH and QPF assuming SETH and QSETH. One straightforward approach is to map the clauses in the k-SAT instance Φ to projectors that act on local qubits. The sum of these local projectors is a k-local Hamiltonian H whose ground subspace encodes the solutions for Φ (see Appendix A for the construction). In this reduction, the number of qubits of H is equal to the number of variables in Φ and therefore directly gives a $\Omega(2^n)$ (resp. $\Omega(2^{n/2})$) lower bound assuming SETH (resp. QSETH). However, the locality of the induced Hamiltonian is precisely k, where k is the parameter of the k-SAT hard instance; this value depends on ε, as characterized by SETH and QSETH. Hence, this approach fails to establish $\Omega(2^n)$ (resp. $\Omega(2^{n/2})$) lower bounds for Hamiltonians with fixed locality (e.g., 3-local).

Another approach is to use known circuit-to-Hamiltonian reductions, in which a verification circuit $U = U_T \cdots U_1$ for a QMA problem (e.g., a verification circuit of size $T = \mathrm{poly}(n)$ for a k-SAT instance) is embedded into a local Hamiltonian H. In particular, for any verification circuit, there are reductions to make the induced Hamiltonian 5-local [34] or 2-local [32], which achieves our purpose of fixed locality. Following the idea by Feynman [21], the standard circuit-to-Hamiltonian construction [34, 32] introduces a clock register, which is used to represent $T+1$ orthogonal clock states. The ground state of the resulting Hamiltonian is the history state

$$\frac{1}{\sqrt{T+1}} \sum_{i=0}^{T} U_i \cdots U_0 \left|\psi\right\rangle_{\mathrm{circuit}} \otimes \left|i\right\rangle_{\mathrm{clock}},$$

where $\left|\psi\right\rangle_{\mathrm{circuit}}$ denotes the initial state of the circuit (including the witness) and $\left|i\right\rangle_{\mathrm{clock}}$ denotes the clock states. However, this standard reduction results in significant size overhead due to the number of additional qubits needed to implement the clock states: in the standard reduction, a unary clock [34, 32] is used, which requires $O(T)$ qubits. Note that the analysis to obtain fine-grained lower bounds is very sensitive to the size overhead between the circuit and the local Hamiltonian. In particular, in order to obtain the desired lower bounds based on (Q)SETH, the clock need to be implementable in $o(n)$ qubits. Although we could use the binary clock [34] to reduce the size of the clock register from $O(T)$ qubits to $O(\log T)$ qubits, the locality would then increase to $O(\log T)$.

In summary, we require a k-SAT-to-LH reduction that simultaneously satisfies the following two conditions – a combination that existing methods fail to achieve:

1. *Size-preserving*: the resulting Hamiltonian acts on $n + o(n)$ qubits, and
2. *Locality-independent*: its locality is a constant independent of k.

1.2 Our results

In this paper, to investigate the fine-grained complexity for QMA-hard problems, we show the first size-preserving circuit-to-Hamiltonian construction. We denote by $\lambda(H)$ the smallest eigenvalue (ground-state energy) of a local Hamiltonian H.

▶ **Theorem 1** (Informal version of Theorem 17). *Let $L = (L_{\mathrm{yes}}, L_{\mathrm{no}})$ be a promise language in* QMA*, and $x \in \{0,1\}^n$ be an input for the language L. Let U_x be a verification circuit acting on N qubits and T be the number of elementary gates of U_x. Then, for any integer $d \geq 1$, there exists a $(d+1)$-local Hamiltonian H_x acting on $N + O(T^{1/d})$ qubits satisfying the following properties;*

- $||H_x|| \leq \mathrm{poly}(n)$,
- *if* $x \in L_{\mathrm{yes}}$, $\lambda(H_x) \leq E_{\mathrm{yes}}$,
- *if* $x \in L_{\mathrm{no}}$, $\lambda(H_x) \geq E_{\mathrm{no}}$,
- $E_{\mathrm{no}} - E_{\mathrm{yes}} \geq \Omega(1)$.

Our construction is the first construction to achieve both the $O(1)$-locality and the $o(T)$-qubit size for the clock register. See Table 1 for a comparison between our construction and the standard constructions.

Table 1 Comparison between our result and standard circuit-to-Hamiltonian constructions. Locality and Size indicate the locality and size of the local Hamiltonian reduced from a quantum circuit of T gates. N denotes the qubit size of verification quantum circuits. Our result holds for any integer $d \geq 1$.

Clock	Locality	Size
Binary [34]	$O(\log T)$	$N + O(\log T)$
Unary [34, 32]	2	$N + O(T)$
Our work (Theorem 17)	$d + 1$	$N + O(T^{1/d})$

We then show k-SAT has a very efficient verification procedure as a unitary circuit (Lemma 23). By combining this with our size-preserving circuit-to-Hamiltonian construction, we obtain the SETH- and QSETH-hardness for the local Hamiltonian problem.

▶ **Theorem 2** (Informal version of Theorem 22). *Assuming SETH (resp. QSETH), estimating the ground-state energy of a 3-local Hamiltonian problem with any constant additive error cannot be solved by any classical algorithm in time* $O^*\left(2^{(1-\varepsilon)n}\right)$ *(resp. by any quantum algorithm in time* $O^*\left(2^{(1-\varepsilon)n/2}\right)$*) for any constant* $\varepsilon > 0$.

Our lower bound for the local Hamiltonian problem matches the performance of the best known classical and quantum algorithms.

▶ Remark 3. Recently, Buhrman et al. [11] constructed classical and quantum algorithms running in time $2^{n(1-\Omega(\delta/M))}$ and $2^{n(1-\Omega(\delta/M))/2}$, respectively, where $M := \sum_i \|H_i\|$, to estimate the ground-state energy of a $O(1)$-local Hamiltonian up to an additive error δ. These upper bounds do not contradict our conditional lower bounds since Theorem 2 is about estimating the ground-state energy within constant additive error, and the upper bounds given by [11] become $O(2^{n(1-1/\mathrm{poly}(n))})$ and $O(2^{n(1-1/\mathrm{poly}(n))/2})$ in this regime.

We further investigate fine-grained complexity of quantum partition functions. We first show the SETH- and QSETH-hardness to estimate QPF with constant relative error.

▶ **Theorem 4** (Informal version of Theorem 26). *Assuming SETH (resp. QSETH), approximating the quantum partition function for 3-local Hamiltonians to within any constant relative error cannot be done by any classical algorithm in time* $O\left(2^{n(1-\varepsilon)}\right)$ *(resp. quantum algorithm in time* $O\left(2^{n(1-\varepsilon)/2}\right)$*) for any constant* $\varepsilon > 0$.

We then show a new algorithm for QPF.

▶ **Theorem 5** (Informal version of Theorem 27). *There exists a quantum algorithm whose running time is* $O^*\left(2^{n/2}\right)$ *to approximate the quantum partition function for a constant-local Hamiltonian on* n *qubits up to a relative error* $1/\mathrm{poly}(n)$.

We now compare our algorithm with the one proposed in [48] and [9]. The running time of their algorithm is $O^*\left(\frac{1}{\delta}\sqrt{\frac{2^n}{Z}}\right)$, which depends on the value of the partition function Z. When the inverse temperature $\beta = O(n^c)$ for some large constant c (i.e., at low temperatures), the running time may exceed $O(2^{n/2})$. In contrast, the running time of our algorithm does not depend on Z: it runs in $O^*(2^{n/2})$ for all $\beta = \mathrm{poly}(n)$.

1.3 Our techniques

Compressed clock state and how to use it for the circuit-to-Hamiltonian construction

To ensure that our reduction works, we need a construction of the clock state associated with the induced Hamiltonian that satisfies the following three conditions:

- **(i)** The clock state uses $o(n)$ of qubits and can encode $\mathrm{poly}(n)$ of computation steps,
- **(ii)** the operation on the clock state needs to be constant-local, with locality is independent of the circuit, and
- **(iii)** the computation process of U is encoded in the ground state of the induced Hamiltonian H.

The unary clock satisfies conditions (ii) and (iii), but does not satisfy condition (i). In [14], a clock construction is proposed that satisfies conditions (i) and (ii). To represent $O(a^d)$ time on an a-qubit register for integers $a \geq d \geq 1$, they use a Johnson graph where the vertices of the Johnson graph (a, d) are the d-element subset of an a-element set and two vertices are adjacent when the intersection of the two vertices (subsets) contains $(d-1)$-elements (see Definition 20 for a formal definition[7]). However, unlike the circuit-to-Hamiltonian constructions [34, 32], the circuit-to-Hamiltonian reduction in [14] encodes the computation in the *time evolution* of H, rather than in its ground state. As a result, there is no guarantee regarding the structure or properties of the ground state in their construction.

We build on the clock-state construction from [14] and a 2-local Hamiltonian construction from Section 5 in [32].[8] By introducing a carefully designed penalty term into H, we enforce that the computation history is stored in the ground state. This yields a novel circuit-to-Hamiltonian reduction that simultaneously preserves the size of the Hamiltonian, maintains constant locality, and ensures that the circuit U is encoded in the ground state of the resulting Hamiltonian H. Our construction satisfies all three requirements (i), (ii) and (iii), overcoming the limitations that neither of the existing reductions can address individually.

Additional techniques to reduce the locality

From the approach above, we have a $d+2$-local Hamiltonian from a quantum circuit whose number of gates is $O(a^d)$. This is because we need a local term operating on $d+1$ qubits in the clock register and 1 qubit in the circuit register for the propagation Hamiltonian.[9]

To reduce the locality further, we introduce a new clock construction that combines the $(a, d-1)$-clock state from [14] and the unary clock on a qubits. Crucial properties of the two clocks are:

---

[7] Not to confuse readers, in the introduction we denote by a the register size of the (a, d)-clock state and by n the input size.

[8] Section 6 of [32] introduces another approach to obtain a 2-local Hamiltonian using the perturbation theory.

[9] We do not need to consider 2-local gates thanks to the technique in Section 5 of [32].

- for the $(a, d-1)$ clock, time can be identified by looking at $(d-1)$ position ($(d-1)$-locality), while increasing or decreasing time by one unit can be done by modifying d positions (d-locality);
- for the unary clock, increasing or decreasing time by one unit can be done by modifying one position (1-locality). Moreover, if time is 0 or a, it can be identified by looking at 1 position (1-locality) (see Table 2 in Section 3.1).

By making the unary clock move periodically from 0 to a and from a to 0, we increase and decrease most of the clock time by d-local operators. By carefully analyzing the locality of all local Hamiltonian, we have a $(d+1)$-local Hamiltonian from a circuit of $O(a^d)$ gates using the $2a$-qubit clock register.

The remaining task is to construct a verification circuit that uses $o(n)$ ancilla qubits for any k-SAT instance whose number of variables is n. To achieve this, we use a counter to record the number of clauses satisfied by a given assignment. The k-SAT formula is satisfied if and only if the value stored in the counter equals m. Since counting up to m requires only $\log m$ qubits, the number of ancilla qubits needed is $n_a = O(\log m) = o(n)$, as required. The number of gates for the verification circuit is $O(n \log^2 n)$ (Lemma 23) and thus we can use $d = 2$ and obtain the SETH- and QSETH-hardness for the 3-local Hamiltonian problem.

Fine-grained complexity and quantum algorithms for the QPF problem

The lower bound for QPF follows directly from the lower bound for LH. When β is a sufficiently large polynomial, the partition function is dominated by low-energy states. If the ground-state energy is at most E_{yes}, then by ignoring all other states except the ground state, we obtain $Z \geq e^{-\beta E_{\text{yes}}}$. Conversely, if the ground-state energy is at least E_{no}, then by treating all eigenstates as having energy E_{no}, we get $Z \leq 2^n e^{-\beta E_{\text{no}}}$. Therefore, by approximating Z, we can distinguish between the two cases and decide the ground-state energy. To ensure this distinction, it suffices that $(1-\delta)e^{-\beta E_{\text{yes}}} > (1+\delta)2^n e^{-\beta E_{\text{no}}}$. In our reduction, we have observed that for any $\delta > 0$, this inequality holds for all sufficiently large n.

To obtain an $O(\sqrt{2^n})$-time quantum algorithm for the QPF with $1/\text{poly}(n)$ relative error, we adopt a technique from the proof of the equivalence between approximating the number of witnesses for a QMA verifier and approximating the quantum partition function [9]. Briefly, the approach involves partitioning the energy spectrum into polynomially many intervals and selecting a representative energy value $\mathcal{E}_j$ for each interval j. By estimating the number of eigenstates M_j within each interval, the partition function Z can be approximated as:

$$\tilde{Z} = \sum_j M_j e^{-\beta \mathcal{E}_j}.$$

We now explain how to count M_j. First, we construct an energy estimation circuit U_{EE} that outputs the corresponding eigenvalue E_p for each eigenstate $|\psi_p\rangle$ of H. In other words, $U_{EE}|\psi_p\rangle = |E_p\rangle|\psi_p\rangle$. The circuit U_{EE} can be implemented using Hamiltonian simulation together with phase estimation. Then, we apply a circuit U_{dec} that decides whether E_p is in the interval j. Let $U_j := U_{dec}U_{EE}$. Together, we have $U_j|\psi_p\rangle = |1\rangle|\phi_p\rangle$ if E_p in the interval j and $U_j|\psi_p\rangle = |0\rangle|\phi_p\rangle$ otherwise, where $|\phi_p\rangle$ is some quantum state corresponding to $|\psi_p\rangle$.

Now, let us suppose that the following two "too-good-to-be-true" assumptions hold:

1. A uniform superposition of all eigenstates can be efficiently prepared, and
2. the energy estimation circuit, U_{EE}, can compute eigenvalues with zero error.

Then, by applying U_j on the state, we obtain

$$U_j \sum_p \frac{1}{\sqrt{N}} |\psi_p\rangle = \sqrt{\frac{M_j}{N}} |1\rangle|\xi_1\rangle + \sqrt{\frac{N-M_j}{N}} |0\rangle|\xi_0\rangle, \tag{1}$$

where $N = 2^n$ is the number of the eigenstates and $|\xi_1\rangle$, $|\xi_0\rangle$ are two quantum states orthogonal to each other. Finally, by using the well-known amplitude estimation algorithm on Equation (1), we can obtain M_j. This algorithms runs in $O(\sqrt{2^n})$ time.

We then remove the two "too-good-to-be-true" assumptions while keeping the $O(\sqrt{2^n})$ runtime. First, directly preparing the uniform superposition of all eigenstates is generally computationally hard since the eigenstates are unknown. To overcome this difficulty, a commonly used technique is to use EPR state (e.g., [48, 33]). We have that uniform superposition over a complete basis, tensored with its complex conjugate, is equivalent to an EPR state. Therefore, by applying U_j to the EPR state, we obtain

$$U_j \sum_p \frac{1}{\sqrt{N}} |\psi_p\rangle |\psi_p^*\rangle = \sqrt{\frac{M_j}{N}} |1\rangle |\xi_1'\rangle + \sqrt{\frac{N - M_j}{N}} |0\rangle |\xi_0'\rangle , \tag{2}$$

where $|\xi_1'\rangle$ and $|\xi_0'\rangle$ are orthogonal quantum states. By applying the amplitude estimation algorithm to the state in Equation (2), we can approximate M_j.

Second, the quantum energy estimation circuit described above introduces an additive error of $1/\mathrm{poly}(n)$, which can significantly impact the approximation of the quantum partition function. Specifically, if the eigenstates are densely concentrated near the boundaries of the intervals in the selected partition (within the precision of the energy estimation), we may "doubly count" these states. Such inaccuracies in state counting translate directly into a constant relative error in the final estimation of the partition function.

To reduce the constant relative error to $1/\mathrm{poly}(n)$, we observe that the "double-counting" issue only arises when the chosen partition unluckily has significant fraction of eigenvalues lie near the boundaries of intervals. Intuitively, by considering polynomially many partitions with mutually disjoint boundaries, any specific eigenstate can be miscounted in at most a constant number of these partitions. Consequently, if we run the algorithm over all such partitions, an averaging argument ensures that at least one set will have a number of double-counted eigenstates that is at most an inverse polynomial fraction of the total. This approach yields an algorithm achieving $1/\mathrm{poly}(n)$ relative error in $O(\sqrt{2^n})$ time.

1.4 Open questions

1. It is known that the 2-local Hamiltonian problem is QMA-hard [32]. **Can we show a size-preserving reduction to 2-local Hamiltonian and its (Q)SETH-hardness?** In our clock construction, we need to touch two clocks and simultaneously apply a 1-qubit gate, so we cannot improve the locality directly. Note that perturbation theory does not seem to work here either: the standard perturbation technique adds an ancilla qubit for each term, and thus does not give a size-preserving construction when applied directly to our 3-local Hamiltonian.
2. Even when we consider geometrically local [45] or specific 2-local terms such as Heisenberg interaction, antiferromagnetic XY interactions [17, 46], 2-local ZX and $ZZXX$ Hamiltonian [7], the problem is still QMA-hard. This leads to a natural question: **Does the (Q)SETH-hardness still hold for geometrically local Hamiltonians whose each local term is restricted to a specific form?**
3. Our relative threshold gap $(E_{\mathrm{no}} - E_{\mathrm{yes}})/M$ in the 3-local Hamiltonian problem is $1/p(n)$ where $M := \sum_{i=i}^{m} \|H_i\|$ and $p(n)$ is a polynomial with large degree. A natural question arises: **when the relative gap increases, for example, $(E_{\mathrm{no}} - E_{\mathrm{yes}})/M = \Omega(1)$, what conditional lower bound can we have under plausible conjectures?** Notably, the algorithm of [11] runs in time $O^*(2^{\frac{n}{2}(1-\Omega(1))})$ when $(E_{\mathrm{no}} - E_{\mathrm{yes}})/M = \Omega(1)$, and therefore having the same lower bound $\Omega(2^{\frac{n}{2}(1-o(1))})$ for the $(E_{\mathrm{no}} - E_{\mathrm{yes}})/M = \Omega(1)$ regime as we have for $(E_{\mathrm{no}} - E_{\mathrm{yes}})/M = 1/p(n)$ in our work is impossible.

4. In this paper, we show the SETH-hardness of the 3-local Hamiltonian problem, which is a quantum analog of MAX-3-SAT, and it is known that there exists a faster exponential time algorithm for MAX-2-SAT, say $O^*(1.732^n)$-time algorithm by [51]. **Can we show the fine-grained complexity of MAX-3-SAT?**
5. In this work, we show that LH can be reduced to QPF, and the reduction is fine-grained even for an arbitrary constant relative error; therefore, we can obtain the $\Omega(2^{n/2})$ lower bound assuming (Q)SETH. Notably, the lower bound holds for any constant relative error and locality at least 3. Along this line, we can ask **whether the lower bound continues to hold when considering 2-local Hamiltonian, geometrical restrictions, and specified local terms, or we can design faster algorithms under these constraints.**
6. Another interesting question is **whether the reverse reduction holds, namely, whether approximating QPF can be reduced to the LH problem.** At first glance, this seems unlikely, as QPF can be interpreted as estimating the dimension of the ground space in the low-temperature regime. However, since we are considering exponential lower bounds, such a reduction could, in principle, be allowed to run in super-polynomial time.
7. Our lower bound for the 3QPF problem holds for $\beta = O(n)$ and $\|H\|$ being a large degree polynomial. A natural question to ask is **what happens as the temperature increases.** In other words, when $\beta\|H\|$ becomes small, does the lower bound still hold in this regime? Or, at what point do we observe a *phase transition*?

1.5 Organization

In Section 2, we give preliminaries for the rest of the paper. In Section 3, we show the size-preserving circuit-to-Hamiltonian construction by introducing a new clock. In Section 4, we show the SETH-hardness of the local Hamiltonian problem. In Section 5, we show the fine-grained complexity to estimate the quantum partition function with relative error.

2 Preliminaries

We refer to [44, 20] for references in quantum computing, [23] for a reference and survey in quantum Hamiltonian complexity.

2.1 General notation

For $n \in \mathbb{N}$, we use $[n]$ to denote the set collecting all the positive integers smaller than or equal to n, that is, $\{1, 2, \ldots, n\}$.

For a bit string $x \in \{0,1\}^*$, we use $int(x)$ to denote the corresponding integer whose binary representation is x. For a nonnegative integer n, we use $bin(n)$ to denote the binary representation of n. For example, $int(01010111) = 87$ and $bin(87) = 01010111$.

For an n-bit string $x \in \{0,1\}^n$ and $i \in [n]$, we use $x[i]$ to denote the ith bit of x. If x is a binary representation of an integer, then $x[1]$ is the most significant bit and $x[n]$ is the least significant bit. Let $S \in [n]$, we use x_S to denote a new string that concatenates x's bits whose indices are in S. For example, if $x = 01010111$ and $S = \{1, 3, 7, 8\}$, then $x_S = 0011$. The Hamming weight of x is denoted by $wt(x)$, that is, $wt(x)$ is the number of 1's in x.

The identity operator is denoted by I. The Kronecker delta $\delta_{i,j}$ is defined by $\delta_{i,j} = 1$ if $i = j$ and $\delta_{jk} = 0$ if $i \neq j$. The indicator function is denoted by $\mathbb{1}_S(i)$, which is defined by $\mathbb{1}_S(i) = 1$ if $i \in S$ and $\mathbb{1}_S(i) = 0$ if $i \notin S$.

The big O star notation $O^*(\cdot)$ hides the polynomial factors in the standard big O notation. For example, $8 \cdot 2^{\frac{n}{2}} \cdot n^7 = O^*(2^{\frac{n}{2}})$. The negligible functions $\mathrm{negl}(\cdot)$ are functions that are smaller than any inverse polynomial: if $\mu(n) \in \mathrm{negl}(n)$ then for all $c \in \mathbb{N}$, there exists $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$, it holds that $\mu(n) < \frac{1}{n^c}$.

2.2 Quantum computation

Quantum register

We use the sans-serif font to denote quantum registers, e.g., A, in, out, anc, clock. Throughout this paper, a register consists of qubits. A qubit may belong to different registers at the same time. For example, a qubit in the output register out of a quantum circuit is also in the ancilla register anc. By an abuse of notation, sometimes we treat the registers as sets. We say $\mathsf{A} \subseteq \mathsf{B}$ if any qubit in A is also a qubit in B. We use $\mathsf{A} \cup \mathsf{B}$ to denote the register that consists of the qubits that are in A or in B, $\mathsf{A} \cup \mathsf{B}$ to denote the portion that belongs to A and B at the same time, and $\mathsf{A} \setminus \mathsf{B}$ to denote the portion that belongs to A but not B. When we specify a system (that can be a quantum circuit or a physical qubit system), we use $\overline{\mathsf{A}}$ to denote the collection of qubits in the system that are not in A. We use $|\mathsf{A}|$ to denote the number of qubits in A. When we specify a register A, we use $\mathsf{A}[\mathrm{i}]$ to denote the ith qubit in A for $i \in [|\mathsf{A}|]$.

We use $|\psi\rangle_{\mathsf{A}}$ to denote the register A is in the state $|\psi\rangle$. For a Hermitian or a unitary operator X, we use X_{A} to denote X acting on the register A. When the registers are specified, the order of the tensor product of the operators is not sensitive. For example, $X_{\mathsf{A}} \otimes Y_{\mathsf{B}} = Y_{\mathsf{B}} \otimes X_{\mathsf{A}}$. We use $X^{\otimes k}$ to denote the kth tensor power of X for $k \in \mathbb{N}$. Namely, $X^{\otimes k} := \underbrace{X \otimes X \otimes \cdots \otimes X}_{k}$.

Quantum circuit

A quantum circuit is a unitary acting on the union of two disjoint register in and anc, which are called input register and ancilla register respectively. Also, there is an output register $\mathsf{out} \subseteq \mathsf{in} \cup \mathsf{anc}$. We require the ancilla register is $|0^{n_a}\rangle_{\mathsf{anc}}$ initially, where $n_a = |\mathsf{anc}|$. That is, if the circuit U takes a quantum state $|\psi\rangle$ as an input, then the final state of the circuit is $U|\psi\rangle_{\mathsf{in}}|0^{n_a}\rangle_{\mathsf{anc}}$. We use the notation $x \leftarrow U(|\psi\rangle)$ to denote the event that we obtain the measurement outcome $x \in \{0,1\}^{|\mathsf{out}|}$ when we measure on the output register out at the end of the circuit U that takes $|\psi\rangle$ as the input state. We have $\Pr[x \leftarrow U(|\psi\rangle)] = \langle\psi|_{\mathsf{in}}\langle 0^{n_a}|_{\mathsf{anc}} U^\dagger (|x\rangle\langle x|_{\mathsf{out}} \otimes I_{\overline{\mathsf{out}}}) U|\psi\rangle_{\mathsf{in}}|0^{n_a}\rangle_{\mathsf{anc}}$.

Universal gate set and elementary gates

A quantum circuit is implemented by a sequence of quantum gates. There is a finite set of one-qubit and two-qubit quantum gates called the universal gate set such that for any unitary U, we can use the gates in the universal gate sets to implement a quantum circuit arbitrarily close to U. We call a member in the universal gate set an elementary gate. We choose $\{\textsc{hadamard}, \pi/8\text{-}\textsc{gate}, \textsc{not}, \textsc{cnot}\}$ as the universal gate set [44]

The $\textsc{not}$ gate acts on a one-qubit register. The truth table of the $\textsc{not}$ gate is defined by $\textsc{not}|x\rangle = |(1+x) \mod 2\rangle$ where $x \in \{0,1\}$.

We also introduce a multi-control-NOT gate.

▶ **Definition 6** (Multi-control-NOT gate and Toffoli gate). *A multi-control gate, denoted by C^kNOT, is a quantum gate that acts on $k+1$ qubits. Let* C *be a k-qubit register and* T *be a one-qubit register. The truth table of the C^kCNOT gate is defined by*

$$C^k\textsc{cnot}|x_1,x_2,\ldots,x_k\rangle_{\mathsf{C}}|x_{k+1}\rangle_{\mathsf{T}} = |x_1,x_2,\ldots,x_k\rangle_{\mathsf{C}}|(x_1x_2\cdots x_k + x_{k+1}) \mod 2\rangle_{\mathsf{T}},$$

where $x_1,x_2,\ldots,x_{k+1}\in\{0,1\}$. We call C *the control qubits and* T *the target qubit.*

For $k=2$, we call C^2NOT a Toffoli gate.

According to the definition of C^kNOT, we have that CNOT is a special case of C^kNOT gate for $k=1$.

▶ Remark 7 (Decompose C^kNOT into Toffolis [44]). A C^kNOT gate can be constructed by $2k-3$ Toffoli gates associated with $k-2$ ancilla qubits. The ancilla qubits are in all zero state initially and stay unchanged at the end of circuit.

▶ Remark 8 (Decompose Toffoli into elementary gates [44]). A Toffoli gate is identical to a circuit that acts on three qubits and is composed of 17 elementary gates.

2.3 Local Hamiltonian problem

A Hamiltonian is an Hermitian operator acting on a quantum register. Let H be a Hamiltonian, we call the eigenvalues of H the energies of the Hamiltonian. We use $\lambda(H)$ to denote the smallest eigenvalue of H. We call $\lambda(H)$ the ground-state energy of H, and we call the corresponding eigenstate(s) the ground state(s) of H. We use $\|H\|$ to denote the operator norm of H, that is, the largest absolute value of the eigenvalues of H.

The k-local Hamiltonian is defined as follows.

▶ **Definition 9** (Local Hamiltonian). *We say a Hamiltonian H acting on n qubits is k-local if H can be written as $H=\sum_{i=1}^{m}H_i$ for $m=\mathrm{poly}(n)$ and for all i, the following holds.*

- *H_i is a Hamiltonian.*
- *H_i non-trivially acts on at most k qubits.*
- *$\|H_i\|\le\mathrm{poly}(n)$.*

We call k the locality of H, and define the local Hamiltonian problem formally.

▶ **Definition 10** (k-local Hamiltonian (k-LH) problem). *The local Hamiltonian problem is a decision problem that asks whether the ground-state energy of a k-local Hamiltonian is greater or less than given energy thresholds.*

- ***Inputs:*** *a k-local Hamiltonian H such that acting on n qubits where $k=O(1)$, and two energy thresholds $E_{\mathrm{yes}},E_{\mathrm{no}}$ satisfying $E_{\mathrm{no}}-E_{\mathrm{yes}}\ge\Omega(1)$.*
- ***Outputs:***
  - *YES, if there exists a quantum state $|\psi\rangle\in\mathbb{C}^{2^n}$ such that $\langle\psi|H|\psi\rangle\le E_{\mathrm{yes}}$.*
  - *NO, if $\langle\psi|H|\psi\rangle\ge E_{\mathrm{no}}$ for all $|\psi\rangle\in\mathbb{C}^{2^n}$.*

In other words, the k-LH problem asks to decide whether $\lambda(H)\le E_{\mathrm{yes}}$ or $\lambda(H)\ge E_{\mathrm{no}}$. We say an algorithm $\mathcal{A}_{\mathrm{LH}}$ solves k-LH$(H,E_{\mathrm{yes}},E_{\mathrm{no}})$ if $\mathcal{A}_{\mathrm{LH}}$ decides $(H,E_{\mathrm{yes}},E_{\mathrm{no}})$ correctly.

2.4 Quantum partition function

Let us introduce the quantum partition function formally.

▶ **Definition 11** (Approximating quantum partition function of k-local Hamiltonian (k-QPF) problem). *The quantum partition function problem is to approximate the partition function of a k-local Hamiltonian up to a multiplicative error under a certain temperature.*

- ***Inputs:*** *a k-local Hamiltonian H acting on n qubits where $k = O(1)$, an inverse temperature $\beta \leq \mathrm{poly}(n)$, and an error parameter $\delta \in (0,1)$.*
- ***Outputs:*** *$\widetilde{Z} \in \mathbb{R}$ such that*

$$(1-\delta)Z \leq \widetilde{Z} \leq (1+\delta)Z, \tag{3}$$

where $Z := \mathrm{tr}[e^{-\beta H}]$.

We say an algorithm $\mathcal{A}_{\mathrm{QPF}}$ solves k-QPF(H,β,δ) if $\mathcal{A}_{\mathrm{QPF}}$ outputs such $\widetilde{Z}$ on the inputs H,β,δ.

From now on, when we use the term QPF or k-QPF, it means to approximate the quantum partition function up to a relative error, instead of to compute the exact value.

2.5 Boolean formula, SAT, SETH and QSETH

A Boolean variable can be assigned to the value 0 or 1. We abbreviate the term Boolean variable as variable. A Boolean formula consists of variables associated with parentheses and logic connectives $\neg$ (NOT), $\vee$ (OR), and $\wedge$ (AND).

Let Φ be a Boolean formula with n variables $x_1, x_2, \ldots, x_n$. We use an an n-bit string $x \in \{0,1\}^n$ to denote an assignment to $x_1, x_2, \ldots, x_n$. The variable x_i for $i \in [n]$ is assigned to the ith bit of x. We use $\Phi(x)$ to denote the value of Φ when $x_1, x_2, \ldots, x_n$ are assigned to x. We say x satisfies Φ if $\Phi(x) = 1$.

For all variables x, it holds that $\neg(\neg x) = x$. Let $x_1, x_2, \ldots, x_n$ be variables and for each $i \in [n]$, let ℓ_i be either x_i or $\neg x_i$. It holds that $\neg(\ell_1 \wedge \ell_2 \wedge \cdots \wedge \ell_m) = \neg\ell_1 \vee \neg\ell_2 \vee \cdots \vee \neg\ell_m$.

We can use a quantum circuit to compute Boolean formulas. For $x \in \{0,1\}$, it holds that

$$\mathrm{NOT}\,|x\rangle = |\neg x\rangle\,. \tag{4}$$

For $x_1, x_2, \ldots, x_k \in \{0,1\}$, it holds that

$$C^k\mathrm{NOT}_{\mathsf{C}\cup\mathsf{T}}\,|x_1, x_2, \ldots, x_k\rangle_{\mathsf{C}}\,|0\rangle_{\mathsf{T}} = |x_1, x_2, \ldots, x_k\rangle_{\mathsf{C}}\,|x_1 \wedge x_2 \wedge \cdots \wedge x_k\rangle_{\mathsf{T}}\,. \tag{5}$$

We define the conjunctive normal form formula as follows.

▶ **Definition 12** (Conjunctive normal form (CNF) formula)**.** *Let Φ be a Boolean formula that consists of n variables $x_1, x_2, \ldots, x_k$. We say Φ is a kCNF formula if Φ is in the form of $m = \mathrm{poly}(n)$ smaller formulas connected by $\wedge$. Each smaller formula contains at most k variables, and the variables are connected by $\vee$.*

To be more precise, $\Phi = \varphi_1 \wedge \varphi_2 \wedge \cdots \wedge \varphi_m$ where $m = \mathrm{poly}(n)$, and for all $i \in [m]$, the following constrains hold.

- *$\varphi_i = (\ell_{i,1} \vee \ell_{i_2} \vee \cdots \vee \ell_{i,k_i})$, where $\ell_{i,p}$ can be x_j or $\neg x_j$ for all $p \in [k_i]$ and $j \in [n]$.*
- *$k_i \leq k$.*
- *For each $j \in [n]$, x_j and $\neg x_j$ do not appear in φ_i at the same time; x_j, $\neg x_j$ appears in φ_i at most once.*

We say φ_i is a clause of Φ for all $i \in [m]$.

▶ **Definition 13** (Satisfiability for k-CNF (k-SAT) problem)**.** *The satisfiability problem is a decision problem that asks whether there is an assignment satisfying the given kCNF formula.*

- ***Inputs:*** *a k-CNF formula Φ that contains n variables. The number of clauses of Φ is $m = \mathrm{poly}(n)$.*
- ***Outputs:***
  - *YES, if there exists an assignment $x \in \{0,1\}^n$ such that $\Phi(x) = 1$.*
  - *No, if $\Phi(x) = 0$ for all $x \in \{0,1\}^n$.*

We say an algorithm $\mathcal{A}_{\mathrm{SAT}}$ solves k-SAT(Φ) if $\mathcal{A}_{\mathrm{SAT}}$ decides Φ correctly.

Though the exact lower bound for k-SAT problem is still unknown, it is widely believed that brute-force search is optimal for classical algorithm and Grover search is optimal for quantum algorithm. The Strong Exponential Time Hypothesis (SETH) [29] states that k-SAT needs roughly 2^n time for large k:

▶ **Conjecture 14** (SETH). *For all ε, there is some $k \geq 3$ such that k-SAT with n variables and $O(n)$ clauses cannot be solved classically in time $O(2^{(1-\varepsilon)n})$.*

A quantum analog was proposed in [1, 12] as Quantum Strong Exponential-Time Hypothesis (QSETH).

▶ **Conjecture 15** (Quantum strong exponential time hypothesis (QSETH) [1, 12]). *For all ε, there is some $k \geq 3$ such that k-SAT with n variables and $O(n)$ clauses cannot be solved quantumly in time $O(2^{(1-\varepsilon)n/2})$.*

In the two conjecture above, we can assume the number of clauses can be assumed to be $O(n)$ via the sparsification lemma [30].

2.6 Fine-grained reduction

Here we introduce the fine-grained reduction.

▶ **Definition 16** (Fine-grained reduction [49]). *Let P and Q be two problems and $\mathcal{A}_Q$ be an oracle that solves Q with probability greater than $\frac{2}{3}$. Let $p(\cdot)$ and $q(\cdot)$ be two non-decreasing functions. We say P is (p, q) reducible to Q if for all ε, there exist ξ, an algorithm $\mathcal{A}_P$, a constant d, and an integer $r(n)$ such that the algorithm $\mathcal{A}_P$ that can black-boxly assess to $\mathcal{A}_Q$ takes an instance of P with size n and satisfies the following.*

1. *$\mathcal{A}_P$ solves P with probability greater than $\frac{2}{3}$.*
2. *$\mathcal{A}_P$ runs in $d \cdot p(n)^{1-\xi}$ time.*
3. *$\mathcal{A}_P$ produces at most $r(n)$ instances of Q adaptively.*
4. *$\sum_{i=1}^{r(n)} (q(n_i))^{1-\varepsilon} \leq d \cdot (p(n))^{1-\xi}$, where n_i is the size of the ith instance of problem Q that is produced by $\mathcal{A}_P$.*

One can also use a more general reduction, *quantum fine-grained reduction* [1]. In a quantum fine-grained reduction, the reduction algorithm $\mathcal{A}_P$ is allowed to query the oracle $\mathcal{A}_Q$ in superposition. In this work, however, Definition 16 is sufficient to establish the lower bounds for both the Local Hamiltonian problem and the approximation of the Quantum Partition Function.

When P is (p, q) reducible to Q and every instance produced by $\mathcal{A}_P$ has size $n + o(n)$, we have that if P cannot be solved within time $O(p(n)^{1-\varepsilon})$ for any ε, then Q cannot be solved within time $O(q(n)^{1-\xi})$ for any ξ.

3 Size-preserving circuit-to-Hamiltonian construction

In this section, we prove the first size-preserving circuit-to-Hamiltonian construction as follows.

▶ **Theorem 17.** *Let $x \in \{0,1\}^n$ be an input and U_x be a* QMA *verification quantum circuit U_x acting on $N = \mathrm{poly}(n)$ qubits, and suppose U_x consists of $T = \mathrm{poly}(n)$ gates. Then, for any integer $d \geq 1$, there exists a $(d+1)$-local Hamiltonian H_x acting on $N + O(T^{1/d})$ qubits and $\|H_x\| \leq \mathrm{poly}(n)$ such that for all $\varepsilon_{\mathrm{err}} \in (0,1)$, the following two hold.*

- *If there exists a quantum proof* $|\psi\rangle$ *such that* $\Pr[U_x \text{ accepts } |\psi\rangle] \geq 1 - \varepsilon_{\text{err}}$*, then* $\lambda(H_x) \leq \varepsilon_{\text{err}}$.
- *If* $\Pr[U_x \text{ accepts } |\psi\rangle] \leq \varepsilon_{\text{err}}$ *for any quantum proof* $|\psi\rangle$*, then* $\lambda(H_x) \geq \frac{1}{2} - \varepsilon_{\text{err}}$.

To prove the theorem above, we first recall the strategy for the circuit-to-Hamiltonian construction in [32]. We then define our clock states and finally prove the statement.

3.1 The construction by Kitaev, Kempe and Regev [32]

In this section, we explain the 2-local construction from [32] with some adaptation.

We consider a promise problem $(L_{\text{yes}}, L_{\text{no}})$ in QMA. Let $U = U_T \cdots U_2 U_1$ be the quantum circuit of the verifier (which depends on the input), where $T = \text{poly}(n)$. We assume that U acts on $N = \text{poly}(n)$ qubits: the first m qubits contain the proof (the last $N - m$ qubits are initialized to $|0\rangle$). Here, each U_i is a 1-qubit gate or a 2-qubit gate. As in [32], we assume that the only 2-qubit gate used by the circuit is the controlled phase gate, each controlled phase gate is preceded and followed by two Z gates[10] Let $T_A \subseteq \{1, \ldots, T\}$ be the set of indices i such that U_i is a 1-qubit gate. Let $T_B \subseteq \{1, \ldots, T\}$ be the set of indices i such that U_i is a 2-qubit gate.

Meta-definition of the Hamiltonian [11]

The Hamiltonian is defined over the Hilbert space $\mathcal{H}_{\text{circuit}} \otimes \mathcal{H}_{\text{clock}}$, where $\mathcal{H}_{\text{circuit}}$ is a Hilbert space of dimension 2^N representing the computation and $\mathcal{H}_{\text{clock}}$ is a Hilbert space of dimension at least $T + 1$ representing the clock. There are $T + 1$ orthonormal unit vectors $|\gamma_0\rangle, \ldots, |\gamma_T\rangle$ in $\mathcal{H}_{\text{clock}}$ that are called the legal clock states. Let $\mathcal{H}_{\text{legal}}$ denote the subspace spanned by the states $|\gamma_0\rangle, \ldots, |\gamma_T\rangle$ and Π_{legal} be the projection into this subspace. Let $\mathcal{H}_{\text{illegal}}$ be its orthogonal subspace. We define the operators

$$\begin{aligned}
\mathcal{S}_t &= |\gamma_t\rangle\langle\gamma_t| \quad \text{for any } t \in \{0, \ldots, T\}, \\
\mathcal{A}_t &= |\gamma_{t+1}\rangle\langle\gamma_t| \quad \text{for any } t \in \{0, \ldots, T-1\}, \\
\mathcal{B}_t &= |\gamma_{t+2}\rangle\langle\gamma_t| \quad \text{for any } t \in \{0, \ldots, T-2\},
\end{aligned}$$

acting on $\mathcal{H}_{\text{clock}}$ that respectively stop the clock, increment the clock by one step and increment the clock by two steps.

Consider the following Hamiltonian:

$$H = \alpha_{\text{in}} H_{\text{in}} + \alpha_{\text{out}} H_{\text{out}} + \alpha_A \sum_{t \in T_A} H_{\text{prop},t} + \alpha_B \sum_{t \in T_B} (H_{\text{qubit},t} + H_{\text{time},t}) + \alpha_{\text{clock}} I \otimes H_{\text{clock}}.$$

The choice of the coefficients α_{in}, α_{out}, α_A, α_B, α_{clock} will be discussed later. The terms H_{in} and H_{out} are defined as follows:

$$H_{\text{in}} = \sum_{i=m+1}^{N} |1\rangle_i\langle 1|_i \otimes \mathcal{S}_0, \qquad H_{\text{out}} = |0\rangle_1\langle 0|_1 \otimes \mathcal{S}_T.$$

For any $t \in T_A$,

$$H_{\text{prop},t} = \frac{1}{2}\Big(I \otimes \mathcal{S}_t + I \otimes \mathcal{S}_{t-1} - U_t \otimes \mathcal{A}_{t-1} - U_t^\dagger \otimes \mathcal{A}_{t-1}^\dagger\Big).$$

[10] Any quantum circuit composed of the standard universal gate set can be compiled to the circuit described above with only linear blowup in circuit size because CNOT $= (I \otimes H)\text{CZ}(I \otimes H)$.

[11] This is a "meta-definition" since it does not define the concrete implementation of the clock states.

For any $t \in T_B$,

$$H_{\text{time},t} = \frac{1}{8} I \otimes \Big(\mathcal{S}_t + 6\mathcal{S}_{t+1} + \mathcal{S}_{t+2} + 2\mathcal{B}_t + 2\mathcal{B}_t^\dagger + \mathcal{A}_t + \mathcal{A}_t^\dagger + \mathcal{A}_{t+1} + \mathcal{A}_{t+1}^\dagger + \mathcal{S}_{t-3} + 6\mathcal{S}_{t-2} + \mathcal{S}_{t-1} + 2\mathcal{B}_{t-1} + 2\mathcal{B}_{t-1}^\dagger + \mathcal{A}_{t-3} + \mathcal{A}_{t-3}^\dagger + \mathcal{A}_{t-1} + \mathcal{A}_{t-1}^\dagger \Big),$$

$$H_{\text{qubit},t} = \frac{1}{2} \Big(-2 |0\rangle\langle 0|_{f_t} - 2 |0\rangle\langle 0|_{s_t} + |1\rangle\langle 1|_{f_t} + |1\rangle\langle 1|_{s_t} \Big) \otimes (\mathcal{A}_{t-1} + \mathcal{A}_{t-1}^\dagger),$$

where f_t and s_t denote the first and second qubits of the controlled phase gate U_t. The term H_{clock} will be chosen later so that the following two conditions are satisfied

$$\|H_{\text{clock}}\| \le \text{poly}(n), \tag{C1}$$

$$\langle \gamma_t | H_{\text{clock}} | \gamma_t \rangle = 0 \quad \text{for any} \quad t \in \{0, \dots, T\}, \tag{C2}$$

$$\langle \varphi | H_{\text{clock}} | \varphi \rangle \ge 1 \quad \text{for any} \quad |\varphi\rangle \in \mathcal{H}_{\text{illegal}}. \tag{C3}$$

Concrete implementation

A straightforward implementation uses a binary clock, as in the original construction by [34]. While requiring only $O(\log T)$ clock qubits (and no H_{clock} term), the locality of the Hamiltonian would become $O(\log n)$. We discuss below how to reduce the locality.

The basic idea is to replace each of the operator $\mathcal{S}_t, \mathcal{A}_t, \mathcal{B}_t$ by another operator $\widetilde{\mathcal{S}}_t, \widetilde{\mathcal{A}}_t, \widetilde{\mathcal{B}}_t$ that has better locality. In order for the implementation to be correct, the following three conditions need to be satisfied.

$$\text{For any} \quad t \in \{0, \dots, T\}, \quad \begin{cases} \widetilde{\mathcal{S}}_t |\gamma_t\rangle = |\gamma_t\rangle \\ \Pi_{\text{legal}} \widetilde{\mathcal{S}}_t |\varphi\rangle = 0 \quad \text{for any } |\varphi\rangle \in \mathcal{H}_{\text{clock}} \text{ orthogonal to } |\gamma_t\rangle. \end{cases} \tag{C4}$$

$$\text{For any} \quad t \in \{0, \dots, T-1\}, \quad \begin{cases} \widetilde{\mathcal{A}}_t |\gamma_t\rangle = |\gamma_{t+1}\rangle \\ \Pi_{\text{legal}} \widetilde{\mathcal{A}}_t |\varphi\rangle = 0 \quad \text{for any } |\varphi\rangle \in \mathcal{H}_{\text{clock}} \text{ orthogonal to } |\gamma_t\rangle. \end{cases} \tag{C5}$$

$$\text{For any} \quad t \in \{0, \dots, T-2\}, \quad \begin{cases} \widetilde{\mathcal{B}}_t |\gamma_t\rangle = |\gamma_{t+2}\rangle \\ \Pi_{\text{legal}} \widetilde{\mathcal{B}}_t |\varphi\rangle = 0 \quad \text{for any } |\varphi\rangle \in \mathcal{H}_{\text{clock}} \text{ orthogonal to } |\gamma_t\rangle. \end{cases} \tag{C6}$$

The first part of each condition says that each $\widetilde{\mathcal{S}}_t, \widetilde{\mathcal{A}}_t, \widetilde{\mathcal{B}}_t$ behaves as $\mathcal{S}_t, \mathcal{A}_t, \mathcal{B}_t$, respectively, on the clock state $|\gamma_t\rangle$. The second part says that while the behavior may differ from the behavior of $\mathcal{S}_t, \mathcal{A}_t, \mathcal{B}_t$ on other states, $\widetilde{\mathcal{S}}_t, \widetilde{\mathcal{A}}_t, \widetilde{\mathcal{B}}_t$ never maps an illegal state or even a legal state other than $|\gamma_t\rangle$ to a legal state.

The analysis based on the projection lemma (Lemma 1 in [32]) shows the following.

▶ **Proposition 18** (Adapted from Section 5 in [32]). *If Conditions* $(C_1) \sim (C_6)$ *are satisfied, coefficients* α_{in}, α_{out}, α_A, α_B, α_{clock} *can be chosen in the range* $[0, \text{poly}(n)]$ *so that the following holds:*

- *If there exists a quantum proof* $|\psi\rangle$ *such that* $\Pr[U_x \text{ accepts } |\psi\rangle] \ge 1 - \varepsilon$, *then* $\lambda(H_x) \le \varepsilon$.
- *If* $\Pr[U_x \text{ accepts } |\psi\rangle] \le \varepsilon$ *for any quantum proof* $|\psi\rangle$, *then* $\lambda(H_x) \ge \frac{1}{2} - \varepsilon$.

Implementation by Kempe-Kitaev-Regev

The implementation in [32] uses the unary encoding 0000, 1000, 1100, 1110, 1111$\cdots$ for the clock states, i.e., uses the T-qubit clock states

$$|\gamma_t\rangle = |1^t 0^{T-t}\rangle$$

for each $t \in \{0,\dots,T\}$, and the operators $\widetilde{\mathcal{S}}_t, \widetilde{\mathcal{A}}_t, \widetilde{\mathcal{B}}_t$ described in the following table:

Table 2 Concrete implementation of operators in [32].

Operator	t	Implementation	Locality
$\mathcal{S}_t$	$t=0$	$\widetilde{\mathcal{S}}_0 = \|0\rangle_1 \langle 0\|_1$	1
	$t \in \{1,\dots,T-1\}$	$\widetilde{\mathcal{S}}_t = \|10\rangle_{t,t+1} \langle 10\|_{t,t+1}$	2
	$t=T$	$\widetilde{\mathcal{S}}_T = \|1\rangle_T \langle 1\|_T$	1
$\mathcal{A}_t$	$t \in \{0,\dots,T-1\}$	$\widetilde{\mathcal{A}}_t = \|1\rangle_{t+1} \langle 0\|_{t+1}$	1
$\mathcal{B}_t$	$t \in \{0,\dots,T-2\}$	$\widetilde{\mathcal{B}}_t = \|11\rangle_{t+1,t+2} \langle 00\|_{t+1,t+2}$	2

The term H_{clock} used in [32] is defined as follows:

$$H_{\mathrm{clock}} = \sum_{1\le i<j\le T} |01\rangle_{ij} \langle 01|_{ij} \ .$$

This implementation satisfies Conditions $(\mathrm{C}_1) \sim (\mathrm{C}_6)$. With this implementation, each term H_{in}, H_{out}, $H_{\mathrm{prop},t}$, $H_{\mathrm{qubit},t}$, $H_{\mathrm{time},t}$, H_{clock} has locality 2. The whole Hamiltonian is thus 2-local.

3.2 (n,d)-clock state from [14] and associated operators to implement LH

In this subsection, we recall an (n,d)-clock state proposed in [14], and introduce and analyze operators to use it for circuit-to-Hamiltonian constructions. The (n,d)-clock state is based on a Johnson graph. The vertices of the Jonson graph are all the subsets of $[n]$ with size $k<n$. Two vertices are adjacent if and only if the subsets share exactly $d-1$ elements.

▶ **Definition 19** (Johnson graph). *Let $n,d \in \mathbb{N}$ and $d<n$. We say a Johnson graph $J(n,d)=(V,E)$ is a graph that satisfies the following requirement.*

- $V = \{S \subseteq [n] : |S| = d\}$.
- $E = \{(S,S') \in V^2 : |S \cap S'| = d-1\}$.

The number of vertices of $J(n,d)$ is $\binom{n}{d}$. In [2], it has been proved that for all n,d, there is a Hamiltonian path[12] in the Johnson graph $J(n,d)$. Also, the proof implicitly construct an algorithm that finds the Hamiltonian path in $O(n^d)$ time.

Now we explain how to construct the clock. A clock state is an n-qubit state. Each clock state $|\gamma_t\rangle$ corresponds to a vertex S_t on a Hamiltonian path in $J(n,d)$. All the clock states are basis states of the computational basis whose number of 1's equals to d. The ith qubit in $|\gamma_t\rangle$ is $|1\rangle$ if and only if i is chosen into S_t. In other words, when we treat the clock register as

[12] A Hamiltonian path is a path that visits every vertex in a graph exactly once. Please do not confuse with the physical quantity (local Hamiltonian) H.

the set $[n]$ and let $\mathsf{S_t}$ be the portion therein corresponding to S_t, the state in $\mathsf{S_t}$ is $|1^d\rangle$. The next clock $|\gamma_{t+1}\rangle$ corresponds to the vertex S_{t+1} that is adjacent to S_t in the Hamiltonian path. To update the $|\gamma_t\rangle$ to $|\gamma_{t+1}\rangle$, we make $d-1$ of 1's unchanged and flip two bits. Hence the update of the clock state can be done with constant local operation. We define the clock state formally as follows.

▶ **Definition 20** ((n,d)-clock state [14]). *A (n,d)-clock state is a collection of n-qubit quantum state $\{|\gamma_t\rangle\}_{t=0}^{T}$ defined by a Johnson graph $J(n,d)=(V,E)$ and $T=\binom{n}{d}-1$. Let $S_0,S_1,\ldots,S_T\in V$ and the sequence $S_0,S_1,\ldots,S_T\in V$ forms a Hamiltonian path in $J(n,d)$. For any $t\in\{0\}\cup[T]$ the state $|\gamma_t\rangle$ is defined by $|\gamma_t\rangle:=\bigotimes_{i\in[n]}|\mathbb{1}_{S_t}(i)\rangle$.*

To use the clock above for the circuit-to-Hamiltonian construction, we introduce several local operators to pause and move time on the clock. For all $t\in 0\cup[T]$, let $\mathsf{S_t}$ be the corresponding register of S_t. For all $t\in[T]$, we define the operator F_t as follows.

$$F_t:=|1\rangle\langle 0|_{\mathsf{S_{t-1}}\setminus\mathsf{S_t}}\otimes|0\rangle\langle 1|_{\mathsf{S_t}\setminus\mathsf{S_{t-1}}}\otimes|1^{d-1}\rangle\langle 1^{d-1}|_{\mathsf{S_t}\cap\mathsf{S_{t-1}}}. \tag{6}$$

Hence we have

$$F_t^\dagger:=|0\rangle\langle 1|_{\mathsf{S_{t-1}}\setminus\mathsf{S_t}}\otimes|1\rangle\langle 0|_{\mathsf{S_t}\setminus\mathsf{S_t}}\otimes|1^{d-1}\rangle\langle 1^{d-1}|_{\mathsf{S_t}\cap\mathsf{S_{t-1}}}. \tag{7}$$

It holds that

$$F_t\,|\gamma_{t'}\rangle=\delta_{t',t-1}\,|\gamma_{t'+1}\rangle\,, \tag{8}$$

and

$$F_t^\dagger\,|\gamma_{t'}\rangle=\delta_{t',t}\,|\gamma_{t'-1}\rangle\,. \tag{9}$$

That is, F_t "forwards" the clock $|\gamma_{t-1}\rangle$ one step, and eliminate all the other state $|\gamma_{t'}\rangle$ where $t'\neq t-1$. Likewise, $F_t^\dagger$ "backwards" the clock $|\gamma_t\rangle$ one step, and eliminate all the other state $|\gamma_{t'}\rangle$ where $t'\neq t$. We have that F_t is $(d+1)$-local. Note that F_t is neither unitary nor Hermitian.

For $t\in\{0\}\cup[T]$, we define the operator P_t as following.

$$P_t:=|1^d\rangle\langle 1^d|_{\mathsf{S_t}}. \tag{10}$$

We have $P_t^\dagger=P_t$. We also have

$$P_t\,|\gamma_{t'}\rangle=\delta_{t',t}\,|\gamma_{t'}\rangle\,. \tag{11}$$

That is, P_t "pauses" the clock $|\gamma_t\rangle$, and eliminate all the other state $|\gamma_{t'}\rangle$ where $t'\neq t$. The operator P_t is d-local.

We have seen that we can "forward", "backword" and "pause" the clock with local operators. To use the (n,d) clock for the circuit-to-Hamiltonian construction, we also want valid clock states $\{|\gamma_t\rangle\}_{t=0}^{T}$ to be the ground state for the Hamiltonian, and give some energy penalty to any other states orthogonal to valid clock states. In order to achieve this, we define

$$H_{\mathrm{stab}}:=\binom{n}{d}H_{>d}+H_{<d}-\left(\binom{n}{d}-1\right)I, \tag{12}$$

where

$$H_{>d}:=\sum_{\mathsf{S}:|\mathsf{S}|=d+1}|1^{d+1}\rangle\langle 1^{d+1}|_{\mathsf{S}}, \tag{13}$$

$$H_{<d}:=\sum_{\mathsf{S}:|\mathsf{S}|=d}\sum_{x\in\{0,1\}^d\setminus\{1^d\}}|x\rangle\langle x|_{\mathsf{S}}, \tag{14}$$

The term $H_{>d}$ gives penalty to the state that contains "too many" 1's, and the term $H_{<d}$ gives penalty to the state that contains "too few" 1's.

▶ **Lemma 21.** *The linear space for the ground states of H_{stab} is spanned by $\{|\gamma_t\rangle\}_{t=0}^{T}$, and for any state $|\phi\rangle$ orthogonal to all $\{|\gamma_t\rangle\}_{t=0}^{T}$, we have*

$$\langle\phi|H_{\mathrm{stab}}|\phi\rangle \geq 1.$$

Proof. Since H_{stab} is diagonalized in the computational basis i.e., each computational state is an eigenstate of H_{stab}. To verify the above statement, we can check $\langle x|H_{\mathrm{stab}}|x\rangle$ for all $x \in \{0,1\}^n$.

We divide the Hilbert space $\mathcal{H}_{\mathrm{clock}}$ into three subspaces $\mathcal{L}_{=d}$, $\mathcal{L}_{>d}$, and $\mathcal{L}_{<d}$ which are defined below.

$$\mathcal{L}_{<d} := \mathrm{Span}(\{w \in \{0,1\}^n : wt(w) < d\}), \tag{15}$$
$$\mathcal{L}_{=d} := \mathrm{Span}(\{y \in \{0,1\}^n : wt(y) = d\}), \tag{16}$$
$$\mathcal{L}_{>d} := \mathrm{Span}(\{z \in \{0,1\}^n : wt(z) > d\}) \tag{17}$$

Note that $\mathcal{L}_{=d}$ is the subspace spanned by the valid clock states $\{|\gamma_t\rangle\}_{t=0}^{T}$.

We first calculate $H_{>d}$ acting on the states in $\mathcal{L}_{=d}$, $\mathcal{L}_{>d}$, and $\mathcal{L}_{<d}$ respectively. For any S with size $d+1$, $|1^{d+1}\rangle\langle 1^{d+1}|_{\mathsf{S}}\,|x\rangle = |x\rangle$ if the subset-string $x_S = 1^{d+1}$. Hence, we have the following.

- For all $w \in \{0,1\}^n$ such that $wt(w) < d$, it holds that $\langle w|\,H_{>d}\,|w\rangle = 0$ because for all $S \subseteq [n]$ with size $d+1$, z_S has at most $d-1$ of 1's.
- For all $y \in \{0,1\}^n$ such that $wt(y) = d$, it holds that $\langle y|\,H_{>d}\,|y\rangle = 0$ because for all $S \subseteq [n]$ with size $d+1$, w_S has at most d of 1's.
- For all $z \in \{0,1\}^n$ such that $wt(z) > d$, it holds that $\langle z|\,H_{>d}\,|z\rangle = r$ where $r \geq 1$, because there is at least one $S \subseteq [n]$ with size $k+1$ such that $y_S = 1^{d+1}$.

Then we calculate $H_{<d}$ acting on the states in $\mathcal{L}_{=d}$, $\mathcal{L}_{>d}$, and $\mathcal{L}_{<d}$ respectively.

Let $w \in \{0,1\}^n$ such that $wt(w) < d$. For any S with size d, we have $\sum_{x\in\{0,1\}^d\setminus\{1^d\}}|x\rangle\langle x|_{\mathsf{S}}\,|w\rangle = |w\rangle$, because there exist exactly one $x \in \{0,1\}^d \setminus \{1^d\}$ such that $w_S = x$. Hence,

$$\sum_{\mathsf{S}:|\mathsf{S}|=d}\sum_{x\in\{0,1\}^d\setminus\{1^d\}}|x\rangle\langle x|_{\mathsf{S}}\,|z\rangle = \binom{n}{d}|z\rangle\,.$$

Let $y \in \{0,1\}^n$ such that $wt(y) = d$. Let $S' \subseteq [n]$ be the subset with size d such that $w_{S'} = 1^d$. For any S with size d such that $S \neq S'$, we have that $\sum_{x\in\{0,1\}^d\setminus\{1^d\}}|x\rangle\langle x|_{\mathsf{S}}\,|y\rangle = |y\rangle$, and for $S = S'$, we have $\sum_{x\in\{0,1\}^d\setminus\{1^d\}}|x\rangle\langle x|_{\mathsf{S}}\,|y\rangle = 0$; Hence,

$$\sum_{\mathsf{S}:|\mathsf{S}|=d}\sum_{x\in\{0,1\}^d\setminus\{1^d\}}|x\rangle\langle x|_{\mathsf{S}}\,|w\rangle = \left(\binom{n_{cl}}{d}-1\right)|w\rangle\,.$$

For $z \in \{0,1\}^n$ such that $wt(z) > d$, since the penalty has been given by $H_{>d}$ already, $H_{<d}\,|z\rangle = p\,|z\rangle$ where $p \geq 0$ is sufficient for us.

Therefore, we have the follows.

- For all $w \in \{0,1\}^n$ such that $wt(w) < d$, it holds that $H_{<d}\,|w\rangle = \binom{n}{d}|w\rangle$.
- For all $y \in \{0,1\}^n$ such that $wt(y) = d$, it holds that $H_{<d}\,|y\rangle = \left(\binom{n}{d}-1\right)|y\rangle$.
- For all $z \in \{0,1\}^n$ such that $wt(z) > d$, it holds that $H_{<d}\,|z\rangle = p\,|z\rangle$ where $p \geq 0$.

As a result, we have the follows.

- For all $w \in \{0,1\}^n$ such that $wt(w) < d$, it holds that $\langle w| H_{\text{stab}} |w\rangle = 1$.
- For all $y \in \{0,1\}^n$ such that $wt(y) = d$, it holds that $\langle y| H_{\text{stab}} |y\rangle = 0$.
- For all $z \in \{0,1\}^n$ such that $wt(z) > d$, it holds that $\langle z| H_{\text{stab}} |z\rangle \geq 1$.

Consequently, if $|\phi\rangle \in \mathcal{L}_{=d}$, then $\langle\phi|H_{\text{stab}}|\phi\rangle = 0$, and if $|\phi\rangle \in \mathcal{L}_{=d}^{\perp} = \mathcal{L}_{<d} \oplus \mathcal{L}_{>d}$, then $\langle\phi|H_{\text{stab}}|\phi\rangle \geq 1$. ◀

3.3 Our new clock states and proof of Theorem 17

We show how to modify the construction from Section 3.1 to obtain a length-preserving construction. The construction is the same as the construction in Section 3.1, except that we are using a new clock, new clock operators, and a new term H_{clock}.

Proof of Theorem 17. If $d = 1$, the statement follows from the construction of [32] in Section 3.1. For the rest of the proof, we assume $d \geq 2$.

To further reduce the locality, we introduce a new clock to count from 0 to T. We assume (without loss of generality) that T is of the form $T = \binom{a}{d-1}(a+1) - 1$ for some a. We use two clocks, each using a qubits. The first clock is the $(a, d-1)$-clock following Definition 20, and the second clock is the unary clock on a qubits. Parameters are slightly different from the previous section, and not to confuse readers, we clarify the operator of the $(a, d-1)$-clock states.

$$F'_t := |1\rangle\langle 0|_{\mathsf{S}_{t-1}\setminus\mathsf{S}_t} \otimes |0\rangle\langle 1|_{\mathsf{S}_t\setminus\mathsf{S}_{t-1}} \otimes |1^{d-2}\rangle\langle 1^{d-2}|_{\mathsf{S}_t\cap\mathsf{S}_{t-1}}. \tag{18}$$

$$P'_t := |1^{d-1}\rangle\langle 1^{d-1}|_{\mathsf{S}_t}. \tag{19}$$

$$H'_{\text{stab}} := \binom{a}{d-1} H'_{>d-1} + H'_{<d-1} - \left(\binom{a}{d-1} - 1\right) I, \tag{20}$$

where

$$H'_{>d-1} := \sum_{\mathsf{S}:|\mathsf{S}|=d} |1^d\rangle\langle 1^d|_{\mathsf{S}}, \tag{21}$$

$$H'_{<d-1} := \sum_{\mathsf{S}:|\mathsf{S}|=d-1} \sum_{x\in\{0,1\}^{d-1}\setminus\{1^{d-1}\}} |x\rangle\langle x|_{\mathsf{S}}, \tag{22}$$

The whole clock thus requires only $2a \leq 2T^{1/d}$ qubits. To represent the order of the $T+1$ representations of the clock, we define a map

$$g\colon \{0,\dots,T\} \to \{1,\dots,\binom{a}{d-1}\} \times \{0,\dots,a\}$$

as $g(t) = (1 + \lfloor t/(a+1)\rfloor, \min(m, 2a+1-m))$ where $m = t \bmod (2a+2)$. We implicitly use g to interpret time steps $t \in \{0,\dots,T\}$ as pairs $(t_1,t_2) \in \{1,\dots,\binom{a}{d}\} \times \{0,\dots,a\}$. We explain the construction of g for $a = 3, d = 1$ in Table 3.

Our implementation of the clock operators is described in Table 4.

The term H_{clock}. The term H_{clock} is defined as follows:

$$H_{\text{clock}} = H'_{\text{stab}} \otimes I + \sum_{0\leq i<j\leq a} I \otimes |01\rangle_{ij}\langle 01|_{ij},$$

where in term, the left part acts on the register corresponding to the first clock and the right part acts on the qubits corresponding to the second clock.

Table 3 Our clock for $T = 11$ (i.e., $a = 3$). The part in red represents the first clock, while the part in blue represents the second clock. The first clock uses the encoding 100, 010, 001. The second clock uses the encoding 000, 100, 110, 111.

t	$g(t) = (t_1, t_2)$	clock state
$t = 0$	$(1, 0)$	$\|100\rangle \otimes \|000\rangle$
$t = 1$	$(1, 1)$	$\|100\rangle \otimes \|100\rangle$
$t = 2$	$(1, 2)$	$\|100\rangle \otimes \|110\rangle$
$t = 3$	$(1, 3)$	$\|100\rangle \otimes \|111\rangle$
$t = 4$	$(2, 3)$	$\|010\rangle \otimes \|111\rangle$
$t = 5$	$(2, 2)$	$\|010\rangle \otimes \|110\rangle$
$t = 6$	$(2, 1)$	$\|010\rangle \otimes \|100\rangle$
$t = 7$	$(2, 0)$	$\|010\rangle \otimes \|000\rangle$
$t = 8$	$(3, 0)$	$\|001\rangle \otimes \|000\rangle$
$t = 9$	$(3, 1)$	$\|001\rangle \otimes \|100\rangle$
$t = 10$	$(3, 2)$	$\|001\rangle \otimes \|110\rangle$
$t = 11$	$(3, 3)$	$\|001\rangle \otimes \|111\rangle$

Table 4 The implementation of clock operations. The part in red acts on the qubits corresponding to the first clock and the part in blue acts on the qubits corresponding to the second clock.

Operator	t_1	t_2	Implementation	Locality
$\mathcal{S}_t$	any	$t_2 = 0$	$P'_{t_1} \otimes \|0\rangle_1 \langle 0\|_1$	d
		$t_2 \in \{1, \ldots, a-1\}$	$P'_{t_1} \otimes \|10\rangle_{t_2,t_2+1} \langle 10\|_{t_2,t_2+1}$	$d+1$
		$t_2 = a$	$P'_{t_1} \otimes \|1\rangle_a \langle 1\|_a$	d
$\mathcal{A}_t$	odd	$t_2 \in \{0, \ldots, a-1\}$	$P'_{t_1} \otimes \|1\rangle_{t_2+1} \langle 0\|_{t_2+1}$	d
		$t_2 = a$ (if $t_1 \neq \binom{a}{d-1}$)	$F'_{t_1+1} \otimes \|1\rangle_a \langle 1\|_a$	$d+1$
	even	$t_2 \in \{1, \ldots, a\}$	$P'_{t_1} \otimes \|0\rangle_{t_2} \langle 1\|_{t_2}$	d
		$t_2 = 0$	$F'_{t_1+1} \otimes \|0\rangle_1 \langle 0\|_1$	$d+1$
$\mathcal{B}_t$	odd	$t_2 \in \{0, \ldots, a-2\}$	$P'_{t_1} \otimes \|11\rangle_{t_2+1,t_2+2} \langle 00\|_{t_2+1,t_2+2}$	$d+1$
		$t_2 = a-1$ (if $t_1 \neq a$)	$F'_{t_1+1} \otimes \|0\rangle_a \langle 1\|_a$	$d+1$
		$t_2 = a$ (if $t_1 \neq a$)	$F'_{t_1+1} \otimes \|1\rangle_a \langle 0\|_a$	$d+1$
	even	$t_2 \in \{2, \ldots, a-2\}$	$P'_{t_1} \otimes \|00\rangle_{t_2-1,t_2} \langle 11\|_{t_2-1,t_2}$	$d+1$
		$t_2 = 1$	$F'_{t_1+1} \otimes \|0\rangle_1 \langle 1\|_1$	$d+1$
		$t_2 = 0$	$F'_{t_1+1} \otimes \|1\rangle_1 \langle 0\|_1$	$d+1$

Analysis of the whole Hamiltonian. Without loss of generality, we assume that

$$U_t = I \text{ for any } t \text{ of the form } t = (t_1, 0) \text{ with } t_1 \text{ even or } t = (t_1, a) \text{ with } t_1 \text{ odd.} \qquad (\star)$$

The assumption can be fulfilled by adding the dummy gates I at such time steps. We add those time steps such that $U_t = I$ to T_A, i.e., we consider the gates $U_t = I$ as 1-qubit gates. The locality of each term of the whole Hamiltonian is represented in the following table (note that we crucially use Assumption $(\star)$ to guarantee that the locality of $H_{\text{prop},t}$ and $H_{\text{qubit},t}$ is $d+1$):

Term	Locality
H_{in}	$1+d = d+1$
H_{out}	$1+d = d+1$
$H_{\text{prop},t}$	$\max\{d+1, 1+d\} = d+1$
$H_{\text{qubit},t}$	$1+d = d+1$
$H_{\text{time},t}$	$\max\{d, d+1\} = d+1$
H_{clock}	d

The whole Hamiltonian is thus $(d+1)$-local. We can easily check that Conditions $(\text{C}_1) \sim (\text{C}_6)$ are satisfied. From Proposition 18, we conclude that the circuit-to-Hamiltonian construction is correct. From the condition $(\star)$, T is $\binom{a}{d-1}a$. Therefore, we can take $a = O(T^{1/d})$ and the number of qubits H_x acting on is $N + O(T^{1/d})$. ◀

4 Fine-grained hardness for local Hamiltonian problem

In this section, we are going to prove the lower bound for the 3-LH problem.

▶ **Theorem 22** (Lower bound for 3-LH). *Assume that SETH (resp. QSETH) holds. Then, for any constant $\xi > 0$, for any classical (resp. quantum) algorithm $\mathcal{A}_{LH}$, there exists a 3-local Hamiltonian H acting on n_H qubits associated with $E_{\text{yes}}, E_{\text{no}}$ satisfying $E_{\text{no}} - E_{\text{yes}} \geq \Omega(1)$ such that $\mathcal{A}_{\text{LH}}$ cannot decide $\lambda(H) \geq E_{\text{no}}$ or $\lambda(H) \leq E_{\text{yes}}$ with probability greater than $2/3$ in $O(2^{n_H(1-\xi)})$ (resp. $O(2^{n_H(1-\xi)/2})$) time.*

We will prove Theorem 22 in Section 4.1. To prove it, we will use the size-preserving circuit-to-Hamiltonian construction (Theorem 17) associated with the following lemma.

▶ **Lemma 23** (A quantum circuit can compute k-CNF). *For any positive integer k and $c > 0$, for all $n \in \mathbb{N}$, for any k-CNF formula Φ that contains n variables and $m = O(n^c)$ clauses, there exists a quantum circuit U_Φ that acts on n input qubits and at most $2c\log n + 2$ ancilla qubits, and U_Φ consists of at most $34c^2n^c\log^2 n + (70k+2)n^c + 35c\log n$ elementary gates such that $U_\Phi |x\rangle_{\mathsf{in}} \otimes |0\rangle_{\mathsf{anc}} = |\Phi(x)\rangle_{\mathsf{out}} \otimes |\psi_x\rangle_{\overline{\mathsf{out}}}$ for all $x \in \{0,1\}^n$, where $|\psi_x\rangle$ is some quantum state depending on x. The construction of U_Φ can be done in $\text{poly}(n)$ time.*[13]

We defer the proof of Lemma 23 to Section 4.2.

[13] It is worth noting that an alternative quantum circuit for computing kCNF appears in [27], Corollary 10. That construction has width $k+2\big(\lceil \log n\rceil + \lceil c\log n\rceil\big)$ and size at most $4\times 3^{\lceil \log_2 n\rceil + \lceil c\log_2 n\rceil - 1}$. However, its circuit size is $O(n^{c+1})$, whereas ours is $O(n^c \log^2 n)$. This gap impacts the locality of the resulting Hamiltonian, which is why Lemma 23 is necessary for our results.

4.1 SETH- and QSETH-hardness of the 3-local Hamiltonian problem

Proof of Theorem 22. Let $\Phi = \varphi_1 \wedge \varphi_2 \wedge \cdots \wedge \varphi_m$ be a kCNF formula defined on n variables $x_1, x_2, \ldots, x_n$, where $m = O(n^c)$ and $c \in [1, 2)$. We construct an algorithm $\mathcal{A}_{\mathrm{SAT}}$ that solves k-SAT(Φ) by using $\mathcal{A}_{\mathrm{LH}}$ as a subroutine.

1. Upon receiving a k-SAT instance Φ, construct a quantum circuit U_Φ from Lemma 23.
2. Construct a Hamiltonian H_{U_Φ} acting on $n_H = n + o(n)$ qubits from U_Φ, and obtain two energy thresholds $E_{\mathrm{yes}}, E_{\mathrm{no}}$ where $E_{\mathrm{no}} - E_{\mathrm{yes}} = \Omega(1)$ by Theorem 17.
3. Run $\mathcal{A}_{\mathrm{LH}}$ on the input H_{U_Φ}. If $\lambda(H_{U_\Phi}) \leq E_{\mathrm{yes}}$, then return YES. If $\lambda(H_{U_\Phi}) \geq E_{\mathrm{no}}$, then return NO.

The running time of $\mathcal{A}_{\mathrm{SAT}}$ is poly(n).

We now show the correctness of $\mathcal{A}_{\mathrm{SAT}}$. By Lemma 23, we have the following:

- If there exists an assignment $x \in \{0,1\}^n$ such that $\Phi(x) = 1$, then $\Pr[1 \leftarrow U_\Phi(|x\rangle)] = 1$.
- If $\Phi(x) \neq 1$ for all assignments $x \in \{0,1\}^n$, then $\Pr[1 \leftarrow U_\Phi(|x\rangle)] = 0$ for all $x \in \{0,1\}^n$.

Combining Theorem 17 and choosing $\varepsilon_{\mathrm{err}} < \frac{1}{4}$, we obtain:

- $||H_{U_\Phi}|| \leq \mathrm{poly}(n_H)$.
- If there exists $x \in \{0,1\}^n$ such that $\Phi(x) = 1$, then $\lambda(H_{U_\Phi}) \leq E_{\mathrm{yes}}$.
- If $\Phi(x) \neq 1$ for all $x \in \{0,1\}^n$, then $\lambda(H_{U_\Phi}) \geq E_{\mathrm{no}}$.
- $E_{\mathrm{no}} - E_{\mathrm{yes}} = \Omega(1)$

Therefore, $\mathcal{A}_{\mathrm{SAT}}$ has the same success probability as $\mathcal{A}_{\mathrm{LH}}$.

Next, we show that $n_H = n + o(n)$ and that H_{U_Φ} is 3-local. Let n_a and g denote the number of ancilla qubits and elementary gates in U_Φ, respectively. By Lemma 23, $n_a = o(n)$, and $g \leq 34c^2 n^c \log^2 n + (70k + 2)n^c + 35c \log n$. For all sufficiently large n, the gate number g is upper-bounded by $35c^2 n^{c'}$, where $c' \in (c, 2)$ is a constant.

By Theorem 17, we can apply $d = 2$ and the number of qubits of H_{U_Φ} is $n_H = n + O(\log n) + o(n) = n + o(n)$ and H_{U_Φ} is 3-local. Furthermore, we have that the threshold gap $E_{\mathrm{no}} - E_{\mathrm{yes}} = \Omega(1)$.

Because $n_H = n + o(n)$, we have $n_H \leq (1 + \eta)n$ where $\eta > 0$ is an arbitrary small constant for all sufficiently large n. For any constant $\xi > 0$, we can take $\xi > \eta > 0$ and a constant $\varepsilon = \xi - \eta + \xi\eta$ such that $\varepsilon > 0$, and it holds that $O(2^{n(1+\eta)(1-\xi)}) \leq O(2^{n(1-\varepsilon)})$ (and $O(2^{\frac{n}{2}(1+\eta)(1-\xi)}) \leq O(2^{\frac{n}{2}(1-\varepsilon)})$).

Assume there exists a classical (resp. quantum) algorithm $\mathcal{A}_{\mathrm{LH}}$ that solves 3-LH($H, E_{\mathrm{yes}}, E_{\mathrm{no}}$) for $E_{\mathrm{no}} - E_{\mathrm{yes}} \geq \Omega(1)$ in $O(2^{n_H(1-\xi)})$ (resp. $O(2^{\frac{n_H}{2}(1-\xi)})$) time for a constant $\xi > 0$. Then we have the classical (resp. quantum) algorithm $\mathcal{A}_{\mathrm{SAT}}$ solves k-SAT defined on n variables in $O(2^{n(1-\varepsilon)})$ (resp. $O(2^{\frac{n}{2}(1-\varepsilon)})$) time for a constant $\varepsilon > 0$. Besides, the reduction is polynomial time. As a result, it violates SETH which is defined in Conjecture 14 (resp. QSETH which is defined in Conjecture 15). By taking a contraposition, we have the statement. ◀

4.2 Efficient verification circuit for SAT

In this section, we show the construction of the quantum circuit U_Φ that computes the formula $\Phi = \varphi_1 \wedge \cdots \wedge \varphi_m$ defined on n variables $x_1, x_2, \ldots x_n$, and $m = O(n^c)$. For each $i \in [m]$, the clause $\varphi_i = \ell_{i,1} \vee \ell_{i,2} \vee \cdots \vee \ell_{i,k_i}$ where for each $p \in [k_i]$, $\ell_{i,p}$ can be either x_j or $\neg x_j$ where $j \in [n]$, and $k_i \leq k$. Let $r := \lceil \log m \rceil$.

Proof of Lemma 23. The ideal is to compute φ_i for each $i \in [m]$ sequentially. We initially set a counter, and increment the counter by one if $\varphi_i(x) = 1$. After compute all φ_i, we check whether the counter is equal to m. Let k_i be the number of variables (x_j or $\neg x_j$) appearing

in φ_i. We have $\Phi(x) = 1$ if and only if the counter equals m. We introduce the ancilla register $\mathsf{anc} = \mathsf{cls} \cup \mathsf{cnt} \cup \mathsf{out}$ where cls is a one-qubit register that temporarily stores the value of $\varphi_i(x)$ for each i, cnt is an r-qubit ($r = \lceil \log m \rceil$) that serves as the counter, and out is a one-qubit register that outputs $\Phi(x)$.

To compute each φ_i, we construct W_i acting on $\mathsf{in} \cup \mathsf{cls}$ such that $W_i |x\rangle_{\mathsf{in}} |0\rangle_{\mathsf{cls}} = |\psi_x\rangle_{\mathsf{in}} |\varphi_i\rangle_{\mathsf{cls}}$, where $|\psi_x\rangle$ is some state depending on x.

To increment the counter by one, we construct a unitary ADDONE acting on cnt such that $\text{ADDONE} |y\rangle_{\mathsf{cnt}} = |bin(int(y) + 1)\rangle_{\mathsf{cnt}}$ for all $y \in \{0,1\}^r \setminus \{1^r\}$. We apply ADDONE only when $\varphi_i(x) = 1$. This is done by letting ADDONE be controlled by the resister cls. Denote the control-ADDONE operator by $C\text{ADDONE}$. Then $C\text{ADDONE}_{\mathsf{cls}\cup\mathsf{cnt}} |0\rangle_{\mathsf{cls}} |y\rangle_{\mathsf{cnt}} = |0\rangle_{\mathsf{cls}} |y\rangle_{\mathsf{cnt}}$ and $C\text{ADDONE}_{\mathsf{cls}\cup\mathsf{cnt}} |1\rangle_{\mathsf{cls}} |y\rangle_{\mathsf{cnt}} = |0\rangle_{\mathsf{cls}} |bin(int(y) + 1)\rangle_{\mathsf{cnt}}$.

After calculating $\varphi_i(x)$ and applying $C\text{ADDONE}$, we apply $W_i^\dagger$ on $\mathsf{in} \cup \mathsf{cls}$ to restore the state to $|x\rangle_{\mathsf{in}} |0\rangle_{\mathsf{cls}}$.

To check whether the counter is equal to m in the final step, we construct a compare operator, denoted by COMPARE, that acts on $\mathsf{cnt} \cup \mathsf{out}$. The operator COMPARE satisfies $\text{COMPARE} |y\rangle_{\mathsf{cnt}} |0\rangle_{\mathsf{out}} = |y\rangle_{\mathsf{cnt}} |\delta_{m,int(y)}\rangle_{\mathsf{out}}$ for all $y \in \{0,1\}^r$.

To sum up, we construct the quantum circuit U_Φ as follows:

$$U_\Phi := \text{COMPARE}(W_m^\dagger C\text{ADDONE} W_m)(W_{m-1}^\dagger C\text{ADDONE} W_{m-1}) \cdots (W_1^\dagger C\text{ADDONE} W_1). \quad (23)$$

Next, we explain how to implement W_i, ADDONE, and COMPARE gates.

Because $\ell_{i_1} \vee \ell_{i_2} \vee \cdots \vee \ell_{i_{k_i}} = \neg(\neg\ell_{i_1} \wedge \neg\ell_{i_2} \wedge \cdots \wedge \neg\ell_{i_{k_i}})$, we can use $C^k\text{NOT}$ together with NOT gates to implement W_i. Let $\mathsf{S_i} \subseteq \mathsf{in}$ be the register defined by $\mathsf{S_i} := \{j \in [n] : x_j \in \varphi_i \text{ or } \neg x_j \in \varphi_i\}$, i.e., $\mathsf{S_i}$ consists of the jth qubits in in such that x_j or $\neg x_j$ in φ_i; and let $\mathsf{R_i} \subseteq \mathsf{in}$ be the register defined by $\mathsf{R_i} := \{j \in [n] : x_j \in \varphi_i\}$, i.e., $\mathsf{R_i}$ consists of the jth qubits in in such that x_j in φ_i, but $\neg x_j$ does not. We construct W_i by $W_i := \text{NOT} \cdot_{\mathsf{cls}} C^{k_i}\text{NOT}_{\mathsf{S_i}\cup\mathsf{cls}} \cdot \text{NOT}_{\mathsf{R_i}}^{\otimes|\mathsf{R_i}|}$.

To implement the ADDONE, observe that for $y \in \{0,1\}^r$, when $int(y)$ is incremented by one, a bit $y[p]$ will be flipped if and only if all the bits with order lower than p are 1's. That is, $y[p]$ will be flipped if and only if for all $p' > p$, $y[p'] = 1$. Hence, ADDONE can be composed of a sequence of multi-control-NOT gates. The qth layer of ADDONE is a $C^{r-q}\text{NOT}$ gate whose control qubits are $q, q+1, \ldots, r$ and whose target is the $q$the qubit. The last layer of ADDONE is a NOT gate acting on the last qubit. To implement the control-ADDONE operation $C\text{ADDONE}$, we let every $C^{r-q}\text{NOT}$ be controlled by register cls. In other words, the qth layer of $C\text{ADDONE}$ is a $C^{r+1-q}\text{NOT}$ gates whose control qubits are $q, q+1, \ldots, r \cup \mathsf{cls}$ and whose target is the qth qubit in cnt.

The COMPARE operator checks whether the value stored in the counter equal to m. That is, COMPARE operator compares each bit in the counter with $bin(m)$, which can be implemented by a $C^r\text{NOT}$ gate. Let $\mathsf{P} \subseteq \mathsf{cnt}$ be defined by $\mathsf{P} := \{j \in [r] : bin(m)[j] = 0\}$. We construct COMPARE by $\text{COMPARE} := C^r\text{NOT}_{\mathsf{cnt}\cup\mathsf{out}} \cdot \text{NOT}_{\mathsf{P}}^{\otimes|P|}$.

We decompose each multi-control-NOT gate into Toffoli gates. For W_i (and $W_i^\dagger$), there is a $C^{k_i}\text{NOT}$ inside, and it can be decomposed into at most $2k$ of Toffoli gates using at most k ancilla bits.

The largest gate in the $C\text{ADDONE}$ is a control $C^r\text{NOT}$ gate, and there are r layers. Hence, $C\text{ADDONE}$ can be decomposed into at most $2r^2$ Toffoli gates. Because the ancillas for the multi-control-NOT gates can be reused, the ancillas required for $C\text{ADDONE}$ is at most r.

The COMPARE operation contains a $C^r\text{NOT}$ gate, which can be decomposed into at most $2r$ Toffoli gates using at most r ancilla qubits.

The ancilla qubits can be reused for the multi-control-NOT gates. Hence, to decomposed U_Φ into Toffoli gates, the number of ancilla qubits required is at most r. Therefore, the total number of ancilla qubits for U_Φ is $|\mathsf{anc}| + r = |\mathsf{cls}| + |\mathsf{cnt}| + |\mathsf{out}| + r = 2\log m + 2 = 2c \log n + 2$.

The total number of Toffoli gates in U_Φ is at most $(4k+2r^2)\cdot m + 2r = 2c^2n^c\log^2 n + 4kn^c + 2c\log n$. We further decompose the Toffoli gates into elementary gates, yielding $34c^2n^c\log^2 n + 68kn^c + 34c\log n$ elementary gates in total. In addition to Toffoli gates, there are at most $2(k+1)\cdot m + r = (2k+2)n^c + c\log n$ of NOT gates. Therefore, the total number of gates is at most $34c^2n^cc^2\log^2 n + (70k+2)n^c + 35c\log n$. ◀

5 Fine-grained complexity of quantum partition function problem

The fine-grain lower bound for the local Hamiltonian problem directly gives us the fine-grain lower bound for the quantum partition function.

5.1 SETH- and QSETH-hardness of quantum partition functions

▶ **Lemma 24** (Fine-grained reduction from k-LH to k-QPF). *For any $T(n) \in \omega(\mathrm{poly}(n))$, k-LH$(H, E_{\mathrm{yes}}, E_{\mathrm{no}})$ is $(T(n), T(n))$ reducible to k-QPF(H, β, δ) in which H acts on n qubits, $\beta \geq \frac{n}{E_{\mathrm{no}}-E_{\mathrm{yes}}}$ and δ satisfies that $\frac{1-\delta}{1+\delta} \geq e^{-0.3n}$.*

We emphasize that in Lemma 24, H in the k-QPF problem and the k-LH problem is the same Hamiltonian.

To better understand Lemma 24, we unpack the underlying fine-grained reduction as follows. Suppose there exist $\xi > 0$ and an algorithm $\mathcal{A}_{\mathrm{QPF}}$ that approximate $\mathrm{tr}[e^{-\beta H}]$ up to relative error δ satisfying $\frac{1-\delta}{1+\delta} \geq e^{-0.3n}$ with probability greater than $2/3$ in $O(T(n)^{1-\xi})$ time, then for any $\varepsilon > 0$, we can use $\mathcal{A}_{\mathrm{QPF}}$ to construct an algorithm $\mathcal{A}_{\mathrm{LH}}$ such that given thresholds $E_{\mathrm{yes}}, E_{\mathrm{no}}$ satisfying $E_{\mathrm{no}} - E_{\mathrm{yes}} \geq n/\beta$, the algorithm $\mathcal{A}_{\mathrm{LH}}$ decides weather $\lambda(H) \geq E_{\mathrm{no}}$ or $\lambda(H) \leq E_{\mathrm{yes}}$ with probability greater than $2/3$ in $O(T(n)^{1-\varepsilon})$ time.

On the other hand, suppose for all $\varepsilon > 0$, for any algorithm $\mathcal{A}_{\mathrm{LH}}$, for infinitely many n there exists a k-local Hamiltonian H_n acting on n qubits and two energy thresholds $E_{\mathrm{yes}}, E_{\mathrm{no}}$ such that $\mathcal{A}_{\mathrm{LH}}$ cannot solve k-LH$(H_n, E_{\mathrm{yes}}, E_{\mathrm{no}})$ with probability greater than $\frac{2}{3}$ in $O(T(n)^{1-\varepsilon})$ time, then for all $\xi > 0$, for all $\delta > 0$, for infinitely many $n \geq n_1$, any algorithm $\mathcal{A}_{\mathrm{QPF}}$ cannot solve k-QPF(H_n, β, δ), where $\beta \geq \frac{n}{E_{\mathrm{no}}-E_{\mathrm{yes}}}$, with probability greater than $2/3$ in $O(T(n)^{1-\xi})$ time. The integer n_1 satisfies that $\frac{1-\delta}{1+\delta} = e^{-0.3n_1}$.

Proof of Lemma 24. Let $(H, E_{\mathrm{yes}}, E_{\mathrm{no}})$ be the k-LH instance, where H is a k-local Hamiltonian acts on n qubits . We are going to construct an algorithm $\mathcal{A}_{\mathrm{LH}}$ that solves k-LH$(H, E_{\mathrm{yes}}, E_{\mathrm{no}})$ by using $\mathcal{A}_{\mathrm{QPF}}$ as a subroutine.

If $(H, E_{\mathrm{yes}}, E_{\mathrm{no}})$ is YES case, then there is at least one eigenstate of H whose energy is lower than or equal to E_{yes}. Hence $Z(\beta) \geq e^{-\beta E_{\mathrm{yes}}}$.

If $(H, E_{\mathrm{yes}}, E_{\mathrm{no}})$ is NO case, then all 2^n number of eigenstates of H have energy higher than or equal to E_{no}. Hence, $Z(\beta) \leq 2^n e^{-\beta E_{\mathrm{no}}} < e^{-\beta E_{\mathrm{no}}+0.7n}$.

When $\beta \geq \frac{n}{E_{\mathrm{no}}-E_{\mathrm{yes}}} \in \mathrm{poly}(n)$, and δ satisfies $\frac{1-\delta}{1+\delta} \geq e^{-0.3n}$, it holds that

$$\frac{1-\delta}{1+\delta} \geq e^{-0.3n} \geq e^{-\beta(E_{\mathrm{no}}-E_{\mathrm{yes}})+0.7n}.$$

Hence we have $(1-\delta)e^{-\beta E_{\mathrm{yes}}} \geq (1+\delta)e^{-\beta E_{\mathrm{no}}+0.7n}$.

Now we present the algorithm $\mathcal{A}_{\mathrm{LH}}$ that solves k-LH$(H, E_{\mathrm{yes}}, E_{\mathrm{no}})$ by using $\mathcal{A}_{\mathrm{QPF}}$.

1. When receiving $H, E_{\mathrm{yes}}, E_{\mathrm{no}}$, calculate δ_0 such that $\frac{1-\delta_0}{1+\delta_0} = e^{-0.3n}$ and $\beta_0 = \frac{n}{E_{\mathrm{no}}-E_{\mathrm{yes}}}$. Set $\beta \geq \beta_0$ and set $\delta \leq \delta_0$.
2. Run $\mathcal{A}_{\mathrm{QPF}}$ on the input H, β, δ, and get the output $\widetilde{Z}$.
3. If $\widetilde{Z} \geq (1-\delta)e^{-\beta E_{\mathrm{yes}}}$, then output YES. If $\widetilde{Z} \leq (1+\delta)e^{-\beta E_{\mathrm{no}}+0.7n}$, then output NO.

When $\mathcal{A}_{\mathrm{QPF}}$ solves k-QPF(H,β,δ) successfully, it is guaranteed that $(1-\delta)Z(\beta) \leq \widetilde{Z} \leq (1+\delta)Z(\beta)$. When the inputs of k-LH is YES case, $\widetilde{Z} \geq (1-\delta)Z(\beta) \geq (1-\delta)e^{-\beta E_{\mathrm{yes}}}$; and when NO case, $\widetilde{Z} \leq (1+\delta)Z(\beta) < (1+\delta)e^{-\beta E_{\mathrm{no}}+0.7n}$.

By the choice of β and δ, it holds that $(1-\delta)e^{-\beta E_{\mathrm{yes}}} \geq (1+\delta)e^{-\beta E_{\mathrm{no}}+0.7n}$. Therefore, when $\mathcal{A}_{\mathrm{QPF}}$ solves k-QPF(H,β,δ) successfully, $\mathcal{A}_{\mathrm{LH}}$ decides k-LH$(H, E_{\mathrm{yes}}, E_{\mathrm{no}})$ correctly. The success probability of $\mathcal{A}_{\mathrm{LH}}$ is the same as $\mathcal{A}_{\mathrm{QPF}}$. Therefore, the requirement 1 in Definition 16 is satisfied.

The algorithm $\mathcal{A}_{\mathrm{LH}}$ queries $\mathcal{A}_{\mathrm{QPF}}$ once in Step 2., and the instance of k-QPF is exactly the input of k-LH. For any ε, we choose $\xi = \varepsilon$. We have $T(n)^{1-\varepsilon} \leq d \cdot T(n)^{1-\xi}$ for any constant $d > 1$. Therefore, the requirement 3 and 4 in Definition 16 are satisfied.

Finally, Step 1. and Step 3. in the algorithm $\mathcal{A}_{\mathrm{LH}}$ run in poly(n) time. Hence, $\mathcal{A}_{\mathrm{LH}}$ runs in poly(n) times, which is less than $d \cdot T(n)^{1-\xi}$ for some constant d because $T(n)$ is super-polynomial. consequently, the requirement 2 is satisfied. This finishes the proof. ◀

▶ Remark 25. The reduction in Lemma 24 is polynomial-time. Since k-LH is QMA-hard, we have that k-QPF is also QMA-hard.

The fine-grain reduction above implies the lower bound for the quantum partition function problem.

▶ **Theorem 26.** *Assume SETH (resp. QSETH) holds. For any $\xi > 0$ and any $\delta \in (0,1)$, for any quantum algorithm $\mathcal{A}_{\mathrm{QPF}}$, for infinitely many n, there exists a 3-local Hamiltonian H acting on n qubits, associated with an inverse temperature $\beta_0 = O(n)$, such that for all $\beta \geq \beta_0$, the algorithm $\mathcal{A}_{\mathrm{QPF}}$ cannot solve* QPF(H,β,δ) *in $O(2^{n(1-\xi)})$ classical time (resp. $O(2^{\frac{n}{2}(1-\xi)})$ quantum time) with probability greater than $2/3$.*

Proof. We prove the theorem by contradiction. Assume that there exists an algorithm $\mathcal{A}_{\mathrm{QPF}}$ that solves the QPF problem. Then, by Lemma 24, we can construct an algorithm $\mathcal{A}_{\mathrm{LH}}$ that solves 3-local Hamiltonian problem, which contradicts Theorem 22.

Suppose there exist $\xi > 0$ and $\delta \in (0,1)$ such that there exist an algorithm $\mathcal{A}_{\mathrm{QPF}}$ and $n_1 \in \mathbb{N}$ satisfying the following: for all $n > n_1$, for all H acting on n qubits, and for arbitrarily large $\beta \geq O(n)$, the algorithm $\mathcal{A}_{\mathrm{QPF}}$ solves QPF(H,β,δ) with probability greater than $2/3$ in $O(2^{n(1-\xi)})$ classical time.

Let $(H, E_{\mathrm{yes}}, E_{\mathrm{no}})$ be an LH instance, where H is a 3-local Hamiltonian acting on n qubits with sufficiently large n (we will specify how large n shall be later), and $E_{\mathrm{no}} - E_{\mathrm{yes}} = \Omega(1)$. Choose β such that $\beta \geq O(n) > \frac{n}{E_{\mathrm{no}} - E_{\mathrm{yes}}}$. By Lemma 24, we construct $\mathcal{A}_{\mathrm{LH}}$, in which we query $\mathcal{A}_{\mathrm{QPF}}$ on the instance (H,β,δ). The algorithm runs in $O(2^{n(1-\xi)})$ time.

Let n_2 be the smallest integer such that $\frac{1-\delta}{1+\delta} \geq e^{-0.3n_2}$. By the hypothesis at the beginning of the proof and by Lemma 24, when $n \geq \max\{n_1, n_2\}$, the algorithm $\mathcal{A}_{\mathrm{LH}}$ decides $(H, E_{\mathrm{yes}}, E_{\mathrm{no}})$ successfully with probability greater than $2/3$.

Note that the reduction works for all H acting on n qubits and $E_{\mathrm{yes}}, E_{\mathrm{no}}$ satisfying $E_{\mathrm{no}} - E_{\mathrm{yes}} = \Omega(1)$ as long as $n \geq \max\{n_1, n_2\}$. We thus obtain a contradiction with Theorem 22. The lower bound of quantum algorithms assuming QSETH follows in the same proof strategy. ◀

5.2 $O^*(2^{\frac{n}{2}})$-time quantum algorithm for k-QPF problem

In this section, we propose an algorithm that estimates the quantum partition function under a polynomial inverse temperature up to an arbitrary inverse polynomial error in $O^*(2^{\frac{n}{2}})$ time. The running time matches our conditional lower bound. The input Hamiltonian H is n-qubit, constant-local, semi-positive and satisfies $\|H\| < 1$.

▶ **Theorem 27** ($O^*(2^{\frac{n}{2}})$ time algorithm for k-QPF). *For all sufficiently large $n \in \mathbb{N}$, for all n-qubit constant-local, semi-definite Hamiltonian H satisfying $\|H\| < 1$, $\beta < \mathrm{poly}(n)$, and $\delta \geq 1/n^c$ for arbitrary positive constant c, there exists a quantum algorithm that solves k-QPF(H, β, δ) in $O^*(2^{\frac{n}{2}})$ time with successful probability $1 - \mathrm{negl}(n)$.*

Tools that are used in the proof of QPF algorithm

Before proving Theorem 27, we first list some tools that will be used in the proof.

The main idea of the algorithm is dividing the energy range $[0, 1)$ into polynomial many disjoint intervals $I_0, I_1, \ldots, I_L$. For each interval I_ℓ, we designate a value $\mathcal{E}_\ell \in I_\ell$ to represent the characteristic energy of I_ℓ. Let M_ℓ be the number of the Hamiltonian's eigenstates whose corresponding energy is in the interval I_ℓ. The approximation of the quantum partition function is

$$\sum_{\ell=0}^{L} M_\ell e^{-\beta \mathcal{E}_\ell}.$$

To estimate the number of eigenstates M_ℓ, there are two key ingredients: (i) energy estimation, and (ii) quantum counting.

Given the description of a local Hamiltonian H and one of H's eigenstate $|\psi_j\rangle$, the energy estimation algorithm outputs its corresponding energy E_j. The energy estimation algorithm can be implemented by phase estimation associated with Hamiltonian simulation.

▶ **Lemma 28** (Phase estimation [44]). *Let U be a n-qubit unitary and $|u_\theta\rangle$ be U's eigenstate satisfying that $U|u_\theta\rangle = e^{i\theta}|u_\theta\rangle$. There exists a quantum circuit $P_{U,r}$, which takes $|u_\theta\rangle$ as the input state and executes control-U operator $O(2^r)$ times with additional $\mathrm{poly}(n)$ number of gates, such that the output state is*

$$|u_\theta\rangle \otimes \sum_{x \in \{0,1\}^r} \alpha_x |x\rangle, \tag{24}$$

where $\alpha \in \mathbb{C}$, satisfying that

$$\sum_{x : |\frac{int(x)\cdot 2\pi}{2^r} - \theta| \leq \frac{2\pi}{2^b}} |\alpha_x|^2 \geq 1 - \frac{1}{2^{r-b}} \tag{25}$$

for any $b < r - 1$.

When we choose $b = r - 2$, we have that the estimation $\widetilde{\theta} := \frac{int(x)\cdot 2\pi}{2^r}$ is in the interval $\theta \pm \frac{2\pi}{2^b}$ with probability greater than $\frac{3}{4}$. We can amplify the probability to $1 - \mathrm{negl}(n)$ by repeating the procedure $\mathrm{poly}(n)$ times and using median of means technique [4, 10].

▶ **Lemma 29** (Confidence amplification of phase estimation). *Following Lemma 28, there exists a quantum circuit $A_{U,r,m}$, which executes $P_{U,r}$ $m = \mathrm{poly}(n)$ times, with additional $\mathrm{poly}(n)$ gates, such that the output state is*

$$|u_\theta\rangle \otimes \sum_{x \in \{0,1\}^r} \alpha'_x |x\rangle |g_x\rangle, \tag{26}$$

where $\alpha' \in \mathbb{C}$, and $|g_x\rangle$ is some state depending on x, satisfying

$$\sum_{x : |\frac{int(x)\cdot 2\pi}{2^r} - \theta| \leq \frac{2\pi}{2^{(r-2)}}} |\alpha'_x|^2 \geq 1 - e^{-\Theta(m)}. \tag{27}$$

Proof. Because after the execution of the phase estimation $P_{U,r}$, the input state $|u_\theta\rangle$ remains unchanged and separable from $\sum_x \alpha_x |x\rangle$, we can reuse the input state. Assume we measure the register that contains $\sum_x \alpha_x |x\rangle$ after each execution of $P_{U,r}$. Let the outcome of each execution be $x_1, x_2, \ldots, x_m$. By Lemma 28, when we choose $b = r-2$, it holds that for each $i \in [m]$ the probability that $\widetilde{\theta_i} := \frac{int(x_i)\cdot 2\pi}{2^r}$ is in the interval $\theta \pm \frac{2\pi}{2^{r-2}}$ is greater than $\frac{3}{4}$.

Because the register that stores $\sum_x \alpha_x |x\rangle$ for each execution are separable, we have that the measurement outcomes of all the execution are independent. Then, by Chernoff bound, the probability that there are more than half of $\widetilde{\theta_1}, \widetilde{\theta_2}, \ldots, \widetilde{\theta_m}$ outside the interval $\theta \pm \frac{2\pi}{2^{r-2}}$ is less than $e^{-\Theta(m)}$. Let $\widetilde{x}$ be the the median of $int(x_1), int(x_2), \ldots, int(x_m)$ and let $\widetilde{\theta} := \frac{\widetilde{x}\cdot 2\pi}{2^r}$. We have the event that $\widetilde{\theta}$ is outside the interval $\theta \pm \frac{2\pi}{2^{r-2}}$ has probability less than $e^{-\Theta(m)}$.

Now we describe the quantum circuit $A_{U,r,m}$. The quantum circuit $A_{U,r,m}$ has m ancilla registers $\mathsf{X}_1, \mathsf{X}_2, \ldots, \mathsf{X}_m$ each of them is r-qubit, and an input register U takes the input state $|u_\theta\rangle$. The quantum circuit $A_{U,r,m}$ executes $P_{U,r}$ m times sequentially on the input $|u_\theta\rangle$ and in the ith execution of $P_{U,r}$, it stores $\sum_x \alpha_x |x\rangle$ in the register X_i. After m times execution, a circuit that computes the median of $\mathsf{X}_1, \mathsf{X}_2, \ldots, \mathsf{X}_m$ is applied. Since $\mathsf{X}_1, \mathsf{X}_2, \ldots, \mathsf{X}_m$ are separable and computing the median is a classical circuit, the distribution of $\widetilde{x}$ is the same regardless of the measurement taking place before or after the computing of median. As a result, Equation Equation (27) holds. ◀

To implement the energy estimation of the local Hamiltonian H, we run phase estimation for $U = e^{-iH}$, which can be implemented by Hamiltonian simulation.

▶ **Lemma 30** (Hamiltonian simulation [41]). *For all sufficiently large $n \in \mathbb{N}$, for all constant-local Hamiltonian $H \in \mathbb{C}^{2^c \times 2^n}$, evolution time $t \in \mathbb{R}^+$, and tolerated error $\varepsilon \in (0,1)$, there exists a quantum circuit $U_{\mathrm{sim}}(H,t,\varepsilon)$ running in $\mathrm{poly}(n,t,\log(1/\varepsilon))$ such that for any quantum state $|\psi\rangle \in \mathbb{C}^{2^n}$, it holds that*

$$\|(U_{\mathrm{sim}}(H,t,\varepsilon) - e^{-iHt} \otimes I)\,|\psi\rangle \otimes |0^{a_{\mathrm{sim}}}\rangle\,\|_2 \leq \varepsilon, \tag{28}$$

where a_{sim} is the number of ancilla qubits of $U_{\mathrm{sim}}(H,t,\varepsilon)$.

Now we formally define the energy estimation.

▶ **Definition 31** (Energy estimation). *Let $H \in \mathbb{C}^{2^n \times 2^n}$ be a constant-local Hamiltonian satisfying that all the eigenvalues of H are in the interval $[0,1)$. We first define an ideal unitary $\widehat{U}_{\mathrm{EE}}(H,\delta E,\eta)$ that estimates the corresponding energy of H's eigenstate $|\psi_j\rangle$. We say $\widehat{U}_{\mathrm{EE}}(H,\delta E,\eta)$ is an ideal energy estimation of the Hamiltonian H if for any H's eigenstate $|\psi_j\rangle$, it holds that*

$$\widehat{U}_{\mathrm{EE}}(H,\delta E,\eta)\,|\psi_j\rangle\,|0^{a_{\mathrm{EE}}}\rangle = |\psi_j\rangle \otimes \sum_{E'} \alpha_{E'}\,|E'\rangle\,|g_{E'}\rangle\,, \tag{29}$$

where a_{EE} is the number of ancilla qubits, $\alpha_{E'} \in \mathbb{C}$, and $|g_{E'}\rangle$ is some state depending on E', satisfying

$$\sum_{E' \in [E_j - \delta E, E_j + \delta E]} |\alpha_{E'}|^2 \geq 1 - \eta, \tag{30}$$

where E_j is the corresponding energy of $|\psi_j\rangle$.

Then, we say a quantum circuit $U_{\mathrm{EE}}(H,\delta E,\eta,\varepsilon)$ is an implementation of energy estimation of H if

$$\|U_{\mathrm{EE}}(H,\delta E,\eta,\varepsilon) - \widehat{U}_{\mathrm{EE}}(H,\delta E,\eta)\| \leq \varepsilon. \tag{31}$$

▶ **Lemma 32** (Efficient energy estimation algorithm). *For all sufficiently large $n \in \mathbb{N}$, For all constant-local, semi-definite Hamiltonian $H \in \mathbb{C}^{2^n \times 2^n}$ satisfying $\|H\| < 1$, and all tolerated error $\delta E = 1/\mathrm{poly}(n)$, there exists a energy estimation implementation $U_{\mathrm{EE}}(H, \delta E, \eta = e^{-n}, \varepsilon = 2^{-n})$ that runs in $\mathrm{poly}(n)$ time.*

Proof. It holds that $e^{-iH}|\psi_j\rangle = e^{-iE_j}|\psi_j\rangle$. We can use phase estimation to obtain E_j. We first consider the ideal unitary $\widehat{U}_{\mathrm{EE}}(H, \delta E, \eta)$. By Lemma 29, we run $A_{U,r,m}$ with $U = e^{-iH}$, $r := \log(\frac{2\pi}{\delta E}) + 2$, and $m = Cn$ where C is a large enough constant, on the input $|\psi_j\rangle$. Let $E' := int(x) \cdot 2\pi/2^r$, we have that Equation Equation (29) and Equation (30) holds with $\eta = e^{-n}$.

Then, we consider the efficient implementation $U_{\mathrm{EE}}(H, \delta E, \eta = e^{-n}, \varepsilon = 2^{-n})$. We replace U in $A_{U,r,m}$ with the implementation of Hamiltonian simulation $U_{\mathrm{sim}}(H, t = 1, \varepsilon/(m \cdot 2^r))$ described in Lemma 30. (To be more precise, we need to implement control-U in $A_{U,r,m}$. Hence, we need to replace the gates in $U_{\mathrm{sim}}(H, t = 1, \varepsilon/(m \cdot 2^r))$ with their controlled version.) There are $m \cdot 2^r$ many of U in $A_{U,r,m}$. By triangular inequality, Equation Equation (31) holds.

For the running time, by Lemma 30, the implementation of $U_{\mathrm{sim}}(H, t = 1, \varepsilon/(m \cdot 2^r))$ takes $\mathrm{poly}(n, 1, \log(/m \cdot 2^r \varepsilon)) = \mathrm{poly}(n, 1, (\log(n) + \delta E + \log(1/\varepsilon)))$ time. We have $\delta E = 1/\mathrm{poly}(n)$ and $\varepsilon = 2^{-n}$, the running time of each implementation of U is $\mathrm{poly}(n)$. There are $m \cdot 2^r = \mathrm{poly}(n)$ many of U in $A_{U,r,m}$. As a result the overall running time of $U_{\mathrm{EE}}(H, \delta E, \eta = e^{-n}, \varepsilon = 2^{-n})$ is $\mathrm{poly}(n)$. ◀

We use quantum counting to estimate the number of eigenstates in a given interval.

▶ **Lemma 33** (Quantum counting). *Let a quantum circuit U taking a n-qubit input state $|\psi\rangle$ satisfy $U|\psi\rangle|0^a\rangle = \sqrt{\frac{M}{N}}|1\rangle|\xi^1\rangle + \sqrt{\frac{N-M}{N}}|0\rangle|\xi^0\rangle$, where $N = 2^n$, $M \leq N$[14], and a is the number of ancilla qubits. For any $c \in \mathbb{N}$, there is a quantum circuit that executes U $O(\mathrm{poly}(n) \cdot 2^{\frac{n}{2}})$ times with additional $O(\mathrm{poly}(n) \cdot 2^{\frac{n}{2}})$ many gates and outputs $\widetilde{M}$ such that $(1 - \Theta(\frac{1}{n^c}))M \leq \widetilde{M} \leq (1 + \Theta(\frac{1}{n^c}))M$ with probability $1 - \mathrm{negl}(n)$.*

In [44], there has already been a proof for the amplitude estimation algorithm that outputs $\widetilde{M}$ satisfying $(1 - \frac{1}{\sqrt{M}})M \leq \widetilde{M} \leq (1 + \frac{1}{\sqrt{M}})M$. To achieve the relative error as small as $\Theta(1/n^c)$, we increase the number of Grover iteration in [44] from $n/2 + 3$ to $n/2 + \Theta(c \log n)$.

To use quantum counting to estimate the number of eigenstates in a given interval, we need to prepare the uniform superposition of all the eigenstates as the input state. However, in general, the eigenstates are unknown. We use the property that EPR state is identical to the uniform superposition of the states and its complex conjugate in any orthogonal basis to overcome this issue.

▶ **Lemma 34.** *For any complete orthogonal basis of n-qubit system $\{|\psi_i\rangle\}_{j=1}^N$,i.e., $\sum_{j=1}^N |\psi_j\rangle\langle\psi_j| = I$ and $\langle\psi_j|\psi_{j'}\rangle = \delta_{j,j'}$ where $N = 2^n$, it holds that*

$$\sum_{j=1}^{N} \sqrt{\frac{1}{N}}|\psi_j\rangle|\psi_j^*\rangle = \sum_{k=1}^{N} \sqrt{\frac{1}{N}}|k\rangle|k\rangle. \tag{32}$$

where $|\psi_j^*\rangle$ is the complex conjugate of $|\psi_j\rangle$.

[14] Actually it requires that $M < N/2$. We can overcome this issue by adding an additional qubit.

Proof. Let $V = \sum_{j,k=1}^{N} V_{k,j}|k\rangle\langle j|$ be a unitary such that $|\psi_j\rangle = V|j\rangle$ for all $j \in [N]$. We have

$$\sum_{j=1}^{N} |\psi_j\rangle \otimes |\psi_j^*\rangle = \sum_{j=1}^{N} \left(\left(\sum_{k=1}^{N} V_{k,j} |k\rangle \right) \otimes \left(\sum_{k'=1}^{N} V_{k',j}^* |k'\rangle \right) \right)$$
$$= \sum_{j=1}^{N} \sum_{k=1}^{N} \sum_{k'=1}^{N} V_{k,j} V_{k',j}^* |k\rangle \otimes |k'\rangle . \tag{33}$$

Because V is a unitary, we have $\sum_{j=1}^{N} V_{k,j} V_{k',j}^* = \delta_{k,k'}$. Equation (33) becomes $\sum_{j=1}^{N} |k\rangle \otimes |k\rangle$. ◀

Proof of Theorem 27

Proof of Theorem 27. Let $N := 2^n$.

For any interval I and $E \in [0,1)$, we define the operator $U_{\mathrm{dec}}(I)|E\rangle|0^a_{\mathrm{dec}}\rangle = |\mathbb{1}_I(E)\rangle|*\rangle$, where a_{dec} is the number of ancilla qubits and $|*\rangle$ is some state. That is, $U_{\mathrm{dec}}(I)|E\rangle|0^a_{\mathrm{dec}}\rangle = |1\rangle|\xi_{E,1}\rangle$ if $E \in I$ and $U_{\mathrm{dec}}(I)|E\rangle|0^a_{\mathrm{dec}}\rangle = |0\rangle|\xi_{E,0}\rangle$ if $E \notin I$, where $|\xi_{E,1}\rangle$ and $|\xi_{E,0}\rangle$ are some states.

We write down the algorithm.

1. Let $L = 4n^c\beta$.
2. For $k \in \{0,1,\ldots,L-1\}$:
   - 2.1 For $\ell \in \{0,1,\ldots,L\}$(If $k=0$, then $\ell \in \{1,2,\ldots,L\}$):
     - 2.1.1 Define the characteristic energy $\mathcal{E}_{k,\ell} := \frac{\ell-1}{L} + \frac{k}{L^2}$, the interval $I_{k,\ell} := [\mathcal{E}_{k,l}, \mathcal{E}_{k,\ell} + \frac{1}{L})$, and $I'_{k,\ell} := [\mathcal{E}_{k,\ell} - 1/2L^2, \mathcal{E}_{k,\ell} + 1/L + 1/2L^2)$.
     - 2.1.2 Define $U_{k,\ell} := U_{\mathrm{dec}}(I'_{k,\ell}) U_{\mathrm{EE}}(H, \delta E = \frac{1}{2L^2}, \eta = e^{-n}, \varepsilon = 2^{-n})$, where U_{EE} is defined in Definition 31.
     - 2.1.3 Prepare the EPR state $|\Psi\rangle := \sum_{k=1}^{N} \sqrt{\frac{1}{N}} |k\rangle|k\rangle$.
     - 2.1.4 Execute quantum counting (Lemma 33) on $U_{k,\ell}|\Psi\rangle|0^{a_U}\rangle$ up to a relative error $\delta_{\mathrm{C}} = \frac{1}{4n^c}$, where $U_{k,\ell}$ acts on the first half of the EPR state and a_U is the number of ancilla qubits. Let $\widetilde{M}_{k,\ell}$ be the output of the quantum counting.
   - 2.2 Let $\widetilde{Z}_k := \sum_{\ell=0}^{L} \widetilde{M}_{k,\ell} e^{-\beta \mathcal{E}_{k,\ell}}$.
3. Let $\widetilde{Z} := \min\{\widetilde{Z}_k\}_{k=0}^{L-1}$.
4. Output $\widetilde{Z}$.

Now we claim the correctness of the algorithm. We define the following notations. Let $I^-_{k,\ell} := [\mathcal{E}_{k,\ell} - 1/L^2, \mathcal{E}_{k,\ell})$, $I^+_{k,\ell} := [\mathcal{E}_{k,\ell} + 1/L, \mathcal{E}_{k,\ell} + 1/L + 1/L^2)$, and $I''_{k,\ell} := I^-_{k,\ell} \cup I_{k,\ell} \cup I^+_{k,\ell}$. Let $Z_{k,\ell} := \sum_{j:E_j \in I_{k,\ell}} e^{-\beta E_j}$ be the partition function that are contributed by the states whose energies are in the interval $I_{k,\ell}$. Also, we define $Z^-_{k,\ell} := \sum_{j:E_j \in I^-_{k,\ell}} e^{-\beta E_j}$ and $Z^+_{k,\ell} := \sum_{j:E_j \in I^+_{k,\ell}} e^{-\beta E_j}$. Let $M_{k,\ell}$ be the number of states whose energies are in the interval $I_{k,\ell}$, $M^-_{k,\ell}$ be the number of states whose energies are in the interval $I^-_{k,\ell}$, and $M^+_{k,\ell}$ be the number of states whose energies are in the interval $I^+_{k,\ell}$.

The quantum partition function is dominated by low energy. The characteristic energy $\mathcal{E}_{k,\ell}$ is lower than all the energy in $I_{k,\ell}$. Hence, $M_{k,\ell} e^{-\mathcal{E}_{k,\ell}}$ gives a upper bound for $Z_{k,\ell}$, i.e.,

$$Z_{k,\ell} \leq M_{k,\ell} e^{-\mathcal{E}_{k,\ell}} \tag{34}$$

for any k, ℓ.

On the other hand, $\mathcal{E}_{k.\ell}$ is greater than all the energy decreased by $1/L$ in $I_{k,\ell}$. Hence, $M_{k,\ell}e^{-\mathcal{E}_{k,\ell}}$ gives a lower bound for $Z_{k,\ell}e^{1/L}$, i.e.,

$$M_{k,\ell}e^{-\mathcal{E}_{k,\ell}} \leq Z_{k,\ell}e^{1/L}. \tag{35}$$

for any k, ℓ.

Similarly, we have

$$M^{-}_{k,\ell}e^{-\beta\mathcal{E}_{k,\ell}} \leq Z^{-}_{k,\ell}, \tag{36}$$

and

$$M^{+}_{k,\ell}e^{-\beta\mathcal{E}_{k,\ell}} \leq Z^{+}_{k,\ell}e^{1/L+1/L^2}. \tag{37}$$

First we consider the ideal case. That is, using $\widehat{U}_{\mathrm{EE}}(H, \delta E, \eta = 0)$ (Equation Equation (29)) instead of U_{EE} in Step 2.1.2. According to Definition 31, if $E_j \in I_{k,l}$, it holds that

$$\widehat{U}_{\mathrm{EE}}(H, \delta E, \eta = 0) \left|\psi_j\right\rangle \left|0^{a_{\mathrm{EE}}}\right\rangle = \left|\psi_j\right\rangle \otimes \sum_{E' \in I'_{k,\ell}} \alpha_{E'} \left|E'\right\rangle \left|g_{E'}\right\rangle;$$

and if $E_j \notin I''_{k,\ell}$, it holds that

$$\widehat{U}_{\mathrm{EE}}(H, \delta E, \eta = 0) \left|\psi_j\right\rangle \left|0^{a_{\mathrm{EE}}}\right\rangle = \left|\psi_j\right\rangle \otimes \sum_{E' \notin I'_{k,\ell}} \alpha_{E'} \left|E'\right\rangle \left|g_{E'}\right\rangle.$$

In other words, letting $\widehat{U}_{k,\ell} := U_{\mathrm{dec}}(I'_{k,\ell})\widehat{U}_{\mathrm{EE}}(H, \delta E, \eta = 0)$, we have

$$\widehat{U}_{k,\ell} \left|\psi_j\right\rangle \left|0^{a_U}\right\rangle = \begin{cases} \left|1\right\rangle \left|\xi^1_j\right\rangle & , \text{ if } E_j \in I_{k,\ell}, \\ \alpha_j \left|1\right\rangle \left|\xi^1_j\right\rangle + \beta_j \left|0\right\rangle \left|\xi^0_j\right\rangle & , \text{ if } E_j \in I^{-}_{k,\ell} \cup I^{+}_{k,\ell}, \\ \left|0\right\rangle \left|\xi^0_{\ell,j}\right\rangle & , \text{ if } E_j \notin I''_{\ell}, \end{cases} \tag{38}$$

where $\alpha_j, \beta_j \in \mathbb{C}$ satisfy $|\alpha_{\ell,j}| + |\beta_{\ell,j}| = 1$, $|\xi^1_j\rangle$, $|\xi^0_j\rangle$ are quantum states depending on j. The coefficients α_j, β_j and the states $\left|\xi^1_j\right\rangle, \left|\xi^0_j\right\rangle$ may depend on k, ℓ.

When $\widehat{U}_{k,\ell}$ acts on the first half of $|\Psi\rangle$, combining Equation (38) and Lemma 34, we have

$$\widehat{U}_{k,\ell} \left|\Psi\right\rangle \left|0^{a_U}\right\rangle = \sqrt{\frac{\widehat{M}_{k,\ell}}{N}} \left|1\right\rangle \left|\xi^1\right\rangle + \sqrt{\frac{N - \widehat{M}_{k,\ell}}{N}} \left|0\right\rangle \left|\xi^0\right\rangle, \tag{39}$$

where $\widehat{M}_{k,\ell} \in [N]$ satisfies

$$M_{k,\ell} \leq \widehat{M}_{k,\ell} \leq M^{-}_{k,\ell} + M_{k,\ell} + M^{+}_{k,\ell}, \tag{40}$$

and $|\xi^1\rangle, |\xi^0\rangle$ are some quantum states.

Let $\widetilde{M}^{ideal}_{k,\ell}$ be the output in Step 2.1.4. when replacing $U_{k,\ell}$ with $\widehat{U}_{k,\ell}$. By the guarantee of quantum counting (Lemma 33), we have that

$$(1 - \delta_{\mathrm{C}})\widehat{M}_{k,\ell} \leq \widetilde{M}^{ideal}_{k,\ell} \leq (1 + \delta_{\mathrm{C}})(\widehat{M}_{k,\ell}). \tag{41}$$

By union bound, the event that Equation Equation (41) holds for all k, ℓ has probability $1 - \mathrm{negl}(n)$.

Now we lower bound $\widetilde{Z}$. Because $M_{k,\ell} \leq \widehat{M}_{k,\ell}$, from Equation Equation (41), we have

$(1-\delta_{\mathrm{C}})M_{k,\ell} \leq \widetilde{M}_{k,\ell}^{ideal}$.

Multiplying by $e^{-\beta\mathcal{E}_{k,\ell}}$ and combining with Equation Equation (34), we have

$$(1-\delta_{\mathrm{C}})Z_{k,\ell} \leq \widetilde{M}_{k,\ell}^{ideal} e^{-\beta\mathcal{E}_{k,\ell}}. \tag{42}$$

For any k, when summing ℓ over $\{0,1,\ldots,L\}$ in Equation Equation (42), we have

$$(1-\delta_{\mathrm{C}})Z \leq \sum_{\ell} \widetilde{M}_{k,\ell}^{ideal} e^{-\beta\mathcal{E}_{k,\ell}}. \tag{43}$$

For the upper bound, because $\widehat{M}_{k,\ell} \leq M_{k,\ell}^{-} + M_{k,\ell} + M_{k,\ell}^{+}$, combining with the guarantee of quantum counting (Equation Equation (41)), we have

$\widetilde{M}_{k,\ell}^{ideal} \leq (1+\delta_{\mathrm{C}})(M_{k,\ell}^{-} + M_{k,\ell} + M_{k,\ell}^{+})$.

Multiplying by $e^{-\beta\mathcal{E}_{k,\ell}}$ and combing with Equation Equation (36) and Equation (37), we have

$$\widetilde{M}_{k,\ell}^{ideal} e^{-\beta\mathcal{E}_{k,\ell}} \leq (1+\delta_{\mathrm{C}})(Z_{k,\ell}^{-} + Z_{k,\ell}e^{-\beta/L} + Z_{k,\ell}^{+}e^{\beta/L+\beta/L^2}). \tag{44}$$

Fixing a k and summing ℓ over $\{0,1,\ldots,L\}$ in Equation Equation (44), we have

$$\sum_{\ell} \widetilde{M}_{k,\ell}^{ideal} e^{-\beta\mathcal{E}_{k,\ell}} \leq (1+\delta_{\mathrm{C}})\Big((\sum_{\ell} Z_{k,\ell}^{-}) + Ze^{\beta/L} + (\sum_{\ell} Z^{+}e^{\beta/L+\beta/L^2})\Big). \tag{45}$$

Consider the term $\sum_{\ell} Z_{k,\ell}^{-} + \sum_{\ell} Z^{+}e^{\beta/L+\beta/L^2}$ and sum over k. That is

$$\sum_{k=0}^{L-1}\Big(\sum_{\ell} Z_{k,\ell}^{-} + \sum_{\ell} Z^{+}e^{\beta/L+\beta/L^2}\Big). \tag{46}$$

Since

$$\bigcup_{k=0}^{L-1}\bigcup_{\ell} I_{k,\ell}^{-} = [-1/L, 0) \cup [0, 1-1/L^2),$$

and all the eigenstates are in the interval $[0,1)$, we have that $\sum_{k=0}^{L-1}\sum_{\ell} Z_{k,\ell}^{-}$ is contributed by the eigenstates in the interval $[0, 1-1/L^2)$. Therefore,

$$\sum_{k=0}^{L-1}\sum_{\ell} Z_{k,\ell}^{-} \leq Z. \tag{47}$$

Similarly, we have

$$\sum_{k=0}^{L-1}\sum_{\ell} Z_{k,\ell}^{+} \leq Z. \tag{48}$$

Plugging Equation Equation (47) and Equation (48) into Equation Equation (46), we have

$$\sum_{k=0}^{L-1}\Big(\sum_{\ell} Z_{k,\ell}^{-} + \sum_{\ell} Z^{+}e^{\beta/L+\beta/L^2}\Big) \leq Z(1+e^{\beta/L+\beta/L^2}). \tag{49}$$

By average argument, there is a k such that

$$\Big(\sum_{\ell} Z^{-}_{k,\ell} + \sum_{\ell} Z^{+} e^{\beta/L+\beta/L^2}\Big) \leq \frac{1}{L} Z(1 + e^{\beta/L+\beta/L^2}). \tag{50}$$

Let $\widetilde{Z}^{ideal}$ be the output of the algorithm when replacing $U_{k,\ell}$ with $\widetilde{U}_{k,\ell}$. By Equation Equation (45) and Equation (50), we have

$$\widetilde{Z}^{ideal} \leq (1+\delta_{\mathrm{C}})\big(\frac{1}{L} + e^{\beta/L} + \frac{1}{L} e^{\beta/L+\beta/L^2}\big) Z. \tag{51}$$

Recall that $\delta_{\mathrm{C}} = \frac{1}{4n^c}$, $L = 4n^c\beta$, and $\beta = \mathrm{poly}(n)$. We have

$$\frac{1}{L} + e^{\beta/L} + \frac{1}{L} e^{\beta/L+\beta/L^2} = e^{\frac{1}{4n^c}} + \frac{e^{\frac{1}{4n^c}(1+\frac{1}{4n^c\beta})}}{4n^c\beta} + \frac{1}{4n^c\beta} \leq 1 + \frac{1}{2n^c}$$

for sufficiently large n. As a result, Equation Equation (51) becomes

$$\widetilde{Z}^{ideal} \leq \big(1 + \frac{1}{4n^c}\big)\big(1 + \frac{1}{2n^c}\big) Z \leq \big(1 + \frac{1}{n^c}\big) Z \tag{52}$$

for sufficiently large n.

Finally, combing Equation Equation (43) and Equation (52), we have

$$(1 - 1/n^c) Z \leq \widetilde{Z}^{ideal} \leq (1 + 1/n^c) Z. \tag{53}$$

Then, we consider the real implementation. There are two steps to change the ideal algorithm to the real implementation. First, we define a intermediate algorithm by replacing $\widehat{U}_{\mathrm{EE}}(H, \delta E, \eta = 0)$ in the ideal algorithm with $\widehat{U}_{\mathrm{EE}}(H, \delta E, \eta = e^{-n})$. By Lemma 33, there are $O(\mathrm{poly}(n) \cdot 2^{\frac{n}{2}})$ many $\widehat{U}_{EE}$ in the algorithm. Also, we have $\|\widehat{U}_{\mathrm{EE}}(H, \delta E, \eta = e^{-n}) - \widehat{U}_{\mathrm{EE}}(H, \delta E, \eta = 0)\| \leq e^{-n}$. By triangular inequity, the distance between the ideal algorithm and the intermediate algorithm is $\mathrm{negl}(n)$.

For the second step, we get the real implementation by replacing $\widehat{U}_{\mathrm{EE}}(H, \delta E, \eta = e^{-n})$ with $U_{\mathrm{EE}}(H, \delta E, \eta = e^{-n}, \varepsilon = 2^{-n})$. We have $\|U_{\mathrm{EE}}(H, \delta E, \eta = e^{-n}, \varepsilon = 2^{-n}) - \widehat{U}_{\mathrm{EE}}(H, \delta E, \eta = e^{-n})\| \leq 2^{-n}$. Again, by triangular inequity, the distance between the intermediate algorithm and the real implementation is $\mathrm{negl}(n)$. As a result, the distance between ideal algorithm and the real implementation is $\mathrm{negl}(n)$. Constituently, the real implementation outputs the correct estimation of quantum partition function with probability $1 - \mathrm{negl}(n)$.

Finally, we analyze the running time. There are $\mathrm{poly}(n)$ of iterations over k and ℓ and each iteration runs in $O(\mathrm{poly}(n) \cdot 2^{\frac{n}{2}})$ time. Therefore the overall running time is $O^*(2^{\frac{n}{2}})$. ◀

References

1. Scott Aaronson, Nai-Hui Chia, Han-Hsuan Lin, Chunhao Wang, and Ruizhe Zhang. On the quantum complexity of closest pair and related problems. In *Proceedings of the 35th Computational Complexity Conference (CCC 2020)*, pages 16:1–16:43, 2020. doi:10.4230/LIPIcs.CCC.2020.16.
2. Brian Alspach. Johnson graphs are Hamilton-connected. *Ars Mathematica Contemporanea*, 6(1):21–23, 2012. doi:10.26493/1855-3974.291.574.
3. Joran van Apeldoorn, András Gilyén, Sander Gribling, and Ronald de Wolf. Quantum SDP-Solvers: Better upper and lower bounds. *Quantum*, 4:230, 2020. doi:10.22331/q-2020-02-14-230.
4. Yosi Atia and Dorit Aharonov. Fast-forwarding of Hamiltonians and exponentially precise measurements. *Nature Communications*, 8(1):1572, 2017. doi:10.1038/s41467-017-01637-7.

5 Arturs Backurs and Piotr Indyk. Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing (STOC 2015)*, pages 51–58, 2015. doi:10.1145/2746539.2746612.

6 Marshall Ball, Alon Rosen, Manuel Sabin, and Prashant Nalini Vasudevan. Average-case fine-grained hardness. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2017)*, pages 483–496, 2017. doi:10.1145/3055399.3055466.

7 Jacob D. Biamonte and Peter J. Love. Realizable hamiltonians for universal adiabatic quantum computers. *Physical Review A*, 78:012352, 2008. doi:10.1103/PhysRevA.78.012352.

8 Fernando GSL Brandão. Entanglement theory and the quantum simulation of many-body physics. *arXiv preprint*, 2008. arXiv:0810.0026.

9 Sergey Bravyi, Anirban Chowdhury, David Gosset, and Pawel Wocjan. Quantum Hamiltonian complexity in thermal equilibrium. *Nature Physics*, 18(11):1367–1370, 2022.

10 Sergey Bravyi, David Gosset, Robert König, and Kristan Temme. Approximation algorithms for quantum many-body problems. *Journal of Mathematical Physics*, 60(3):032203, March 2019. doi:10.1063/1.5085428.

11 Harry Buhrman, Sevag Gharibian, Zeph Landau, François Le Gall, Norbert Schuch, and Suguru Tamaki. Beating the natural Grover bound for low-energy estimation and state preparation. *Physical Review Letters*, 135:030601, 2025. doi:10.1103/29qw-bssx.

12 Harry Buhrman, Subhasree Patro, and Florian Speelman. A framework of quantum strong exponential-time hypotheses. In *Proceedings of the 38th International Symposium on Theoretical Aspects of Computer Science (STACS 2021)*, pages 19:1–19:19, 2021. doi:10.4230/LIPIcs.STACS.2021.19.

13 Yanlin Chen, Yilei Chen, Rajendra Kumar, Subhasree Patro, and Florian Speelman. QSETH strikes again: Finer quantum lower bounds for lattice problem, strong simulation, hitting set problem, and more. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2025)*, pages 6:1–6:24, 2025. doi:10.4230/LIPIcs.APPROX/RANDOM.2025.6.

14 Nai-Hui Chia, Kai-Min Chung, Yao-Ching Hsieh, Han-Hsuan Lin, Yao-Ting Lin, and Yu-Ching Shen. On the impossibility of general parallel fast-forwarding of Hamiltonian simulation. In *Proceeding of the 38th Computational Complexity Conference (CCC 2023)*, pages 33:1–33:45, 2023. doi:10.4230/LIPIcs.CCC.2023.33.

15 Anirban N Chowdhury, Rolando D Somma, and Yiğit Subaşı. Computing partition functions in the one-clean-qubit model. *Physical Review A*, 103(3):032422, 2021.

16 Stephen Cook. The complexity of theorem proving procedures. In *Proceedings of the 3rd ACM Symposium on Theory of Computing (STOC 1971)*, pages 151–158, 1971. doi:10.1145/800157.805047.

17 Toby Cubitt and Ashley Montanaro. Complexity classification of local Hamiltonian problems. *SIAM Journal on Computing*, 45(2):268–316, 2016. doi:10.1137/140998287.

18 Alexander M. Dalzell, Nicola Pancotti, Earl T. Campbell, and Fernando G. S. L. Brandão. Mind the Gap: Achieving a Super-Grover Quantum Speedup by Jumping to the End. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC 2023)*, pages 1131–1144, 2023. doi:10.1145/3564246.3585203.

19 Roee David, Karthik C. S., and Bundit Laekhanukit. On the complexity of closest pair via polar-pair of point-sets. *SIAM Journal on Discrete Mathematics*, 33(1):509–527, 2019. doi:10.1137/17M1128733.

20 Ronald de Wolf. Quantum computing: Lecture notes. *arXiv preprint arXiv:1907.09415*, 2019. arXiv:1907.09415.

21 Richard P Feynman. Quantum mechanical computers. *Foundations of physics*, 16(6):507–531, 1986.

22 Yimin Ge, Jordi Tura, and J. Ignacio Cirac. Faster ground state preparation and high-precision ground energy estimation with fewer qubits. *Journal of Mathematical Physics*, 60(2):022202, February 2019. doi:10.1063/1.5027484.

23 Sevag Gharibian, Yichen Huang, Zeph Landau, and Seung Woo Shin. Quantum Hamiltonian complexity. *Foundations and Trends in Theoretical Computer Science*, 10(3):159–282, 2015. doi:10.1561/0400000066.

24 András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC 2019)*, pages 193–204, 2019. doi:10.1145/3313276.3316366.

25 Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC 1996)*, pages 212–219, 1996. doi:10.1145/237814.237866.

26 M. B. Hastings. A Short Path Quantum Algorithm for Exact Optimization. *Quantum*, 2:78, 2018. doi:10.22331/q-2018-07-26-78.

27 Cupjin Huang, Michael Newman, and Mario Szegedy. Explicit lower bounds on strong quantum simulation. *IEEE Transactions on Information Theory*, 66(9):5585–5600, 2020. doi:10.1109/TIT.2020.3004427.

28 Jeremy Ahrens Huang, Young Kun Ko, and Chunhao Wang. On the (classical and quantum) fine-grained complexity of log-approximate CVP and Max-Cut. *arXiv preprint arXiv:2411.04124*, 2024. doi:10.48550/arXiv.2411.04124.

29 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. *Journal of Computer and System Sciences*, 62(2):367–375, 2001. doi:10.1006/jcss.2000.1727.

30 Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *Journal of Computer and System Sciences*, 63(4):512–530, 2001. doi:10.1006/JCSS.2001.1774.

31 C. S. Karthik and Pasin Manurangsi. On closest pair in euclidean metric: Monochromatic is as hard as bichromatic. *Combinatorica*, 40(4):539–573, 2020. doi:10.1007/s00493-019-4113-1.

32 Julia Kempe, Alexei Y. Kitaev, and Oded Regev. The complexity of the local Hamiltonian problem. *SIAM Journal on Computing*, 35(5):1070–1097, 2006. doi:10.1137/S0097539704445226.

33 Alex Kerzner, Vlad Gheorghiu, Michele Mosca, Thomas Guilbaud, Federico Carminati, Fabio Fracas, and Luca Dellantonio. A square-root speedup for finding the smallest eigenvalue. *Quantum Science and Technology*, 6(4):045025, 2024. doi:10.1088/2058-9565/ad6a36.

34 Alexei Yu. Kitaev, Alexander H. Shen, and Mikhail N. Vyalyi. *Classical and Quantum Computation*, volume 47 of *Graduate Studies in Mathematics*. American Mathematical Society, 2002.

35 Emanuel Knill and Raymond Laflamme. Power of one bit of quantum information. *Physical Review Letters*, 81(25):5672, 1998.

36 Jacek Kuczyński and Henryk Woźniakowski. Estimating the largest eigenvalue by the power and Lanczos algorithms with a random start. *SIAM Journal on Matrix Analysis and Applications*, 13(4):1094–1122, 1992. doi:10.1137/0613066.

37 Cornelius Lanczos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *Journal of Research of the National Bureau of Standards*, 45(4):255–282, 1950. doi:10.6028/jres.045.026.

38 Leonid Levin. Universal search problems. *Problems of Information Transmission*, 9(3):265–266, 1973.

39 Lin Lin and Yu Tong. Near-optimal ground state preparation. *Quantum*, 4:372, 2020. doi:10.22331/q-2020-12-14-372.

40 Sixue Liu. Chain, generalization of covering code, and deterministic algorithm for k-SAT. In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, pages 88:1–88:13, 2018. doi:10.4230/LIPIcs.ICALP.2018.88.

41 Guang Hao Low and Isaac L. Chuang. Optimal Hamiltonian simulation by quantum signal processing. *Phys. Rev. Lett.*, 118:010501, 2017. doi:10.1103/PhysRevLett.118.010501.

42 John M. Martyn, Zane M. Rossi, Andrew K. Tan, and Isaac L. Chuang. Grand unification of quantum algorithms. *PRX Quantum*, 2:040203, 2021. doi:10.1103/PRXQuantum.2.040203.

43 Jesper Nederlof. An invitation to "fine-grained complexity of NP-complete problems". *arXiv preprint arXiv:2601.05044*, 2026. doi:10.48550/arXiv.2601.05044.

44 Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.

45 Roberto Oliveira and Barbara M. Terhal. The complexity of quantum spin systems on a two-dimensional square lattice. *Quantum Information and Computation*, 8(10):900–924, 2008. doi:10.26421/QIC8.10-2.

46 Stephen Piddock and Ashley Montanaro. The complexity of antiferromagnetic interactions and 2D lattices. *Quantum Information and Computation*, 17(7-8):636–672, 2017. doi:10.26421/QIC17.7-8-6.

47 David Poulin and Pawel Wocjan. Preparing ground states of quantum many-body systems on a quantum computer. *Physical Review Letters*, 102(13):130503, 2009. doi:10.1103/PHYSREVLETT.102.130503.

48 David Poulin and Pawel Wocjan. Sampling from the Thermal Quantum Gibbs State and Evaluating Partition Functions with a Quantum Computer. *Phys. Rev. Lett.*, 103:220502, November 2009. doi:10.1103/PhysRevLett.103.220502.

49 Virginia Vassilevska Williams. Hardness of easy problems: Basing hardness on popular conjectures such as the strong exponential time hypothesis. In *Proceedings of the 10th International Symposium on Parameterized and Exact Computation (IPEC 2015)*, pages 17–29, 2015. doi:10.4230/LIPIcs.IPEC.2015.17.

50 Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity. In *Proceedings of the International Congress of Mathematicians: Rio de Janeiro 2018*, pages 3447–3487. World Scientific, 2018.

51 Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theoretical Computer Science*, 348(2):357–365, 2005. doi:10.1016/j.tcs.2005.09.023.

A A trivial fine-grained reduction to $k(\varepsilon)$LH from $k(\varepsilon)$SAT

▶ **Theorem 35** (Lower bound of $k(\varepsilon)$LH). *Assuming SETH (resp. QSETH), for any $\varepsilon > 0$, there is k (depending on ε) such that for any algorithm, there exists $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$, there is a Hamiltonian H acting on n qubits, associated $E_{\mathrm{yes}}, E_{\mathrm{no}}$ satisfying $E_{\mathrm{no}} - E_{\mathrm{yes}} \geq \Omega(1)$ such that $\mathrm{LH}(H, E_{\mathrm{yes}}, E_{\mathrm{no}})$ cannot be solved in $O(2^{n(1-\varepsilon)})$ classical time (resp. $O(2^{\frac{n}{2}(1-\varepsilon)})$ quantum time).*

Proof. We construct a Hamiltonian H from a k-SAT instance $\Phi = \varphi_1 \wedge \varphi_2 \wedge \cdots \wedge \varphi_m$.

Let $S_i \subseteq [n]$ be the set collecting the index j such that x_j or $\neg x_j$ appears in φ_i, and let $\mathsf{S_i}$ be the corresponding register. Let $y_i \in \{0,1\}^{|S_i|}$ be the assignment to the variables appearing in φ_i such that $\varphi_i(y_i) = 0$. Because φ_i is in the disjunctive form, y_i is unique. Let $H := \sum_{i=1}^{m} H_i$, where $H_i := I - |y_i\rangle\langle y_i|_{\mathsf{S_i}}$ for all $i \in [m]$. It holds that $H_i |y_i\rangle = 0$ if $\varphi_i(y_i) = 1$ and $H_i |y_i\rangle = |y_i\rangle$ if $\varphi_i(y_i) = 0$. Each H_i acts non-trivially on at most k qubits. Hence, H is k-local.

If there exists an assignment $x \in \{0,1\}^n$ satisfying Φ, then $H|x\rangle = 0$. Otherwise, $\lambda(H) \geq 1$. Hence, solving $\mathrm{LH}(H, E_{\mathrm{yes}} = 0, E_{\mathrm{no}} = 1)$ can decide k-SAT. The construction of H takes $\mathrm{poly}(n)$ time. Hence, if $\mathrm{LH}(H, E_{\mathrm{yes}} = 0, E_{\mathrm{no}} = 1)$ is solved in $O(2^n(1-\varepsilon))$ time, then k-SAT(Φ) is also solved in $O(2^{n(1-\varepsilon)})$ classical time, which violates SETH. The argument assuming QSETH follows in the same proof strategy. ◀